

Bangor University

DOCTOR OF PHILOSOPHY

How attention influences the emotional evaluation of complex objects

Westoby, Nikki

Award date:
2008

Awarding institution:
Bangor University

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?


Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

How Attention Influences the Emotional Evaluation of Complex Objects

Nikki Westoby, BSc, MSc

This thesis is submitted in partial fulfilment of the requirements for the degree of Doctor of
Philosophy, completed in the School of Psychology, Bangor University.



**MISSING
PAGES
REMOVED ON
INSTRUCTION
FROM THE
UNIVERSITY**

Acknowledgments

First, I would like to thank Jane Raymond. Her supervision, support and encouragement extend well beyond the realms of this thesis. For helping me develop my research skills, for the opportunities in consumer research, for your mum's lemon squares, and for introducing me to Ronan, I will always be grateful to you.

In the lab, I wish to acknowledge Mark Fenske, who helped me find my feet in the early days and to Madge Jackson who spurred me on after the non-significant findings. Also, to the latest reincarnation of the Raymond lab, especially Carys, Helena and Jen. For being an amazing supportive network that I could call on for anything, for finding references, for fighting with the printer on my behalf (and that's just in the last week!), thank you.

Outside of the university, I want to thank my friends Kirsty and Liz for their continued friendship and understanding, and for not taking it personally when I get lost in my work for months and forget to call. You both mean the world to me. Also, my landlady, Anne, for being my second mum, and for making me feel at home for the last six years.

Finally my family, Mum, Dad, Paul and Ronan. How do I even begin? You have all been so wonderful to me. To Mum and Dad for giving me the most amazing start in life, and for continuing to support me in so many ways. For teaching me to read and love books, for years of driving me to flute lessons, tap dancing, ballet, ice-skating, golf... no matter how crazy life might have been, for sending TV programs I missed through the post, and even the odd bit of E-Prime programming, I will never be able to thank you enough. To Paul, and Sammy, for reminding me what's important and for keeping me grounded. And finally to Ronan, who quite simply has made me happier than I've ever been. I love you.

Dedication

I dedicate the first page of this thesis to Mum and Dad. Without your unwavering support, I would never have had the opportunity to start a PhD.

I dedicate the last page to Ronan. The end of this is the beginning of our new life together.

The rest of the pages are up for grabs.

Table of Contents

Abstract	1
Chapter 1: An introduction to attention, emotion, and their interactions	2
Perceiving objects	7
Categorisation	9
Types, tokens and repetition blindness	12
Attention	14
Visual search	15
Inhibition of object features	18
Inhibition of locations	19
Inhibition of return	20
Inhibition of objects	22
Inhibition of object identity	23
Retrieval of attention processes from memory	23
Long-term IOR	24
Negative priming	26
Emotion	28
Some early theories of emotion (as ‘feelings’)	29
Emotional evaluation	33
Intuitive appraisal	34
Affective priming	36
Is attention necessary for emotional evaluation?	39
Reflective appraisal	42
The results of emotional evaluation	44
The role of emotions in decision making	45
Emotions enhance perception	48
Emotions as cues to action	50
Emotions enhance memory	54
Emotion biases attention	56
Neuropsychology	57
Behavioural evidence	57
Do negative stimuli bias attention more than positive stimuli?	59
Arousal theories of attentional capture	61
Biasing emotional evaluation	66
Classical conditioning	67
Evaluative conditioning	68
Affective priming of neutral stimuli	69
Mere exposure	70
Emotion is biased by attention	74
Neurological evidence for an attention-emotion link	75
Distractor devaluation	79

Chapter 2: An introduction to, and general methods for, an oddball search paradigm	91
General methods	96
Chapter 3: Pilot studies of the oddball search paradigm	102
Experiment 1: Oddball search pilot study	103
Method	103
Results	107
Discussion	111
Experiment 2: The effect of a lag following an oddball search	112
Method	113
Results	114
Discussion	118
Chapter 3 discussion	119
Chapter 4: Token effects on evaluation in an oddball search paradigm	122
Experiment 3: Token distractor devaluation effects	123
Method	125
Results	129
Discussion	133
Chapter 5: Type effects on evaluation in an oddball search paradigm	135
Experiment 4: Type distractor devaluation effects	136
Method	137
Results	139
Discussion	143
Experiment 5: Type effects of faces in an oddball search paradigm	144
Method	144
Results	147
Discussion	153
Generalisation of mere exposure	153
Distractor devaluation and effort	155
Chapter 6: Distractor devaluation following behavioural inhibition	160
Experiment 6a: Devaluation of no-go faces - baseline	169
Method	170
Results	171
Discussion	172
Experiment 6b: Devaluation of no-go faces	174
Method	175
Results	179
Discussion	193
Distractor devaluation of conjunction no-go's	193
No distractor devaluation for featural no-go's	194
Effort	198

Flattening versus devaluation	199
Why were negative faces devalued more than positive faces?	200
Scene versus object recognition	204
Novelty preference	205
Chapter 7: Devaluation of own- versus other-race faces	207
Experiment 7a: Devaluation of own- and other-race faces – Race task	211
Method	211
Results	215
Discussion	218
Experiment 7b: Devaluation of own- and other-race faces – Gender task	219
Method	220
Results	221
Discussion	223
Chapter 7 discussion	223
Chapter 8: Devaluation of own- versus other-age faces	225
Experiment 8: Generalisation of distractor devaluation to own- and other-age group faces	228
Method	230
Results	234
Discussion	241
Chapter 9: Generalisation of distractor devaluation in branded products	245
Experiment 9: Generalisation of distractor devaluation in branded products	247
Method	248
Results	252
Discussion	255
Chapter 10: General discussion	258
Attention and emotion	259
Summary of findings: The individuation hypothesis	260
Summary of findings: The generalisation hypothesis	262
Inhibition of categories	266
Structural versus semantic object representations	268
Other generalisation findings	270
Stereotyping	273
The relative contribution of distractor devaluation generalisation	276
Outstanding issues and further research	278
Conclusions	284
References	287

Abstract

Humans are not always privy to the mechanisms that drive their decisions, choices, or preferences. In situations where ability or motivation to evaluate an object is lacking, heuristics can be applied to the situation allowing past experience to influence current decision-making. The objective of this thesis was to explore whether the distractor devaluation effect (Raymond et al., 2003) might be one such heuristic. In a series of nine experiments I demonstrated that inhibiting an object could have detrimental effects on the emotional evaluation of another, similar, object. In Experiments 3, 6 and 8 I replicated the distractor devaluation at the individual (inhibited object) level. In Experiment 4 I found evidence that distractor devaluation could generalise to objects of same basic-level category as the inhibited item. In Experiments 7 and 8 I used in- and out-group faces to show that members of a subordinate category can also be devalued. In Experiment 9 I used branded products and demonstrated that distractor devaluation can generalise within superordinate categories of objects. To account for these findings, and in line with the 'devaluation-by-inhibition' hypothesis, I propose that in cases where the initial to-be-inhibited object is not successfully individuated, an inhibitory tag can be associated instead with a whole category representation. This inhibitory tag can be subsequently used to guide the evaluation of a previously unseen member of that category (the generalisation effect). I also propose that the situational demands at the time of inhibition determine the extent of this generalisation process.

Chapter 1

An Introduction to Attention, Emotion, and their Interactions

Life is full of choices. But the choice we make at one point in time may not be the same choice we make, in the same situation, at a different point in time. What drives our preference for one choice over another? What forces are at work to change our preferences? In deciding where to sit on a crowded train, or who to vote for in the upcoming election, how do we decide which available seat, or candidate, we prefer?

Two systems lie at the heart of this decision making process; both act in concert to prioritise objects for consideration and help us make our choice. The first of these systems is a selective attention system that facilitates the processing of task-relevant objects and suppresses processing of irrelevant distracting objects (Kastner & Ungerleider, 2000). Theories of attention attempt to describe how an object might gain priority for processing in order to activate a response. But where they fall short is in the link between the encoding stage and the response elicited. Humans are not robots: we do not act on an object simply because it affords an action. We act on an object (or do not) because it benefits us to do so. Between processing an object and acting on it, we predict its value, and the reward to be gained.

Therefore, the second system at work is an emotion system that evaluates object representations in terms of current and future goals (Ortony, Clore & Collins, 1988). These systems work together to minimise conflict and promote rapid responses. For example, when stepping onto a crowded train, the selective attention system must engage to search for the empty seats, and the emotional system will then tell us which person to sit next to by evaluating who looks the most trustworthy.

Below, I describe each of these systems in greater detail and present evidence that brain structures common to the two systems are selectively activated during selective

attention and emotional evaluation tasks (Armony & Dolan, 2002; Bush, Luu, & Posner, 2000; Vuilleumier, Armony, Driver, & Dolan, 2001). Following this I present some of the evidence describing how emotional stimuli produce specific effects on selective attention tasks by attracting and holding attention relative to non-emotional stimuli (Eastwood, Smilek & Merikle, 2001; Fox, Russo, Bowles & Dutton, 2001). I then present recent work that describes the reciprocal effect: the effect of selective attention on emotional evaluation.

Until recently the latter effect was unexplored. However there is an abundance of research indicating that perceptual experience can alter emotional evaluation. For example repetition (Bornstein, 1989; Zajonc, 2001), brightness (Reber, Winkielman & Schwarz, 1998), and prior presentation of other emotionally salient stimuli (Murphy & Zajonc, 1993), have been shown to make evaluations of stimuli more positive. No study had manipulated the participant's attentional state allowing the effect of attention on emotional responses to be explored. In 2003 Raymond and colleagues (Raymond, Fenske & Tavassoli, 2003) began work to explore the possibility that prior inhibition to an object will devalue that object upon later emotional evaluation.

My thesis stands on the shoulders of this work. Raymond et al. (2003) found that if an object is inhibited, it will be emotionally devalued on subsequent evaluation. They proposed the 'devaluation-by-inhibition' theory to account for their findings. They posit that an inhibitory tag will be stored with the distractor's representation in memory. On subsequent evaluation, this inhibitory tag is interpreted as an emotional devaluation. I pose the following questions: Is it only the inhibited object that gets devalued, or are similar objects also devalued? When are similar objects subjected to this devaluation process, and when are they not? I propose that what occurs during encoding, specifically during

categorisation of the inhibited object, will affect whether that object, or one similar, will be devalued.

Categorisation principals suggest that an object can be encoded as a unique item (token categorisation). Alternatively, an item may be encoded as belonging to a group of items, but not encoded as a unique exemplar (type categorisation). For example if I see a dog in the park, there are a number of ways I may encode him. I may encode him as: a) an animal (superordinate category); b) a dog (basic-level category); c) a collie (subordinate category) or; d) Laird (token exemplar). I propose that the level of categorisation an inhibited object reaches during encoding determines how the inhibitory tag is stored, and subsequently how this tag is applied to an emotional evaluation.

Imagine you are in a park searching for your child who just ran off to play on the swings. A dog ('Laird') is barking at you to come and play 'Fetch'. This dog is interfering with your search for the child and must be ignored. After the child is found, how will you emotionally evaluate any dogs that want you to play Fetch in the future? I propose that if the original distracting dog was recognised as 'Laird', the inhibitory tag would have been stored with 'Laird's' representation, and only Laird will be devalued. So if you evaluate Laird he will be devalued, if you evaluate Daisy she will not. Alternatively, if the original dog was recognised as a 'collie' but not specifically as Laird, the category of 'collie' will receive the inhibitory tag. So, if you evaluate Laird he will be devalued, if you evaluate Daisy (another collie) she will also be devalued. However next door's dog Butch (a Boxer) will not be devalued. Alternatively, if your dog recognition skills are particularly bad, the barking dog may only have been categorised as a 'dog', and so the inhibitory tag will have been stored with the category 'dog'. Thus Laird, Daisy and Butch will all be devalued on subsequent evaluation.

A theory like this is particularly important when applied to the issue of stereotyping. Stereotyping literature suggests that, to make evaluation easier, classes of people are grouped and heuristics applied to the whole group (Fiske & Neuberg, 1990). In this case, can we find evidence for devaluation to a whole group of people who share something in common (e.g. race) with a previously ignored person? This may be especially true if the ignored person is not recognised as a unique exemplar when encountered. If members of another race “all look alike to me”¹, then they are unlikely to be categorised as a token exemplar. As such, the inhibitory tag may be stored with a whole group of people (e.g. Asian people) and any member of that group to be emotionally evaluated will be devalued. Given the compulsion to apply heuristics to person evaluation, is devaluation-by-inhibition at work proliferating and compounding stereotypes of minority groups that are frequently ignored?

To address these questions, this thesis presents a series of nine experiments using two different paradigms: oddball search and go / no-go. Experiments 1 and 2 establish a new paradigm for investigating the distractor devaluation effect: the oddball search paradigm. Experiments 3, 6 and 8 demonstrate distractor devaluation at the individual level. Experiment 4 finds evidence for distractor devaluation for stimuli that belong to the same basic-level category as the inhibited item. In Experiments 7 and 8 I use face stimuli to show that when a face is categorised to the individual level (in-group faces), distractor devaluation affects only that face. When a face is not recognised as an individual (out-group faces), members of the same subordinate category are also devalued. In Experiment 9 I use

¹ From Malpass (1981) who noted that this was the general reaction of White Americans and White Europeans when confronted with faces of Asians, African Americans or Black Africans.

branded products and investigate whether distractor devaluation effects can generalise within superordinate categories (i.e. between basic-level categories).

Perceiving Objects

Upon presentation of an object, how do we recognise it as a dog, a chair, a motorbike or our dad? Three perceptual routes have been proposed to lead to the formation of an object representation, these are: featural, configural and holistic.

Featural processes reflect the analysis of small structural and surface features of an object. Structural features of an object can be two-dimensional lines, vertices and edges (Tarr & Bülthoff, 1995; Edelman, 1998). Alternatively, they may be more complex three-dimensional parts (Marr & Nishihara, 1978; Marr, 1982). The colour and texture of an object are surface features that are also analysed featurally. Configural processes refer to the analysis of the relationships between features (structures or surfaces). Holistic processing involves the analysis of more abstract information such as the shape of the object, and does not rely on fine-grained visual detail (Peterson & Rhodes, 2003). Therefore in order to recognise a dog, the colour and texture of its fur may be analysed featurally, along with the shape of its ears; configural analysis may focus on the distance between the ears and the ratio of body and leg length; and the silhouette of the head, body and tail may be incorporated into holistic analysis.

How these three perceptual processing routes combine to result in object recognition has been the subject of much theorising. Two theories that dominate are Recognition-by-Components (RBC; Biederman, 1987) and Global Precedence (Navon, 1977).

RBC can be interpreted as a combination of featural and configural processes. The theory proposes that object representations are composed of small units called 'geons'. Geons are three-dimensional shapes such as cylinders, cones, blocks and wedges. Biederman defined 36 qualitatively different geons. Information about the shape of a geon, and its spatial relation to other geons in the object representation is interpreted and leads to object recognition. A limitation of RBC is that the analysis it proposes is not discrete enough to identify one collie, for example, from another. The two dogs are composed of the same geons at largely the same spatial relationships (as is a Boxer dog). Accurate discrimination may require more fine-grained analysis of structural and surface details (e.g. fur colour).

RBC theory does not account for the findings of Navon (1977) who described a 'global precedence' in object recognition. Navon presented participants with stimuli similar to the figures shown below (Figure 1). He found that the larger letters (in this case 'E') were reported faster than the smaller letters ('H'). He concluded that perceptual processes are temporally organized so that they proceed from global (holistic) structuring towards more and more fine-grained (featural) analysis.

```

HHHHHHHHHH
H
H
H
H
HHHHHHHH
H
H
H
H
HHHHHHHHHH

```

Figure 1: An example of a 'Navon' figure. In this example the global letter is 'E' and the local letter is 'H'.

Each of these perceptual routes (featural, configural and holistic) might be useful at different stages of object recognition, especially if as Navon suggested, processing proceeds from holistic to featural encoding. I will return to this shortly.

Categorisation

Categorisation theory suggests that once an object representation is formed, it is compared to pre-existing templates of categories. A decision is then made as to which category the object belongs to, and this decision process directs encoding into memory and information about the meaning of the object to the perceiver. On the basis of the comparison process, there must be a method of deciding which category an object belongs to. If the object representation is an exact match to a category representation, then this decision should be relatively straightforward: the object is a member of that category. But how does the system cope with novel object representations that it has not encountered before and so will not be an exact match to an existing representation? Also, how does the system cope with uniqueness, i.e. an object that cannot belong to more than one mutually exclusive category (if an object is a dog, it cannot also match the category 'cat')? Three decision rules have been proposed.

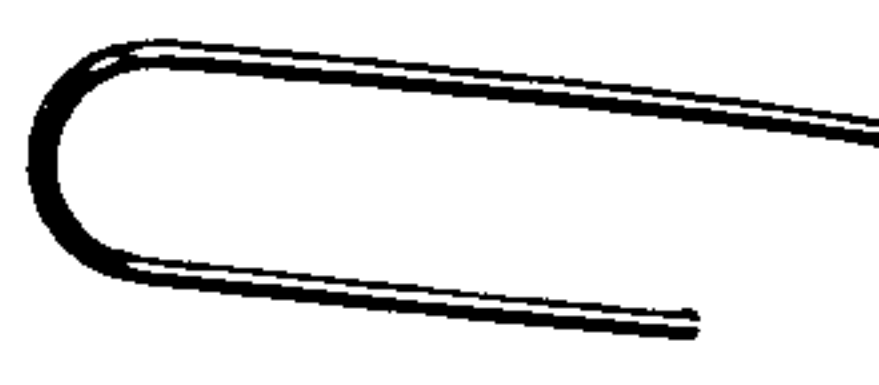
Threshold rules propose that object-category comparisons are assigned a critical value. If the current comparison exceeds that value, the object is assigned to the category. The object will be assigned to however many categories exceed the critical value. This process will recognise a novel object (as no object-category comparison will reach the critical value) and will create a new category to accommodate it. However, the rule does not recognise uniqueness as more than one category can exceed threshold.

Maximum (best-fit) rules choose whichever category exhibits the closest representation to the object. This process will recognise uniqueness, but will not cope with a novel item (it will assign it to the closest category).

Finally, the maximum-over-threshold rule combines the best of both of the above processes. It sets a threshold for the comparison value below which objects will be recognised as novel, and above which the object will be assigned to the category with the highest value.

Any rule of categorisation must be able to cope with the following: that an object does not necessarily belong to only one category, it may belong to many. For example, an object can be an animal, a dog, a collie, and Laird simultaneously. These categories are largely hierarchical (see Figure 2).

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.



Rosch and colleagues (Rosch, Mervis, Gray, Johnson & Boyes-Braem, 1976) proposed three levels of categorisation: superordinate (animal), basic-level (dog) and subordinate (collie). They proposed that most people first recognise objects at an intermediate level of categorisation, at the basic-level. Basic-level categories were defined in terms of three criteria: shape similarity, similar motor interactions and common attributes. Shape similarity was believed to be the most important of these; basic-level categories are the highest-level categories whose members have similar shapes. Superordinate level categories can be very dissimilar (a dog compared to a bird or a fish). As the levels move down the hierarchy, members become even more similar (a collie compared to another collie). In Figure 2 above, I have added a fourth level: individuals, or token exemplars. This is a category that only one member can belong to.

If shape is the most important criteria for categorisation, and most people first categorise objects by their basic-level, it follows that basic-level categorisation is achieved by holistic perceptual processing. As people begin to make more fine-grained categorisations (e.g. it is a collie, it is Laird) they must discriminate more fine-grained visual details of the object. Therefore subordinate and individual level categorisation may be achieved by configural and featural perceptual processing. This fits nicely with Navon's view of global precedence. The initial holistic processing would lead to basic-level categorisation, subsequent featural and configural processing would lead to subordinate, and perhaps individual, categorisation. Therefore, one might conclude that before you can recognise Laird, you must first recognise that it is a dog.

This notion of holistic processing leading to basic-level categorisation and featural processing leading to subordinate or individual categorisation, is supported by the findings of Collin and McMullen (2005). They presented participants with object images in three

spatial frequency conditions: low, high and unfiltered (see Figure 3). Participants were asked to verify subordinate, basic-level and superordinate names of object images that were presented at each spatial frequency. Low spatial frequencies are thought to emphasise global and configural object details, and high spatial frequencies are thought to emphasise featural object details (Fink et al., 1997). The unfiltered images acted as a control condition. Collin and McMullen found that more errors to name subordinate level objects were made in the low spatial frequency condition compared to the high spatial frequency and the unfiltered conditions. This suggested that the featural details of the objects were important in making subordinate level categorisations.

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.

Types, Tokens and Repetition Blindness

Progression from basic-level to subordinate and individual level encoding is mirrored in the findings of Repetition Blindness (RB) studies. RB is a phenomenon in which two identical items presented in a rapid serial visual presentation (RSVP) are encoded as being only one instance of the item (Kanwisher, 1987). Therefore, recall is typically of only one item occurring, not both. RB is usually cited as a failure of

tokenisation. Tokenization (or individuation) refers to the process by which an event is recognized as being a distinct occurrence, as opposed to the process of merely identifying an item (token versus type information).

Relevant to the current research, is the finding that this effect can also occur for items that are not completely identical. If the items share a commonality, they are also susceptible to RB (Bavelier, 1994, 1999; Bavelier & Potter, 1992; Buttle, Ball, Zhang, & Raymond, 2005; Kanwisher & Potter, 1990; Kanwisher, Yin, & Wojciulik, 1999; MacKay & Miller, 1994). For example the number '8' and the word 'ate' are visually distinct, but are pronounced in the same way. If their phonology is encoded before tokenization then they will be susceptible to RB. What is particularly noteworthy is that RB has been shown to occur for phonologically similar items (e.g. the words sun/son) and semantically similar items (e.g. pictures of airplane / helicopter; Bavelier & Potter, 1992; Kanwisher, Yin, & Wojciulik, 1999). Moreover, Buttle, Ball, Zhang, and Raymond (2005) demonstrated that different brands from the same product category can also produce semantic repetition blindness effects, while Bavelier (1994) has shown that mixed picture/word formats can produce RB (e.g. the word duck and a picture of a duck).

It could be argued that in the above cases, the two items were encoded as members of a category, but not as individuals. Thus, we could say that token (individual) encoding succeeds type (category) encoding.

In the remainder of this thesis I will sometimes use the terms 'type' to refer to categories (either superordinate, basic-level or subordinate) and 'token' to refer to individuals. Where possible, when discussing the effects found, I will speculate as to the category level at which effects may be occurring.

Attention

At any given time, our environment is bombarding our senses with information. For example, you may feel a stone in your shoe, at which point the telephone may ring, shortly before you smell your toast burning. The opportunities presented to us are vast, yet our capacity to process the information presented is limited (Broadbent, 1958; Kahneman, 1973). How do we decide whether to deal with the shoe, the telephone or the toast first? Selective attention is the cognitive mechanism that prioritises our percepts and actions, determines the content of our consciousness and ‘chooses’ what information proceeds to working memory. Selection involves two processes, the first is the facilitation of relevant information, and the second is the rejection, or inhibition of irrelevant information. Objects compete for our attention within and between all modalities. However, it is well beyond the scope of this thesis to review attentional mechanisms in all modalities and so, I concentrate my review on visual attention mechanisms and capacities.

Visual attention has been likened to a mental ‘spotlight’ (Posner, 1980) that ‘illuminates’ a spatial location: objects that fall within its ‘beam’ are enhanced and processed more efficiently than objects at unattended locations. Behaviourally, this results in more accurate responses and faster responding to objects at attended versus unattended locations. This ‘spotlight’ can also focus on an object, as opposed to a location. Duncan (1984) found that when participants were required to make a single judgement about two objects (at a single location), compared to two judgements about a single object, errors increased. The two objects occupied the same location, thus a purely location-based account of attention would predict no cost of attending these two objects. This demonstrated that attention can also be object-based. Allocating attention to an object or location can enhance the perceptual processes of whatever is in the ‘beam’. Spatial resolution is enhanced

(Yeshurun & Carrasco, 1998, 2000; Carrasco, Williams & Yeshurun, 2002) as is contrast sensitivity (Carrasco, Penpeci-Talgar & Eckstein, 2000; Lu & Doshier, 2000; Cameron, Tai & Carrasco, 2002; Carrasco, Ling & Read, 2004; Liu, Pestilli & Carrasco, 2005).

The biased-competition model (Desimone & Duncan, 1995) and similar accounts of visual attention (Cave, 1999) propose selection as an outcome of competitive interactions between visual inputs: when one input is favoured, another is lost. Object representations compete for control through reciprocal suppression. The relative strength of each object is determined by the bottom-up salience of the stimuli and the top-down attentional control settings of task context and goals. Selection occurs when competitive interactions resolve in favour of a single object.

Visual Search

Visual search paradigms neatly illustrate the problem of selection in multiple object arrays.

The standard task for participants in visual search experiments is to state whether or not a designated object (a target) is present in an array of objects (distractors) that differ from it. When the target and distractors differ on a single dimension such as colour or shape, this is referred to as feature search. Alternatively, when the target must be discriminated from the distractors on multiple dimensions such as colour and shape, this is referred to as conjunction search. While error rates to find the target are typically low (given unlimited exposure time), the reaction times (RTs, time taken to respond “yes” to a target-present trial or “no” to a target-absent trial) provide a measure of search. Key to this is the manipulation of set-size (the number of distractors in the array plus the target). The general observation is that with feature search RT is independent of set-size, leading to flat

search slopes if RT is plotted against set size. But with conjunction search RTs increase linearly with the number of distractor items in the array, leading to steep slopes (Treisman & Gelade, 1980).

A classical interpretation (but see Townsend, 1971, 1976, 1990, for a different view) is that a flat search slope reflects parallel search, whereby attention is distributed and all the items are analysed simultaneously. The target item is easily identified due to its bottom-up salience by virtue of being different to the distractors on a basic visual feature (e.g. colour). A steep slope in contrast reflects a serial search, where focal attention is directed towards each item which is analysed sequentially with respect to feature combinations. The search self-terminates once the target is located (Treisman & Gelade, 1980). However, the data do not support such a strong dichotomy and more integrated theories (such as The Guided Search Model; Wolfe, Cave & Franzel, 1989; Wolfe, 1994) are built around the idea of a continuum between the two processes. They emphasise excitatory mechanisms that capitalise on the results of a parallel, preattentive² stage that establishes basic features (Wolfe et al., 1989; Wolfe, 1994), stimulus saliency (e.g. Cave & Wolfe, 1990), or similarity among stimuli (e.g. Duncan & Humphreys, 1989) in a bottom-up fashion. This is

² Here I wish to clarify my views on attention versus awareness. While attention in my view can occur outside of consciousness, awareness refers to a stage where an object has reached an individual's consciousness. Many theorists propose the idea of a 'pre-attentive' and 'attentive' stage of processing. I am sceptical about whether an object can be processed pre-attentively and suggest that 'pre-awareness' would be more accurate. Attention can be described as distributed (across a wide location, many objects, or many features) or focal (directed towards fewer, perhaps one, locations, objects or features). A limited capacity model would suggest that processing resources are distributed across any objects that are within the realm of distributed attention, but that an object within the realm of focal attention will benefit from greater processing resources allocated to it. Perhaps pre-attentive and attentive processing is better thought of as distributed and focal attention.

then matched with top-down strategic influences to ensure that targets receive the priority for processing needed for successful recognition and appropriate response.

However, such theories alone fail to explain satisfactorily how each item is encoded and either accepted or rejected as either target or distractor. What is the fate of the distractor stimuli? Are they encoded and inhibited, or is this search amnesic? An inhibitory tag on items (and locations) that had already been searched and rejected as targets would lead to a more efficient search, avoiding re-inspection of previously searched items.

Some accounts of behavioural performance in search tasks explicitly posit that interference from distractors is reduced by the application of inhibition to locations containing stimuli with task-irrelevant features (Treisman, 1993; Treisman & Sato, 1990). A number of computational models of search also implement distractor inhibition (e.g., Cave, 1999; Deco & Zihl, 2001; Heinke & Humphreys, 2003; Humphreys & Müller, 1993; Mozer & Sitton, 1998; Tsotsos et al., 1995), and several behavioural RT studies of distractor interference have proposed inhibition to account for their results (e.g., Cave & Zimmerman, 1997; Caputo & Guerra, 1998; Cepeda, Cave, Bichot & Kim, 1998; Mounts, 2000a, 2000b; Cutzu & Tsotsos, 2003).

In addition, single-unit recordings from the monkey inferior temporal cortex during visual search tasks have shown that cells sensitive to distractors showed discharge inhibition about 200 ms after search displays were presented (Chelazzi, Miller, Duncan & Desimone, 1993). Similarly, single-unit recordings from monkey frontal eye field cells showed that target selection during a typically highly efficient pop-out search task also used distractor suppression (Bichot & Schall, 2002). Together, these studies provide support for the idea that distractor inhibition may operate during visual search.

Inhibition of Object Features

What level of object representation might inhibition be applied to? Literature describing the Stroop effect suggests that inhibition can be applied to object features.

Features of an object may not seem like a natural unit for selection: humans interact with whole objects and not isolated features of them. However, in different tasks constituent features of an object may differentially provide information, depending on their relevance to current goals, and it is important that we are able to focus on the important information and ignore information that is not helpful in guiding our actions. For example, when picking a shirt to wear in the morning, I may want to ignore the features that tell me it is white, has long sleeves and blue stripes, and instead attend to the features that tell me that it has a big stain down the front and that I should not choose this particular shirt.

Two or more features (e.g. stain, long sleeves) within a single object (e.g. shirt) may compete for an individual's attention. This has been formally studied using the Stroop paradigm (Stroop, 1935). In these tasks, observers are required to indicate the colour in which letter strings are presented. Observers typically are quicker to respond to the ink colour (e.g. green) when the word is either congruent (the word green) or neutral (a letter string e.g. XXXX), than when the word is incongruent (e.g. the word blue). In the case of incongruent trials, the irrelevant feature of the word meaning competes for attention with the relevant feature, the ink colour and must be inhibited in order to respond correctly.

Studies using other paradigms also support the ability to select and inhibit features as opposed to whole objects. Flanker interference tasks typically require participants to respond to a dimension of a centrally presented target while ignoring two distractors presented to either side of the target (the 'flankers'; Eriksen & Eriksen, 1974). It has been

shown that if the distractors do not match the target on the to-be-reported feature error rates are high and RTs are long, compared to if the distractors do not match the target on an irrelevant feature dimension (Cohen & Shoup, 1997; Maruff, Danckert, Camplin & Currie, 1999; Remington & Folk, 2001).

Inhibition of Locations

Spatial attention and its movement has been studied using the Posner paradigm (Posner, Nissen & Ogden, 1977). In this task, participants must press a button as soon as they detect a brief flash of light (the target). The light is presented to either the right or the left of fixation. Critically, a cue is presented prior to the target. The cue is either a cross, or a right-pointing or left-pointing arrow. A cross gives participants no information as to the location of the upcoming target. A right-pointing arrow informs participants that on 80% of the trials the target will appear to the right; a left-pointing arrow informs participants that on 80% of the trials the target will appear on the left. If the arrow correctly informs the participant of the location of the upcoming target, this is a valid trial. If it does not, this is an invalid trial. Figure 4 depicts the procedure of the attentional cuing paradigm.

RT to detect the target is fastest on valid trials, suggesting that participants shift their attention to the correct location and are ready to respond to the target. RT to detect the target is slowest on invalid trials. In this case, participants have incorrectly shifted their attention and must re-orient to the correct location of the target.

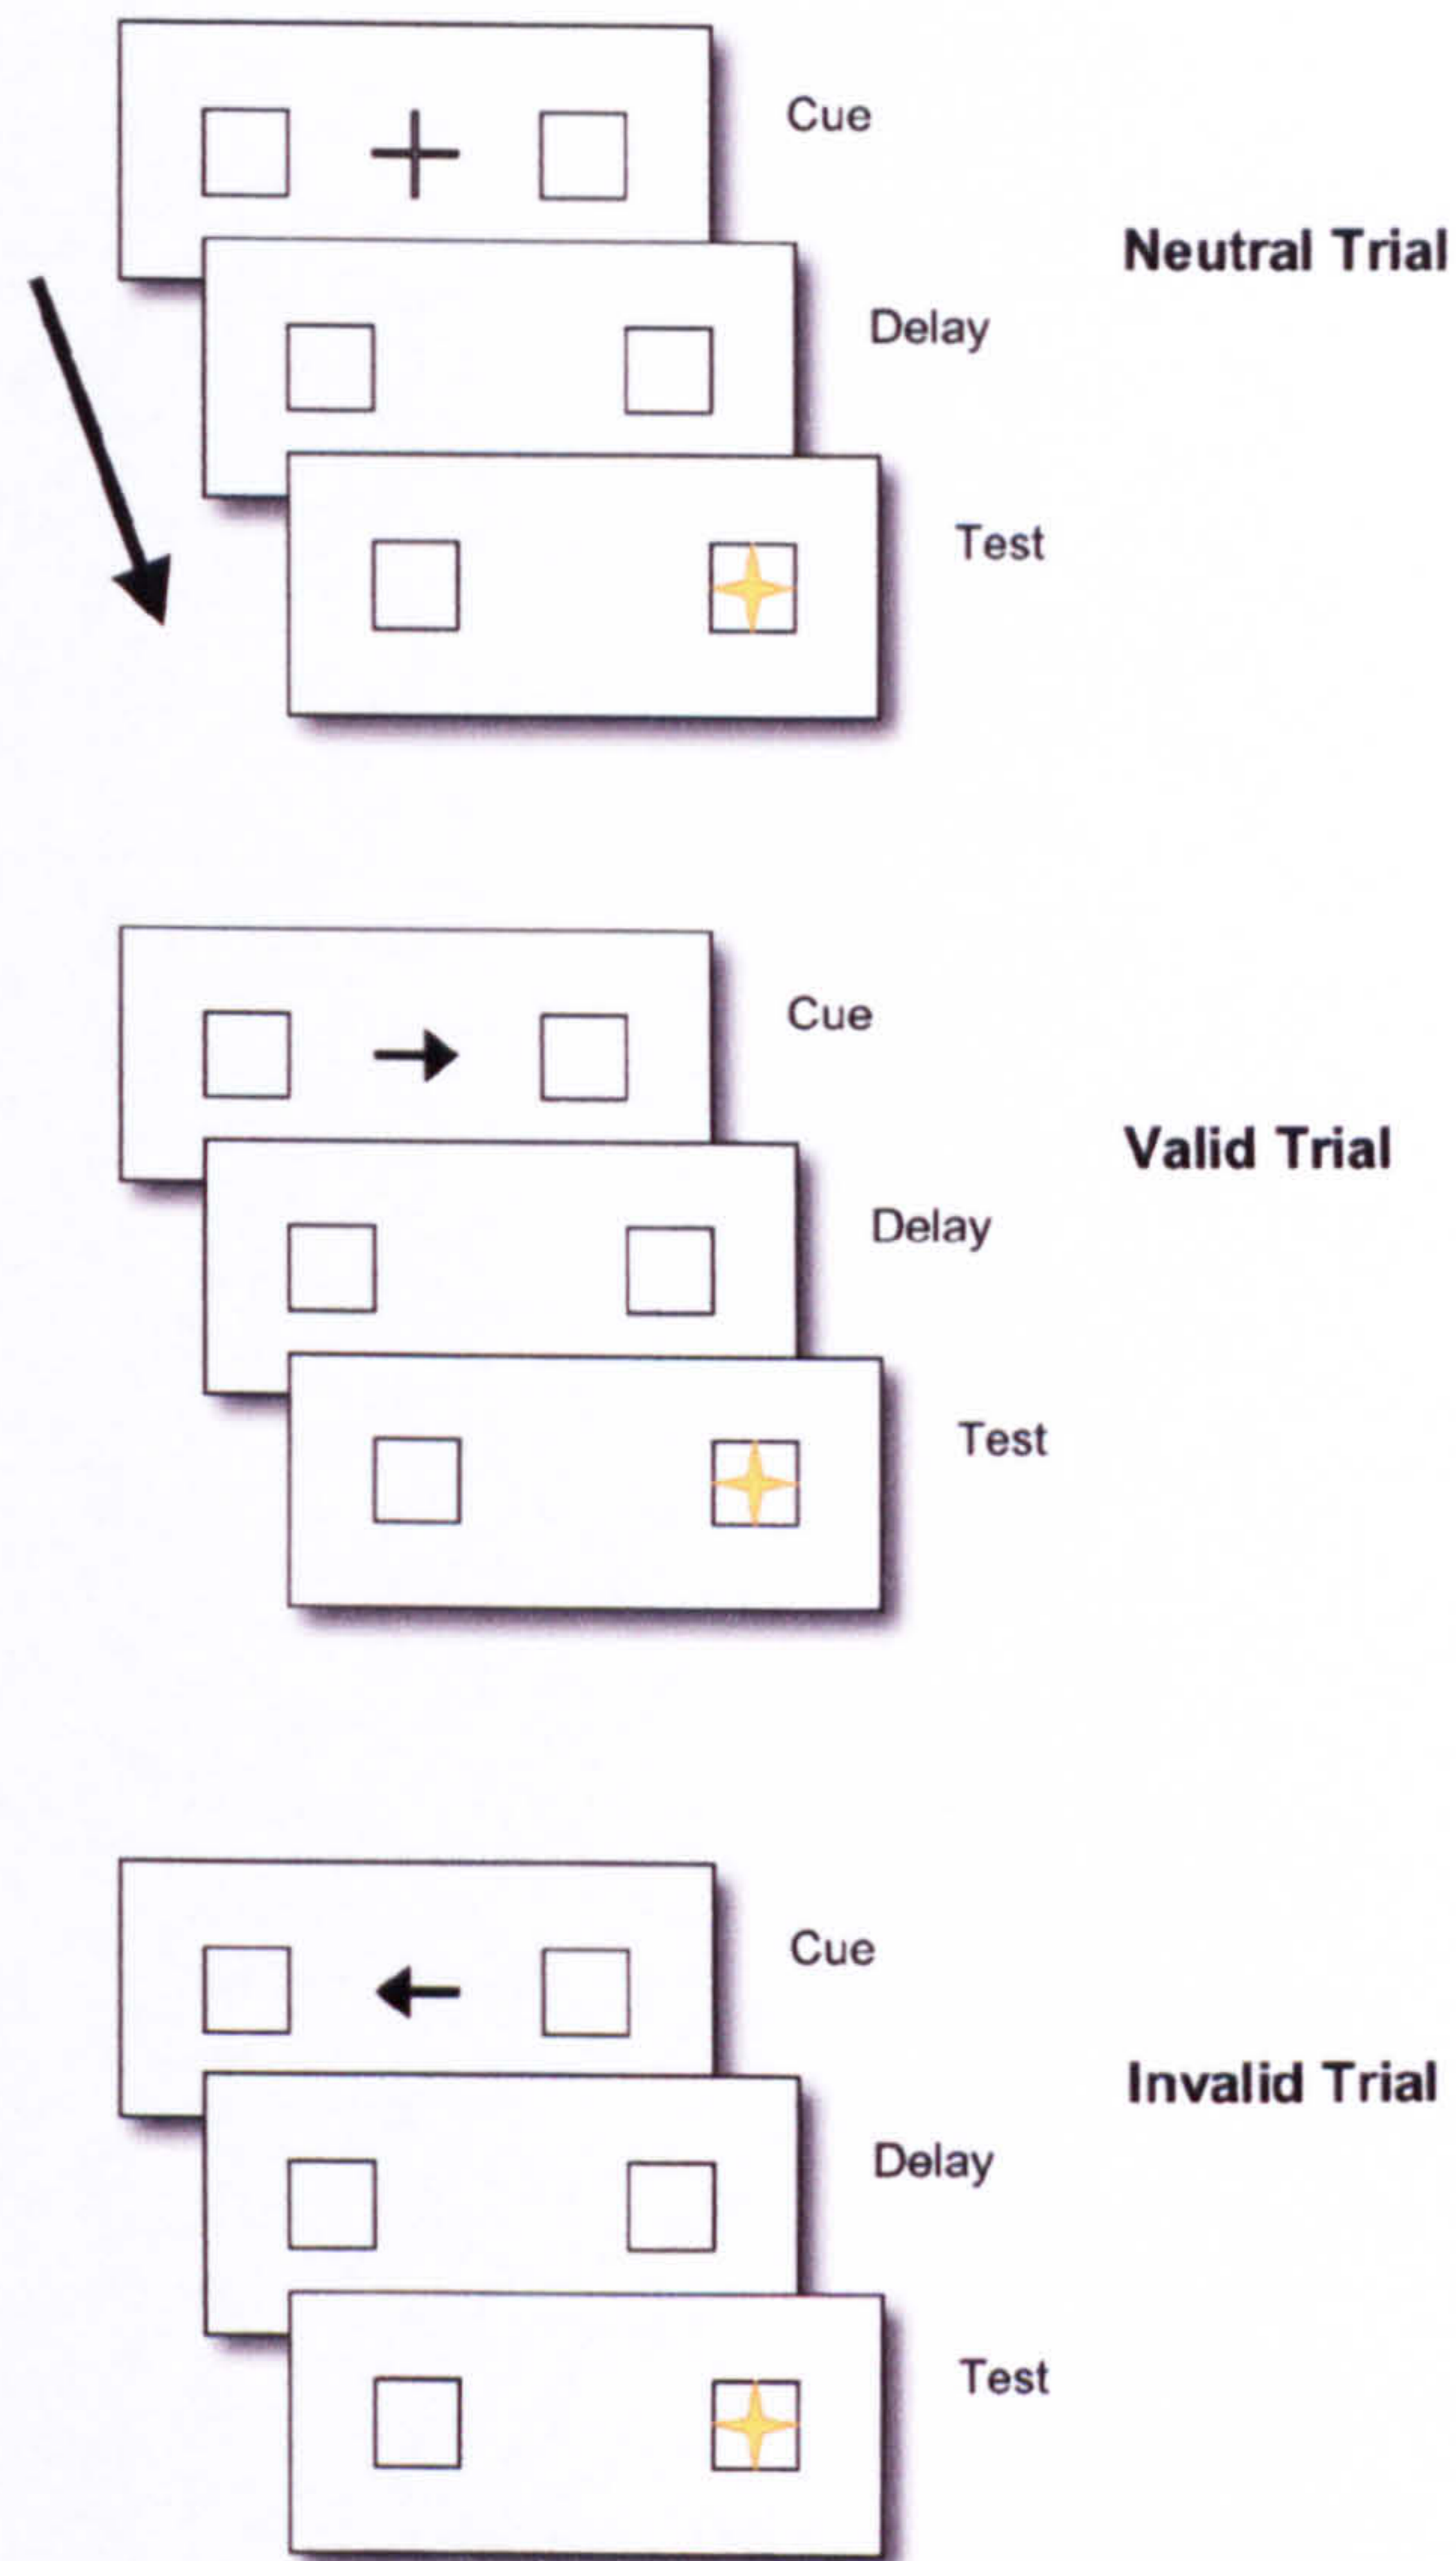


Figure 4: The attentional cuing paradigm. An arrow pointing left or right is used to cue participants to attend there for a target. On valid trials, the target appears in the cued location. On invalid trials, it occurs in the uncued location. On neutral trials, a cross appears at fixation.

Inhibition of Return

Inhibition of return studies developed the cuing paradigm above to investigate the role of inhibition in spatial attention. The critical difference is that in contrast the endogenous cue used in the cuing task (which allows participants to shift voluntarily their attention), these paradigms present an exogenous cue (a brief peripheral light flash) that automatically captures attention (see Figure 5). Following the exogenous cue the detection

of a target at the same location is either unchanged or enhanced relative to a no-cue baseline condition if the temporal interval between the cue and the target is about 300 ms or less. If this interval is more than 300 ms on the other hand, target detection is impaired (Posner & Cohen, 1984). This slowing of responding has been coined the ‘inhibition of return’ (IOR).




Figure 5: The three boxes in each display represent the potential locations of the target. The exogenous cue is represented by the darker outlined boxes.

IOR has been shown to occur at the location where a previously detected target had been presented and so is not due to response inhibition of the original cue, which must itself not be responded to (Maylor & Hockey, 1985). Further, IOR does not appear to rely on covert attentional orienting as IOR has been demonstrated at equal magnitude to each of two (and up to six) separate simultaneously cued locations, compared to a single cued location (Posner & Cohen, 1984; Snyder & Kingstone, 2001).

Klein (1988) reasoned that IOR reflects the inhibition of previously attended items and can serve as a ‘foraging facilitator’ useful as a within-trial memory mechanism in visual

search to promote the selection of new (previously unsampled) stimuli. To test this, participants were required to perform a conjunction search task and, following each display, had to detect the presence of a small probe target that could appear at either a previously occupied or unoccupied location. The detection of the probe at the previously occupied location was slower than when the probe appeared at the previously unoccupied location. In contrast, participants who performed a single-feature search task in which the RT-display size function was relatively flat demonstrated no such effect, indicating that attention had not moved serially to each item. It was argued that, in the serial conjunction search task, an attentional mechanism (inhibition) had acted to lower the probability that previously examined locations were re-examined and so had increased the efficiency of the search. It should be noted though that Wolfe and Pokorny failed to replicate this effect (Wolfe & Pokorny, 1990), although it has since been successfully replicated (Müller & von Mühlénen, 2000; Takeda & Yagi, 2000),

Inhibition of Objects

In addition to features and locations, evidence has also been found for an object-centred IOR. Tipper and colleagues (Tipper, Driver & Weaver, 1991) demonstrated that RTs to detect a target were delayed if the target appeared on an object that was presented at the location of a peripheral cue, and that then moved location. The inhibition did not remain at the location where the stimulus was cued originally, but instead followed the stimulus as it moved. They suggested that inhibition could be associated with an object-based representation and so did not operate exclusively on spatial locations. This effect was replicated and extended by Tipper, Weaver, Jerreat and Burak (1994) who demonstrated that both environmental and object-based inhibitory mechanisms in IOR can coexist. Other studies have found greater IOR when an object, relative to an empty location, was cued

(Jordan & Tipper, 1998; Leek, Reppa & Tipper, 2003; Reppa & Leek, 2006) and that this effect lasted longer (Paul & Tipper, 2003).

Inhibition of Object Identity

Inhibition can also be associated with higher-level representations, such as the object's identity. To investigate whether IOR could be exhibited for object identity, Grison and colleagues devised a cuing task that presented upright or inverted faces as the cue and/or target (Grison, Paul, Kessler & Tipper, 2005). IOR was greater when both cue and target faces were upright than when cue and/or target faces were inverted. Because physically there was no difference between the four orientations with respect to object or location representations, the results are explained by inhibition of object identity. Upright faces were more easily recognised in the cue sequence and inhibition was applied to that information. Likewise, in the target sequence an upright face made recognition easy, allowing prior inhibition to impact processing of the target. Large IOR effects in the upright-upright condition were thus elicited because inhibition had been applied to several representations: location, object file and object identity. In the remaining three conditions however, inverted faces were harder to recognise, and inhibition did not affect identity processing. So, IOR effects were reduced because inhibition was applied to fewer representations: location and object file.

Retrieval of Attention Processes from Memory

If, as the devaluation-by-inhibition theory of distractor devaluation suggests, prior inhibition towards an object results in its subsequent emotional devaluation, a system must exist that remembers the object was initially inhibited. An inhibitory tag must be stored with the object's representation in memory.

A link between attentional and memorial processes has been established. As the amount of time that a picture is viewed in a RSVP stream increases, recognition memory improves (Potter & Levy, 1969). Thus, the more time dedicated to processing, and presumably attending, the object, the better the memory for that object. Beyond this, it has been shown that selectively attending to a picture improved memory for that picture, even when presentation duration was shorter than for other pictures (Intraub, 1984). Extending this finding to a visual search environment, Williams and colleagues (Williams, Henderson & Zacks, 2005) found that even though participants were not instructed to memorise items during search, targets were remembered better than distractors. Further, distractors that shared features in common with the target were remembered better than distractors that did not share features with the target. This supports the notion that during feature search, distractors can be inhibited en masse in parallel (and are therefore poorly encoded into memory), but that during conjunction search, distractors are searched and rejected serially, allowing encoding into memory.

But note that while increased attention and looking-time have been found to have enhancing effects on memory, the effect of inhibition on memory has not been established: eye movement data collected in this study revealed that targets were viewed for a longer period of time than related distractors, so it cannot be said that inhibition per se was the cause of poorer memory.

Long-Term IOR

While attention has been shown to improve memory, it is also now becoming clear that humans can retain information regarding attentional processes in memory. Such encoding would be beneficial in interrupted search tasks for example. If a search episode

had to be paused, before the target was identified, applying 'tags' to already searched-through items (as suggested in IOR) that can be reactivated would prevent the entire search from having to be reinitiated. Items that had already been rejected as distractors need not be identified again. Such a mechanism would require inhibitory tags to be encoded into memory, with the object's representation, in such a way that could be re-instantiated on subsequent presentation.

The governing view of attentional mechanisms such as the inhibition found in IOR has been that they are transitory in nature. It has been suggested that the inhibition may exist for approximately three seconds before it decays (Samuel & Kat, 2003). This view has been challenged though, and it has recently been shown that specific states in the attention network can be encoded into memory, to be later retrieved and used, enabling past experiences to guide current behaviours.

In an exogenously-cued IOR sequence, IOR was found to exist up to thirteen minutes, and hundreds of trials after a face was initially encountered (Tipper, Grison, & Kessler, 2003). A face was presented to the left and right of fixation. Participants were instructed to locate as fast as possible a green stimulus when it flashed on one of the faces. The cue occurred only on a few trials, in which a red stimulus flashed onto one of the faces. Participants were told to ignore a red flash. The red flash oriented attention to the face, and when withdrawn, activated IOR to that face. Even though face processing was incidental (the face itself was task-irrelevant), regular IOR effects of the face in the order of 2 seconds after cue presentation were found. Tipper et al went on to demonstrate its persistence. In one study, faces were not seen again for about three minutes, with 40 trials intervening these exposures. In a second study, faces were not reencountered until about 13 minutes and 100 intervening trials later. IOR effects for cued faces were observed in both instances.

These effects provide evidence that inhibition can be associated with the identity of the object that attention was drawn to by a sudden-onset, task irrelevant cue. Furthermore, when the object was subsequently encountered, part of the object-recognition process retrieved the prior history of cuing, and thus the inhibition.

This inhibition retrieval appeared to be an unconscious process. When asked to consciously recall which face in a pair had been cued, participants performed worse than chance (Kessler & Tipper, 2004). While apparently counter-intuitive, this can be explained by a response-bias away from the inhibited cued face.

Negative Priming

The mechanisms involved in negative priming and IOR are thought, by many, to be similar. Negative priming studies demonstrate another instance of prior inhibition impacting subsequent tasks.

Negative priming involves inhibitory carryover effects from one trial to another. A distractor presented in one trial (the prime trial) becomes the target on the next (the probe trial). The typical finding is that processing of this target is delayed relative to targets that are unrelated to previous trials (Tipper, 1985). Houghton and Tipper (1994) proposed an inhibitory theory of negative priming and suggested that in order to attend to a target, the distractor must be inhibited. Stimuli that match a target template receive excitatory feedback, and those that do not receive inhibitory feedback (in the case of the distractor). When a previously ignored distractor appears as a target on a subsequent trial, the previous inhibition in its representation causes responses to it to be slower and less accurate than new items.

Negative priming effects demonstrate that, not only does the attentional state of the prime trial affect responding to an immediate probe trial, but that it can also influence a probe trial that occurs several seconds and several trials later (Tipper, Weaver, Cameron, Brehaut & Bastedo, 1991). It is possible that a prime trial will impact a probe trial even 30 days later (DeSchepper & Treisman, 1996). However, the notion that inhibitory mechanisms are at the heart of negative priming is not universally accepted (e.g. Neill, Valdes, Terry & Gorfein, 1992; Milliken, Joordens, Merikle & Seiffert, 1998).

Episodic retrieval theories of negative priming (Neill et al., 1992) conversely hold that a distractor will be processed with a 'do not respond' tag. When a distractor is later presented as a target, it is the 'do not respond' tag that slows reaction times. While inhibition theories emphasise processing involved in encoding of the prime trial, episodic retrieval theories emphasise memory retrieval processes in the probe trial.

Thus far, I have discussed how an object might be perceived and categorised, and how once an object representation is formed, attentional (inhibitory) processes that were directed to the object might be stored in memory with it. But this is only one side of the distractor devaluation coin. Theories of attention attempt to describe how an object might gain priority for processing in order to activate a response. But, humans are not robots: we do not act on an object simply because we are presented with an opportunity to do so. We act on an object (or do not) because it benefits us to do so. Between processing an object and acting on it, we predict its value, and the reward to be gained. While it is true that in the majority of attentional studies correctly acting on objects will hold an element of reward for their participants, namely in implicit approval or achievement, what must be considered is the reward value of the object themselves and how this is evaluated. How might an inhibitory tag be used by an emotional evaluation system to lead to distractor devaluation?

And if generalisation of distractor devaluation is possible, then how might a second object ‘use’ an inhibited object’s tag? Can evaluations of one object affect the evaluation of another object?

I now turn your attention to the literature on emotional processing and evaluation, before returning to attention to describe how the two mechanisms might interact.

Emotion

Emotions are the states of mind that we know best and remember with the most clarity. But they are the states of mind that, it can be argued, we know least about. Their timeframe can be subject to huge variations; they can be caused by situations, stimuli, or memories that are obvious or completely unknown. We can be surprised by our emotions, suggesting we do not have total insight into them. They may occur spontaneously, suggesting we have limited control over them. What causes unconscious processes to give rise to conscious experiences of emotions? What drives our preference for one object compared to another? How may an object acquire an affective dimension? A plethora of literature exists making attempts to describe the nature and mechanisms of emotion. Indeed emotion research has enjoyed somewhat of a renaissance in recent times. However it is not my intention to exhaustively review current theories of emotion. In fact I think at this point it would be prudent to outline what I will not review here.

I wish to draw a distinction between emotions as ‘feelings’, i.e. experienced body states and emotions as ‘affects’ i.e. a like or dislike of an object or event. While the early theories of emotion (discussed briefly below) may have been prompted by the former, I am interested in the latter and its implications.

Emotions can be described as social communicators: a facial expression can convey information in a very short period of time. However, I will only discuss facial expressions in terms of stimuli to their observer. Emotional feelings can also be thought of as outcomes of an action, as rewards or punishers. While this is highly relevant to the mechanisms involved in decision making and choice (the reward of choosing a great lipstick is the feeling of happiness), I will not be discussing emotions in terms of outcomes.

The aspects of emotion I will instead focus on are emotional evaluation of a stimulus (object or event) and its impact on behaviours and on subsequent processing of the stimuli (particularly attention). Theories of appraisal and motivation are of particular relevance here. I will then proceed to discuss how emotional evaluations might be altered by prior perceptual experience, with particular regard for the role of attention.

Some Early Theories of Emotion (as 'Feelings')

James (1884) consolidated the first major theory of emotion in his attempt to identify the processes that intervene the occurrence of a stimulus and the feeling that is produced in reaction. The 'Feedback Theory' suggested that emotions feel different to other states of mind due to the accompanying bodily responses found in emotional reaction. These reactions lead to internal sensation, the difference between which, and the accompanying bodily response, cause the different feelings of emotions. The physiological reactions to the stimulus are returned to the brain and are interpreted as different emotions. This account would lead to the proposition that we are sad, because we cry.

Cannon (1915) conceived of emotions as "emergency reactions". According to 'Fight or Flight Theory', the flow of blood is redistributed to body areas involved in emergency situations so that energy will reach the critical muscles and organs necessary for

reacting in a given emergency situation. This was believed to be mediated by the autonomic nervous system (ANS), a network of neural cells and fibres that control the activity of the internal organs and glands in response to brain signals. A result of activation of the sympathetic division of the ANS was the characteristic bodily signs of emotional arousal (increase in heart rate, increased perspiration on the palms). It was thought that the sympathetic division of the ANS, regardless of the way in which it was activated, or the reason for the activation, reacted in a uniform way. This would lead to identical bodily reactions to all emotional states, thus all states would have the same ANS signature and so the James Feedback Theory could not be correct. In addition, Cannon found that the ANS responses were too slow to fit into James' theory as conscious emotional reactions to stimuli had already occurred. Cannon proposed a role though for the interpretation of bodily responses in giving emotional reactions their sense of intensity. James and Cannon do appear to agree however, in the proposition that emotional states feel different to other conscious states due to bodily responses.

Schachter and Singer (Schachter & Singer, 1962) assumed that physiological responses in emotion inform the brain that a state of heightened arousal is occurring. The 'Cognitive Arousal Theory (CAT)' of emotion claimed that the type of emotional experience that a person will feel depends on the interpretation of social cues available in the situation, assuming that the reason for the arousal was not immediately obvious. A knowledge of what type of emotions are usually experienced in the current situation also contributes to identification of emotional state by the person. Thus, cognitive interpretations determine emotional feelings.

This was, in part, reinforced by the work of Vallins and Ray (Vallins & Ray, 1967) who gave subjects inaccurate information about how their body was responding to a

situation. It was found that in order for physiological activity to contribute to an emotional experience, the activity has to be cognitively represented; it is the cognitive representation of the arousal and not the arousal itself that generates feeling through the interactions with thoughts regarding the situation. So, “feelings” could be thought of as cognitive interpretations of situations.

The above theories attempt to describe how an emotional reaction might be interpreted once it has occurred, but they are unsatisfactory because they do not explain how these emotional states arise. In order for an emotion to occur, a stimulus must be recognised (not necessarily consciously) as warranting an emotional response. More recent theories of emotional experience suggest that instead of a peripheral (bodily) control of emotional experience, we should view emotions as ‘action dispositions’ that occur once a stimulus has been interpreted (Frijda, 1986), and conscious experience of this action readiness is what constitutes ‘emotion’.


Emotion researchers are currently involved in a debate regarding the ‘building blocks’ of emotion. The ‘basic emotion’ approach argues that certain categories of emotion, Fear, Rage, Lust, Separation Distress, Play/Social Affection, Nurturance and Disgust, are biologically basic and inherited reflex-like modules that cause a distinct behavioural and physiological pattern (Ekman, 1972; Panksepp, 1998, although note that Panksepp does not include ‘disgust’). The ‘basic emotion’ approach has fallen out of favour though as neurological studies provide increasingly inconsistent data in the form of patterns of activation during these emotional states.

The ‘dimensional’ approach to emotion is more popular and coincides with the notion of emotions as action dispositions. This approach argues that anger, sadness, fear etc.

are categories that characterise elaborate responses derived from more fundamental biological properties such as preservative and protective (Konorski, 1967), positive and negative (Watson & Tellegen, 1985), valence (pleasure / displeasure) and arousal (high activation / low activation; Russell & Barrett, 1999), or appetitive and aversive (Dickinson & Dearing, 1979; Lang, Bradley & Cuthbert, 1990). The appetitive and aversive dimension, known under the umbrella of 'motivation' has received particular attention. This dimension combines the separate properties of valence and arousal, and corresponds to prototypical action dispositions of 'approach' and 'withdraw'.

Under the motivational approach, emotions can be defined by where they fall in two-dimensional space, where one dimension is pleasure and the other dimension is arousal. A highly arousing, highly pleasurable stimulus will result in 'joy', an unarousing, unpleasurable stimulus will result in 'sadness'. Figure 6 illustrates the location of emotions in this two-dimensional space model.

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.



Lang (1995) proposes that the primary dimension is the valence dimension, and emotions will correspond to stimuli valence. Arousal reflects variations in the activation. From an evolutionary perspective, a simple valence-response system would be an advantage, promoting fast responses. Categorising emotions along the motivation dimension then appears sensible; indeed, appetitive stimuli are believed to promote 'approach' actions, whereas aversive stimuli promote 'avoid' actions.

I have briefly discussed above what might constitute an emotional experience. But how do these emotional experiences arise? How is a stimulus evaluated in order to determine whether it warrants an emotional response or an action disposition? I now turn to the literature regarding emotional evaluation, specifically the theory of Appraisal.

Emotional Evaluation

Arnold (Arnold, 1960) proposed two stages of emotional evaluation in her Appraisal theory. The first, intuitive appraisal, is an unconscious process that produces an affect towards a stimulus (a like or a dislike). The second, reflective appraisal, is a conscious process that evaluates the stimulus in terms of the self (how does this relate to me and my current goals or concerns) and produces an emotional reaction (a feeling). Note that the two stages have been given different names by different authors. 'Implicit' or 'intrinsic' appraisals refer to 'intuitive'; 'explicit' or 'motivational' appraisals are 'reflective'.

Much debate has occurred in the literature, often interpreted as an attack on Appraisal theory. I argue that much of this has occurred due to poorly defined terms. Arnold did not use the terms 'intuitive' and 'reflective' (these were coined in a review by Kappas, 2006) and instead grouped them under the banner of 'appraisal'. While Lazarus (Lazarus & Alfert, 1964; Lazarus, 1982, 1984, 1999) acknowledged the existence of the intuitive stage,

his work focussed on and highlighted the reflective stage in which 'cognition' is key (cognitive primacy). Zajonc (1980) then, in his 'Preferences Need No Inferences' paper misinterpreted appraisal theory as a purely 'cognitive' process and argued for an automatic, unconscious evaluation stage (affective primacy), neglecting to recognise that this is what Arnold had originally argued for (see Kappas, 2006 for more on this debate).

Despite the divergence of naming conventions, the distinction between the two levels has been largely supported (e.g. Scherer, 1984; Frijda, 1986; Ortony, Clore & Collins, 1988; Ellsworth, 1991).

Intuitive Appraisal

An intuitive appraisal (according to Arnold) is independent of the momentary goal state of the organism and will result in an affect (not an emotion). It can be conceptualised as stored valence: a valence tag stored in memory with the object's representation that may be retrieved upon encountering the object.

The notion of intuitive appraisal is, in my opinion, supported by the work of Zajonc (1980). He advocated an affective primacy hypothesis predicated by the notions that evaluations can be automatic (Bargh, 1997). People 'feel' without rationalisation (Haidt, 2001) and organisms appear to be evolutionarily prepared to respond quickly to threats (Öhman, 1997, discussed later). The affective primacy hypothesis posits that emotion is independent from cognition.

The Mere Exposure (Zajonc, 1980) effect was proposed to champion affective primacy. After repeated exposures to a stimulus, participants prefer the repeated stimulus compared to a novel stimulus, even if the stimulus was presented outside of awareness.

Thus the preference was not dependent on conscious identification and categorisation, and Zajonc presented this as evidence of cognition-free emotional processing. I will discuss Mere Exposure findings in more detail later.

Also in support of a cognition-free account, LeDoux (1996) proposed a sub-cortical 'low-route' to emotional evaluation that allows stimulus processing without cortical influence. Specifically, he proposed that the visual cortex (required for object identification and categorisation) was not necessary for emotional evaluation. In rat studies, when the visual cortex is lesioned, subsequent learning that a light would be paired with a shock was unimpaired. This led him to propose that such learning was mediated by the sensory thalamus, which projects directly to the amygdala, which in turn has connections to the motor cortex allowing for rapid responses to stimuli without conscious identification. While the existence of this route in humans, and the validity of using learning of light detection in rats as analogy for complex human behaviour have been challenged (Dolan, 2000), studies of the amygdala provide useful data on the processes of emotional evaluation. Although a neurological argument for emotion without cognition is flawed (indeed the anatomical areas necessary for cognition and emotion are highly interconnected; Halgren, 1992; Ghashghaei & Barbas, 2002), these studies suggest that emotional evaluation can occur without conscious awareness.

The neuropsychological case of patient GY supports this. GY displayed blindsight due to damage of the left occipital cortex that produced a right hemianopia. Despite being unable to 'see' a face presented in the right visual field (the faces were outside of the patient's conscious awareness), GY was able to guess the emotional expression of the face at above-chance levels (de Gelder, Vroomen, Poutois & Weiskrantz, 1999). In a follow-up study, fMRI revealed that GY's amygdala responded preferentially to fearful faces, even

when these faces were presented in the right visual field (Morris, de Gelder, Weiskrantz & Dolan, 2001).

Thus, it appears that there is a neuroscientific basis for intuitive appraisal, and conscious awareness need not be involved in emotional evaluation. Behaviourally, affective priming studies are perhaps the most compelling evidence for intuitive appraisal and its impact on evaluations of subsequent stimuli.

Affective Priming

Affective priming is a variant of classic priming research. Initially developed using adjectives as stimuli, Fazio and colleagues (Fazio, Sanbonmatsu, Powell & Kardes, 1986) presented a probe word to participants and asked them to evaluate the adjective as either “good” or “bad”. Prior to probe presentation, an affectively congruent or incongruent prime word was presented (within awareness levels). The latency of probe evaluation was found to be short for congruent pairs and long for incongruent pairs (the word “cockroach” presented as a prime will lead to fast evaluation of “disgusting” as a probe word but slow evaluation of “appealing” as a probe word). It was concluded that the participants’ attitude toward the prime word had been automatically activated, and interfered with probe evaluation, even though the prime word was irrelevant to the participants’ task. Presentation of an attitude object as a prime appeared to activate associated evaluations and facilitate judgements of an affectively related probe stimulus.

Affective priming was only found at short stimulus onset asynchronies (SOA) between prime and probe (about 300 ms, also Hermans, De Houwer & Eelen, 1994; De Houwer, Hermans & Eelen, 1998; Hermans, Spruyt, & Eelen, 2003). At longer SOAs, the effect vanished, as (presumably) participants were able to regain control over their

'automatic' evaluation tendencies. Importantly, the effect has also been found when the prime words are presented subliminally (Greenwald, Klinger & Liu, 1989; Greenwald, Draine & Abrams, 1996; Wittenbrink, Judd & Park, 1997), supporting the idea of non-conscious intuitive appraisal processes.

The generality of this effect has been shown to extend to non-word primes whose 'valence' was just previously learned by the participants (De Houwer, et al., 1998), black and white line drawings of objects (Giner-Sorolla, Garcia & Bargh, 1999), and full colour object images as primes (Fazio, 1993; Fazio, Jackson, Dunton & Williams, 1995; Hermans et al., 1994). The effect has also been extended to a modification of the probe stimulus. Instead of presenting an adjective, evaluation of valenced nouns and objects as probes has also been shown to result in an affective priming effect (Greenwald et al., 1989, 1996; Hermans et al., 1994, 1998).

For the purposes of supporting intuitive appraisal, evaluative affective priming effects (particularly when found to occur for primes presented below the threshold of consciousness) are successful. Interesting though, are the possible mechanisms underlying the competition and facilitation effects. Several researchers have pointed out the possibility that the competition and facilitation effects occur at the level of response selection, much like in a Stroop task. The prime may lead to a response readiness: response to a subsequent congruent probe will be facilitated as the response pathway is already activated, but response to a subsequent incongruent probe will be slowed as the activated pathway has to be inhibited and reinitiated. However, a modification of the evaluative affective priming task, the naming affective priming task, speaks against this explanation of results.

In the naming task, participants simply read the probe word aloud as quickly as possible. Although results have been mixed, in some cases it has been found that affectively congruent primes can facilitate naming of probe words, and affectively incongruent primes slow the naming of probe words (Hermans et al., 1994; Bargh, Chaiken, Raymond & Hymes, 1996). Such results point to a 'spreading of activation hypothesis' in which the evaluation of the prime object is activated in memory and as a result, less additional activation is required by a congruent probe.

The mixed results of the naming task are puzzling. As described above, some researchers have found effects that mirror those found in the evaluation task, whereas others have either failed to find an effect (De Houwer et al., 1998; Klauer, 1998, Klauer & Musch, 2001; De Houwer & Hermans, 1999), or have found a reverse priming effect (Glaser & Banaji, 1999). Such results may be due to differential treatment of the prime word by the participants. If participants are instructed to actively ignore (inhibit) the prime, it may be that the associated evaluations and response pathways of the prime are also inhibited, which would lead to a facilitation of the incongruent probe. However, the papers rarely report the explicit instructions given to the participants regarding the prime, and so the role of attention is difficult to interpret. A failure to find the effect at all may be due to a mixed attempt by the participants to ignore the prime, or it may be that the primes were not polarised sufficiently on an evaluative scale. Again, this is difficult to interpret without specific experimental manipulations that speak to these possibilities. However it leads us to an important question: Is attention necessary for emotional evaluation?

Is Attention Necessary for Emotional Evaluation?

Although the studies reviewed above support the notion that intuitive appraisal can occur very quickly, and without conscious intention or awareness, these studies do not address whether such evaluation can occur without some level of attention devoted to the object under evaluation. Studies that do address this issue provide mixed answers: some suggest that emotional evaluation can occur pre-attentively; others suggest that some measure of attentional resources are required.

Vuilleumier and colleagues (Vuilleumier, Armony, Driver & Dolan, 2001) manipulated spatial attention by having participants fixate on a central cue and compare either two faces or two houses presented eccentrically. On each trial, participants focussed their attention to whatever was presented to the left and right of fixation, a pair of faces or a pair of houses. Above and below fixation, a pair of pictures from the other category was presented and was to be ignored by the participants (on other trials the vertically presented pair was to be attended, and the horizontally presented pair was to be ignored). On each trial, the participants performed a matching task: they had to report whether the pictures in the attended pair were the same or not. Of the face pairs, some depicted neutral faces, and some depicted fearful faces. fMRI activation to the fearful pair was larger than the neutral pair on attend-house trials. This led the authors to conclude that the emotionality of the faces was being evaluated outside of attention.

In an object-based attention study (as opposed to spatial attention) Anderson et al. (Anderson, Christoff, Panitz, De Rosa & Gabrieli, 2003) used 'Double-Exposure' images in order to control for spatial attention. Images contained both a face (of different emotional expressions) and a building, which were semi-transparent and overlaid. Participants were

required to make a male or female judgement (attend to faces) or an outside or inside judgement (attend to buildings). fMRI revealed that amygdala activation was similar for both attended and unattended fearful or neutral faces, suggesting that attention was not required for emotional evaluation.

However in a study that employed a very similar paradigm to Vuilleumier et al., but that investigated the effects using ERPs rather than fMRI, Holmes, Eimer and colleagues (Holmes, Vuilleumier, & Eimer, 2003; Eimer, Holmes, & McGlone, 2003) found results that challenge this conclusion. Participants were required to match two faces or two houses that were presented to the left and right or above and below fixation. When faces were attended, both early and late ERP components were modulated by facial expression (fearful versus neutral). However, both early and late component differentials were eliminated when faces were unattended (during attend-houses trials), suggesting that attention is required for emotional evaluation (at least they did not find ERP evidence in favour of attention-free emotional evaluation).

Pessoa and colleagues (Pessoa, McKenna, Gutierrez & Ungerleider, 2002) also challenge the attention-free account. Participants were required to focus on faces of different gender (which could have either fearful or neutral expressions), or on bars of different orientation. All stimuli (faces and bars) were displayed at the same time. In some trials, participants were asked to judge whether a central face was male or female, and in other trials participants were asked to judge whether two peripheral bars had the same orientation or not. The authors attempted to make the bar-orientation task as difficult as possible in order that most, if not all, attentional resources were consumed leaving only small amounts, if any, available for processing of the unattended faces. During the gender task, fearful faces evoked greater neural activity than the neutral faces. This difference was

not found during the bar-orientation task, suggesting that for emotional evaluation to occur, the faces must be attended.

This conclusion was supported by Silvert and colleagues (Silvert et al., 2007). In their experiment, participants were simultaneously presented with four peripheral pictures. Two of the pictures were faces (either both neutral or both fearful) and the other two pictures were houses. The pictures were slightly tilted. The participants' task was to match two of the pictures (defined by their position) for either the orientation of the tilt, or for the identity. The identity task produced lower accuracy, longer RTs and poorer recognition performance than the orientation task. This, in addition to patterns of activation from fMRI data suggested that more attention was required to perform the identity task. In the orientation task (low attentional load) ignored fearful faces led to greater right amygdala activation than ignored neutral faces. This difference was not evident however when participants performed the identity task (high attentional load). The authors concluded that processing of emotional peripheral faces in the amygdala requires attention.

The role of attention in an evaluative priming task has been investigated by Musch and colleagues (Musch & Klauer, 2001) who manipulated the locational uncertainty of the probe word. Primes and probes were presented simultaneously in different locations. In the focused attention condition, a cue signalled the location of the probe; in the distributed attention condition, the cue was uninformative. Affective priming was eliminated in the focused attention condition, but found in the distributed attention condition, suggesting that in order to be evaluated and affect probe evaluation, a prime must have a small amount of attentional processing dedicated to it. Similarly, in an experiment in which the prime and probe words varied on two dimensions, colour and valence, trials in which participants had to report the probe's colour did not result in affective priming, but trials where the

participants reported probe valence did. The valence of the prime only interfered with probe evaluation when it was task-relevant, further suggesting a role for attention. Further to this, it has been found that even if prime valence is not task-relevant, but attention is directed towards valence as a dimension, priming effects are preserved (Spruyt, De Houwer, Hermans & Eelen, 2007). This suggests that the affective dimension of the prime need not be task-relevant to be activated, but will exert an effect as long as a small amount of attention is directed towards it.

Together, the studies presented above do not provide sufficient evidence for us to be able to conclude that emotional evaluation can occur without attention. In studies claiming to provide evidence of emotional affects without attention, it can be argued that a small amount of attentional resources were actually available for emotional processing. It appears that emotional evaluation will occur only if sufficient attentional resources allow. However, the emotional dimension of the stimulus need not be the subject of focal attention (i.e. task-relevant): it appears that emotional evaluation can occur if there are sufficient attentional resources 'left-over' from the participant's main task (i.e. evaluation may occur even if emotion is task-irrelevant).

Reflective Appraisal

In the above affective priming studies, reflective appraisal of the prime stimulus need not have occurred in order to impact evaluation of the probe stimulus. Moreover, the prime object is not goal-relevant yet it impacts the probe object. This demonstrates that conflicts during evaluation can occur. A system that is able to deal with these conflicts and evaluate goal-relevance is necessary in order to produce a coherent, sensible response. This is the role of reflective appraisal.

Reflective appraisal is not necessary for an emotion (in the form of affect) to occur, but it has a modulatory role. Firstly, it might function to evaluate stimuli in terms of the self. As Damasio puts it, “the current context may play a role and enhance or reduce the competence of the stimulus.” (Damasio, 2004, p. 51). Objects subjected to reflective appraisal are evaluated in terms of current goals and concerns, in relation to the self. Objects that enhance goal attainment are evaluated positively; objects that obstruct goal attainment lead to negative evaluations (Scherer, 1988). Similarly, reflective appraisal has been described as a comparison process between the encountered stimulus and a current desired state: a match will lead to positive evaluations, a mismatch to negative (Frijda, 1986). For example, chocolate cake might be desirable to a hungry person, but aversive to someone who has eaten too much.

Emotional stimuli may also differ in complexity. Although some stimuli may be simple to evaluate (the word ‘murder’ is negative), others may elicit both positive and negative reactions in intuitive appraisal (the word ‘abortion’ may have both positive and negative aspects). Leaving evaluation to intuitive appraisal alone would leave us ambivalent: we need to recruit a more flexible process. In such cases, the reflective appraisal system may be recruited in order to further process and resolve the conflict.

Variants of the affective priming procedure demonstrate the occurrence of reflective appraisal. Typically, the prime is an intrinsically neutral stimulus, but its goal-relevance is manipulated. It is therefore known as a motivational prime. Motivational primes can indicate reward or loss, or they can help or hinder goal achievement. For example, in a computer-based study participants engage in a ‘game’ wherein a blue stimulus will win them 10 points. If an intrinsically neutral (‘XXXX’) motivational prime word is presented in blue, it will facilitate evaluation of a subsequent probe word that is ‘good’. The same

motivational prime will slow the evaluation of a negative word as 'bad'. The 'gain' colour is randomly varied between trials, and so the authors argue that the colour does not receive an intuitive valence 'tag' as this would lead to not affective priming on the subsequent trial in which 'yellow' objects may signify gain (Moors & De Houwer, 2001).

In a slightly more elegant study, Rothermund and colleagues (Rothermund, Wentura & Bak, 2001) asked participants to report as quickly as possible a target letter that was printed in light grey. A distractor letter was also presented. The target or distractor could be one of four letters, each with an associated motivational valence. One letter indicated 'danger' meaning that, in their environmental context, points would be lost if the target was named too slowly. One letter indicated 'chance', meaning that points could be gained if the target was named within time. The other two letters were neutral, meaning that no points could be gained or lost. When danger or chance letters were presented as targets, RT to name them was faster than for neutral letters presented as targets, suggesting that the letters' motivational relevance had an impact (and thus reflective appraisal had occurred). Further, when the 'chance' or 'danger' letters were presented as distractors, they both slowed responding to the target, compared to neutral letters, again signalling the occurrence of reflective appraisal.

The Results of Emotional Evaluation

I have attempted to outline above how an emotional stimulus might be evaluated to produce an emotional response (either affect or feeling). But what is the point of an emotional response? While a feeling may be a reward or punisher in itself, it could also be viewed as a behavioural cue. I discuss now how an emotional reaction might: a) be involved in decision making; b) enhance perception; c) act as a motivational cue to efficiently guide

actions; d) enhance memory; and e) guide focal attention and awareness towards relevant stimuli.

The Role of Emotions in Decision Making

The act of decision making is more than a simple cost-benefit analysis. Humans are not equipped with unlimited time, knowledge or information-processing capacity.

Decisions, particularly those which involve high risks or rewards, are made not only upon a calculation of expected utility based on explicit knowledge of outcomes, but also depend critically upon emotions.

The 'somatic-marker hypothesis' suggests that emotions bias decisions towards choices that maximise reward and minimise punishment. Using the Iowa Gambling Task, Damasio (1994) found that in a group of patients (with lesions in the ventromedial prefrontal cortex, vmPFC) decisions were often made that were detrimental to their well-being. In the Iowa Gambling Task (Bechara, Damasio, Damasio & Anderson, 1994), participants are presented with four decks of cards. They are told that each time they choose a card they will win some money. Every so often, however, choosing a card causes them to lose some money. The goal of the game is to win as much money as possible. Every card drawn will earn the participant a reward (a large reward for decks A and B; a smaller reward for decks C and D). Occasionally, a card will also have a penalty (Decks A and B have a larger overall penalty than decks C and D). Thus, A and B are "bad decks", and C and D are "good decks", because Decks A or B will lead to net losses, and Decks C or D will lead to net gains. After about 40 or 50 selections, healthy participants are fairly good at sticking to the good decks. However, Damasio's patients appeared to be unable to learn from their mistakes and repeated decisions that led to negative consequences. These patients

also demonstrated flat affect, and an inability to react to emotional situations. This pattern of results could not be explained by impairments in comprehension or expression, attention or memory, which were found to be intact relative to controls.

Also using the Iowa Gambling Task, it was found that control subjects elicit large skin-conductance responses (SCRs) prior to making risky decisions. The SCR was absent in the patients, demonstrating an absence of anticipatory emotional responses. It was suggested that this region (vmPFC) was necessary for anticipating the emotional impact of future rewards and punishments (Bechara & Damasio, 2005).

The above account suggests that anticipating future emotion is helpful in guiding decisions. However a current emotion may be detrimental to decision-making. According to the 'affect-as-information' hypothesis (Schwarz & Clore, 1983) people use emotions as information, just as they use any other criterion. In doing so, they attempt to determine the informational value of their affective reactions to the judgment at hand. If they believe that their feelings are a sound basis for judgment, they use them in making their decisions. If they believe that these feelings are irrelevant, they exclude them from consideration. However, if people lack the ability or motivation to fully consider the issue at hand, affect may automatically be used (Petty, Schumann, Richman & Strathman, 1993; Clore, Schwarz & Conway, 1994). In situations where the affect is relevant, this automatic process will aid effective decision making. Thus, in situations in which the affective reaction is irrelevant, the automatic process will lead to an inappropriate attribution, and will hinder effective decision making. In such cases, negative moods promote elaborative processing and lead to over-estimations of risk, while positive moods promote automatic processing and may lower perceptions of risk (Isen & Means, 1983; Johnson & Tversky, 1983).

To demonstrate this effect, a telephone survey of life satisfaction was conducted on two days. On the first day, the weather was warm and sunny, on the second it was cold and rainy. Participants interviewed on the first day reported greater life satisfaction than participants interviewed on the second day. The interviewees' mood induced by the weather was misattributed to their judgements of life satisfaction. However, if the participants were first asked about the weather, the difference disappeared. This suggested that when participants are able to correctly attribute their mood to its actual stimulus, judgments made regarding a secondary stimulus are unaffected (Schwarz & Clore, 1983).

If reflective appraisal is indeed a match between a desired state and a current state (Frijda, 1986), a negative emotion indicates that one's current goals are not being met. Therefore, this may be a cue to engage in more elaborate processing, whereas a positive emotion indicates success and allows 'coasting'. In one such demonstration of the effect of mood on processing, participants were induced into either a positive or a negative mood. They then responded to a global-local perception task. The task involved a triangle made of small squares and a square made of small triangles. Participants were required to select which one matched an 'original' picture: a square made of squares or a triangle made of triangles. Participants in a sad mood selected the picture that matched the original in its small shapes: participants in a happy mood selected the picture that matched on its large shape. Participants in a sad mood adopted a more local approach than participants in a happy mood, who attended globally (Gasper & Clore, 2002). Consistent with the 'affect-as-information' hypothesis, when participants were directed to the cause of their emotions prior to the perceptual task, the difference between sad and happy participants disappeared.

It appears, then, that emotions can aid decision making in order to maximise rewards and minimise punishments, although emotions may potentially hinder successful decision

making if the cause of the emotion is misattributed. Emotion can also affect processing strategy: a negative emotion can be a cue to engage in more detailed processing, such as attending to the local details of a stimulus. But can emotion directly alter perception?

Emotions Enhance Perception

Emotion has been found to bias perceptual judgements. For example, mood can alter judgements of physical reality. When participants who were played sad music were asked about the incline of a hill they were expecting to climb, they tended to over-estimate the angle (Proffitt, 2006).

The emotional valence of stimuli can also alter perceptual judgements. Meier and colleagues (Meier, Robinson, & Clore, 2004) asked participants to evaluate words with a positive or negative meaning presented in a black or white font colour. Participants were faster and more accurate to evaluate positive words that were presented in white than black, and faster and more accurate to evaluate negative words that were presented in black than white. Here, they suggested that the font colour activated 'metaphoric associations' where "good guys wear white". The teams found similar metaphoric demonstrations for vertical position (finding that the "sunny side is up"; Meier & Robinson, 2004; Crawford, Margolies, Drake & Murphy, 2006) and stimulus size (where a "Big Mac is a Good Mac"; Meier, Robinson & Caven, 2008). In all of these studies, they demonstrated that the metaphoric congruency of the perceptual features aided stimulus valence evaluation.

In a further study, the same team demonstrated that perceptual judgement responses can be biased in favour of these metaphoric associations (Meier, Robinson, Crawford, & Ahlvers, 2007). When presented a uniformly grey box and asked to judge its brightness, participants judged the same box as lighter following a positive evaluation of an unrelated

stimulus, and as darker following a negative evaluation of an unrelated stimulus. In a second study, participants were presented positive and negative words that varied on font brightness. Subsequent to evaluation of the word, the participants were required to match the brightness of the font colour to one of five comparison standard boxes. Positive evaluations led to lighter matching responses and negative evaluations to darker matching responses (even though the words “light” or “dark” were not mentioned in this task). In one further study, the team removed the requirement to evaluate the words prior to perceptual judgement. The words were presented as primes, with a short SOA, in an affective priming-like task. Even though no explicit evaluation of the primes were required, they altered the evaluation of a subsequent perceptual probe that had to be evaluated as light or dark. Probes that followed positive primes were evaluated as lighter than probes that followed negative primes.

While it cannot actually be said that perception was biased by the emotion in these studies (as it may be that only the response has been biased), they hint at a link between emotion and perception, in that an emotional word can bias the report of a perceptual judgement. But can an emotional stimulus actually enhance perceptual processing, compared to a neutral stimulus? Phelps and her colleagues (Phelps, Ling & Carrasco, 2006) believe that it can. Participants were presented with upright and inverted faces with fearful or neutral expressions. The face was presented at fixation for 75 ms. Participants then had to indicate if one of four subsequently presented Gabor patches was tilted to the left or the right; Gabor patches were displayed for 40 ms and were of contrasts between 2% and 20%. An interval of 165 ms between face presentation onset and Gabor patch offset ensured that participants allocated only covert attention to the task (which it is claimed would have kept attention constant for the fearful versus neutral trials). The orientation discrimination task

was used as performance on the task improves as stimulus contrast increases (Carrasco et al., 2000; Cameron et al., 2002). Participant accuracy at detecting the orientation of the Gabor patch was larger if the Gabor patch was preceded by a fearful face compared to a neutral face. This effect was not present for inverted faces (in which emotional detection is more difficult; McKelvie, 1995). The authors argued that the level of contrast needed to perform the orientation discrimination task was lower when the stimuli were preceded by a fearful face than a neutral face, i.e. the mere presence of an emotional stimulus had enhanced contrast sensitivity. While this is an impressive claim, further studies need to be conducted to ensure that the effect of emotion on perception really was independent of attentional mediation, and did not reflect an effect of arousal, rather than emotion per se.

Emotions as Cues to Action

Behaviours quintessential to the survival of any species are the four 'F's': feeding, fighting, fleeing and reproducing. While these behaviours may be limited in describing the scope of human behaviours, their acute importance is incontrovertible. An external stimulus relevant to one of these behaviours can reflexively produce physiological reactions (i.e., without our conscious intention). For example, a delicious-looking apple might cause us to salivate (although this requires learning). On the other hand, stimuli relevant to the final three behaviours (e.g. an angry person, a snarling bear, a potential mate) can cause our hearts to pump, getting oxygen circulated to the muscles needed to implement the necessary actions. Studies investigating the effect of emotional stimuli on physiological arousal have found that both skin conductance (Öhman, 1986; Esteves, Dimberg & Öhman, 1994; Dimberg & Öhman, 1996) and levels of stress hormones (van Honk et al., 2000) increase subsequent to presentation of emotional faces, even if presented below the threshold of awareness.

Reflexes have also been shown to be potentiated by emotional stimuli. The startle eyeblink reflex (a blink response to a sudden auditory stimulus) is larger when participants view threatening pictures (Vrana, Spence & Lang, 1988; Lang, Davis & Öhman, 2000). Spinal reflexes are also mediated by emotional stimuli (Bonnet, Bradley, Lang & Requin, 1995; Both, Everaerd & Laan, 2003). For example the 'foot kick' reflex to a hammertap on the heel tendon is larger (measured using electromyogram, EMG) when viewing both positive and negative pictures, compared to neutral pictures.

While it is clear that emotionally valent stimuli elicit physiological changes and potentiate reflexes, the issue of whether they can automatically (i.e. without conscious awareness of them) elicit actions is under debate.

Arnold (1960) and Frijda (1986) proposed that intuitive appraisal of an emotional stimulus will result in an 'action disposition', to approach or avoid. Much like a cup might provide a 'to be grasped at the handle' affordance, so a snarling dog might provide a 'to be distanced from' affordance. The action disposition might not necessarily reach awareness following intuitive appraisal: it is only a 'readiness'. The studies presented above are consistent with this view. In fact the interface between emotion and action has been further demonstrated in a transcranial magnetic stimulation (TMS) study of the motor cortex. Hajcak and colleagues (Hajcak et al., 2007) found an increase in the activity of the motor cortex when participants were viewing emotionally valent pictures, compared to neutral pictures. They did not, however, control for motor action in the images. Their results may thus be due to the activation of mirror neurons which are premotor neurons that fire both when an animal acts and when the animal observes the same action performed by another (especially conspecific) animal (Rizzolatti et al., 1988; Rizzolatti & Arbib, 1998). It is possible that the emotional pictures simply depicted more people, and thus more action.

Following reflective appraisal, Frijda (1986) proposed that awareness of action readiness includes the felt urge to approach or flee, and that the experience of action readiness is essential in emotion experience. This concurs with the 'motivational' view of emotion described earlier. Once a stimulus has been intuitively appraised, it may receive further processing from an appetitive or aversive system. This suggests that its action disposition has been assessed as either approach or avoid, in order to gain entry into either the appetitive or aversive system.

Can emotion automatically (without conscious awareness) cause action? Studies of physiological arousal and reflexes showed that these reactions can occur automatically following an emotional stimulus, but can actions? LeDoux's subcortical 'Low-Route' is again of interest here (LeDoux, 1996). While it may not be able to bypass cognition to result in an emotional feeling, can it bypass cognition (awareness) to result in an action? I.e., can we run away from a snake without knowing what we are running from?

Emotional appraisal of stimuli can prime approach and withdraw responses in humans. In an experiment conducted by Chen and Bargh (1999), participants were instructed to evaluate a target word by moving a lever either toward their body or away from their body, depending on the valence of the word. Surprisingly, the authors corresponded a pull of the lever toward the body with approach responses and pushing the lever away with avoid responses (even though the stimuli were presented in front of the participants in which a lever pull would have meant putting more distance between the stimulus and the lever, and a lever push would have meant putting less distance between the stimulus and the lever). Regardless, participants were faster to pull the lever than push it for positively valenced stimuli, and were faster to push the lever than pull it for negatively valenced stimuli. The authors interpreted this result as an automatic activation of approach

responses to positive stimuli and avoid responses to negative stimuli. The 'automaticity' of this can be questioned though. The words were presented inside conscious awareness. Further, the very use of words as stimuli requires cognition as the words must be read in order to determine their valence.

A further piece of evidence against the automatic activation of actions is that these effects appear to be mediated by goals and intentions. Markman and Brendl (2005) demonstrated that when presented with their own name on the computer screen, participants are faster to move positive words towards it, than away from it. This occurred whether the direction of movement was towards or away from their own body. In a study using faces, Rotteveel and Phaf (Rotteveel & Phaf, 2004) required participants to push or pull a lever in order to categorise them as either positive or negative, or as male or female. In the valence categorisation task, positive faces primed approach responses and negative faces primed avoid responses. However in the gender task, there was no interaction of response with valence, indicating that when emotion is an irrelevant dimension (i.e. it is not attended), positive and negative stimuli do not prime actions, and thus the link between emotion and action is not automatic.

The evidence for a link between emotion and action without conscious awareness is tenuous and suggests that it is not an automatic process. The priming of actions by emotion provides us with one more interesting piece of evidence though. While facial expressions of 'fear' and 'anger' may both be aversive, fear has been shown to prime avoid responses and anger to prime approach responses (Wacker, Heldmann & Stemmler, 2003). This dichotomy of responses between two stimuli of the same valence lends further support to the motivational view of emotion processing, as opposed to a simple valence view.

Emotions Enhance Memory

The majority of studies to investigate emotional enhancement of memory have found that emotionally charged real-life events (D'Argembeau, Comblain & Van der Linden, 2003; Talarico, LaBar & Rubin, 2004), film clips (Christianson & Loftus, 1991; Burke, Heuer & Reisberg, 1992; Guy & Cahill, 1999), and lists of words and pictures (Rubin & Friendly, 1986; Bradley, Greenwald, Petry & Lang, 1992; Ochsner, 2000; Kensinger & Corkin, 2003) are all remembered better than neutral counterparts. Most studies focus on explicit (declarative) recall of emotional events, but implicit memory has also been shown to benefit from emotional arousal. In a repetition priming paradigm (where the probe is identical to the prime), previously presented taboo word probes are named faster than neutral word counterparts (Thomas & LaBar, 2005).

Some authors suggest that negative events are remembered better than are positive events. In tasks requiring participants to indicate whether they vividly remember an event or simply know that it occurred, negative events tended to be 'remembered' better than positive ones (Ochsner, 2000). Conversely, recognition for facial identity has been shown to be better for faces that were originally seen with happy compared to angry expressions (D'Argembeau, Van der Linden, Comblain, & Etienne, 2003). Other studies have demonstrated no difference in memory for positive and negative events (Bradley et al., 1992; Hamann, Cahill & Squire, 1997; Hamann, Monarch & Goldstein, 2000).

While it seems intuitively obvious that emotional events are remembered better than bland events, some studies actually reveal reduced memory for emotional events. This has been repeatedly demonstrated: for a negative event, elements that are centrally tied to the emotional item (e.g. a gun) are remembered, but at the expense of peripheral event details

(e.g. the mugger's face; Loftus, 1979; Kensinger, Garoff-Eaton & Schacter, 2007). Some posit that this occurs because emotions trigger binding mechanisms that link an emotional event to salient contextual features. In a Stroop task containing neutral and taboo words, participants were required to name the font colour. In a subsequent surprise memory task, participants remembered the colour of the taboo words better than the colour of the neutral words (MacKay et al., 2004). In a follow-up study, taboo words did not enhance memory for the location the words were presented in, compared to neutral words (MacKay & Ahmetzanov, 2005). The authors suggest that only details relevant to the task were enhanced in memory due to the binding mechanism triggered by the emotion. The enhanced memory for central detail is echoed in studies finding mood congruent enhancements of memories. People in a positive mood remember positively valenced stimuli better than negatively valenced stimuli (Matt, Vásquez & Campbell, 1992). Studies of people with depression and anxiety show that these people tend to remember more negative details than positive ones (Mathews & Bradley, 1983; Williams, Watts, MacLeod & Mathews, 1988; Mathews, Mogg, May & Eysenck, 1989).

There are a lot of questions left to be answered regarding whether memory is enhanced by emotion. It has been suggested that these effects occur due to increased rumination of emotional events: that emotional events are not initially encoded any better than neutral events, but that repeated retrieval of information strengthens memory (Guy & Cahill, 1999). Also, it may be that the memory itself is not enhanced for an emotional event, but that confidence in the memory is enhanced instead (Ochsner, 2000; Yonelinas, 2002; Rotello, Macmillan, Reeder & Wong, 2005). Furthermore, 'liking' a stimulus may lead to a misattribution of liking to remembering the stimulus (Monin, 2003). Attractive faces are more often called 'familiar' than unattractive faces. This may be because attractive faces are

often more 'average' than unattractive faces (Langlois & Roggman, 1990), in which case they are more likely to conform to a mentally-constructed prototype image (Rosch, 1978) which people misinterpret as prior exposure (Solso & McCarthy, 1981; Bomba & Siqueland, 1983). Alternatively, average faces may be more 'fluent' to process, and this ease of processing may be misinterpreted as memory (Jacoby & Kelley, 1987; Jacoby, Kelley, Brown & Jasechko, 1989). I will return to 'fluency' theories later.

One last point: Levine and Pizarro (Levine & Pizarro, 2004) point out that the vast majority of studies have focused on emotional arousal or valence, without much regard for discrete emotional states. If, as suggested by appraisal theories, emotions are responses to changes in goals, then different emotions will reflect different motivations, goals and actions. Therefore, what may be remembered in one negative emotion (angry) may not be the same as another negative emotion (fear). This is especially interesting when we remember their differential effect on approach and avoid actions.

Clearly, attention may also be a mediating factor. If attention is drawn to emotional events or features, then they are more likely to be processed and encoded than neutral events or features. Thus, whether emotion has a direct effect on memory is called into question. As I will now outline below, an emotional stimulus is very compelling to attention.

Emotion Biases Attention

Although the ability to evaluate emotional significance independently from attention has been called into question (e.g. Vuilleumier et al., 2001; Pessoa et al., 2002; Silvert et al., 2007) a great deal of research suggests that awareness (and perhaps use of attentional resources) can be biased by emotional stimuli.

This idea is not new. In the 1950s the “cocktail party effect” (Moray, 1959) demonstrated that awareness could be captured by one’s own name, even when embedded in a stream of otherwise ignored information. This can be considered as early evidence for the role of personal meaning in guiding attention (Robinson, 1998). More recently behavioural, rather than anecdotal, evidence has been put forward in favour of this effect. Evidence comes from several sources: neurological evidence, visual search, flanker tasks, dot probe tasks, cueing and attentional blink. I will discuss each of these, and their various contributions to the theory that emotion biases attention, in turn.

Neuropsychology

Patients with unilateral neglect typically have right-hemisphere damage and are therefore unable to attend to stimuli presented in the left visual field. Extinction refers to the finding that patients are particularly bad at attending to stimuli presented in the left visual field when stimuli are also presented in the right visual field. The stimuli are thought to compete for attention, and the stronger right visual field will extinguish any stimuli presented in the weaker left visual field. Importantly, this extinction is less likely to occur if an emotional stimulus (rather than a neutral stimulus) is presented in the left visual field (Vuilleumier & Schwartz, 2001a, 2001b; Fox, 2002). It is suggested that the emotionality of the stimulus draws attentional resources, allowing a stimulus that would otherwise have been extinguished to gain access to awareness.

Behavioural Evidence

Dot probe tasks (MacLeod, Mathews & Tata, 1986) have been used to assess emotional biases of attention using ‘normal’ participants. In these tasks, a pair of pictures is displayed in separate locations. One of the pair is a neutral picture; the other is an emotional

picture. The pictures are presented for a short amount of time, at which point they disappear to be replaced by a small dot (the probe) presented at the location of one of the previously presented pictures. Participants are instructed to locate the dot as quickly as possible. Systematically faster responding to the dot probe presented at the location previously occupied by the emotional picture indicates that attention was previously allocated to that location, and that emotional information engaged attention (Mogg & Bradley, 1999).

Likewise the flanker paradigm (described above) provides evidence for an attentional bias by emotion. In an emotional flanker task (Fenske & Eastwood, 2003), targets are either positive or negative faces, and participants are instructed to report the valence of the target face. Flankers (distractors) can be emotionally congruent (e.g. an angry target flanked by angry distractors), emotionally incongruent (e.g. an angry target flanked by happy distractors) or neutral (e.g. an angry target flanked by distractors to which no response is required). RTs to targets with incongruent distractors are longer than RTs to targets with congruent distractors. This indicates that the distractors could not be ignored. Interestingly, RTs to an incongruent trial in which the flankers are negative are longer than RTs to an incongruent trial in which the flankers are positive. This suggested that negative stimuli draw more attentional resources than positive stimuli, a result that is echoed in the emotional visual search literature.

An emotional variant of the visual search task was developed by Hansen and Hansen (1988). Search arrays were composed of faces: one face in the array (the target) had a different emotional expression than the remainder (the distractors). It was found that participants were faster to locate an angry target in an array of happy distractors than to locate a happy target in an array of angry distractors. Hansen and Hansen concluded that angry faces draw more attention than happy faces. However, the methodology was flawed

in that the effect may have been a result of the distractors, rather than the target. For example, perhaps it is quicker to search through and reject happy faces as distractors than angry faces (Hampton, Purcell, Bersine, Hansen & Hansen, 1989; Koster, Crombez, Van Damme, Verschuere, & De Houwer, 2004). There have also been inconsistent results (Hampton et al., 1989; Suzuki & Cavanagh, 1992), failures to replicate (Nothdurft, 1993; White, 1995) and findings that low-level visual artefacts may have been responsible for the effect (Purcell, Stewart & Skov, 1996).

To remedy this criticism, Eastwood and colleagues (Eastwood, Smilek & Merikle, 2001) presented either an angry or happy schematic face target in an array of neutral schematic face distractors. It was found that increasing the set size had a lesser impact (a shallower slope) on the angry targets than the happy targets. This suggested that angry faces attract attention more efficiently than happy faces (see also Fox et al., 2000). The effect disappeared when the faces were inverted, supporting the concept that emotion biases attention, rather than physical features of the stimuli. Similar results were found with error rate data (Öhman, Lundqvist, & Esteves, 2001): for friendly face targets, participants made more errors as set size increased, this effect was not evident for angry face targets. These effects have been replicated using carefully controlled photos of real faces (Horstmann & Bauland, 2006) and pictures of animals (Öhman, Flykt & Esteves, 2001).

Do negative stimuli bias attention more than positive stimuli?

Further research suggests that not only do angry faces attract attention, but they also hold attention for longer relative to positive and neutral faces (Fox, Russo, Bowles & Dutton, 2001). Fox and colleagues (2001) used a spatial-cuing paradigm in which a cue was presented to the left or right of fixation, and after a short delay was followed by a target

circle that was presented on the same side or on the opposite side as the cue. The participants were required to locate the target. Responses to targets are facilitated when they appear on the same side as the cue (validly-cued trials) and are slowed when they appear on the opposite side to the cue (invalidly-cued trials). In this study, positive, negative or neutral schematic faces were presented as cues. Critically, the emotional valence of the cue affected RTs on invalidly-cued trials but not on validly-cued trials. RTs on invalidly-cued trials were slowest when the cue was negative. The authors concluded that an angry face cue held attention, slowing disengagement from the location of the invalid cue required in order to respond to the target presented on the opposite side. This study has been replicated using threat-related scenes as cues (Yiend & Mathews, 2001).

The studies presented above all suggest that negative stimuli bias attention, but do not suggest much of a role for positive stimuli in biasing attention. This is a view held by many researchers in the field. There are three schools of thought regarding the differential bias of attention by positive and negative stimuli: the categorical negativity hypothesis, the evolutionary threat hypothesis and the arousal hypothesis.

The categorical negativity hypothesis proposes that stimuli are unconsciously evaluated as either positive or negative. Stimuli that are evaluated as negative attract attention because the detection of negative stimuli is more critical to survival than the detection of positive stimuli (Pratto & John, 1991). This theory proposes that initial evaluation of stimulus valence will result in only a categorical evaluation of stimuli positivity or negativity: the mechanism does not provide information about the degree of pleasantness or unpleasantness along the valence dimension. Therefore, a mildly negative stimulus would attract attention to the same degree as a highly negative stimulus. However Mogg and colleagues (Mogg et al., 2000) found that extreme negative pictures attracted

more attention than mildly negative pictures. Also, in some cases, positive stimuli have been shown to attract attention (Buodo, Sarlo & Palomba, 2002; Anderson, 2005). Pratto (1994) for example found that the negativity bias disappeared when a list of positive words included arousing stimuli.

The evolutionary threat hypothesis is of a similar vein to the categorical negativity hypothesis. It posits that the mechanism has evolved to detect stimuli that at some point presented a threat to human survival. Thus, the evolutionary threat hypothesis makes slightly different predictions about the types of stimuli that will capture attention, in that they should be limited to stimuli that signalled a threat to survival during evolution (e.g. angry faces, spiders or snakes). However some studies have failed to find evidence that pictures of spiders or snakes attract more attention than neutral pictures (Kindt & Brosschot, 1997; Thorpe & Salkovskis, 1998), even if the participants are phobic of spiders and snakes (Lavy, Van den Hout & Arntz, 1993; Merckelbach, Kenemans, Dijkstra & Schouten, 1993; Kindt & Brosschot, 1999). More compelling against the evolutionary threat hypothesis though, is the finding referred to earlier that positive stimuli can sometimes attract attention.

Arousal Theories of Attentional Capture

Finally, the arousal hypothesis falls in line with the motivational approach to emotion (Lang, 1995). Recall that the motivational approach suggests that stimuli are evaluated along two dimensions: valence and arousal. While some responses to stimuli vary with the valence of this stimulus (e.g. affect is determined by stimulus valence), attention is believed to vary with the arousal level of the stimulus. More arousing stimuli will attract more attention. This theory accounts for why negative stimuli more often capture attention than positive stimuli (because they are generally more arousing), and why positive stimuli

have sometimes been shown to capture attention (if they are arousing enough). The most compelling evidence in favour of this theory over the categorical negativity or evolutionary threat hypotheses come from emotional variations of the attentional blink (AB) paradigm.

The traditional AB paradigm is comprised of a RSVP stream of items presented in a single location at a rapid speed (often 10 items per second). When participants search the stream for two targets, report accuracy for the first target (T1) is typically high. Report accuracy for the second target (T2) varies as a function of the lag between T1 and T2. If T2 is presented approximately 500 ms after T1, report accuracy is high, whereas if T2 is presented less than 500 ms after T1 report accuracy suffers (the AB; Raymond, Shapiro & Arnell, 1992). In contrast, if the RSVP stream consists of only one target, report accuracy is uniformly high, as distractor items do not deplete the attentional resources needed to perform the target task. The AB is usually explained in terms of T1 consuming attentional resources for a period of approximately 500 ms, during which time T2 does not receive sufficient attentional resources in order for it to be processed and consolidated into memory (Chun & Potter, 1995; Shapiro, Arnell & Raymond, 1997; Jolicoeur, 1998, 1999).

There have been several emotional variants of the AB paradigm, each manipulating a different item in the RSVP stream. Anderson (2005) and Keil and Ihssen (2004) manipulated T2. Both studies demonstrated that when T2 was an emotionally arousing word, the AB was dramatically reduced compared to when T2 was a neutral word. This suggests that the emotional T2 receives preferential attentional processing. The attenuated AB for emotionally arousing words was not found in a study of patients who had their left amygdala resected (Anderson & Phelps, 2001). This supports the claim that the attenuated emotional AB in normal participants was a result of the emotional salience of T2 words.

Mathewson, Arnell and Mansfield (2008) manipulated the emotional arousal of T1. They observed that the magnitude of the AB increased as a function of T1 arousal. When taboo or sexual words were presented as T1, the AB was larger than when positive, negative or neutral unarousing words were presented. The arousal level, and not the valence of T1 predicted the size of the AB. Arousing T1 words captured more attention, leaving less attentional resources 'left-over' for the processing of T2.

Arnell, Killman and Fijavz (2007) did not alter targets, but instead manipulated the emotional arousal of distractor items in the RSVP stream. Specifically, they presented either an arousing or an unarousing distractor item before T1 (there was no T2 in the stream). Distractors could be emotionally neutral, positive, negative, sexual, threatening or anxious. All distractors were pre-rated for both arousal and valence. The participants' task was to report a colour word (e.g. the word brown, all targets were presented in the same colour font). The authors found that when an unarousing critical distractor was presented before T1, report accuracy was high regardless of the valence of the item. Conversely, when an arousing critical distractor was presented before T1, report accuracy suffered regardless of valence. The AB magnitude varied as a function of distractor arousal, not valence. A highly arousing critical distractor captured attention and prevented detection of the target. This finding is not limited to words: similar effects were found using arousing versus unarousing pictures rather than words as critical distractors by Most and colleagues (Most, Smith, Cooter, Levy & Zald, 2007).

The view that emotional arousal, and not valence is critical in biasing attention I believe explains the varied findings of mood studies of emotional effects on attention. If we view arousal as an indication of how personally relevant a stimulus is, we can speculate that

different stimuli are differentially relevant to different individuals, and are thus differentially arousing.

For example, a stimulus depicting threat may be deemed highly relevant by an anxious individual, and will thus be highly arousing. Therefore we could predict that threatening stimuli will capture more attention in high-anxiety participants than in low-anxiety participants. Indeed this is the case. In a dot-probe task, participants in high and medium anxiety groups were slower to respond to targets when invalidly-cued by threatening faces compared to neutral faces. There was no difference between threatening and neutral invalidly-cued trials in the low anxiety group (Bradley, Mogg & Millar, 2000). In further support of this, the high and medium anxiety groups did not show a bias for happy invalidly-cued trials (nor did the low anxiety group), suggesting that attention was not drawn to a stimulus that was not arousing. Likewise in search tasks, threatening words and pictures are more likely to capture attention in anxious participants than in non-anxious participants (Mathews & MacLeod, 1994; Yiend & Mathews, 2001). In emotional AB tasks, only high anxious participants show attenuated AB by fearful face stimuli. Low anxious participants show normal AB effects: their attention is not captured by fearful faces (Fox, Russo, Bowles & Dutton, 2001).

Angry stimuli may be particularly arousing to angry participants. Indeed van Honk and colleagues (van Honk, Tuiten, de Haan, van den Hout & Stam, 2001) found that participants who scored high on trait anger showed an attentional bias for angry faces in a pictorial version of the emotional Stroop. Participants who scored low on trait anger showed no bias to angry faces. In a second experiment, participants were divided first into groups based on trait anger, and then into groups based on trait anxiety levels. Again, when the participants were divided into high and low trait anger groups, only the high anger group

showed the bias to angry faces. When the participants were subsequently grouped by trait anxiety, there was no relation between anxiety group and attentional bias to anger. Thus, angry faces are arousing to angry participants and so capture attention.

These mood studies suggest that attention will be captured by stimuli that are personally relevant to the individual, and are thus arousing to the individual. We can go one step further here. It is not just mood studies that support the arousal view: motivational studies of attention capture demonstrate the same effect.

To smokers, a picture of a cigarette is a relevant, highly arousing stimulus. To the majority of non-smokers (who have never smoked) it is not. In a probe task, smokers were faster to respond to a cue that replaced a smoking-related picture than a cue that replaced a neutral picture. This suggests that the smoking-related picture had captured smokers' attention. In the non-smokers, there was no difference in RT to detect a cue that replaced a smoking-related versus a neutral picture (Mogg, Bradley, Field & De Houwer, 2003). Eye movement data collected in the same study revealed that smokers, relative to non-smokers, oriented quicker and looked longer at smoking-related pictures compared to neutral pictures. The highly relevant, arousing pictures both captured and held the smokers' attention. This supports the arousal view of attention capture and directly goes against the negativity-bias view. The cigarettes were arousing to the smokers, but would have been negative to the non-smokers. Therefore, supporters of the negativity-bias view would have predicted cigarettes to capture non-smokers' attention.

Likewise, Dalglish (1995) found that compared to 'normal' participants, bird-watchers displayed an attentional bias towards bird-related words. To a bird-watcher, the

word kestrel is arousing and captures attention; to a non-bird-watcher, it is not, and does not.

These studies combine to provide strong support for the arousal hypothesis of attentional bias. Further, they reinforce the notion of reflective appraisal. An individual will evaluate a stimulus in terms of personal goals and concerns. Threat-related pictures are relevant to anxious individuals and will be evaluated as such. Pictures of cigarettes are relevant to participants who are stuck in a testing room and are counting down the minutes until their next nicotine fix, and thus are relevant and arousing and will capture their attention. Bird-watchers' attention is captured by bird-related words, and erotic words and pictures attenuate the attentional blink of university undergraduates. We are all motivated and aroused by different stimuli, and as such there will be individual differences in the bias of attention. This ensures that stimuli that are of maximum relevance to us are the stimuli that are processed in our limited-capacity attentional system and are thus most likely to reach awareness and be acted upon appropriately.

Biasing Emotional Evaluation

As has been demonstrated, an emotional reaction to an object will make interaction with, or response to that object more efficient. While some objects might be innately valent (i.e. a snake) and their valence might be learned (a snake is dangerous), a system that allows a previously encountered neutral stimulus to be encoded in memory as emotionally valent (or a 'tag' to be interpreted by and speed intuitive appraisal) would promote more efficient responding on future encounters. Here I discuss evidence in favour of such a system in the form of: a) classical conditioning; b) evaluative conditioning; c) affective priming of neutral stimuli; d) mere exposure; and e) distractor devaluation studies.

Classical Conditioning

Classical conditioning (or Pavlovian conditioning) refers to a set of experimental procedures in which the experimenter arranges a contingency between stimuli by presenting those stimuli independent of the individual's behaviour. In the original study, an initially neutral stimulus (a bell) was paired with a biologically relevant unconditioned stimulus (US; food) that normally elicits a reflexive or unconditioned response (UR; salivation). As a result of the pairings, the bell became a conditioned stimulus (CS) that was capable of evoking salivation as a conditioned response (CR). Pavlov (1927) argued that the CR developed because an association had formed between a representation of the CS and one of the US. In this way, novel neutral stimuli, through associative pairing, could control response mechanisms.

Further, it has been suggested that the CS can evoke a representation of an emotion (such as fear if the CS has been paired with a shock, or excitement at the expectation of reward). It is proposed that the emotional valence of the CS is 'tagged' to its representation. This is supported by transreinforcer blocking demonstrations in which the presence of a CS that has previously been paired with a shock can prevent conditioning to a CS paired with the absence of an expected food reward (Dearing & Dickinson, 1979). The individual does not learn the second CS in the presence of the first CS because it predicts the same (aversive) US. The two reinforced stimuli share no commonalities other than their aversiveness, and therefore the blocking is believed to be a result of an association between the CS and its affective valence.

Classical conditioning has thus been proposed as a mechanism that leads to attitude formation (Eagly & Chaiken, 1993). However, in modern learning theories, classical

conditioning is considered an instance of signal learning, a higher order cognitive learning mechanism that allows the individual to make predictions about significant events in the environment (Rescorla & Wagner, 1972; Mackintosh, 1983). What is learned by the individual is not the valence of the stimulus, but the 'if-then' relationship between the CS and the US. It has been put forward that attitude formation towards a stimulus does not refer to the prediction of events, rather the affective meaning the stimulus acquires in the context of pleasant or unpleasant outcomes (Cacioppo, Marshall-Goodell, Tassinary & Petty, 1992). While the efficient prediction of events is important in guiding decision making processes, it may not be that classical conditioning allows the stimulus to acquire an affective 'tag' *per se*.

Evaluative Conditioning

Evaluative conditioning addresses this issue and has been described as the learning of likes and dislikes, i.e. the acquisition of preferences. Evaluative conditioning refers to the transfer of affect from an US to a CS as the result of a learning procedure. The paradigm was first described by Levey and Martin (Levey & Martin, 1975; Martin & Levey, 1978) and was developed by Baeyens and colleagues (Baeyens, Eelen & Van den Bergh, 1990). The paradigm consists of three sequential phases. In the baseline phase, participants rate the valences of the entire stimulus set. On the basis of this, stimuli are grouped into liked, disliked and neutral categories. In the learning phase, the neutral stimuli are paired with either a liked or a disliked stimulus and are presented to participants. In the test phase the stimulus set is rated for a second time. Previously neutral stimuli that were paired with liked stimuli are subsequently rated more positively, and previously neutral stimuli that were paired with disliked stimuli are subsequently rated more negatively. In contrast to the classical conditioning procedure, the CS (the formerly neutral picture) does not acquire a

predictive value, but instead attains the affective qualities of the US. This effect is usually described in terms of the formation of an association between the representations of the US and the CS.

Interestingly, and unlike classical conditioning, evaluative conditioning can occur without participants possessing awareness of the learning contingencies (Baeyens et al., 1990; De Houwer, Hendrickx & Baeyens, 1997; Hammerl & Grabitz, 2000). Indeed evaluative conditioning can be reduced (and even inverted) when participants are aware versus unaware of the pairings (Baeyens, Heremans, Eelen & Crombez, 1993; Hammerl & Grabitz, 2000; Walther, 2002; Fulcher, 2002). However, attention to the pairs (but not necessarily consciously) is necessary for learning to occur: dividing attention attenuates learning (Nissen & Bullemer, 1987; Reber & Squire, 1994, 1998; Field & Moore, 2005).

Affective Priming of Neutral Stimuli

Very similar to evaluative conditioning is the affective priming paradigm in which an emotional prime precedes a neutral probe. In much the same way that the previously neutral CS acquires the valence of the US, the neutral probe will 'acquire' the valence of the emotional prime.

Niedenthal (1990) provided the first demonstration of this effect. Participants were asked to form an impression of a cartoon character. They were given a series of adjectives and were asked if each matched with their impression of the cartoon character, 'yes' or 'no'. Prior to the presentation of each cartoon character, a face depicting either joy or disgust or a neutral expression was presented as a subliminal prime. Participants in the disgust condition matched more negative adjectives to the cartoon than participants in the joy condition. Thus, the emotionality of the prime faces had been misattributed to the cartoons. However,

this study can be thought of as a more general investigation of mood on emotional evaluation. This was a between-subject design, so a participant would have always been primed with either a disgust, a joy or a neutral face. Thus, it may have been that a mood was induced in the participants, which the cartoons were devalued as a result of.

In a within-subjects study, Murphy and Zajonc (1993) presented participants with Chinese ideographs that were preceded by 4 ms subconscious primes. The primes were faces expressing either positive or negative emotions. When preceded by a positive facial expression, the ideographs were judged more positively than when preceded by a negative facial expression. Even though participants were not aware of the presence of the face prime, its emotionality had been evaluated, and this was misattributed as an emotional evaluation of the Chinese ideograph. This effect has since been replicated (Murphy, Monahan & Zajonc, 1995; Rotteveel, de Groot, Geutkens & Phaf, 2001).

Although it cannot be said that the probe's valence 'tag' has been altered in these cases, the affect of the ideograph was influenced by the prior presentation of an emotional face. Further research needs to address whether this effect remains on subsequent presentations of the probe without the preceding prime, in order to assert whether the emotional 'tag' has been altered or whether the effect is merely due to a misattribution that only lasts for one trial.

Mere Exposure

Unreinforced repeated exposure to an affectively neutral and unfamiliar stimulus results in a positive affective judgement. This is the mere exposure effect (Zajonc, 1968). Experiments using a wide range of stimuli (e.g. Chinese ideographs, 'Turkish' words, line drawings and faces) and procedures (e.g. forced-choice preference judgments, likeability

ratings, pleasantness ratings, behavioural indices of preference and self-reports of mood) have demonstrated this phenomenon (Harrison, 1977; Bornstein, 1989). Repeated pre-exposure has not only been found to enhance preference judgements, but also to elicit psychophysiological responses such as smiling (measured using facial electromyography; Winkielman & Cacioppo, 2001; Harmon-Jones & Allen, 2001).

The mere exposure effect is robust, but it sometimes fails to appear, and seems to rely on several boundary conditions. For example, pre-exposure has the strongest influence on preference when stimuli are initially presented for relatively short durations (Bornstein & D'Agostino, 1992) with low pre-exposure frequency (Van den Bergh & Vrana, 1998) and when stimuli are complex (Cox & Cox, 1988, 2002; Bornstein, Kale & Cornell, 1990).

Fluency theories of the mere exposure effect suggest that a previous encounter with a stimulus will enhance the processing fluency and this will serve as a basis for preference (Seamon, Brody & Kauff, 1983; Jacoby, Kelley & Dywan, 1989; Bornstein & D'Agostino, 1994; Whittlesea, 1993; Whittlesea & Price, 2001). Enhanced processing fluency is defined as the ease with which information can be processed, as reflected by the speed and ease with which a stimulus is perceived.

If processing fluency due to prior exposure is key to mere exposure effects, Schwarz and colleagues (Reber & Schwarz, 2001; Winkielman, Schwarz, Fazendeiro & Reber, 2003; Schwarz, 2004; Reber, Schwarz & Winkielman, 2004) predicted that any other variables that facilitate fluent processing should likewise increase preference (even without pre-exposure). Several studies have supported this idea: variables such as figure-ground contrast, clarity of stimulus presentation, symmetry and presentation duration have been shown to ease processing and increase preference judgements (Reber et al., 2004).

Three major fluency theories have been proposed. Several authors suggest that fluency, and therefore previous exposure evokes the experience of processing ease that itself is positively valenced and is subsequently interpreted as stimulus quality (Reber, Winkielman & Schwarz, 1998; Phaf & Rotteveel, 2005). This is supported by Monahan and colleagues who found that following repeated exposures, self-reported mood is elevated in participants (Monahan, Murphy & Zajonc, 2000). The authors continued to demonstrate that even novel stimuli unrelated to the exposed stimuli benefit from following a repeated-exposure block of trials. They concluded that the positive mood state induced by fluency is interpreted as stimulus quality. However, stimuli that were actually exposed were preferred to novel stimuli, suggesting that the fluency of the stimulus itself must play a role. A positive mood state cannot account for the whole effect.

The perceptual fluency / attributional framework accounts for this, and also the finding that mere exposure effects are often larger when pre-exposure occurs subliminally, rather than supraliminally (Bornstein, 1989). Many researchers have demonstrated that mere exposure effects occur even when participants are unable to consciously recall or recognise previously presented stimuli (Kunst-Wilson & Zajonc, 1980; Seamon, Brody & Kauff, 1983; Seamon, Marsh & Brody, 1984; Bonanno & Stillings, 1986; Bornstein, Leone & Galley, 1987). Mere exposure effects do not rely on conscious awareness, and thus the learning processes can apparently be unconscious, involving implicit rather than explicit memory. Bornstein and D'Agostino (1992, 1994) draw a distinction between implicit and explicit memory effects in mere exposure. They suggest that, in the latter case, participants realise that their performance may be affected by their prior exposure to the stimuli and in some cases 'correct' their initial interpretation of the resulting perceptual fluency. However, when stimuli are presented subliminally, there is less opportunity for the participants to

become 'aware' that stimuli have been previously presented. Participants are likely then to misattribute fluency to a preference for the stimuli.

A problem to this account of fluency though, is that increasing the number of exposures to a point increases the size of the mere exposure effect (Zajonc, 1968; Lee, Sundberg & Bernstein, 1993; Kruglanski, Freund & Bar-Tal, 1996). As the number of exposures increases we would expect that the likelihood of stimuli recognition would increase, which would increase correct attributions of fluency to pre-exposure and thus reduce the mere exposure effect. This is not the case.

Whittlesea and Price (2001) highlighted this problem, and extended it. Note that they pointed out that it was difficult to understand why the perceptual fluency was used as a basis for affective but not recognition judgements. It has been proposed that there are two bases for making recognition judgments: retrieval of contextual detail and a feeling of familiarity. This is the two-factor theory of recognition (Jacoby & Dallas, 1981). The familiarity component is presumed to rely on the same fluency as mere exposure, therefore increased fluency might be expected to benefit both preference and recognition judgements. To explain why this is not the case, Whittlesea and Price (2001) proposed that, during preference tasks participants may be adopting non-analytic strategies (relying on fluency), and during recognition tasks they may adopt analytic strategies (relying on the recall of detail). They demonstrated that when participants were encouraged to use a non-analytic strategy recognition performance increased, and when they were encouraged to use an analytic strategy, preference judgements decreased. This analytic / non-analytic view of perceptual fluency accounts for the mere exposure findings well.

Whichever theory we choose to support, the phenomenon of mere exposure demonstrates that the affective valence of an item can be biased. This can be accomplished simply by the prior presentation of the stimulus, even if this is subliminal. Mere exposure and fluency accounts are limited though. They propose only a role for enhancement of preferences. No mere exposure demonstrations suggest a role for attention, nor do they manipulate attentional state toward the stimulus during pre-exposure. If mere exposure findings rely on successful encoding of stimuli (as presumably they do), might effects be increased by an increase of attention to the stimulus during pre-exposure? If attention is eliminated, might mere exposure effects be eliminated? Further to this, if a stimulus is inhibited during pre-exposure, might mere exposure effects be reversed? Fluency accounts do not suggest a role for attention or inhibition; in fact one is left with the impression that 'any exposure is good exposure'. Distractor devaluation studies speak to this issue and demonstrate that affective judgements can not only be influenced by prior exposure, but also by the attentional state of the stimulus during prior exposure.

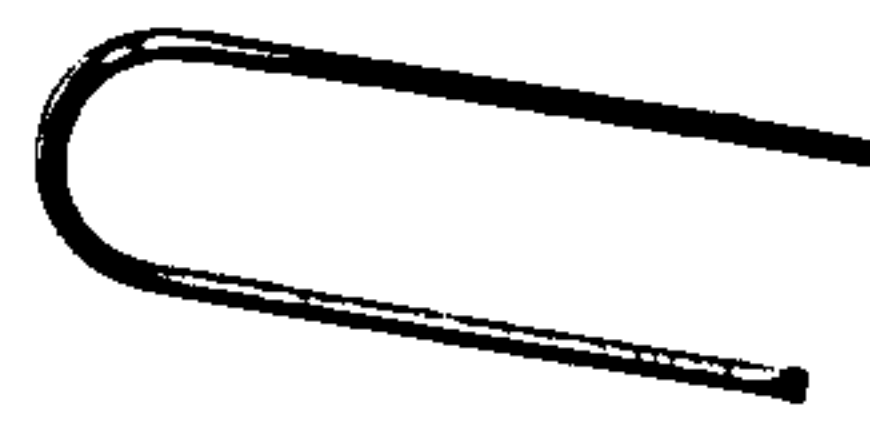
Emotion is Biased by Attention

I have established above that there are links between attention and emotion: emotion can bias attentional resources for example. I have also described how an attentional state can be stored with an object's representation in memory. Can this be used to bias future emotional evaluation? Neurological evidence suggests that the brain areas responsible for attentional and emotional processes are highly interconnected and so may support such a system. Below I briefly review some of this neuropsychological evidence.

Neurological Evidence for an Attention-Emotion Link

Can we evaluate more than one object at a time? Common sense suggests not. The attentional system may engage to ensure that only a single, highly relevant object is put forward for evaluation. The following diagram (Figure 7) is taken from Palermo and Rhodes (2007). It describes the face perception and attention systems and is an amalgam of the findings from several influential papers. While Palermo and Rhodes primarily focus on the systems involved in emotional face detection, the diagram illustrates the connections between

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.



The model presented in Figure 7 is a highly simplified representation of the regions implicated in emotion and attention processing, and the connections between them. Areas in red represent regions involved in emotion analysis, as described by Adolphs (2002). Areas shaded blue represent the fronto-parietal cortical network involved in spatial attention, as described by Hopfinger, Buonocore and Mangun (2000). Less relevant to the current

discussion are the bevelled-edge rectangles (indicating the core system for face perception, Haxby et al., 2000), and areas shaded in yellow (identity and associated semantic information, Adolphs, 2002). In the diagram, solid lines indicate cortical pathways and dashed lines represent the subcortical route for rapid and/or coarse emotional expression processing.

While the diagram above is by no means definitive, it provides a nice place to start. The amygdala has received the lion's share of attention in emotion research. Indeed I have already briefly discussed the amygdala and whether it can 'recognise' an emotion without attentional processing. However, below I will focus my discussion on the prefrontal cortex (PFC) specifically, the anterior cingulate cortex (ACC), as these are the areas that point to a reciprocal interaction between attention and emotion systems (as highlighted in the diagram above in blue and red). Note that I will not be discussing the superior temporal sulcus as this is implicated in face perception (and other bodily expression of emotion), and is not of primary interest to a more general discussion of attention and emotion.

It is largely believed that the PFC maintains the representations of goals and the means to achieve them (Miller & Cohen, 2001). This appears to be particularly true in cases where a situation is ambiguous, in which the PFC facilitates elicitation of task-relevant responses despite conflict by potentially stronger alternatives, by sending bias signals to other areas of the brain. Anticipation involving emotional experience associated with anticipated choice, and affect-guided planning of behaviours, are the crux of adaptive, emotion-based decision making, and it has been consistently shown to be impaired in patients with lesions of ventromedial PFC (Damasio, 1994).

Specifically, left-sided PFC regions are believed to be involved in approach-related, appetitive goals, and right-sided PFC regions are thought to be critical in the maintenance of goals that require behavioural inhibition and withdrawal in circumstances involving strong alternatives. Recent neuroimaging studies suggest that the representation of rewards and punishers are represented in the orbital and ventral frontal cortex, and that different areas may differentially emphasise reward versus punishment (Kawasaki et al., 2001; O'Doherty et al., 2001). Further, left-sided medial regions of the orbitalfrontal cortex (OFC) have been shown to be particularly responsive to rewards. Conversely, lateral right-sided regions are particularly responsive to punishers (O'Doherty et al., 2001).

A widely held viewpoint on the relationship between emotion and cognition posits that certain brain regions within the PFC are responsible for either cognitive or emotional tasks. For example, Drevets and Raichle (1998) found that the dorsolateral prefrontal cortex (dlPFC) and the dorsal ACC were more active in cognitive tasks than in emotional tasks. Conversely, the OFC, the ventral ACC, and the amygdala were more active during emotional, versus cognitive tasks.

Exploring this dichotomy in the ACC, researchers have identified two subdivisions: the 'affect subdivision' incorporates rostral and ventral areas of the ACC; the 'cognitive subdivision' incorporates dorsal regions of the ACC (Devinsky et al., 1995; Vogt et al., 1992, 1995; Whalen et al., 1998). The affect subdivision has widespread connections to the limbic and paralimbic regions (amygdala, nucleus accumbens, OFC, periaqueductal grey, anterior insula, and autonomic brainstem motor nuclei, for example). It is thought to be involved in the regulation of visceral and autonomic responses to stressful behavioural and emotional experiences, emotional expression and social behaviour. The cognitive subdivision, on the contrary, is thought to be critical in response selection and processing of

cognitively demanding information, owing to its high connectivity with the dlPFC, posterior cingulate, parietal cortex, supplementary motor area and spinal cord.

Little is currently known regarding the interactions between the cognitive and affective subdivisions of the ACC. However, it may have a more integrative role than a strictly cognitive versus emotional account would suggest. Several authors have proposed that the affective subdivision may integrate salient emotional and cognitive information and, subsequently, attentional resources within the cognitive subdivision are modulated accordingly (Mega, Cummings, Salloway & Malloy, 1997; Mayberg, 1997; Mayberg et al., 1999; Pizzagalli et al., 2001).

It is a widely held belief that the ACC acts as a bridge between attention and emotion (Devinsky, Morrell, & Vogt, 1995; Ebert & Ebmeier, 1996; Mayberg, 1997; Vogt, Nimchinsky, Vogt & Hof, 1995). It has been described as the point of integration for visceral, attentional and affective information (Thayer & Lane, 2000), and is involved in assessing and responding to the behavioural significance of environmental stimuli.

The ACC has been emphasised as being especially involved in conflict monitoring (Carter, Botvinick & Cohen, 1999; Carter et al., 2000; Miller & Cohen, 2001). It is posited that the ACC is involved in an evaluation capacity, reflecting the degree of conflict that a task elicits. Conflict occurs under conditions wherein two or more possible task-related decisions compete or interfere with each other. The “cognitive monitoring hypothesis” conjectures that the cognitive subdivision of the ACC monitors conflicts between brain regions. If the signal of competition is detected, this output signals the need for controlled processing. The representation and maintenance of task demands necessary for the control

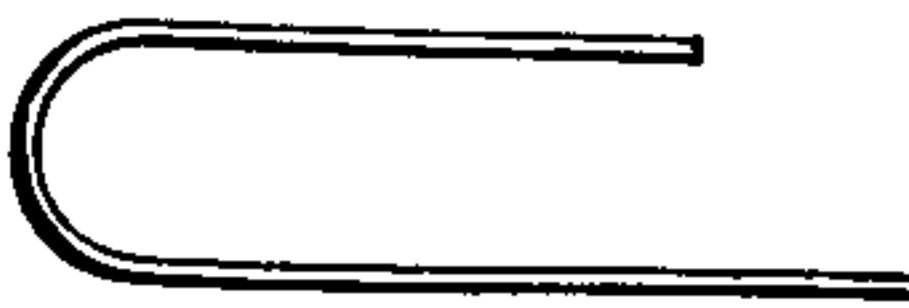
and inhibition required in this process, and the increase in neural activity required for brain regions involved in the conflict is assumed to be critically located in the dlPFC.

In addition, ACC activation has been associated with occurrences of a requirement for the effortful emotional regulation in situations where behaviour is failing to achieve a desired outcome. Further, activation has also been shown to be elicited in affect-inducing non-normative contexts, i.e. most laboratory settings (Bush et al., 2000; Ochsner & Barrett, 2001).

Evidence from studies of the PFC and ACC support reciprocal connections between attention and emotion systems (Compton, et al., 2003). I have already discussed above the behavioural evidence in support of emotional system influence on the attentional system, and now I will focus on behavioural evidence for the reciprocal effect, specifically, studies of 'distractor devaluation'.

Distractor Devaluation

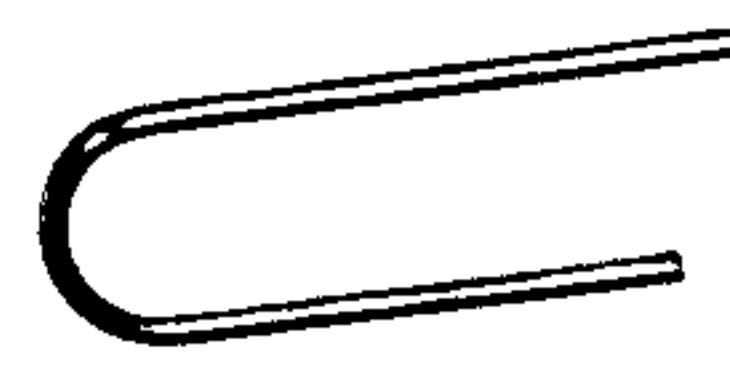
Raymond, Fenske and Tavassoli (2003) demonstrated that selectively attending to or ignoring a stimulus could impact its later affective evaluation (see Fenske & Raymond, 2006 for a review). In an attention task, participants were presented with two abstract images, one composed of circles, one of squares. The participants' task was to locate either the circle or square pattern image. In a subsequent evaluation task, participants were presented with an abstract image and were required to rate it on a three-point scale as either 'cheery' or 'dreary'. The to-be-rated images were either previously attended (target) items, previously ignored (distractor) items, or novel items not presented in the attention task. Figure 8 describes the stimuli and procedure of this experiment.



Third Party Material excluded from digitised copy.
Please refer to original text to see this material.

The two rating scales ('cheery' and 'dreary') were used to ensure that participant evaluations reflected evaluative emotional tone rather than a bias to respond to whatever was being held in mind. Ratings from participants who used the negatively valenced scale were reversed and combined with ratings from participants who used the positively valenced scale. The mean ratings across both scales showed that ignored distractors from the previous attention task were significantly devalued (i.e. were liked less) compared to attended targets from the previous attention task, and to novel items not seen in the attention task.

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.



The important finding of Raymond et al. (2003) was that previously ignored images were evaluated more negatively than previously attended and previously unseen images (whose evaluations did not differ; see Figure 9). Attention only impacted ratings of the prior distractors, thus it was the act of ignoring rather than the act of attending that had affective consequences for the exposed item. Visual attention had been shown to alter emotional evaluation to previously neutral images.

To account for their findings, Raymond et al. (2003) proposed the 'devaluation-by-inhibition' hypothesis. They proposed that when a distractor competes for a response, attentional inhibition is applied to the distractor and is stored with the object's representation in memory. This is consistent with the views of Tipper et al. (2003). When

the previously ignored item is subsequently re-encountered, the inhibition is reinstated and, when applied to an evaluative task, leads to an affective devaluation.

The team (Fenske, Raymond & Kunar, 2004) extended the findings into a visual search paradigm. A complex, temporally segmented search task known as preview search (Watson & Humphreys, 1997, 2000) was used. In preview search, the target is defined by feature conjunction (e.g. colour and shape) and each distractor shares one of these features. Critically, the search array is temporally segregated on half of the trials (preview trials). On these trials, a subset of the distractors (the preview set) that share the task-relevant feature with the target are presented 1000 ms prior to the remaining distractors and the target. The remaining distractor (the search set) share the task-irrelevant feature with the target and so, if the previewed distractors are efficiently excluded (inhibited) from the search prior to search set onset, the remaining search task is an easy feature-based search. On no-preview trials, all items in the search array are presented simultaneously, resulting in a difficult conjunction search task. The preview effect (benefit) is where search slopes are flatter in the preview condition compared to the no-preview condition (i.e. the search is more efficient when some distractor items are previewed).

Using a preview search task, Fenske et al. (2004) found more efficient search (faster RTs) on preview trials, indicating that previewed distractors had been efficiently inhibited. In a subsequent evaluation task, previewed distractors were liked less than non-previewed distractors. Consistent with the devaluation-by-inhibition hypothesis, the team found that the RT benefit of the preview set corresponded with a more negative rating of these items.

This result highlighted the problem with straightforward fluency theory in accounting for this effect. Fluency theory was unable to account for distractors being rated

as more negative than novel items. Conventional fluency theory would suggest that, because they had been pre-exposed, distractors ought to be rated more positively than unseen items. More problematic for this theory, is that Fenske et al. (2003) found that distractors which had been displayed for a second longer (preview set) were liked less than distractors which were displayed for less time (search set). Fluency theory predicts that the longer an item is displayed, the more fluent (and so more liked) it will become. This simple orthogonal relationship clearly needs to be revised. That is not to say that fluency has no place: it just needs to accommodate a role for attention, and allow for negativity, not just positivity. Perhaps inhibiting an item causes it to be less fluent, leading to devaluation?

The team (including myself: Raymond, Fenske & Westoby, 2005) then asked several important questions. In Experiment 1 we asked: 1) can devaluation be found in simple search tasks, 2) do devaluation effects vary with set size, 3) would the proximity of distractors to targets impact the size of the devaluation effect? We predicted that we would find devaluation of distractors in simple search, that this would be irrespective of set size, and that distractors originally presented near to the target would be devalued more than distractors far from the target. This latter effect was expected as distractors near to targets are believed to receive more inhibition than far-distractors (Mounts, 2000a, 2000b), because they compete more strongly for responding. It is thought that the visual system adaptively engages a ring of attentional inhibition that surrounds an attended item (centre-surround theory: Mounts, 2000a, 2000b; also Cave & Zimmerman, 1997; Caputo & Guerra, 1998; Cepeda, Cave, Bichot & Kim, 1998; Bahcall & Kowler, 1999; Slotnick, Hopfinger, Klein & Sutter, 2002; Slotnick, Schwarzbach & Yantis, 2003; Cutzu & Tsotsos, 2003).

In Experiment 1 we presented participants with arrays of 4, 8 or 16 unique Mondrian patterns and asked them to locate a target as soon as possible. The target was always

defined by colour alone, and distractors always shared a colour (different to the target). A few seconds later, participants evaluated, on a five-point scale, the just-seen target or one of the distractors. In Experiment 1, the to-be-rated stimuli were presented in the same location that they had occupied in the search array. Figure 10 describes the procedure of Experiment 1.

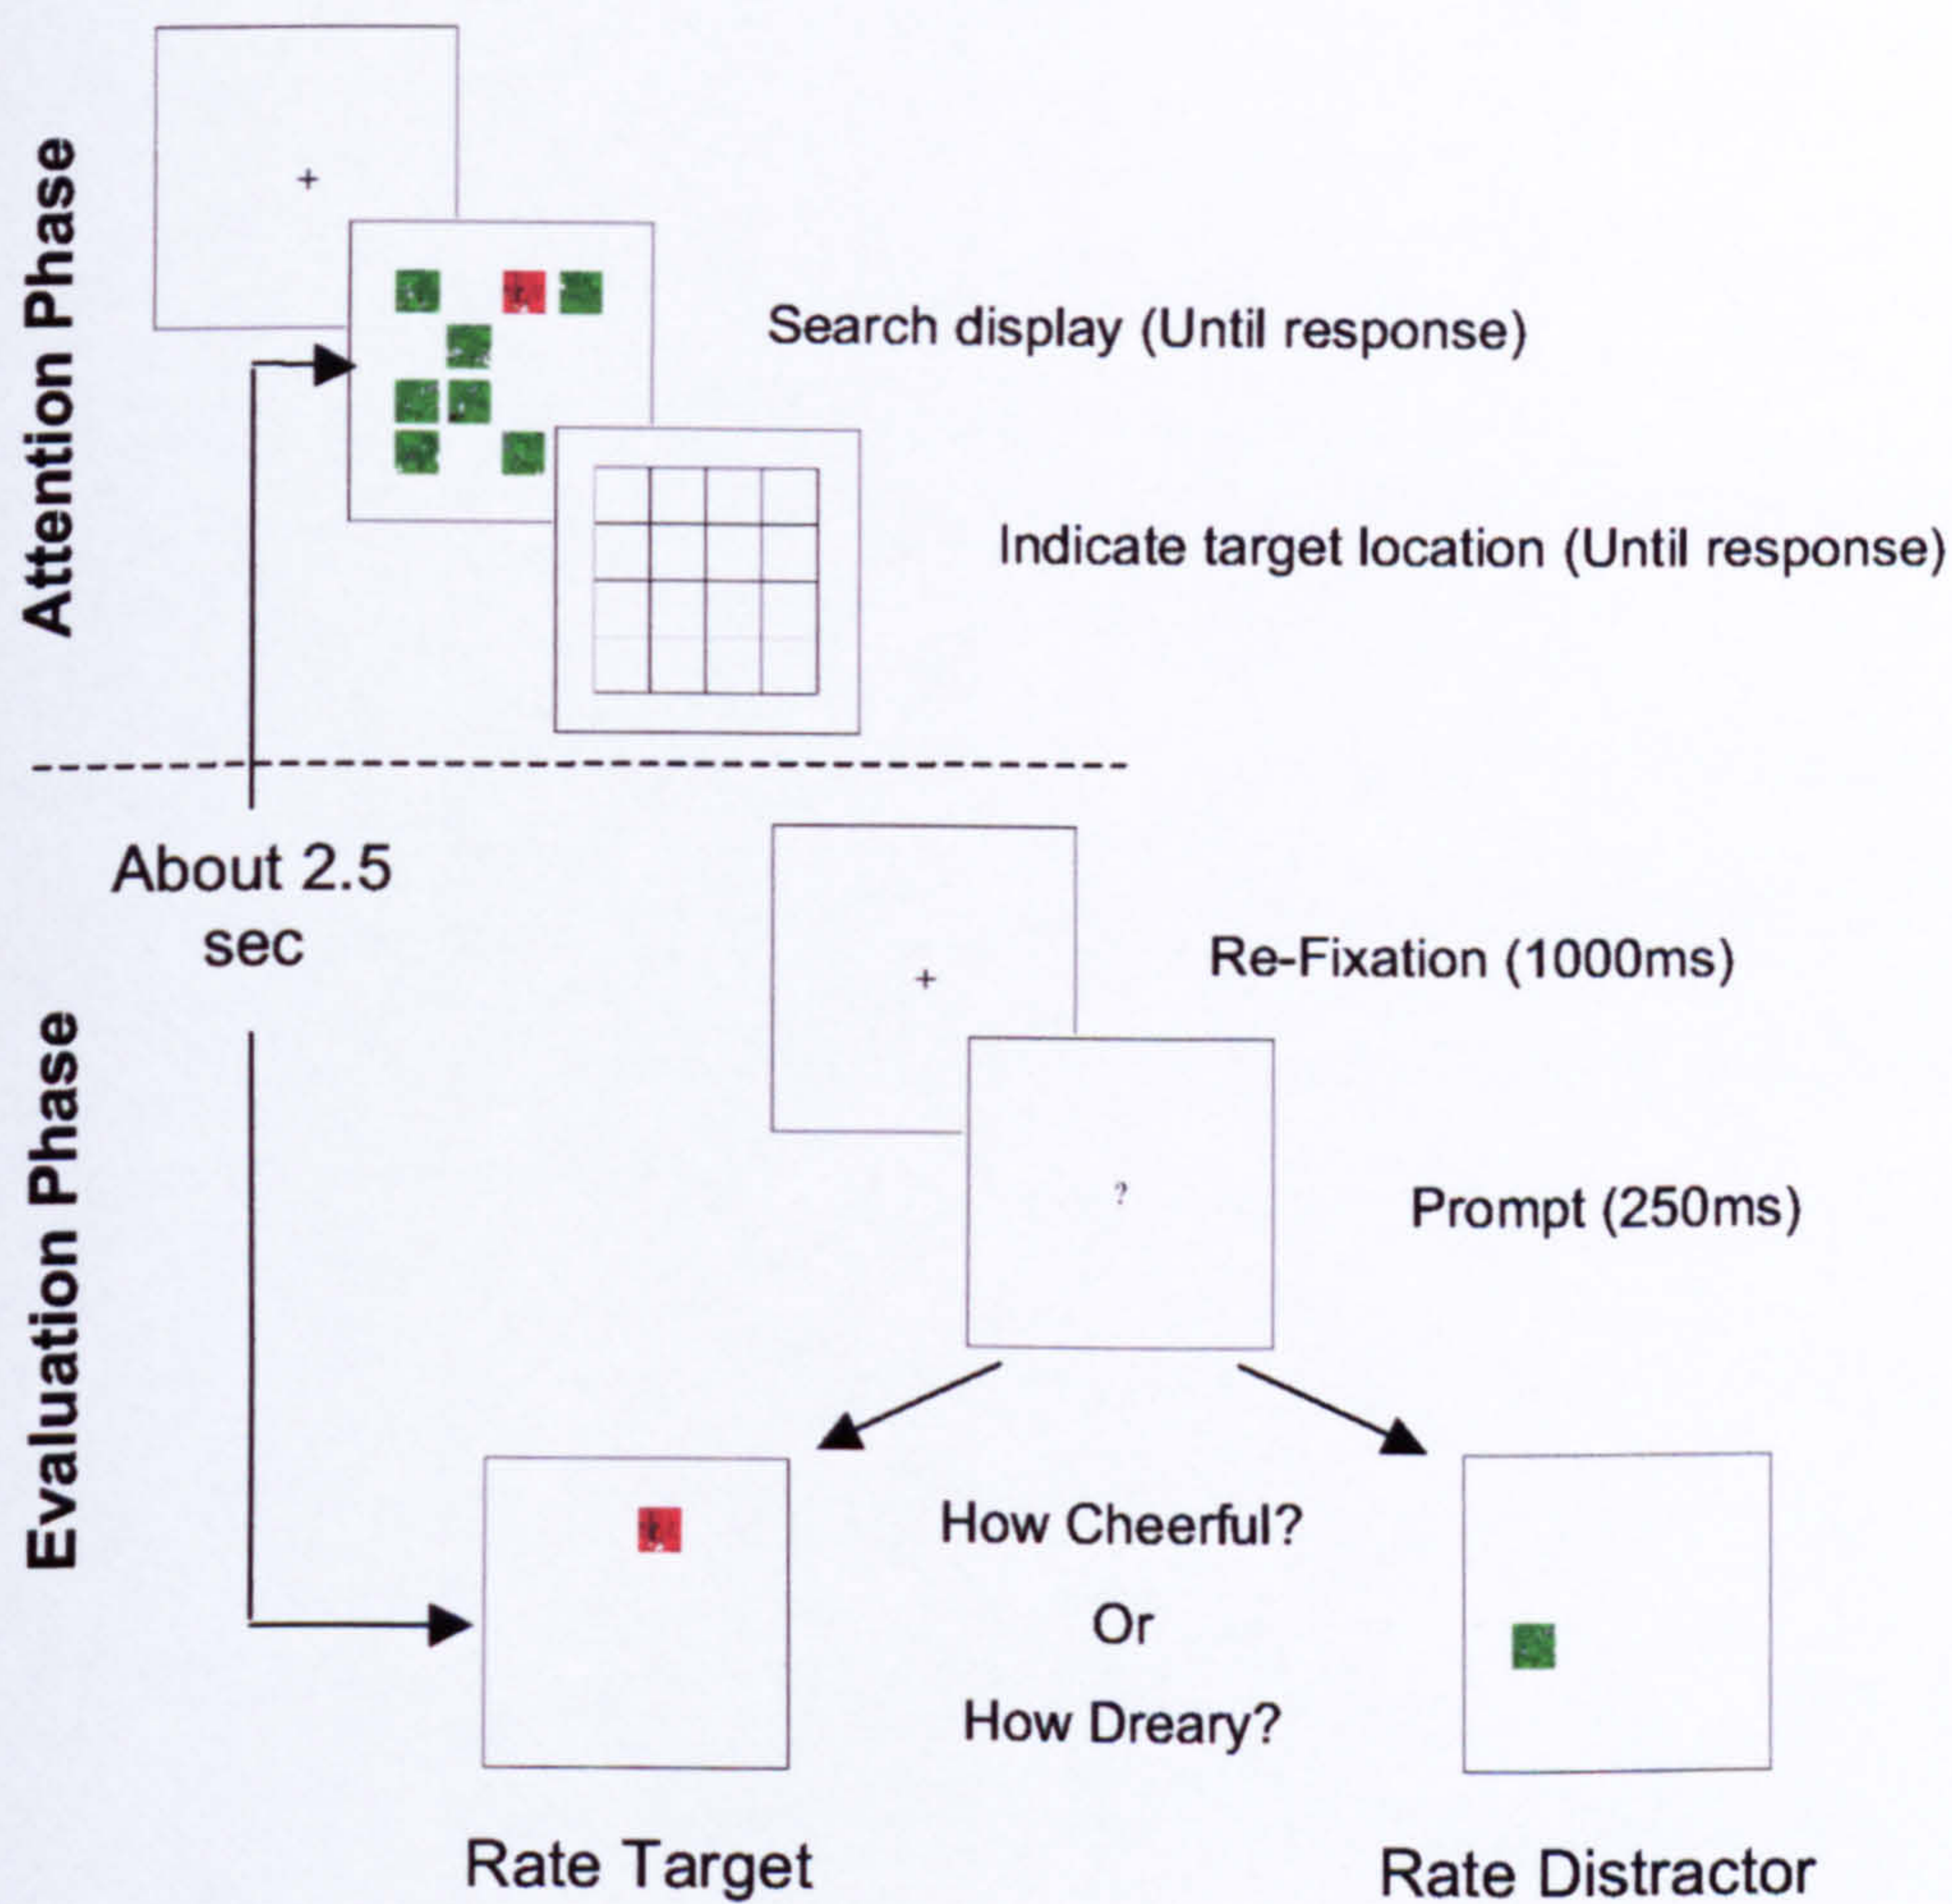


Figure 10: An example of a trial from Experiment 1

The first important finding from Experiment 1 was that distractors were liked less than targets. This effect did not vary with set size. Also, distractors presented near to targets were liked less than distractors presented far from targets (see Figure 11).




Figure 11: Group mean ratings for distractor stimuli just previously seen near or far from the target. Error bars indicate ± 1 standard error of the mean. The hatched area represents ± 1 standard error of the mean target rating.

That we found a devaluation effect is problematic for saliency theories of visual search (i.e. the guided search model), which do not hold a role for inhibition in visual search. Our results suggested that inhibition is indeed applied to distractors during visual search, as revealed by devaluation of distractors relative to targets. Our results also support the centre-surround theory of inhibition surrounding an attended item. Distractors close to the target received more inhibition, and so were devalued more. This effect is also interesting as it suggests that devaluation-by-inhibition is not a binary effect. It is not all-or-nothing: it is graded depending on the amount of inhibition an object receives.

In Experiment 2, we asked another important question: is distractor devaluation location-based or object-based? Because, in Experiment 1, we presented the to-be-rated items in their original location, we were unable to disentangle the relative contributions of location- and object-based inhibitory effects. Therefore, in Experiment 2, we presented the to-be-rated item in the centre of the display for evaluation.

In this experiment, we found no difference in the ratings between targets and distractors. Because the only element that had changed between Experiments 1 and 2 was the location of the presentation of the to-be-rated item, we concluded that the location-based inhibition must have been important for the elicitation of our devaluation effect found in Experiment 1. Object-based inhibition alone was not strong enough to elicit the effect in Experiment 2. This was surprising, as in the Raymond et al. study (2003), the to-be-rated item was also evaluated at a different location than that occupied in the prior attention task, yet a significant distractor devaluation effect was found there. There were two likely explanations for this. First, in the Raymond et al. (2003) study, the distractor was always seen in close proximity to the target, which may have exacerbated the devaluation effect (as seen in our near-far difference in Experiment 1). Second, the to-be-rated item in the Raymond et al. (2003) study was presented very close to its original location (as items were presented at either side of fixation in the attention task). This suggested that the magnitude of the location shift might have been a factor. We examined these two possibilities in our data.

As in Experiment 1, distractors presented near to the target were rated significantly more negatively than distractors presented far from the target. The difference between near and far distractors was in the same order as the difference found in Experiment 1, but in Experiment 1 distractors were rated more negatively overall (see Figure 12).




Figure 12: Group mean ratings for near (black bars) and far (grey bars) distractors obtained for the ‘cheery’ group in each experiment. Error bars indicate ± 1 standard error of the mean. The horizontal dashed line indicates the mean target rating.

In addition to this, we performed a location-shift analysis. We compared evaluations of items that had been presented at peripheral locations in the search array to those that had been presented centrally (which required less of a shift to be evaluated at fixation). As expected, there was an attention by location interaction: the distractor devaluation effect was larger for items requiring less of a location shift.

Taken together, these results suggest that distractor devaluation effects depend on both object- and location-based mechanisms. Although our location-shift analysis suggested that effects are larger when combined with location-based inhibition, it is evident from the near-far distractor analysis that object-based effects also occur.

In Experiment 3, we wanted to determine whether distractor devaluation could be found with more meaningful stimuli, as opposed to the abstract patterns used previously. Second, we wanted to determine whether a slow, effortful search would produce devaluation. Third, we wanted to further investigate the object-based devaluation effect by removing the task-relevant search feature from the stimuli at evaluation.

In Experiment 3, we asked participants to search for and locate a face target defined by a conjunction of gender and tint. In the visual search array, greyscale faces were tinted either blue or yellow: targets were yellow men for half of the participants and blue men for the remaining participants. Distractors were male faces tinted in the non-target colour and female faces tinted in the target colour. Set sizes were reduced to 3, 5 or 7 items. After the search task a previously seen male target or a previously seen male distractor was rated for 'trustworthiness'. We chose to have participants rate this attribute rather than attractiveness as it more clearly specifies that the judgements were to be made about how the participant evaluates the face, not how people in general might evaluate the face. To-be-rated items were presented at fixation, in greyscale without the colour tint. In this experiment, not only was prior location information removed from the to-be-rated item, its previous task-relevant colour information was also removed. This provided a strong test of an object-based devaluation effect.

In Experiment 3, we found a marginally significant difference between targets and distractors: the distractor devaluation effect was elicited for faces evaluated in a different location (see Figure 13). Devaluation effects were found for faces, and data strongly supported object-based distractor devaluation effects. Further to this, finding distractor devaluation after the removal of the task-relevant colour information from the to-be-rated face (and thus changing the object presented for evaluation), supports identity-based devaluation. However, a role for the location-based component of the devaluation effect was preserved in Experiment 3: we again found a location-shift effect (items originally presented in central locations generated a larger devaluation effect than those presented in peripheral locations). Again also, the near-far effect was found: distractors originally

presented near the target were devalued more than distractors originally presented far from the target.




Figure 13: Group mean ratings of trustworthiness of faces seen as targets (white bars), near distractors (black bars), and far distractors (grey bars). Error bars indicate ± 1 standard error of the mean.

The above ‘distractor devaluation’ experiments combined to demonstrate that attention can influence emotion. The interaction between the two systems is not unidirectional. This is in concord with the proposal that an attentional ‘tag’ with an object’s representation in memory would allow effective, efficient affective evaluation to guide swift decision making and behaviour. ‘Devaluation-by-inhibition’ allows items that were previously harmful to performance to be emotionally recognised as such, in order to allow for more efficient behaviour on subsequent tasks involving that item.

Here I begin my thesis. In the three experiments of Fenske et al. (2005), we found that distractor devaluation is not just a location-based effect, it is also an object-based effect. This is key to the series of experiments I present in my thesis. Recall that my primary interest is in establishing whether a generalisation of the distractor devaluation effect can be found to be responsible for effects of stereotyping affect toward an individual to the whole

group that the individual belongs to. In other words, can the 'attentional tag' that is stored with an object's representation also be accessed to represent another object?

Before (but also during) my study of group effects, I explore some of the boundary conditions of the distractor devaluation effect. Several questions were unanswered from the initial studies of the effect reviewed above. I raise the following issues: First, does devaluation require effortful inhibition, or is non-selection sufficient?

I also investigate whether distractor devaluation can result from two different kinds of inhibitory processes. The first is the inhibition applied to the representations of distractor objects during the attentional selection of a target object, as discussed in this chapter. I study this in the first five experiments of this thesis using the oddball search paradigm. The second inhibition type is the deselection of a previously selected item, and the inhibition of the action associated with it. I discuss this in further detail in Chapter 6 and study whether distractor devaluation can result from this type of inhibition using a go / no go task in the final four experiments. Assuming that both processes do result in distractor devaluation, would one type of inhibition be stronger or longer lasting than another, leading to a difference in the distractor devaluation effects?

Also, importantly, does the effect reflect 'devaluation' or is it actually a 'flattening' of affect? Finally, how does devaluation relate to memorial processes, and the ability to encode an item as an individual exemplar?

Chapter 2

An Introduction to, and General Methods for, an Oddball Search Paradigm

The 'distractor devaluation' experiments using visual search and preview search paradigms discussed in Chapter 1 demonstrated that, following exposure in a search array, distractors are rated more negatively than targets. The authors concluded that this effect was due to the attentional state applied to the rated items during exposure in the search task: distractors were rated more negatively than targets because they had been inhibited. But there is an alternative explanation. Exposure effects may explain the difference in ratings between former targets and distractors.

Consider the visual search task of Raymond et al. (2005). Participants were first informed which category the target item would belong to (e.g., a 'blue male'). Participants thus hold a target template that has been pre-defined before the trial begins. On presentation of the search array, participants then search for the item that matches their target template. Once the target is found, the search is ended. In a serial search strategy, on average the target will be found after about half of the items have been searched. Therefore, half of the distractors will not have been looked at. If we suggest a fluency account of the findings, then targets will be liked more than distractors because there is only a 50% chance that a given distractor has been fixated by the participant. Therefore targets will be rated more positively than distractors because there is a greater chance that they have had the opportunity to be processed thereby becoming more fluent.

However a fluency account is unlikely because it cannot explain why 'near' distractors are rated more negatively than 'far' distractors. In fact, a fluency account would predict the opposite pattern of results as distractors presented close to the target are likely to be processed more than distractors presented far from the target.

Another criticism that the visual search studies may face is that distractors, even if fixated and processed, need not be held in memory. If an item does not match the target template (as a distractor does not) it can be rejected. Neither its representation nor the attentional state used to analyse it need be remembered. Even if this amnesic search requires multiple rejections of the same distractor, it can still be forgotten because the target will be stumbled upon eventually. Indeed there is some evidence for this. Recall that in the Raymond et al. (2005) the 'devaluation effect' was larger for items presented in their original location than for items moved to the centre of the display. It would appear that the location of the distractor has been remembered (which would aid in a search strategy), allowing 'distractor devaluation' effects to manifest. However the object representation of the distractor alone does not appear to have been encoded in memory to the same degree: when the object moved location, smaller 'distractor devaluation' effects were found.

Therefore, even if the distractors have the same potential for processing as the target and we do not accept fluency accounts as an explanation, we may have to concede that the results may be a result of a failure to store distractors in memory. I aimed to develop a procedure to address these two issues. Distractors and targets must be equally exposed and processed, and both must demand memory. The use of a pre-defined target template is the critical problem for the traditional visual search paradigm. Not so for the oddball search paradigm.

The first series of experiments I present use variations on the oddball search paradigm. In an oddball search paradigm, there is no pre-defined target template: participants do not know prior to the trial which category the target will belong to. In a three-item search display one item belongs to a different category than the

remaining two items. Once the search array is displayed, each item in the display must be fixated, categorised, remembered and compared against the other two items in order for a participant to decide which item is the 'odd one out'. The oddball search paradigm therefore requires that targets *and* distractors are processed and remembered. A difference in subsequent emotional evaluation between targets and distractors can therefore be confidently attributed to attentional differences during exposure in the search array.

Experiments 1 and 2 aimed to pilot the oddball paradigm and to investigate whether the 'distractor devaluation' effect diminished following a lag. If the effect did not persist following a delay, or intervening trials, then it would not be possible to attribute the effect to the reinstatement of attentional processes encoded with the object representation in long-term memory. It may only be a transient effect that relies on short-term memorial processes. Perhaps the effect allows for emotional mediation of the task at hand, but does not allow for mediation of future tasks, by biasing report selection, and not encoding. Finding 'distractor devaluation' following a lag would support the long-term attentional encoding account of 'devaluation-by-inhibition'.

Experiments 1 and 2 also sought to establish an important theoretical extension to the 'distractor devaluation' effect. Could the effect generalise to everyday objects? Previous experiments have demonstrated the effect using abstract stimuli and faces. These are two specialised sub-sets of objects. The abstract stimuli would only have been encountered by participants in the experimental environment, therefore no object representation would have existed for them prior to the experiment. Faces are also a 'special' kind of stimuli. Humans can process faces extremely efficiently (Farah, Wilson, Drain & Tanaka, 1998) and it has been

suggested that faces can automatically capture attention (White, 1995; Cauquil, Edmonds & Taylor, 2000; Critchley et al., 2000; Lavie, Ro & Russell, 2003, but see Jackson & Raymond, 2006). Face processing has been suggested to occur even when faces are ignored. An N170 (an Evoked-Response Potential component indexing face recognition) can be activated (Liu, Higuchi, Marantz & Kanwisher, 2000; Eimer, 2000) and fusiform face area (FFA) activation elicited (Downing, Liu & Kanwisher, 2001) even when faces are task-irrelevant (but see Pessoa, 2005; Silvert et al., 2007). Extending the findings to everyday objects (in this case animals and bottles) would suggest that the 'distractor devaluation' effect is general, and could be potentially used to unconsciously guide our preferences and behaviours in the 'real world'.

Experiments 1 and 2 required participants to select the target (the 'oddball') from a three-item display and report its type. A display consisted of either faces, animals or bottles. In a face array, the target could be either male or female, and the two distractors were the remaining type. In an animal array, the items were birds and fish, and in a bottle array the items were bottles that were filled with cleaning fluids (cleaning bottles) and bottles that were filled with beverages (drinking bottles). These categories were chosen for the following reasons: faces are an example of objects we are experts at processing; to categorise animal as a bird or a fish is easy; and the bottles are ambiguous. Following this, either immediately or after a lag, one of the items from the array (either the target or one of the distractors) was presented and participants were required to emotionally evaluate the items on a 5-point scale. The critical comparison was the difference in ratings between former targets and former distractors. The General Methods describes this in more detail and forms the basis of

all the experiments using the oddball search paradigm (Experiments 1-5; Chapters 3-4).

General Methods

Participants

All participants were recruited from the participant panel of Bangor University. All had normal or corrected to normal vision and informed consent was obtained.

Apparatus

A Pentium-4 computer, running E-Prime 1.0, recorded data and presented stimuli on a 55.9 cm monitor (100 Hz, 1024 x 768 resolution). The viewing distance was 70 cm.

Stimuli

Stimuli were digital colour photographs taken from the internet and consisted of 1311 images. Images depicted male faces, female faces, birds, fish, bottles of drinking products or bottles of cleaning products (214 images in each category with additional images in each category for practice trials).

Face images were frontal views of adults, with neutral or smiling expressions and visible hair, neck and eyes. Animal images were also full colour images of visually distinct birds and fish. Both face and animal images were presented with their original background (i.e., most pictures of birds depict the bird sitting in a tree). Bottle images were of foreign (to British participants) brands with writing obscured

so it was not possible for participants to read any information on the bottle image.

Bottle images were presented on a uniform white background. Every image subtended 5.7° in height. Face images were 3.4° in width, animal images were 5.7° in width, and bottle image widths varied from 2.3° to 4.0° .

Images were presented in an array of three (see Figure 14). The three images were presented equidistant from fixation (approximately 5° from image centre to fixation), with the 'top' image positioned 5° directly above fixation, and the 'left' and 'right' images positioned below to its left and right.



Figure 14: An example face, animal and bottle array, depicting the arrangement of the images in each array, not to scale.


Procedure

There were two phases to each trial: the attention phase, and the evaluation phase. Participants first saw an instruction screen informing them of the upcoming trial category (faces, animals, or bottles), and what two response keys represented. This instruction screen remained displayed until the participants pressed a key to initiate the attention phase of the trial.

After seeing a 500 ms fixation cross, participants saw three images and were required to name the category of the target (the oddball) using one of two response keys, described later. This is the attention phase.

Following this, the participants saw a blank screen displayed for 500ms, and a second instruction screen was displayed which asked participants to rate the upcoming image on a 5-point scale. This instruction screen introduced a variable, self-controlled interval which makes this experiment different from its predecessors. This experiment is more ecologically valid than previous experiments, accomplished by this relaxed approach to timing.

Once participants pressed a key to initiate the second part of the trial, the evaluation phase began. Following a fixation cross of 500ms, a single image from the attention phase was displayed in the centre of the screen for 250ms and replaced by a “?” which remained displayed until the participants rated the image they had seen. This item will be referred to as the to-be-rated (TBR) item. Note that due to the timing of the second instruction screen, the time between the attention array and the evaluation of the TBR item could vary considerably. See Figure 15 for a single trial example.



Instructions 1: 6 Possible Displays

Face Array	Animal Array	Bottle Array	
What is the sex of the Oddball? Male Female	What animal is the Oddball? Bird Fish	What kind of bottle is the Oddball? Drinking Cleaning	
What is the sex of the Oddball? Female Male	What animal is the Oddball? Fish Bird	What kind of bottle is the Oddball? Cleaning Drinking	Response keys switched.

Figure 15: 1 trial procedure example for Experiments 1 and 2. Each trial begins with an attention phase and is followed by an evaluation stage. The six displays represent the different instruction screens that the participant may encounter during a trial.

Design

In a face category array with two males and one female, the female face is the target and the two male faces are the distractors. Target type was balanced within

category array (e.g. on 50% of face trials, a female face was the target) and presented in a pseudo-random order. The target appeared in each of the positions equally often.

On 50% of the trials the TBR item was the previous target, and on remaining trials it was one of the previous distractors. The position that the TBR item had previously occupied in the oddball search array was counterbalanced.

Response keys were aligned vertically; for half of the subjects, the top key represented 'Male', 'Bird' or 'Drinking Bottle', and the bottom key represented 'Female', 'Fish', or 'Cleaning Bottle'. For the other half of the participants, response key assignments were switched (see Figure 15, page 102). Half of the participants evaluated the TBR item for how affectively 'good' it was (1 = not very good, 5 = very good), and half of the participants evaluated the TBR item for how affectively 'bad' it was (1 = not very bad, 5 = very bad). Participants were instructed to base their evaluations on their first response to the TBR item. Response key assignment and evaluation scale valence were counterbalanced across participants, making 4 groups, each with 25% of participants per experiment assigned to them. The participants were given a sheet of paper with their key assignments and evaluation scale to refer to throughout the experiment.

Each block began with a face category search array, then animal, then bottle and continued in that sequence for 36 trials at which point there was a rest period (see Figure 16). There were four blocks, making 144 trials in total (48 face trials, 48 animal trials and 48 bottle trials). Prior to the main body of the experiment, participants saw 9 practice trials (3 of each type).




Figure 16: The trial sequence within a block. The figure shows only the oddball search array and evaluation displays per trial: instruction, fixation and blank screens are omitted. One face trial was first presented (both search and evaluation), then one animal trial and finally one bottle trial. This sequence was repeated for a total of 36 trials (12 face, 12 animal, and 12 bottle trials). There were four of these blocks in the experiment.

Data Analysis

Evaluations of TBR items from trials in which the target was incorrectly identified, or identified with a RT greater than 6000 ms or less than 300 ms were excluded from further analysis. RTs were averaged for each participant for each condition. Trials in which the RT was greater than 5 standard deviations (SD) of their own average were then excluded. In addition, ratings made more than 3000 ms after the TBR item had been presented were removed.

Chapter 3

Pilot Studies of the Oddball Search Paradigm

Experiment 1: Oddball search pilot study

The aim of Experiment 1 was to provide a primary exploration into the effects of attention in an oddball search paradigm and sought to show that ‘distractor devaluation’ could be found using an oddball search paradigm. The experiment also addressed whether there was an effect of response scale valence, which I relied on in subsequent experiments.

Method

Participants

Eighteen participants (14 females, mean age = 19.0 years) from Bangor University volunteered to take part.

Apparatus and Stimuli

As in the General Methods.

Design and Procedure

The critical difference in this experiment from the General Methods is that the TBR item presented following an oddball search array was not always drawn from that array. In each block, the TBR item was presented at a different lag (lag 0, lag 1, lag 4 or lag 7). Each block contained 12 trials per array category, beginning with a face category array. The block order was counterbalanced across participants.

Lag 0 is described in the General Methods where the TBR item comes from the just seen oddball search (Figure 16, page 101). In a lag 1 block, participants saw

one additional oddball search array before they were required to evaluate the original TBR item (see Figure 17).

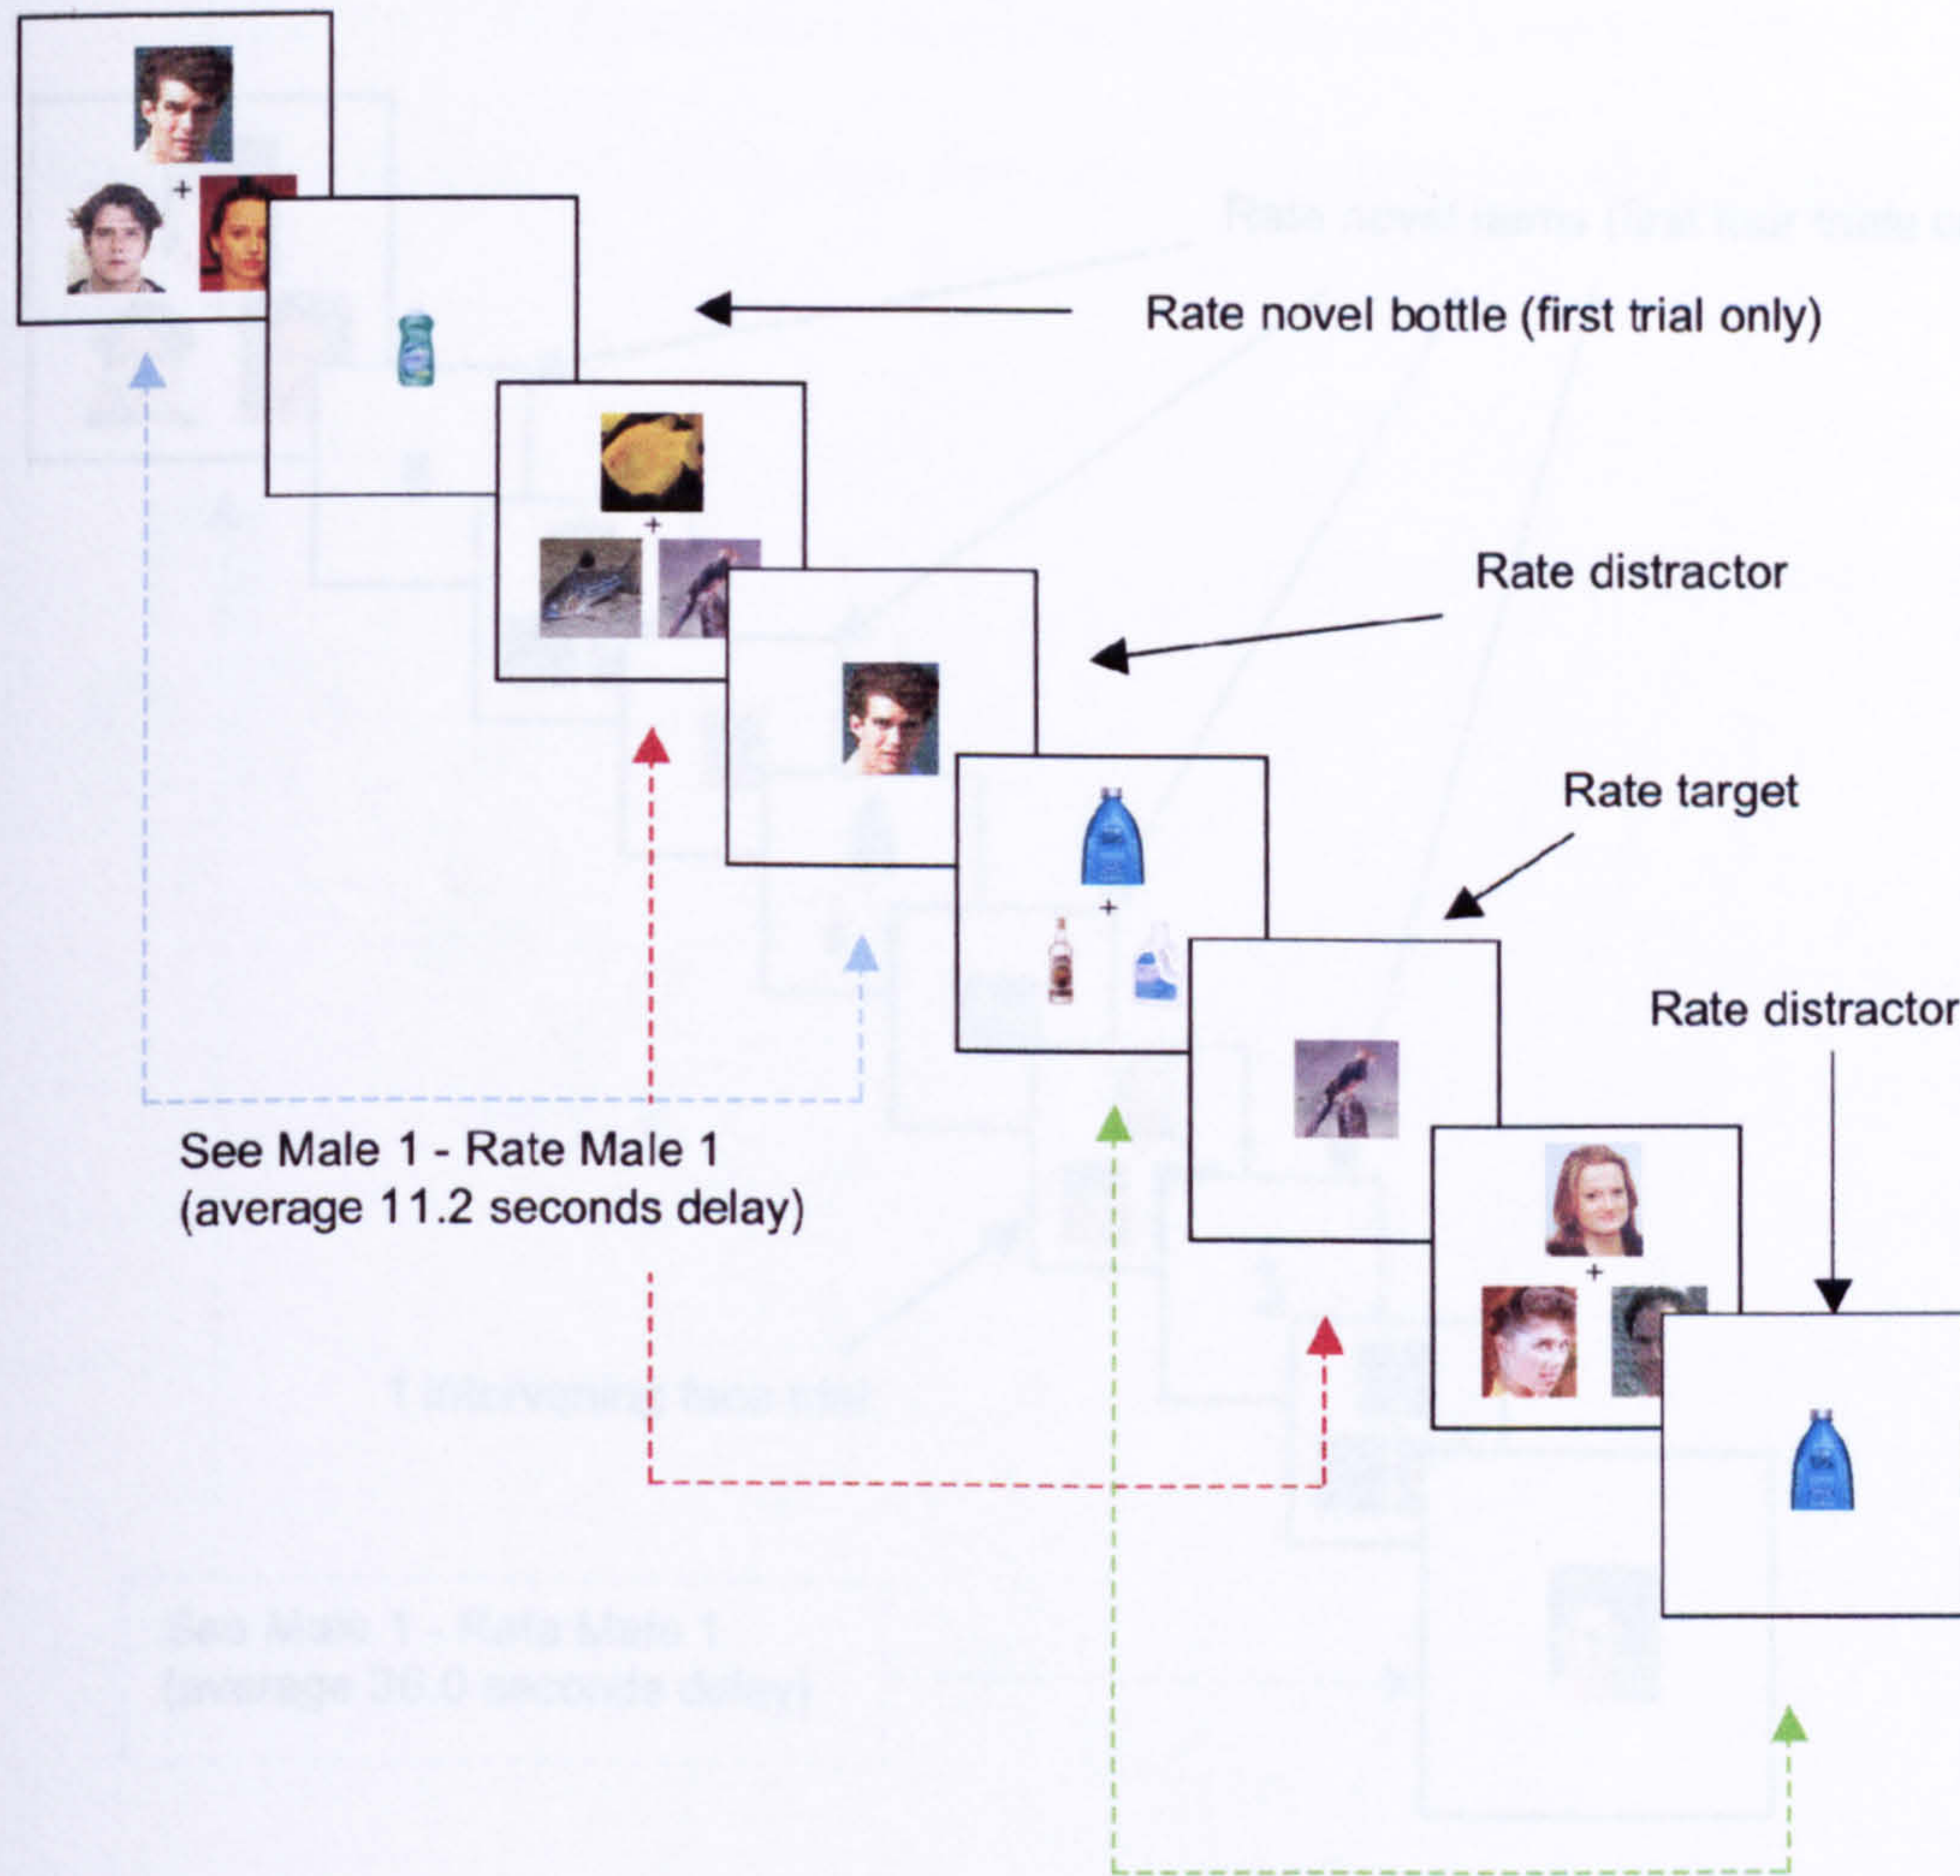


Figure 17: Trial sequence for lag 1 block. The figure shows only the oddball search array and evaluation displays per trial: instruction, fixation and blank screens are omitted. One face trial was first presented (both search and evaluation), then one animal trial and finally one bottle trial. This sequence was repeated for a total of 36 trials (12 face, 12 animal, and 12 bottle trials).

Note that this ‘lagging’ required one additional trial to be added to the end of the block so that the final TBR item in the block could also be seen following an oddball search array. This block thus contained 37 trials. A novel bottle was presented in place of a first TBR item in the first trial and its evaluation was not included in the analysis.

Similarly, the lag 4 block presented the TBR item 4 trials after it was originally seen (Figure 18). This required 4 supplemental trials at the end of the block and 4 unseen rated items at the beginning. This block contained 40 trials.




Figure 18: Trial sequence for lag 4 block. The figure shows only the oddball search array and evaluation displays per trial: instruction, fixation and blank screens are omitted. One face trial was first presented (both search and evaluation), then one animal trial and finally one bottle trial. This sequence was repeated for a total of 36 trials (12 face, 12 animal, and 12 bottle trials).

Note that intervening the exposure and evaluation of a TBR item, there is one additional same category search array and evaluation (in Figure 18 above, two face arrays are shown). The first is the oddball search array that the TBR item shown in the final display comes from, and the second is the intervening same category search array.

The lag 7 block required 7 supplemental trials at the end of the block and 7 novel items to rate at the beginning leading to 43 trials in the block (see Figure 19).

In the lag 7 block, between the exposure in the search array and the subsequent evaluation of a TBR item there are two intervening same category search arrays.




Figure 19: Trial sequence for lag 7 block. The figure shows only the oddball search array and evaluation displays per trial: instruction, fixation and blank screens are omitted. One face trial was first presented (both search and evaluation), then one animal trial and finally one bottle trial. This sequence was repeated for a total of 36 trials (12 face, 12 animal, and 12 bottle trials).

Results

Data Analysis

Ratings from participants using the affectively 'bad' response scale were reversed (a score of 5 becomes a score of 1) to match the affectively 'good' scale.

Removal of 'errors' to detect the target (as defined in the General Methods) excluded 11.63% of evaluations. Removal of 'errors' to rate the TBR item excluded a further 7.67% of evaluations from the analysis.

Finally, two participants were removed because their accuracy to correctly identify the target in either the face, animal or bottle oddball search arrays was lower than 85%. This may seem a high cut-off, but the average accuracy to locate the target was 95.37% with a SD of 3.27%. One participant had used the affectively 'good' response scale, the other had used the affectively 'bad' response scale. The results of Experiment 1 are thus based on 16 participants with an average of 20.44 data points per participant, per category, per attention condition (across four lags).

Oddball Search Performance

After participant removal, the mean accuracy to identify the target in the oddball search array was 96.27% (SD = 2.04%). Accuracy (see Table 1) differed significantly depending on the category of the search array, $F(2, 15) = 13.48, p = .001$. It was harder for participants to correctly identify a target in a bottle array than either a face array, $t(1, 15) = 5.57, p < .001$, or an animal array, $t(1, 15) = 4.48, p < .001$. The accuracy to identify target faces and animals did not differ ($p > .1$).

Table 1: The mean accuracy and SD to identify the target in face, animal and bottle oddball search arrays.

Array Category	Mean	SD
Faces	97.79%	1.93%
Animals	97.53%	2.66%
Bottles	93.27%	3.31%

The accuracy did not differ across lags (i.e. between blocks; $p < 1$), nor did lag interact with array category ($p > .1$).

The mean RT (see Table 2) to correctly identify the target was 2717 ms ($SD = 601$ ms). This did not interact with the category of the array ($p > .1$).

Table 2: The mean RT and SD to identify the target in face, animal and bottle oddball search arrays.

Array Category	Mean (ms)	SD (ms)
Faces	2666	703
Animals	2707	585
Bottles	2780	584

The RT to correctly identify the target also did not interact with the lag ($p > .1$), nor did this interact with array category.

Rating Data

A 2 x 2 x 3 x 4 repeated measures ANOVA with response scale (good, bad) as a between-participant factor and attention (target, distractor), category (face, animal, bottle) and lag (0,1,4,7) as within-participant factors was performed. There were main effects of response scale, $F(1, 15) = 8.30, p = .01$, category, $F(2, 15) = 16.01, p < .001$, and lag, $F(3, 15) = 4.16, p < .05$, but no significant main effect of attention ($p > .1$) and no significant interactions between any factors (all $p > .1$).

Response Scale

There was a between-participant main effect of response scale (see Table 3): participants using the affectively 'bad' scale had rated TBR items 0.44 points more positively on the 5-point scale, $t(1, 15) = 2.97, p = .01$. Critically, the effect of response scale did not interact with any other variable (all $p > .1$), and so response scale was collapsed.

Table 3: Mean ratings (on a 5-point scale) and SDs for all items evaluated by participants using the affectively 'good' and affectively 'bad' response scales.

Group	<i>M</i>	<i>SD</i>
Good Scale	3.02	.35
Bad Scale	3.45	.22

Attention

Disappointingly, the main effect of attention on subsequent ratings was non significant ($p > .1$). Targets ($M = 3.14, SD = 0.47$) were not rated differently than distractors ($M = 3.16, SD = 0.42$).

Lag

There was a significant main effect of lag (see Figure 20). T-tests revealed that TBR items rated just after they were exposed in the search array (at lag 0) were rated as more positive than TBR items rated after a lag of 1 trial, $t(1, 15) = 2.97, p = .01$, or a lag of 4 trials, $t(1, 15) = 2.65, p < .05$. TBR items rated after no lag were also rated more positively than TBR items rated after a lag of 7 trials, but not significantly so ($p > .1$). There were no other significant differences between the lags (all $p > .05$), nor did lag significantly interact with attention ($p > .1$).




Figure 20: Mean ratings for all items at each lag. Error bars represent ± 1 Standard Error (SE) of the mean.

Category

There was a significant main effect of category: bottles were rated less positively than both faces, $t(1, 15) = 5.95, p < .001$, and animals $t(1, 15) = 3.34, p < .005$. Ratings of faces and animals did not significantly differ ($p > .1$; see Table 4). There was no significant interaction of category with attention or lag (both $p > .1$).

Table 4: Mean ratings and SD for faces, animals and bottles.

Category	<i>M</i>	<i>SD</i>
Faces	3.64	0.18
Animals	3.36	0.13
Bottles	2.70	0.12

Each category was then explored separately and a 2 x 2 ANOVA on stimulus type (e.g. males, females) and attention was performed. Consistent with earlier

findings, there was no main effect of attention within any category (all $p > .1$). There was a main effect of stimulus type within each category (see Table 5 for a summary), but this did not interact with attention in any case.

Table 5: Summary of the main effects of stimulus type within category, and the mean ratings and SD per stimulus type.

Stimulus Category	Main Effect of Type	Stimulus Type	<i>M</i>	<i>SD</i>
Face	$F(1, 15) = 21.59, p < .001$	Male	3.33	0.22
		Female	3.96	0.14
Animal	$F(1, 15) = 6.74, p < .05$	Birds	3.51	0.15
		Fish	3.20	0.13
Bottles	$F(1, 15) = 73.84, p < .001$	Drinking	3.49	0.18
		Cleaning	1.90	0.13

Discussion

Experiment 1 provided a primary exploration into the effects of attention in an oddball search paradigm. Results were disappointing. There was no evidence of prior attentional states influencing subsequent emotional evaluations.

There are a number of reasons why I might have found a null result. Maybe distractor devaluation is not a consequence of an oddball search. Or, perhaps too many conditions were creating too much variance in the data, thereby disguising distractor devaluation effects hidden under the surface.

For example, there was a main effect of category (although this did not interact with attention). These category effects may be diluting attention effects. If bottles are rated near the ‘floor’ of the scale for example, there is not much scope for a ‘devalued’ distractor bottle to be rated as lower on the scale.

There was also a main effect of lag on evaluation. Figure 20 suggests that items rated just after they had been seen in the oddball search array are rated more positively than items rated after a lag. This could be thought of as a fluency effect. Items that were just seen may be more fluent than items that have to be retrieved from memory following a lag. That lag did not interact with attention suggests that ‘distractor devaluation’ effects do not survive over time. However, a lack of power in Experiment 1 may simply have meant that I was unable to statistically show a lag difference in attention effects.

Importantly, Experiment 1 found that there was no significant main effect of response scale, and that the valence of the response scale did not interact with the effects of attention. Because of this, in all subsequent experiments participants used only affectively positive response scales. This was done in order to reduce variance between participant ratings that may have been diluting ‘distractor devaluation’ effects in Experiment 1.

Experiment 2: The effect of a lag following an oddball search

In Experiment 1 there were many conditions; this reduced the number of trials per condition and thus reduced the power of any potentially statistical differences. The aim of Experiment 2 was to increase the power of effect sizes and to probe more deeply into the effect of a lag on the attention effect. In Experiment 2 the basic structure of Experiment 1 was preserved, in order to replicate the findings as closely as possible, but the number of lag conditions was reduced to two (lag 0 and lag 7).

Method

Participants

Sixteen participants (12 females, mean age = 21.4 years) from Bangor University volunteered to take part.

Apparatus and Stimuli

As in the General Methods.

Design and Procedure

The procedure of the experiment was almost identical to that of Experiment 2 with the following exceptions. All participants rated items on an affectively 'good' 5-point scale. In this experiment, participants were given the verbal instruction that they should use the whole scale for each type of object they were going to evaluate. They were instructed not to reserve a certain portion of the scale for category A, and another portion of the scale for category B and so on. So, female faces should be rated on a scale of 1 to 5, as should male faces and cleaning bottles etc.

In Experiment 2, only two lag blocks were used: lags 0 and 7. Participants saw two blocks of each lag. The lag blocks were presented together, i.e. participants saw both blocks of one lag, then both blocks of the other lag. Block order was counterbalanced across participants.

Note that in a lag 7 block, the same category search trial preceding the TBR item followed could be matched or mismatched to the original oddball trial. For example, in Figure 19 on page 106 the TBR item is a male face. When this male was

originally seen in the search array it was a distractor. However, in the face search array that precedes the rating of this male, a male face is a target. Is this male evaluated based on the attentional state when it was pre-exposed in the original oddball search array (original target, OT, or original distractor, OD), or does the rating depend on whether it belongs to a class of stimuli that has just been a target or distractor (just previous target, JPT, or just previous distractor, JPD)?

Results

Data Analysis

Removal of 'errors' to detect the target (as defined in the General Methods) excluded 14.46% of evaluations. Removal of 'errors' to rate the TBR item excluded a further 7.67% of evaluations from the analysis.

The results of Experiment 2 are thus based on 16 participants with an average of 20.53 data points per participant, per category, per attention condition (across two lags).

Oddball Search Performance

The mean accuracy to identify the target in the oddball search array was 92.36% ($SD = 4.94\%$). The accuracy differed significantly depending on the category of the items in the search array (Table 6), $F(2, 15) = 21.57, p < .001$. It was harder for participants to correctly identify a target in a bottle array than either a face array, $t(1, 15) = 5.42, p < .001$, or an animal array, $t(1, 15) = 5.29, p < .001$. The accuracy to identify target faces and animals did not differ ($p > .1$).

Table 6: The mean accuracy and SD to identify the target in face, animal and bottle search arrays.

Array Category	<i>M</i>	<i>SD</i>
Faces	95.31%	3.44%
Animals	94.01%	6.22%
Bottles	87.76%	6.76%

The mean RT to correctly identify the target was 2871 ms ($SD = 558$ ms).

There was a significant main effect of category on RT (Table 7), $F(2, 15) = 23.79$, $p < .001$. Participants were slower to correctly identify the target in a bottle category search array than in either a face category, $t(1, 15) = 6.75$, $p < .001$, or an animal category $t(1, 15) = 5.20$, $p < .001$, search array. RT for face and animal oddball search arrays did not differ ($p > .1$).

Table 7: The mean RT and SD to identify the target in face, animal and bottle oddball search arrays.

Array Category	<i>M</i> (ms)	<i>SD</i> (ms)
Faces	2658	646
Animals	2776	578
Bottles	3231	561

Rating Data

A 2 x 2 x 3 repeated measures ANOVA with attention (target, distractor), lag (0, 7) and category (face, animal, bottle) as within-participant factors was conducted. The ANOVA revealed a marginally significant main effect of attention, $F(1, 15) = 3.93$, $p = .066$, but non significant main effects of lag and category (both $p > .1$). Lag, however, interacted with both attention, $F(1, 15) = 4.80$, $p < .05$, and category, $F(1,$

15) = 6.93, $p < .01$. There were no other statistically significant interactions (all $p > .1$)

Attention Effects

The main effect of attention on subsequent ratings was marginally significant: distractors ($M = 3.33$, $SD = 0.09$) were rated as more negative than targets ($M = 3.43$, $SD = 0.10$). The attention effect interacted with lag: distractors were rated more negatively than targets at both lags 0 and 7, but this difference was only significant at lag 0, $t(1, 15) = 2.69$, $p < .05$; lag 7: $p > .1$ (see Figure 21).




Figure 21: The effect of attention at lags 0 and 7. Error bars represent ± 1 SE of the mean.

The above results describe the results of original targets and distractors (OT and OD). For the lag 7 blocks only, the attention effect was investigated for just previous targets and distractors (JPT and JPD). TBR items following trials in which

that type of item was a target (JPT) were compared to items that followed trials in which that type of item was a distractor (JPD). JPTs were originally seen as targets on 50% of trials, and originally seen as distractors on 50% of trials. Likewise JPDs were originally seen as targets on 50% of trials, and originally seen as distractors on 50% of trials.

There was no significant main effect attention between JPTs or JPDs, nor did this interact with the category (both $p > .1$).

Category Effects

The main effect of category was non significant: faces, animals and bottles were rated as equally positive. Separate ANOVAs comparing the stimulus types within each category revealed that only males and females were rated differently, $F(1, 15) = 26.67, p < .001$. Females were rated more positively than males, $t(1, 15) = 5.17, p < .001$. Birds and fish, and cleaning and drinking bottles were not rated differently (both $p > .1$).

Category effects did interact with lag however (see Figure 22). Faces were rated more positively when they were rated after a lag, $t(1, 15) = 3.17, p < .01$. Conversely, bottles were rated more positively when they were rated immediately following exposure in the oddball search array, $t(1, 15) = 3.29, p < .01$. Animals were rated equally after a long lag as at no lag ($p > .1$).

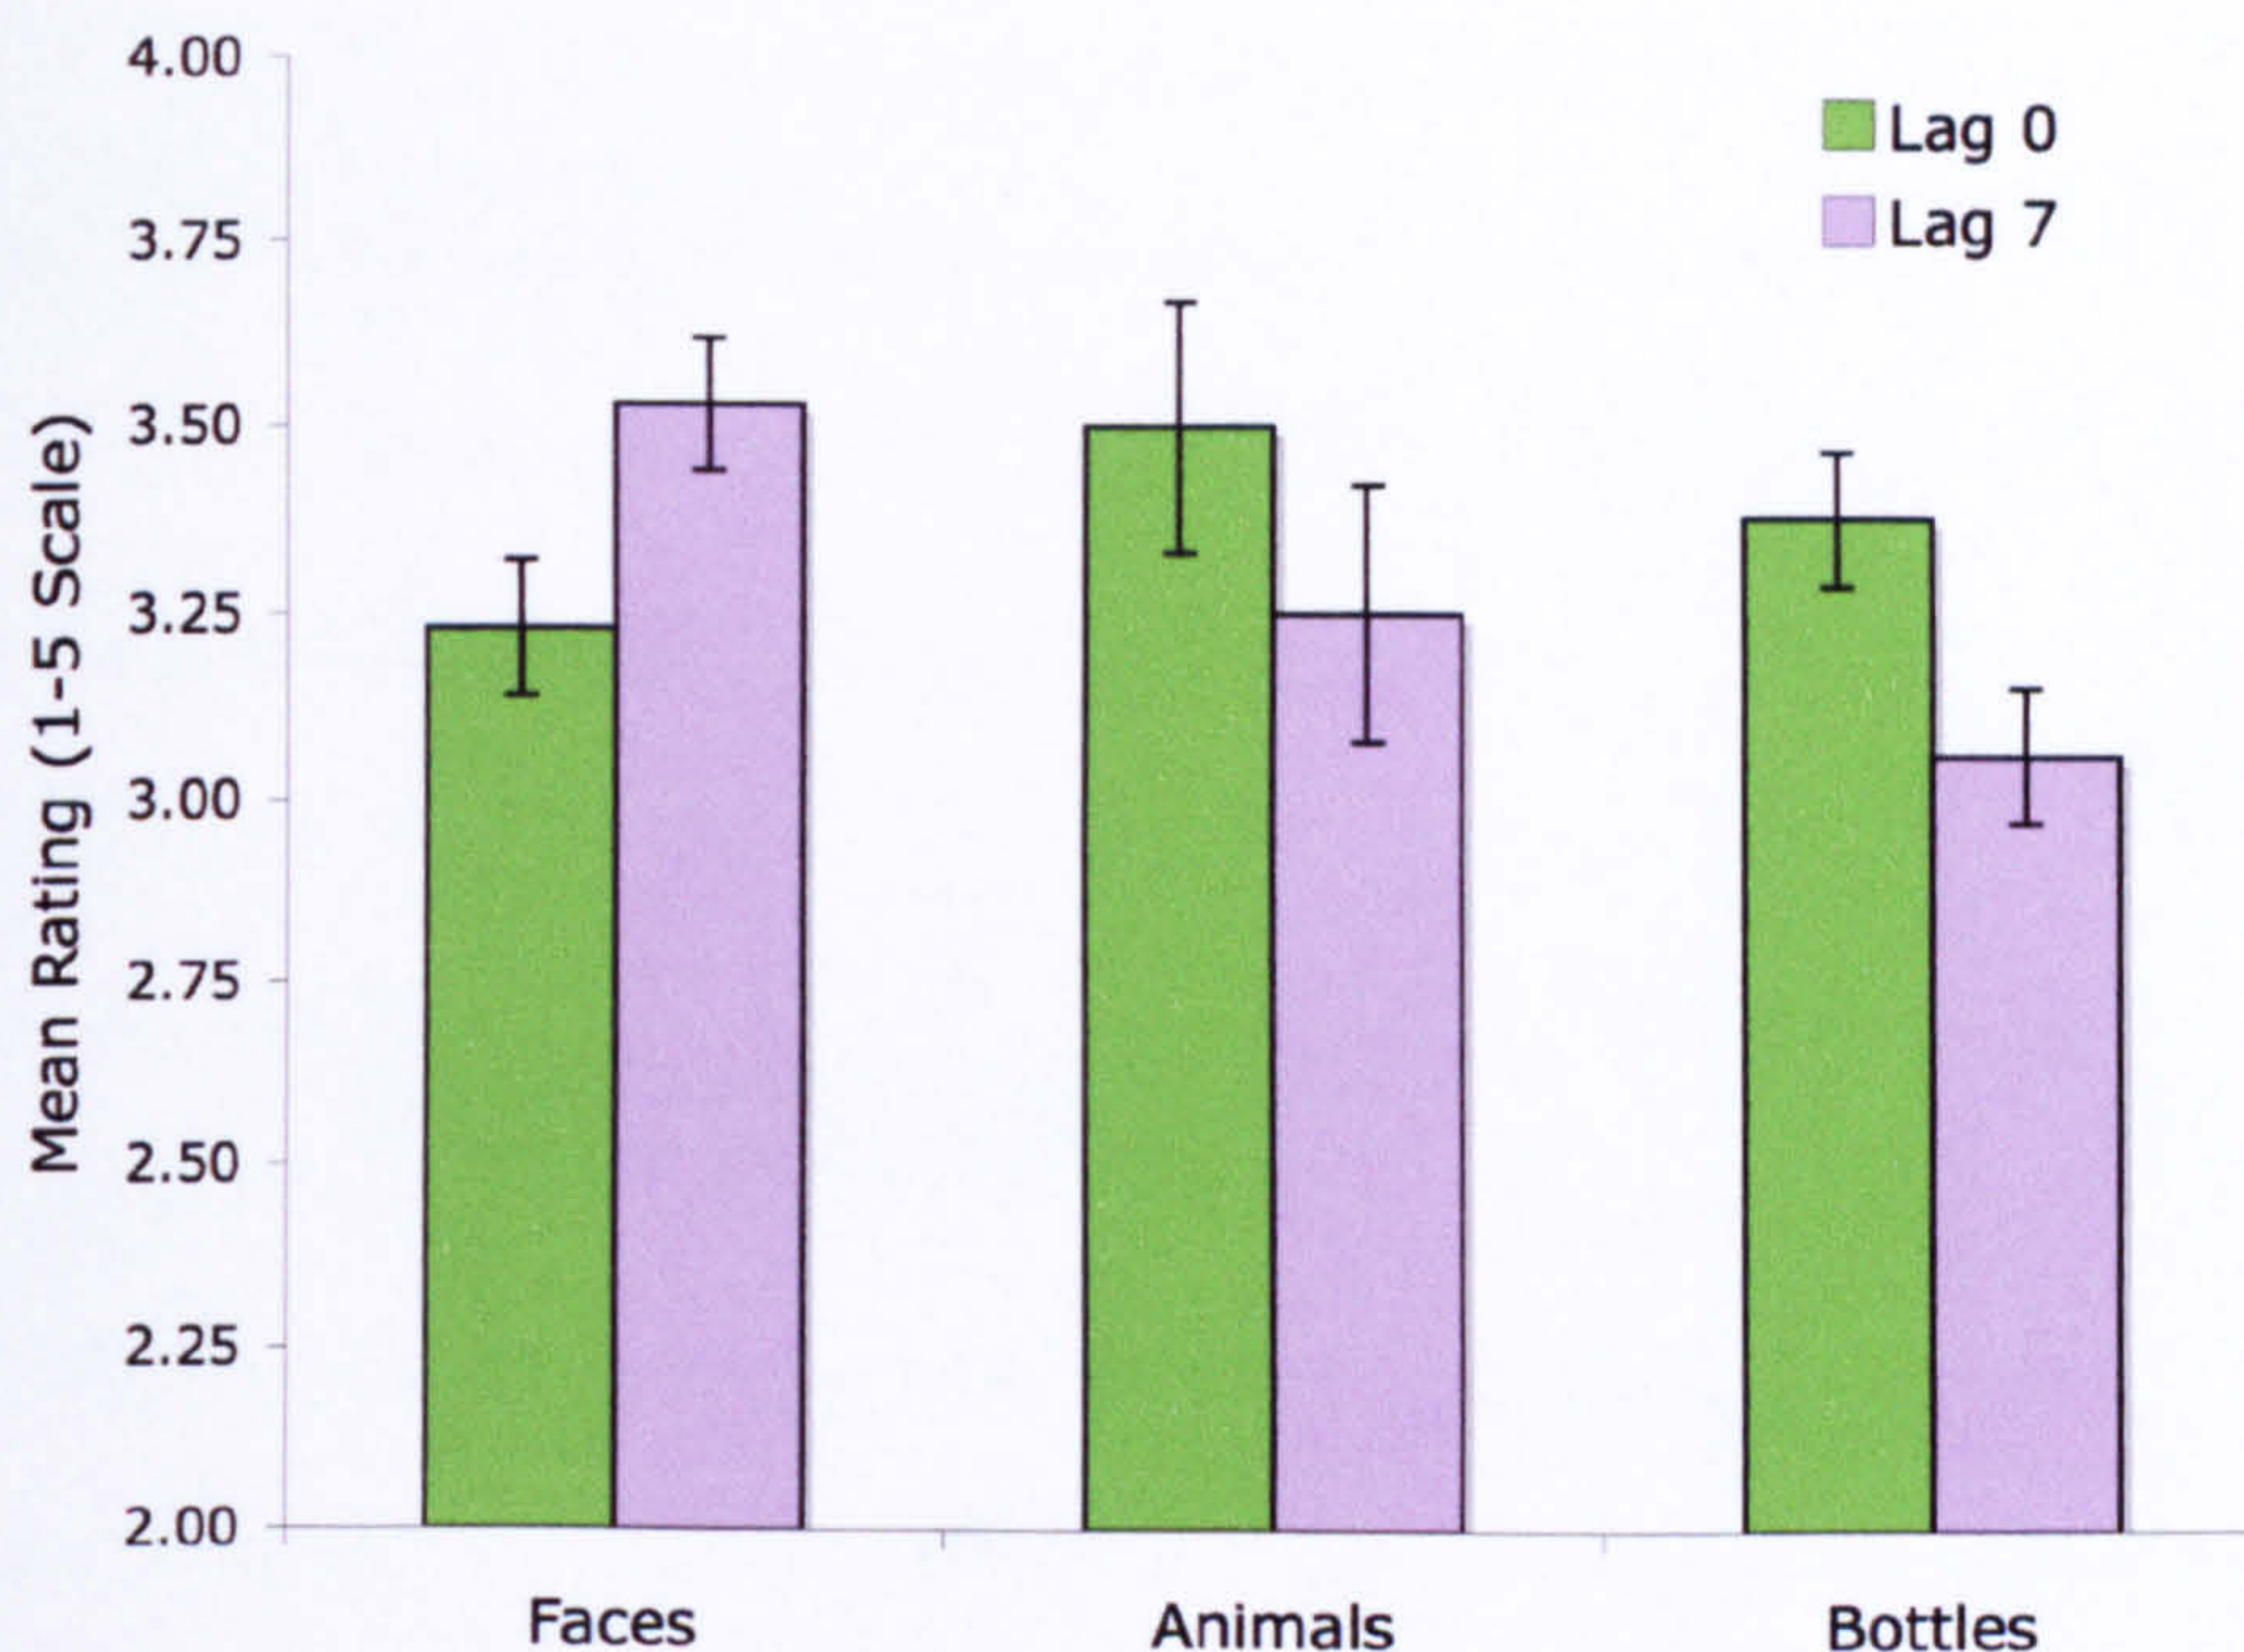


Figure 22: Ratings of faces, animals and bottles at lags 0 and 7. Error bars represent ± 1 SE of the mean.

Discussion

Experiment 2 found distractor devaluation effects where Experiment 1 failed. Distractor devaluation can be found using an oddball search paradigm. The results of Experiment 2 suggest that the strength of the distractor devaluation effect may weaken after a lag. At a long lag (in this experiment after approximately one minute) it disappeared.

Type attention effects were not found in Experiment 2. The evaluation of an item did not depend on whether it followed a same category trial in which the item was target-like or distractor-like.

Interestingly, the different categories appeared to reveal different mere exposure-like effects. Both animals and bottles were evaluated as more positive after a short lag (compared to a long lag), but faces were preferred after a long lag (compared to the short lag). This results suggests that the fluency trace of animals and bottles decays relatively quickly, but it is preserved for longer for faces. This may be due to our expertise with faces, and the evolutionary criticality of remembering faces. If a face has been encoded strongly then the fluency to that representation might be longer lasting than an animal (for example), whose fluency trace may fade quickly as the representation was weak. However why mere exposure should increase for a face after a long lag is puzzling.

Chapter 3 Discussion

In Experiment 2 I have demonstrated that 'distractor devaluation' can be found following an oddball search paradigm. The criticisms of the previous distractor devaluation experiments using traditional visual search paradigms have, in my opinion, been laid to rest. The distractor devaluation effects cannot be attributed to a distractor simply not having been fixated or analysed by the participants. In order to correctly identify the target in an oddball search array, participants must categorise each object in the oddball search array. Therefore effects cannot be explained by greater exposure of the target relative to distractors. Equally, category information of each object in the oddball search array must be held in working memory in order to respond correctly. Therefore distractor devaluation effects cannot be explained by an immediate 'forgetting' of distractors. The results of Experiment 2 lend further support to the 'devaluation-by-inhibition' hypothesis. The difference in ratings of targets and

distractors can be attributed to the attention directed towards objects during pre-exposure.

Experiment 2 provides the first evidence that these effects may generalise to everyday objects. Previous experiments used abstract art images and faces as stimuli. Here, I have replicated the findings of Raymond et al. (2005; Experiment 3) by demonstrating 'distractor devaluation' effects with faces, but I have also extended the effects to include animals, and even inanimate objects: bottles.

Experiments 1 and 2 provide a somewhat unclear picture as to the effect of a lag on 'distractor devaluation' effects. In Experiment 2, an ANOVA did show an interaction of attention with lag. As can be seen in Figure 21 on page 116 and the accompanying t-tests distractor devaluation was present at lag 0, but had disappeared by lag 7. Taken at face value, these results suggest that 'distractor devaluation' is a short-term effect that does not persist over time. However a closer inspection of the lag 7 attention effects in Figure 21 show that the attention effect (though non-significant) is trending in the direction of distractor devaluation. Perhaps there was not enough power in Experiments 1 and 2 to demonstrate 'distractor devaluation' at longer lags to a statistically significant level.

One variable that may have been particularly diluting the effects in Experiments 1 and 2 is the category of the item being evaluated. Target bottles were identified more slowly and with less accuracy than either face targets or animal targets. This suggests that drinking and cleaning bottles were harder to discriminate between than male and female faces, or birds and fish. Further, Experiment 1 showed a large main effect of category in the ratings. In Experiment 2 I managed to eliminate

significant main effects of category by encouraging the participants to use the whole scale to rate every category of object. However, there was still a significant difference in the ratings of male and female faces. Category also significantly interacted with lag. Faces were rated as more positive after a lag, whereas animals and bottles were rated as more positive when rated just after pre-exposure. This effect is curious, and may again be due to noise in the data. While the effect did not significantly interact with the primary variable of interest (i.e. attention) in Experiments 1 and 2 this, and the other differences between faces, animals and bottles, suggest that ratings of categories may be behaving in different ways. Investigating the attention effects in categories separately may shed light on some interesting variations and provide explanations for some of the data found in Experiments 1 and 2.

Chapter 4

Token Effects on Evaluation in an Oddball Search Paradigm

My aim in the next three experiments was to explore further the findings of Experiments 1 and 2. In particular, in Experiments 1 and 2 the categories being rated appeared to be behaving in different ways. For example the accuracy to identify the target in a bottle search array was lower than in face or animal arrays, and it also took longer. Further, in the rating data, faces and animals behaved differently in terms of main effects and also in how they interacted with lag.

I hypothesised that much of the difference in attention effects was due to the ability to which a stimulus could be individuated. If an object is encoded as an individuated item (a token), then attentional states active during search could potentially be encoded with that object's representation in memory. These could then be reinstated and 'devaluation-by-inhibition' could occur. If, on the other hand, an object is not individuated but is only encoded as an occurrence of a category (a type), then attentional states may not be encoded with the individual object representation and therefore item specific distractor devaluation may not emerge.

The aims of Experiment 3 were: first, to demonstrate that distractor devaluation occurs when the inhibited object has been successfully individuated and second, to investigate whether token distractor devaluation effects can persist following a lag.

Experiment 3: Token distractor devaluation effects

In order to investigate distractor devaluation effects of tokenised stimuli, in Experiment 3 I used participant ratings of bottle stimuli only. In order to categorise a

bottle as either a drinking or cleaning bottle, an analysis of bottle shape, colour, design and function must typically occur. Therefore categorisation requires individuation of the bottle.

In Experiments 1 and 2, accuracy to identify the target in bottle search arrays was low (compared to face and animal target identification accuracy). This was accompanied by longer RTs to identify the bottle target, compared to faces and animals. Thus, this category was more difficult, and we can postulate that effortful inhibition was required in order to suppress responding to the distractor bottles in favour of the target bottle.

I predicted that this category of stimuli would be successfully individuated and therefore strong distractor devaluation effects would be found. I also predicted that these effects would persist. Half of the participants rated the bottles after a short lag (about 11 seconds), and half of the participants rated the bottles after a long lag (about a minute). I chose to make this a between participant design in order to maximise the number of trials within a participant to each lag in order to reduce the variance in the data and increase the power of the effects.

Another way in which I reduced the variance was to have participants rate only one type of bottle. Half of the participants therefore rated drinking bottles, and half rated cleaning bottles. In this way, each participant would use the whole scale to rate their type of stimulus, thereby increasing the potential for distractor devaluation effects to emerge.

Participants in Experiment 3 were also asked to rate 'novel' bottles (i.e. drinking or cleaning bottles that had not been exposed in the oddball search array). These succeeded a trial in which that type of item had been either a target or a distractor (nJPT or nJPD). I hypothesised that if old (exposed in the search array) bottles were successfully individuated, then attentional encoding would be mapped onto token object representations. Therefore, I did not expect the attentional state directed towards an old stimulus to have implications for a subsequent novel stimulus of the same type. Therefore no attentional effects on the ratings of novel bottles was expected.

Method

Participants

Forty participants (30 females, mean age = 20.13 years) from Bangor University volunteered to take part.

Apparatus

As described in the General Methods.

Stimuli

The bottle and animal stimuli described in the General Methods were used in this experiment. The animal stimuli were used in a filler search array trial only, and participants were never asked to affectively evaluate animals.

Procedure

The critical differences between Experiment 3 and the previous experiments are participants were only required to rate one type of stimulus (e.g. drinking bottles) and, participants were required to rate items that they had not seen in a search array.

For a participant rating drinking bottles, following a bottle search array, the participant saw a novel drinking bottle (novel TBR item) and was required to rate it. A filler search array of animals was presented and then the drinking bottle seen in the original search array ('old' TBR item) was rated (a lag of 1, see Figure 23).

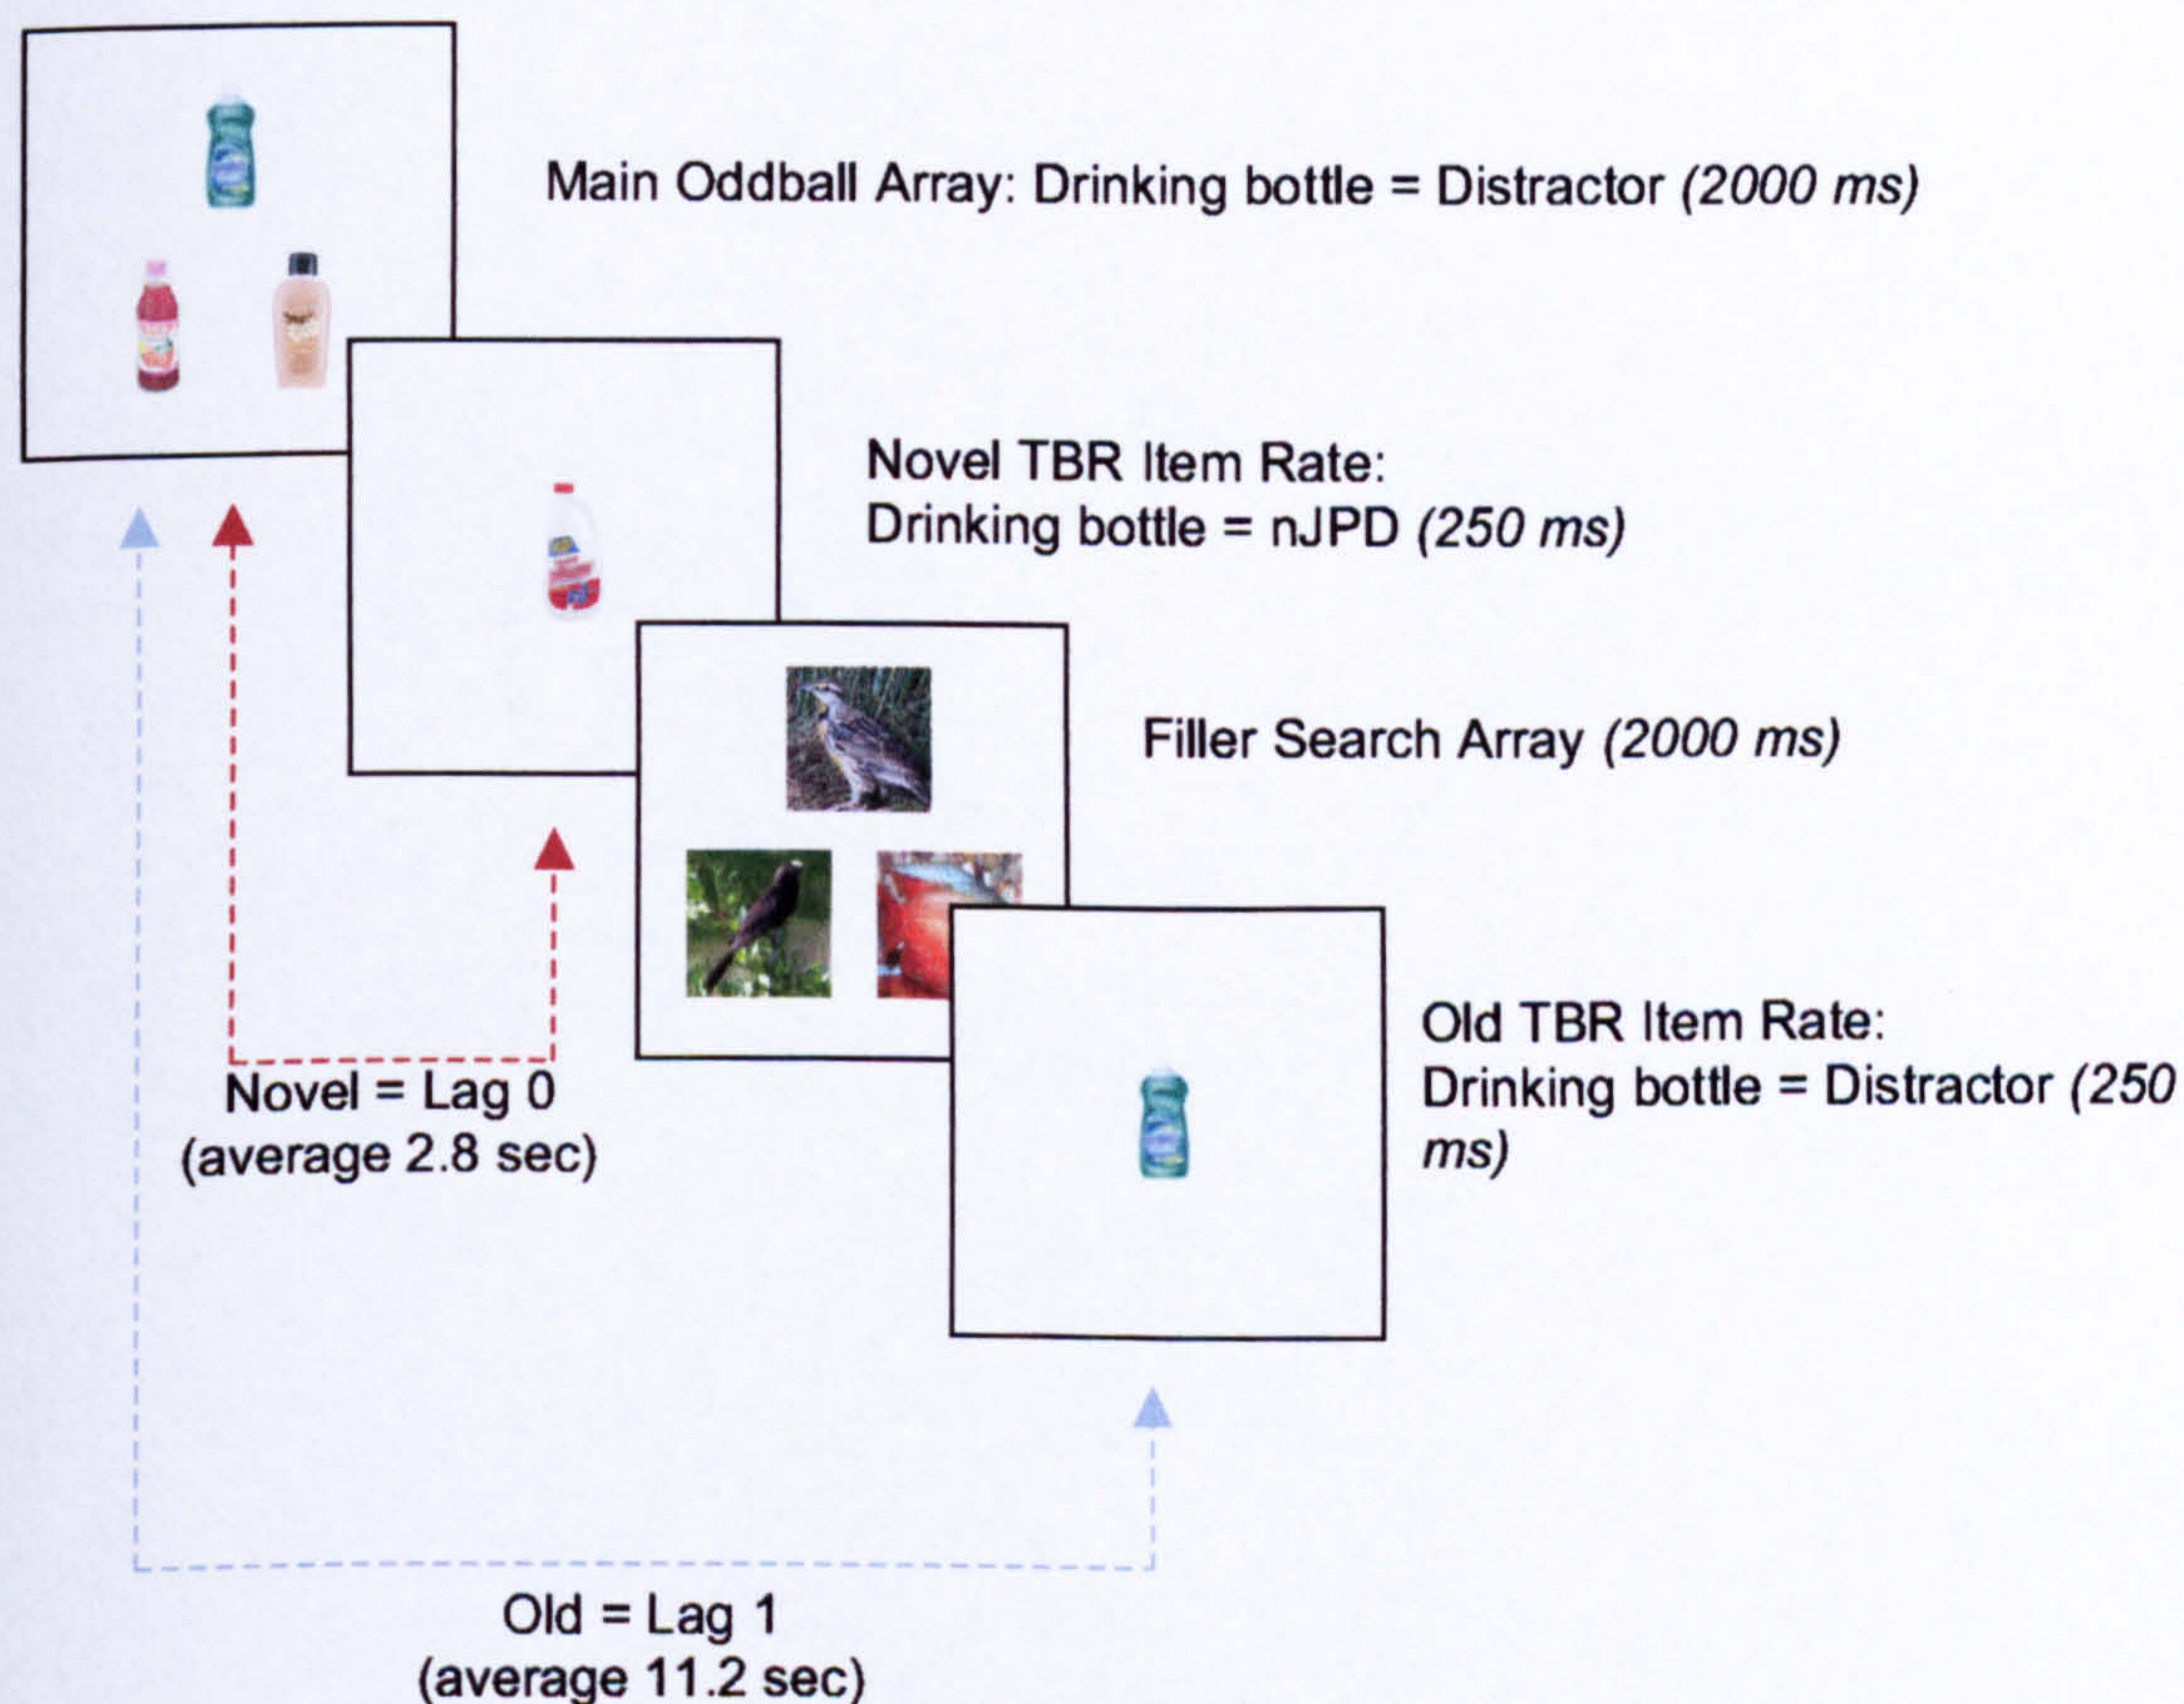


Figure 23: Trial sequence for a participant rating drinking bottles at lag 1 (60 trials).

Half of the participants saw the old TBR item at lag 1 and half saw it at lag 7 (see Figure 24). A novel TBR item was always presented at lag 0. Table 8 describes the participant groups.




Figure 24: Trial sequence for a participant rating drinking bottles at lag 7 (63 trials).

All participants rated items on an affectively 'good' 5-point scale. Prior to the main body of the experiment, participants saw 6 practice trials to ensure satisfactory performance.

Design

50% of the participants rated drinking bottles, and 50% rated cleaning bottles. Of the participants who were rating drinking bottles, half rated the drinking bottles pre-exposed in the search array after a lag of 1 trial, and half rated them after a lag of 7 trials.

Table 8: Participant groups.

	Rate Drinking Bottles	Rate Cleaning Bottles
Lag 1	n = 10	n = 10
Lag 7	n = 10	n = 10

The lag 1 condition contained 60 trials (60 bottle search arrays, 60 novel TBR bottle rates, 60 animal (filler) search arrays and 60 old TBR bottle rates). Of the novel TBR bottles, half were target-like and half were distractor-like. Of the old TBR bottles, half were originally seen as targets (OT) and half were seen as distractors (OD).

As with previous experiments, a lag 7 condition required additional trials. In order to make up the 60 old TBR bottle rates and 60 novel TBR bottle rates, 6 supplemental items were required for rating at the beginning of the experiment (rates from these items were not included in the analysis). Also, 3 extra bottle search arrays and 3 extra animal search arrays were required at the end of the block.

Note that, for participants at lag 7 only, the bottle search array that the old TBR bottle followed could be matched or mismatched to the bottle oddball search array which the old TBR bottle comes from. For example, in the trial presented in Figure 24 on page 127, the TBR item is a old drinking bottle. When this drinking bottle was originally seen in a search array it was a distractor (OD). However, in the bottle search array just previous to the rating of this drinking bottle, a drinking bottle is a target (JPT). Is this drinking bottle rated based on the attentional state directed towards it when pre-exposed in the original search array (OT or OD), or does the rating depend on whether it belongs to a class of stimuli that has just been a target or distractor (JPT or JPD)?

The novel TBR bottle (for participants at both lags) could follow a bottle search array in which that type of item was a target, or a distractor (novel just-previous target, nJPT, or novel just-previous distractor, nJPD). Half of the novel TBR bottles were nJPT and half were nJPD.

Results

Data Analysis

Removal of 'errors' to detect the target excluded 26.38% of evaluations. Removal of 'errors' to rate the TBR item excluded a further 5.23% of evaluations from the analysis. Finally, evaluations of old TBR bottles were only included if they followed filler search trials in which the target animal was correctly identified. This excluded a further 8.95% of the old TBR evaluations from the analysis.

Two participants were removed from the analysis. One participant was removed because his accuracy to correctly identify the target bottle was only 55%. The second participant was removed because he was not using the rating scale correctly. Of the two participants who were removed, one had rated old TBR items at lag 1, the other had rated old TBR items at lag 7. The results of Experiment 3 are thus based on 38 participants with an average of 42.37 old TBR item evaluations, and 42.29 novel TBR item evaluations.

Search Performance

The mean accuracy to identify the target in the bottle search array was 83.99% ($SD = 7.89\%$). The mean accuracy to identify the target in the filler trials, the animal search array, was 91.27% ($SD = 6.86\%$). The mean RT to correctly identify a target bottle was 3260 ms ($SD = 528$ ms). Neither the bottle or animal search accuracy, nor the RT to identify the bottle target differed between the 'lag 1' and 'lag 7' participant groups (all $p > .05$).

Rating Data

A 2 x 2 x 2 repeated measures ANOVA with exposure (old, novel) and attention (target, distractor) as within-group factors and lag (1, 7) as a between-group factor was first conducted. The ANOVA revealed no significant main effects, and only one marginal interaction: a three-way interaction of exposure, lag and attention, $F(2, 37) = 3.32, p = .08$ (all other p 's $> .05$). Ratings of old and novel bottles were next considered separately.

Old TBR Item Evaluations

Evaluations of old TBR bottles were considered in a repeated measures ANOVA with attention (target, distractor) as a within-group factor and lag (1, 7) as a between-group factor. The main effect of attention on subsequent ratings was non significant ($p > .1$). However, a marginally significant interaction of attention with lag, $F(1, 37) = 3.29, p = .08$, prompted further analysis. When attention effects were examined within each lag group, significant effects emerged. Importantly, distractors were rated more negatively than targets when evaluated after a lag of 1 trial, $t(1, 18) = 2.26, p < .05$. This was absent at lag 7 ($p > .1$). The main effect of lag on 'seen' ratings was non significant ($p > .1$). The first four bars of Figure 25 show the interaction of lag and attention for old TBR items.




Figure 25: Attention effects for old TBR bottles at lags 1 and 7 and novel TBR bottles (all presented at no lag). Error bars represent ± 1 SE of the mean.

Novel TBR Item Evaluations

Figure 25 also displays the mean ratings of novel TBR items (nJPT and nJPD are all lag 0 items and so ratings are collapsed across group). An ANOVA revealed that the main effects of attention and lag on novel ratings were non significant, and there was no significant interaction (all $p > .1$). The ratings of nJPT and nJPD TBR items did not differ significantly.

When old TBR target and distractor bottles were compared to novel TBR nJPT and nJPD bottles in a repeated measures ANOVA, the main effects of exposure (old, novel) and attention (target, distractor) were non significant (both $p > .1$). There was also no significant interaction ($p > .1$).

Evaluations by Attention in Immediately Preceding Trial

Evaluations of old items were compared in a paired samples t-test according to whether they were JPT or JPD in the bottle search trial immediately preceding the evaluation. This was conducted only for participants in the lag 7 group. For participants rating items at lag 7, JPT were not rated differently than JPD ($p > .1$). Lag 7 participants rated JPTs as, on average, 2.90 ($SD = 0.45$) and distractor-like items as 2.99 ($SD = 0.39$) on the 'good' affective scale.

Category Effects

The main effect of bottle type was significant, $F(1, 37) = 5.39, p < .05$. Drinking bottles were evaluated as more positive than cleaning bottles, $t(1, 37) =$

2.65, $p < .05$. This effect did not interact with attention in either the old TBR items or the novel TBR items (both $p > .1$). See Table 9 for bottle type ratings.

Table 9: Mean ratings and SD per stimulus type.

Stimulus Type	<i>M</i>	<i>SD</i>
Drinking bottles	3.13	0.28
Cleaning bottles	2.83	0.41

Discussion

Experiment 3 successfully replicated the distractor devaluation effect. Items that were inhibited during pre-exposure were subsequently devalued on affective evaluation. Like Experiments 1 and 2, this effect diminished over time in Experiment 3. After a lag of 1 minute (and several intervening trials) the distractor devaluation effect disappeared. The distractor devaluation effect may be transient: a mechanism to guide only short-term behaviours. After a period, the attention effect may re-set in favour of established affective attitudes towards the object in question.

I predicted that in Experiment 3 I would find attention effects in evaluations of token stimuli. I hypothesised that bottle stimuli must be tokenised (i.e. individuated) in order to successfully categorise them as either drinking or cleaning bottles. This appeared to be the case: only bottles that were inhibited in the oddball search arrays were subsequently devalued. Attention effects did not generalise to bottles of the same type as the distractor bottles. A novel bottle that was of the same type as the inhibited bottle did not get devalued, even though it was evaluated before the old distractor.

In addition to this, novel bottles and old bottles evaluated after a long lag were not influenced by the attentional state of same-type bottles in the preceding trial. Even though novel bottles have no attentional state associated with them, and old bottles that are not rated immediately appear to have lost attentional effects, neither 'acquire' the attentional effects of the previous trial.

The results of Experiment 3 support one side of the individuation hypothesis: for a successfully individuated object, attentional inhibition during pre-exposure will result in an affective devaluation of the evaluation of that object. It appears that the attentional state active during pre-exposure has been encoded with individual object representations. Therefore, distractor devaluation effects appear to affect 'token' exemplars of objects. 'Type' exemplars of inhibited objects do not appear to be subject to 'devaluation-by-inhibition'. This begs the question: "What is the effect of attention on subsequent affective evaluation if an object has *not* been successfully individuated?"

Chapter 5

Type Effects on Evaluation in an Oddball Search Paradigm

In Experiment 3 I demonstrated token-specific distractor devaluation effects following object individuation. But what are the effects of attention on subsequent affective evaluations if objects are not successfully individuated? If an object is not individuated but is only encoded as an occurrence of a category (a type), then attentional states may not be encoded with the individual object representation. Instead, the attention state active during pre-exposure may be encoded with the whole category representation. Upon subsequent presentation then, do we see devaluation of any item belonging to the distractor type, even if the evaluated item is novel? I hypothesised that when object individuation does not occur distractor devaluation will generalise to items of the same type as the previously inhibited object.

Experiment 4: Type distractor devaluation effects

Consider a bird. When you see one in the garden any number of cues may tell you that it is, in fact a bird. It has a wing, it has a beak, it is chirping, it came into your garden by flying. Importantly, it is not necessary to individuate bird A from bird B in order to be able to classify them both as “birds”. This is in contrast to the bottle stimuli in Experiment 3, which were ambiguous. The bottle stimuli’s category had to be inferred from thorough analysis of each bottle, promoting individuation and token effects. In the case of animals, the presence of a wing or a fin indicates to participants the animal’s category. The search task can be completed without further need for individuation, promoting type effects.

In Experiment 4 I follow the exact methodology of Experiment 3, with the exception that the evaluated items were animals. I hypothesised that the animals

would not be individuated and therefore group effects of attention on subsequent affective emotional evaluation would emerge. I predicted that distractor devaluation would generalise to any animal of the same type as the animal just inhibited in a search array.

Method

Participants

Forty participants (36 females, mean age = 20.15 years) from Bangor University volunteered to take part.

Apparatus

As described in the General Methods.

Stimuli

The animal and bottle stimuli described in the General Methods were used in this experiment. Bottles were used as filler stimuli only, and were never evaluated by participants.

Procedure

As described in Experiment 3. The critical difference between Experiment 4 and Experiment 3 is that participants were required to rate animals. Like Experiment 3, participants were required to rate animals that they had not seen in a search array.

For a participant rating birds, following an animal search array, the participant saw a novel bird (novel TBR item) and was required to rate it. A filler search array of bottles was presented and then the bird seen in the original oddball task array was rated (OT or OD TBR item at a lag of 1, see Figure 26).

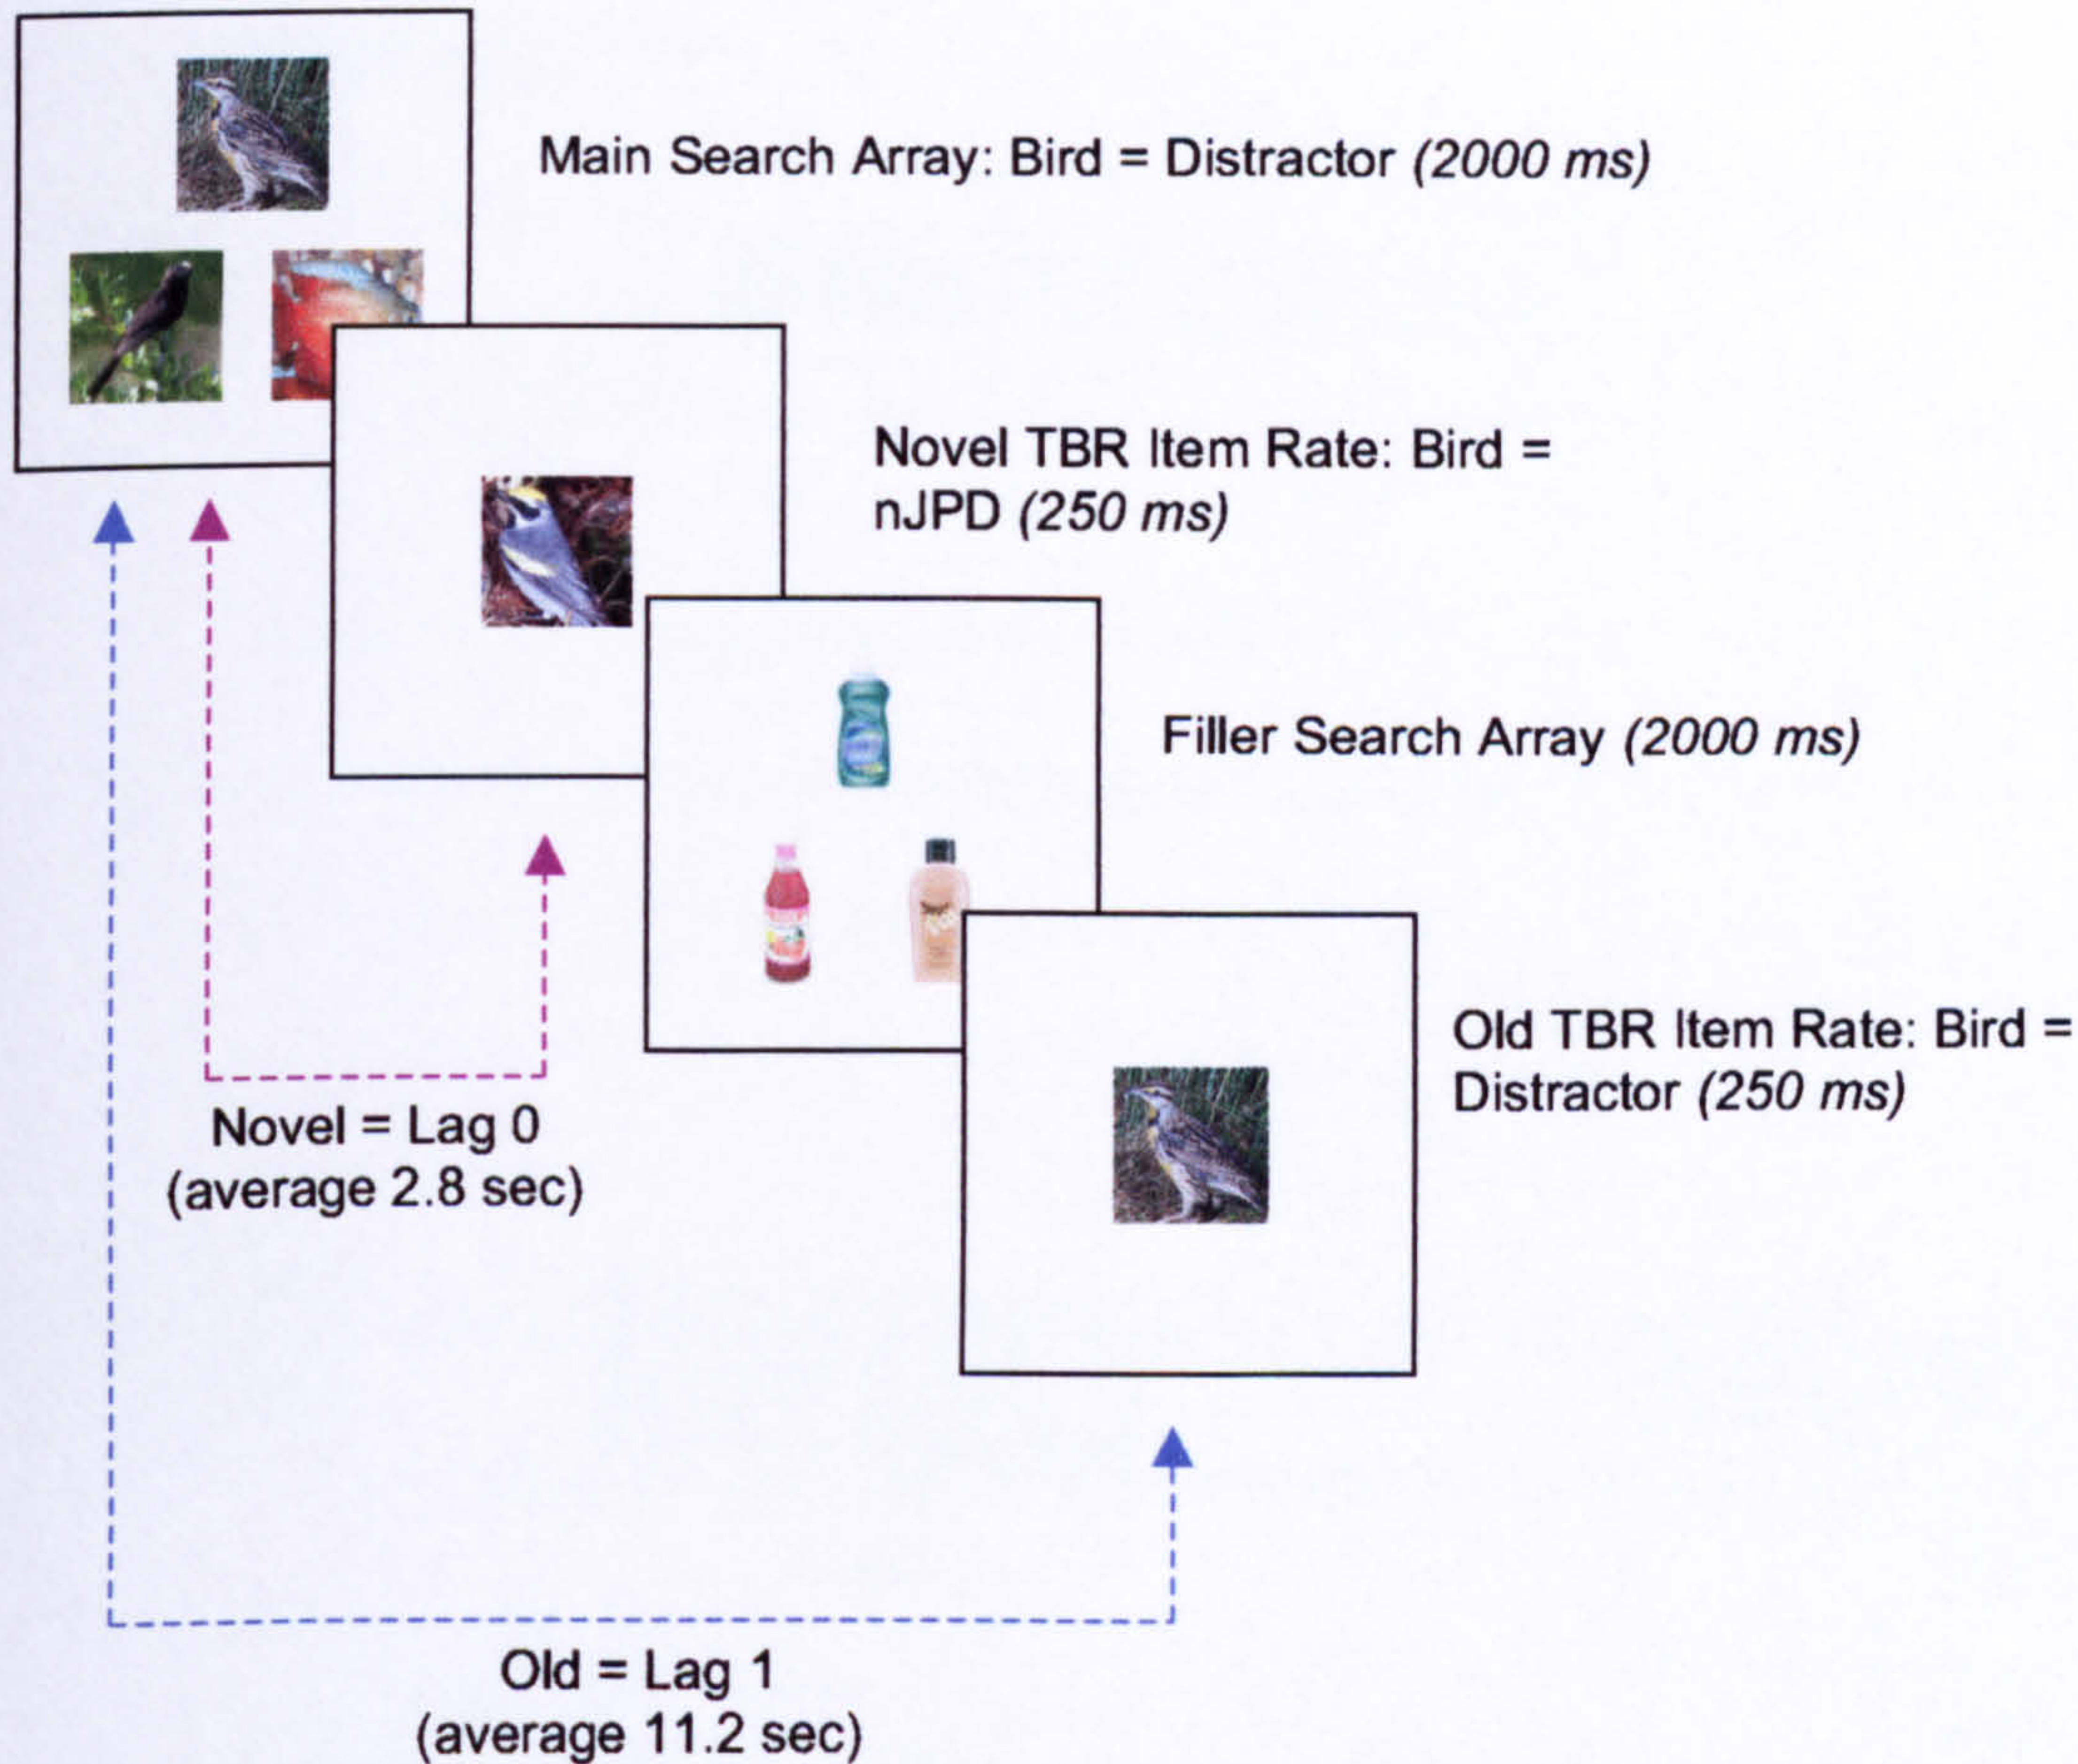


Figure 26: Trial sequence for a participant rating birds at lag 1 (60 trials).

Half of the participants saw the old TBR item at lag 1 and half saw it at lag 7 (as in Figure 24, page 127). A novel TBR item was always presented at lag 0. The table below (Table 10) describes the participant groups.

All participants rated items on an affectively 'good' 5-point scale. Prior to the main body of the experiment, participants saw 6 practice trials to ensure satisfactory performance.

Design

Half of the participants rated birds, and half rated fish (their filler search array contained bottles).

Table 10: Participant groups.

	Rate Birds	Rate Fish
Lag 1	n = 10	n = 10
Lag 7	n = 10	n = 10

The lag 1 condition contained 60 trials (60 animal search arrays, 60 novel TBR animal rates, 60 bottle search arrays and 60 old TBR animal rates). Of the novel TBR items, half were nJPT and half were nJPD. Of the old TBR items, half were OT and half were OD.

As in Experiment 3, for participants at lag 7 only, the animal search array that the TBR item followed could be matched or mismatched to the original search array (JPT or JPD).

Results

Data Analysis

Removal of ‘errors’ to detect the target excluded 6.96% of evaluations. Removal of ‘errors’ to rate the TBR item excluded a further 3.87% of evaluations from the analysis. Finally, evaluations of old TBR animals were only included if they followed filler search trials in which the target bottle was correctly identified. This excluded a further 13.11% of the old TBR evaluations from the analysis.

The results of Experiment 4 are thus based on 40 participants with an average of 53.78 old TBR item evaluations, and 53.55 novel TBR item evaluations.

Oddball Search Performance

The mean accuracy to identify the target in the animal search array was 96.46% ($SD = 3.27\%$). The mean accuracy to identify the target in the filler trials, the bottle search array, was 86.33% ($SD = 6.44\%$). The mean RT to correctly identify a target animal was 2709 ms ($SD = 519$ ms). Neither the animal or bottle search accuracy, nor the RT to identify the animal target differed between the ‘lag 1’ and ‘lag 7’ participant groups (all $p > .1$).

As can be seen from Table 11, participants in Experiment 4 were faster and more accurate at categorising and reporting the target in their search arrays than participants in Experiment 3. This suggests that to categorise an animal was easier than to categorise a bottle. The performance data, then, support the idea that in the bottle oddball search experiment (Experiment 3) participants were individuating and then categorising stimuli, but in the current experiment, participants are first categorising stimuli.

Table 11: Experiments 3 and 4 search performance. Mean accuracy and RT for main search arrays only. Numbers are *M* and (*SD*).

Performance Measure	Experiment 3 (Bottles)	Experiment 4 (Animals)
Accuracy	83.99% (7.89%)	96.46% (3.27%)
RT	3260 ms (528 ms)	2709 ms (519 ms)

Rating Data

A 2 x 2 x 2 repeated measures ANOVA with exposure (old, novel) and attention (target, distractor) as within-group factors and lag (1, 7) as a between-group factor was first conducted. The ANOVA revealed no significant main effects, and only one interaction, exposure by attention, $F(1, 39) = 6.40, p < .05$ (all other p 's $> .05$). Ratings of old and novel bottles were next considered separately.

Old TBR Item Evaluations

A repeated measures ANOVA with attention (OT, OD) as a within group factor and lag (1, 7) as a between group factor was performed. There were no statistically significant main effects of attention or lag, nor was there an interaction (all $p > .1$). The evaluations of animals did not vary significantly as a function of whether they had previously been seen as a target or as a distractor (see Figure 27).




Figure 27: Attention effects for old TBR animals at lags 1 and 7, and novel TBR animals (all presented at no lag). Error bars represent ± 1 SE of the mean.

Novel TBR Item Evaluations

A repeated measures ANOVA with exposure (old, novel) and attention (target, distractor) as within-group factors revealed no significant main effects of exposure or attention (both $p > .1$). There was, however, an interaction, $F(1, 39) = 4.17, p < .05$. Importantly, a t-test revealed that nJPT were rated as significantly more positive than nJPD, $t(1, 39) = 2.07, p < .05$. Distractor devaluation has generalised to the novel stimuli. There was no significant difference between participants in the lag 1 group and those in the lag 7 group in their ratings of novel TBR animals ($p > .1$). Figure 27 displays the mean ratings of nJPT and nJPD (as all are lag 0 items, data is collapsed across group).

Evaluations by Attention in Immediately Preceding Trial

Evaluations of old items were compared in a paired samples t-test according to whether they were JPT or JPD. There were no significant main effects or interactions between attention in the just previous trial or lag (all $p > .1$). Ratings of old TBR animals were not influenced by the attention directed towards same-type animals in the just previous search trial. Distractor devaluation generalisation effects did not persist to a lag of 1 trial.

Category Effects

The main effect of animal category was non significant ($p > .1$). Birds and fish were evaluated as equally positive. This did not interact with attention in either the old TBR items or the novel TBR items (both $p > .1$, see Table 12).

Table 12: Mean ratings and SD per stimulus type.

Stimulus Type	<i>M</i>	<i>SD</i>
Birds	2.89	0.36
Fish	2.95	0.36

Discussion

In Experiment 4 I made a very important finding: distractor devaluation can *generalise* to items similar to an inhibited object. nJPD were devalued. It appears that during animal categorisation in the search trials animals were not individuated. Therefore inhibition appears to have been mapped, not onto a token representation, but onto type representations: whole categories of objects. Thus we can posit that the next item evaluated from that category reinstantiates this inhibition and gets devalued. This occurred even though the item evaluated had never been seen before; that it was of the same type as the distractor was sufficient for it to be devalued.

Surprisingly, this effect did not survive the filler search trial in order to affect ratings of old animals. The old TBR items evaluated at lag 1 did not get devalued. This firmly suggests that: a) inhibition was not mapped onto token representations as the distractor devaluation effects found in Experiment 3 were not replicated; and b) the ‘devaluation-by-inhibition’ of the category is transient as they did not affect the old TBR items evaluated after the filler search trial. This ‘transience’ may have been a result of the experimental design: in one trial ‘inhibition’ may have been mapped onto the category representation, but in the following trial ‘attention’ may have been mapped onto the category representation. That the attentional state encoded with the

category representation needed to be repeatedly updated may have caused this effect to have been adaptively transient.

Experiment 5: Type effects of faces in an oddball search paradigm

In Experiment 5, the experimental design of Experiments 3 and 4 was replicated using face stimuli. I expected that, due to our expertise with human faces, these stimuli would be easily individuated by participants. I thus expected that token distractor devaluation effects, and not type distractor devaluation effects, would result from using face stimuli in this paradigm. This was, however, not the case.

Method

Participants

Forty participants (34 females, mean age = 20.30 years) from Bangor University volunteered to take part.

Apparatus

The same apparatus were used as in the General Methods.

Stimuli

In this experiment, a computer-generated face set was used (GenHead 1.2; Genemation Limited, 2002-2004). This program generated faces that were all Caucasian, between 20 – 30 years, with no facial expression (facial expression weights set to 0) and with no hair or other features visible. See Figure 28 for female and male exemplars of Genemation faces. Faces were presented on a uniform black

background. 175 female faces and 175 male faces were produced. The bottle stimuli used in Experiments 1-4 were used for the filler search trials in this experiment.



Figure 28: Genemation Faces, Female and Male exemplars.

Design and Procedure

The procedure of this experiment was identical to that of Experiments 3 and 4. The only difference is that in Experiment 5, the main search trials displayed the Genemation faces, and the filler search trials displayed the bottle stimuli. Half of the participants saw the old TBR item at lag 1 and half saw it at lag 7. A novel TBR item was always presented at lag 0.

Half of the participants rated males, and half rated females. All participants rated items for ‘trustworthiness’ on an affectively positive 5-point scale (where a score of 1 means ‘untrustworthy’ and a score of 5 means ‘very trustworthy’). Table 13 describes the participant groups.

Prior to the main body of the experiment, participants saw 6 practice trials to ensure satisfactory performance.

Table 13: Participant groups.

	Rate Males	Rate Females
Lag 1	n = 10	n = 10
Lag 7	n = 10	n = 10

The lag 1 condition contained 60 trials (60 face search arrays, 60 novel TBR face rates, 60 bottle search arrays and 60 old TBR face rates; see Figure 29). Of the novel TBR items, half were nJPT and half were nJPD. Of the old TBR items, half were OT and half were OD.

As in Experiments 3 and 4, for participants at lag 7 only, the face array that the TBR item followed could be matched or mismatched to the original face search array. Therefore half of the old TBR items were JPT and half were JPD.




Figure 29: Trial sequence for a participant rating males, 'seen' TBR items evaluated at lag 1 (60 trials).

Results

Data Analysis

Removal of 'errors' to detect the target excluded 14.37% of evaluations. Removal of 'errors' to rate the TBR item excluded a further 1.34% of evaluations from the analysis. Finally, evaluations of old TBR faces were only included if they followed filler search trials in which the target bottle was correctly identified. This excluded a further 15.84% of the old TBR evaluations from the analysis.

Two participants were excluded from the analysis. The first participant's RT to identify a target face was over 2.5 SD slower than the average RT. This participant was rating female faces at lag 7. The second participant was removed because she was performing the search task quickly, but inaccurately and appeared to be engaging in a speed/accuracy trade-off. This participant had approximately half the data points contributing to the analysis than the other participants. This participant was rating male faces at a lag of 7. The results of Experiment 5 are thus based on 38 participants with an average of 50.82 old TBR item evaluations, and 50.18 novel TBR item evaluations.

Oddball Search Performance

The mean accuracy to identify the target in the face search array was 91.40% ($SD = 4.56\%$). The mean accuracy to identify the target in the filler trials, the bottle search array, was 84.39% ($SD = 8.08\%$). The mean RT to correctly identify a target face was 2721 ms ($SD = 494$ ms). Neither the face or bottle oddball search accuracy, nor the RT to identify the face target differed between the 'lag 1' and 'lag 7' participants (all $p > .1$).

Rating Data

A 2 x 2 x 2 repeated measures ANOVA with exposure (old, novel) and attention (target, distractor) as within-group factors and lag (1, 7) as a between-group factor was first conducted. The ANOVA revealed a marginally significant main effect of exposure, $F(1, 37) = 3.91, p = .056$. Also, a significant main effect of attention, $F(1, 37) = 9.08, p < .01$. There was also one interaction, lag by attention, $F(1, 37) = 4.91, p < .05$ (all other p 's $> .05$).

Exposure

When OT and OD faces were compared to nJPT and nJPD faces in the ANOVA, the main effect of exposure (old, novel) was (marginally) significant. Old TBR faces were rated more positively than novel TBR faces, $t(1, 37) = 1.96$, $p = .057$, see Figure 30. Lag group did not interact with effects of exposure ($p > .1$).




Figure 30: Ratings of 'seen' and novel faces collapsed across lag group and attention condition. Error bars represent ± 1 SE of the mean.

To elucidate the other findings of the global ANOVA, the ratings of old and novel bottles were next considered separately.

Old TBR Item Evaluations

A repeated measures ANOVA with attention (target, distractor) as a within group variable and lag (1, 7) as a between group variable revealed no significant main effects (both $p > .1$), but a marginally significant interaction of attention and lag, $F(1, 37) = 3.53$, $p = .07$. After a lag of 1 trial, OD faces were rated more *positively* than OT

faces (this was a marginally significant effect: $t(1, 19) = 1.92, p = .07$). After a lag of 7 trials, OT faces and OD faces were not rated significantly differently ($p > .1$). The first four bars of Figure 31 illustrate this.




Figure 31: Attention effects for old TBR faces at lags 1 and 7 and novel TBR faces (all presented at no lag). Error bars represent ± 1 SE of the mean.

Novel TBR Item Evaluations

A repeated measures ANOVA with attention (nJPT, nJPD) as a within group factor and lag (1, 7) as a between group factor found that there was no significant difference between participants in the lag 1 group and those in the lag 7 group in their ratings of novel TBR faces, nor did this interact with attention (both $p > .1$). However, the main effect of attention on novel ratings was significant, $F(1, 37) = 14.16, p = .001$. nJPD TBR faces were rated as more *positive* than nJPT TBR faces, $t(1, 37) = 3.80, p = .001$. Figure 31 displays the mean ratings of nJPT and nJPD faces (all lag 0 items, collapsed across group).

Evaluations by Attention in Immediately Preceding Trial

Evaluations of old TBR items were compared according to whether they were JPT or JPD in the face search trial immediately preceding the evaluation. JPD faces were rated more *positively* than JPT faces, $t(1, 37) = 2.95, p < .01$. This did not interact with lag group ($p > .1$). It appears that the attention state of the preceding trial, and not the attentional state directed towards the face during search determines subsequent emotional evaluation. This ‘attention effect’ did not interact with exposure of the faces ($p > .1$): distractor-like faces were rated as more positive than target-like faces in both old (JPD, JPT) and novel (nJPD, nJPT) faces.

Figure 32 displays the data shown above in Figure 31 with the critical difference that the old TBR items for lag 7 participants are now displayed according to just previous trial attention (JPT, JPD), and not original (OT, OD) attention. Note that, at lag 1, the just previous trial *is* the original trial, and so these bars are labelled as before. The data for the novel TBR faces is also re-plotted.




Figure 32: Mean ratings of target-like and distractor-like faces, defined by the attentional state of the category in the preceding trial. Error bars represent ± 1 SE of the mean.

Category Effects

The main effect of face category was non significant ($p > .1$). Males and females were evaluated as equally positive. This effect did not interact with attention in either the old TBR items (original attention or just previous trial attention) or the novel TBR items (all $p > .1$). See Table 14 for mean ratings of male and female faces.

Table 14: Mean ratings and SD per stimulus type.

Stimulus Type	<i>M</i>	<i>SD</i>
Males	2.89	0.30
Females	2.96	0.22

Discussion

Generalisation of Mere Exposure

The surprising finding of this experiment was the consistent discovery of generalised *mere exposure* effects. In Experiment 5, old faces were rated as more trustworthy than novel faces. This is a classic demonstration of mere exposure. This result in itself is impressive, as with the majority of mere exposure studies, to see an effect requires multiple presentations of a stimulus (Bornstein, 1989). Here, I found mere exposure with only a single presentation of a face in a search array.

Above and beyond this, though, is the finding that a face belonging to the category just seen as distractors will be rated as more positive than a face belonging to the category just seen as a target. This is the case whether the face to be evaluated is old or novel. Further, old faces evaluated after a lag (of about one minute and several intervening trials) are not rated according to the attentional state active in the trial in which they were exposed, but are evaluated according to whether they belong to the target or distractor category of the just previous trial.

Why do I interpret this as mere exposure? Because in the oddball search task, the target is defined by being the odd-one-out. There is only one exemplar from the target category, but two exemplars from the distractor category. The distractor category has received twice the 'exposure' of the target category. To my knowledge, this is the first demonstration of a category-based generalisation of mere exposure.

Other demonstrations of generalised mere exposure have found the effect to generalise to different orientations of the same stimulus (Seamon & Delgado, 1999). This is important because it implies that the perceptual fluency that a stimulus benefits from is tolerant of a different representation of the same object. Extending this finding, researchers have used artificial grammar strings to demonstrate the occurrence of generalised mere exposure (Gordon & Holyoak, 1983; Manza, Zizak & Reber, 1998). In these studies, participants are exposed to a series of consonant strings. These strings are called “grammatical” because each one conforms to the rules of grammar. Following exposure, participants affectively evaluate a set of novel strings, half of which conform to the rules of grammar, and half of which do not. Participants tend to give higher ratings to novel strings that conform to the rules of grammar, than to those which do not. This finding has been interpreted as evidence that the positive affect produced by exposure to a set of stimuli generalises to previously seen stimuli that are structurally related to the exposed stimuli.

Of most relevance to the current results, are the findings of Rhodes, Halberstadt and Brajkovich (2001). In their experiments, Rhodes and colleagues first exposed participants to a set of faces, and then required participants to affectively evaluate the exposed faces, new faces, or composite faces (created by averaging two ‘exposed’ faces into a new image). They found that both exposed and composite faces were liked more than novel faces. They interpreted this as generalisation of mere exposure to novel faces that are structurally related to the exposed faces. Like the Seamon and Delgado (1999) study described previously, this could be interpreted as testing the boundaries of the tolerance of the mere exposure effect. Composite faces

were, in effect, half an exposed face. It appears that this was enough to allow perceptual fluency to act on affective evaluations.

Importantly, the Rhodes et al. study did not require participants to act on faces during exposure, they were instructed only to study each face. In the current study, participants were required to perform a task during exposure: namely, to categorise the faces. Encouraging this categorisation, and controlling the number of exposures per category, I believe led to the current finding that mere exposure can generalise to novel faces of the most exposed category.

Recall, that in Chapter 1 I discussed that Monahan, Murphy and Zajonc (2000) demonstrated that even novel stimuli unrelated to the exposed stimuli benefit from following a repeated-exposure block of trials. They concluded that a positive mood state is induced by fluency and is interpreted as stimulus quality, regardless of the item under evaluation. Here, I have evidence that directly contradicts this interpretation. Distractor-like faces were rated as more positive than target-like faces. If generalised mere exposure were due to a free-floating generic positive affect, one would have to predict no difference in the ratings of distractor- and target-like novel items. This is not the case, and therefore I suggest that the generalised mere exposure to novel faces in this instance was related to category exposure during the search task.

Distractor Devaluation and Effort

The finding of generalised mere exposure for faces is interesting, but unexpected. Why did I not find generalised distractor devaluation? If distractor devaluation is a result of an inhibitory tag associated with either a token or a type,

then perhaps this inhibitory tag was simply not stored with either representation in memory.

Kiss et al. (2007) recently suggested that distractor devaluation only results from *effortful* inhibition. In their study, evaluation followed a simple two-item search task. During search, ERP measures of the N2Pc component were taken. The N2Pc is believed to reflect attentional selection difficulty. They found that distractors which were rated negatively had come from search trials in which a large N2Pc was elicited, compared to distractors which were rated positively. In contrast, no such difference in N2Pc magnitude was found between search trials in which positively versus negatively evaluated targets originated. The authors concluded that distractor devaluation is closely linked to selection difficulty, and will only result when the inhibition of a distractor was effortful.

Therefore, perhaps in the current study performing an oddball search on the basis of gender was too easy. Perhaps to inhibit a distractor was not effortful.

Meta-analyses of Experiments 3-5 do not provide clear insight into this possibility. A MANOVA with independent variables of experiment category (bottles, animals, faces) and lag group (1, 7), and dependent variables of accuracy to locate the target and RT to correctly locate the target was performed. As expected there was no significant main effect of lag ($p > .1$). There was, though, a main effect of experiment category, $F(2, 115) = 24.16, p < .001$. There was no significant interaction of lag with category ($p > .1$). Further analyses of RT and accuracy effects (collapsed by lag

groups) were performed and are depicted in Figure 33 (RTs) and Figure 34 (accuracies).

Paired t-tests on the RTs to correctly identify the target in search revealed that a bottle target took significantly longer to identify than either an animal target, $t(1, 77) = 4.64, p < .001$, or a face target, $t(1, 75) = 4.59, p < .001$. There was no significant difference in the RT to identify an animal target or a face target ($p > .1$). From this data, it appears that identifying a bottle target was the most difficult task, but there is no evidence to suggest that identifying a target face was any easier than identifying a target animal.




Figure 33: Meta-analysis of the RTs to correctly locate a target by participants in Experiments 3, 4 and 5.

Similarly, paired t-tests were performed on the accuracy to identify a target in the three experiments. Accuracy to identify a target face was found to be greater than

accuracy to identify a target bottle, $t(1, 75) = 5.01, p < .001$, but less than accuracy to identify a target animal, $t(1, 77) = 5.65, p < .001$.

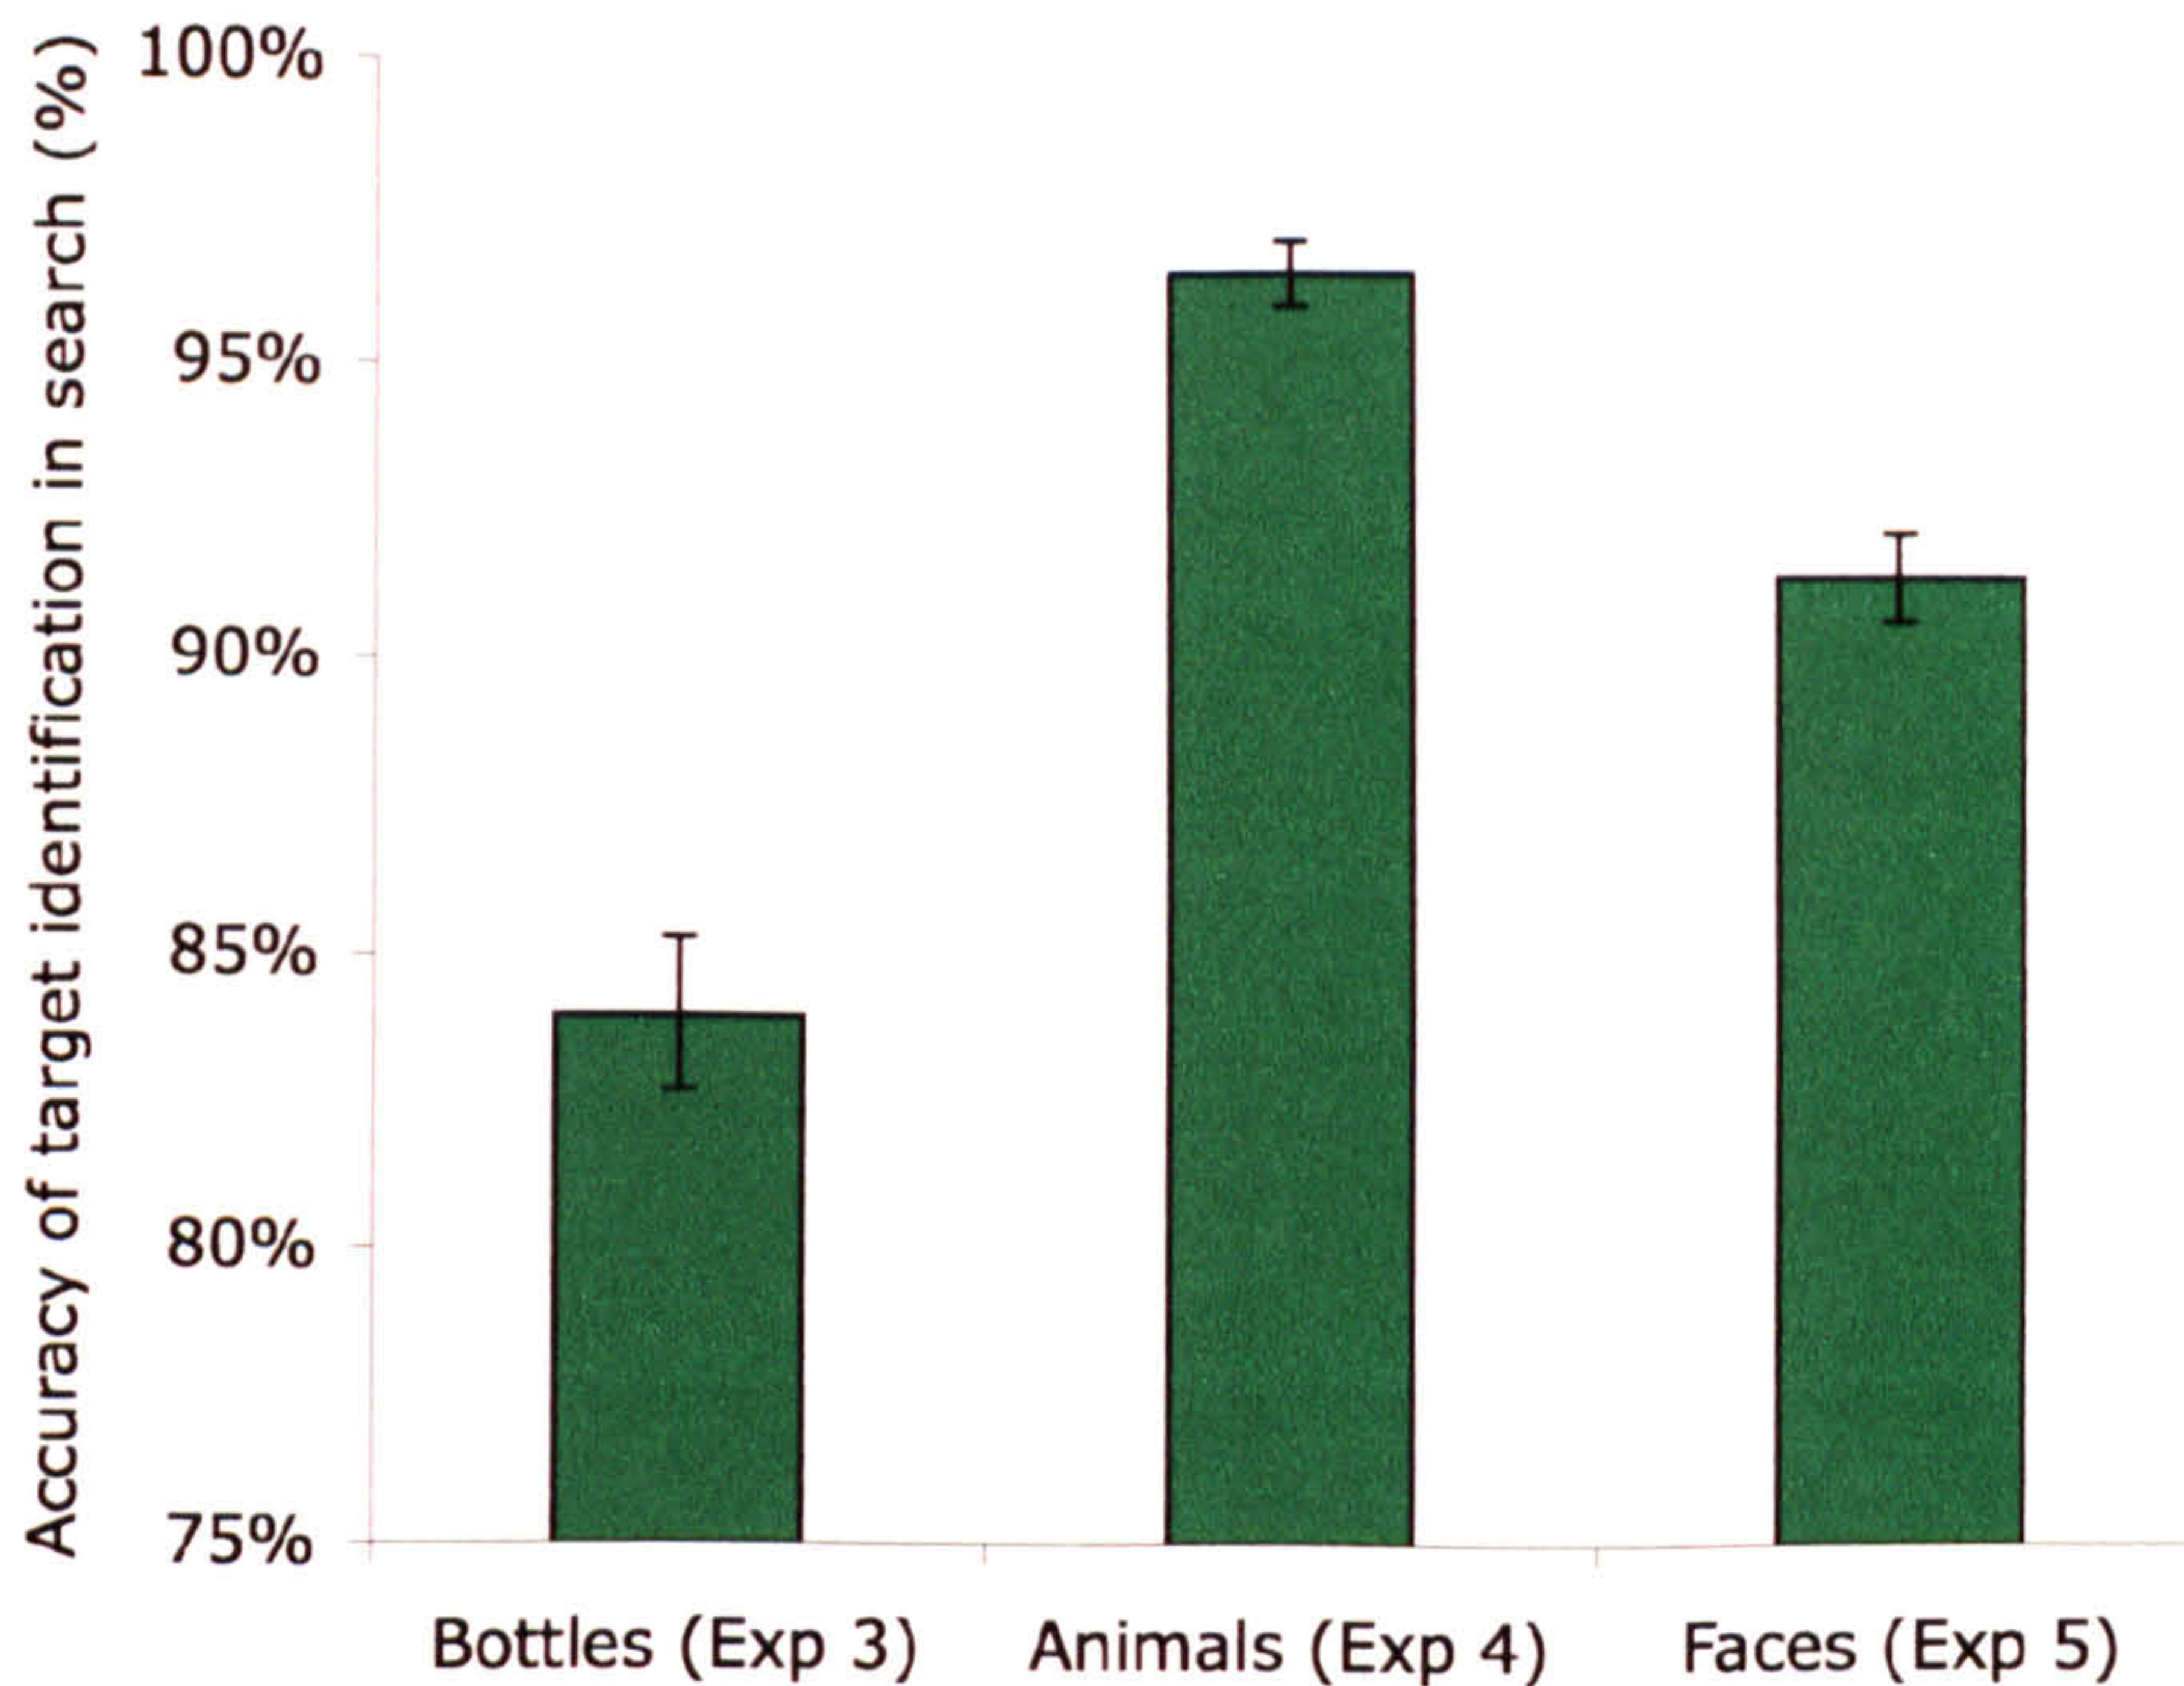


Figure 34: Meta-analysis of the accuracy to locate the target in oddball search by participants in Experiments 3, 4 and 5.

Taken at face value, we would have to conclude that discriminating between male and female faces was in fact harder than discriminating between birds and fish (which was performed more accurately). Why, then, did I find distractor devaluation for a task which required less effort and no distractor devaluation for a task which required more? Possibly the meta-analysis results of task-difficulty are not showing the full picture. It is possible that, due to the homogeneous nature of the Genemation face stimuli, participants were simply more bored in this experiment compared to the animals experiment in which stimuli were much more varied, had more colour, were presented on backgrounds and depicted animals of different species. Perhaps an

increase in mistakes reflects boredom and a loss of concentration, rather than task-difficulty.

That is not to say that I am arguing against the occurrence of distractor devaluation for faces per se; I only argue against it in the context of this experiment and using these face stimuli. For example, Raymond et al. (2005, Experiment 3) found distractor devaluation for faces following search. However, in that experiment the search was for a conjunction of gender and colour mask. This task was probably more difficult than the current experiment, and therefore devaluation effects were seen there. Further studies using the N2Pc as a measure of attentional effort may elucidate the findings of Experiment 5.

Whatever the reason for it, the discovery of mere exposure effects in the oddball paradigm illuminated a potential confound: in a given trial, there are more instances of the distractor category than of the target category. If ratings are influenced by both inhibition and mere exposure of a category, then any potential distractor devaluation effects (from the inhibition of the distractor category) may be diluted by increased mere exposure (from double exposures of the distractor category compared to target category). Therefore, for the remaining experiments of this thesis I developed a new paradigm to test distractor devaluation. All subsequent experiments in this thesis use variations of a go / no-go task.

Chapter 6

Distractor Devaluation Following Behavioural Inhibition

Thus far, my thesis has focussed on one set of inhibition processes:

interference control. As outlined in Chapter 1, we can assume that organisms have a limited capacity for processing stimuli; the presence of a mechanism that enables selection of the most relevant information for processing will be more efficient than a system in which no prioritising occurs. The suppression of task-irrelevant locations, features and objects enables task-relevant locations, features and objects to achieve priority for consciousness and action. I have demonstrated that such interference control can result in distractor devaluation upon subsequent emotional evaluation of an inhibited object.

Behavioural inhibition, on the other hand, refers to the suppression of an action evoked by a location, a feature or an object. The ability to inhibit responses is an important component of cognitive control. Imagine you are driving down the road when someone pulls out in front of you causing you to brake suddenly. The action of using our right foot to press the brake pedal is good: it will stop us crashing. However, when dining with friends in a restaurant, it is polite not to begin eating until your companions have also received their meals. Inappropriate actions such as eating before it is polite to do so must be inhibited.

Failures of inhibitory control are characteristic of a wide range of neuropsychiatric disorders: attention deficit hyperactivity disorder, obsessive-compulsive disorder, and schizophrenia, for example (Barkley, 1997; Hershey et al., 2004). In the general population, inhibitory control problems are related to high risk for substance abuse and other dangerous behaviours (Finn, Justus, Mazas & Steinmetz, 1999; Swann, Bjork, Moeller & Dougherty, 2002).

To measure behavioural inhibition, two behavioural paradigms are commonly employed: the stop signal paradigm and the go / no-go task. In the stop signal paradigm (Logan & Cowan, 1984; Logan, 1994) participants perform a choice reaction task. On a random selection of the trials a 'stop' signal instructs participants to withhold their response (a stop signal trial). On trials with a short interval between the presentation of the stimulus and the presentation of the stop signal, participants typically withhold their response easily. On trials with a long delay between stimulus and stop signal presentations, participants are typically unable to withhold their response. Logan and Cowan (1984) proposed the 'horse-race model' to account for these findings. They proposed that two processes occur, simultaneously and independently: a 'go' process and a 'stop' process. If the stop process is completed before the go process, the response is successfully withheld. On the other hand if the go process finishes first, participants will execute their response.

More relevant to the upcoming experiments of this thesis is the go / no-go task. The task requires a rapid decision about whether to respond to a particular stimulus. In a simple version, participants are instructed to press a key as fast as possible if they see a letter ('go'), but if the letter is an 'X', they must not respond ('no-go'). The letter X appears on approximately 20% of trials and so the prepotent response is to press the key (as this strategy leads to a fast response). Problems with behavioural inhibition are measured by responses to a no-go trial (i.e. false alarms).

Can the inhibitory state of an action be encoded with an object's representation in order to guide future evaluations? The team from Bangor (including myself: Fenske, Raymond, Kessler, Westoby & Tipper, 2005) used a variant of the go / no-go task to answer this question. Whether we would successfully find 'distractor

devaluation' following an action inhibition was by no means clear-cut. Research with brain-injured and normal participants suggested that no single brain circuit controls all kinds of inhibition (e.g. Miyake, Friedman, Emerson, Witzki & Howerter, 2000; Hamilton & Martin, 2005). Even within response inhibition, debate continues as to the contributions of the supplementary motor area, the dlPFC, the ACC, and the vmPFC, each of which appears to vary according to the response required (or to be inhibited) and the complexity of the task (e.g. Dias, Robbins & Roberts, 1997; Mostofsky et al., 2003; Clark, Cools & Robbins, 2004).

However, interference control and behavioural inhibition appear to share some common processes. For example, Freitas and colleagues recently demonstrated that the past attentional state of a cue can impact performance on a subsequent go / no-go task (Freitas, Azizian, Leung & Squires, 2007). In this experiment, participants were presented with two digits and were required to report the digit presented in a target colour (target), and ignore the digit not presented in the target colour (distractor). Targets and distractors were then presented in a go / no-go task. The previously attended and ignored digits were presented inside a circle or a square. If the digit was in a circle, the participants were required to make a response (go trial), if the digit was presented inside a square participants were to withhold a response (no-go trial). Figure 35 demonstrates this procedure.

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.

Go trials which contained prior targets were responded to faster than go trials which contained prior distractors. ERPs were also collected. The P3 component, believed to be indicative of response control, was modulated by the prior attentional state of the stimuli. As you would expect, the P3 component was larger (had a higher amplitude) on no-go trials than go trials. Interestingly, the authors found that this difference was only significant for previous targets, and not previous distractors. The authors concluded that inhibiting a response to a previously selected target requires additional response control. While Freitas et al. found no evidence for response inhibition having been encoded in memory, as measured by ERP components, we (Fenske et al., 2005) did, using evaluations of prior no-go stimuli.

In our study, participants were presented with pairs of faces, presented to the left and right of fixation. A cue was briefly flashed over one of the pair, either a red or a green semi-opaque oval. If a green cue was presented, participants had to press


either a left or right response key to indicate which face the cue had flashed on (go trial). If a red cue was presented, participants had to withhold a response (no-go trial). For half of the participants a green oval was a no-go cue and a red oval was a go cue.

After a filler trial, participants were re-presented with the pair of faces (without the cues) and were asked one of four questions. Two questions were perceptual in content: “Which background is lighter?”, and “which background is darker?”. Two questions were emotional in content: “Which person is more trustworthy?”, and “which person is less trustworthy?”. Participants used the left and the right response keys to select a face in answer to the question (see Figure 36).

By contrasting the responses to the positively and negatively valenced emotional questions, we were able to demonstrate the negative emotional impact of inhibition to a face. Following no-go trials, the faces that were uncued (not inhibited) were chosen more often as “more trustworthy”. Faces that were associated with no-go cues (inhibited) were chosen more as “less trustworthy”.

Inhibition specifically affected emotional responses; judging which background was “lighter” or “darker” was unaffected by the cue. Evaluations of faces presented in go trials were unaffected by presentation of the go cue. This supported the “devaluation-by-inhibition” hypothesis: it was the inhibition of a no-go cued face that had caused an emotional devaluation, the attention of a go cued face did not affect subsequent ratings. Figure 37 illustrates these results.

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.



Third Party Material excluded from digitised copy.
Please refer to original text to see this material.

The remaining experiments in this thesis use the go / no-go paradigm. In Experiment 6, my aim was to investigate several important questions, as yet unanswered, regarding distractor devaluation effects. First, and critically, does distractor devaluation truly reflect devaluation, or do they actually reflect a 'flattening' of evaluations? In all experiments demonstrating attention effects to this point, positive stimuli were used. By 'positive' I mean stimuli that were evaluated on average as more affectively positive than the mid-point of the scale being used. It is possible that all effects reflect a regression of evaluations towards the mean, and not

devaluation. By comparing attention effects for both positively and negatively valenced stimuli in Experiment 6, I was able to directly address this issue. If distractor devaluation occurred for positive stimuli, but ‘distractor up-valuation’ occurred for negative stimuli, we would have to conclude that prior inhibition results in emotional flattening. Such a finding would require a modification of the ‘devaluation-by-inhibition’ hypothesis.

My second aim in Experiment 6 was to investigate the correlations between attention, emotion, and memory. Do distractor devaluation effects occur because distractors are not explicitly remembered, compared to targets? In Chapter 1 I discussed that increasing attention to a stimulus has been demonstrated to improve memory for that stimulus (Intraub, 1984; Williams, Henderson & Zacks, 2005). Conversely, the ‘devaluation-by-inhibition’ hypothesis *requires* that distractors are encoded into memory successfully (along with their associated inhibition). Are stimuli that are more likely to be remembered also more likely to elicit devaluation effects?

Third, I aimed to determine how the effort required to inhibit a stimulus affects resulting evaluations. Experiment 5 suggested that if a stimulus is particularly easy to inhibit, it will not be devalued on subsequent evaluation. In Experiment 6 I investigated this in a different way. I directly compared the evaluations of distractors that were defined by a simple cue feature (colour) to those of distractors defined by a conjunction of features (stimulus type and colour).

Levels of Processing theory is an influential theory of memory proposed by Craik and Lockhart (1972). The model proposed that information could be processed

in a number of different ways and the durability or strength of the memory trace was a direct function of the depth of processing involved. Moreover, depth of processing was postulated to fall on a shallow to deep continuum. Shallow processing (e.g., processing words based on their phonemic and orthographic components) leads to a fragile memory trace that is susceptible to rapid forgetting. On the other hand, deep processing (e.g., semantic or meaning based processing) results in a more durable memory trace.

In Experiment 6 I proposed that a no-go stimulus defined by a featural cue would require only shallow processing, and as such only a fragile memory would exist. Conversely, a no-go stimulus defined by a conjunction of featural cue and stimulus type would require deep processing, and would therefore create a stronger memory trace for the stimulus. I expected, therefore, that an associated inhibitory tag would be more likely to be stored with a conjunction no-go in memory and that this would lead to greater devaluation effects, compared to featural no-gos.

Experiment 6 is in two parts: in Experiment 6a I collected baseline evaluations of the stimuli used in Experiment 6b. In Experiment 6b go / no-go tasks were performed using face and scene stimuli. Subsequent emotional evaluations and memory scores were collected and compared for positive and negative stimuli.

Experiment 6a: Devaluation of No-Go faces - Baseline

Experiment 6a served to provide baseline ratings for each of 1496 stimuli to be used in Experiment 6b. These baseline ratings were used to divide stimuli

according to valence to allow for subsequent comparison of negative and positive stimuli and investigation of the 'flattening' hypothesis.

Method

Participants

Two groups of five participants (Group 1: 3 females, mean age = 25.8 years; Group 2: 4 females, mean age = 37.6 years) from Bangor University volunteered to take part. All had normal / corrected to normal vision and informed consent was obtained.

Apparatus

A Pentium-4 computer, running E-Prime 1.0, recorded data and presented stimuli on a 55.9 cm monitor (100 Hz, 1024 x 768 resolution). The viewing distance was 70 cm.

Stimuli

Stimuli were digital colour photographs taken from the internet and consisted of 816 faces and 680 scenes. Face images (height 5.7°) were frontal views of Caucasian adults, with neutral / smiling expressions and visible hair, neck and eyes. Half were female, half male. Scene images varied in size (minimum height 5.5°, maximum height 12.9°) and were neutral images of British and American houses. Half were exterior scenes, half interior.

Design and Procedure

The experiment, approximately 45 min long, required subjects to rate, on a scale of 1 to 7, a series of images presented one at a time in the centre of the screen. Group 1 rated 816 face images for 'trustworthiness' where 1 indicated very untrustworthy and 7 very trustworthy. Group 2 rated 680 scene images for 'liking' where 1 indicated a strong dislike and 7 a strong like for the scene. Images were displayed until response, with a 0 ms inter-trial-interval. Participants were given a paper copy of the scale to refer to if required.

Results

Evaluations and RT to rate each of the images were averaged across participants. Items were ranked according to their average rating: a higher rank given to a high rating. If items had the same average rating, the higher rank was given to the item with the fastest RT to rate. This was calculated separately for faces and scenes. Ranked items were subsequently divided into thirds, the lowest third hereafter referred to as negative stimuli, the next as neutral, and the highest third as positive stimuli. The face stimulus set consisted of 272 negative, 272 neutral and 272 positive items. The scene stimulus set consisted of 227 negative, 226 neutral and 227 positive items.

Mean ratings for the negative, neutral, and positive faces were 3.26 ($SD = .56$), 4.42 ($SD = .24$) and 5.25 ($SD = .32$) respectively. The scenes had mean ratings of 2.79 (negative; $SD = .46$), 3.83 (neutral; $SD = .26$) and 4.84 (positive; $SD = .43$). The mean RT to rate face stimuli was 2366 ms ($SD = 772$ ms) and the mean RT to

rate scene stimuli was 1837 ms ($SD = 457$). RTs did not differ significantly across stimuli valence for either face or scene stimuli. Within faces, the negative class contained 69% male faces, the neutral 47%, and the positive 33%, indicating that female faces were on the whole judged as more trustworthy than male. Within scenes, the negative class contained 24% external scenes, the neutral 48%, and the positive 78%, indicating that external scenes were liked more than internal.

Discussion

The mean rating for the neutral stimuli was, for both faces and scenes, approximately only 2 SD different from the mean ratings of negative or positive stimuli of the same type. As Figure 38 demonstrates, this leads to a high probability of overlap between the distributions of the scores, which in turn would lead to compromised data if comparisons using the neutral classes were used. Conversely, the mean ratings of the negative and positive stimuli are approximately 5 SD different, thus leading to a high confidence that comparisons between negative and positive stimuli reflect two truly different groups of stimuli.

For the comparison of negative and positive stimuli in Experiment 6b, the ratings of neutral stimuli were excluded from analysis. This allowed for a clear and confident comparison of negative versus positive stimuli.

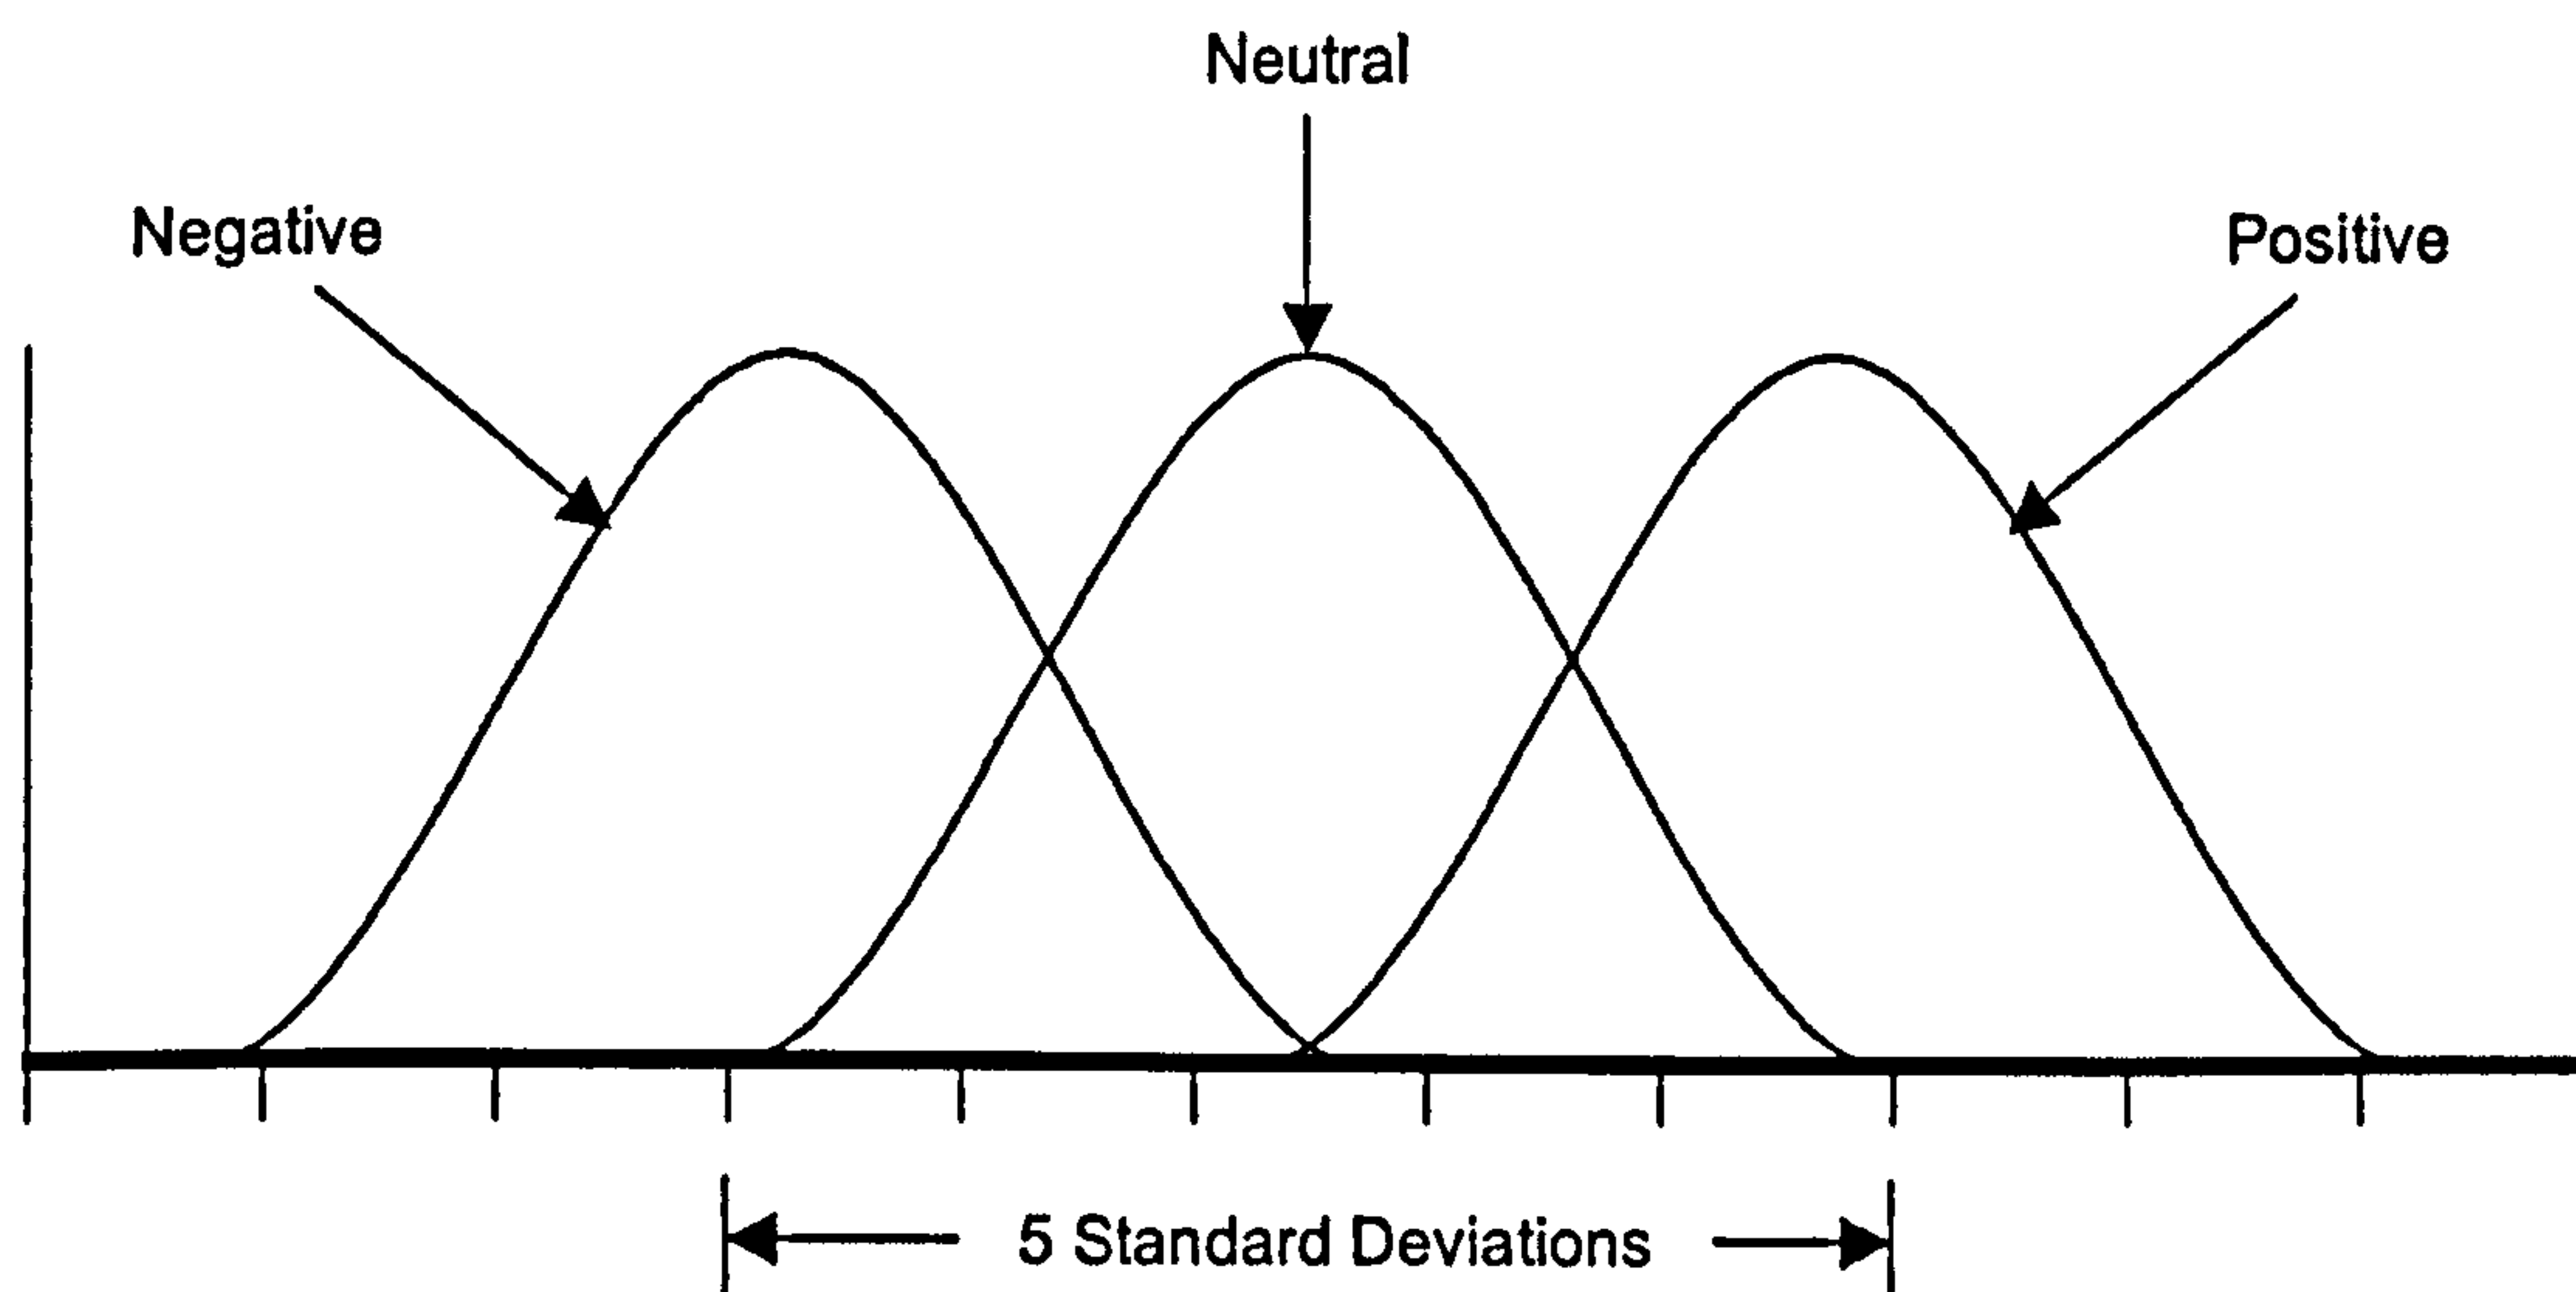


Figure 38: A diagrammatic representation of the overlap between negative, neutral and positive stimuli. Each mark on the x-axis represents one standard deviation. By removing neutral stimuli from the analysis, two sets of stimuli (negative and positive) with virtually no overlap remain in the analysis.

In order to investigate devaluation versus evaluative flattening of distractors in Experiment 6b, the attention effect was compared for negative and positive stimuli associated with a no-go cue. The mean ratings of negative and positive faces (as determined by Experiment 6a) are 3.26 and 5.25 respectively. The mean ratings of negative and positive scenes are 2.79 and 4.84. Assuming that the 7-point scale is used linearly by the participants (as they are instructed to do), an evaluation of 4.00 reflects a neutral evaluation. Both negative faces and scenes fall below this threshold and both positive faces and scenes fall above this threshold.

Therefore, both the 'devaluation-by-inhibition' and the 'flattening' hypothesis would predict that positive stimuli associated with a no-go cue (inhibited) will be devalued. For positive items, prior no-go stimuli should be rated as more negative than previously unseen stimuli according to both accounts.

The two hypotheses hold two different predictions for the evaluations of negative stimuli. The 'devaluation-by-inhibition' account predicts that negative stimuli associated with a no-go cue should be rated as more negative than previously unseen stimuli. Conversely, the 'flattening' account predicts that negative stimuli associated with a no-go cue should be rated as more *positive* than previously unseen stimuli: the ratings should trend towards a score of 4.00.

Experiment 6b: Devaluation of No-Go faces

Participants in Experiment 6b performed a go / no-go task. In the go / no-go task either a face or a scene image was presented one at a time at fixation. Go and no-go images were defined by their stimulus type (interior scene / exterior scene or female / male) and by a salient featural cue ('X' / 'O' letter presented on a scene or a blue frame / green frame displayed around a face). Two kinds of no-gos were presented. The first were featural no-gos, which were of the same stimulus type as the go but differed on the featural cue (i.e. if go was a male in a blue frame, a no-go was a male in a green frame). The second kind was a conjunction no-go, which had the same featural cue as the go but were of a different stimulus type (i.e. if go was a male in a blue frame, a no-go was a female in a blue frame). Go stimuli were presented on two thirds of trials, and so I hypothesised that the prepotent response would be to respond.

Items from the go / no-go task were subsequently presented, along with novel counterparts, for either emotional evaluation or as a memory test. I hypothesised that withholding a response to a featural no-go would be easy and only shallow processing

of the stimulus would be required, and thus I did not expect to observe distractor devaluation for featural no-gos. I hypothesised that withholding a response to a conjunction no-go would be more effortful, requiring deeper processing, and that distractor devaluation effects would emerge as a result.

Method

Participants

64 students from Bangor University participated in exchange for course credit or money. All participants gave informed consent and had normal / corrected to normal vision. Participants were assigned to one of four experimental groups (group 1a: 9 females, mean age 21.94 years; group 1b: 9 females, 19.25 years; group 2a: 9 females, 19.94 years; group 2b: 8 females, 19.81 years).

Apparatus

The apparatus used in this experiment was the same as that used in Experiment 6a.

Stimuli

The same face and scene images were used as in Experiment 6a. Two types of cue were used in conjunction with the images. When presented with a face, the cue was a 0.3° thick frame surrounding the image. The frame could be blue or green, depending on the experimental condition. When presented with a scene, a letter cue was used. Letter cues (height 0.8°) of an 'X' or an 'O' were overlaid in the centre of

the scene image, were always aqua and presented in Arial font (see Figure 39 for an example).



Figure 39: An example scene image and cue (not to scale).

Design and Procedure

The experiment lasted for approximately 30 minutes and consisted of 8 blocks, between which participants were allowed to take a break. Each block had three phases of 48 trials. Phase 1 was the attention task, phase 2 was the evaluation task, and phase 3 was the memory task.

Phase 1 contained interleaved face and scene trials (24 face trials, 24 scene trials) always starting with a face trial. The participants' task in this phase was to press the space bar to a go image, and withhold a response to a no-go image. Go and no-go images were defined by a conjunction of cue type (face: blue / green frame; scene: X / O) and image type (face: male / female; scene: interior / exterior).

There were three attention conditions: go (16 face, 16 scene trials), featural no-go (in which the image type matches the go stimulus; 4 face and 4 scene trials), and conjunction no-go (in which the cue type, not the image type, matches the go stimulus; 4 face and 4 scene trials). These conditions are defined in the table below (Table 15). Note that groups 1a and 2a, and groups 1b and 2b are the same at this stage. Conditions were randomly selected within each block.

Table 15: Go and no-go conditions definitions by group

	Face go / no-go task		Scene go / no-go task	
	1a and 2a	1b and 2b	1a and 2a	1b and 2b
Go	Blue Male	Blue Female	Interior O	Exterior O
Featural No-Go	Green Male	Green Female	Interior X	Exterior X
Conjunction No-Go	Blue Female	Blue Male	Exterior O	Interior O

Images were presented in the centre of the screen for 1500 ms with a fixation cross ITI of 1000 ms. See Figure 40 for an example trial sequence. Participants had to respond to a go stimulus within 1500 ms or it was counted as an error trial. If a participant made an error (slow or absent go trial response, or a failure to withhold a response to a no-go trial), an error tone sounded immediately as feedback.




Figure 40: Phase 1 trial sequence.

Following phase 1, participants evaluated 48 face images (phase 2). 24 of these were the faces seen in phase 1 (presented in the order of their previous exposure), and the further 24 (randomly interspersed) were novel faces. The images were presented without featural cues (frames) and the novel faces were selected to match the gender biases of the seen faces (20 go- and featural no-go-like, 4 conjunction no-go-like). Groups 1a and 1b evaluated the faces for ‘Trustworthiness’ and rated them on a 7-point scale (1 = Very untrustworthy; 7 = Very trustworthy). Groups 2a and 2b evaluated the faces based on their memory for them. Participants had a choice of three responses when presented with an image: “It’s a new image”, “maybe I remember it”, or “I know I remember it”.

In Phase 3 participants evaluated 48 scene images. Like the faces, 24 of these were seen in phase 1 (presented in the order of their previous exposure), and the

further 24 (randomly interspersed) were novel scenes. Scenes were presented without cues (letters) and the novel scenes were chosen to match the scene type biases of the seen scenes (20 go- and featural no-go-like, 4 conjunction no-go-like type). Groups 1a and 1b evaluated the scenes based on their memory for them and groups 2a and 2b evaluated the scenes based on their liking for them (where 1 indicates a strong dislike and 7 indicates a strong like).

Table 16 describes the phase conditions for each group.

Table 16: Summary of phase conditions for each group

Go	Faces Task		Scenes Task	
	Rate	Memory	Rate	Memory
Blue Male / Interior O	Group 1a	Group 2a	Group 2a	Group 1a
Blue Female / Exterior O	Group 1b	Group 2b	Group 2b	Group 1b

Results

Data Analysis

Error trials (failure to respond to, or being too slow to respond to a go trial, or failure to withhold a response to a no-go trial) were removed from the analysis. This eliminated 3.21% of data. Anticipatory errors were removed (an anticipatory error defined as responding to a go trial faster than 300 ms after stimulus onset). This eliminated a further 0.20% of data. Trials in which participants responded 3 *SD* faster or slower than their own mean RT to respond to that trial type (either faces or scenes) were removed. This removed a further 0.83% of data. Evaluations of stimuli that were made slower than 3000 ms were not included in the analysis. This excluded a further

3.87% of ratings of items presented in phase 1. Ratings of novel stimuli (which were not subjected to the previous data eliminations) that took longer than 3000 ms were also removed and excluded 4.33% of data.

Two participants were removed from the analysis as they failed to withhold a response to a conjunction no-go on more than 65% of trials. These two participants were in group 1b. One participant from group 1a was removed from the analysis as their mean RT to respond to a go trial was over 4 *SD* slower than other participants performing the same task. One further participant was removed in order to balance the n's of groups 1a and 1b. This participant was removed simply by virtue of being the last person tested in group 1b. One participant was removed from group 2a as they failed to respond to over 5% of go trials. This seems low, but the accuracy over go trials was exceedingly high: an error rate of 5% was over 8 *SD* from the mean accuracy of participants performing the same task. A final participant was removed from group 2b in order to balance the n's of groups 2a and 2b. Again, this person was the last person tested in group 2b.

The reported data is thus based on 58 participants: 14 in groups 1a and 1b, and 15 in groups 2a and 2b. Evaluations are based on an average of 122 go items, 31 featural no-go items, 27 conjunction no-go items, 155 go- and featural no-go-like novels and 31 conjunction no-go-like novels per participant.

Phase 1: Go / No-Go Responses

The accuracy to respond to a go stimulus, or to withhold a response to a no-go stimulus did not significantly differ between face and scene go / no-go tasks ($p > .1$).

For both face and scene go trials accuracy was near ceiling performance. The accuracy to respond to a go trial was significantly greater than the accuracy to withhold a response to a featural no-go trial, $t(1, 57) = 3.43, p = .001$. The accuracy to withhold a response to a featural no-go trial was in turn greater than the accuracy to withhold a response to a conjunction no-go trial, $t(1, 57) = 8.13, p < .001$. This suggests that it is more difficult to withhold a response to a conjunction no-go trial than to a featural no-go trial, and we can postulate that more inhibition is required to prevent a response to the former (see Figure 41).

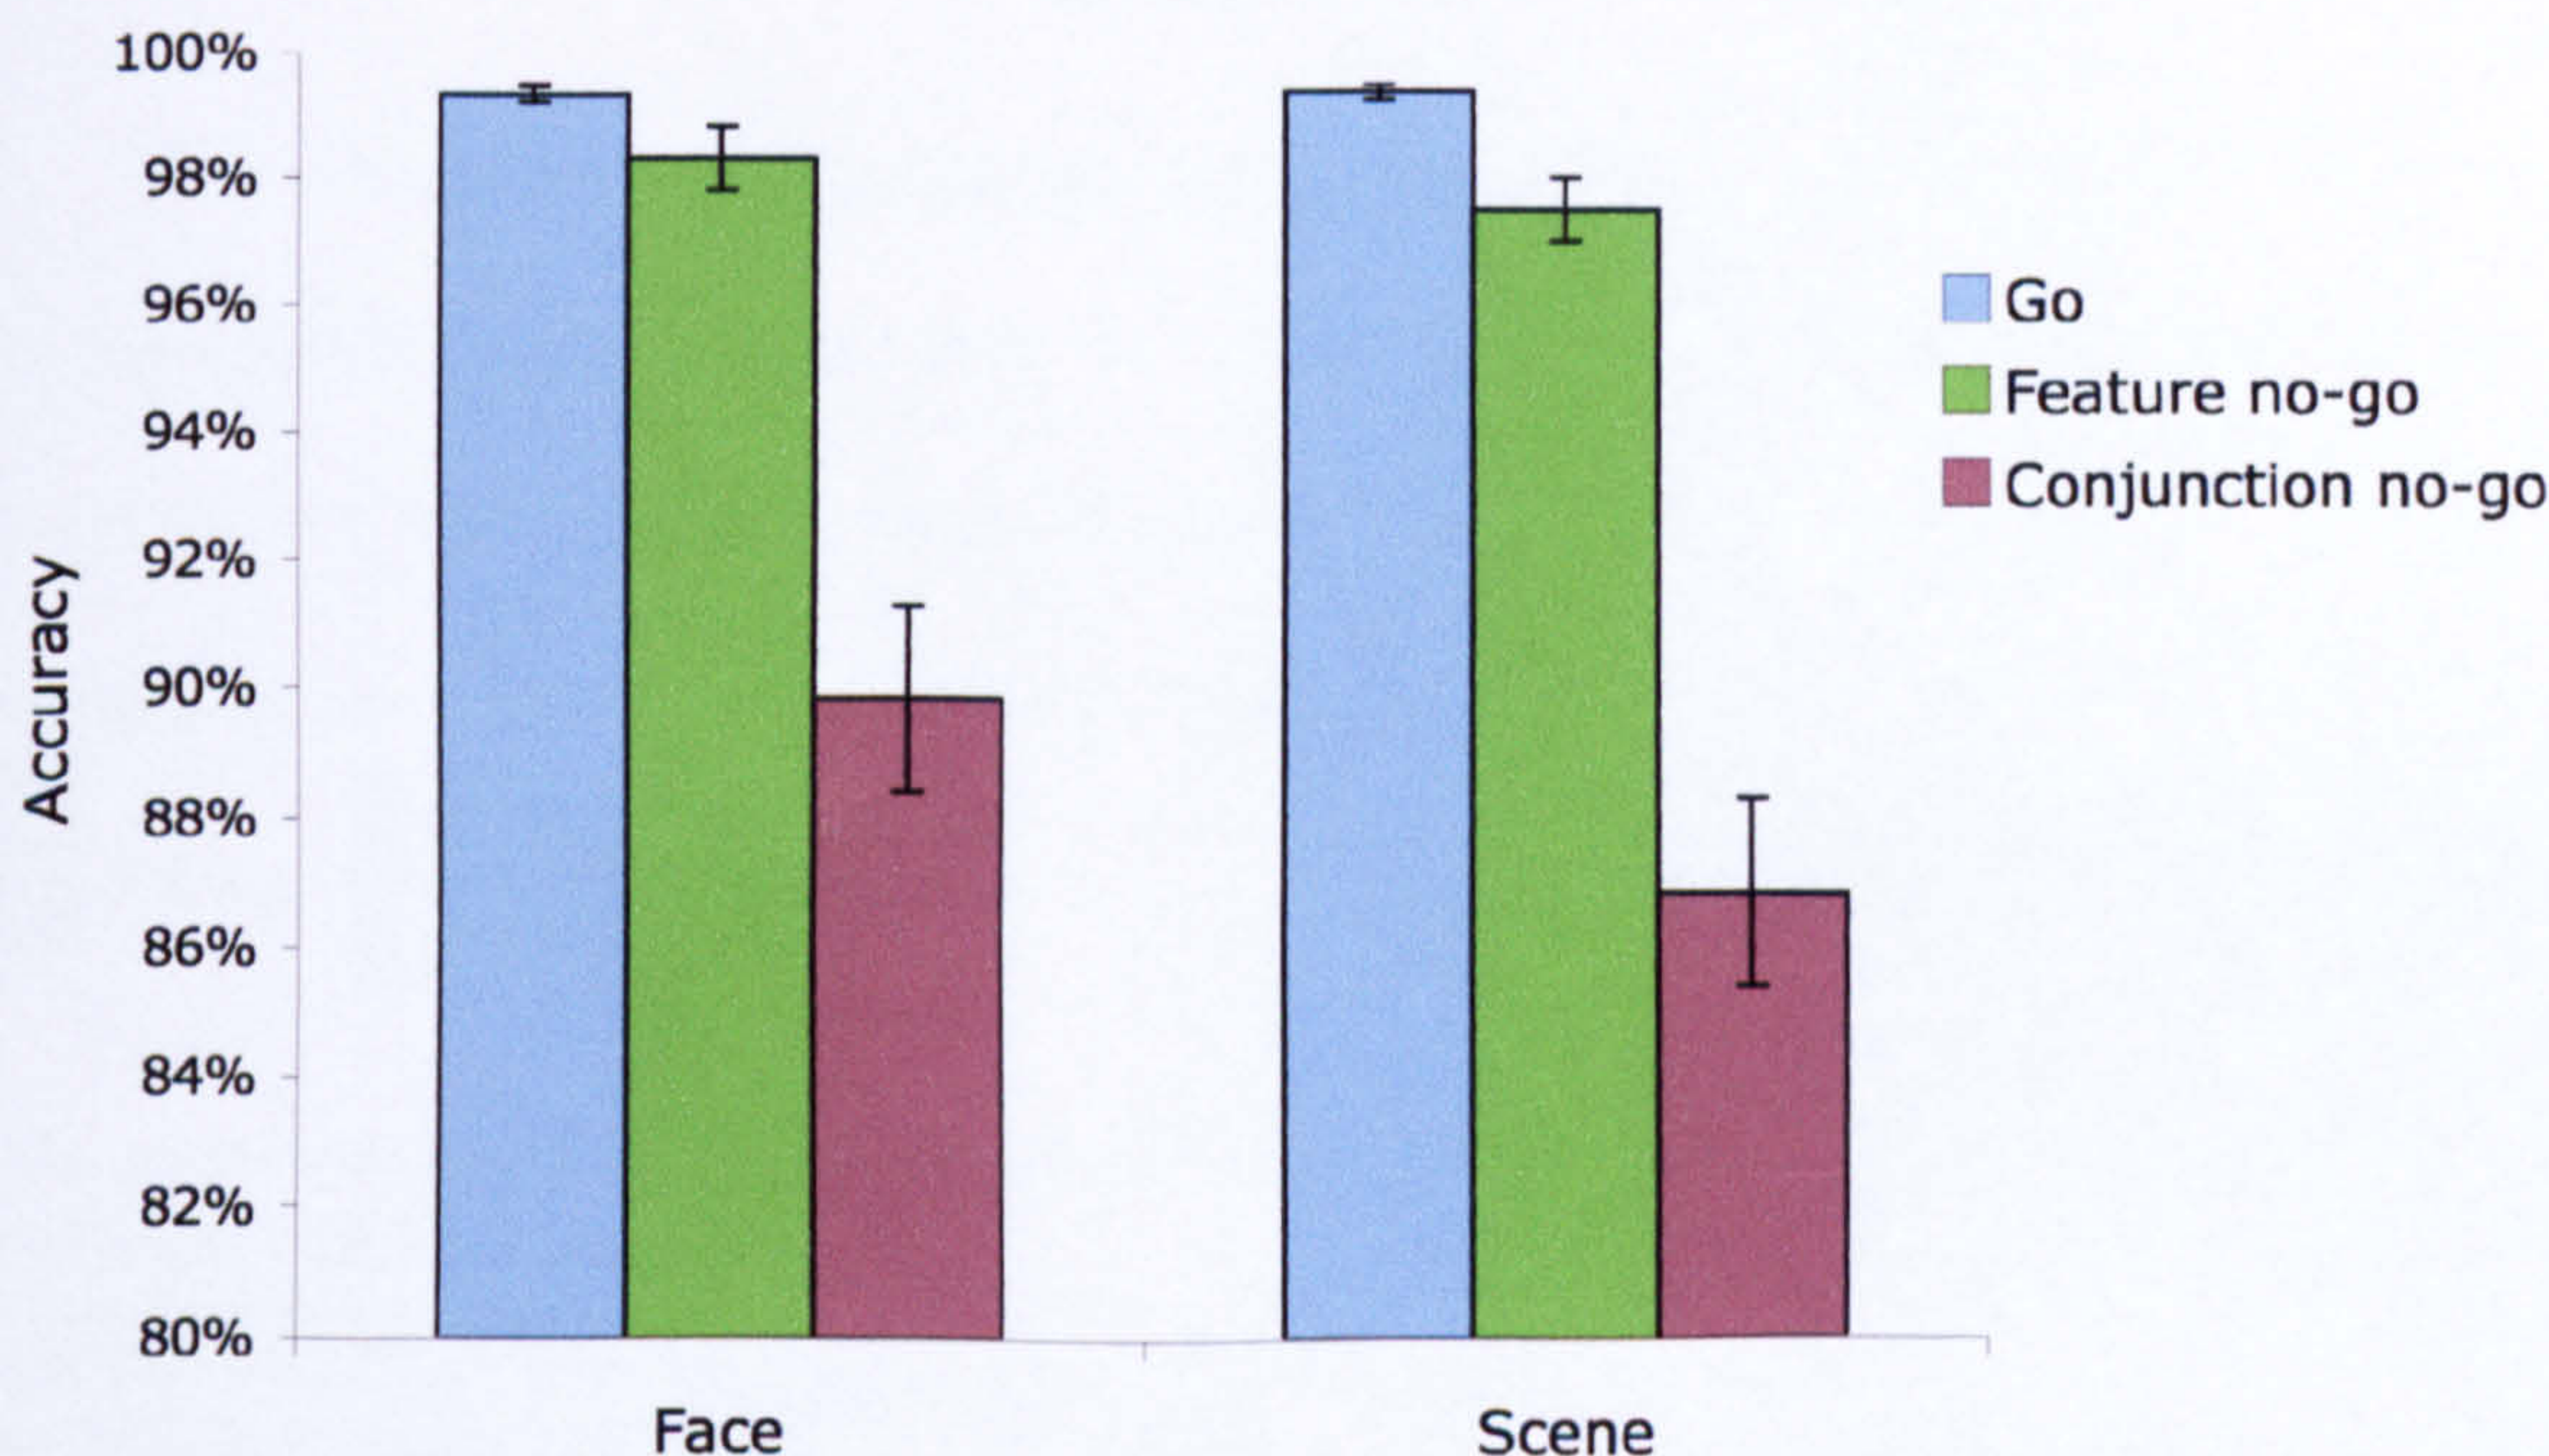


Figure 41: Mean accuracy to respond (or withhold a response) to the attention conditions in a face and a scene go / no-go sequence. Error bars represent ± 1 SE of the mean.

There were no differences between any of the groups in the error rates: faces (both males and females) and scenes (both interior and exterior) were affected equally by attentional conditions (all $p > .1$). An ANOVA confirmed a significant main effect of attention (go, featural no-go, conjunction no-go), $F(2, 57) = 72.01, p < .001$, but

no significant main effect of task type (faces, scenes) group or an interaction of these with attention (all $p > .1$)

Similarly, while the RT to respond to a stimulus in a go trial was significantly longer in a scene task ($M = 669$ ms, $SD = 72$ ms) than a face task ($M = 596$ ms, $SD = 69$ ms), $t(1, 57) = 3.54$, $p = .001$, there was no difference between the mean RT to respond to male and female go stimuli ($p > .1$), or interior and exterior scene go stimuli ($p > .1$).

Ratings

Raw ratings are presented in Table 17. The average raw ratings of male faces were significantly more negative than the average raw ratings of female faces, $t(1, 27) = 6.09$, $p < .001$. Likewise, the average raw ratings of interior scene stimuli were significantly more negative than the average raw ratings of exterior scene stimuli, $t(1, 29) = 3.20$, $p < .01$.

Table 17: Average raw ratings of Males, Females, Interior and Exterior Scenes. Table displays M and SD .

	M	SD
Male	4.27	0.89
Female	4.67	0.73
Exterior	4.25	0.83
Interior	3.39	0.80

Because of these category differences, subsequent results are presented in terms of the effect of exposure. That is, results are expressed as ratings from exposed stimuli per condition, minus the ratings of their novel counterparts. Ratings of go- and featural no-go-like novel stimuli are subtracted from the ratings of both go and

featural no-go stimuli. Ratings of conjunction no-go-like novels are subtracted from the ratings of conjunction no-go stimuli. Therefore, a positive rating difference indicates that the pre-exposed stimuli are rated more positively than unexposed stimuli (up-valuation), and any negative rating difference indicates that pre-exposed stimuli are rated more negatively than unexposed stimuli (devaluation).

Attention Effects on Emotional Evaluation

All results reflect rating difference as described above. An initial ANOVA using attention (go, featural no-go, conjunction no-go) as a within-group factor and stimulus type (face, scene) as a between-group factor was conducted. The ANOVA found that neither main effect was statistically significant (both $p > .05$), but that the interaction of attention and stimulus type approached significance, $F(1, 57) = 2.63, p = .08$. The effect of attention was thus explored separately for each stimulus type, using t-tests.

The important finding is that faces seen as conjunction no-gos in the go / no-go task are subsequently rated as less trustworthy than their novel counterparts. This effect approached significance, $t(1, 27) = -1.76, p = .09$. Both faces seen as go stimuli and faces seen as featural no-go stimuli were rated as significantly more positive than their novel counterparts, go: $t(1, 27) = 3.41, p < .01$, feature no-go: $t(1, 17) = 3.87, p = .001$.

It appears that no-gos that are inhibited on the basis of the gender of the face were devalued (but not significantly so). No-gos inhibited on the basis of a featural cue (frame colour) were rated as more positive than baseline, as were go faces. This difference between featural and conjunction no-go faces was significant, $t(1, 27) =$

2.98, $p < .01$, as was the difference between conjunction no-go and go faces, $t(1, 27) = 3.00$, $p < .01$. Featural no-go and go faces were not rated differently ($p > .1$).

No-go scene stimuli (both conjunction and featural) were not rated differently than baseline (both $p > .1$). Go scene stimuli were rated more positively than their novel counterparts, $t(1, 29) = 2.11$, $p < .05$. Go, featural no-go and conjunction no-go scenes were not rated differently (all $p > .1$). See Figure 42 for these rating difference results.




Figure 42: Mean attention effect per go / no-go condition. Scores reflect ratings of stimuli seen in the go / no-go task, minus the average rating of novel counterparts. A score above zero indicates that stimuli were evaluated as more positive following exposure in the go / no-go task; a score below zero indicates that stimuli were rated as more negative following exposure in the go / no-go task. Error bars represent ± 1 SE of the mean.

Devaluation versus Flattening

Trials in which a neutral stimulus was presented were removed from the analysis, eliminating a further one third of trials. As a result, rating and memory data are based on, on average, 81 go trials, 21 featural no-go trials, 19 conjunction no-go trials, 103 go- and featural no-go-like novel trials and 21 conjunction no-go-like novel trials. Half of the trials in each attention condition were negative stimuli, and half were positive stimuli.

The mean ratings for novel stimuli were first examined to ensure consistency of stimuli evaluation between Experiments 6a and 6b. Faces that participants from Experiment 6a had defined as positive were rated, by participants in Experiment 6b, as significantly more positive than faces Experiment 6a participants had defined as negative, $t(1, 27) = 8.97, p < .001$. Likewise, scenes that participants from Experiment 6a had defined as positive were rated, by participants in Experiment 6b, as significantly more positive than scenes Experiment 6a participants had defined as negative, $t(1, 29) = 6.93, p < .001$.

An ANOVA confirmed that this main effect of valence was significant for both faces, $F(1, 27) = 81.18, p < .001$, and scenes, $F(1, 29) = 51.98, p < .001$, and did not interact with group (male, female; interior, exterior; both $p > .1$). Thus, the stimulus valence categories could be further examined with confidence.

A repeated measures ANOVA on the rating differences of faces with valence (positive, negative) and attention (go, featural no-go, conjunction no-go) as within-group factors was performed. There was a main effect of attention, $F(2, 27) = 7.46, p$

< .01. However, the main effect of valence, and the interaction of valence with attention, were not (both $p > .1$). Further investigation of this data was conducted using t-tests.

The critical finding of this analysis is that both positive *and* negative conjunction no-go faces were devalued relative to their novel counterparts. In fact the devaluation of the negative conjunction no-go faces was larger than that of positive conjunction no-go faces. The devaluation effect of positive faces was non significant ($p > .1$), but the devaluation effect of negative faces was significant, $t(1, 27) = -2.67$, $p < .05$ (see Figure 43). This rules out the 'flattening' hypothesis as an explanation of the distractor devaluation effects found in experiments thus far.

Further investigation of the differences between attention effects in positive and negative faces revealed that for negative faces, go and featural no-go stimuli were rated as more trustworthy than their novel counterparts, go: $t(1, 27) = 2.29$, $p < .05$, featural no-go: $t(1, 27) = 2.27$, $p < .05$. For positive faces, the attentional state of a face during exposure in the go / no-go task did not affect their subsequent ratings relative to novel counterparts (all $p > .1$). However, paired t-tests revealed that the differences between negative and positive faces were non-significant (all $p > .1$). Therefore, positive faces were affected by attention during pre-exposure in the same way as negative faces, but to a lesser extent.




Figure 43: Mean attention effect per go / no-go condition for negative and positive faces. Scores reflect ratings of stimuli seen in the go / no-go task, minus the average rating of novel counterparts. A score above zero indicates that stimuli were evaluated as more positive following exposure in the go / no-go task; a score below zero indicates that stimuli were rated as more negative following exposure in the go / no-go task. Error bars represent ± 1 SE of the mean.

A repeated measures ANOVA on the rating differences of scenes with valence (positive, negative) and attention (go, featural no-go, conjunction no-go) as within-group factors was also performed. The main effects of attention, valence, and the interaction of attention with valence were all non-significant (all $p > .1$). Ratings of positive and negative scene stimuli exposed during the go / no-go task did not significantly differ relative to their novel counterparts (all $p > .1$).

Effort

The argument for devaluation-by-inhibition of the no-go stimulus is given further weight by a more in-depth analysis of the ratings for negative no-go faces. We can assume that the RT to respond to a go stimulus is an index of effort: a fast participant finds discrimination between go and no-go stimuli relatively easy, and a slow participant finds discrimination between a go and a no-go stimulus comparatively more difficult. It is likely, then, that a slow participant requires greater effort to inhibit a response to a no-go face, and then may be evidenced in their ratings to no-go faces.

Indeed, the mean RT to respond to a go face was compared to mean rating differences for featural no-go faces, and a significant correlation was found, Pearson's $r = -.656, p < .001$. The faster a participant is to make a response, the more likely they are to show mere exposure effects for featural no-go faces (see Figure 44). Conversely, the slower a person is to respond to a go face, the more likely they are to devalue a featural no-go face relative to novel counterparts. No such correlation was found for conjunction no-go faces ($p > .1$; see Figure 45) or for featural and conjunction no-go positive faces (both $p > .1$).




Figure 44: Mean rating differences for negative featural no-go faces. Scores reflect ratings of stimuli seen in the go / no-go task, minus the average rating of novel counterparts. A score above zero indicates that stimuli were evaluated as more positive following exposure in the go / no-go task; a score below zero indicates that stimuli were rated as more negative following exposure in the go / no-go task.




Figure 45: Mean rating differences for negative conjunction no-go faces.

Novel Ratings

The following analyses were conducted using the raw ratings. Interestingly, there was a main effect of novel typicality (typical novels being go- and featural no-go-like, and representing 83.33% of novel stimuli, atypical novels being conjunction no-go-like and representing 17.67% of novel stimuli). Stimuli belonging to an infrequently presented stimulus class (be it male, female, interior scenes or exterior scenes) are rated more positively than those of a frequently seen class, $F(1, 57) = 6.22, p < .05$ (see Figure 46). This effect did not interact with group ($p > .1$).




Figure 46: Mean raw ratings of typical and atypical novel faces and scenes. Error bars represent ± 1 SE of the mean.

Memory Data

The data presented here are expressed as d' prime scores. 'Know' responses were used in the analysis. A false alarm is defined as responding 'Know' to a novel

stimulus and a hit is defined as responding 'Know' to a pre-exposed stimulus. In cases where either hit or false alarm rate were either 1 or 0 (and so z scores could not be calculated), scores of 0 were thus converted to 0.001, and scores of 1 were converted to 0.999. Corrected scores were subsequently transformed into z scores, and the z scores of false alarms were subtracted from the z scores of hits. Therefore, the higher the score, the better the memory.

Memory for go, featural no-go and conjunction no-go faces and scenes were analysed. A repeated measures ANOVA with attention (go, featural no-go, conjunction no-go) and stimulus type (face, scene) as factors found a significant main effect of attention, $F(2, 57) = 16.06, p < .001$; the main effect of stimulus type and the interaction of stimulus type with attention were non significant (both $p > .1$). Featural no-go stimuli were remembered more poorly than either go stimuli, $t(1, 57) = -5.22, p < .001$, or conjunction no-go stimuli, $t(1, 57) = -3.96, p < .001$. Memory for go stimuli and conjunction no-go stimuli did not differ ($p > .1$). Overall, memory for faces and scenes did not differ ($p > .1$). Figure 47 displays the memory scores for faces and scenes in each attention condition.




Figure 47: Memory scores (*d Prime*) for faces and scenes as a function of their go / no-go condition. A high score indicates good memory. Error bars represent ± 1 SE of the mean.

Memory scores for positive and negative stimuli were compared in a repeated measures ANOVA with valence and attention as within-group factors. This was done separately for faces and scenes. For the faces, there was a main effect of valence, $F(1, 27) = 43.69, p < .001$, but no significant main effect of attention or an interaction of the two factors (both $p > .1$), see Figure 48. Negative faces were remembered better than positive faces, $t(1, 27) = 4.40, p < .001$. For the scene stimuli, there was no difference in the memory scores for positive and negative scene stimuli ($p > .1$), but the main effect of attention was preserved, $F(1, 29) = 5.03, p = .01$. There was no significant interaction between these two factors ($p > .1$).




Figure 48: Memory scores (d Prime) for negative and positive faces. A high score indicates good memory. Error bars represent ± 1 SE of the mean.

Discussion

Distractor devaluation of conjunction no-go's

Once again, the distractor devaluation effect has been replicated. Faces which must not be responded to, and are defined as no-go's by their gender, are rated as more negative than their novel counterparts. Faces which are go stimuli, and must be responded to, are subsequently rated as more positive than their novel counterparts, suggesting that a mere exposure effect has also occurred. Curiously, faces which are defined as no-go's by their frame colour are not subsequently devalued relative to novel faces.

No distractor devaluation for featural no-go's

From the error rate data, we can see that withholding a response to a conjunction no-go is more difficult than withholding a response to a feature no-go. Conjunction no-go faces get mistaken for go faces more often than do featural no-go faces. By implication, this means that participants are engaging in a 'frame-based' strategy. This is the dimension that go faces and featural no-go faces differ on, and that go faces and conjunction no-go faces are the same on. This strategy is sensible, as colour 'pops out', responding on the basis of this dimension allows for fastest go responses.

If participants reject a stimulus based on the frame, there is no need for the participant to ever even look at a face that is in a green frame. Therefore inhibition may not be tied to the face identity. Shallow processing only may be required, and this may be the reason that on a subsequent encounter with that face, no distractor devaluation effects emerge.

However, in order to successfully withhold a response to a conjunction no-go, the face *must* be looked at and categorised, i.e. deeply encoded. It is the gender of the face which defines it as a no-go. Therefore when participants reject a face based on gender, inhibition is tied with identity of the face, and distractor devaluation effects result.

This theory is supported by the memory data. Conjunction no-go faces are remembered better than featural no-go faces. This suggests that in the case of conjunction no-go faces, the face itself has been looked at and encoded, whereas for

featural no-go faces this is less likely to have been the case. Attention (and subsequently inhibition) may have been drawn to face identity in conjunction no-go conditions, but not in featural no-go conditions. Go faces are also remembered better than featural no-go faces. To decide that a face is a go face the participant must also look at the identity, as it shares its frame colour with the conjunction no-go faces. Conjunction no-go faces are likely to have been remembered better than go faces simply because there were fewer exemplars of the category.

However, although go faces are remembered better than featural no-go faces, featural no-go faces are rated as more positive than go faces. There may be several reasons for this, I will speculate on three. First, processing duration may be longer for no-go stimuli. After deciding that a face is a go face, participants may look down to their response key while making their response (no action is required in the case of a no-go). This may decrease exposure time to a go face and result in less fluency, and subsequently less mere exposure. Second, it may be that after the decision has been made to respond to a go face, the participants switch their attention to making the response, and therefore have less attentional resources for further processing of the go face, compared to a featural no-go face which is free to be further processed for the full duration of exposure. This again may lead to less fluency, and less mere exposure effect. Third, it may be that this difference is highlighting the dichotomy between analytic and non-analytic evaluation as put forward by Whittlesea and Price (2001). It may be that having to engage in a gender decision about a face engages analytic processing, which results in an increased memory for the face (an analytic evaluation), However, passive viewing of the face (as may be occurring for featural

no-go faces) may allow non-analytic processing to prevail, impacting future affective evaluations (which are non-analytic).

What is clear is that a featural no-go face must have been encoded to some extent, as it benefits from mere exposure. However, it appears to have been encoded without an inhibitory tag. Therefore it does not subsequently suffer from ‘devaluation-by-inhibition’.

In this experiment, the frame was not presented at evaluation. If a featural no-go is indeed rejected on the basis of its featural cue, perhaps inhibition is in some way tied to the feature of the frame, rather than the face itself. Perhaps if the faces had been presented for evaluation with their frames, distractor devaluation would have occurred for featural no-go faces. Indeed, this is supported by recent findings from our Lab. Goolsby et al. (in press; Experiment 1³) presented participants with a pair of colour-tinted faces. Participants were required to search for a target face, defined by its colour tint. A greyscale version of the previous target, previous distractor or a novel face was presented for evaluation. The authors found that the previous distractor faces were not devalued when presented without the colour cue. They concluded that the inhibition had been stored with the colour tint feature in memory, and that devaluation would only result if that feature were present at evaluation. This is in line with the findings of no devaluation for featural no-go’s in the current experiment.

³ I will discuss the findings of Goolsby et al. (in press) in further detail in the General Discussion (Chapter 10).

Why was this not the case in our previous experiment (Fenske et al., 2005)?

There, faces were solely rejected on the basis of featural cues (red or green semi-opaque masks). Participants never had to engage in any face-specific categorisation in order to learn that a face was a no-go. Why is it that distractor devaluation occurred for the featural no-go faces in the Fenske et al. experiment (which were also presented for evaluation in greyscale), and not in the current experiment?

There are three possible explanations for this. First, in the Fenske et al. study, we allowed participants less time to respond to go face (1650 ms compared to 2000 ms in the current study). It may have been that participants were more urgent to respond to go faces, making the pre-potent response to press a key even stronger. When a no-go cue appeared over a face, participants may have had to work harder (and thus apply more inhibition) in order to overcome this pre-potent 'respond' reflex.

The second possibility is that in the Fenske et al. experiment, faces were displayed for 1000 ms without a cue, 200 ms with the cue, and then a further 1000 ms without the cue. Displaying the faces without the go or no-go cues during the go / no-go task may have allowed, and even encouraged face identity encoding. Cues 'flashed' onto the faces, and so were not a part of the face representation. In the current study, conversely, featural no-go faces never appeared without their featural cues (frames). Therefore, there was no opportunity to encode the face prior to featural cue onset and the resulting rejecting of the face. Furthermore, the face and the frame may have been encoded as a whole object, therefore when participants were presented with only the face for evaluation the 'inhibitory tag' may not have been associated with the particular object representation. It may have been the combination of the face

plus the frame that would have reinstated the inhibition, which would have led to distractor devaluation.

Finally, in the Fenske et al. experiment, participants were presented with a pair of faces: one to the left, and one to the right of fixation. On presentation of the featural cue, then, participants would have had to orient their attention (either covertly or overtly, exogenously or endogenously) to the face with the cue. This increased attentional activity may have required a reactive increase in inhibition in order to withhold the response to a no-go cue. This increase in resultant inhibition may have subsequently led to distractor devaluation. In contrast, in the present study, all faces are presented singly at fixation. In this case, there is no need for an orientation of attention. This may have made the processes of inhibition easier, and weaker. Therefore, inhibitory tags stored with face representations may also have been weaker, resulting in weaker distractor devaluation effects.

Effort

It is possible to further speculate on the effect of effort in this experiment. Specifically, whether an increase in the effort required to inhibit a no-go stimulus increased the resultant distractor devaluation. As previously discussed, more effortful inhibition results in greater distractor devaluation effects (Kiss et al., 2007). In the present experiment, I used participants' mean RT to respond to a go face as a measure of overall 'effort'. A fast RT may be the result of an efficient use of the frame colour as a cue to reject a stimulus for response and classify it as a no-go item. Therefore, little or no inhibition is required for stimuli identity in a featural no-go condition. A slowing of RTs though, suggests participants are finding rejection of no-go's more

effortful, and participants with longer RTs may be applying greater inhibition to featural no-go faces. This would lead to devaluation even in the case of these faces which have cues for rejection unrelated to the stimulus itself.

Indeed this appears to be this case. Participants with fast RTs are more likely to show mere exposure effects for featural no-go faces, but participants with slow RTs are more likely to show distractor devaluation effects for featural no-go faces. It appears that slow participants are unable to efficiently reject featural no-go faces on the basis of their frame colour, and are engaging in a more effortful processes involving inhibiting the face itself. Therefore the inhibition appears to be encoded with the featural no-go face representation, causing subsequent devaluation on reinstatement. However, in a go / no-go task, there is no RT *not* to respond. Therefore we have to infer effortful inhibition based on a whole-participant strategy. A trial-by-trial RT to devaluation correlation would be a safer measure of this effect, however the current results suggest that this is the case.

Flattening versus Devaluation

This experiment, I believe, has convincingly demonstrated that distractor devaluation effects are indeed the result of devaluation, and not a flattening of affect. Until now, this issue had not been addressed. In the current experiment, faces that were both positive *and* negative were devalued as a result of being a no-go in a go / no-go task. Faces that had been evaluated as negative by a different group of participants were rated as more negative when they had been seen as a conjunction no-go in a go / no-go task. This is as predicted by the ‘devaluation-by-inhibition’ hypothesis. A ‘flattening’ hypothesis would conversely predict that faces that had

been pre-rated as negative would be rated as more positive when they had been seen as a conjunction no-go.

In the current experiment, a 7-point Likert scale was used for evaluation. Assuming that the scale was used linearly by the participants (as they had been instructed to use it) a score of 4.00 would represent the mid-point. However, let us assume for the sake of argument that 4.00 was not the mid-point and that participants had anchored the scale more negatively, so that now a lower score, say 2.00 represents an affectively neutral evaluation. We could argue that this is the reason that both positive and negative faces were 'devalued'. However, the fact that negative faces actually demonstrated *more* devaluation following inhibition than positive faces argues directly against this. Even if the anchor point had been 2.00, a flattening account would expect that devaluation effects would increase the further away from 2.00 a face had started. A face that was originally evaluated as 6.00 has further to go to the 'mid-point' than face that was originally rated at 3.00. We would thus expect that the 6.00 face would show more devaluation. This was not this case, and as such the current results provide very strong support for the 'devaluation-by-inhibition' hypothesis.

Why were negative faces devalued more than positive faces?

Comparing emotionally valenced as opposed to emotionally arousing stimuli in their resulting 'distractor devaluation' effects was not the point of using positive and negative stimuli in this experiment. The stimuli were split in order to test the 'flattening' hypothesis. However, that negative faces seen as conjunction no-go's in

the go / no-go task were devalued more than positive faces is interesting, and there are several possible explanations for this which I will now speculate on.

The first is that negative faces may have attracted more attentional resources than did positive faces. As discussed in Chapter 1, emotionally salient stimuli capture and hold attentional resources relative to emotionally neutral stimuli. Several researchers have suggested that this occurs to a greater degree for negative stimuli than for positive stimuli (e.g. Eastwood, Smilek & Merikle, 2001; Fox, Russo, Bowles & Dutton, 2001). The Evolutionary Threat, and the Negativity Bias hypotheses were put forward to account for the findings. It is possible that such mechanisms are at work in the current study. Perhaps when a negative face (relative to a positive face) was presented as a conjunction no-go, it initially attracted more attentional resources, which resulted in greater attentional effort required to inhibit the face once it was classified as a no-go. The resulting inhibitory tag with the object representation may have been stronger (relative to positive faces) thereby resulting in greater distractor devaluation. However, as discussed in Chapter 1, it seems more likely that arousal, and not negativity per se is responsible for greater attentional capture (e.g. Arnell, Killman & Fijavz, 2007). While it is likely that, overall, negative faces are indeed more arousing than positive faces, no 'arousal' evaluations were collected in the current study. It is more than likely that many faces that were positive were also high in arousal, and many that were negative were low in arousal. Without arousal scores to complement the valence scores, it is unwise to speculate on this further. In addition to this, and arguing against this account of the findings, negative featural no-go's were 'up-valued' following exposure in the go / no-go task more than were positive featural no-go faces. If attention had been initially attracted more to

negative versus positive faces, negative faces should have been inhibited more, resulting in less mere exposure effect on subsequent evaluation.

All participants in this study were instructed to evaluate the faces based on how trustworthy they considered them to be. However, many studies have shown that there is a high correlation between 'trustworthiness' ratings and 'attractive' ratings. The 'beauty premium' describes that physical attractiveness has several benefits (Hamermesh & Biddle, 1994). Attractive people receive more favourable treatment in hiring and promotion (Landy & Sigall, 1974; Dipboye, Arvey & Terpstra, 1977). Attractive people are rated as more intelligent, more extraverted and more socially skilled (Dion, Berscheid & Walster, 1972). Attractive people are consistently accredited with more desirable traits than unattractive people, and this is likely to include 'trustworthiness' (although to my knowledge no studies have investigated this).

Assuming then that 'trustworthy' faces are also 'attractive', an aspect of attractive faces becomes particularly important. As discussed in Chapter 1, attractive faces may be 'average', (Langlois & Roggman, 1990) and more likely to conform to a mentally constructed prototype face (Rosch, 1978). Faces may be judged as more attractive because they represent better examples of face categories, and are therefore easier to classify (Johnston & Ellis, 1995). Extending this, attractive faces may be 'attractive' because they are easier to classify, and are better examples of male and female categories (O'Toole et al., 1998). Indeed, Hoss and colleagues demonstrated that RT and accuracy to identify the sex of faces was faster and more accurate for attractive, compared to unattractive faces (Hoss, Ramsey, Griffin & Langlois, 2005).

If gender discrimination is easier for attractive faces, we can postulate that in the current study it was easier for positive faces. In this case, a positive conjunction no-go would have been easier to categorise and reject for responding than a negative conjunction no-go. A negative conjunction no-go would have been more difficult to classify by gender, and may have competed for a response more strongly, thereby requiring more inhibition. If this is the case, then the stronger inhibitory tag stored with the object's representation would have resulted in the stronger distractor devaluation effects found for negative faces in Experiment 6.

A second implication of 'averageness' theory is that an average face may be more poorly remembered. I discussed in Chapter 1 that participants may misinterpret an attractive face as prior exposure when probed in a memory test: participants are more likely to respond that an attractive, but previously unseen face, has been seen before, than an unattractive face (Solso & McCarthy, 1981; Bomba & Siqueland, 1983). In the current study, by using *d* primes which take into account the false alarm rates as well as the hit rates, I was able to determine that memory for positive (and therefore attractive faces) was poorer than memory for negative (unattractive) faces. If unattractive faces are less likely to conform to an 'average', this suggests that unattractive faces will be better individuated, and thus better encoded in memory. Extending this, we can postulate that in turn, the inhibitory tag associated with an unattractive face will be better encoded, and will thus be more likely to lead to 'devaluation-by-inhibition' on a subsequent encounter with the unattractive face.

This idea links back to the 'individuation hypothesis'. An object must be successfully individuated if inhibition is to be stored with the object's representation

in memory. However the 'individuation' effects found for negative faces are only implied. This is a post hoc test of individuation; an a priori, within category test will be a more convincing test on which to base individuation theory (I address this in Experiments 7 and 8).

Scene versus object recognition

The findings above suggest that memory for an object is an important step in the 'devaluation-by-inhibition' process. However memory is not everything. Scene stimuli were remembered, and even demonstrated the same pattern of results as faces in that conjunction no-go's were remembered better than go's, which were remembered better than featural no-go's. Why is it that no distractor devaluation effects were found for scene stimuli?

Compared to our knowledge of object processing, what is known about scene processing is relatively limited. Biederman (1988) suggested that his 'geon' structural model of object recognition might be extended to scene recognition. He proposed that 3-D primitives with a larger spatial scale than those used to represent objects could represent scene specific information independently from object information. Thus, a scene might be represented as a large object, in which case we might have expected to have found distractor devaluation. If the inhibitory tag could have been associated with a single 'scene' representation in memory, it may have subsequently led to 'devaluation-by-inhibition'. However, such a representation of a scene seems unlikely.

Essentially, a scene can be thought of as a collection of objects rather than as an object in its own right. However, it is more than this; a scene describes object 'context'. Scene categorisation is possible with very brief visual presentations (Potter & Levy, 1969; Potter, 1975, 1976; Biederman, Mezzanotte & Rabinowitz, 1982; Schyns & Oliva, 1994; Intraub, 1997; Oliva & Schyns, 1997, 2000). This is taken as evidence for fast underlying mechanisms that could be performed simultaneously or even precede object identification.

Scene processing also appears to take place in a distributed system which is largely separated from the object system. Specifically, the parahippocampal area is thought to be responsible for processing the spatial layout of a scene (Epstein & Kanwisher, 1998; Epstein, Graham & Downing, 2003). It is thought to mediate, in conjunction with the retrosplenial cortex, both spatial and non-spatial contextual processing (Bar & Aminoff, 2003). As such, it appears that a scene is not represented in memory in such a way that allows encoding of inhibition with its representation. Instead, it may be that inhibition must be encoded with an object in a scene, and not the context in which it appears.

Novelty preference

In Experiment 6 'typical' novels (both faces and scenes) were rated more negatively than 'atypical' novels. Participants prefer a stimulus class that is presented infrequently throughout the experiment. The typical class of stimuli was presented five times more often throughout the experiment than the atypical class. It may be that the class of stimuli seen most frequently throughout the experiment becomes associated with a general negative affect as participants start to lose interest in the

experiment. This demonstrates the importance of comparing ratings of 'seen' stimuli to their novel counterparts. By using this comparison, we can assess the effects of attention in the go / no-go task, rather than the effects of exposure.

This atypicality preference appears to be in contrast to the group mere exposure effects found for novel faces in Experiment 5. However, there is an important difference between the two effects. In the oddball task, the stimulus group exposures varied within each trial on a ratio of 2 distractors to 1 target. After each trial we might expect a 'boost' in ratings of novel stimuli belonging to the stimulus group with more exemplars exposed (distractor-like novels). Importantly, across the course of the oddball experiments, the group exposures were matched. There was no reason for a 'negative affect' to become associated with one stimulus class or another. This is in stark contrast to the current experiment in which typical stimuli are seen across the whole experiment five times more than atypical stimuli.

Chapter 7

Devaluation of Own- versus Other-Race Faces

In Experiments 7a and 7b I wanted to take advantage of the Other Race Effect (ORE) in order to investigate individuation versus categorisation in an a priori way, within a single category (i.e. faces). The ORE refers to the well-studied phenomenon that faces of a different race are more poorly recognised than own-race faces, an effect that has been widely replicated (Malpass & Kravitz, 1969; Cross, Cross & Daly, 1971; Chance & Goldstein, 1981; Anthony, Cooper & Mullen, 1992) and is robust across racial groups and research paradigms (Meissner & Brigham, 2001).

Two theories have been proposed to account for the ORE. Perceptual-expertise models (Meissner & Brigham, 2001) posit that racial segregation leads perceivers to have differential expertise in processing faces of their own race versus faces of another race (although see Ng & Lindsay, 1994, for a variation on this hypothesis). This difference in expertise leads to differences in recognition accuracy. Because perceivers have had less opportunity for processing other-race faces than own-race faces, they are less expert at distinguishing between other-race faces.

The mechanisms by which differential expertise might lead to poorer recognition accuracy are a matter of debate. One line of thinking is that the lack of contact leads to a lack of expertise with the dimensions on which other-race faces actually vary (MacLin & Malpass, 2001). Another is that lower levels of expertise with other-race faces elicit less holistic processing (which in the case of face perception is believed to be a hallmark of efficient processing) and more feature-based processing, compared to own-race face perception (Rhodes, Brake, Taylor & Tan, 1989; Michel, Rossion, Han, Chung & Caldara, 2006).

The perceptual-expertise hypothesis has received empirical support. For example, Sangrigoli, Pallier, Argenti, Ventureyra and de Schonen (2005) found that Korean individuals, who had been adopted as children by Caucasian families, showed a reversal of the ORE by adulthood. To these individuals, Caucasian faces were treated as own-race faces and Korean faces as other-race faces.

However, differential expertise alone does not appear to be sufficient to account for the ORE. For example, if participants are instructed prior to the experiment that they are likely to show a racial bias at recognition, and that they should attend to the individuating characteristics of faces, the ORE can be eliminated (Hugenberg, Miller & Claypool, 2007). Also, MacLin and Malpass (2001, 2003) found that adding either Latino- or Black-stereotypic hairstyles to racially ambiguous Latino-Black faces was sufficient to eliminate the ORE in Latino and Black observers. Moreover, Michel, Corneille and Rossion (2007) demonstrated that identical racially ambiguous faces were processed more holistically after they had been categorised as own-race, compared to when they had been categorised as other-race.

This research has contributed to the development of the social-cognitive model of the ORE. This model posits that the mechanisms of individuation and categorisation (Sporer, 2001) are differentially applied to own- and other-race faces. The crux of social-cognitive theory is that people tend to think categorically about out-group members (Bodenhausen, Macrae & Hugenberg, 2003). Therefore, according to social-cognitive models, the ORE is due to differences in social

cognitions elicited when processing in-group and out-group members (see Sporer, 2001 for a review, discussed further in Chapter 8).

Social-cognitive models also vary in the underlying mechanisms proposed to account for the ORE. For example, Levin's (1996, 2000) feature-selection model posits that there is an asymmetrical search for features in own- versus other-race faces. Thus, although people tend to encode the individuating features of own-race faces, they tend to encode race-specifying features (e.g. skin-tone) of other-race faces, at the expense of individuating information (see also MacLin & Malpass, 2001, 2003). Alternatively, categorisation of a face as an out-group member might reduce the motivation to further process the face, leading to weaker encoding of the individuating features of other-race, compared to own-race faces (Rodin, 1987).

In the following experiment, I aimed to take advantage of this ORE to investigate distractor devaluation for individuated (own-race) versus categorised (other-race) faces. I hypothesised that own-race faces, which were likely to be individuated, would be devalued on a one-to-one basis, as an inhibitory tag could be successfully stored with the face's representation in memory. On the other hand, I predicted that other-race faces, which were likely to be categorised but not individuated, would receive whole-category devaluation. I hypothesised that, in the case of other-race faces, an inhibitory tag would not be able to be stored with an individual face in memory, and would instead be stored with the category representation.

I was also interested to investigate the level to which distractor devaluation would generalise. Would generalisation occur at a basic (task-relevant) level of categorisation, or would generalisation only extend to members of the same subordinate category as the inhibited faces?

Experiment 7a: Devaluation of Own- and Other-Race Faces – Race Task

In Experiment 7a I used four categories of faces (Asian: male/female and Caucasian: male/female) and asked participants to engage in a go/no-go task using race as the go cue but only ever giving them no-go stimuli that differed from go stimuli on both race and gender. I then had them evaluate faces seen in the go / no-go task and also faces of all four categories that they had never seen before.

Method

Participants

Sixteen Asian (mean age, 21.3 years, 12 females) and sixteen Caucasian (mean age, 19.0 years, 13 females) Bangor University students participated in exchange for course credits or £5. The participants gave informed consent prior to participating and were naïve to the aims of the experiment.

Apparatus and Stimuli

The experiment was conducted on a Pentium III computer with a 33 cm colour monitor (100 Hz, 1024 x 768 resolution) running E-Prime 1.1 software (Schneider, Eschman & Zuccolotto, 2002). The stimuli appeared on a uniform white field at a 70 cm average viewing distance. Fixation appeared in black 18-point Arial font.

Items in the go / no-go sequence were full-colour headshot photos of young adults (approximately aged 18-30). Photos were taken from copyright-free webpages on the Internet and cropped where necessary. Faces subtended 5.7° in height, 3.4° in width and had neutral or smiling expressions.

Stimuli were of four types: Asian females (AF); Asian males (AM); Caucasian females (CF); and Caucasian males (CM), see Figure 49 for examples of each.



Asian Females
(AF)



Asian Males
(AM)



Caucasian Females
(CF)



Caucasian Males
(CM)

Figure 49: Types of stimuli used.

Procedure and Design

The general procedure and the sequence of trial events are shown in Figure 50. There were two phases in each block of trials. Phase 1 was the go / no-go task; phase 2 was the evaluation task.

In the go / no-go phase, 10 faces were presented individually and sequentially, each for 1500 ms (one block). The task was to depress a key as quickly as possible when one of the stimuli matched a pre-defined 'go' category and to withhold responding when it did not (no-go category). For each block, the 'go' category was defined by race (e.g. Asian), but race and gender were fully redundant with each other

(unbeknownst to participants) because the no-go category shared neither race nor gender with the go category. For example, if the go category was Asian, all the Asian faces in the block were female (or male), and all the no-go faces were male (or female).

Participants were told to go (respond) to Asian faces in 4 blocks, and to Caucasian faces in another 4 blocks of trials. Stimuli were displayed for 1500 ms regardless of response. Audio feedback was given for errors (failure to respond to a go; failure to withhold a response to a no-go; a response that was > 1500 ms). There were 5 go and 5 no-go faces in each block (presented in a pseudo random order). There were 8 blocks in total: the presentation order was counterbalanced.

After each phase 1 block participants viewed another sequence of faces, this time rating each for Trustworthiness on a 7-point scale (where 1 is very untrustworthy, and 7 is very trustworthy). Each face was displayed in the centre of the screen and remained displayed until a response was made. In each phase 2 block, 30 faces were presented: 5 faces in each of 6 conditions. Condition 1, Go faces (all the go faces seen in the preceding phase 1 block); condition 2, No-go faces (all the no-go faces seen in the preceding phase 1 block); condition 3, Two-match Go novels (novel faces from the go category in the immediately preceding phase 1 block); condition 4, Two-match No-go novels (novel faces from the no-go category in the immediately preceding phase 1 block); condition 5, One-match Go novels (novel faces that match the go category in the task-relevant dimension only, i.e. race); condition 6, One-match No-go novels (novel faces that match the no-go category in the task-relevant dimension only, i.e. race). Table 18 depicts these conditions and the number of trials

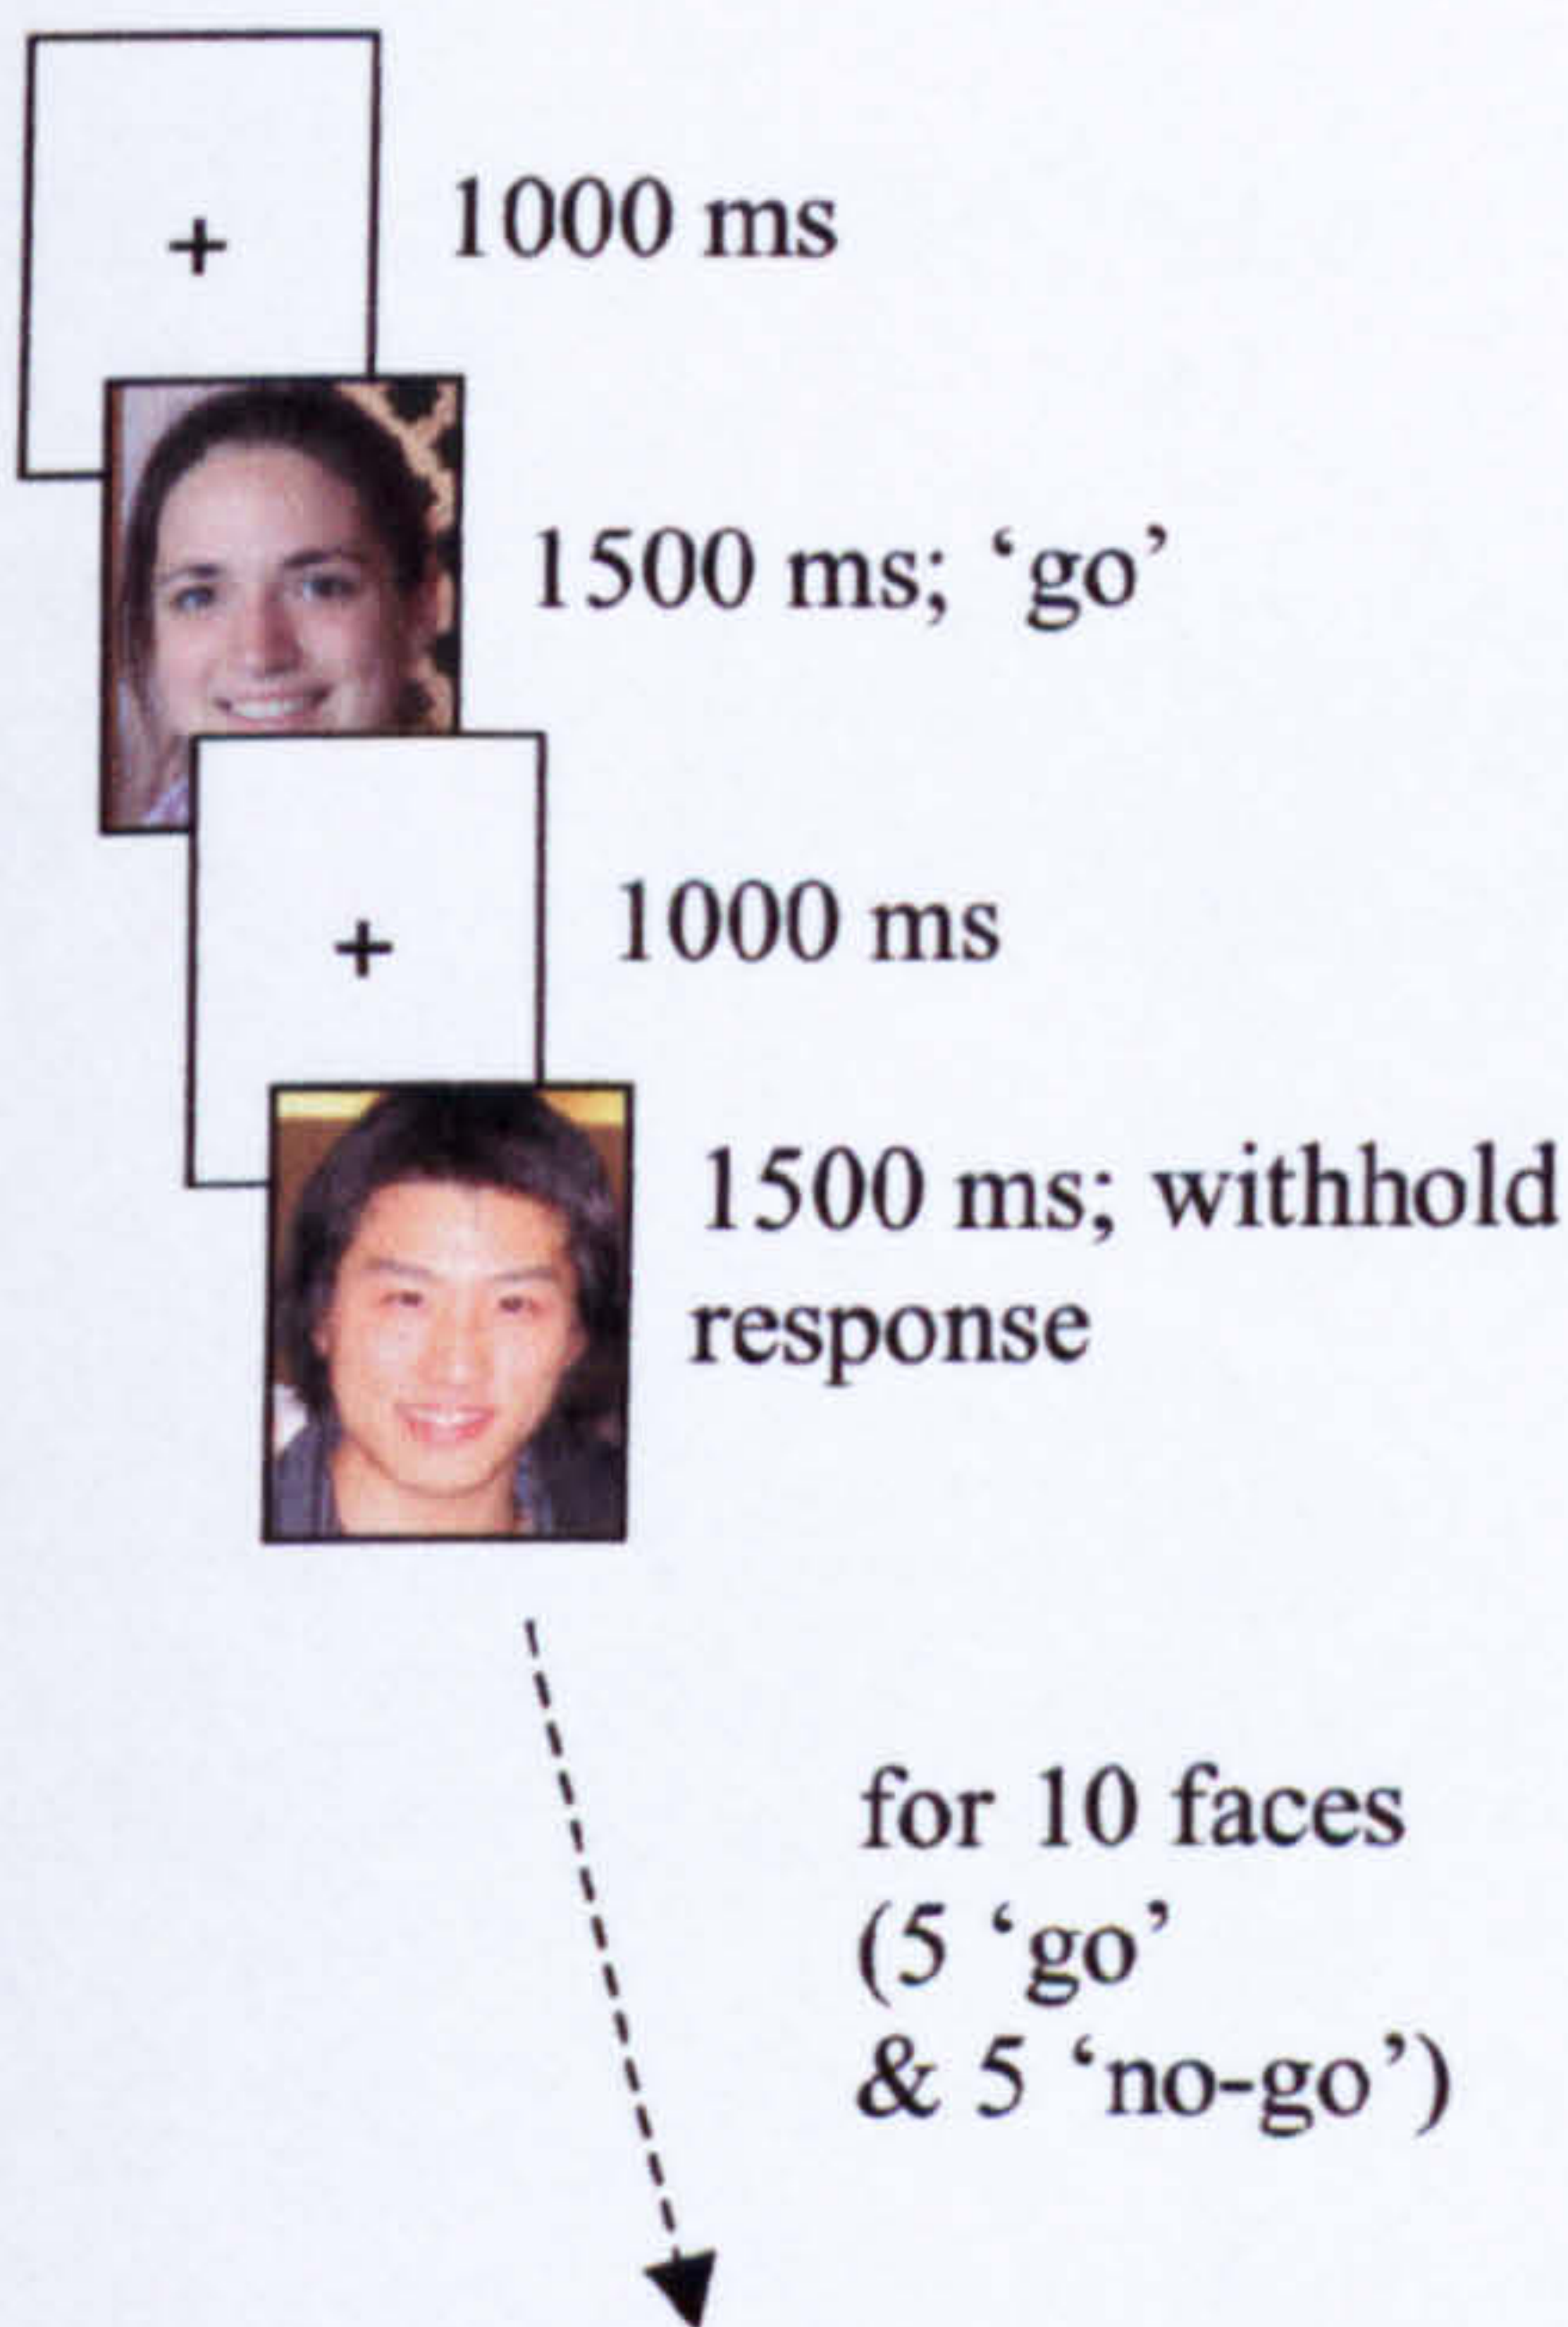
in each, per block. Faces previously seen in phase 1 were presented for evaluation in the order in which they appeared in phase 1, with novel faces randomly interspersed.

Table 18: Phase 2 stimuli

	Seen in phase 1		Novel	
	Go	No-Go	Go	No-Go
Match race & gender	5	5	5	5
Match race only	-	-	5	5

The sequence of events in a block is described in Figure 50.

Phase 1: Go/No-Go Task
"Go = Caucasian Females"



Phase 2: Evaluation Task
All faces displayed until an evaluation response is made.

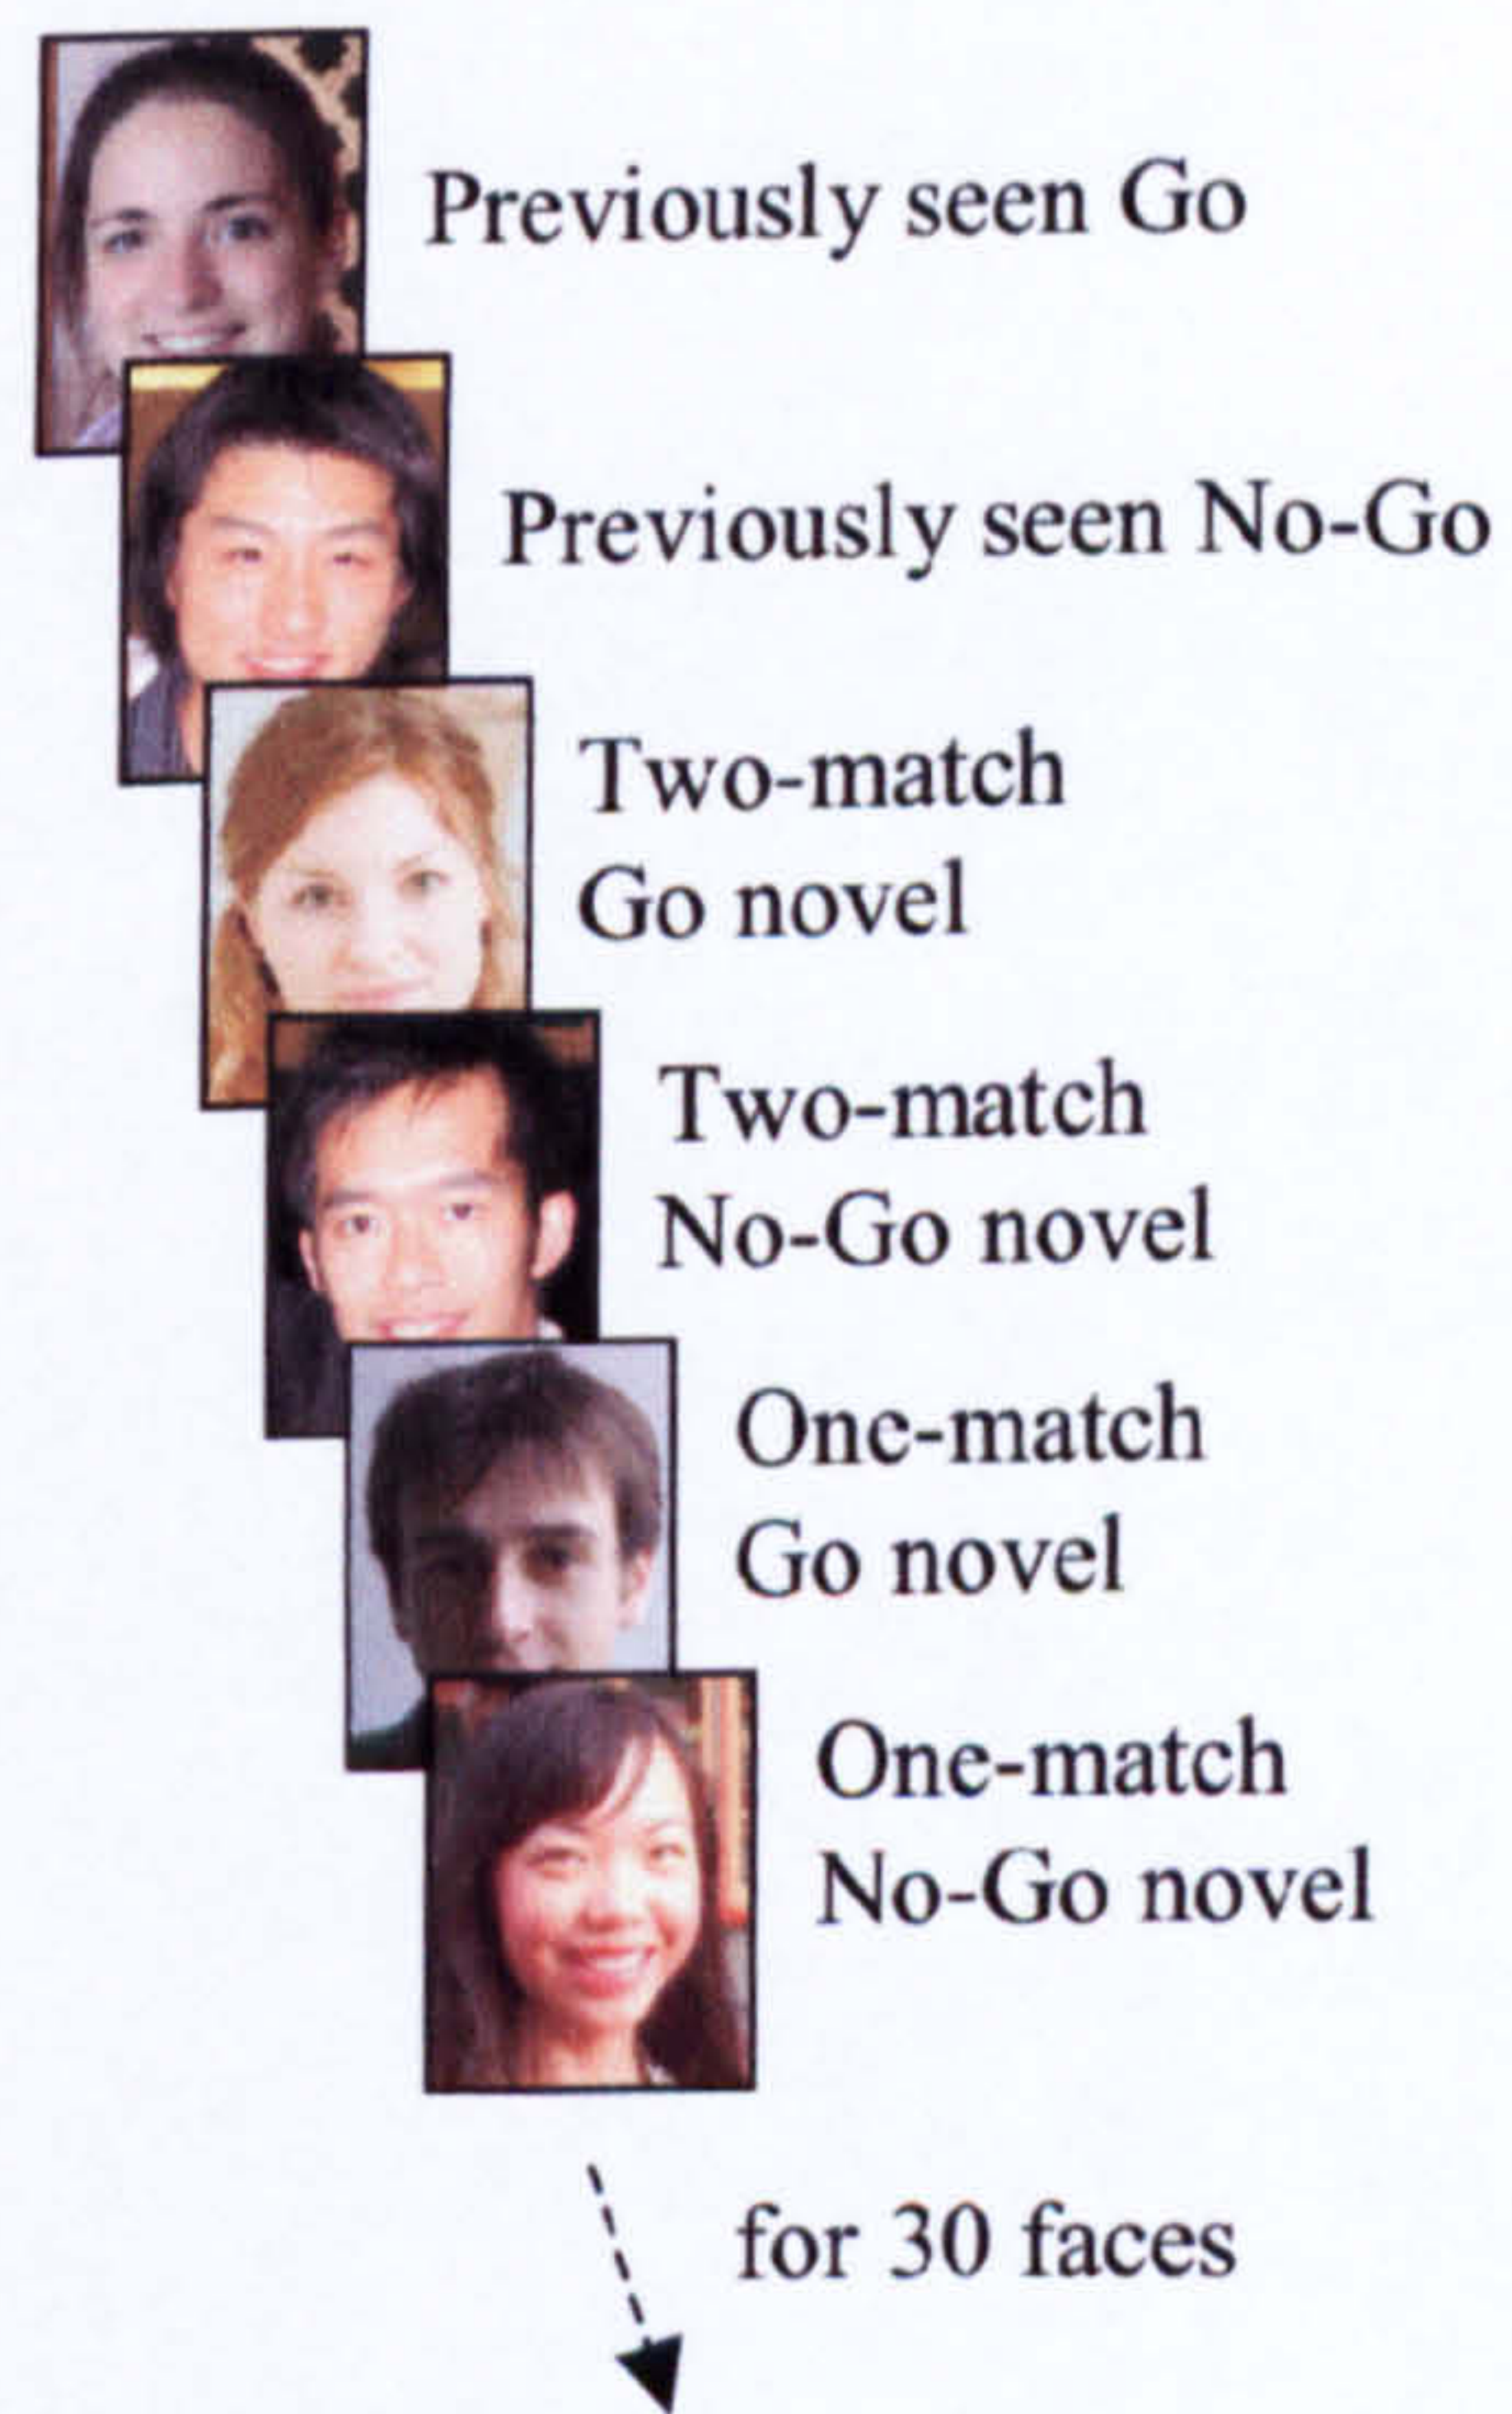


Figure 50: Sequence of events for experimental trials in one block, where 'go' is defined as Caucasian female (and by default 'no-go' is defined as Asian male). Phase 2 depicts each of the 6 evaluation conditions.

Data Analysis

Phase 2 ratings from trials in which a phase 1 error was made were eliminated (2% of data). Rates to faces were compared like-for-like. For example, CF faces when they were seen as 'go' stimuli in one block were compared to CF faces seen as 'no-go' faces in another block. There were no significant main effects or interactions with sex of the face ($p > .1$) and so results are collapsed for male and female faces. Rates to faces were analysed by own- versus other-race. Own-race faces are defined as Asian faces for Asian participants and as Caucasian faces for Caucasian participants.

Results

Own- v Other-Race Effects

A repeated measures ANOVA with race (own-, other-race), exposure type (seen, novel two-match, novel one-match) and attention (go, no-go) as within group factors was conducted. There were no significant main effects of race or attention, nor any two-way interactions (all $p > .1$). There was a main effect of exposure, $F(2, 31) = 4.75, p < .05$: faces exposed in the attention phase were rated as more trustworthy than faces that were two-match novel faces, $t(1, 31) = 2.73, p = .01$, or one-match novel faces, $t(1, 31) = 2.99, p < .01$. The two novel types were not rated differently ($p > .1$). Also, there was a significant race by exposure by attention interaction, $F(1, 31) = 3.67, p < .05$, and so ratings of own- and other-race faces were analysed further separately.

T-tests confirmed that, for own-race faces, previously seen no-go faces are devalued relative to their go counterparts, $t(1, 31) = 2.08, p < .05$. In addition, and

unexpectedly, own-race faces that match the race and gender of the previously seen no-go face are also devalued relative to their go counterparts, $t(1, 31) = 2.09, p < .05$ (see Figure 51). Other-Race faces showed no evidence of a devaluation effect for either previously seen, or novel faces (all $p > .1$).




Figure 51: Mean trustworthiness ratings of own race faces (on a scale of 1 to 7 where 1 = very untrustworthy and 7 = very trustworthy). Plotted for 'go' and 'no go' conditions by exposure type. Error bars represent ± 1 SE.




Figure 52: Mean trustworthiness ratings of other race faces (on a scale of 1 to 7 where 1 = very untrustworthy and 7 = very trustworthy). Plotted for 'go' and 'no go' conditions by exposure type. Error bars represent ± 1 SE.

Between Group Differences

Although accuracies averaged across go and no-go trials did not differ between Asian and Caucasian participants ($p > .1$), when analysing go trials alone Caucasian participants were significantly more accurate ($M = 99.2$, $SD = 1.5$) than Asian participants ($M = 96.6$, $SD = 2.9$), $t(1, 31) = 3.30$, $p < .005$. This would appear to be a speed/accuracy trade-off as Caucasian participants were also slower to respond to a go face ($M = 637$ ms, $SD = 83$ ms) than their Asian counterparts ($M = 616$ ms, $SD = 75$ ms), although not significantly so ($p > .1$).

There was a main effect of participant race in the ratings, $F(1, 31) = 4.29$, $p < .05$, and an interaction of participant race with stimulus race, $F(1, 31) = 40.85$, $p < .001$. This was apparently a mere reflection of: a) the more trusting nature of the

Caucasian participants (*mean rating* = 4.42, *SD* = 1.50) compared to the Asian participants (*mean rating* = 3.99, *SD* = 1.34), and b) the greater ‘trustability’ of the Caucasian stimuli, which both Asian and Caucasian participants rated as more trustworthy than the Asian stimuli, $t(1, 31) = 6.48, p < .001$. As the latter result is a manifest of a biased stimuli set (Caucasian faces chosen from the Internet were simply ‘nicer’ than Asian faces chosen), and I have no desire to guess why our Caucasian participants were in a more trusting mood than our Asian participants and, critically, as neither of these effects interacted with the variables of interest (exposure, attention, own- versus other-race), participant race was collapsed in favour of the gross own- versus other-race variable.

Discussion

Prior attention directed towards an other-race face during exposure had no observable effect on subsequent ratings of trustworthiness for that face. This reinforces the idea that in order for inhibition to result in distractor devaluation, the inhibited face must be individuated and remembered. In addition, no generalised distractor devaluation effects were observed for other-race faces. Why this was the case is puzzling, and I will discuss this in the Chapter 7 Discussion section.

Unexpectedly, type effects emerged in own-race faces where I was expecting token effects. Figure 51 shows that distractor devaluation effects are of the same magnitude for old faces as for two-match novel faces. This firmly suggests that the inhibitory tag was applied to the category, and not to individual faces. This may have been a result of the categorisation task. Just like the ORE can be removed by

encouraging participants to individuate (e.g. Hugenberg, Miller & Claypool, 2007), here I appear to have induced the opposite effect by encouraging participants to categorise own-race faces, making them less individuated by participants.

Interestingly, distractor devaluation generalisation of own-race faces has been restricted on a second dimension that was not task-relevant i.e. gender. Even though the task was to inhibit faces based on race, these two variables were mutually exclusive, and the ‘devaluation-by-inhibition system’ seems to have detected this and adapted distractor devaluation generalisation accordingly. Basic-level categorisation in the race task could be defined as task-relevant categorisation, i.e., Caucasian v Asian (also, Black, Latino etc. races not presented in this experiment). However, the data indicate specificity of distractor devaluation generalisation to subordinate categories (Caucasian females, Caucasian males) that were probably directed by the design of the experiment.

Experiment 7b: Devaluation of Own- and Other-Race Faces – Gender

Task

Experiment 7a was replicated with the following exception: go and no-go stimuli were defined on the basis of the gender of the face stimuli, and not the race. Recall that race and gender were fully redundant within a phase 1 block. So, if an Asian Male was a go face, a Caucasian Female must have been a no-go face. In Experiment 7a, the Asian Male was a go face because he was Asian. In Experiment 7b, he is now a go face because he is Male. The instructions given to participants were the only difference between Experiments 7a and 7b.

Method

Participants

Sixteen Caucasian (mean age, 18.9 years, 14 females) Bangor University students participated in exchange for course credits or £5. The participants gave informed consent prior to participating and were naïve to the aims of the experiment.

Apparatus and Stimuli

As in Experiment 7a.

Procedure and Design

The general procedure matches that of Experiment 7a, with the following exception: for each block, the go category was defined by gender (e.g. Male). Again, the no-go category shared neither race nor gender with the go category.

Participants were told to go (respond) to Male faces in 4 blocks, and to Female faces in another 4 blocks of trials.

In each phase 2 block, 30 faces were presented: 5 faces in each of 6 conditions. Condition 1, Go faces (all the go faces seen in the preceding phase 1 block); condition 2, no-go faces (all the no-go faces seen in the preceding phase 1 block); condition 3, Two-match Go novels (novel faces from the go category in the immediately preceding phase 1 block); condition 4, Two-match No-go novels (novel faces from the no-go category in the immediately preceding phase 1 block); condition 5, One-match Go novels (novel faces that match the go category in the task-relevant

dimension only, i.e. gender); condition 6, One-match No-go novels (novel faces that match the no-go category in the task-relevant dimension only, i.e. gender). Table 19 depicts these conditions and the number of trials in each, per block. Faces previously seen in phase 1 were presented for evaluation in the order in which they appeared in phase 1, with novel faces randomly interspersed.

Table 19: Phase 2 stimuli

	Seen in phase 1		Novel	
	Go	No-Go	Go	No-Go
Match gender & race	5	5	5	5
Match gender only	-	-	5	5

Data Analysis

Phase 2 ratings from trials in which a phase 1 error was made were eliminated (2% of data). Rates to faces were compared like-for-like. For example, CF faces when they were seen as 'go' stimuli in one block were compared to CF faces seen as 'no-go' faces in another block. There were no significant main effects or interactions with sex of the face ($p > .1$) and so results are collapsed for male and female faces. Rates to faces were analysed by own- versus other-race. As only Caucasian participants were tested in Experiment 7b, own-race in this case means Caucasian faces and other-race faces are Asian faces.

Results

Own- v Other-Race Effects

A repeated measures ANOVA with race (own-, other-race), exposure type (old, one-match novel, two-match novel) and attention (go, no-go) as within group

factors was performed. The ANOVA revealed a main effect of race, $F(1, 15) = 18.38$, $p = .001$. This was a replication of the finding in Experiment 7a that the Caucasian faces were more trustworthy than the Asian faces (see Table 20).

Table 20: Mean ratings and SD for Caucasian (own-race) and Asian (other-race) faces.

	<i>M</i>	<i>SD</i>
Caucasian (Own-race)	4.57	0.55
Asian (Other-race)	4.04	0.37

The ANOVA also revealed a significant main effect of exposure type, $F(1, 15) = 5.26$, $p < .05$. The difference between the two novel exposure types (one-match and two-match) was non-significant and so these variables were collapsed (along with race and attention) to reveal that old faces ($M = 4.54$, $SD = .62$) are rated as more trustworthy than novel faces ($M = 4.35$, $SD = .39$; mere exposure), $t(1, 15) = 2.26$, $p < .05$. This is consistent with Experiment 7a. There were no other significant main effects or interactions (all $p > .1$).

Task Difficulty Effects (meta-analysis of Experiment 7a and 7b)

The accuracy and RT data from the Caucasian participants in Experiment 7a were compared to those of the Caucasian participants in Experiment 7b. Note that the participants saw exactly the same stimuli in the same experimental procedure. The only difference between the two groups was their go / no-go definitions. Participants in Experiment 7a discriminated go and no-go stimuli by the race of the face; participants in Experiment 7b discriminated them by the gender of the face.

There were no differences between the two groups in the accuracy to perform the tasks ($p > .1$). However, participants in Experiment 7b were significantly faster to perform the go / no-go task than participants in Experiment 7a, $t(1, 31) = 2.65, p < .05$. Participants discriminating the faces by gender were faster ($M = 557$ ms, $SD = 88$ ms) than participants discriminating the faces by race ($M = 637$ ms, $SD = 83$ ms).

Discussion

Here, while prior exposure to faces caused more positive subsequent evaluations of them, prior attention directed to faces during the go / no-go task had no observable effect on subsequent ratings of trustworthiness. That mere exposure effects were found suggests that exposed faces are in some way encoded into memory (as they were more fluent); an inhibitory tag does not appear to have been stored with the faces' representations. This was true of both own- and other-race faces.

To perform this task on the basis of gender (Experiment 7b) was easier than to perform the task on the basis of race (Experiment 7a), as evidenced by the RT differences. This reinforces the idea put forward in Chapter 5 that in order for prior inhibition to result in distractor devaluation, the inhibition must be effortful (Kiss et al., 2007). This also suggests that, even though it was possible for participants in Experiment 7a to perform the task on the basis of gender, they did not, and instead performed the task on the basis of race as instructed.

Chapter 7 Discussion

While I was not expecting generalisation of distractor devaluation to occur for own-race faces, the results still addressed the issue of the level to which distractor

devaluation would generalise. Interestingly, distractor devaluation generalisation was restricted to subordinate category members. Members of the same basic-level (task-relevant) category that belonged to different subordinate categories were not devalued. Even though the task was to inhibit faces based on race, race and gender were mutually exclusive. The face features were encoded to a more sophisticated level than required by the task, and generalisation of distractor devaluation appears to be sensitive to this.

Why no distractor devaluation generalisation was found for other-race faces is puzzling. Perhaps the measures of the experiment were not sensitive enough to have observed the effect. It may be that an inhibitory tag spread over a whole category is weak and needs several instances to accumulate. In Experiment 7a the to-be-ignored category changed from block to block. This perhaps diluted any effects I might have otherwise observed if the to-be-ignored category remained constant throughout the experiment. I addressed this in Experiment 8 (Chapter 8).

Chapter 8

Devaluation of Own- versus Other-Age Faces

Using the phenomenon of the ORE as a foundation, Sporer (2001) developed the In-group / out-group model of face processing (see Figure 53 for a diagrammatic representation). The model proposes that on initial encounter with a face, one of two alternative processes is activated:

1. When confronted with an in-group face (a face that is of the same social category as the observer), default automatic processing commences with configural (holistic) encoding, characteristic of expert processing of a normal upright face.
2. When confronted with an out-group face (a face that is of a different social category to the observer), perception of an out-group characterisation cue triggers categorisation before other, more typical face processing strategies commence.

So, when a face conforms to a cultural default (e.g. White British), the dimensions along which in-group and out-group members differ will not become salient and therefore will not be processed (Levin, 1996). This is the default mode. But, when an obvious out-group characterisation cue is present and is immediately detected (e.g. skin-tone, hair colour), characterisation is automatically triggered.

This general model of in-group / out-group face processing thus applies to face characterisations other than race. It has the advantage of being able to account for, not only the ORE, but also differences in recognition ability between any categories of faces (e.g. an own-age bias).

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.

An own-age bias in the ability to identify and remember faces has been demonstrated. Several investigators have examined whether participants show superior recognition of faces belonging to their own age group, relative to faces belonging to a different age group. Wright and Stroud (2002) studied this using young (18-25 year old) and middle-aged (35-55 year old) men. Participants viewed four videotapes of a 21, 23, 48 or 51 year old perpetrator stealing either a car or a television. In a second similar experiment, participants were 18-33 or 40-55 year old men. Across both experiments, the younger participants were more likely to correctly

identify the younger perpetrators. The middle-aged participants were more likely to correctly identify the older perpetrators (but not statistically significantly so). These data suggest that an own-age bias in face recognition exists, at least for younger participants.

Bäckman (1991) investigated the own-age bias in young and elderly adults and found evidence for the bias persisting into old age. Accuracy in face recognition decreases with age, but Bäckman demonstrated that the deficit is reduced when elderly adults are asked to recognise faces of their own age. Elderly adults were more accurate to identify faces of their own age relative to faces of a different age.

Taking this evidence of greater memory for own-age faces, I wondered whether the distractor devaluation effect would be larger for own-age distractor faces compared to different-age distractor faces. If inhibition is stored with an object's representation in memory, and is responsible for the distractor devaluation effect, then faces which are more likely to be remembered (faces which belong to the same age group as the participants) should be devalued more than faces which are less likely to be remembered (faces which belong to a different age group than the participants).

Experiment 8: Generalisation of Distractor Devaluation to Own- and Other-Age Group Faces

In Experiment 8 I wanted to investigate whether out-group generalisation effects could be found (I failed to find them in Experiment 7a). In order to maximise the chances of finding such an effect, the to-be-ignored category was not altered

during the course of the experiment. Pre- and post-test evaluations were collected in order to explore the gross effect the experiment had on the to-be-ignored category.

In this experiment, I also wanted to replicate the subordinate category generalisation effects found in Experiment 7a, and see whether it would extend to a more subtle category distinction than race (i.e. age).

In Experiment 8 young adults (18-25 year olds) were exposed to young faces (approximately 20-30 years old) and elderly faces (approximately 50-65 years old⁴) in the go / no-go task. For one group of participants, distractor faces were young (same-age). For a second group of participants, distractor faces were elderly (different-age). Participants were not required to judge the age of the faces, nor were they told that faces belonged to different age groups.

I hypothesised that the first group of participants would be better able to individuate and remember their distractor faces and so I predicted that the first group of participants would demonstrate a larger distractor devaluation effect than the second group of participants.

The design of Experiment 8 closely replicates that of Experiment 6b. Blue and green frames were used and only the conjunction no-go faces were biased for age (go and featural no-go face conditions consisted of half young and half elderly faces). In

⁴ While the 50-65 year old age group is not elderly, I use the term 'elderly' in this context in order to avoid confusion. An 'old' face in previous experiments indicated that the face had been exposed during the attention task (as opposed to a 'novel' face). By using 'elderly' here, I can contrast this with 'young' faces and avoid confusion.

this way, age biases were made more subtle in an attempt to see if the subordinate category generalisations of Experiment 7a were a result of the more obvious race / gender redundancies of the experimental design.

Method

Participants

Two groups of 16 participants (Group 1: 13 females, mean age = 21.3 years; Group 2: 11 females, mean age = 23.3 years) from Bangor University volunteered to take part. All had normal / corrected to normal vision and informed consent was obtained.

Apparatus

A Pentium-4 computer, running E-Prime 1.0, recorded data and presented stimuli on a 55.9 cm monitor (100 Hz, 1024 x 768 resolution). The viewing distance was 70 cm.

Stimuli

Stimuli were digital colour photographs taken from the internet and consisted of 816 faces and 680 scenes. Face images (height 5.7°) were frontal views of Caucasian adults, with neutral or smiling expressions and visible hair, neck and eyes. Half were female, half male; half were photographs of older adults (approximately 50 – 65 years; elderly faces) and half were photographs of younger adults (approximately 18 – 30 years; young faces). The face stimuli thus consisted of four pools: 204 elderly

females, 204 young females, 204 elderly males and 204 young males. Throughout the experiment all face images were selected randomly from one of these four pools.

Scene images varied in size (minimum height 5.5° , maximum height 12.9°) and were neutral images of British and American houses. Half were exterior scenes, half interior.

Like Experiment 6b, two types of cue were used in conjunction with the images. When presented with a face, the cue was a 0.3° thick frame surrounding the image. The frame could be blue or green, depending on the experimental condition. When presented with a scene, a letter cue was used. Letter cues (height 0.8°) of an 'X' or an 'O' were overlaid in the centre of the scene image, were always aqua and presented in Arial font.

Design and Procedure

The experiment lasted for approximately 30 minutes. Participants began with a pre-experiment evaluation block in which they were asked to evaluate faces for trustworthiness. 40 faces (presented in a random order) were evaluated by each participant, 10 faces from each face pool. Face images were presented in the centre of the screen, without a frame, and remained displayed until participants made a response. Following a response, the next face was presented without an inter-trial-interval.

Participants then proceeded to the main experimental portion of the experiment which consisted of 8 blocks, between which participants were allowed to


take a break. Each block had three phases of 48 trials. Like Experiment 6a, Phase 1 was the attention task, phase 2 was the evaluation task, and phase 3 was the memory task.

Phase 1 contained 24 face trials with a go / no-go task. The participants' task in this phase was to press the space bar to a go image, and withhold a response to a no-go image. Go and no-go images were defined by a conjunction of face gender and frame colour. A go trial was a male in a blue frame (16 trials), a featural no-go was a male in a green frame (4 trials), and a conjunction no-go was a female in a blue frame (4 trials).

Critical to the design of this experiment was the age bias of each of these conditions. Go trials contained 50% elderly males and 50% young males. Featural no-go trials contained 50% elderly males and 50% young males. It was the conjunction no-go trials in which the age of the stimuli were biased. Group 1 saw 100% elderly females, and group 2 saw 100% young females. Therefore the critical comparison is between the evaluations of conjunction no-go female distractors made by group 1 participants (different-age faces) and group 2 participants (same-age faces).

Scene trials served as filler trials and contained 16 go trials (interior 'O'), 4 featural no-go trials (interior 'X') and 4 conjunction no-go trials (exterior 'O').

Images were presented in the centre of the screen for 1500 ms with a fixation cross ITI of 1000 ms. Participants had to respond to a go stimulus within 2000 ms or it was counted as an error trial. If a participant made an error (slow or absent go trial response, or a failure to withhold a response to a no-go trials), an error tone sounded immediately as feedback. See Figure 54 for the phase 1 trial sequence.



... until 24 face and 24 filler trials

Figure 54: Phase 1 trial sequence. The atypical distractor female face presented (a female in a blue frame) is young: this is a trial from group 2.

Following phase 1, participants evaluated 48 face images (phase 2). 24 of these were old (the 24 faces seen in phase 1, presented in the order of their previous exposure), and the further 24 were novel (previously unseen faces randomly interspersed with the old faces). The images were presented without cues and the novel faces were selected to match the gender biases of the old faces (20 males, 4

females). Of these novel faces, 50% belonged to the elderly face group and 50% belonged to the young face group, for both male and female faces. Faces were evaluated for 'trustworthiness' and rated on a 7-point scale (1 = Very untrustworthy; 7 = Very trustworthy).

Phase 3 served as a filler task between blocks. Participants evaluated 48 scene images. Like the faces, 24 of these were seen in phase 1 (presented in the order of their previous exposure), and the further 24 (randomly interspersed) were novel scenes. Scenes were presented without cues and the novel scenes were chosen to match the scene type biases of the old scenes (20 interior, 4 exterior). Participants evaluated the scenes based on their memory for them. Participants had a choice of three responses when presented with the scene: It's a new scene; maybe I remember it; or I definitely remember it.

After 8 blocks of the main experimental portion, participants post-rated a further 40 faces (presented in a random order), 10 faces from each face pool. Face images were presented in the centre of the screen, without a frame, and remained displayed until participants made a response. Following a response, the next face was presented without an inter-trial-interval.

Results

Data Analysis

Data from face trials only was analysed; scene stimuli served purely as filler trials to aid memory consolidation for faces between trials, and as such will be disregarded from this point.

Data from one participant was removed as on over 50% of their atypical no-go trials they failed to withhold a response. One further participant was removed because they were not using the scale correctly. Thus remaining in the analysis were 15 group 1 participants and 15 group 2 participants.

The same data trimming method as in Experiment 6 was applied. As a result, 7.03% of old item evaluations and 3.68% of novel item evaluations were removed. Results are thus based on an average of 134 go items, 32 featural no-go items, 28 conjunction no-go items, 164 go- and featural no-go-like novel items and 33 conjunction no-go-like novel items.

Experimental Phase 1: Attention trials

Error rates varied between the attention conditions. Go trials had a mean error rate of 1.64% ($SD = 2.02\%$) there was no difference between the groups ($p > .1$) or the age of the stimulus ($p > .1$). Featural no-go trials had an error rate of 2.29 % ($SD = 3.84\%$), with no difference between the groups ($p > .1$) or stimulus age ($p > .1$). Conjunction no-go trials had a mean error rate of 12.81% ($SD = 10.29\%$). This is significantly different from the error rates of the featural no-go trials, $t(1, 29) = 4.98$, $p < .001$, but did not interact with group ($p > .1$).

Further, error rates for the conjunction no-go trials for group 1 were higher than for group 2 (marginal difference: $t(1, 29) = 1.91$, $p = .066$), indicating that it is harder to distinguish an old female from a male than it is to distinguish a young female from a male in a no-go task.

The mean RT to respond to a face in a go trial was 650 ms ($SD = 73$ ms).

There was no difference between groups or between stimulus age (both $p > .1$).

Experimental Phase 2: Rating Data

Group 1 and group 2's ratings of novel stimuli were subjected to a 4 x 2 repeated measures ANOVA with stimulus type (young female, young male, elderly female, elderly male) as a within group factor and group (group 1, bias elderly females; group 2, bias young females) as the between group factor. There was no significant difference between the groups in the rating of novel stimuli ($p > .1$). *T*-tests revealed that this was true for each stimulus type (all $p < .1$; see Table 21). Ratings of novel stimuli were thus used as a baseline to describe the ratings of old stimuli exposed in the attention phase.

Table 21: Mean ratings (and SD) for novel faces by groups 1 and 2. There were no significant differences between the ratings of the groups for any stimulus type.

Group	Stimulus Type			
	Young Female	Young Male	Elderly Female	Elderly Male
1: Bias Elderly Female	4.61 (0.21)	3.98 (0.19)	4.94 (0.17)	4.42 (0.17)
2: Bias Young Female	4.65 (0.33)	3.97 (0.20)	4.82 (0.24)	4.38 (0.19)

All evaluation results are presented in terms of the effect of exposure. That is, results are expressed as the difference between the ratings of old faces (exposed in phase 1) per condition, minus the ratings of their novel age-matched counterparts. Note that for comparison with featural no-go trials, novel female ratings are only included if they belong to the same age category as that of the featural no-go. So, only elderly female novels are analysed for group 1, and only young female novels are

analysed for group 2. Therefore, a positive rating difference indicates that the old pre-exposed stimuli are rated more positively than new, previously unexposed stimuli (a mere exposure effect). Conversely, any negative rating difference indicates that old pre-exposed stimuli are rated more negatively than unexposed stimuli (devaluation).

Attention Effects on Emotional Evaluation

The rating differences of the go and featural no-go conditions were first analysed (all male stimuli, no age biases). A 2 x 2 repeated measures ANOVA with attention (go, featural no-go) as the within-group factor and group (group 1, group 2) as the between-group factor revealed no statistically significant rating differences between attention conditions or groups (both $p > .1$), nor was there a significant interaction between the two ($p > .1$; see Figure 55).




Figure 55: Mean attention effect per go / no-go condition. Scores reflect ratings of stimuli seen in the go / no-go task, minus the average rating of novel counterparts. A score above zero indicates that stimuli were evaluated as more positive following exposure in the go / no-go task; a score below zero indicates that stimuli were rated as more negative following exposure in the go / no-go task. Error bars represent ± 1 SE of the mean.

Ratings of male faces by group 1 were affected by the attention condition at pre-exposure. Faces seen as a featural no-go or as a go were trusted more than their novel counterparts (featural no-go, $t(1, 14) = 2.61, p < .05$; go, $t(1, 14) = 2.54, p < .05$). Group 2's ratings of featural no-go faces and go faces were not significantly different from baseline (both $p > .1$).

T-tests comparing ratings of conjunction no-go's to the baseline (novel) revealed that ratings of conjunction no-go's by group 2 (young females) replicated the devaluation effect. Pre-exposed faces in this condition were liked less than their novel counterparts, $t(1, 14) = -2.32, p < .05$. Conversely, the ratings of conjunction no-go's by group 1 (elderly females) were not significantly different to their novel

counterparts ($p > .1$). Own-age conjunction no-go's were devalued, but other-age conjunction no-go's were not (see Figure 56).




Figure 56: Biased conjunction no-go faces. Error bars represent ± 1 SE of the mean.

Pre-rates and Post-rates

The mean pre-experiment trustworthiness ratings for each set of face stimuli are shown in Table 22. There was no difference between the groups to rate any of these face categories (all $p > .1$).

Table 22: Mean and standard deviation pre-rating of each face category.

	Elderly Female	Young Female	Elderly Male	Young Male
<i>M</i>	5.02	4.54	4.51	4.07
<i>SD</i>	.17	.14	.16	.14

The mean pre-rate was subtracted from the mean post-rate to give a difference score, where a positive rating difference indicates the face category is liked more after the experiment, and a negative rating difference indicates the face category is liked less after the experiment.

There was no difference in the pre- and post-rates of male faces (young or elderly) for either group ($p > .1$). Female faces however, were affected by exposure during the experiment (see Figure 57). T-tests revealed that the two unbiased conditions were not different from each other (group 1 young females versus group 2 elderly females: $p > .1$). Biasing elderly versus young faces in the experiment did have an effect (group 1 old females versus group 2 young females, $t(1, 29) = 3.30, p < .01$). The devaluation of the elderly female faces by participants who saw that category as conjunction no-go's was significant (post-rate compared to pre-rate, $t(1, 14) = -2.40, p < .05$). This is key evidence for devaluation generalising to a whole out-group subordinate category. The other conditions do not significantly differ in their pre- to post-ratings (all $p > .1$).




Figure 57: Pre- and post-experiment rating differences of elderly and young females by each group. Error bars represent ± 1 SE of the mean.

Discussion

In accordance with Experiment 6b, only no-go's which were rejected on the basis of their gender (and not the frame colour) were subsequently devalued. This reinforces the idea that subsequent devaluation depends on the inhibited feature being present at evaluation (Goolsby, et al., in press).

The important finding of the current experiment though is that inhibited faces which were individuated by participants (own-age faces) were subsequently devalued on a one-to-one basis (individuated distractor devaluation), but inhibited faces which were merely categorised by participants (other-age faces) were devalued on a whole category basis (distractor devaluation generalisation).

This generalisation of distractor devaluation has, like in Experiment 7a, remained with the subordinate level category (gender *and* age) and has not been applied to the basic-level (task-relevant) category (gender). This is evident from the ratings of group 1 who devalue elderly females as a result of the whole experiment, but not young females. Once again, the 'devaluation-by-inhibition system' appears to be adaptive, recognising that only elderly females were being inhibited during the experiment, and thus not applying the devaluation rule to young females. What is particularly incredible about this is that the age biases within the experiment were very subtle. Conjunction no-go faces made up only 16.67% of trials in the experiment. In a block of 24 trials, only four were conjunction no-go trials. If the experiment had not been age biased, two of these would have been young faces and two would have been elderly faces. Therefore, in this biased experiment, only two faces in a whole block of 24 trials have been drawn from a different age group (8.33% of trials). This was a very subtle biasing procedure, but was enough to induce generalisation to the subordinate category only.

Note that in the current experiment I did not run an equivalent group of older participants (from whom elderly faces would be own-age and young faces would be other-age). This was for two reasons: first, recruiting older participants would have been impractical and costly; but more importantly, I did not expect older participants to demonstrate an exactly inverted pattern of results compared to young participants.

Findings of an own-age bias in older adults have been inconsistent. While it is commonly observed that younger adults recognise younger faces better than older faces, older participants typically fail to show an effect of age of face (Bartlett &

Leslie, 1986; Rodin, 1987; Fulton & Bartlett, 1991). Several explanations have been put forward to account for this asymmetry. Based on the perceptual-expertise model of own-group face recognition, Bartlett and Leslie (1986) proposed that there is a development of expertise for faces of various ages throughout the lifespan. Therefore, young adults have greater knowledge of young compared to older faces. However, older adults, who have themselves been young, are equally knowledgeable of, and are equally good at remembering, young compared to older faces.

An alternative hypothesis proposes that motivational factors account for this difference. Rodin (1987) proposed that certain attributes of strangers may serve as 'disregard' cues, i.e. signals that the stranger is unsuitable to the observer's social purposes. Thus, particularly in Western society, advancing age is often associated with a decline in social status, and therefore, for people who are not themselves old, old age serves as a common disregard cue. So, younger people's difficulties in remembering older faces are due to their lack of interest in older people. However, the equivalent cue is not applied by older persons to younger faces, and thus there is no difference in older people's recollection of older versus younger faces.

Even if I had assumed that no such own-age asymmetry exists, another reason why I did not expect older participants to display equivalent inverted results compared to younger participants was in the potential differential application of the 'inhibitory tag' to evaluation by older participants. Age-associated changes in emotion processing have been demonstrated to occur. For example, cross-sectional and longitudinal studies have shown that older adults experience relatively less negative affect than younger adults (Gross et al., 1997; Carstensen, Pasupathi, Mayr

& Nesselroade, 2000; Charles, Reynolds & Gatz, 2001; Mroczek, 2001). Older adults are less likely to attend to negative than to neutral or positive pictures (Charles, Mather & Cartensen, 2003), and older adults remember less emotionally negative information and more emotionally positive information compared with younger adults (Levine & Bluck, 1997; Mather & Johnson, 2000; Charles, Mather & Carstensen, 2003; Fung & Carstensen, 2003; Mather & Carstensen, 2003). Based on this evidence, it is therefore likely that the degrees to which an inhibitory tag would be interpreted as negative by older and younger adults, would be different.

Chapter 9

Generalisation of Distractor Devaluation in Branded Products

A branded product can be thought of as an individual, much like a person. This idea was first introduced about 50 years ago and stemmed from research of how consumers perceive products and brands. Levy (1959) conceptualised a brand as a complex symbol that incorporates consumers' motives, feelings, logic and attitude. As a consequence, consumers think of brands not just as a collection of features and benefits, but instead as richly complex entities with different personalities, public personas and other symbolic qualities and implications.

When confronted with a branded product, categorisation initially occurs at the basic level, much like other objects. In the case of products, the basic level has been identified as categories such as tea or sports cars, which are largely alike in terms of shape and function. Subordinate levels might divide tea into herbal or caffeinated, and sports cars into racing or commercial. These categories share even more features with each other, without yet having brand labels. Individual exemplars of teas and sports cars might include Twinings Infusions (herbal tea), PG Tips (caffeinated tea), McLaren F1 (racing sports car) or Lamborghini (commercial sports car). Superordinate category identification is believed to be largely inferred from schemas following basic-level categorisation. For example, product X is a beverage (superordinate) because product X is a juice (basic; Rosch, Mervis, Gray, Johnson & Boyes-Braem, 1976; Lingle, Altom & Medin, 1984). Figure 58 demonstrates the typical category hierarchy structure using the example of beverages.

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.

Experiment 9: Generalisation of Distractor Devaluation in Branded Products

In Experiments 7a and 8 I demonstrated that distractor devaluation could generalise to categories. Importantly, I demonstrated that this generalisation could transfer down the categorical hierarchy from basic-level to subordinate categories. In Experiment 9 I used branded products and asked the following important question: can distractor devaluation generalisation transfer *up* the categorical hierarchy from basic-level to superordinate level categories?

In Experiment 9, participants performed a simple go / no-go task. Go and no-go products were defined as either edible or inedible products (superordinate category). I hypothesised that in order to perform this task, products must be first

identified at the basic-level (e.g. spaghetti). Subsequently, participants must use pre-existing knowledge about this product to decide whether it belongs to either the edible or the inedible superordinate category. Therefore, if no-go stimuli are defined as 'edible', then a packet of spaghetti must be inhibited. I hypothesised that if the inhibitory tag is transferred from the basic-level category of spaghetti, to the superordinate level of 'edible', then any subsequently presented edible product will be devalued, i.e., can inhibiting spaghetti devalue baked beans?

Method

Participants

Thirty-six participants (27 females, mean age = 20.9 years) from Bangor University volunteered to take part. All had normal / corrected to normal vision and informed consent was obtained.

Apparatus

A Pentium-4 computer, running E-Prime 1.0, recorded data and presented stimuli on a 55.9 cm monitor (100 Hz, 1024 x 768 resolution). The viewing distance was 70 cm.

Stimuli

Stimuli were digital colour photographs taken from the internet and consisted of 60 inedible branded products and 60 edible branded products. All products were unfamiliar to British participants. Of the inedible products, there were 12 in each of five categories: washing up liquids, soap, laundry powder, toilet paper, and shampoo.

Of the edible products, there were also 12 in each of five categories: baked beans, spaghetti, orange juice, tuna, and butter. Image sizes varied between 5.5° and 7.5° in height.

Images were degraded (slightly blurred with poor resolution). This was done in order to limit individuation of products and encourage categorisation instead. Foreign (to British participants) products were also used for this reason: products were not easily identifiable by participants and therefore a brand name could not be easily generated.

Prior to the experiment, an example of each basic-level category (not used in the experiment) was presented to participants and labelled with their basic-level category name. This was done in order to encourage efficient categorisation at the basic-level. Figure 59 depicts an idea of what the participants were presented with and an example stimulus from each category.



Figure 59: Examples of each category.

Design and Procedure

The experiment lasted for approximately 20 minutes and consisted of 6 blocks, between which participants were allowed to take a break. Each block had two phases: phase 1 was the attention task, and contained 8 trials; phase 2 was the evaluation task and contained 16 trials.

In phase 1, the participants' task was to press the space bar to a go image, and withhold a response to a no-go image. For 3 blocks, go images were edible products, and no-go's were inedible, for the other 3 blocks go images were inedible products, and no-go's were edible products. The go category alternated between blocks and participants were given instructions before each block informing them of which

category was to be the go category. Half the participants started with edible products as go, and half started with inedible products as go.

Within a block, 5 trials were go trials, and 3 trials were no-go trials (see Figure 60). Trial conditions were randomly presented within the block and images were randomly selected from the image pools.

Images were presented in the centre of the screen for 1500 ms with a fixation cross ITI of 1000 ms. Participants had to respond to a go stimulus within 2000 ms or it was counted as an error trial. If a participant made an error (slow or absent go trial response, or a failure to withhold a response to a no-go trials), an error tone sounded immediately as feedback.




Figure 60: Phase 1 trial sequence. Images not to scale.

Following phase 1, participants evaluated 16 product images (phase 2). 8 of these were the 8 products seen in phase 1 (presented in the order of their previous exposure), and the further 8 (randomly interspersed) were novel products (5 from the go category, 3 from the no-go category). Products were evaluated simply on quality. A 7-point scale was used (1 = It's a very low-quality product; 7 = It's a very high-quality product). Images were presented in the centre of the screen until response.

Results

Participant Removal

One participant was removed from the analysis because his accuracy on the go / no-go task was only 61.11%. Thus remaining in the analysis were 35 participants.

Phase 1: Go / No-Go Responses

Error trials (failure to respond to, or too slow to respond to a go trial, or failure to withhold a response to a no-go trial) were removed from the analysis. This eliminated 5.18% of data. Error rates varied between the attention conditions and between superordinate product categories. Go trials had a mean error rate of 4.29% ($SD = .51\%$). This varied significantly between superordinate product categories, $t(1, 34) = 3.60, p = .001$. For blocks in which edible products were the go category, error rates were 1.90% ($SD = .65\%$); for blocks in which inedible products were the go category, error rates were 6.67% ($SD = 1.00\%$). Error rates for no-go trials were not different between superordinate product categories ($p > .1$) and were 6.67% ($SD = .97\%$).

Similarly, the RT to respond to a stimulus in a go trial was significantly longer for inedible products ($M = 781$ ms, $SD = 16$ ms) than edible products ($M = 715$ ms, $SD = 12$ ms), $t(1, 34) = 4.51, p < .001$.

Phase 2: Rating Data

Ratings from edible and inedible trials were analysed separately. Ratings of novel products were divided into 4 conditions. Products could be go-like or no-go-like (if edible products were go, a novel edible product would be go-like). Further, products could belong to a present (in that block in phase 1) category group, or an absent category group. For example, if 'Spaghetti A' was a go product in phase 1, then 'Spaghetti B' would be a present go-like novel in phase 2. If no spaghetti had been seen in phase 1 though, Spaghetti B would be an absent go-like novel. This division led to an average of 26 old go rates, 15 old no-go rates, 19 category-present go-like novel rates, 8 category-present no-go-like novel rates, 9 category-absent go-like novel rates, and 9 category-absent no-go-like novel rates, per participant. Of these, half were rates of edible products, and half were rates of inedible products. Any condition, per participant, for which there were less than 4 valid trials was excluded from the analysis, excluding a further 1.9% of data.

Attention Effects on Emotional Evaluation

The mean ratings per condition, per participant were analysed in a repeated measures ANOVA with product type (edible, inedible), exposure (old, present-category novels, absent-category novels), and attention (go, no-go) as within group factors. There were no interactions of any factors (all $p > .1$).

There was a main effect of product type, $F(1, 34) = 6.89, p < .05$, where edible products ($M = 3.83, SD = .10$) were judged to be of lower quality than the inedible products ($M = 3.98, SD = .08$).

The exposure condition also produced a main effect, $F(2, 34) = 4.47, p < .05$. Paired t-tests revealed that old products were judged to be of higher quality than novel (collapsed across present and absent category) products, $t(1, 34) = 2.47, p < .05$. Moreover, present-category novel products were judged as better quality than absent-category novel products, $t(1, 34) = 2.63, p < .05$. Here, I have found another example of mere exposure generalising to categories (see Figure 61).

The attention (go, no-go) directed towards the superordinate product categories during phase 1 was also significant, $F(1, 34) = 5.58, p < .05$: products that belonged to the superordinate category that was just inhibited were thought to be of lower quality than products that belonged to the superordinate category that was just responded to. This effect did not interact with exposure condition (old, present category novels, absent category novels), and as can be seen in Figure 61, the devaluation effect is of equal magnitude for each of the exposure conditions. Therefore it can be concluded that distractor devaluation generalised to the superordinate level categories.

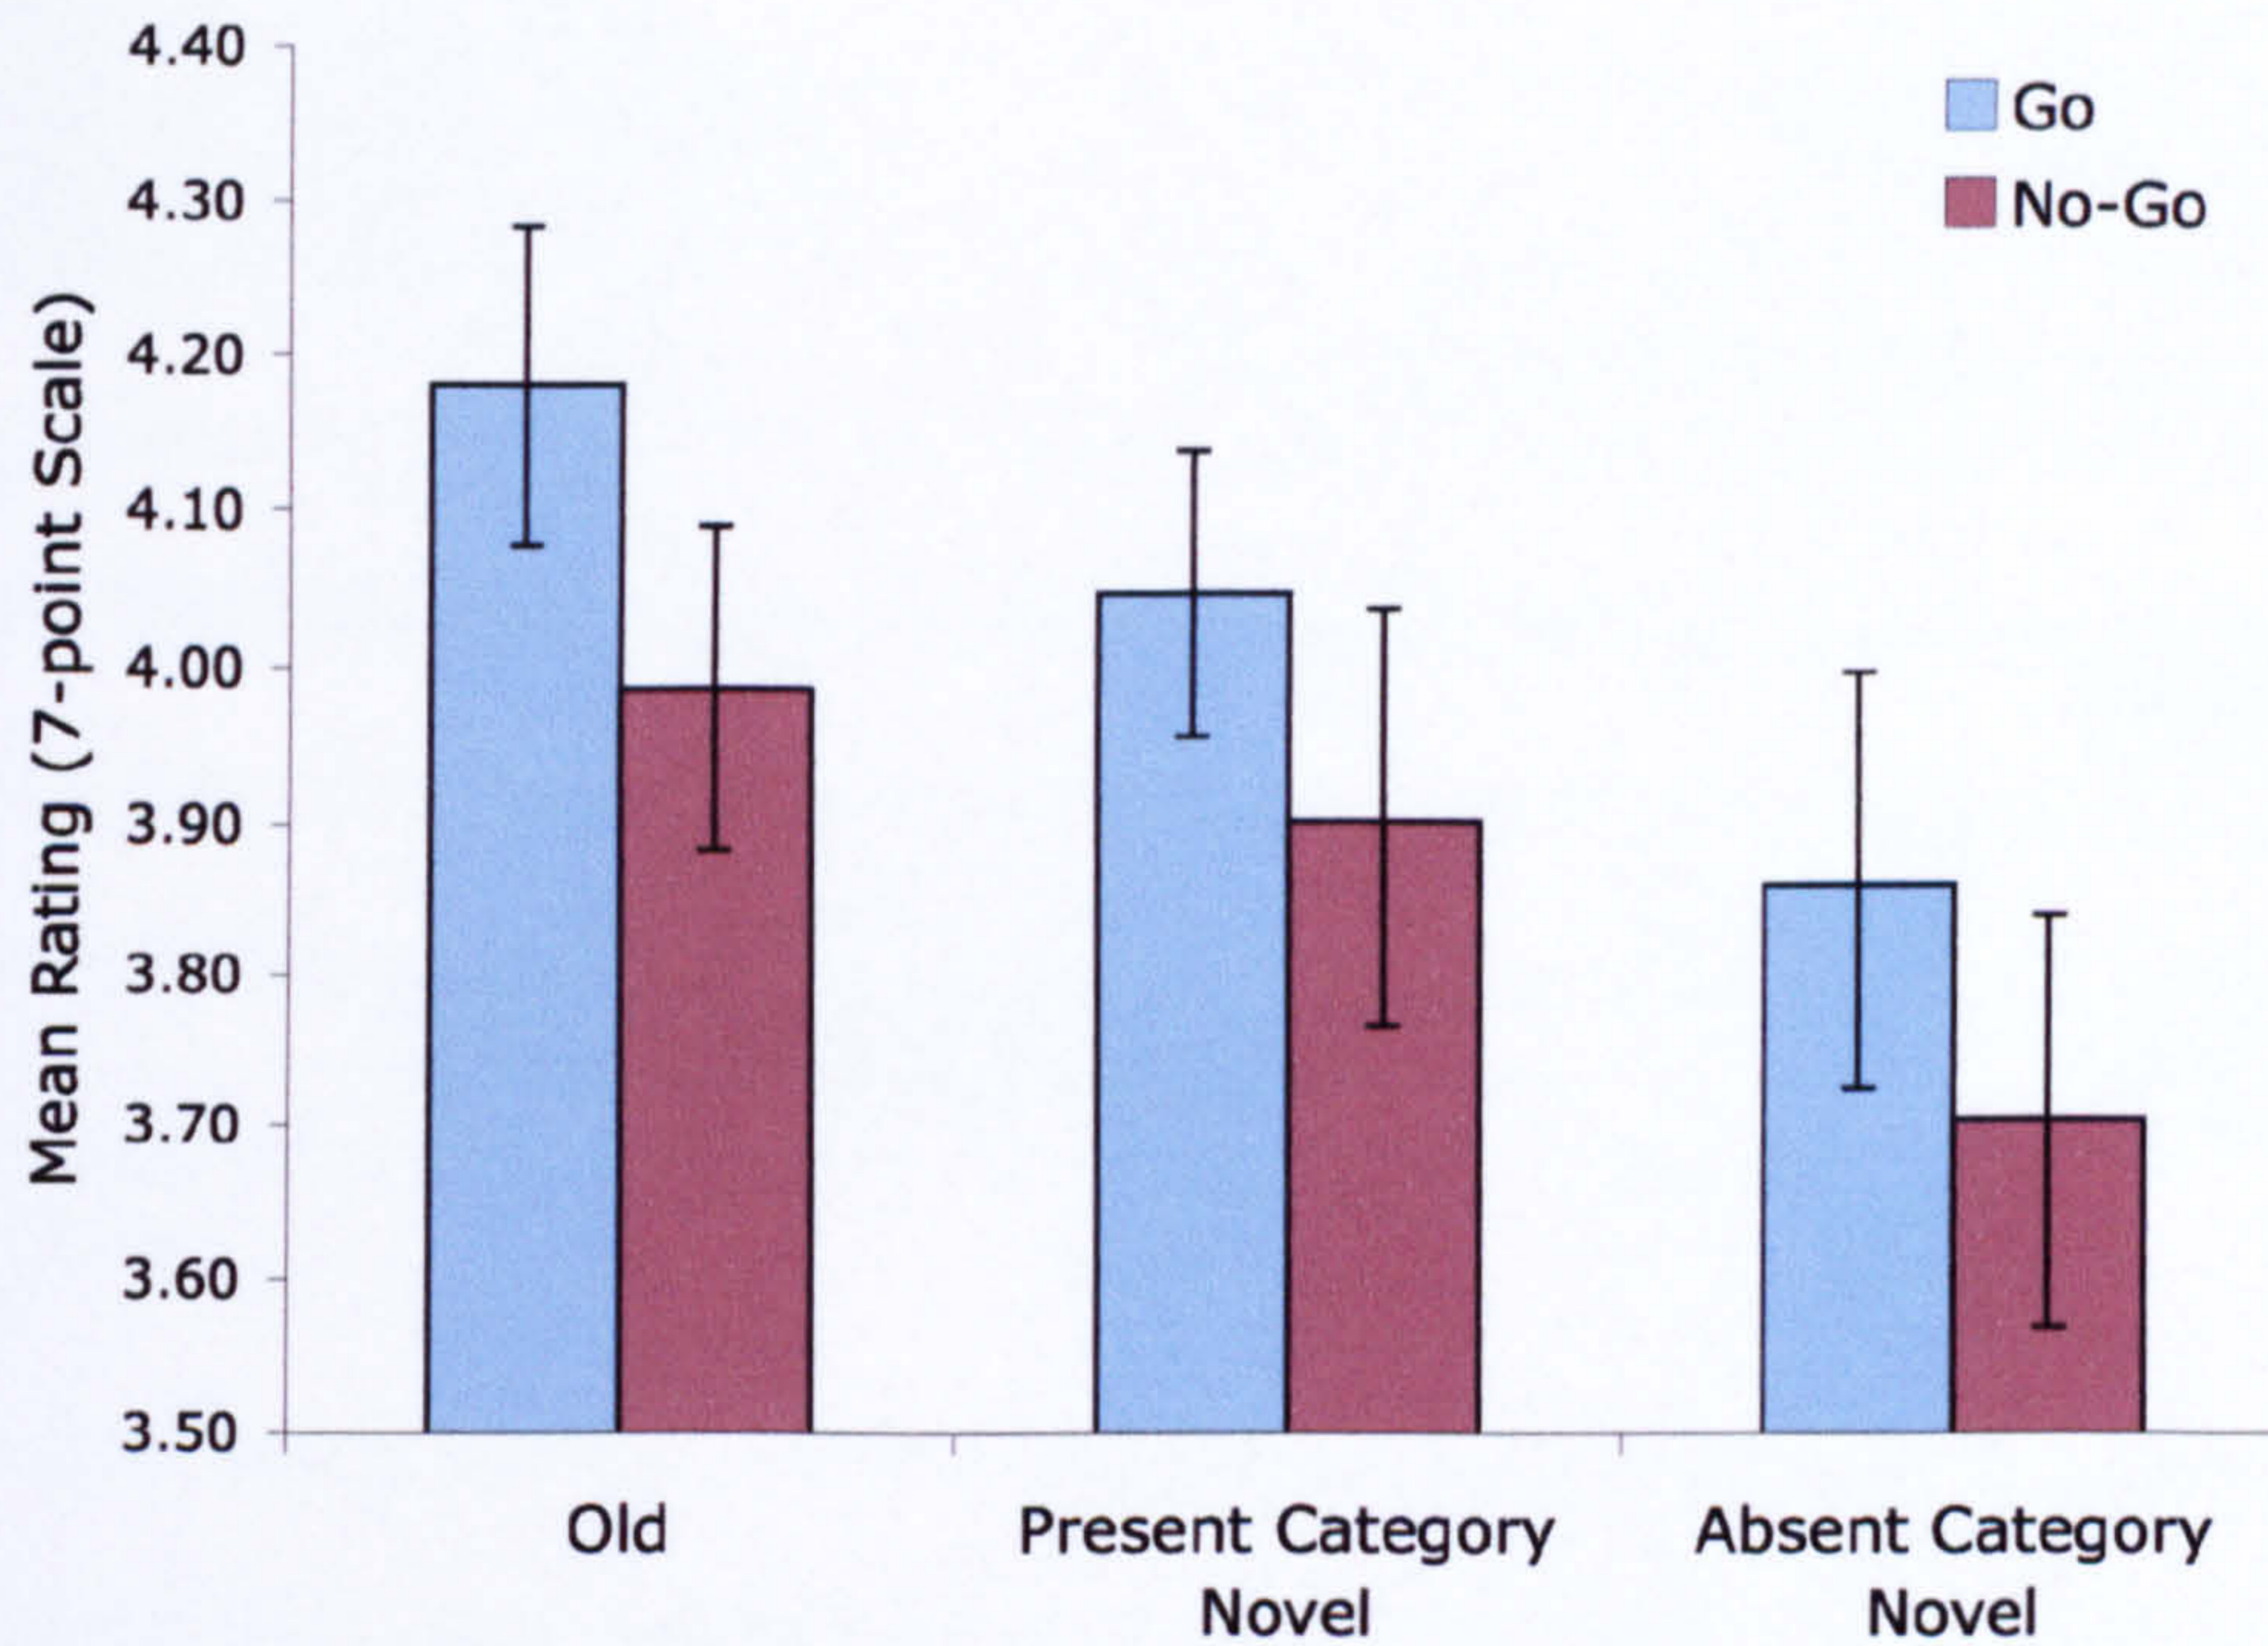


Figure 61: Mean rating for each exposure and attention condition (collapsed across edible and inedible products). Error bars represent ± 1 SE of the mean.

Discussion

In Experiment 9, I have once again demonstrated that the ‘devaluation-by-inhibition system’ can be adaptive to situational needs. In this case, it appears that the inhibitory tag has moved *up* the categorical hierarchy from the basic level to be placed with the superordinate level category. Therefore, inhibiting a packet of spaghetti could result in the devaluation of a tin of baked beans.

However categorisation, while not necessary to successfully perform the task, has occurred at both basic and individual levels. This is evident from the mere exposure effects. Indeed, the other remarkable finding of Experiment 9 is the finding once again of a category-based mere exposure effect.

In this experiment, old products were subsequently judged to be of higher quality than novel products. This is a basic mere exposure finding. However, a comparison of the two types of novel products is where the categorical mere exposure is found. If, for example, Spaghetti A was presented during the go / no-go task, a novel spaghetti (e.g. Spaghetti B) will be judged to be of higher quality than a can of baked beans (if no can of baked beans was present in the go / no-go task). Mere exposure appears to have generalised to the basic-level categories.

The finding that old products are rated more positively than present-category novel products suggests that the effect is not due to participants mistaking one product for another. Products have been individuated, and old token stimuli benefit from this additional fluency. I can therefore think of two possible mechanisms by which this categorical mere exposure might be taking place.

First, fluency for the whole category has developed. Processing one member of a category causes all other members of that category to be more easily processed in the future, perhaps by activating a category template, resulting in mere exposure.

Second, an auditory mere exposure effect has occurred. Prior to categorisation of a product as either edible or inedible, the participants may have named the products at their basic-level category ("it's a soap, that's inedible"). The word 'soap' may have then become fluent, and a second exemplar of a 'soap' may have also activated the name 'soap' (a phonologically identical repetition) and benefited from this naming fluency. Such an effect has not been widely studied. Mere exposure effects for musical sequences have been found (Bradley, 1971; Halpern & O'Connor, 2000;

Szpunar, Schellenberg & Pliner, 2004). Mere exposure effects for heard words has also been found, but not if the words were of low arousal and positive valence, which one might argue the categories in the current experiment are (Bruce, Harman & Turner, 2007⁵). As such, the mere exposure effects found in the current experiment are interesting, but theories as to the underlying mechanisms causing them are at this point speculative.

⁵ The Abstract only of this paper was available.

Chapter 10

General Discussion

But then, the people became terribly afraid and anxious. For lo! the Cognitive Miser had become transformed, by the magic of further research, into the Cognitive Monster. No longer did the creature use simplifying categories and stereotypes by choice or strategy, their use had become an addition—uncontrollable, not a matter of choice at all—and the creature's Will was powerless to do anything else. 'We must do something!' cried the people of Social Psychology. 'We must slay the monster!' And so their heroes came forth.

Bargh (1999, p. 361)

As the quotation from Bargh (1999) colourfully illustrates, humans are not always privy to the mechanisms that drive their decisions, choices, or preferences. Stereotyping theory posits that under conditions of low processing ability or motivation, schemas, or heuristics can be applied to the situation allowing past experience to influence the evaluation of previously unencountered objects.

The objective of this thesis was to explore the extent to which the distractor devaluation effect (Raymond et al., 2003) might be used as a heuristic in such situations. In a series of nine experiments, I demonstrated that inhibiting an object could have detrimental effects on the emotional evaluation of another, similar, object. I propose that in cases where the initial to-be-inhibited object is not successfully individuated, an inhibitory tag can be associated instead with a whole category representation. This inhibitory tag can be subsequently used to guide the evaluation of a previously unseen member of that category (the distractor devaluation generalisation effect). I also propose that the extent to which this generalisation affects similar objects depends on the situational demands at the time of inhibition.

Attention and Emotion

In Chapter 1, I described how a selective attention system facilitates the processing of task-relevant objects and suppresses processing of irrelevant distracting objects (Kastner & Ungerleider, 2000). I then proceeded to describe an emotion system that evaluates object representations in terms of current and future goals (Ortony, Clore & Collins, 1988). I demonstrated that these systems can work together to minimise conflict and promote rapid responses; for example, emotional stimuli produce specific effects on selective attention tasks by attracting and holding attention

relative to non-emotional stimuli (Eastwood, Smilek & Merikle, 2001; Fox, Russo, Bowles & Dutton, 2001). I then presented recent work that describes the reciprocal effect: the effect of selective attention on emotional evaluation.

Raymond and colleagues (Raymond, Fenske & Tavassoli, 2003) found that if an object were inhibited, it would be emotionally devalued on subsequent evaluation. They proposed the 'devaluation-by-inhibition' theory to account for their findings. They posited that an inhibitory tag would be stored with the distractor's representation in memory. On subsequent devaluation, this inhibitory tag is interpreted as an emotional devaluation.

The following questions were of interest to me: Is it only the inhibited object that gets devalued, or are similar objects also devalued? When are similar objects subjected to this devaluation process, and when are they not? I proposed that what occurs during encoding, specifically during categorisation of the inhibited object, will affect whether that object, or one similar, will be devalued.

Summary of Findings: The Individuation Hypothesis

In Experiment 3, I used the newly developed oddball search paradigm (Experiments 1 and 2) to replicate the original distractor devaluation findings. That is, when an object, in this case a bottle containing drinking or cleaning fluid, was a distractor in a search array, that object was devalued on subsequent evaluation, compared to an object that had been a target. Novel objects that matched the category of the distractor were not devalued. I argued that due to the nature of the stimuli, which were ambiguous, each object had to be individuated before it could be

responded to as a target or rejected as a distractor. This result was particularly striking because a novel stimulus evaluation intervened the presentation of the old stimulus at search and at evaluation. The inhibitory tag did not affect this intervening evaluation, and instead was 'carried over' from the search task to the evaluation task to affect only the specific bottle which had been inhibited.

Likewise in Experiment 6, using a go / no-go paradigm, I found that only objects presented as distractors were devalued: novel objects that matched the distractor category were not. This experiment used positive and negative faces. The largest devaluation effects were found for the negative faces. I argued that a face that is average would be more likely to be judged as positive. Therefore a negative face is less average, less likely to conform to a prototype representation of a face, and therefore more likely to be individuated. Recognition of the faces subsequent to the go / no-go task was tested in a second group of participants. The negative faces were remembered better than the positive faces. This lent weight to the individuation hypothesis, suggesting that negative faces were better individuated than positive faces, and thus were more likely to be encoded with the inhibitory tag, resulting in larger distractor devaluation effects for these faces.

Finally, in Experiment 8 I also found that distractor faces that were more likely to be individuated by participants were subsequently devalued more than distractor faces that were less likely to be individuated by participants. I achieved this using own- and other-age faces. Own-age faces are more likely to be remembered by participants than other-age faces. In Experiment 8, own-age faces that had been presented as no-go stimuli were devalued compared to novel no-go like faces.

Conversely, other-age faces that had been presented as no-go stimuli were not devalued compared to novel no-go like faces.

The findings of Experiments 3, 6 and 8 unite to support the ‘individuation hypothesis’ I presented in Chapter 1. I proposed that distractor devaluation will only result if an inhibited object has been successfully individuated. This is because, in line with the ‘devaluation-by-inhibition’ hypothesis, the inhibitory tag can only be stored with the object’s representation if it has been encoded as a unique object. Therefore, an inhibited object that has been successfully stored with its inhibitory tag will be devalued, but an inhibited object which has not been successfully stored with its inhibitory tag will not be devalued on subsequent evaluation.

This begs the question: What is the fate of the inhibitory tag if it cannot be stored with an individuated object’s representation? And how is this used on subsequent evaluation?

Summary of Findings: The Generalisation Hypothesis

I proposed that in cases where an inhibited object has not been individuated, it would instead have been only categorised. I theorised, then, that an inhibitory tag would instead be placed with the whole category representation, rather than with the individual’s representation. In Experiments 4, 8 and 9, I found evidence that could support this ‘generalisation hypothesis’.

In Experiment 4, I used stimuli that did not need to be individuated in order to be identified as a target or distractor, i.e., animals. The target and distractors were

either birds or fish, and the presence of a wing or a fin, for example, could determine these objects' category membership. In this case, animals that were seen as prior distractors were not devalued relative to animals that were seen as prior targets. However, critically, a novel animal that was evaluated between the search and evaluation of the old animal, *was* devalued, if it matched the distractor category. This suggested that the inhibitory tag had been encoded in such a way that it was accessible by this novel animal. I theorised that the tag had been associated with the distractor animal category, rather than with the individual animal that had been inhibited.

In Experiment 8, contrary to the findings of the own-age faces, other-age faces demonstrated generalisation. Because other-age faces were less likely to be individuated and remembered, I proposed that participants who had to inhibit these faces would place the inhibitory tag with the category representation instead of with individual representation. This proposal was supported by the data. Following the go / no-go task, other-age no-go faces were not devalued relative to no-go-like other-age novel faces. However, pre- and post-experiment evaluations revealed that repeated inhibition of the other-age category had devalued final evaluations, relative to initial evaluations. This demonstrated that the whole inhibited category had been devalued, supporting generalisation.

A further piece of evidence came from Experiment 9 in which participants performed a go / no-go task on the basis of whether branded products were edible or inedible. In this case, any subsequently evaluated product, old or novel, which matched the no-go category, was devalued. This devaluation was of the same

magnitude for old (actually inhibited) and novel products. There was no additional devaluation effect for individual items that had actually been inhibited, relative to similar items. This suggested that all devaluation effects were a result of an inhibitory tag on the inhibited category's representation, thus supporting the 'generalisation hypothesis'.

What was particularly interesting about the findings of Experiments 4, 8 and 9, was that each generalised distractor devaluation effect appeared to be acting at a different level of categorisation, depending on the task demands.

In Chapter 1 I outlined how categorisation is structured hierarchically. Categories range from superordinate, to basic-level, down to subordinate. Members belonging to categories at lower levels in the hierarchy share more features in common than those at higher levels in the hierarchy.

In Experiment 4, the generalisation effect appeared to be operating at the basic-level category. The task demanded that objects be categorised as birds or fish. Birds and fish are basic level categories. They can be further divided into robins, or sparrows (or angel fish or salmon), which would be the subordinate level. Conversely, both birds and fish belong to the superordinate category 'animals'. As a result of the oddball search task, any member of the basic-level inhibited category was devalued.

In Experiment 8, generalisation appeared to be acting at the subordinate level. The task demanded only that 'females' (in a blue frame) were inhibited. However, for a group of participants all of these females were "elderly". While the go / no-go task did not specify age as a factor, nor was age ever brought to the attention of the

participants, distractor devaluation only generalised to elderly females. Young females were not subject to the generalisation effect. For these stimuli, the basic-level category is female (or male), and the subordinate category is elderly female (or young female). Generalisation only occurred within the subordinate category.

In Experiment 9 however, the distractor devaluation generalisation effect appeared to be operating at the superordinate level. In order to perform the task, participants had to first identify the products at their basic level (e.g. spaghetti, soap), and then infer the superordinate category membership of the products (edible or inedible), and respond to or inhibit the products according to this latter categorisation. Because the task was defined by the superordinate category, it appears that this is the level of representation at which the inhibitory tag was stored. Any member of the inhibited superordinate category that was subsequently evaluated was devalued, even if it was not presented during the task, or did not even belong to a basic-level category that was inhibited during the task.


It appears, then, that the generalisation effect is adaptive to situational demands. The inhibitory tag appears to be able to be placed with whichever level of categorisation is most appropriate according to the situation at hand. Even when the subordinate category bias was very subtle, and not relevant to the task, as in Experiment 8 (age bias), the inhibitory tag appears to have been placed with this subordinate category, and other subordinate categories are not then 'tarred with the same brush'. The 'devaluation-by-inhibition' system appears, then, to use this inhibitory tag accordingly at evaluation.

Inhibition of Categories

Is inhibition at a category-level supported in the literature?

In an IOR experiment, Morgan, Paul and Tipper (2005) found no evidence of inhibition generalising to other category members. In their IOR sequences, either identical objects, objects in the same basic-level category, or different objects were presented (see Figure 62). On the first two presentations, attention was oriented to either the left or the right of these objects using an exogenous cue. On the third presentation, participants responded to a cue presented on either the left or the right of the object. Large IOR effects were found when the objects were identical, but IOR for categories of objects were of the same magnitude as that for different objects. The authors concluded that IOR cannot be associated with an object's category.

Third Party Material excluded from digitised copy.
Please refer to original text to see this material.



However, in a modified IOR paradigm using words, Fuentes et al. (Fuentes, Vivas & Humphreys, 1999), found evidence to support the inhibition of categories. In their series of experiments, a prime word (e.g. dog) was followed by an unrelated intervening stimulus (e.g. sea). Subsequently, a stimulus that was either related (e.g. cat) or unrelated (e.g. finger) to the prime word was presented as a target. Target words were interspersed with target non-words and RTs for participants to judge the target as either a word or a non-word were collected. The authors found that responses to unrelated targets were made more quickly than to related targets when the intervening word belonged to a new category (thus shifting attention away from the primed category). However, when attention was not disengaged from the primed category, the opposite result was found and RTs to related rather than unrelated words were faster. The authors argued that the results reflected a semantic equivalent of spatial IOR, and that the slowing of RTs reflected inhibition to the primed category.

Closely replicating the design of the experiments of Fuentes et al. (1999), Weger and Inhoff (2006) extended these findings and suggested that category-based IOR would only occur if the item pool were homogeneous. They found that when a heterogeneous item pool was used, category-based IOR was attenuated. This supports the individuation versus categorisation hypotheses I have presented. When an item can be individuated (belongs to a heterogeneous item pool), inhibition is likely to be stored with the individual object's representation. On the other hand, if the item is not successfully individuated, and is only categorised (belongs to a homogeneous item pool) the inhibitory tag will be stored with the category's representation.

Structural versus semantic object representations

Why did Morgan et al. (2005) find no category-based IOR? I suggest it is because, in their task (in contrast to the experiments presented in this thesis), neither the identity nor the category of the object was relevant to the task. Participants only had to respond to the exogenous cue presented over the object. As such, the inhibitory tag may have only been placed with a structural representation of the object, rather than a semantic representation of the object.

There is evidence that different forms of object representation can exist for different properties of objects. Categorisation of objects and recognition of specific objects are different processes that rely on separate brain systems. For example, amnesic patients with severely impaired recognition memory can perform as well as control participants in categorisation tasks (Knowlton & Squire, 1993; Kolodny, 1994; Squire & Knowlton, 1995; Reed, Squire, Patalano, Smith & Jonides, 1999). Brain imaging studies have revealed that the processes of object categorisation and recognition of specific objects have different temporal characteristics (Curran, Tanaka & Weiskopf, 2002). Finally, it has been shown that different brain areas are responsible for structural and semantic processes (Gerlach, Law, Gade & Paulson, 2000).

Therefore, it is possible that due to the task demands of the experiments presented in this thesis (and those of Fuentes et al., 1999, and Weger & Inhoff, 2006) the inhibitory tag was associated with semantic representations (representations of the meaning of the object), thus allowing categories to be inhibited. However in the study by Morgan et al. (2005) the inhibition may have only been associated with a structural

representation (a representation of the object's physical features), and therefore no category effects were observed.

This structural versus semantic object representation difference may explain the mere exposure versus generalised distractor devaluation effects found in the current series of experiments. For example, in Experiment 9, generalised distractor devaluation effects were found for the superordinate categories of edible or inedible. However, mere exposure effects were observed for old compared to new products (Figure 61, page 323). I have already discussed that the mere exposure findings of the present category novels may reflect a fluency for the repeated exposures of the name of the category (e.g. during the go / no-go task, the category name 'spaghetti' may have been activated, and thus the novel spaghetti may benefit from this phonological repetition). However, this cannot explain why old items receive mere exposure benefits over and above those of the present category novels as both instances would have activated the name 'spaghetti'.

Structural fluency versus semantic inhibition may explain this difference. It may be that two representations were created for the item exposed in the go / no-go task. The first may have been a structural representation, which upon repetition of the identical stimulus led to fluency, and therefore enhanced liking. This would occur for both go and no-go (inhibited) objects. However, due to the need to categorise objects in the task, a semantic representation may have also been activated. It may be that an inhibitory tag (in the case of no-go items) was placed with the semantic representation. In Experiment 9 the semantic representation may have been the superordinate category; under different task requirements, it may be any level of

categorisation, or the individual object semantic representation. Therefore, on re-presentation of the previously inhibited object, the structural representation will be fluent, and therefore increase liking, but the semantic representation will be associated with the inhibitory tag and will lead to distractor devaluation.

While explaining these results by means of two separate processes working together violates the law of parsimony, it appears that this dual-influence on evaluation is supported by both the literature, and the results presented here.

Other Generalisation Findings

Now that I have established my theories and have demonstrated (in my opinion) the validity of category-based inhibition leading to generalised distractor devaluation, my findings, and theories, must be considered in relation to others in the field. There are three key papers whose results my theories need to be able to accommodate. These are: Raymond, Fenske and Tavassoli (2003); Zhou, Wan and Fu (2007); and Goolsby, et al. (in press).

In the first paper of its kind, Raymond and colleagues (Raymond, Fenske, Tavassoli, 2003) demonstrated distractor devaluation. However, in the second experiment of this paper, they also demonstrated the occurrence of distractor devaluation generalising to same-type novel stimuli. Recall in Chapter 1 I discussed that Raymond et al. used colourful Mondrian (abstract pattern) stimuli. Stimuli were composed of squares or circles, and it was this composition that determined whether the Mondrian was a target or a distractor in a given trial. In their second experiment, Raymond et al. presented novel stimuli for evaluation that were also composed of

squares and circles. The magnitude of the distractor devaluation effect was the same for novel items that belonged to the to-be-ignored category as it was for old items. I argue that in their experiments, the Mondrian stimuli were difficult to individuate, and as such the inhibitory tag was not placed with the individual to-be-ignored item, but was placed instead with the category representation of the to-be-ignored stimulus type. The results of Raymond et al., to my mind, support the generalisation hypothesis.

Zhou and colleagues, (Zhou, Wan & Fu, 2007) also demonstrated generalisation of the distractor devaluation effect using artificial grammar strings: a unique approach that our lab has never taken. They used a similar artificial grammar procedure as that described in Chapter 5. In their experiment, participants were required first to memorise strings of letters that all followed the same rule (attended string). During this phase, a second string of letters that always followed a different rule was to be ignored (ignored string). In the second stage, participants were required to categorise novel letter strings as either following the same grammar rule as the attended string, or not. Past studies have repeatedly shown that people can simultaneously extract complex rules from both the attended and the ignored strings, even without the conscious attempt to do so (Tanaka, Kiyokawa, Yamada & Shigemasa, 2006). In a final phase, participants were presented with two strings of letters, and had to indicate which string they preferred. The string could conform to the attended rule, the ignored rule, or a new rule. The authors found that novel strings of letters were liked less if they followed the grammar rule of the previously inhibited strings of letters, compared to strings of letters that followed the new rule. The authors argued therefore, that distractor devaluation could generalise to novel stimuli.

This finding, again, supports the generalisation hypothesis. Letter strings had to be categorised according to the grammar rule, which may have encouraged categorisation and discouraged individuation. A novel letter string belonging to the inhibited grammar category was subsequently devalued.

Finally, Goolsby and colleagues (Goolsby et al., in press) recently collected data that also support the generalisation hypothesis. In Experiments 2 and 3 of their paper, participants were required to select the target face based on its gender, and report the colour of the overlying tint. Therefore, gender was the basic-level category, and the task demands required participants to attend or inhibit the faces on the basis of these categories. This, I believe, encouraged categorisation rather than individuation of the faces. Further encouraging categorisation was the requirement to report the colour tint overlaying the faces; once a face had been categorised, no further analysis of the face was required and attention had to then be directed to the colour tint. Finally, the Experiments of Goolsby et al. used the Genemation face stimuli I described in Experiment 5 (Chapter 5). These stimuli are highly homogenous (as they are generated from a limited number of 'parent' faces, configural details do not vary greatly) and the use of these may have reinforced the categorisation rather than individuation of these faces.

The authors found that previously inhibited faces were liked less than previously attended faces (classic distractor devaluation). However, novel faces that belonged to the inhibited gender category were devalued compared to attended gender category faces. This devaluation was of the *same magnitude* for old and novel faces, in both experiments. Therefore, I believe the above findings support the generalisation

hypothesis. Faces were categorised, not individuated, and this resulted in distractor devaluation generalisation to the basic-level category of the inhibited gender.

Further mirroring the findings of this thesis, the data of Goolsby et al. indicate that old faces received a boost in ratings compared to novel faces, independent of prior attention. I argue that this fits with the semantic representation devaluation / structural representation fluency hypothesis I put forward previously. Old faces may have been represented structurally, and these individual structural representations benefited from mere exposure. However, category (gender) was represented semantically, and the inhibitory tag was stored with this representation, which is why distractor devaluation generalised to the whole inhibited category.

Stereotyping

How might a mechanism of generalised distractor devaluation be at work in the real world? A theory like this is particularly important when applied to the issue of stereotyping.

Fiske and Taylor (1984) described the social perceiver as a 'cognitive miser'. They proposed that humans are rarely motivated to thoroughly evaluate each individual they encounter, and mental work is streamlined by the activation of category-based knowledge structures (stereotypes). In 1991, the authors updated the metaphor by suggesting that perceivers are 'motivated tacticians' who have multiple cognitive strategies which are employed based on goals, motives and needs (Fiske & Taylor, 1991). Therefore, a social perceiver could be thought of as an 'efficiency expert' (Macrae, Milne & Bodenhausen, 1994; Sherman, Lee, Bessenoff & Frost,

1998). Given basic information-processing limitations and a challenging social environment, perceivers streamline and simplify the demands of person perception by categorical thinking (Allport, 1954; Bruner, 1957; Brewer, 1988; Fiske & Neuberg, 1990; Fiske, 1998; Bodenhausen & Macrae, 1998; Brewer & Feinstein, 1999; Fiske, Lin & Neuberg, 1999; Macrae & Bodenhausen, 2000). By thinking of others based on the social categories that they belong to (e.g. race, gender, age), perceivers can apply stereotype-based knowledge (believed to reside in long-term memory; Tajfel, 1969; Bodenhausen, Macrae & Garst, 1998; Bodenhausen, Macrae & Sherman, 1999), to the current perception process.

This stereotype application may be automatic and uncontrollable (see the quote from Bargh, 1999, on page 326). In the past, attitudes were seen as conscious evaluations based on a considerable amount of weighing of pros and cons in any given situation. However, Fazio and colleagues (Fazio, Sanbonmatsu, Powell & Kardes, 1986) demonstrated that mere perception of an object is sufficient to automatically activate attitudes about it: attitudes that are prepared to guide future behaviour.

However, it is important to note that these automatically activated attitudes are malleable, and distractor devaluation generalisation might be one way in which stereotypical attitudes are altered.

Consider the scenario: you run to the supermarket because you are out of an essential ingredient for tonight's dinner. Time is of the essence: the kids need to be picked up from school in ten minutes. The supermarket has just rearranged its layout

and you have no idea where the key ingredient is now located. You must ask an employee for help, but whom do you ask?

The first employee you see is a lady in her sixties. She is immediately categorised as an 'old person' (especially since she is an out-group member compared to yourself), and a stereotype is activated. You are in a rush and old people are 'slow', 'inefficient' and 'prone to confusion'. You therefore ignore this employee and go in search of a younger one, believing this to be the faster strategy.

The act of ignoring this older employee will, according to the generalisation hypothesis, have placed an inhibitory tag with the category of old people.

The next employee you encounter is another old lady. This second employee is also categorised as 'old' and stereotypes are again activated. However, this time, the inhibitory tag is also a factor. It is reinstated and results in affective devaluation of the second employee. Having rejected one category member in the past, this member suffers an even worse affective evaluation, through no fault of her own.

As such, the distractor devaluation generalisation mechanism appears to be at work, compounding and proliferating stereotypes. If a perceiver holds pre-existing negative stereotypes about a category of people, these may serve as 'disregard cues'. However, ignoring these people will lead to a subsequent emotional devaluation that will reinforce the 'disregard cues' and make matters worse in the future. This system may be responsible for increasing negative attitudes between groups of people. Any group of people which are ignored on the basis of their category membership (old

people, racial or ethnic minorities, the disabled) may suffer from additional devaluation as a result. This is surely a bad thing. Why would such a system have evolved?

The assumption is that all stereotypes are negative and inaccurate. While many stereotypes are harmful and unfounded (black people are dangerous, blondes are dumb), it is important to remember that stereotypes evolved for a reason. In some cases, they are *useful* heuristics on which to base our decisions. For example, people who wear glasses may read more than people who do not, and this may be the reason for their deteriorated eyesight. Perhaps people who wear glasses really are more intelligent. Likewise, older people do perform slower in tasks (e.g. Salthouse, 1996), and have reduced memory (e.g. Balota, Dolan & Dulchek, 2000). Perhaps trusting our initial evaluation, that we should find a younger employee to help us at the supermarket, will lead us to be on time to pick up the children from school. The distractor devaluation generalisation mechanism, while in many cases may compound negative and inaccurate stereotypes, is probably not maladaptive in nature. In the majority of cases it may be at work improving our efficiency at rejecting unsuitable choices or decisions. It is difficult to understand why, if it were wholly maladaptive, such a system would have evolved.

The Relative Contribution of Distractor Devaluation Generalisation

I am by no means suggesting that the generalisation of the distractor devaluation effect is wholly responsible for affective evaluation. Other, more stable, aspects of faces, and other stimuli, and the goal states active at the time of evaluation, will be the largest contributors to evaluation. However, throughout the series of

experiments presented in this thesis, devaluation-by-inhibition, and the generalisation of it, has had a small, but consistent influence on evaluations.

By converting distractor devaluation effects found in each experiment into effect sizes, it is possible to compare the results of each experiment. In the oddball series of experiments, a 5-point evaluation scale was used. In the go / no-go series of experiments, a 7-point evaluation scale was used. Results are thus compared by expressing the effect sizes in terms of percentage of scale change as a result of inhibition. Table 23 displays the effect sizes (the size of the distractor devaluation effect) for conditions in which individuated (old) objects were devalued. The mean effect size is 2.5% with a SD of only 0.8%, effect sizes ranged from 3.8% to 1.6%. Likewise, Table 24 displays the effect sizes for conditions in which a generalised distractor devaluation effect was observed. Effect sizes ranged between 3.0% and 1.9% ($M = 2.5$, $SD = 0.6$). The distractor devaluation effect sizes observed in each condition, for each experiment are small, but remarkably consistent, varying by only 2 or 3 percent across participants, stimuli and experimental designs.

This is with the notable exception of the generalisation of the distractor devaluation in Experiment 8, which was a particularly large effect. This result reflects the difference in the ratings of elderly females from pre- to post-experiment test, by participants who had to inhibit elderly females throughout the experiment. The large effect size was thus probably a result of long-term inhibition directed to a single category over the course of the whole experiment. This is particularly interesting, because it suggests that an inhibitory tag is not binary 'all-or-nothing', but graded.

Repeated inhibition appears to strengthen the inhibitory tag associated with a category's representation.

Table 23: Individuation distractor devaluation effects

Baseline	Experiment	Condition	Effect Size (% of scale)
Target	2 Oddball Pilot	Distractors	2.0
Target	3 Oddball Bottles	Distractors	3.8
Target	7 No-Go Races	Old Distractors, Own-race	2.9
Target	9 No-Go Brands	Old Distractors	2.7
Distractor-like Novel	6 No-Go	Conjunction No-go, Positive and Negative Faces	1.6
Distractor-like Novel	8 No-Go Ages	Old Conjunction Distractors, Own-age	1.9

Table 24: Generalisation distractor devaluation effects

Baseline	Experiment	Condition	Effect Size (% of scale)
Target	4 Oddball Animals	Distractor-like Novel	3.0
Target	7 No-Go Races	No-go-like Novel, Two-match, Own race	2.9
Target	9 No-Go Brands	No-go-like Novel, Present Category	1.9
Target	9 No-Go Brands	No-go-like Novel, Absent Category	2.0
<i>Pre-experiment Evaluation</i>	<i>8 No-Go Ages</i>	<i>No-go-like post ratings, Other-Age, Whole experiment effect</i>	<i>6.4</i>

Outstanding Issues and Further Research

The conundrum of scientific research is that every effect found will, more than likely, raise more questions than it answers. I will now outline a few interesting theoretical questions that have occurred to me as a result of my thesis work, and that

might prove to be fruitful sources of further research. The first questions relate to overall distractor devaluation mechanisms that might be at work during both individual and generalised distractor devaluation. The later questions relate specifically to the mechanisms that might be at work for distractor devaluation generalisation to occur.

The first issue relates back to the notions of intuitive versus reflective appraisal processes as put forward by Arnold (1960). We can assume that inhibition affects subsequent evaluation unconsciously. That is, evaluators have no explicit memory for the attentional state active during prior exposure of the object (Kessler & Tipper, 2004). Therefore, it follows that distractor devaluation may have its influence in intuitive appraisal processes. However, in these, and to my knowledge all studies of distractor devaluation, participants are required to evaluate objects reflectively, i.e., at a conscious level, using a rating scale or forced-choice method. Is the requirement of reflective evaluation diluting the devaluation effect?

When participants are required to evaluate objects at a conscious level, they may consider many more aspects of the to-be-rated item, and the situation that it is presented in, than if they were required to only evaluate items intuitively. For example, Rotteveel et al. (Rotteveel, de Groot, Geuskens & Phaf, 2001) reported that the influence of an emotional stimulus on subsequent judgements was stronger if the emotional stimulus was presented below the level of conscious awareness. Similarly, it may be that an inhibitory tag may exert its strongest influence if conscious factors do not enter the equation. Such conscious factors may be a desire of the participant not to appear racist, or ageist for example, which may result in participants engaging

in a strategy to 'manage' or regulate their evaluations of stimuli. If evaluations were instead accessed at the intuitive level, in which presumably fewer factors are at play, the influence of an inhibitory tag might be proportionately larger, resulting in larger distractor devaluation effects. Presenting the previously inhibited object as an affective prime in a priming task, for example, could be one method by which to assess intuitive evaluations of previously inhibited stimuli.

The second interesting issue raised is whether distractor devaluation effects (especially generalisation effects) are determined by the attentional state active during the first exposure of an object (or category), or are they updated if the attentional requirements towards that object (or category) change? Does the object's (or category's) representation 'remember' the first tag, or is it wiped in favour of the most recent attentional state? Or, do attentional tags accumulate, and the net 'attention / inhibition tag' is interpreted by evaluation processes?

The findings of this thesis suggest that distractor devaluation generalisation effects are adaptable to situational demands (for example distractor devaluation generalisation can be specific to subordinate categories even if the task demands only specify basic-level inhibition). Is this flexibility evident in an updating mechanism? Experiment 9 speaks to this issue. In this experiment, the superordinate category to be inhibited swapped from edible to inedible between blocks. It was not the first block of trials that determined distractor devaluation generalisation, but instead the generalisation effect was evident for each block, suggesting an updating process was at work. This suggests that attentional (or inhibitory) tags may be wiped in favour of the most recent version.

However, the largest distractor devaluation generalisation effects were found for the other-age faces in Experiment 8. This suggests that over the course of the experiment, the inhibitory tag was reinforced, and the generalisation effect accumulated. So perhaps swapping the to-be-inhibited category between blocks in Experiment 9 was diluting generalisation effects, as conflicting signals had to be interpreted.

Related to this issue is whether distractor devaluation occurs at encoding or retrieval of an inhibited object. Does the act of inhibiting an object devalue it, or is it the reinstatement of the inhibitory tag that is subsequently translated as devaluation?

Brain imaging studies might elucidate this question. For example, perhaps amygdala activity to previously neutral objects is modulated by a sudden onset cue that directs the perceiver to either respond to, or inhibit a response to that object. If this is the case, we might expect that distractor devaluation occurs at the encoding stage.

I now turn your attention to the mechanisms that might be at work during generalisation. Specifically, the issue of how inhibitory tags might be associated with categories, and how other members of that category might access them.

How do inhibitory tags 'know' which category to associate themselves with? An old woman might be, 'old', a 'woman', or a 'supermarket employee'. The experiments presented in this thesis suggest that task demands active at the point of categorisation determine how an object is categorised. This 'contextual categorisation' is demonstrated nicely by Figure 63.

TAE CAT

Figure 63: An example of contextual categorisation. Adapted from Palmer, p. 429.

In Figure 63, the perceived identity of the letters is strongly influenced by the letters surrounding it. Normal perceivers will read the words as THE CAT. The figure is never interpreted as TAE CHT, THE CHT, or TAE CAT, even though each of these possibilities is equally compatible. The central letters of these two words are identical, but the context in which they appear determines their categorisation as either an H or an A.

However, what is the fate of the other, many, potential categories? After 'old' is chosen to be the most suitable category for a person, what happens to the categories of 'woman' or 'supermarket employee'? One viewpoint is that these competing categorisations are inhibited during the category selection process (Bodenhausen & Macrae, 1998; Macrae, Bodenhausen & Milne, 1995). That is, inappropriate categorisations are removed as potential choices through a process of spreading inhibition (Neumann & DeSchepper, 1992). As predicted by my results in this thesis, perceivers' motivational states also seem to play an influential role in the active inhibition of competing categories. Sinclair and Kunda (1999) demonstrated that after participants received favourable feedback from a black doctor, stereotypes relating to the category 'black' became significantly less accessible in their minds, but stereotypes relating to the category 'doctor' became significantly more accessible. When motivated to view a black doctor in a favourable light, participants inhibited the

category 'black' and activated the category 'doctor'. They did just the reverse, however, when the black doctor provided negative feedback, and were thus motivated to view him unfavourably.

If these other potential categories are inhibited, does this mean that the inhibitory tag that was activated as a result of ignoring our old female supermarket employee becomes associated exclusively with the category 'old', and not with the categories 'female' or 'supermarket employee'? That is, we can assume that distractor devaluation will generalise to other members of the 'old' category, but critically, if our old lady is presented in the context of other supermarket employees, will another employee of the supermarket be devalued? Can distractor devaluation generalisation cross categories depending on the context at *evaluation*?

This will presumably be dependent on how the inhibitory tag is stored in memory. With an individual representation, is it in a way that is accessible to other members of that category, so it 'jumps' from one individual to another? Or, rather than an 'inhibitory tag', is it more accurately an 'inhibitory blanket' across a whole category? If the latter is true, then it follows that only members that belong to the original to-be-ignored category will be devalued. However, if the former is true, then changing the context at evaluation might mean that the nature of the generalisation effect is changed. Inhibiting our old supermarket lady might devalue any old people if she is presented for evaluation amongst old and young people; but any supermarket employees might be devalued if she is presented for evaluation amongst supermarket employees and taxi drivers.

It is my belief that the 'blanket theory' is more likely to be accurate. The former theory requires the inhibitory tag to be encoded with an individual exemplar, in order that it can be accessed by multiple category members, dependent on context. In order for this to be this case, some level of individuation must have occurred at the time of inhibition. Therefore I would predict, in line with the individuation hypothesis, that only that exemplar would be subject to distractor devaluation and no generalisation effects would occur.

Conclusions

In this thesis I have attempted to firmly establish the existence of the distractor devaluation generalisation effect. The series of experiments I presented suggested that this effect is rooted in the nature of encoding at the time of stimulus presentation. According to my generalisation hypothesis, generalisation effects emerge when the initial to-be-inhibited stimulus is not encoded as a unique token exemplar. Instead, only type information is encoded, and as such any subsequent stimulus of the same type is prone to devaluation. I hypothesised that in such cases, attentional (specifically inhibitory) information is associated with the category representation, and as such it is accessible by, and will have affective consequences for, any member of the category presented for subsequent evaluation.

This individuation versus generalisation hypothesis was supported by my experiments. Object classes that were likely to be individuated demonstrated token distractor devaluation effects, while object classes that were unlikely to be individuated demonstrated type distractor devaluation effects. Further to this, when stimuli had to be processed deeply in order to be rejected as a distractor (i.e.

conjunction no-gos), distractor devaluation effects were large. However stimuli that could be processed minimally in order to be rejected (featural no-gos) showed no resultant distractor devaluation effects. This suggested that in order for distractor devaluation effects to manifest, a strong object representation must be encoded into memory, in order for an inhibitory tag to be successfully associated with it.

The two paradigms used in this thesis draw on two different types of inhibitory process. In the oddball experiments, inhibition was applied to the distractor object representation during attentional selection. In the go / no-go experiments, inhibition was applied to the motor response associated with a distractor object. Both types of inhibition resulted in distractor devaluation effects. However, it appears that stronger distractor devaluation effects were found in the go / no-go experiments compared to oddball experiments. In the oddball experiments, distractor devaluation did not survive a lag. Only distractors that were evaluated after a short delay (a few seconds) were devalued. However, in the go / no-go experiments, evaluations did not occur until several seconds and up to a minute (and several intervening trials) later. This suggests that the inhibitory tag created by motor inhibition is stronger than that of object representation inhibition, resulting in larger distractor devaluation effects.

This links to effort. It may be harder to withhold a response to a singly presented distractor than not to select a distractor from a multiple object array. If this is the case, a stronger inhibitory tag might result from the former scenario. This was echoed in the correlation found in Experiment 6 in which participants who were struggling to perform the task efficiently demonstrated larger distractor devaluation effects. A larger inhibitory tag appears to mean a larger distractor devaluation effect.

The findings presented combined to suggest that the distractor devaluation, and generalisation mechanisms have the following important characteristics: First, devaluation effects are a result of the effortful inhibition of an object; distractors must compete for attention. Devaluation effects may only emerge once the strength of an 'inhibitory tag' associated with an object or category has reached a threshold level. Related to this is that, second, the inhibitory tags associated with objects and categories may be graded, and repeated inhibition to a particular category may reinforce and strengthen the inhibitory tag associated with it. This in turn may lead to greater devaluation effects. It appears that this inhibitory tag must then be associated with a successfully encoded representation, which requires deep processing of the distractor object. The inhibitory tag will be associated with either a token, or a type object representation, depending on whether the object has been successfully individuated. If the distractor was individuated, token distractor devaluation effects will result. If the object was not individuated, the categorisation of an object appears to be a flexible process that can adapt to task demands. As such, the inhibitory tag may be associated with the most appropriate category in a given situation, and this will be reflected in the generalisation effects that manifest.

References

- Adolphs, R. (2002). Recognizing emotion from facial expressions: Psychological and neurological mechanisms. *Behavioral and Cognitive Neuroscience Reviews, 1*(1), 21–62.
- Allport, G. W. (1954). *The nature of prejudice*. Oxford, England: Addison Wesley.
- Anderson, A. K. & Phelps, E. A. (2001). Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature, 411*(6835), 305-309.
- Anderson, A. K. (2005). Affective Influences on the Attentional Dynamics Supporting Awareness. *Journal of Experimental Psychology: General, 134*(2), 258-281.
- Anderson, A. K., Christoff, K., Panitz, D., De Rosa, E. & Gabrieli, J. D. E. (2003). Neural Correlates of the Automatic Processing of Threat Facial Signals. *Journal of Neuroscience, 23*(13), 5627-5633.
- Anthony, T., Copper, C. & Mullen, B. (1992). Cross-racial facial identification: A social cognitive integration. *Personality and Social Psychology Bulletin, 18*(3), 296-301.
- Armony, J. L. & Dolan, R. J. (2002). Modulation of spatial attention by fear-conditioned stimuli: An event-related fMRI Study. *Neuropsychologia, 40*, 817-826.
- Arnell, K. M., Killman, K. V. & Fijavz, D. (2007). Blinded by emotion: Target misses follow attention capture by arousing distractors in RSVP. *Emotion, 7*(3), 465-477.
- Arnold, M. (1960). *Emotion and personality*. New York: Columbia University Press.

- Bäckman, L. (1991). Recognition memory across the adult life span: The role of prior knowledge. *Memory & Cognition*, 19(1), 63-71.
- Baeyens, F., Eelen, P. & Van den Bergh, O. (1990). Contingency awareness in evaluative conditioning: A case for unaware affective-evaluative learning. *Cognition & Emotion. Special Issue: Evaluative conditioning*, 4(1), 3-18.
- Baeyens, F., Heremans, R., Eelen, P. & Crombez, G. (1993). Hidden-covariation detection and imagery ability. *European Journal of Cognitive Psychology*, 5(4), 435-456.
- Bahcall, D. O. & Kowler, E. (1999). Attentional interference at small spatial separations. *Vision Research*, 39(1), 71-86.
- Balota, D. A., Dolan, P. O. & Duchek, J. M. (2000). Memory changes in healthy older adults. In E. Tulving & F. I. M. Craik (Eds.), *The Oxford Handbook of Memory*. New York: Oxford University Press.
- Bar, M. & Aminoff, E. (2003). Cortical analysis of visual context. *Neuron*, 38, 347-358.
- Bargh, J. A. (1997). The automaticity of everyday life. *Advances in social cognition*, 10, 1-61.
- Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology*. New York: Guilford.
- Bargh, J. A., Chaiken, S., Raymond, P. & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology*, 32(1), 104-128.

- Barkley, R. A. (1997). Behavioral inhibition, sustained attention, and executive functions: Constructing a unifying theory of ADHD. *Psychological Bulletin*, 121(1), 65-94.
- Bartlett, J. C. & Leslie, J. E. (1986). Aging and memory for faces versus single views of faces. *Memory & Cognition*, 14(5), 371-381.
- Bavelier, D. & Potter, M. C. (1992). Visual and phonological codes in repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, 18(1), 134-147.
- Bavelier, D. (1994). Repetition blindness between visually different items: The case of pictures and words. *Cognition*, 51(3), 199-236.
- Bavelier, D. (1999). Role and nature of object representations in perceiving and acting. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli*. Cambridge, MA, US: MIT Press.
- Bechara, A. & Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games and Economic Behavior. Special Issue on Neuroeconomics*, 52(2), 336-372.
- Bechara, A., Damasio, A. R., Damasio, H. & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1-3), 7-15.
- Bichot, N. P. & Schall, J. D. (2002). Priming in macaque frontal cortex during popout visual search: Feature-based facilitation and location-based inhibition of return. *Journal of Neuroscience*, 22(11), 205-219.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94(2), 115-147.

- Biederman, I. (1988). Aspects and extension of a theory of human image understanding. In Z. Pylyshyn (Ed.), *Computational Processes in Human Vision: An Interdisciplinary Perspective*. Norwood, NJ: Ablex Publishing Corporation.
- Biederman, I., Mezzanotte, R. J. & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology*, 14(2), 143-177.
- Bodenhausen, G. V. & Macrae, C. N. (1998). Stereotype activation and inhibition. In R. S. Wyer, Jr, (Ed.), *The Automaticity of Everyday Life*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Bodenhausen, G. V., Macrae, C. N. & Garst, J. (1998). Stereotypes in thought and deed: Social-cognitive origins of intergroup discrimination. In C. Sedikides, J. Schopler & C. A. Insko (Eds.), *Intergroup Cognition and Intergroup Behavior*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Bodenhausen, G. V., Macrae, C. N. & Hugenberg, K. (2003). Activating and inhibiting social identities: Implications for perceiving the self and others. In G. Bodenhausen & A. J. Lambert (Eds.), *Foundations of Social Cognition*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Bodenhausen, G. V., Macrae, C. N. & Sherman, J. S. (1999). On the dialectics of discrimination: Dual processes in social stereotyping. In S. Chaiken & Y. Trope (Eds.), *Dual-Process Theories in Social Psychology*. New York, NY: Guilford Press.
- Bomba, P. C. & Siqueland, E. R. (1983). The nature and structure of infant form categories. *Journal of Experimental Child Psychology*, 35(2), 294-328.

- Bonanno, G. A. & Stillings, N. A. (1986). Preference, familiarity and recognition after repeated brief exposures to random geometric shapes. *American Journal of Psychology*, 99, 403-415.
- Bonnet, M., Bradley, M. M., Lang, P. J. & Requin, J. (1995). Modulation of spinal reflexes: Arousal, pleasure, action. *Psychophysiology*, 32(4), 367-372.
- Bornstein, R. F. & D'Agostino, P. R. (1992). Stimulus recognition and the mere exposure effect. *Journal of Personality and Social Psychology*, 63(4), 545-552.
- Bornstein, R. F. & D'Agostino, P. R. (1994). The attribution and discounting of perceptual fluency: Preliminary tests of a perceptual fluency/attributional model of the mere exposure effect. *Social Cognition*, 12(2), 103-128.
- Bornstein, R. F. (1989). Exposure and affect: Overview and meta-analysis of research, 1968-1987. *Psychological Bulletin*, 106(2), 265-289.
- Bornstein, R. F., Kale, A. R. & Cornell, K. R. (1990). Boredom as a limiting condition on the mere exposure effect. *Journal of Personality and Social Psychology*, 58(5), 791-800.
- Bornstein, R. F., Leone, D. R. & Galley, D. J. (1987). The generalizability of subliminal mere exposure effects: Influence of stimuli perceived without awareness on social behavior. *Journal of Personality and Social Psychology*. Special Issue: Integrating personality and social psychology, 53(6), 1070-1079.
- Both, S., Everaerd, W. & Laan, E. (2003). Modulation of spinal reflexes by aversive and sexually appetitive stimuli. *Psychophysiology*, 40(2), 174-183.

- Bradley, B. P., Mogg, K. & Millar, N. H. (2000). Covert and overt orienting of attention to emotional faces in anxiety. *Cognition & Emotion*, 14(6), 789-808.
- Bradley, I. L. (1971). Repetition as a factor in the development of musical preferences. *Journal of Research in Music Education*, 19(3), 295-298.
- Bradley, M. M., Greenwald, M. K., Petry, M. C. & Lang, P. J. (1992). Remembering pictures: Pleasure and arousal in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(2), 379-390.
- Brewer, J. D. (1988). Theories of Race and Ethnic Relations. *Patterns of Prejudice*, 22(1), 52-53.
- Brewer, M. B. & Feinstein, A. S. H. (1999). Dual processes in the cognitive representation of persons and social categories. In S. Chaiken & Y. Trope (Eds.), *Dual-Process Theories in Social Psychology*. New York, NY: Guilford Press.
- Broadbent, D. E. (1958). *Perception and communication*. New York: Pergamon Press.
- Bruce, A. J., Harman, M. J. & Turner, M. T. (2007). Reactions to Heard Words: Varying Emotional Content and Frequency of Exposure. *Journal of General Psychology*, 134(1), 67-82.
- Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, 64(2), 123-152.
- Bullier, J. (2001). Feedback connections and conscious vision. *Trends in Cognitive Sciences*, 5(9), 369-370.

- Buodo, G., Sarlo, M. & Palomba, D. (2002). Attentional resources measured by reaction times highlight differences within pleasant and unpleasant, high arousing stimuli. *Motivation & Emotion*, 26(2), 123-138.
- Burke, A., Heuer, F. & Reisberg, D. (1992). Remembering emotional events. *Memory & Cognition*, 20(3), 277-290.
- Bush, G., Luu, P. & Posner, M. I. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Sciences*, 4(6), 215-222.
- Buttle, H., Ball, C. K., Zhang, J. & Raymond, J. E. (2005). Semantic association of brand images at the implicit level: Evidence from repetition blindness. *Applied Cognitive Psychology*, 19(9), 1199-1210.
- Cacioppo, J. T., Marshall-Goodell, B. S., Tassinary, L. G. & Petty, R. E. (1992). Rudimentary determinants of attitudes: Classical conditioning is more effective when prior knowledge about the attitude stimulus is low than high. *Journal of Experimental Social Psychology*, 28(3), 207-233.
- Cameron, E. L., Tai, J. C. & Carrasco, M. (2002). Covert attention affects the psychometric function of contrast sensitivity. *Vision Research*, 42(8), 949-967.
- Cannon, W. B. (1915). *Bodily changes in pain, hunger, fear and rage: An account of recent researches into the function of emotional excitement*. New York, NY: D Appleton & Company.
- Caputo, G. & Guerra, S. (1998). Attentional selection by distractor suppression. *Vision Research*, 38(5), 669-689.
- Carrasco, M., Ling, S. & Read, S. (2004) Attention alters appearance. *Nature Neuroscience*, 7, 308-313.

- Carrasco, M., Penpeci-Talgar, C. & Eckstein, M. (2000). Spatial covert attention increases contrast sensitivity across the CSF: Support for signal enhancement. *Vision Research*, 40(10-12), 1203-1215.
- Carrasco, M., Williams, P. E. & Yeshurun, Y. (2002). Covert attention increases spatial resolution with or without masks: Support for signal enhancement. *Journal of Vision. Special Issue: Vision and color meeting*, 2(6), 467-479.
- Carstensen, L. L., Pasupathi, M., Mayr, U. & Nesselroade, J. R. (2000). Emotional experience in everyday life across the adult life span. *Journal of Personality and Social Psychology*, 79(4), 644-655.
- Carter, C. S., Mac Donald A. M., Ross, L. L., Stenger, V. A., Noll, D. & Cohen, J. D. (2000). Parsing executive processes: strategic versus evaluative functions of the anterior cingulate cortex. *Proceedings of the National Academy of Sciences of the U.S.A.*, 97, 1944-48.
- Carter, C. S., Botvinick, M.M. & Cohen, J.D. (1999). The contribution of the anterior cingulate cortex to executive processes in cognition. *Reviews in the Neurosciences*, 10, 49-57.
- Cauquil, A. S. v., Edmonds, G. E. & Taylor, M. J. (2000). Is the face-sensitive N170 the only ERP not affected by selective attention? *Neuroreport*, 11(10), 2167-2171.
- Cave, K. R. & Wolfe, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, 22(2), 225-271.
- Cave, K. R. & Zimmerman, J. M. (1997). Flexibility in spatial attention before and after practice. *Psychological Science*, 8(5), 399-403.

- Cave, K. R. (1999). The FeatureGate model of visual selection. *Psychological Research/Psychologische Forschung. Special Issue: Vision and visual cognition*, 62(2-3), 182-194.
- Cepeda, N. J., Cave, K. R., Bichot, N. P. & Kim, M.-S. (1998). Spatial selection via feature-driven inhibition of distractor locations. *Perception & Psychophysics*, 60(5), 727-746.
- Chance, J. E. & Goldstein, A. G. (1981). Depth of processing in response to own- and other-race faces. *Personality and Social Psychology Bulletin*, 7(3), 475-480.
- Charles, S. T., Mather, M. & Carstensen, L. L. (2003). Aging and emotional memory: The forgettable nature of negative images for older adults. *Journal of Experimental Psychology: General*, 132(2), 310-324.
- Charles, S. T., Reynolds, C. A. & Gatz, M. (2001). Age-related differences and change in positive and negative affect over 23 years. *Journal of Personality and Social Psychology*, 80(1), 136-151.
- Chelazzi, L., Miller, E. K., Duncan, J. & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, 363(6427), 345-347.
- Chen, M. & Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin*, 25(2), 215-224.
- Christianson, S.-Å. & Loftus, E. F. (1991). Remembering emotional events: The fate of detailed information. *Cognition & Emotion*, 5(2), 81-108.

- Chun, M. M. & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, 21(1), 109-127.
- Clark, L., Cools, R. & Robbins, T. W. (2004). The neuropsychology of ventral prefrontal cortex: Decision-making and reversal learning. *Brain & Cognition*, 55(1), 41-53.
- Clore, G. L., Schwarz, N. & Conway, M. (1994). Affective causes and consequences of social information processing. In R. S. Wyer, Jr. & T. K. Srull (Eds.), *Handbook of social cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Cohen, A. & Shoup, R. (1997). Perceptual dimensional constraints in response selection processes. *Cognitive Psychology*, 32(2), 128-181.
- Collin, C. A. & McMullen, P. A. (2005). Subordinate-level categorization relies on high spatial frequencies to a greater degree than basic-level categorization. *Perception & Psychophysics*, 67(2), 354-364.
- Compton, R. J., Banich, M. T., Mohanty, A., Milham, M. P., Herrington, J., Miller, G. A., et al. (2003). Paying attention to emotion: An fMRI investigation of cognitive and emotional Stroop tasks. *Cognitive, Affective & Behavioral Neuroscience*, 3(2), 81-96.
- Cox, D. S. & Cox, A. D. (1988). What does familiarity breed? Complexity as a moderator of repetition effects in advertisement evaluation. *Journal of Consumer Research*, 15(1), 111-116.

- Cox, D. S. & Cox, A. D. (2002). Beyond first impressions: The effects of repeated exposure on consumer liking of visually complex and simple product designs. *Journal of the Academy of Marketing Science*, 30(2), 119-130.
- Craik, F. & Lockhart, R. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Thinking and Verbal Behavior*, 11, 671-684.
- Crawford, L. E., Margolies, S. M., Drake, J. T. & Murphy, M. E. (2006). Affect biases memory of location: Evidence for the spatial representation of affect. *Cognition & Emotion*, 20(8), 1153-1169.
- Critchley, H., Daly, E., Phillips, M., Brammer, M., Bullmore, E., Williams, S., et al. (2000). Explicit and implicit neural mechanisms for processing of social information from facial expressions: A functional magnetic resonance imaging study. *Human Brain Mapping*, 9(2), 93-105.
- Cross, J. F., Cross, J. & Daly, J. (1971). Sex, race, age, and beauty as factors in recognition of faces. *Perception & Psychophysics*, 10(6), 393-396.
- Curran, T., Tanaka, J. W. & Weiskopf, D. M. (2002). An electrophysiological comparison of visual categorization and recognition memory. *Cognitive, Affective & Behavioral Neuroscience*, 2(1), 1-18.
- Cutzu, F. & Tsotsos, J. K. (2003). The selective tuning model of attention: Psychophysical evidence for a suppressive annulus around an attended item. *Vision Research*, 43(2), 205-219.
- Dalgleish, T. (1995). Performance on the emotional Stroop task in groups of anxious, expert, and control subjects: A comparison of computer and card presentation formats. *Cognition & Emotion*, 9(4), 341-362.

- Damasio, A. R. (1994). The brain binds entities and events by multiregional activation from convergence zones. In H. Gutfreund & G. Toulouse (Eds.), *Biology and Computation: A physicist's choice*. River Edge, NJ: World Scientific Publishing.
- Damasio, A. R. (2004). Emotions and Feelings: A Neurobiological Perspective. In A. Manstead, N. Frijda & A. Fischer, (Eds.), *Feelings and emotions: The Amsterdam symposium. Studies in emotion and social interaction*. New York, NY: Cambridge University Press
- Damasio, A., Everitt, B. J. & Bishop, D. (1996). The Somatic Marker Hypothesis and the Possible Functions of the Prefrontal Cortex. *Philosophical Transactions: Biological Sciences*, 351(1346) 1413-1420.
- D'Argembeau, A., Comblain, C. & Van der Linden, M. (2003). Phenomenal characteristics of autobiographical memories for positive, negative, and neutral events. *Applied Cognitive Psychology*, 17(3), 281-294.
- D'Argembeau, A., Van der Linden, M., Comblain, C. & Etienne, A.-M. (2003). The effects of happy and angry expressions on identity and expression memory for unfamiliar faces. *Cognition & Emotion*, 17(4), 609-622.
- de Gelder, B., Vroomen, J., Pourtois, G. & Weiskrantz, L. (1999). Non-conscious recognition of affect in the absence of striate cortex. *Neuroreport*, 10(18), 3759-3763.
- De Houwer, J. & Hermans, D. (1999). Nine attempts to find affective priming of pronunciation responses: Effects of SOA, degradation, and language. Paper presented at the Tagung der Fachgruppe Sozialpsychologie, Kassel, Germany.

- De Houwer, J., Hendrickx, H. & Baeyens, F. (1997). Evaluative learning with "subliminally" presented stimuli. *Consciousness & Cognition*, 6(1), 87-107.
- Dearing, M. F. & Dickinson, A. (1979). Counterconditioning of shock by a water reinforcer in rabbits. *Animal Learning & Behavior*, 7(3), 360-366.
- Deco, G. & Zihl, J. (2001). Top-down selective visual attention: A neurodynamical approach. *Visual Cognition*, 8(1), 119-140.
- DeSchepper, B. & Treisman, A. (1996). Visual memory for novel shapes: Implicit coding without attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(1), 27-47.
- Desimone, R. & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193-222.
- Devinsky, O., Morrell, M. J. & Vogt, B. A. (1995). Contributions of anterior cingulate cortex to behaviour. *Brain*, 118(1), 279-306.
- Dias, R., Robbins, T. W. & Roberts, A. C. (1997). Dissociable forms of inhibitory control within prefrontal cortex with an analog of the Wisconsin Card Sort Test: Restriction to novel situations and independence from "on-line" processing. *Journal of Neuroscience*, 17(23), 9285-9297.
- Dickinson, A. & Dearing, M. F. (1979). Appetitive-aversive interactions and inhibitory processes. A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of Learning and Motivation*. Hillsdale, NJ: Lawrence Erlbaum Associates
- Dimberg, U. & Öhman, A. (1996). Behold the wrath: Psychophysiological responses to facial stimuli. *Motivation & Emotion. Special Issue: Facial Expression and Emotion*, 20(2), 149-182.

- Dion, K., Berscheid, E. & Walster, E. (1972). What is beautiful is good. *Journal of Personality and Social Psychology*, 24(3), 285-290.
- Dipboye, R. L., Arvey, R. D. & Terpstra, D. E. (1977). Sex and physical attractiveness of raters and applicants as determinants of resumé evaluations. *Journal of Applied Psychology*, 62(3), 288-294.
- Dolan, R. J. (2000). Functional neuroimaging of the amygdala during emotional processing and learning. In J. P. Aggleton (Ed.), *The amygdala: a functional analysis*. Oxford: Oxford University Press.
- Downing, P., Liu, J. & Kanwisher, N. (2001). Testing cognitive models of visual attention with fMRI and MEG. *Neuropsychologia. Special Issue: Imaging selective attention in the human brain*, 39(12), 1329-1342.
- Drevets, W. C. & Raichle, M. E. (1998). Reciprocal suppression of regional cerebral blood flow during emotional versus higher cognitive processes: Implications for interactions between emotion and cognition. *Cognition & Emotion. Special Issue: Neuropsychological perspectives on affective and anxiety disorders*, 12(3), 353-385.
- Duncan, J. & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, 96(3), 433-458.
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113(4), 501-517.
- Eagly, A. H. & Chaiken, S. (1993). *The psychology of attitudes*. Orlando, FL: Harcourt Brace Jovanovich College Publishers.

- 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.
- Ebert, D. & Ebmeier, K. P. (1996). The role of the cingulate gyrus in depression: From functional anatomy to neurochemistry. *Biological psychiatry*, 39(12), 1044-1050.
- Edelman, S. (1998). Representation is representation of similarities. *Behavioral and Brain Sciences*, 21(4), 449-498.
- Eimer, M. (2000). Effects of face inversion on the structural encoding and recognition of faces: Evidence from event-related brain potentials. *Cognitive Brain Research*, 10(1-2), 145-158.
- Eimer, M., Holmes, A. & McGlone, F. P. (2003). The role of spatial attention in the processing of facial expression: An ERP study of rapid brain responses to six basic emotions. *Cognitive, Affective & Behavioral Neuroscience*, 3(2), 97-110.
- Ekman, P. (1972). Universals and cultural differences in facial expressions of emotion. In J. Cole (Ed.), *Nebraska Symposium on Motivation 1971*. Lincoln, NE: University of Nebraska Press.
- Ellsworth, P. (1991). Some implications of cognitive appraisal theories of emotion. In K. T. Strongman (Ed.), *International Review of Research on Emotion*. New York: Wiley.
- Epstein, R. & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392(6676), 598-601.

- Epstein, R., Graham, K. & Downing, P. (2003). Viewpoint-specific scene representations in human parahippocampal cortex. *Neuron*, 37, 865–876.
- Eriksen, B. A. & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, 16(1), 143-149.
- Esteves, F., Dimberg, U. & Öhman, A. (1994). Automatically elicited fear: Conditioned skin conductance responses to masked facial expressions. *Cognition & Emotion*, 8(5), 393-413.
- Farah, M. J., Wilson, K. D., Drain, M. & Tanaka, J. N. (1998). What is "special" about face perception? *Psychological Review*, 105(3), 482-498.
- Fazio, R. H. (1993). Variability in the likelihood of automatic attitude activation: Data reanalysis and commentary on Bargh, Chaiken, Govender, and Pratto (1992). *Journal of Personality and Social Psychology*, 64(5), 753-758.
- Fazio, R. H., Jackson, J. R., Dunton, B. C. & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, 69(6), 1013-1027.
- Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C. & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, 50(2), 229-238.
- Fenske, M. J. & Eastwood, J. D. (2003). Modulation of Focused Attention by Faces Expressing Emotion: Evidence From Flanker Tasks. *Emotion*, 3(4), 327-343.

- Fenske, M. J. & Raymond, J. E. (2006). Affective Influences of Selective Attention. *Current Directions in Psychological Science, 15*(6), 312-316.
- Fenske, M. J., Raymond, J. E. & Kunar, M. A. (2004). The affective consequences of visual attention in preview search. *Psychonomic Bulletin & Review, 11*(6), 1055-1061.
- Fenske, M. J., Raymond, J. E., Kessler, K., Westoby, N. & Tipper, S. P. (2005). Attentional Inhibition Has Social-Emotional Consequences for Unfamiliar Faces. *Psychological Science, 16*(10), 753-758.
- Field, A. P. & Moore, A. C. (2005). Dissociating the effects of attention and contingency awareness on evaluative conditioning effects in the visual paradigm. *Cognition & Emotion, 19*(2), 217-243.
- Fink, G. R., Halligan, P. W., Marshall, J. C., Frith, C. D., Frackowiak, R. S. J. & Dolan, R. J. (1997). Neural mechanisms involved in the processing of global and local aspects of hierarchically organized visual stimuli. *Brain, 120*(10), 1779-1791.
- Finn, P. R., Justus, A., Mazas, C. & Steinmetz, J. E. (1999). Working memory, executive processes and the effects of alcohol on Go/No-Go learning: Testing a model of behavioral regulation and impulsivity. *Psychopharmacology. Special Issue: Impulsivity, 146*(4), 465-472.
- Fiske, S. T. & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology*. New York, NY: Academic Press.

- Fiske, S. T. & Taylor, S. E. (1984). *Social cognition (1st ed.)*. Reading, MA: Addison-Wesley.
- Fiske, S. T. & Taylor, S. E. (1991). *Social cognition (2nd ed.)*. New York: McGraw-Hill.
- Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. T. Gilbert, S. T. Fiske & Gardner, L. (Eds.), *The Handbook of Social Psychology*. New York, NY: McGraw-Hill.
- Fiske, S. T., Lin, M. & Neuberg, S. L. (1999). The continuum model: Ten years later. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology*. New York: Guilford.
- Fox, E. (2002). Processing emotional facial expressions: The role of anxiety and awareness. *Cognitive, Affective & Behavioral Neuroscience*, 2(1), 52-63.
- 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.
- Freitas, A. L., Azizian, A., Leung, H.-C. & Squires, N. K. (2007). Resisting recently acted-on cues: Compatibility of Go/NoGo responses to response history modulates (frontal P3) event-related potentials. *Psychophysiology*, 44(1), 2-10.
- Frijda, N. H. (1986). *The emotions. Studies in emotion and social interaction*. New York, NY: Cambridge University Press.

- Fuentes, L. J., Vivas, A. B. & Humphreys, G. W. (1999). Inhibitory mechanisms of attentional networks: Spatial and semantic inhibitory processing. *Journal of Experimental Psychology: Human Perception and Performance*, 25(4), 1114-1126.
- Fulcher, E. P. (2002). Neurons with attitude: A connectionist account of human evaluative learning. In S. Moore & M. Oaksford (Eds.), *Emotional Cognition: From Brain to Behaviour*. Amsterdam: John Benjamin Publishing.
- Fulton, A. & Bartlett, J. C. (1991). Young and old faces in young and old heads: The factor of age in face recognition. *Psychology & Aging*, 6(4), 623-630.
- Fung, H. H. & Carstensen, L. L. (2003). Sending memorable messages to the old: Age differences in preferences and memory for advertisements. *Journal of Personality and Social Psychology*, 85(1), 163-178.
- Gasper, K. & Clore, G. L. (2002). Attending to the big picture: Mood and global versus local processing of visual information. *Psychological Science*, 13(1), 34-40.
- Genemation Limited (2002-2004). *GenHead 1.2*.
- Gerlach, C., Law, I., Gade, A. & Paulson, O. B. (2000). Categorization and category effects in normal object recognition: A PET study. *Neuropsychologia*, 38(13), 1693-1703.
- Ghashghaei, H. T. & Barbas, H. (2002). Pathways for emotion: Interactions of prefrontal and anterior temporal pathways in the amygdala of the rhesus monkey. *Neuroscience*, 115(4), 1261-1279.

- Giner-Sorolla, R., Garcia, M. T. & Bargh, J. A. (1999). The automatic evaluation of pictures. *Social Cognition, 17*(1), 76-96.
- Glaser, J. & Banaji, M. R. (1999). When fair is foul and foul is fair: Reverse priming in automatic evaluation. *Journal of Personality and Social Psychology, 77*(4), 669-687.
- Goolsby, B., Raymond, J. E., Silvert, L., Kiss, M., Fragopanagos, N., Taylor, J. G., Eimer, M., Nobre, A. C. & Shapiro, K. L. (in press). Affective consequences of effortful selection of faces.
- Gordon, P. C. & Holyoak, K. J. (1983). Implicit learning and generalization of the "mere exposure" effect. *Journal of Personality and Social Psychology, 45*(3), 492-500.
- Greenwald, A. G., Draine, S. C. & Abrams, R. L. (1996). Three cognitive markers of unconscious semantic activation. *Science, 273*(5282), 1699-1702.
- Greenwald, A. G., Klinger, M. R. & Liu, T. J. (1989). Unconscious processing of dichoptically masked words. *Memory & Cognition, 17*(1), 35-47.
- Grisson, S., Paul, M. A., Kessler, K. & Tipper, S. P. (2005). Inhibition of object identity in inhibition of return: Implications for encoding and retrieving inhibitory processes. *Psychonomic Bulletin & Review, 12*(3), 553-558.
- Gross, J. J., Carstensen, L. L., Pasupathi, M., Tsai, J., Götestam Skorpen, C. & Hsu, A. Y. C. (1997). Emotion and aging: Experience, expression, and control. *Psychology & Aging, 12*(4), 590-599.
- Gur, R. E., McGrath, C., Chan, R. M., Schroeder, L., Turner, T., Turetsky, B. I., et al. (2002). An fMRI study of facial emotion processing in patients with schizophrenia. *American Journal of Psychiatry, 159*(12), 1992-1999.

- Guy, S. C. & Cahill, L. (1999). The role of overt rehearsal in enhanced conscious memory for emotional events. *Consciousness & Cognition*, 8(1), 114-122.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108(4), 814-834.
- Hajcak, G., Molnar, C., George, M. S., Bolger, K., Koola, J. & Nahas, Z. (2007). Emotion facilitates action: A transcranial magnetic stimulation study of motor cortex excitability during picture viewing. *Psychophysiology*, 44(1), 91-97.
- Halgren, E. (1992). Emotional neurophysiology of the amygdala within the context of human cognition. In J. P. Aggleton (Ed.) *The amygdala: a functional analysis*. Oxford: Oxford University Press.
- Halpern, A.R. & O'Connor, M. G. (2000). Implicit memory for music in Alzheimer's disease. *Neuropsychology*, 14, 391-397.
- Hamann, S. B., Cahill, L. & Squire, L. R. (1997). Emotional perception and memory in amnesia. *Neuropsychology*, 11(1), 104-113.
- Hamann, S. B., Monarch, E. S. & Goldstein, F. C. (2000). Memory enhancement for emotional stimuli is impaired in early Alzheimer's disease. *Neuropsychology*, 14(1), 82-92.
- Hamermesh, D. S. & J. E. Biddle (1994). Beauty and the Labor Market. *The American Economic Review*, 84(5), 1174-1194.
- Hamilton, A. C. & Martin, R. C. (2005). Dissociations among tasks involving inhibition: A single-case study. *Cognitive, Affective & Behavioral Neuroscience*, 5(1), 1-13.

- Hammerl, M. & Grabitz, H.-J. (2000). Affective-evaluative learning in humans: A form of associative learning or only an artifact? *Learning & Motivation*, 31(4), 345-363.
- Hampton, C., Purcell, D. G., Bersine, L., Hansen, C. H. & et al. (1989). Probing "pop-out": Another look at the face-in-the-crowd effect. *Bulletin of the Psychonomic Society*, 27(6), 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(6), 917-924.
- Harmon-Jones, E. & Allen, J. J. B. (2001). The role of affect in the mere exposure effect: Evidence from psychophysiological and individual differences approaches. *Personality and Social Psychology Bulletin*, 27(7), 889-898.
- Harrison, A. A. (1977). *Mere exposure*. San Diego: Academic Press.
- Haxby, J. V., Hoffman, E. A. & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, 4(6), 223-233.
- Heinke, D. & Humphreys, G. W. (2003). Attention, spatial representation, and visual neglect: Simulating emergent attention and spatial memory in the selective attention for identification model (SAIM). *Psychological Review*, 110(1), 29-87.
- Hermans, D., De Houwer, J. & Eelen, P. (1994). The affective priming effect: Automatic activation of evaluative information in memory. *Cognition & Emotion*, 8(6), 515-533.

- Hermans, D., Spruyt, A. & Eelen, P. (2003). Automatic affective priming of recently acquired stimulus valence: Priming at SOA 300 but not at SOA 1000. *Cognition & Emotion*, 17(1), 83-99.
- Hermans, D., Van den Broeck, A. & Eelen, P. (1998). Affective priming using a color-naming task: A test of an affective-motivational account of affective priming effects. *Zeitschrift fur experimentelle Psychologie: Organ der Deutschen Gesellschaft fur Psychologie*, 45(2), 136-148.
- Hershey, T., Black, K. J., Hartlein, J., Braver, T. S., Barch, D. M., Carl, J. L., et al. (2004). Dopaminergic modulation of response inhibition: An fMRI study. *Cognitive Brain Research*, 20(3), 438-448.
- 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(2), 174-184.
- Hopfinger, J. B., Buonocore, M. H. & Mangun, G. R. (2000). The neural mechanisms of top-down attentional control. *Nature Neuroscience*, 3(3), 284-291.
- Horstmann, G. & Bauland, A. (2006). Search Asymmetries With Real Faces: Testing the Anger-Superiority Effect. *Emotion*, 6(2), 193-207.
- Hoss, R. A., Ramsey, J. L., Griffin, A. M. & Langlois, J. H. (2005). The role of facial attractiveness and facial masculinity/femininity in sex classification of faces. *Perception*, 34(12), 1459-1474.
- Houghton, G. & Tipper, S. P. (1994). A model of inhibitory mechanisms in selective attention. In D. Dagenbach & T. Carr (Eds.), *Inhibitory Mechanisms in Attention Memory on Language*. San Diego: Academic Press.

- Houwer, J. D., Hermans, D. & Eelen, P. (1998). Affective and identity priming with episodically associated stimuli. *Cognition & Emotion*, 12(2), 145-169.
- Hoyer, W. D. & MacInnis, D. J. (1997). *Consumer Behaviour*. Boston: Houghton Mifflin Company.
- Hugenberg, K., Miller, J. & Claypool, H. M. (2007). Categorization and individuation in the cross-race recognition deficit: Toward a solution to an insidious problem. *Journal of Experimental Social Psychology*, 43(2), 334-340.
- Humphreys, G. W. & Müller, H. J. (1993). SEArch via Recursive Rejection (SERR): A connectionist model of visual search. *Cognitive Psychology*, 25(1), 43-110.
- Intraub, H. (1984). Conceptual masking: The effects of subsequent visual events on memory for pictures. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1), 115-125.
- Intraub, H. (1997). The representation of visual scenes. *Trends in Cognitive Sciences*, 1(6), 217-222.
- Isen, A. M. & Means, B. (1983). The influence of positive affect on decision-making strategy. *Social Cognition*, 2(1), 18-31.
- Jackson, M. C. & Raymond, J. E. (2006). The role of attention and familiarity in face identification. *Perception & Psychophysics*, 68(4), 543-557.
- Jacoby, L. L. & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General*, 110(3), 306-340.
- Jacoby, L. L. & Kelley, C. M. (1987). Unconscious influences of memory for a prior event. *Personality and Social Psychology Bulletin*, 13(3), 314-336.

- Jacoby, L. L., Kelley, C. M. & Dywan, J. (1989). Memory attributions. In E. Henry, L. Roediger, III, E. Fergus & I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving*. Hillsdale: Erlbaum.
- Jacoby, L. L., Kelley, C., Brown, J. & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology*, 56(3), 326-338.
- James, W. (1884). What is an emotion? *Mind*, 9, 188-205
- Johnson, E. J. & Tversky, A. (1983). Affect, generalization, and the perception of risk. *Journal of Personality and Social Psychology*, 45(1), 20-31.
- Johnston, R. A. & Ellis, H. D. (1995). Age effects in the processing of typical and distinctive faces. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 48A(2), 447-465.
- Jolicœur, P. (1998). Modulation of the attentional blink by on-line response selection: Evidence from speeded and unspeeded Task-sub-1 decisions. *Memory & Cognition*, 26(5), 1014-1032.
- Jolicœur, P. (1999). Dual-task interference and visual encoding. *Journal of Experimental Psychology: Human Perception and Performance*, 25(3), 596-616.
- Jordan, H. & Tipper, S. P. (1998). Object-based inhibition of return in static displays. *Psychonomic Bulletin & Review*, 5(3), 504-509.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice Hall.
- Kanwisher, N. G. & Potter, M. C. (1990). Repetition blindness: Levels of processing. *Journal of Experimental Psychology: Human Perception and Performance*, 16(1), 30-47.

- Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, 27(2), 117-143.
- Kanwisher, N., Yin, C. & Wojciulik, E. (1999). Repetition blindness for pictures: Evidence for the rapid computation of abstract visual descriptions. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli*. Cambridge, MA, US: MIT Press.
- Kappas, A. (2006). Appraisals are direct, immediate, intuitive, and unwitting... and some are reflective. *Cognition & Emotion. Special Issue: Magda B. Arnold's contributions to emotion research and theory*, 20(7), 952-975.
- Kastner, S. & Ungerleider, L. G. (2000). Mechanisms of Visual Attention in the Human Cortex. *Annual Review of Neuroscience*, 23, 315-341.
- Kawasaki, H., Adolphs, R., Kaufman, O., Damasio, H., Damasio, A. R., Granner, M., et al. (2001). Single-neuron responses to emotional visual stimuli recorded in human ventral prefrontal cortex. *Nature Neuroscience*, 4(1), 15-16.
- Keil, A. & Ihssen, N. (2004). Identification Facilitation for Emotionally Arousing Verbs During the Attentional Blink. *Emotion*, 4(1), 23-35.
- Kensinger, E. A. & Corkin, S. (2003). Memory enhancement for emotional words: Are emotional words more vividly remembered than neutral words? *Memory & Cognition*, 31, 1169-1180.
- Kensinger, E. A., Garoff-Eaton, R. J. & Schacter, D. L. (2007). Effects of emotion on memory specificity: Memory trade-offs elicited by negative visually arousing stimuli. *Journal of Memory and Language*, 56(4), 575-591.

- Kessler, K. & Tipper, S. P. (2004). Retrieval of implicit inhibitory processes: The impact of visual field, object-identity, and memory dynamics. *Visual Cognition, 11*(8), 965-995.
- Kindt, M. & Brosschot, J. F. (1997). Phobia-related cognitive bias for pictorial and linguistic stimuli. *Journal of Abnormal Psychology, 106*(4), 644-648.
- Kindt, M. & Brosschot, J. F. (1999). Cognitive bias in spider-phobic children: Comparison of a pictorial and a linguistic spider Stroop. *Journal of Psychopathology and Behavioral Assessment, 21*(3), 207-220.
- Kiss, M., Goolsby, B. A., Raymond, J. E., Shapiro, K. L., Silvert, L., Nobre, A. C., et al. (2007). Efficient attentional selection predicts distractor devaluation: Event-related potential evidence for a direct link between attention and emotion. *Journal of Cognitive Neuroscience, 19*(8), 1316-1322.
- Klauer, K. C. & Musch, J. (2001). Does sunshine prime loyal? Affective priming in the naming task. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 54A*(3), 727-751.
- Klauer, K.C. (1998). Affective priming. *European Review of Social Psychology, 8*, 67-103.
- Klein, R. (1988). Inhibitory tagging system facilitates visual search. *Nature, 334*(6181), 430-431.
- Knowlton, B. J. & Squire, L. R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science, 262*(5140), 1747-1749.

- Kolodny, J. A. (1994). Memory processes in classification learning: An investigation of amnesic performance in categorization of dot patterns and artistic styles. *Psychological Science, 5*(3), 164-169.
- Konorski, J. (1967). *Integrative activity of the brain*. Chicago: University of Chicago Press.
- Koster, E. H. W., Crombez, G., Van Damme, S., Verschuere, B. & De Houwer, J. (2004). Does Imminent Threat Capture and Hold Attention? *Emotion, 4*(3), 312-317.
- Kruglanski, A. W., Freund, T. & Bar-Tal, D. (1996). Motivational effects in the mere-exposure paradigm. *European Journal of Social Psychology, 26*(3), 479-499.
- Kunst-Wilson, W. R. & Zajonc, R. B. (1980). Affective discrimination of stimuli that cannot be recognized. *Science, 207*(4430), 557-558.
- Landy, D. & Sigall, H. (1974). Beauty is talent: Task evaluation as a function of the performer's physical attractiveness. *Journal of Personality and Social Psychology, 29*(3), 299-304.
- Lang, P. J. (1995). The emotion probe: Studies of motivation and attention. *American Psychologist, 50*(5), 372-385.
- Lang, P. J., Bradley, M. M. & Cuthbert, B. N. (1990). Emotion, attention, and the startle reflex. *Psychological Review, 97*(3), 377-395.
- Lang, P. J., Davis, M. & Öhman, A. (2000). Fear and anxiety: Animal models and human cognitive psychophysiology. *Journal of Affective Disorders. Special Issue: Arousal in anxiety, 61*(3), 137-159.

- Langlois, J. H. & Roggman, L. A. (1990). Attractive faces are only average. *Psychological Science, 1*(2), 115-121.
- Lavie, N., Ro, T. & Russell, C. (2003). The role of perceptual load in processing distractor faces. *Psychological Science, 14*(5), 510-515.
- Lavy, E. H., Van den Hout, M. & Arntz, A. (1993). Attentional bias and spider phobia: *Conceptual and clinical issues. Behaviour Research and Therapy, 31*(1), 17-24.
- Lazarus, R. S. & Alfert, E. (1964). Short-circuiting of threat by experimentally altering cognitive appraisal. *The Journal of Abnormal and Social Psychology, 69*(2), 195-205.
- Lazarus, R. S. (1982). Thoughts on the relations between emotion and cognition. *American Psychologist, 37*(9), 1019-1024.
- Lazarus, R. S. (1984). On the primacy of cognition. *American Psychologist, 39*(2), 124-129.
- Lazarus, R. S. (1999). *The cognition-emotion debate: A bit of history*. New York, NY: John Wiley & Sons Ltd.
- LeDoux, J. E. (1996). *The emotional brain: The mysterious underpinnings of emotional life*. New York, NY: Simon & Schuster.
- Lee, M. A., Sundberg, J. L. & Bernstein, I. H. (1993). Concurrent processes: The affect-cognition relationship within the context of the "mere exposure" phenomenon. *Perception & Psychophysics, 54*(1), 33-42.
- Leek, E. C., Reppa, I. & Tipper, S. P. (2003). Inhibition of return for objects and locations in static displays. *Perception & Psychophysics, 65*(3), 388-395.

- Levey, A. B. & Martin, I. (1975). Classical conditioning of human “evaluative” responses. *Behaviour Research and Therapy*, *13*, 221–226.
- Levin, D. T. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(6), 1364-1382.
- Levin, D. T. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General*, *129*(4), 559-574.
- Levine, L. J. & Bluck, S. (1997). Experienced and remembered emotional intensity in older adults. *Psychology and Aging*, *12*(3), 514-523.
- Levine, L. J. & Pizarro, D. A. (2004). Emotion and memory research: A grumpy overview. *Social Cognition*, *22*(5), 530-554.
- Levy, S. J. (1959). Symbols for sale. *Harvard Business Review*, *37*, 117-124.
- Lingle, J. H., Altom, M. W. & Medin, D. L. (1984). Of cabbages and kings: Assessing the extendability of natural object concept models to social things. In R. S. Wyer, Jr. & T. K. Srull (Eds.), *Handbook of social cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Liu, J., Higuchi, M., Marantz, A. & Kanwisher, N. (2000). The selectivity of the occipitotemporal M170 for faces. *Neuroreport*, *11*(2), 337-341.
- Liu, T., Pestilli, F. & Carrasco, M. (2005) Transient attention enhances perceptual performance and fMRI response in human visual cortex. *Neuron*, *45*, 469-477.

- Loftus, E. F. (1979). *Eyewitness testimony*. Cambridge, MA: Harvard University Press.
- Logan, G. D. & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*(3), 295-327.
- Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language*. San Diego: Academic Press.
- Lu, Z.-L. & Doshier, B. A. (2000). Spatial attention: Different mechanisms for central and peripheral temporal precues? *Journal of Experimental Psychology: Human Perception and Performance*, *26*(5), 1534-1548.
- MacKay, D. G. & Ahmetzanov, M. V. (2005). Emotion, Memory, and Attention in the Taboo Stroop Paradigm: An Experimental Analogue of Flashbulb Memories. *Psychological Science*, *16*(1), 25-32.
- MacKay, D. G. & Miller, M. D. (1994). Semantic blindness: Repeated concepts are difficult to encode and recall under time pressure. *Psychological Science*, *5*(1), 52-55.
- MacKay, D. G., Shafto, M., Taylor, J. K., Marian, D. E., Abrams, L. & Dyer, J. R. (2004). Relations between emotion, memory, and attention: Evidence from taboo Stroop, lexical decision, and immediate memory tasks. *Memory & Cognition*, *32*(3), 474-488.
- Mackintosh, N. J. (1983). *Conditioning and associative learning*. New York: Oxford University Press.
- MacLeod, C., Mathews, A. & Tata, P. (1986). Attentional bias in emotional disorders. *Journal of Abnormal Psychology*, *95*(1), 15-20.

- MacLin, O. H. & Malpass, R. S. (2001). Racial categorization of faces: The ambiguous race face effect. *Psychology, Public Policy, and Law*, 7(1), 98-118.
- MacLin, O. H. & Malpass, R. S. (2003). The ambiguous-race face illusion. *Perception*, 32(2), 249-252.
- Macrae, C. N. & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology*, 51, 93-120.
- Macrae, C. N., Bodenhausen, G. V. & Milne, A. B. (1995). The dissection of selection in person perception: Inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology*, 69(3), 397-407.
- Macrae, C. N., Milne, A. B. & Bodenhausen, G. V. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology*, 66(1), 37-47.
- Malpass, R. S. & Kravitz, J. (1969). Recognition for faces of own and other race. *Journal of Personality and Social Psychology*, 13(4), 330-334.
- Malpass, R. S. (1981). Training in face recognition. In G. M. Davies, H. D. Ellis & J. W. Shepherd (Eds.), *Perceiving and remembering faces*. London: Academic Press.
- Manza, L., Zizak, D. & Reber, A. S. (1998). Artificial grammar learning and the mere exposure effect: Emotional preference tasks and the implicit learning process. In M. A. Stadler & P. A. Frensch (Eds.), *Handbook of Implicit Learning*. Thousand Oaks, CA: Sage.
- Markman, A. B. & Brendl, C. M. (2005). Constraining Theories of Embodied Cognition. *Psychological Science*, 16(1), 6-10.

- Marr, D. & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London: Series B*, 200, 269-294.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.
- Martin, I. & Levey, A. B. (1978). Evaluative conditioning. *Advances in Behaviour Research and Therapy*, 1, 57-101.
- Maruff, P., Danckert, J., Camplin, G. & Currie, J. (1999). Behavioral goals constrain the selection of visual information. *Psychological Science*, 10(6), 522-525.
- Mather, M. & Carstensen, L. L. (2003). Aging and attentional biases for emotional faces. *Psychological Science*, 14(5), 409-415.
- Mather, M. & Johnson, M. K. (2000). Choice-supportive source monitoring: Do our decisions seem better to us as we age? *Psychology & Aging*, 15(4), 596-606.
- Mathews, A. & Bradley, B. P. (1983). Mood and the self-reference bias in recall. *Behaviour Research and Therapy*, 21(3), 233-239.
- Mathews, A. & MacLeod, C. (1994). Cognitive approaches to emotion and emotional disorders. *Annual Review of Psychology*, 45, 25-50.
- Mathews, A., Mogg, K., May, J. & Eysenck, M. (1989). Implicit and explicit memory bias in anxiety. *Journal of Abnormal Psychology*, 98(3), 236-240.
- Mathewson, K.M., Arnell, K.M. & Mansfield, C. (2008). Capturing and holding attention: The impact of emotional words in rapid serial visual presentation. *Memory & Cognition*, 36(1), 182-200.

- Matt, G. E., Vásquez, C. & Campbell, W. K. (1992). Mood-congruent recall of affectively toned stimuli: A meta-analytic review. *Clinical Psychology Review, 12*(2), 227-255.
- Mayberg, H. S. (1997). Limbic-cortical dysregulation: A proposed model of depression. *Journal of Neuropsychiatry & Clinical Neurosciences, 9*(3), 471-481.
- Mayberg, H. S., Liotti, M., Brannan, S. K., McGinnis, S., Mahurin, R. K., Jerabek, P. A., et al. (1999). Reciprocal limbic-cortical function and negative mood: Converging PET findings in depression and normal sadness. *American Journal of Psychiatry, 156*(5), 675-682.
- Maylor, E. A. & Hockey, R. (1985). Inhibitory component of externally controlled covert orienting in visual space. *Journal of Experimental Psychology: Human Perception and Performance, 11*(6), 777-787.
- McKelvie, S. J. (1995). Emotional expression in upside-down faces: Evidence for configurational and componential processing. *British Journal of Social Psychology, 34*(3), 325-334.
- Mega, M. S., Cummings, J. L., Salloway, S. & Malloy, P. (1997). The limbic system: An anatomic, phylogenetic, and clinical perspective. *Journal of Neuropsychiatry & Clinical Neurosciences, 9*(3), 315-330.
- Meier, B. P. & Robinson, M. D. (2004). Why the Sunny Side Is Up. Associations Between Affect and Vertical Position. *Psychological Science, 15*(4), 243-247.

- Meier, B. P., Robinson, M. D. & Caven, A. J. (in press). Why a big mac is a good mac: Associations between affect and size. *Basic and Applied Social Psychology*.
- Meier, B. P., Robinson, M. D. & Clore, G. L. (2004). Why Good Guys Wear White. Automatic Inferences About Stimulus Valence Based on Brightness. *Psychological Science, 15*(2), 82-87.
- Meier, B. P., Robinson, M. D., Crawford, L. E. & Ahlvers, W. J. (2007). When "light" and "dark" thoughts become light and dark responses: Affect biases brightness judgments. *Emotion, 7*(2), 366-376.
- Meissner, C. A. & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychology, Public Policy, and Law, 7*(1), 3-35.
- Merckelbach, H., Kenemans, J. L., Dijkstra, A. & Schouten, E. (1993). No attentional bias for pictorial stimuli in spider-fearful subjects. *Journal of Psychopathology and Behavioral Assessment, 15*(3), 197-206.
- Michel, C., Corneille, O. & Rossion, B. (2007). Race categorization modulates holistic encoding. *Cognitive Science, 31*, 911-924.
- Michel, C., Rossion, B., Han, J., Chung, C.-S. & Caldara, R. (2006). Holistic Processing Is Finely Tuned for Faces of One's Own Race. *Psychological Science, 17*(7), 608-615.
- Miller, E. K. & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24*, 167-202.

- Milliken, B., Joordens, S., Merikle, P. M. & Seiffert, A. E. (1998). Selective attention: A reevaluation of the implications of negative priming. *Psychological Review*, *105*(2), 203-229.
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H. & Howerter, A. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology*, *41*(1), 49-100.
- Mogg, K. & Bradley, B. P. (1999). Orienting of attention to threatening facial expressions presented under conditions of restricted awareness. *Cognition & Emotion*, *13*(6), 713-740.
- Mogg, K., Bradley, B. P., Field, M. & De Houwer, J. (2003). Eye movements to smoking-related pictures in smokers: Relationship between attentional biases and implicit and explicit measures of stimulus valence. *Addiction*, *98*(6), 825-836.
- 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.
- Monahan, J. L., Murphy, S. T. & Zajonc, R. B. (2000). Subliminal mere exposure: Specific, general, and diffuse effects. *Psychological Science*, *11*(6), 462-466.
- Monin, B. (2003). The Warm Glow Heuristic: When Liking Leads to Familiarity. *Journal of Personality and Social Psychology*, *85*(6), 1035-1048.

- Moors, A. & De Houwer, J. (2001). Automatic appraisal of motivational valence: Motivational affective priming and Simon effects. *Cognition & Emotion*, *15*(6), 749-766.
- Moray, N. (1959) Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly Journal of Experimental Psychology*, *11*, 55-60.
- Morgan, H. M., Paul, M. A. & Tipper, S. P. (2005). Inhibition of return can be associated with object identity but not with object category. *European Journal of Cognitive Psychology*, *17*(4), 499-520.
- Morris, J. S., de Gelder, B., Weiskrantz, L. & Dolan, R. J. (2001). Differential extrageniculostriate and amygdala responses to presentation of emotional faces in a cortically blind field. *Brain*, *124*(6), 1241-1252.
- Most, S. B., Smith, S. D., Cooter, A. B., Levy, B. N. & Zald, D. H. (2007). The naked truth: Positive, arousing distractors impair rapid target perception. *Cognition & Emotion*, *21*(5), 964-981.
- Mostofsky, S. H., Schafer, J. G. B., Abrams, M. T., Goldberg, M. C., Flower, A. A., Boyce, A., et al. (2003). fMRI evidence that the neural basis of response inhibition is task-dependent. *Cognitive Brain Research*, *17*(2), 419-430.
- Mounts, J. R. W. (2000a). Evidence for suppressive mechanisms in attentional selection: Feature singletons produce inhibitory surrounds. *Perception & Psychophysics*, *62*(5), 969-983.
- Mounts, J. R. W. (2000b). Attentional capture by abrupt onsets and feature singletons produces inhibitory surrounds. *Perception & Psychophysics*, *62*(7), 1485-1493.

- Mozer, M. C. & Sitton, M. (1998). Computational modeling of spatial attention. Pashler, H. (Ed.), *Attention*. Hove, England: Psychology Press/Erlbaum (UK) Taylor & Francis.
- Mroczek, D. K. (2001). Age and emotion in adulthood. *Current Directions in Psychological Science*, 10(3), 87-90.
- Müller, H. J. & von Muhlenen, A. (2000). Probing distractor inhibition in visual search: Inhibition of return. *Journal of Experimental Psychology: Human Perception and Performance*, 26(5), 1591-1605.
- Murphy, S. T. & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology*, 64(5), 723-739.
- Murphy, S. T., Monahan, J. L. & Zajonc, R. B. (1995). Additivity of nonconscious affect: Combined effects of priming and exposure. *Journal of Personality and Social Psychology*, 69(4), 589-602.
- Musch, J. & Klauer, K. C. (2001). Locational uncertainty moderates affective congruency effects in the evaluative decision task. *Cognition & Emotion*, 15(2), 167-188.
- Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, 9(3), 353-383.
- Neill, W. T., Valdes, L. A., Terry, K. M. & Gorfein, D. S. (1992). Persistence of negative priming: II. Evidence for episodic trace retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(5), 993-1000.

- Neumann, E. & DeSchepper, B. G. (1992). An inhibition-based fan effect: Evidence for an active suppression mechanism in selective attention. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 46(1), 1-40.
- Ng, W.-J. & Lindsay, R. C. L. (1994). Cross-race facial recognition: Failure of the contact hypothesis. *Journal of Cross-Cultural Psychology*, 25(2), 217-232.
- Niedenthal, P. M. (1990). Implicit perception of affective information. *Journal of Experimental Social Psychology*, 26(6), 505-527.
- Nissen, M. J. & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19(1), 1-32.
- Nothdurft, H. (1993). Faces and facial expressions do not pop out. *Perception*, 22(11), 1287-1298.
- Ochsner, K. N. & Barrett, L. F. (2001). A multiprocess perspective on the neuroscience of emotion. In T. J. Mayne & G. Bonnano (Eds.), *Emotion: Current Issues and Future Directions*. New York, NY: Guilford Press.
- Ochsner, K. N. (2000). Are affective events richly recollected or simply familiar? The experience and process of recognizing feelings past. *Journal of Experimental Psychology: General*, 129(2), 242-261.
- O'Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J. & Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*, 4(1), 95-102.
- Öhman, A. (1986). Face the beast and fear the face: Animal and social fears as prototypes for evolutionary analyses of emotion. *Psychophysiology*, 23(2), 123-145.

- Öhman, A. (1997). As fast as the blink of an eye: Evolutionary preparedness for preattentive processing of threat. In P. J. Lang, R. F. Simons & M. T. Balaban (Eds.), *Attention and orienting: Sensory and motivational processes*. Mahwah, NJ: Lawrence Erlbaum Associates Publishers.
- 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., Lundqvist, D. & Esteves, F. (2001). The face in the crowd revisited: A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology*, 80(3), 381-396.
- Oliva, A. & Schyns, P. G. (1997). Coarse blobs or fine edges? Evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive Psychology*, 34(1), 72-107.
- Oliva, A. & Schyns, P. G. (2000). Diagnostic colors mediate scene recognition. *Cognitive Psychology*, 41(2), 176-210.
- Ortony, A., Clore, G. L. & Collins, A. (1988). *The cognitive structure of emotions*. New York, NY: Cambridge University Press.
- O'Toole, A. J., Defenbacher, K. A., Valentin, D., McKee, K., Huff, D. & Abdi, H. (1998). The perception of face gender: The role of stimulus structure in recognition and classification. *Memory & Cognition*, 26(1), 146-160.
- Palermo, R. & Rhodes, G. (2007). Are you always on my mind? A review of how face perception and attention interact. *Neuropsychologia*, 45(1), 75-92.
- Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.

- Panksepp, J. (1998). *Affective neuroscience: The foundations of human and animal emotions*. New York, NY: Oxford University Press.
- Paul, M. A. & Tipper, S. P. (2003). Object-based representations facilitate memory for inhibitory processes. *Experimental Brain Research*, *148*, 283-289.
- Pavlov, I. (1927). *Conditioned reflexes*. New York, NY: Oxford University Press.
- Pessoa, L. (2005). To what extent are emotional visual stimuli processed without attention and awareness? *Current Opinion in Neurobiology*, *15*(2), 188-196.
- Pessoa, L., McKenna, M., Gutierrez, E. & Ungerleider, L.G. (2002). Neural processing of emotional faces requires attention. *Proceedings of the National Academy of Sciences*, *99*, 11458-11463.
- Peterson, M. A. & Rhodes, G. (2003). *Perception of faces, objects, and scenes: Analytic and holistic processes. Advances in visual cognition*. New York, NY: Oxford University Press.
- Petty, R. E., Schumann, D. W., Richman, S. A. & Strathman, A. J. (1993). Positive mood and persuasion: Different roles for affect under high- and low-elaboration conditions. *Journal of Personality and Social Psychology*, *64*(1), 5-20.
- Phaf, R. H. & Rotteveel, M. (2005). Affective modulation of recognition bias. *Emotion*, *5*(3), 309-318.
- Phelps, E. A., Ling, S. & Carrasco, M. (2006). Emotion Facilitates Perception and Potentiates the Perceptual Benefits of Attention. *Psychological Science*, *17*(4), 292-299.

- Pizzagalli, D., Pascual-Marqui, R. D., Nitschke, J. B., Oakes, T. R., Larson, C. L., Abercrombie, H. C., et al. (2001). Anterior cingulate activity as a predictor of degree of treatment response in major depression: Evidence from brain electrical tomography analysis. *American Journal of Psychiatry*, *158*(3), 405-415.
- Posner, M. I. & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. Bouwhuis (Eds), *Attention and performance*. London: Erlbaum.
- Posner, M. I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology*, *32*(1), 3-25.
- Posner, M.I., Nissen, M.J. & Ogden, W.C. (1977). Attended and unattended processing modes: The role of set for spatial location. In H.L. Pick & I. J. Saltzman (Eds.), *Modes of perceiving and processing information*. Hillsdale, N.J.: Erlbaum.
- Potter, M. C. & Levy, E. I. (1969). Recognition memory for a rapid sequence of pictures. *Journal of Experimental Psychology*, *81*(1), 10-15.
- Potter, M. C. (1975). Meaning in visual search. *Science*, *187*(4180), 965-966.
- Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory*, *2*(5), 509-522.
- Pratto, F. & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, *61*(3), 380-391.
- Pratto, F. (1994). Consciousness and automatic evaluation. In P. Niedenthal & S. Kitayama (Eds.), *The heart's eye: Emotional influences on perception and attention*. San Diego, CA: Academic Press.

- Proffitt, D. R. (2006). Embodied perception and the economy of action. *Perspectives on Psychological Science, 1*(2), 110-122.
- Purcell, D. G., Stewart, A. L. & Skov, R. B. (1996). It takes a confounded face to pop out of a crowd. *Perception, 25*(9), 1091-1108.
- Raymond, J. E., Fenske, M. J. & Tavassoli, N. T. (2003). Selective attention determines emotional responses to novel visual stimuli. *Psychological Science, 14*(6), 537-542.
- Raymond, J. E., Fenske, M. J. & Westoby, N. (2005). Emotional devaluation of distracting patterns and faces: A consequence of attentional inhibition during visual search? *Journal of Experimental Psychology: Human Perception and Performance, 31*(6), 1404-1415.
- Raymond, J. E., Shapiro, K. L. & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18*(3), 849-860.
- Reber, P. J. & Squire, L. R. (1994). Parallel brain systems for learning with and without awareness. *Learning & Memory, 1*(4), 217-229.
- Reber, P. J. & Squire, L. R. (1998). Encapsulation of implicit and explicit memory in sequence learning. *Journal of Cognitive Neuroscience, 10*(2), 248-263.
- Reber, R. & Schwarz, N. (2001). The hot fringes of consciousness: Perceptual fluency and affect. *Consciousness & Emotion, 2*(2), 223-231.
- Reber, R., Schwarz, N. & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review, 8*(4), 364-382.

- Reber, R., Winkielman, P. & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science, 9*(1), 45-48.
- Reed, J. M., Squire, L. R., Patalano, A. L., Smith, E. E. & Jonides, J. (1999). Learning about categories that are defined by object-like stimuli despite impaired declarative memory. *Behavioral Neuroscience, 113*(3), 411-419.
- Remington, R. W. & Folk, C. L. (2001). A dissociation between attention and selection. *Psychological Science, 12*(6), 511-515.
- Reppa, I. & Leek, E. C. (2006). Structure-based modulation of inhibition of return is triggered by object-internal but not occluding shape features. *The Quarterly Journal of Experimental Psychology, 59*(11), 1857-1866.
- Rescorla, R. A. & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black & W. F. Prokasy (Eds), *Classical Conditioning II: Current Research and Theory*. New York: Appleton Century Crofts.
- Rhodes, G., Brake, S., Taylor, K. & Tan, S. (1989). Expertise and configural coding in face recognition. *British Journal of Psychology, 80*(3), 313-331.
- Rhodes, G., Halberstadt, J. & Brajkovich, G. (2001). Generalization of mere exposure effects to averaged composite faces. *Social Cognition, 19*(1), 57-70.
- Rizzolatti, G. & Arbib, M. A. (1998). Language within our grasp. *Trends in Neurosciences, 21*(5), 188-194.
- Rizzolatti, G., Camarda, R., Fogassi, L., Gentilucci, M., Luppino, G. & Matelli, M. (1988). Functional organization of inferior area 6 in the macaque monkey: II. Area F5 and the control of distal movements. *Experimental Brain Research, 71*, 491-507.

- Robinson, M. D. (1998). Running from William James' bear: A review of preattentive mechanisms and their contributions to emotional experience. *Cognition & Emotion, 12*(5), 667-696.
- Rodin, M. J. (1987). Who is memorable to whom: A study of cognitive disregard. *Social Cognition, 5*(2), 144-165.
- Rosch, E. (1978). Principles of Categorization. In E. Rosch and B. Lloyd (Eds.), *Cognition and Categorization*. Hillsdale, NJ: Lawrence Erlbaum.
- Rosch, E., Mervis, C. B., Gray, W., Johnson, D. & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 3*, 382-439.
- Rotello, C. M., Macmillan, N. A., Reeder, J. A. & Wong, M. (2005). The remember response: Subject to bias, graded, and not a process-pure indicator of recollection. *Psychonomic Bulletin & Review. Special Issue: Memory strength and recency judgments, 12*(5), 865-873.
- Rothermund, K., Wentura, D. & Bak, P. M. (2001). Automatic attention to stimuli signalling chances and dangers: Moderating effects of positive and negative goal and action contexts. *Cognition & Emotion, 15*(2), 231-248.
- Rotteveel, M. & Phaf, R. H. (2004). Automatic Affective Evaluation Does Not Automatically Predispose for Arm Flexion and Extension. *Emotion, 4*(2), 156-172.
- Rotteveel, M., de Groot, P., Geuskens, A. & Phaf, R. H. (2001). Stronger suboptimal than optimal affective priming? *Emotion, 1*(4), 348-364.
- Rubin, D. C. & Friendly, M. (1986). Predicting which words get recalled: Measures of free recall, availability, goodness, emotionality, and pronunciability for 925 nouns. *Memory & Cognition, 14*(1), 79-94.

- Russell, J. A. & Barrett, L. F. (1999). Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology*, 76(5), 805-819.
- Salthouse, T. A. (1996). The processing-speed theory of adult age differences in cognition. *Psychological Review*, 103(3), 403-428.
- Samuel, A. G. & Kat, D. (2003). Inhibition of return: A graphical meta-analysis of its time course and an empirical test of its temporal and spatial properties. *Psychonomic Bulletin & Review*, 10(4), 897-906.
- Sangrigoli, S., Pallier, C., Argenti, A. M., Ventureyra, V. A. G. & de Schonen, S. (2005). Reversibility of the Other-Race Effect in Face Recognition During Childhood. *Psychological Science*, 16(6), 440-444.
- Schachter, S. & Singer, J. (1962). Cognitive, social, and physiological determinants of emotional state. *Psychological Review*, 69(5), 379-399.
- Scherer, K. R. (1984). On the nature and function of emotion: A component process approach. In K. R. Scherer & P. Ekman (Eds.), *Approaches to emotion*. Hillsdale, NJ: Erlbaum.
- Scherer, K. R. (1988). Criteria for emotion-antecedent appraisal: A review. In V. Hamilton, G. H. Bower & N. H. Frijda (Eds.), *Cognitive perspectives on emotion and motivation*. Boston: Kluwer.
- Schneider, W., Eshman, A. & Zuccolotto, A. (2002). *E-prime user's guide*. Pittsburgh: Psychology Software Tools Inc.
- Schwarz, N. & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology*, 45(3), 513-523.

- Schwarz, N. (2004). Metacognitive Experiences in Consumer Judgment and Decision Making. *Journal of Consumer Psychology, 14*(4), 332-348.
- Schyns, P. G. & Oliva, A. (1994). From blobs to boundary edges: Evidence for time- and spatial-scale-dependent scene recognition. *Psychological Science, 5*(4), 195-200.
- Seamon, J. G. & Delgado, M. R. (1999). Recognition memory and affective preference for depth-rotated solid objects: Part-based structural descriptions may underlie the mere exposure effect. *Visual Cognition, 6*(2), 145-164.
- Seamon, J. G., Brody, N. & Kauff, D. M. (1983). Affective discrimination of stimuli that are not recognized: Effects of shadowing, masking, and cerebral laterality. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*(3), 544-555.
- Seamon, J. G., Marsh, R. L. & Brody, N. (1984). Critical importance of exposure duration for affective discrimination of stimuli that are not recognized. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*(3), 465-469.
- Shapiro, K. L., Arnell, K. M. & Raymond, J. E. (1997). The attentional blink. *Trends in Cognitive Sciences, 1*(8), 291-296.
- Sherman, J. W., Lee, A. Y., Bessenoff, G. R. & Frost, L. A. (1998). Stereotype efficiency reconsidered: Encoding flexibility under cognitive load. *Journal of Personality and Social Psychology, 75*(3), 589-606.
- Silvert, L., Lepsien, J., Fragopanagos, N., Goolsby, B., Kiss, M., Taylor, J. G., Raymond, J. E., Shapiro, K. L., Eimer, M. & Nobre, A. C. (2007).

- Influence of attentional demands on the processing of emotional facial expressions in the amygdala. NeuroImage, 38(2), 257-366.*
- Sinclair, L. & Kunda, Z. (1999). Reactions to a Black professional: Motivated inhibition and activation of conflicting stereotypes. *Journal of Personality and Social Psychology, 77(5), 885-904.*
- Slotnick, S. D., Hopfinger, J. B., Klein, S. A. & Sutter, E.E. (2002). Darkness beyond the light: attentional inhibition surrounding the classic 'spotlight'. *Neuroreport, 13, 773-778.*
- Slotnick, S. D., Schwarzbach, J., and Yantis, S. (2003). Attentional inhibition of visual processing in human striate and extrastriate cortex. *NeuroImage, 21, 1772-1780.*
- Snyder, J. J. & Kingstone, A. (2001). Inhibition of return at multiple locations in visual search: When you see it and when you don't. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 54A(4), 1221-1237.*
- Solso, R. L. & McCarthy, J. E. (1981). Prototype formation of faces: A case of pseudo-memory. *British Journal of Psychology, 72(4), 499-503.*
- Sporer, S. L. (2001). Recognizing faces of other ethnic groups: An integration of theories. *Psychology, Public Policy, and Law, 7(1), 36-97.*
- Spruyt, A., De Houwer, J., Hermans, D. & Eelen, P. (2007). Affective Priming of Nonaffective Semantic Categorization Responses. *Experimental Psychology, 54(1), 44-53.*

- Squire, L. R. & Knowlton, B. J. (1995). Memory, hippocampus, and brain systems. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences*. Cambridge, MA: MIT Press.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18(6), 643-662.
- Suzuki, S. & Cavanagh, P. (1992). Facial expression as an emergent feature in visual search. *Investigative Ophthalmology and Visual Science*, 33, 1355.
- Swann, A. C., Bjork, J. M., Moeller, F. G. & Dougherty, D. M. (2002). Two models of impulsivity: Relationship to personality traits and psychopathology. *Biological Psychiatry*, 51(12), 988-994.
- Szpunar, K. K., Schellenberg, E. G. & Pliner, P. (2004). Liking and Memory for Musical Stimuli as a Function of Exposure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2), 370-381.
- Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of Social Issues*, 25, 79-97.
- Takeda, Y. & Yagi, A. (2000). Inhibitory tagging in visual search can be found if search stimuli remain visible. *Perception & Psychophysics* 62(5), 927-934.
- Talarico, J. M., LaBar, K. S. & Rubin, D. C. (2004). Emotional intensity predicts autobiographical memory experience. *Memory & Cognition*, 32(7), 1118-1132.
- Tanaka, D., Kiyokawa, S., Yamada, A. & Shigemasa, K. (2006). Investigating the role of selective attention in implicit learning using overlapping letter strings. Paper presented at the 28th Annual Conference of the Cognitive Science Society. Vancouver, British Columbia, Canada.

- Tarr, M. J. & Bülthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or by multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception and Performance*, 21(6), 1494-1505.
- Thayer, J. F. & Lane, R. D. (2000). A model of neurovisceral integration in emotion regulation and dysregulation. *Journal of Affective Disorders. Special Issue: Arousal in anxiety*, 61(3), 201-216.
- Thomas, L. A. & LaBar, K. S. (2005). Emotional arousal enhances word repetition priming. *Cognition & Emotion*, 19(7), 1027-1047.
- Thorpe, S. J. & Salkovskis, P. M. (1998). Studies on the role of disgust in the acquisition and maintenance of specific phobias. *Behaviour Research and Therapy*, 36(9), 877-893.
- Tipper, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 37A(4), 571-590.
- Tipper, S. P., Driver, J. & Weaver, B. (1991). Object-centred inhibition of return of visual attention. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 43A(2), 289-298.
- Tipper, S. P., Grison, S. & Kessler, K. (2003). Long-term inhibition of return of attention. *Psychological Science*, 14(1), 19-25.
- Tipper, S. P., Weaver, B., Cameron, S., Brehaut, J. C. & Bastedo, J. (1991). Inhibitory mechanisms of attention in identification and localization tasks: Time course and disruption. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(4), 681-692.

- Tipper, S. P., Weaver, B., Jerreat, L. M. & Burak, A. L. (1994). Object-based and environment-based inhibition of return of visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, 20(3), 478-499.
- Townsend, J. T. (1971). A note on the identifiability of parallel and serial processes. *Perception & Psychophysics*, 10(3), 161-163.
- Townsend, J. T. (1976). Serial and within-stage independent parallel model equivalence on the minimum completion time. *Journal of Mathematical Psychology*, 14(3), 219-238.
- Townsend, J. T. (1990). Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science*, 1(1), 46-54.
- Treisman, A. & Sato, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance*, 16(3), 459-478.
- Treisman, A. (1993). The perception of features and objects. In A. Baddeley & L. Weiskrantz (Eds.) *Attention: Selection, awareness, and control: A tribute to Donald Broadbent*. New York, NY: Clarendon Press/Oxford University Press.
- Treisman, A. M. & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97-136.
- Tsotsos, J. K., Culhane, S., Wai, W., Lai, Y., Davis, N. & Nufo, F. (1995). Modelling visual attention via selective tuning. *Artificial Intelligence*, 78, 507-545.

- Vallins, S. & Ray, A.A. (1967). Effects of cognitive dissonance on avoidance behaviour. *Journal of Personality and Social Psychology*, 7, 345-350.
- Van den Bergh, O. & Vrana, S. R. (1998). Repetition and boredom in a perceptual fluency/attributional model of affective judgements. *Cognition & Emotion*, 12(4), 533-553.
- van Honk, J., Tuiten, A., de Haan, E., van den Hout, M. & Stam, H. (2001). Attentional biases for angry faces: Relationships to trait anger and anxiety. *Cognition & Emotion*, 15(3), 279-297.
- van Honk, J., Tuiten, A., van den Hout, M., Koppeschaar, H., Thijssen, J., de Haan, E. & Verbaten, R. (2000). Conscious and preconscious selective attention to social threat: Different neuroendocrine response patterns. *Psychoneuroendocrinology*, 25(6), 577-591.
- Vogt, B. A., Finch, D. M. & Olson, C. R. (1992). Functional heterogeneity in cingulate cortex: The anterior executive and posterior evaluative regions. *Cerebral Cortex. Special Issue: Cingulate cortex*, 2(6), 435-443.
- Vogt, B. A., Nimchinsky, E.A., Vogt, L.J. & Hof, P.R. (1995). Human cingulate cortex: Surface features, flat maps, and cytoarchitecture. *Journal of Comparative Neurology*, 359, 490-506.
- Vrana, S. R., Spence, E. L. & Lang, P. J. (1988). The startle probe response: A new measure of emotion? *Journal of Abnormal Psychology*, 97(4), 487-491.
- Vuilleumier, P. & Schwartz, S. (2001a). Beware and be aware: Capture of spatial attention by fear-related stimuli in neglect. *Neuroreport*, 12(6), 1119-1122.
- Vuilleumier, P. & Schwartz, S. (2001b). Emotional facial expressions capture attention. *Neurology*, 56(2), 153-158.

- Vuilleumier, P., Armony, J. L., Driver, J. & Dolan, R. J. (2001). Effects of attention and emotion on face processing in the human brain: An event-related fMRI study. *Neuron*, 30(3), 829-841.
- Wacker, J., Heldmann, M. & Stemmler, G. (2003). Separating emotion and motivational direction in fear and anger: Effects on frontal asymmetry. *Emotion*, 3(2), 167-193.
- Walther, E. (2002). Guilty by mere association: Evaluative conditioning and the spreading attitude effect. *Journal of Personality and Social Psychology*, 82, 919-934.
- Watson, D. & Tellegen, A. (1985). Toward a consensual structure of mood. *Psychological Bulletin*, 98(2), 219-235.
- Watson, D. G. & Humphreys, G. W. (1997). Visual marking: Prioritizing selection for new objects by top-down attentional inhibition of old objects. *Psychological Review*, 104(1), 90-122.
- Watson, D. G. & Humphreys, G. W. (2000). Visual marking: Evidence for inhibition using a probe-dot detection paradigm. *Perception & Psychophysics*, 62(3), 471-481.
- Weger, U. W. & Inhoff, A. W. (2006). Semantic inhibition of return is the exception rather than the rule. *Perception & Psychophysics*, 68(2), 244-253.
- Whalen, P. J., Bush, G., McNally, R. J., Wilhelm, S., McInerney, S. C., Jenike, M. A., et al. (1998). The emotional counting Stroop paradigm: A functional magnetic resonance imaging probe of the anterior cingulate affective division. *Biological Psychiatry*, 44(12), 1219-1228.

- White, M. (1995). Preattentive analysis of facial expressions of emotion. *Cognition & Emotion, 9*(5), 439-460.
- Whittlesea, B. W. A. & Price, J. R. (2001). Implicit/explicit memory versus analytic/nonanalytic processing: Rethinking the mere exposure effect. *Memory & Cognition, 29*(2), 234-246.
- Whittlesea, B. W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*(6), 1235-1253.
- Williams, C. C., Henderson, J. M. & Zacks, R. T. (2005). Incidental visual memory for targets and distractors in visual search. *Perception & Psychophysics, 67*(5), 816-827.
- Williams, J. M. G., Watts, F. N., MacLeod, C. & Mathews, A. (1988). *Cognitive psychology and emotional disorders*. Oxford, England: John Wiley & Sons.
- Winkielman, P. & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: Psychophysiological evidence that processing facilitation elicits positive affect. *Journal of Personality and Social Psychology, 81*(6), 989-1000.
- Winkielman, P., Schwarz, N., Fazendeiro, T. A. & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion*. Mahwah, NJ: Lawrence Erlbaum.
- Wittenbrink, B., Judd, C. M. & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology, 72*(2), 262-274.
- Wolfe, J. M. & Pokorny, C. W. (1990). Inhibitory tagging in visual search: A failure to replicate. *Perception & Psychophysics, 48*(4), 357-362.

- Wolfe, J. M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, 1(2), 202-238.
- Wolfe, J. M., Cave, K. R. & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 15(3), 419-433.
- Wright, D. B. & Stroud, J. N. (2002). Age differences in lineup identification accuracy: People are better with their own age. *Law and Human Behavior*, 26(6), 641-654.
- Yeshurun, Y. & Carrasco, M. (1998). Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, 396(6706), 72-75.
- Yeshurun, Y. & Carrasco, M. (2000). The locus of attentional effects in texture segmentation. *Nature Neuroscience*, 3(6), 622-627.
- Yiend, J. & Mathews, A. (2001). Anxiety and attention to threatening pictures. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 54A(3), 665-681.
- Yonelinas, A. P. (2002). Components of episodic memory: The contribution of recollection and familiarity. In A. Baddeley, J. P. Aggleton & M. A. Conway (Eds.), *Episodic memory: New directions in research*. New York, NY: Oxford University Press.
- Zajonc, R. B. (1968). Attitudinal Effects of Mere Exposure. *Journal of Personality and Social Psychology*, 9(2), 1-27.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35(2), 151-175.

Zajonc, R. B. (2001). Mere exposure: A gateway to the subliminal. *Current Directions in Psychological Science*, 10(6), 224-228.

Zhou, H., Wan, L. & Fu, X. (2007). Generalized "stigma": Evidence for devaluation-by-inhibition hypothesis from implicit learning. In, A Paiva, R. Prada. & R. W. Picard (Eds.), *ACII, LNCS 4738*, Heidelberg.