

Salience asymmetries in the Implicit Association Test

Author: Chang, Betty

Publication Date: 2008

DOI: https://doi.org/10.26190/unsworks/7259

License:

https://creativecommons.org/licenses/by-nc-nd/3.0/au/ Link to license to see what you are allowed to do with this resource.

Downloaded from http://hdl.handle.net/1959.4/41524 in https:// unsworks.unsw.edu.au on 2024-05-03



SALIENCE ASYMMETRIES IN THE IMPLICIT ASSOCIATION TEST

BETTY CHANG

PLEASE TYPE

THE UNIVERSITY OF NEW SOUTH WALES Thesis/Dissertation Sheet

Surname or Family name: Chang

First name: Betty

Abbreviation for degree as given in the University calendar: $\ PhD$

School: Psychology

Faculty: Science

Other name/s: Pei Ing

Title: Salience asymmetries in the Implicit Association Test

Abstract 350 words maximum: (PLEASE TYPE)

The Implicit Association Test (IAT) is the most popular indirect measure of attitudes in social psychology. It has be suggested that salience asymmetries are a non-associative contaminant of the IAT that threatens the accurate assessment of attitudes. Salience asymmetries in the IAT are claimed to correspond with visual search asymmetries, and differences in targ familiarity. In this thesis, I propose that processing fluency is the common mechanism underlying both visual search asymmetries, familiarity, or process fluency most reliably corresponds with salience asymmetry effects in the IAT.

The first series of experiments revealed that processing fluency is a better predictor of salience asymmetry effects ir. IAT than is visual search asymmetry (Chapter 2). In Chapter 3, a novel method was developed to distinguish between the eff of valence and salience in the IAT. Using this method, I demonstrated that the effects of salience in the IAT are consistent wi fluency account of salience asymmetries. Familiarity was also shown to produce salience asymmetry effects in the IAT (Cha₁ 4), which is also consistent with the fluency account. When fluency and familiarity were set against each other in Chapter 5, was processing fluency, rather than familiarity, that predicted salience asymmetry effects in the IAT. Although processing fluency is a good predictor of salience asymmetries, the results of Chapter 6 reveal that the fluency account cannot explain al examples of salience asymmetries in the IAT.

The data presented here are consistent with the view that the more fluently processed target category is compatible w the pleasant attributes on the grounds of salience asymmetries. The current experiments suggest that when there are valence differences between the target categories, salience asymmetries can potentially distort IAT effects. When the positive target category is more salient, salience asymmetries appear to increase IAT effects. In contrast, when the negative target category is more salient, salience asymmetries appear to decrease IAT effects. However, further evidence is required to determine how th effects of salience and valence combine in the IAT.

Declaration relating to disposition of project thesis/dissertation

I hereby grant to the University of New South Wales or its agents the right to archive and to make available my thesis or dissertation in whole o part in the University libraries in all forms of media, now or here after known, subject to the provisions of the Copyright Act 1968. I retain all property rights, such as patent rights. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertat

I also authorise University Microfilms to use the 350 word abstract of my thesis in Dissertation Abstracts International (this is applicable to docto theses only).

The University recognises that there may be exceptional circumstances requiring restrictions on copying or conditions on use. Requests for restriction for a period of up to 2 years must be made in writing. Requests for a longer period of restriction may be considered in exceptional circumstances and require the approval of the Dean of Graduate Research.

Dote of completion: 14-10-2008

Salience Asymmetries in the Implicit Association Test

Betty Chang

2008

Thesis submitted in partial fulfillment of the requirements of the degree of Doctor of

Philosophy at the University of New South Wales

Declaration of Originality

I hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, or substantial proportions of material which have been accepted for the award of any other degree or diploma at UNSW or any other educational institution, except where due acknowledgement is made in the thesis. Any contribution made to the research by others, with whom I have worked at UNSW or elsewhere, is explicitly acknowledged in the thesis. I also declare that the intellectual content of this thesis is the product of my own work, except to the extent that assistance from others in the project's design and conception or in style, presentation and linguistic expression is acknowledged.

Date .03/12/2008

COPYRIGHT STATEMENT

'I hereby grant the University of New South Wales or its agents the right to archive and to make available my thesis or dissertation in whole or part in the University libraries in all forms of media, now or here after known, subject to the provisions of the Copyright Act 1968. I retain all proprietary rights, such as patent rights. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

I also authorise University Microfilms to use the 350 word abstract of my thesis in Dissertation Abstract International (this is applicable to doctoral theses only).

I have either used no substantial portions of copyright material in my thesis or I have obtained permission to use copyright material; where permission has not been granted I have applied/will apply for a partial restriction of the digital copy of my thesis or dissertation.'

Signed	Br	
Date	08/12/2008	

AUTHENTICITY STATEMENT

'I certify that the Library deposit digital copy is a direct equivalent of the final officially approved version of my thesis. No emendation of content has occurred and if there are any minor variations in formatting, they are the result of the conversion to digital format.'

Signed		
Date	03/12/2008	

Acknowledgements

My deepest gratitude goes to Chris Mitchell, for his unfailing commitment, patience, and assistance in this project. He has taught me the value of independent thinking, scientific rigor, and clarity – qualities which I hope to bring to my future work.

I would like to extend my appreciation to Bill von Hippel, for guiding me through the mysterious world of social psychology. His feedback was always prompt and considered, and his knowledge of social psychology verged on the encyclopedic.

Special thanks are due to the following people for their support. To my academic siblings, Oren Griffiths and Yvonna Lavis, for enriching my PhD experience through their camaraderie. Thanks to Jee Hyun Kim, for always providing a sensible and empathic perspective. I am also grateful to Simon Byrne, for believing in me, and encouraging my passion for data analysis.

Abstract

The Implicit Association Test (IAT) is the most popular indirect measure of attitudes in social psychology. It has been suggested that salience asymmetries are a non-associative contaminant of the IAT that threatens the accurate assessment of attitudes. Salience asymmetries in the IAT are claimed to correspond with visual search asymmetries, and differences in target familiarity. In this thesis, I propose that processing fluency is the common mechanism underlying both visual search asymmetries and familiarity. Several experiments were conducted to determine whether visual search asymmetries, familiarity, or processing fluency most reliably corresponds with salience asymmetry effects in the IAT.

The first series of experiments revealed that processing fluency is a better predictor of salience asymmetry effects in the IAT than is visual search asymmetry (Chapter 2). In Chapter 3, a novel method was developed to distinguish between the effects of valence and salience in the IAT. Using this method, I demonstrated that the effects of salience in the IAT are consistent with a fluency account of salience asymmetries. Familiarity was also shown to produce salience asymmetry effects in the IAT (Chapter 4), which is also consistent with the fluency account. When fluency and familiarity were set against each other in Chapter 5, it was processing fluency, rather than familiarity, that predicted salience asymmetry effects in the IAT. Although processing fluency is a good predictor of salience asymmetries, the results of Chapter 6 reveal that the fluency account cannot explain all examples of salience asymmetries in the IAT.

The data presented here are consistent with the view that the more fluently processed target category is compatible with the pleasant attributes on the grounds of salience asymmetries. The current experiments suggest that when there are valence differences between the target categories, salience asymmetries can potentially distort IAT effects. When the positive target category is more salient, salience asymmetries appear to increase IAT effects. In contrast, when the negative target category is more salient, salience asymmetries appear to decrease IAT effects. However, further evidence is required to determine how the effects of salience and valence combine in the IAT.

Table of Contents

Declaration of Originality	ii
Acknowledgements	iii
Abstract	iv
Table of Contents	vi
List of Figures	xi
Chapter 1. Introduction	. 13
1.1. The Implicit Association Test	. 14
1.2. Validity of the IAT	. 17
1.2.1. Correspondence of the IAT with behavioral measures	. 17
1.2.2. Correspondence of the IAT with neuroimaging measures	. 18
1.2.3. Contextual effects in the IAT	. 20
1.3. Associations in the IAT	. 22
1.3.1. The associative account	. 22
1.3.2. The task-set switching account	24
1.4. Similarity in the IAT	25
1.4.1. The preference for familiarity account	. 28
1.4.2. The salience asymmetry account	28
1.5. The effect of salience asymmetries in the IAT	. 31
1.5.1. Salience asymmetries in the race IAT	31
1.5.2. Salience asymmetries in the British/Foreign IAT	. 33
1.6. The theoretical basis of salience asymmetries in the IAT	34
1.6.1. Visual search asymmetry as a predictor of salience asymmetry	. 35
1.6.2. Linguistic markedness as a predictor of salience asymmetry	37
1.6.3. Processing fluency as a predictor of salience asymmetry	. 39
1.7. The current research	. 40
Chapter 2. Determinants of salience asymmetry: visual search asymmetries ve	rsus
processing fluency	. 43
2.1. Experiments 1a-1b	. 46
2.1.1. Method	. 47
2.1.2. Results	50

2.2. Experiments 1c-1d 52
2.2.1. Method 52
2.2.2. Results and Discussion 53
2.3. Experiments 2a-2c 55
2.3.1. Method 55
2.3.2. Results
2.3.3. Discussion
2.4. Experiment 3
2.4.1. Method
2.4.2. Results
2.5. General Discussion
Chapter 3. Distinguishing between the effects of valence and salience in the IAT:
the development of a new procedure 71
3.1. Overview of research methods
3.2. Experiment 4a
3.2.1. Method
3.2.2. Results
3.3. Experiment 4b
3.3.1. Method
3.3.2. Results
3.4. Experiment 5a
3.4.1. Method
3.4.2. Results
3.4.3. Discussion
3.5. Experiment 5b 89
3.5.1. Method 89
3.5.2. Results and Discussion
3.6. General Discussion
Chapter 4. Familiarity in the Implicit Association Test
- 4.1. Experiment 6a 102
4.1.1. Method 103
4.1.2. Results

4.2. Experir	nent 6b 105
4.2.2	1. Method 105
4.2.2	2. Results and Discussion
4.3. Experir	nents 7a 109
4.3.7	1. Method
4.3.2	2. Results and Discussion
4.4. Experin	nent 7b 113
4.4.	1. Method
4.4.2	2. Results and Discussion
4.5. Experin	nents 8a 117
4.5.	1. Method
4.5.2	2. Results and Discussion 118
4.6. Experin	ment 8b 120
4.6.	1. Method 121
4.6.2	2. Results and Discussion 121
4.7. Genera	l Discussion 124
Chapter 5. Social	categories in the Implicit Association Test
Chapter 5. Social 5.1. Experin	categories in the Implicit Association Test
Chapter 5. Social 5.1. Experin 5.1.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132
Chapter 5. Social 5.1. Experin 5.1. 5.1.	categories in the Implicit Association Test130ments 9a-9b1311. Method1322. Results and Discussion133
Chapter 5. Social 5.1. Experin 5.1. 5.1. 5.1.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137
Chapter 5. Social 5.1. Experin 5.1. 5.1. 5.2. Experin 5.2.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138
Chapter 5. Social 5.1. Experin 5.1. 5.1. 5.2. Experin 5.2. 5.2.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138
Chapter 5. Social 5.1. Experin 5.1. 5.2. Experin 5.2. 5.2. 5.2. 5.2.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138 3. Discussion 143
Chapter 5. Social 5.1. Experin 5.1. 5.2. Experin 5.2. 5.2. 5.2. 5.2. 5.2.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138 3. Discussion 143 ments 10a-10b 144
Chapter 5. Social 5.1. Experin 5.1. 5.2. Experin 5.2. 5.2. 5.2. 5.2. 5.3. Experin 5.3.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138 3. Discussion 143 ments 10a-10b 144
Chapter 5. Social 5.1. Experin 5.1. 5.2. Experin 5.2. 5.2. 5.2. 5.3. Experin 5.3. 5.3.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138 3. Discussion 143 ments 10a-10b 144 1. Method 144 2. Results 144
Chapter 5. Social 5.1. Experin 5.1. 5.2. Experin 5.2. 5.2. 5.3. Experin 5.3. 5.3. 5.3.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138 3. Discussion 143 ments 10a-10b 144 1. Method 144 2. Results 144 3. Discussion 144 1. Method 144
Chapter 5. Social 5.1. Experin 5.1. 5.2. Experin 5.2. 5.2. 5.3. Experin 5.3. 5.3. 5.3. 5.3. 5.3. 5.3.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138 3. Discussion 143 ments 10a-10b 144 1. Method 144 2. Results 145 3. Discussion 148 ment 10c-10d 149
Chapter 5. Social 5.1. Experin 5.1. 5.2. Experin 5.2. 5.2. 5.3. Experin 5.3. 5.3. 5.3. 5.4. Experin 5.4.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138 3. Discussion 143 ments 10a-10b 144 1. Method 145 3. Discussion 148 ment 10c-10d 149 1. Method 150
Chapter 5. Social 5.1. Experin 5.1. 5.2. Experin 5.2. 5.2. 5.3. Experin 5.3. 5.3. 5.3. 5.4. Experin 5.4. 5.4.	categories in the Implicit Association Test 130 ments 9a-9b 131 1. Method 132 2. Results and Discussion 133 ments 9c-9d 137 1. Method 138 2. Results 138 3. Discussion 143 ments 10a-10b 144 1. Method 144 2. Results 145 3. Discussion 145 3. Discussion 148 ment 10c-10d 149 1. Method 150 2. Results and Discussion 150

Chapter 6. Mere acceptance and processing fluency in t	he Implicit Association
Test	
6.1. Experiment 11a	
6.1.1. Method	
6.1.2. Results and Discussion	
6.2. Experiment 11b	
6.2.1. Method	
6.2.2. Results	
6.2.3. Discussion	
6.3. Experiment 12a	
6.3.1. Method	
6.3.2. Results	
6.3.3. Discussion	
6.4. Experiment 12b	
6.4.1. Method	
6.4.2. Results	
6.4.3. Discussion	
6.5. General Discussion	
Chapter 7. General Discussion	
7.1. Processing fluency as a predictor of salience asymr	netry 189
7.2. Evaluating the assumptions of the split IAT	195
7.2.1. The assumption of equal salience	
7.2.2. The assumption of equal valence	
7.2.3. The assumption that the split IAT is a val	id measure of IAT effects
7.3. The relationship between salience asymmetries a	nd processing fluency in
the IAT	
7.3.1. How is salience determined in classificati	on? 203
7.3.2. How do salience asymmetries influence p	rocessing fluency? 206
7.3.3. When salience and fluency do not corresp	oond 208
7.4. How do salience asymmetries combine with valence	ce in the IAT? 210
7.5. Implications of salience asymmetries for other IA	T effects 214
7.5.1. The race IAT effect	

7.5.2. Category-level	IAT effects	
7.6. Concluding remarks		

References	. 222
Appendix A: Target stimuli for Experiments 12a-12b	. 236
Appendix B: Summary tables of statistical analysis	. 237
Appendix C: Raw reaction time data	259

List of Figures

Figure		Page
2.1	Elephant stimuli and cow stimuli.	49
2.2	Mean reaction times for the elephant/word IAT of Experiment 2a.	60
2.3	Mean reaction times for the cow/word IAT of Experiment 2b	61
2.4	Mean reaction times for the elephant/cow IAT of Experiment 3.	65
3.1	Mean reaction times for the flower/insect IAT of Experiment 4a.	77
3.2	Mean reaction times for the flower/insect split IAT of Experiment 4b.	81
3.3	Mean reaction times for the Go/Nogo IAT of Experiment 5a.	87
3.4	Mean reaction times for the Go/Nogo split IAT of Experiment 5b.	91
4.1	Mean reaction times for the insect/nonword IAT of Experiment 4a.	104
4.2	Mean reaction times for the insect/nonword split IAT of Experiment 6b.	107
4.3	Mean reaction times for the familiar/unfamiliar paintings IAT of Experiment 7a.	113
4.4	Mean reaction times for the familiar/unfamiliar paintings split IAT of Experiment 7b.	115
4.5	Mean reaction times for the even/odd number IAT of Experiment 8a.	119
4.6	Mean reaction times for the even/odd number split IAT of Experiment 8b.	122
5.1	Mean reaction times for the race IATs of Experiment 9a with faces and Experiment 9b with names.	135
5.2	Mean reaction times for the race split IATs of Experiment 9c with faces and Experiment 9d with names.	140
5.3	Mean reaction times for the age IATs of Experiments 10a with faces and 10b with names.	147
5.4	Mean reaction times for the age split IATs of Experiment 10c with faces and Experiment 10d with names.	152

Figure		Page
6.1	Examples of the stimuli used in the mere acceptance experiment of C.J. Mitchell (2004, Experiment 2).	162
6.2	Mean reaction times for the mere acceptance IAT of Experiment 11a.	165
6.3	Mean reaction times for the mere acceptance split IAT of Experiment 11b.	168
6.4	Mean reaction times for the Row/Column mere acceptance IAT of Experiment 12a.	175
6.5	Mean reaction times for the Row/Column mere acceptance split IAT of Experiment 12b.	180

Chapter 1. Introduction

Gordon Allport (1935) described attitudes as 'the most distinctive and indispensable concept in contemporary social psychology'. Unfortunately for social psychologists however, attitudes are not always easy to measure. Firstly, attitudes may not be stable across contexts, making it difficult to obtain reliable results. Secondly, self-report questionnaires may be inaccurate because people are reluctant to express certain attitudes, particularly those that are socially undesirable, such as prejudiced or deviant thoughts (Crowne & Marlowe, 1960; Paulhus, 1984). Thirdly, people may be unable to report their attitude if it is unconscious or inaccessible (Nisbett & Wilson, 1977; Wilson & Dunn, 2004). Sigmund Freud (1926/1958) recognized this problem when he developed the free association technique in an attempt to uncover thoughts that were believed to reside in the subconscious. In free association, people are asked to respond with anything which comes into their minds, often in response to a key word, such as 'mother'. This method assumes that memories are arranged in an associative network, and that by following the links in the network through the association of thoughts, the person will eventually come across the psychologically critical memory associated with the concept, 'mother'.

Nowadays, many attitude measures adopt a principle similar to that on which Freud's free association technique is based. They conceive of an attitude as an association between a target object and a valence concept (e.g. Fazio, 1995; Greenwald et al., 2002). For example, if one associates a target object such as 'flowers' with positive concepts such as 'fragrant', 'beautiful', and 'lovely', this indicates that one has a positive attitude towards flowers. In these attitude measures, emotional words serve as associative 'triggers' that influence speeded responding to target concepts. The most widely used tasks of this sort are the Implicit Association Test (Greenwald, McGhee, & Schwartz, 1998), the affective priming task (Fazio, Sanbonmatsu, Powell, & Kardes, 1986), and the Extrinsic Affective Simon Task (De Houwer, 2003). These indirect measures were developed to overcome the problem of social desirability associated with self-report measures. However, it appears that indirect measures are not immune to self-presentation strategies (for reviews see De Houwer, 2006; Gawronski, LaBel, & Peters, 2007). This thesis focuses on the most prominent of these measures, the Implicit Association Test, and examines its underlying mechanisms.

1.1. The Implicit Association Test

The Implicit Association Test (IAT) is the most popular indirect attitude measure in social psychology. It has been used to measure an extensive range of attitudes, including social prejudice (e.g. Rudman, Greenwald, Mellott, & Schwartz, 1999), consumer attitudes (Maison, Greenwald, & Bruin, 2004), self-esteem (Greenwald & Farnham, 2000), and animal phobias (Teachman, Gregg, & Woody, 2001). Although the IAT is only ten years old, there are over 400 papers that feature this measure. This flurry of IAT activity has been likened to the Copernican revolution spurred on by Galileo's findings in the sixteenth and seventeenth centuries (Kester, 2001).

The IAT rests on the assumption that similar concepts are associated in a semantic network, and that associated concepts can prime one another through spreading activation. Thus, if one has a positive attitude toward a particular target concept (e.g. flowers), then the target stimulus should more readily prime positive attributes (e.g. beautiful, lovely, fragrant) than negative attributes (e.g. disgusting, putrid, vile). The IAT gauges the attitude

to a target concept by measuring how quickly a target stimulus is classified with pleasant or unpleasant attributes in two speeded binary classification tasks. Stimuli are two target categories (e.g. names of flowers, and names of insects) that are classified with two other categories: a pleasant attribute category consisting of pleasant words, and an unpleasant attribute category consisting of unpleasant words. A procedural overview of the IAT is presented in Table 1.1. In the first task, the target categories are classified using separate keys (e.g. flower = left key; insect = right key). In the second task, the attributes are classified in the same way using the same keys as in the first task (e.g. pleasant = left key; unpleasant = right key). In the third task, the first and second tasks are combined, so that flowers and pleasant words share one response key (e.g. left key), and insects and unpleasant words share another response key (e.g. right). The fourth task is the same as the first task, except that the key assignments to the target categories are reversed. The fifth task is the same as the third task, except that the key assignments to the target categories (but not the attributes) are swapped so that flowers and unpleasant words now share one response key, and insects and pleasant words share another response key. The order of the two combined classification tasks (Tasks 3 and 5) is counterbalanced, and so is the key assignment of targets and attributes.

When people perform the flower/insect IAT, the conventional finding is that when flowers are paired with pleasant attributes in the combined classification task, responding is faster and more accurate than on the combined classification task in which flowers are paired with unpleasant attributes (Greenwald et al., 1998, Experiment 1). This is known as an IAT effect for flowers. This pattern of results may indicate that either pleasant attributes are more easily associated with flowers than with insects, or that unpleasant attributes are more easily associated with insects than with flowers, or both. On the basis of this finding, a positive attitude toward flowers relative to insects is inferred. Differences in the size of IAT effects are thought to reflect differences in attitudinal strength, with a larger IAT effect indicating a more extreme preference for one category over the other. This interpretation is supported by the finding that young boys show a smaller IAT effect for flowers than young girls (Baron & Banaji, 2006), and entomologists also show a weaker preference for flowers compared to a control group (Citrin & Greenwald, 1998; cited in Lane, Banaji, Nosek, & Greenwald, 2007).

Table 1.1

Classification task	Left key assignment	Right key assignment
Task 1	Flower	Insect
Task 2	Pleasant	Unpleasant
Task 3	Flower	Insect
	Pleasant	Unpleasant
Task 4	Insect	Flower
Task 5	Insect	Flower
	Pleasant	Unpleasant

Overview of the Implicit Association Test

Note. The categories represented are those in the flower/insect IAT. For half the participants, the category assignment is reversed in Task 1, and Tasks 4 and 5 precede Tasks 2 and 3.

1.2. Validity of the IAT

1.2.1. Correspondence of the IAT with behavioral measures

The IAT is popular because it is a very robust measure. It produces large effect sizes, and even people who know that it assesses attitudes still show a reliable IAT effect. This has been taken to demonstrate its resistance to self-presentation concerns. Despite its popularity however, the validity of the IAT as a measure of attitudes is still under question (e.g. Arkes & Tetlock, 2004; Blanton, Jaccard, Gonzales, & Christie, 2006; Gawronski, et al. 2007; Rothermund & Wentura, 2001, 2004; Kinoshita & Peek-O'Leary, 2005, 2006; De Houwer, Beckers, & Moors, 2007). One particular IAT effect that has attracted much attention is the race IAT effect, in which people find it easier to classify white (Caucasian) people with pleasant attributes, and black (African American) people with unpleasant attributes. This effect has been attributed to an implicit prejudice against blacks, but some have criticized this conclusion. For example, Arkes and Tetlock (2004) have suggested that the race IAT effect may reflect cultural stereotype norms rather than personal attitude, or that it is caused by other negative emotions or cognitions that are not necessarily prejudiced.

In response to the concern that the race IAT may not measure racial prejudice, some researchers have investigated the predictive validity of the measure. McConnell and Leibold (2001) demonstrated that race IAT performance predicted behavioral measures of racial discrimination. Larger IAT effects for white were correlated with a higher level of participant friendliness and comfort in the presence of a white experimenter than a black experimenter. This behavior was characterized by, among other things, longer speaking

time, more smiling, and fewer speech errors. However, Kinoshita and Peek-O'Leary (2005) argue that such behavior may reflect greater unfamiliarity with black people, rather than prejudice. More difficult to reconcile with the unfamiliarity argument is Hugenberg and Bodenhausen's (2003) finding that prowhite bias in the IAT was associated with a lowered threshold for detecting hostility on black, but not white, faces. Similarly, people who exhibited larger race IAT effects were more likely to judge racially ambiguous faces with hostile expressions as black (Hugenberg & Bodenhausen, 2004).

In other areas, IAT scores and behavior do not always correspond. Karpinski and Hilton (2001) showed that in an IAT where the target categories were apples versus candy bars, people found it easier to classify apples with pleasant attributes, and candy bars with unpleasant attributes. However, when given a choice over an apple or a candy bar, they were more likely to select the candy bar. This supports the idea that performance on the IAT may sometimes reflect environmental associations – that is, apples are known to be 'good' for the health, whereas candy bars are known to be 'bad' for the health – rather than any personal attitude toward the targets. Other evidence that supports the argument against the IAT as an attitude measure comes from the neuroimaging literature.

1.2.2. Correspondence of the IAT with neuroimaging measures

Another way to measure attitude is by using physiological methods, such as neuroimaging (e.g. Breiter et al., 1996). One brain structure that has been implicated in emotional processing is the amygdala (Davis, 1997; LeDoux, 1996). Based on this, Phelps and colleagues (Phelps et al., 2000; Phelps, Cannistraci, & Cunningham, 2003) hypothesized that amygdala activation is correlated with performance on the race IAT. However, studies comparing race IAT performance and amygdala activation have shown an inconsistent correspondence between the two measures (Phelps et al., 2000; Phelps et al., 2003). Phelps et al. (2000, Experiment 1) used an fMRI to show that the amygdala of white people was more activated when they viewed novel pictures of black people than when they viewed novel pictures of white people, and this difference was correlated with IAT effects. However, when people viewed pictures of positively evaluated famous black and white faces (e.g. Michael Jordan, John F. Kennedy), amygdala activity was the same for both black and white faces, even though there was still an IAT effect for white. From this we see that IAT effects do not always correspond with amygdala activation.

In other circumstances, performance on the IAT remains the same even when there are differences in amygdala activation. This is illustrated by a study in which Phelps et al. (2003) compared the race IAT performance of patients with bilateral amygdala damage against those of normal controls. Previous research has shown that amygdala damage impairs the ability to judge whether an individual appears approachable or trustworthy (Adolphs, 1998). Compared to normal controls, patients with amygdala damage are more likely to rate pictures of individuals as being trustworthy and approachable, even when normal controls rate the same pictures as being untrustworthy and unapproachable. Because damage to the amygdala appears to impair social evaluation, if the IAT does indeed measure attitude, then we would expect patients with amygdala damage to show a smaller race IAT effect than normal controls. However, the two groups actually produced equal effect sizes on the race IAT. This result further supports the idea that the IAT does not always reflect affective evaluation (as indexed by amygdala activation). Instead, it may be that the IAT assesses semantic associations between targets and attributes. Consistent with this idea, studies have shown that IAT performance is susceptible to contextual manipulations of target and attribute associations.

1.2.3. Contextual effects in the IAT

Many studies have shown that IAT performance is subject to contextual influences, indicating that it does not always measure long-term associations between target and valence concepts in memory (Blair, Ma, & Lenton, 2001; Dasgupta & Greenwald, 2001; Gawronski & Bodenhausen, 2005; Karpinski & Hilton, 2001; Lowery, Hardin, & Sinclair, 2001; J.P. Mitchell, Nosek, & Banaji, 2003; Wittenbrink, Judd, & B. Park, 2001; see Blair, 2002 for a review). For example, Blair et al. (2001) demonstrated that performance on a gender IAT can be moderated by experimentally-induced mental imagery. In their study, participants were asked to imagine a woman who was weak (stereotype condition), a woman who was strong (counterstereotype condition), or a gender-neutral topic (e.g. holiday). Following this, they completed a gender IAT in which they classified male and female names with 'strong' (e.g. durable) and 'weak' (delicate) attributes. Those in the counterstereotypical condition showed less stereotyping on the IAT (i.e., there was less compatibility between male names and strong attributes, and between female names and weak attributes) compared with the stereotypical and control condition. These results indicate that varying target and attribute associations through activating specific representations can alter responding in the IAT.

Contextual manipulations in the IAT are not necessarily accompanied by corresponding changes in explicit attitude measures, which is potentially problematic for the validity of the IAT as a measure of attitude. For example, Dasgupta and Greenwald (2001) were able to influence responses on the race IAT and the age IAT by first presenting participants with pictures of either admired black and disliked white individuals (e.g., Michael Jordan, Charles Manson) or admired white and disliked black individuals (e.g., Tom Hanks, Mike Tyson). Participants then completed the race IAT, and self-report attitude measures of the relevant target groups. Participants who were exposed to admired black exemplars and disliked white exemplars had smaller IAT effects for white, compared to when they were exposed to admired white and disliked black exemplars. However, the explicit evaluations of the target groups generally did not differ between the two conditions. Similar results were obtained on an age IAT where the target categories were old and young people. These results suggest that the IAT may have been measuring short-term associative changes that are independent of the more stable object-valence associations presumed to reside in long-term memory.

It may also be the case that the IAT does not actually measures attitude, stable or otherwise. This idea is supported by research from De Houwer et al. (2007) showing that attitudes toward novel stimuli can be faked on the IAT. In their study, participants were asked to imagine that a researcher had given them information about the fictitious social groups 'Niffites' and 'Luupites'. In one condition, they were to imagine that the researcher had told them that Niffites were good and Luupites were bad. In the other condition, they were to imagine that the researcher had told them that Luupites were good and Niffites were bad. Participants were instructed that, although the information they were given was false, the researcher wanted them to respond in a manner that conformed to the expectations of the researcher (i.e. to show that the participants were influenced by the description). Half the participants in each condition, and the other half were asked to behave as the researcher expected (consistent faking condition), and the other half were asked to behave in the *opposite* way to what the researcher expected (inconsistent faking condition). IAT performance was in line with the faking instructions, such that participants in the consistent faking condition showed an IAT effect for the group that they imagined being told was good, whereas those in the inconsistent faking condition showed the reverse effect. These results show that performance on the IAT can be influenced by task demands, rather than actual attitude.

In summary, it appears that the IAT does not always correspond with behavioral, neuropsychological, and explicit measures of attitude. The question then becomes, what else is the IAT measuring? Some researchers have shown that other than affective and semantic associations, the IAT measures other types of similarity between the target and attribute categories. These types of similarity include perceptual similarity (Mierke & Klauer, 2003; De Houwer, Geldof, & De Bruycker, 2005), and similarity based on selective attention (Rothermund & Wentura, 2001, 2004). In the following sections, I first describe how associations are thought to operate in the IAT, and then compare these accounts to theories claiming that the IAT can measure other types of similarity between categories.

1.3. Associations in the IAT

1.3.1. The associative account

According to the account proposed by Greenwald et al. (2000), the IAT measures relative associative strengths between target and attribute categories. In the case of an IAT with valence attributes, these associations are thought to be affective in nature. In other IATs, such as those which measure gender stereotypes using the attribute categories of 'weak' and 'strong', or 'arts' and 'science', these associations are assumed to be semantic in nature. It is possible that the associations measured by an IAT with valence attributes may also be semantic. Therefore, unless specified, the term 'association' and its variants will be used hereafter to indicate associations that may be affective and/or semantic in nature. Greenwald, Nosek, Banaji, and Klauer (2005) take a 'theory-uncommitted' view of the concept of association in the IAT. They suggest that the term 'association' may describe, among other things, similarity between two concepts. However, if the IAT measures associations based on similarity that is neither affective nor semantic, this means that the IAT effect is not a reliable measure of attitudes, as it may be contaminated by variables unrelated to attitude. Therefore, in order to use the IAT as an attitude measure, it is necessary to clarify the circumstances under which IAT performance is based on affective and/or semantic similarity, and when it may be based on other types of similarity.

To identify the type of similarity that influences IAT effects, one should first understand how associations are measured in the IAT. However, the associative account proposed by Greenwald et al. (2000) does not clarify this issue. Although Greenwald et al. (1998) compared the IAT to an evaluative priming measure of attitude, they did not specify a mechanism underlying the operation of the IAT. It cannot be assumed that the IAT operates by a similar mechanism to priming measures. In evaluative priming, an attitude concept (e.g. flowers) is thought to automatically activate its corresponding evaluation (positive). Affective priming is based on the idea that it is easier to evaluate a positive word (e.g. lovely) as being positive when it is preceded by an affectively congruent prime (e.g. flower), than when it is preceded by an affectively incongruent prime (e.g. insect). This is because the positive attribute node receives activation from the closely associated flower node sooner than from the distally associated insect node. Mierke and Klauer (2001) point out that the spreading activation account of priming cannot adequately explain IAT performance. This is because in the IAT, participants respond to both target categories (e.g. flower and insect) and both pleasant and unpleasant attribute categories in the same task (Tasks 3 and 5 in Table 1.1). This means that flowers should activate pleasant attributes equally easily in both tasks, and unpleasant attributes with equally difficulty in both tasks (and vice versa for insects). However, the general finding is that both targets and attributes are more difficult to classify in only one of the tasks, when flowers share a key with unpleasant attributes, and insects share a key with pleasant attributes. Because the only variable that differs between the tasks is the target and attribute key assignment, any theory of the IAT must explain how this factor affects IAT performance. One theory that explains the role of key assignment in the IAT is the task-set switching account of Mierke and Klauer (2001, 2003; Klauer & Mierke; 2005).

1.3.2. The task-set switching account

The task-set switching account claims that when there is an overlap between the features of targets and attributes, participants may classify the categories on the basis of this shared feature, rather than on the nominal features of each category. Thus, when the positive target and positive attribute categories share a key, and when the negative target and negative attribute categories share another key, the combined classification task can be reduced to a single valence categorization task, which leads to faster responding. This strategy cannot be adopted when a positive and a negative category share a key, and thus participants must classify all four categories separately. The difficulty of the task is thereby

increased, requiring executive control processes to identify and switch to the appropriate task-set. This results in a task-switching cost that slows down reaction times.

In an IAT with flowers and insects as the target categories, Mierke and Klauer (2001) assessed the task-switching cost by measuring how reaction times varied depending on the preceding trial. Reaction times were longer when the previous trial was of a different type to the current trial (e.g. a target classification trial preceded by an attribute classification trial), compared to when the previous trial was the same (e.g. a target classification trial preceded by a target classification trial). As predicted by their account, this task-switching cost was greater when the same key was assigned to both a positive and negative category (e.g. flowers and negative attributes shared the same key, and insects and positive attributes shared the same key). This implies that it took more effort to make a mental switch between classifying target and attribute categories when they were affectively incongruent compared to when they were affectively congruent. Thus, rather than classifying the targets and attributes based on their nominal, category-specific features, this result suggests that participants perform the IAT by recoding target and attribute categories with respect to their shared valence features.

1.4. Similarity in the IAT

Mierke and Klauer (2003, Experiment 1) have shown that recoding in the IAT can be based on factors other than affective or semantic similarity, which challenges the validity of the IAT as an attitude measure. Rather than using standard target and attribute categories in their study, stimuli were classified on the basis of size ('small' vs. 'large') or color ('red' vs. 'blue'). Objects that were to be classified in terms of size were neither red nor blue. However, all red objects were small, and all blue objects were large. Thus red objects were perceptually similar to the small objects, and blue objects were perceptually similar to the large objects. This similarity influenced IAT performance, with responses being faster when red objects were classified with small objects, and blue objects were classified with large objects, compared to when the pairings were reversed. There was also a larger task-switching cost in the condition in which the categories sharing the same key were perceptually incongruent, compared to when they were perceptually congruent. This suggests that participants classified all four categories in terms of size or color in the former condition, but that they classified the categories primarily in terms of the similar feature (size) in the latter condition. Thus, it appears that recoding strategies in the IAT are not limited to the dimension of valence, but can also include other features, such as perceptual similarity.

The type of similarity that is measured by the IAT seems to depend on the particular dimension that is focused upon. This principle is demonstrated in a study by De Houwer et al. (2005). Using an IAT in which the categories were coins versus snakes, and pizzas versus rivers, De Houwer et al. (2005, Experiment 2) manipulated whether participants' attention was first drawn to the perceptual similarity between the categories (i.e. coins and pizzas are both round, and snakes and rivers are both long and winding), or the functional similarity between the categories (i.e. snakes and pizzas are both edible, and coins and rivers are both inedible). Participants who first judged whether the category items were round or winding were faster to respond when the perceptually similar stimuli shared the same key (ie. coins shared a key with pizzas, and snakes shared a key with rivers). Conversely, participants who first judged whether the category items were edible or inedible tended to find it easier to classify the two edible items together (snakes and pizzas)

and the two inedible items together. Although this last finding was not reliable, the judgment type that participants were primed with (i.e. round/winding vs. edible/inedible) did show a significant interaction with the direction of the IAT effect. This result demonstrates that the IAT is sensitive to perceptual and functional similarity. Moreover, it appears that similarity in the IAT is flexible, as the same IAT can produce different results based on different types of similarity. From this, De Houwer et al. (2005) concluded that the IAT can measure various dimensions of similarity in addition to the evaluative basis of similarity proposed by Greenwald and his colleagues.

Another source of similarity in the IAT may be based on selective attention. Rothermund and Wentura (2001, 2004) suggest that some targets and attributes may command more attention than others in the IAT, and those that are similar in their ability to command attention are compatible in the IAT. Rothermund and Wentura refer to this ability to command attention as 'salience'. One IAT effect claimed by Rothermund and Wentura to be affected by salience differences is the insect/nonword IAT effect by Brendl, Markman and Messner (2001, Experiment 2). In this IAT where insects and nonwords are the target categories, people find it easier to classify insects with pleasant attributes, and nonwords with unpleasant attributes. This result is problematic for the associative account of the IAT for two reasons. Firstly, in Brendl et al.'s (2001) study, participants' self-reports showed that insects and their associates were considered to be more negative than nonwords and their associates. This indicated that insects were more strongly associated with negative attributes than were nonwords. Secondly, given that nonwords are novel, they should have no prior associations. Therefore, the associative account of the IAT would predict that insects should be compatible with unpleasant attributes, and nonwords should be equally compatible with both pleasant and unpleasant words, resulting in an IAT effect for

27

nonwords over insects. Two accounts have been proposed to explain the insect/nonword IAT effect, the preference for familiarity account, and the salience asymmetry account.

1.4.1. The preference for familiarity account

Brendl et al. (2001) suggested that the IAT effect for insects over nonwords reflects a preference for familiar items. This claim is based on evidence from the mere exposure effect showing that repeated presentation of a neutral stimulus increases liking for that stimulus (Zajonc, 1968, 2001; Kunst-Wilson & Zajonc, 1980; Mandler, Nakamura, & Van Zandt, 1987; Bornstein & D'Agostino, 1992). If Brendl et al.'s insect/nonword IAT effect is due to a preference for the familiar category of insects, then this is consistent with the idea that the IAT is a measure of affective similarity between target and attribute items. The critical difference between this account and the associative account, however, is that when familiarity is preferred, the affect engendered by the target items is not based on associative links with valence concepts. As such, preference for a target category would not reflect the interindividual differences that are assumed to underlie attitudes such as social prejudice. Instead, in this case, the experience of affect seems to be generated online, by differences in familiarity.

1.4.2. The salience asymmetry account

Rothermund and Wentura (2001, 2004) proposed a non-affective explanation of Brendl et al.'s (2001) insect/nonword IAT effect. They explained the effect in terms of salience asymmetries. Their argument is that IAT effects may be at least partly driven by differences in the salience of stimulus categories. They suggest that items within the target and attribute categories differ in their relative salience, so that, in the example above, nonwords may be more salient than insects, and unpleasant attributes may be more salient than pleasant attributes. This may lead people to simplify the IAT combined classification task by looking only for the more salient category of the respective pairs, i.e. the nonword and unpleasant categories. Thus when the more salient categories share one response key, and the less salient categories share another response key, participants can classify the items in terms of their relative salience. This allows them to press one key when a salient item is presented, and another key when a less salient item is presented. This strategy cannot be used when a more salient target category shares a key with a less attribute salient category (and vice versa), and therefore participants must resort to classifying all four categories separately. In this way, an IAT effect may be caused by compatibility between target and attribute categories that is based on salience rather and valence. This is similar to the recoding principle invoked in the task-set switching account of Mierke and Klauer (2001, 2003; Klauer & Mierke; 2005), however unlike that account, salience compatibility does not depend on shared similarity in features, but rather, shared similarity in salience.

In support of the salience asymmetry account of IAT effects, Rothermund and Wentura (2001, 2004) showed that IAT effects can be reversed by selectively increasing the salience of target and attribute categories. They demonstrated this effect in an age IAT with young and old target categories, in which the target stimuli were old names such as 'Walter' and young names such as 'Patrick'. In the standard age IAT, people find it easier to classify young stimuli with pleasant attributes, and old stimuli with unpleasant attributes (Nosek, Banaji, & Greenwald, 2002). In Rothermund and Wentura's (2001) Experiment 2, salience was manipulated using a Go/Nogo task performed prior to the IAT. In this task participants were required to respond to only one of the categories within the target and attribute dimensions (the Go categories), whilst making no response to the remaining target or attribute category (the Nogo categories). For example, in one condition of an old/young IAT, the Go categories were young names and unpleasant attributes, and the Nogo categories were old names and pleasant attributes. Rothermund and Wentura (2004) claimed that this encourages people to search for items from the Go categories (in this case, young names and unpleasant attributes) which establishes an attentional focus toward those categories, and makes them more salient than the Nogo categories. The targets and attributes were then classified in the usual IAT. It was found that any combination of target and attribute (e.g. young and unpleasant) could be easily classified together in the IAT if they were both made the focus (i.e. were Go items) in a previous Go/Nogo task. Thus, it appears that salience asymmetries in the IAT can override target-attribute associations in some circumstances. As such, salience asymmetries represent the greatest threat to the construct validity of the IAT as a measure of attitudes. It is for this reason that the salience asymmetry account will be the focus of this thesis.

Salience asymmetries may also account for other effects in the IAT, such as the contextual effects found by Blair et al. (2001). In Blair et al.'s study, when participants visualized a strong woman prior to performing a gender stereotype IAT, this would have primed them with the concepts of 'woman' and 'strong'. Thus, when performing the IAT immediately afterwards, the target category 'woman' and the attribute category 'strong' would likely to have been more salient than the categories 'man' and 'weak'. This may make women more compatible with strong attributes, and men more compatible with weak attributes, based on salience asymmetries.

1.5. The effect of salience asymmetries in the IAT

Familiarity and valence are argued to be the major sources of salience asymmetries. Differences in familiarity characterize many target pairings in the IAT, particularly those that compare attitudes to ingroups and outgroups, because ingroups are usually more familiar than outgroups. Both Rothermund and Wentura (2001, 2004), and Kinoshita and Peek-O'Leary (2005, 2006) have argued that the familiar target category is compatible with positive attributes, and the unfamiliar target category is compatible with negative attributes, on the basis of salience asymmetries. Thus, if one of the target categories is more salient than the other, any IAT effect for the ingroup may be influenced by a non-attitudinal factor that may inflate the IAT effect.

1.5.1. Salience asymmetries in the race IAT

One particular IAT effect that is claimed to be affected by salience asymmetries is the race IAT effect (Kinoshita & Peek-O'Leary, 2005; Rothermund & Wentura, 2004). More specifically, Kinoshita and Peek-O'Leary claim that the race IAT effect is entirely due to a salience compatibility between the familiar white category and positive attributes, and between the unfamiliar black category and negative attributes. Thus, they suggest that the race IAT effect does not indicate a preference for white over black.

In reaction to criticism that the race IAT effect reflects familiarity with the target categories rather than evaluative associations, many experimenters have sought to demonstrate that IAT effects still exist when familiarity is controlled for. For instance, Dasgupta, McGhee, Greenwald, and Banaji (2000) tried to statistically equate familiarity

between black and white by regressing IAT effects onto a differential measure of familiarity (measuring the speed with which black and white names were discriminated from pseudonames). At the point that the familiarity advantage was zero, the regression equation yielded a positive intercept, indicating a significant IAT effect when familiarity was equated between the black and white names. From this, Dasgupta et al. concluded that familiarity did not contribute to the race IAT effect. However, Rothermund and Wentura (2004) and Kinoshita and Peek-O'Leary (2005) note that there was also a positive slope to the regression equation, indicating that greater familiarity with white names was associated with a larger IAT effect. This suggests that familiarity did play a role in modulating the race IAT effect. However, because familiarity was not independently manipulated in this study, one cannot say whether familiarity had a causal influence on IAT performance.

Other studies have tried to equate the familiarity of race categories by using the faces of unfamiliar people as the target exemplars (Dasgupta et al., 2000), or by matching the familiarity of black and white exemplar names (Dasgupta et al., 2000; Ottaway, Hayden, & Oakes, 2001; see Rudman et al., 1999 for related findings involving Christian/Jewish names, Old/Young names, and American/Soviet names). Because race IAT effects remain under these circumstances, these researchers claim that familiarity does not influence the race IAT effect. However, both Rothermund and Wentura (2004) and Kinoshita and Peek-O'Leary (2005) point out that these manipulations do nothing to change differences in familiarity between the target categories as a whole. That is, nonblack people find white people as a group to be more familiar than black people as a group. They argue that it is the familiarity of the target categories, rather than the familiarity of the individual exemplars, that influences the race IAT effect.
1.5.2. Salience asymmetries in the British/Foreign IAT

Rothermund and Wentura (2001, 2004) posit that differences in target category familiarity may also have mediated the British/Foreign IAT effect demonstrated by De Houwer (2001). In this experiment, British participants classified the target categories British and Foreign with positive and negative attributes. Each target category consisted of the names of six public figures, three of whom were positively evaluated (e.g. British – Princess Diana, Foreign – Mahatma Ghandi), and three of whom were negatively evaluated (e.g. British – Margaret Thatcher, Foreign – Adolf Hitler). Participants were faster to classify British with positive attributes, and Foreign with negative attributes. De Houwer interpreted this result as showing that in the IAT, performance is based on the valence of the category, rather than of the individual exemplars that comprise a category.

Alternatively, Rothermund and Wentura (2004) suggest that De Houwer's (2001) category level effect may have been caused by salience asymmetries based upon the British category being more familiar to British participants than the Foreign category. That is, the British category was compatible with positive attributes because it was similarly salient (compared with the Foreign and negative categories). This idea is supported by an experiment in which Rothermund and Wentura (2004, Experiment 4) showed that a familiar category consisting of negative stimuli may still be compatible with positive attributes in the IAT. In their experiment, participants classified the target categories of known and unknown people with 'good' and 'bad' attributes. In one condition, the 'known' category consisted of people that were negatively evaluated. Across both conditions, known people were compatible with good attributes, and unknown people with

bad attributes. Furthermore, the positive and negative known categories yielded IAT effects of similar magnitude (135ms for known-positive, and 170ms for known-negative). Rothermund and Wentura interpreted their finding as supporting the salience asymmetry account, because the familiar (known) target category was compatible with positive (good) attributes, regardless of whether the target exemplars were positive or negative. Note, however, that this explanation is also consistent with Brendl et al.'s (2001) claim that familiar items are preferred in the IAT. Therefore, Rothermund and Wentura's knownunknown IAT effect does not provide unequivocal evidence that familiar items are differentially salient to unfamiliar items in the IAT.

1.6. The theoretical bases of salience asymmetries in the IAT

Although Kinoshita and Peek-O'Leary (2005, 2006) agree with Rothermund and Wentura (2001, 2004) that familiar is compatible with positive, and unfamiliar is compatible with negative, they disagree with respect to which ends of the valence and familiarity dimensions they consider to be more salient in the IAT. Rothermund and Wentura consider that unfamiliar and negative stimuli are more salient than familiar and positive stimuli, whereas Kinoshita and Peek-O'Leary argue that positive and familiar are the more salient categories in the IAT. This discrepancy in perspective is caused by the different theoretical positions of the two groups. Rothermund and Wentura generally conceive of salience asymmetries in the IAT in terms of visual search asymmetries. In the visual search literature, unfamiliar and negative stimuli are considered to be more salient. In contrast, Kinoshita and Peek-O'Leary adopt a definition of salience asymmetries from the psycholinguistic literature, which considers the positive and familiar dimension to be more salient. Each of these potential predictors of salience asymmetry will be discussed in more detail in the next section.

It should be noted that Rothermund and Wentura's (2001, 2004) salience asymmetry account does not hinge on whether familiar/unfamiliar or positive/negative are the more salient categories in the IAT. However, the relative salience within each target and attribute pairing is an important distinction when interpreting particular IAT effects. For instance, if a target stimulus is salient for a particular reason, then Rothermund and Wentura would predict that it should be compatible with negative attributes, whereas Kinoshita and Peek-O'Leary (2005, 2006) would predict the reverse effect, that it should be compatible with positive attributes. Therefore understanding which types of categories are more salient will allow us to interpret whether a given IAT effect is consistent or inconsistent with a salience asymmetry effect. The following section will evaluate Rothermund and Wentura's claim that visual search asymmetries are a predictor of salience asymmetries in the IAT. This will be contrasted with Kinoshita and Peek-O'Leary's assertion that salience asymmetries in the IAT follow linguistic principles. A third potential predictor of salience asymmetries in the IAT will also be suggested.

1.6.1. Visual search asymmetry as a predictor of salience asymmetry

As mentioned previously, Rothermund and Wentura (2001, 2004) consider that unfamiliar and negative items are more salient in the IAT than familiar and positive items respectively. These assumptions are based on independent evidence of visual search asymmetries between stimuli that differ in familiarity and valence. Search asymmetries occur when it is easier to find stimulus A among many examples of stimulus B than it is to

35

find stimulus B among many examples of stimulus A (Treisman & Souther, 1985). A number of visual search studies have shown that unfamiliar targets (e.g. a mirror-reversed 'N') are more readily detected among familiar distractors (e.g. 'N') than vice versa (e.g. Strayer & Johnston, 2000; Wang, Cavanagh, & Green, 1994). Similarly, negative items (e.g. angry faces) are detected more quickly among an array of positive items (e.g. happy faces), than vice versa (Fox, Russo, Bowles, Pichler, & Dutton, 2000; Ohman, Flykt, & Esteves, 2001a). This search asymmetry has been taken as evidence that negative and unfamiliar stimuli capture attention more effectively (Tipples, Young, Quinlan, Broks, & Ellis, 2002), or hold attention to a greater extent (Fox, Russo, & Dutton, 2002), than positive and familiar items.

In support of the relationship between visual search asymmetries and salience asymmetries, Rothermund and Wentura (2004, Experiments 1a, 1b, 1d, and 1e) demonstrated that categories that were more quickly detected on a visual search task were also compatible in the IAT. For example, first they showed that one old name was more quickly detected among three young names than vice versa, and that one multi-colored string was detected more quickly among three single-colored strings than vice versa. This was interpreted to suggest that old names and multi-colored stings commanded greater attentional resources, and thus were more salient than young names and single-colored strings respectively. Based on this result, Rothermund and Wentura predicted that in an IAT in which young and old names were classified with single- and multi-colored strings, the more salient categories (old names and multi-color strings) would be compatible in the IAT. This hypothesis was supported: participants performed 37ms faster in the condition in which old names were classified with multi-colored strings and young names with singlecolored strings, than when the pairings were reversed. It seems reasonable to assume that old names are not affectively or semantically related to multi-colored strings, nor are young names similarly related to single-colored strings. Therefore, this IAT effect is likely to have resulted from salience asymmetries within the target and attribute pairings.

In the same set of experiments by Rothermund and Wentura (2004, Experiments 1a, 1b, 1d, and 1e), words versus nonwords were also placed in a visual search task, as were good versus bad words. It was found that nonwords and bad words were more quickly detected amongst words and good words respectively, than vice versa. Old and young names were then classified in an IAT with either word versus nonword attributes, or good versus bad attributes. The categories that were detected more quickly in their respective visual search tasks (i.e. nonwords and bad words) were found to be compatible with old names in the IAT. From this, Rothermund and Wentura (2004) concluded that the categories that are detected more quickly in visual search (old names, multi-colored strings, nonwords and bad words) are also more salient in the IAT. Thus, they recommended that the visual search task could be used to directly assess salience asymmetries that may contribute to IAT effects.

1.6.2. Linguistic markedness as a predictor of salience asymmetry

Kinoshita and Peek-O'Leary's (2005, 2006) conceptualisation of salience asymmetries is based on the principles of psycholinguistics, where it is the positive and the familiar that is deemed to be more salient (e.g. Greenberg, 1966; Clark, 1973). This principle is apparent in the concept of linguistic markedness (see Greenberg, 1966), which describes how language is often expressed has having an unmarked aspect (that is, the basic, canonical form), and a marked aspect (the non-basic, less natural form). Positivity is considered to be unmarked, because people will use the default form, "How good is it?" when evaluating whether a stimulus is good or bad, rather than "How bad is it?" which implies that a stimulus is already judged to be bad (e.g. Lyons, 1977). Similarly, positive words are often transformed into negative words by adding a prefix, (e.g. unhappy, unsafe, dissatisfied), whereas the reverse is not true of negative words. These phenomena suggest that stimuli are classified using the positive dimension as the referent category; items are evaluated as to whether they are good, or not good, rather than whether they are bad, or not bad. The idea that positivity is the default referent category implies that positive words should be more readily processed than negative words. Consistent with this idea, people respond to positive words slightly more quickly than negative words (in lexical decision and naming tasks) when the stimuli are matched on frequency and word length (Estes & Adelman, in press; see Unkelbach, Fiedler, Bayer, Stegmüller & Danner, in press, for a similar result involving German words).

Likewise, familiarity is more salient because old is unmarked with respect to new (Clarke, 1973). When inquiring about the age of an item, the default form of the question is 'How *old* is it?' rather than 'How *new* is it?' Familiar words are also responded to more quickly than unfamiliar words on a lexical decision task, (Balota & Chumbley, 1984; Whaley, 1978). This shows that familiar words are more readily accessible than unfamiliar words, suggesting that the familiar is the default, unmarked form of the familiarity dimension.

1.6.3. Processing fluency as a predictor of salience asymmetry

Although Rothermund and Wentura (2001, 2004) and Kinoshita and Peek-O'Leary (2005, 2006) both advocate the salience asymmetry account of the IAT, they argue that these salience asymmetries correspond with two different phenomena: visual search asymmetries and linguistic markedness. One aspect common to both these phenomena is differences in processing fluency. Processing fluency can be defined as the speed with which a stimulus is processed. The categories that are considered to be more salient by Rothermund and Wentura, that is, unfamiliar and negative items, tend to be less fluently processed on lexical decision and naming tasks. The categories that are considered to be more salient by Kinoshita and Peek-O'Leary, that is, familiar and positive, tend to be more fluently processed on the same tasks. Therefore, both of these accounts are consistent with the idea that in the IAT, the more fluently processed target and attribute categories are more easily categorized together, and the less fluently processed target and attribute categories are more easily categorized together. Where the two accounts differ is in the categories that are considered to be the focus of attention. According to Rothermund and Wentura, people tend to focus on the less fluently processed categories in the IAT, whereas Kinoshita and Peek-O'Leary argue that people tend to focus on the more fluently processed categories.

A critical step towards identifying and managing salience asymmetry effects is to clarify more precisely the nature and potential sources of these effects. For example, if salience asymmetries in the IAT correspond with visual search asymmetries, the visual search task would be an appropriate measure of salience asymmetries in the IAT. However, if differences in target familiarity are responsible for salience asymmetry effects, then it would be much more difficult to control for salience asymmetries by equating category familiarity. Alternatively, processing fluency may be responsible for salience asymmetry effects. If so, then tasks that measure this aspect may be used as a diagnostic tool to detect salience asymmetries in the IAT. Moreover, if salience asymmetries can be controlled for by equating processing fluency between the categories, then this method may be used to decontaminate the IAT.

1.7. The current research

Although the IAT has the potential to be a very useful tool for measuring attitudes, it also yields results that do not appear to have an associative basis. The current research aims to resolve this problem by examining the nature of these non-associative contaminants. In doing so, it will address four main issues involving salience asymmetries in the IAT.

The first issue concerns the variables that predict salience asymmetries in the IAT. In Chapter 2, Rothermund and Wentura's (2001, 2004) account of salience asymmetries in the IAT in terms of visual search asymmetry will be compared to a processing fluency account. Experiments will be presented in which search asymmetry and fluency are manipulated independently, and the dimension along which categories are compatible in the IAT will be examined.

The second issue concerns whether the separate contributions of valence and salience in the IAT can be measured. Chapter 3 trials a modified version of the IAT and tests whether it is able to discriminate between valence differences and salience asymmetries in the IAT.

The third issue involves the role of familiarity in the IAT. In Chapter 4, experiments will be presented in which the familiarity of the target categories is experimentally manipulated to examine whether familiarity produces valence-based effects in the IAT (as suggested by Brendl et al., 2001), or whether it produces effects based on salience asymmetries (as suggested by the salience asymmetry account).

The fourth issue concerns the role of salience asymmetries in IATs that assess attitudes toward social categories, such as the race IAT and the age IAT. Salience asymmetries are assumed to inflate IAT effects for ingroups over outgroups (Rothermund & Wentura, 2001, 2004; Kinoshita & Peek-O'Leary, 2005). To verify whether this is indeed the case, in Chapter 5, I test whether white/black and young/old target categories produce valence and/or salience effects in the IAT.

To anticipate the results, the data from Chapters 2-5 support the fluency account of salience asymmetries in the IAT, showing that the more fluently processed target category is the more salient target in the IAT. However, there is one IAT effect that appears to contradict this account – the mere acceptance effect. The mere acceptance effect (C.J. Mitchell, 2004) describes an IAT effect in which rule-conforming stimuli are compatible with positive attributes in the IAT. For example, in C.J. Mitchell's study (Experiment 1), participants in the 'Flight'/'No Flight' condition were instructed to classify target items according to whether they can fly (e.g. balloon, arrow) or not (e.g. kitten, zipper). Those in the 'Teeth'/'No Teeth' condition classified the same items according to whether they had teeth (e.g. kitten, zipper) or not (e.g. balloon, arrow). Results revealed that the target category that conformed to the given rule (Flight/Teeth) was more easily classified with positive attributes in the IAT. In one mere acceptance effect (C.J. Mitchell 2004, Experiment 2), it appears that the less fluently processed target category is the salient target

in the IAT. This finding contradicts the results of the IAT studies presented in the thesis thus far, which show that the salient target category is the more fluently processed target. Therefore, in Chapter 6, I investigated the underlying mechanism of the mere acceptance effect in an attempt to elucidate the nature of salience asymmetries in the IAT.

Although salience asymmetries appear to threaten the construct validity of the IAT, there is still much about salience asymmetries that need to be clarified. Understanding the nature of non-associative contaminants in the IAT will be fruitful in two important ways. Firstly, it may shed light on the mechanisms underlying previous IAT effects claimed to be non-associative in nature, such as the race IAT effect, contextual IAT effects, the insect/nonword IAT effect, and De Houwer's (2001) category-level IAT effect. Secondly, this information may allow us to 'decontaminate' the IAT by controlling for salience asymmetries. This should, in turn, increase the accuracy of attitude measurement.

Chapter 2. Determinants of salience asymmetry: visual search asymmetries versus processing fluency

In Chapter 1, it was argued that salience asymmetries are an important issue in IAT research. To investigate the mechanisms behind salience asymmetries in the IAT, it is first necessary to clarify the predictors of these effects. Chapter 2 examines two accounts of salience asymmetry in the IAT: Rothermund and Wentura's (2001, 2004) account drawn from ideas in the visual search literature, and an account of salience asymmetry in terms of processing fluency. These two accounts of salience will be pitted against each other in a series of experiments. Differences between the processing fluency account and Kinoshita and Peek-O'Leary's (2005, 2006) account of salience asymmetries will be explored in Chapters 4 and 5.

To review, Rothermund and Wentura (2001, 2004) consider that salience asymmetries in the IAT align with visual search asymmetries. Thus, categories that are more readily detected in visual search – in particular, unfamiliar items (Strayer & Johnston, 2000; Wang et al., 1994) and negative items (Fox et al., 2000; Ohman et al., 2001a) – are also more salient in the IAT, making them easier to classify together. However, there are two unresolved issues with this argument. The first issue concerns the mechanism underlying search asymmetries that result from manipulations of familiarity. Although it is easier to find novel targets among familiar distractors than vice versa, this effect is not thought to be due to the novelty of the target, but to the familiarity of the distractor. Much research has shown that all kinds of targets are detected more quickly among familiar distractors, regardless of whether the targets themselves are familiar or novel (Shen & Reingold, 2001; Malinowski, & Hübner, 2001; Wolfe, Alvarez, Wong, & Klempen, 2000, cited in Wolfe, 2001). For example, in Wolfe et al.'s (2000) study, the targets were either V (familiar) or inverted V (unfamiliar), and the distractors consisted of A or inverted A. When averaged across distractor conditions, participants were faster to find both the familiar and unfamiliar target (V or inverted V) among familiar distractors, than they were to find both targets among unfamiliar distractors. Because familiar stimuli are more easily processed than unfamiliar stimuli, this suggests that the ease of distractor rejection is at the heart of the search asymmetry. This makes it difficult to apply conclusions from visual search asymmetry effects to salience asymmetries in the IAT, because there are no distractor stimuli present in IAT trials, and each stimulus is categorized individually.

The second issue involves the nature of search asymmetries between different types of affective stimuli. Rothermund and Wentura (2001, 2004) assume that negative items are more salient than positive items because search for a negative target among positive distractors is quicker than the reverse. They cite a study by Fox et al. (2000) showing that angry faces are detected more quickly among happy faces than vice versa, and another by Ohman et al. (2001a), showing that people are faster to find threatening stimuli (spiders and snakes) among non-threatening stimuli (flowers and mushrooms) than the reverse. However, other studies reveal that affective search asymmetries do not fall along this simple positive/negative dichotomy. Ohman, Flykt and Esteves (2001b) showed that angry faces are detected more quickly than sad or scheming faces, both when the targets are presented among neutral distractors, and when are were presented among other emotional faces. This led Ohman et al. (2001b) to argue that the 'angry advantage' is due to threat rather than negative emotion. Examining this emotional search asymmetry in more detail, Williams, Moss, Bradshaw, and Mattingley (2005) had participants search for angry. happy, sad or fearful faces among neutral face distractors. When Williams et al. (2005)

directly compared the size of the search asymmetries between the different emotional conditions, they showed that both angry and happy faces were detected more quickly among neutral faces than were sad or fearful faces. This demonstrates that affective search asymmetries do not depend on the positive/negative valence of the stimuli, but the specific emotion that is shown.

As mentioned in Chapter 1 (Section 1.6.3), a variable that may account for salience asymmetry effects between categories that differ in familiarity and valence is processing fluency. Familiar stimuli are generally more fluently processed than unfamiliar stimuli (Balota & Chumbley, 1984; Whaley, 1978), and positive words are also processed more fluently than negative words (Estes & Adelman, in press). In the IAT, participants may notice these processing fluency differences when classifying the two target categories in the initial target classification task (Task 1 in Table 1.1), and the two attribute categories in the following attribute classification task (Task 2 in Table 1.1). They may then use this fluency asymmetry as a cue in deciding which categories to focus on. In this way, participants may find it easier to classify the two more fluently classified target and attribute categories together (familiar and positive), and also the two less fluently classified target and attribute categories together (unfamiliar and negative).

In the experiments presented in this chapter, visual search asymmetries and processing fluency were manipulated independently to determine which dimension was responsible for salience asymmetry effects in the IAT. This led to the creation of two different category pairs that varied orthogonally along these dimensions. To anticipate the results of Experiments 1a-1d, in one pair of categories (upright elephants vs. inverted elephants), the category which was more quickly detected on a visual search task (inverted elephants) was processed more slowly (less fluently) on a separate binary classification task compared with the other category (upright elephants). In another pair of categories (big cows vs. small cows) the category which was more easily detected in the visual search task (big cows) was classified more quickly on a separate binary classification task than the other category (small cows). According to Rothermund and Wentura (2001, 2004) items that are detected more quickly in visual search (inverted elephants and big cows) are more salient than their distractors, and thus should be compatible with other salient categories in the IAT. If this is correct, then inverted elephants and big cows should behave similarly in the IAT. However, if processing fluency underlies salience asymmetries in the IAT, then we would expect the two categories that are more fluently processed (upright elephants and big cows) to produce similar effects in the IAT.

2.1. Experiments 1a-1b

Experiments 1a and 1b were conducted to verify that the stimuli used in the following experiments exhibited the expected visual search asymmetries. Visual search asymmetries have been previously demonstrated with stimuli that vary in familiarity (e.g. Strayer & Johnston, 2000; Wang et al., 1994; Shen & Reingold, 2001; Malinowski, & Hübner, 2001; Wolfe et al., 2000, cited in Wolfe, 2001) and size (Treisman & Gormican, 1988). For example, Wolfe et al. (2000, Experiment 4, cited in Wolfe, 2001) showed that search for an inverted elephant silhouette (unfamiliar stimulus) among upright elephant silhouettes (familiar stimuli) was faster than the reverse. Very similar stimuli were used here to show a familiarity search asymmetry. Treisman and Gormican (1988, Experiment 1) demonstrated that a big item (an 8mm line) also was detected more quickly among smaller items (6.5mm or 5mm lines) than vice versa. Based on this principle, I created two categories that differed only in size: big cows and small cows. It was expected that people would be quicker to process displays in which there was one big cow among multiple small cows, than vice versa.

2.1.1. Method

Participants

First-year psychology students from the University of New South Wales volunteered to participate in the experiment in exchange for course credit. Twelve students participated in Experiment 1a (elephant visual search task), and 8 students participated in Experiment 1b (cow visual search task).

Stimuli and Apparatus

The experiment was run on an IBM-compatible PC using Inquisit version 2.0 software by Millisecond. The same apparatus was used in all of the experiments reported in this thesis. The stimuli in Experiment 1a consisted of four elephant silhouette drawings that measured 24.6mm x 17.5mm. Two elephants were upright, they belonged to the category of 'live' elephants. The other two elephants were inverted, they belonged to the category of 'dead' elephants. In each category, one elephant faced left, and the other faced right. The stimuli in Experiment 1b were four colored line drawings of cows on a background of grass. The cows had identical features, except that two of the cows measured 27.8mm x 20.1mm, and the other two measured 44.6mm x 33.6mm. All cows were presented on a background of green dots that was 49.5mm x 38.1mm. The elephant and cow stimuli are presented in Figure 2.1. For half the participants, the instruction indicated that the two smaller cows belonged to the 'Zif' category (denoting one fictitious breed of cow), and the two larger cows belonged to the 'Wug' category (denoting another fictitious breed of cow);

this assignment was reversed for the remaining participants. One cow from each category faced left, and the other cow faced right. Cows were referred to by their breed, rather than by their size, because 'big' is the linguistically unmarked, default category when referring to size. According to Kinoshita and Peek-O'Leary (2005, 2006) and Rothermund and Wentura (2001, 2004), linguistic markedness may be a source of salience asymmetries in the IAT. Thus, if the stimulus categories were defined by unmarked (big) and marked (small) labels, it may be that the source of salience asymmetries could stem from the stimulus labels, rather than from the stimulus features themselves (e.g. the visual search asymmetry they incur, or their processing fluency). Linguistic markedness as a source of salience asymmetries is an important issue in later experiments, thus these labeling measures were adopted in Experiment 1b to maintain consistency of the stimuli across experiments. It could be argued that 'live' elephants are also linguistically unmarked relative to 'dead' elephants, because live objects are more familiar than dead objects. However, because familiarity is proposed to be an important component of visual search asymmetries and salience effects in the IAT, it was considered necessary to examine categories that differ in familiarity, and by implication, linguistic markedness associated with familiarity as well.

Procedure

Both the elephant visual search task and the cow visual search task followed the same format. Participants first received 40 randomized trials in which they classified the stimuli as belonging to either the live or dead category in the case of elephants (Experiment 1a), or the Zif or Wug category in the case of cows (Experiment 1b). The feedback that

participants received for incorrect responses in Experiment 1b allowed them to learn which size of cow belonged to which breed. The purpose of this task was to allow participants to differentiate between the two categories of stimuli, which was a necessary requirement for the following visual search trials. In each classification trial, a stimulus was presented onscreen, and participants pressed the left hand 'a' key if it belonged to one category, and the right-hand '5' key if it belonged to the other category. The key-category assignment was counterbalanced between participants in both experiments. During the task, the category labels were presented on the side of the screen that corresponded with the response-assignment of the categories. Stimuli were presented onscreen until a response was made. Incorrect responses received the feedback of 'WRONG RESPONSE' presented in red font at the centre of the screen for 200ms.



Figure 2.1. Elephant stimuli (top panel) and cow stimuli (bottom panel). The stimuli from the elephant categories differ in their upright/inverted orientation, whereas stimuli from the cow categories differ in size. Only half of the stimulus set is presented here, the remaining stimuli were mirror-reversed versions of the above exemplars.

Following the binary classification task, participants completed the visual search phase. A fixation cross appeared for 500ms at the start of each trial, followed by a visual search array. In each array, eight stimuli were joined to form an outline of a rectangle, each side consisting of three stimuli. Participants pressed one key (the I or E key) if all the stimuli belonged to the same category (a 'same' trial), and another key (again the I or E) if one of the stimuli belonged to a different category (a 'different') trial. This constituted four different trial conditions for the two categories: two categories (e.g. live/dead elephants) x two response types (same/different). In other words, the conditions could be described as: all live elephants (or big cows), all dead elephants (or small cows), a majority of live elephants (or big cows), and a majority of dead elephants (or small cows). There were 12 practice trials and 192 test trials, with 48 randomized test trials in each of the four conditions, presented in a randomized order. The stimulus that was the odd-one-out on the different trials appeared an equal number of times (six times) in each of the eight positions of the stimulus array. When the search array appeared, the 'same' and 'different' response labels were presented on the side of the screen that corresponded with the assigned response key.

2.1.2. Results

Data reduction

One participant was excluded in Experiment 1b for committing 16 errors in at least one of the conditions (33.3% of all test trials). In the test trials of all the experiments reported in this thesis, erroneous responses were omitted from the analysis (5.6% in Experiment 1a, 5.5% in Experiment 1b), as were those that were 3.5 standard deviations above the mean in each condition (1.0% in Experiment 1a, 0.7% in Experiment 1b).

Visual search task analysis

T-tests were conducted to compare the mean reaction times and errors for each condition in Experiments 1a and 1b. In Experiment 1a, when there was an odd-elephant-out, participants found it easier to detect an inverted elephant among upright elephants (1087ms), than vice versa (1147ms), t(11) = 2.72, p < .05. Reponses for the odd-one-out inverted elephant were also more accurate ($M_{error} = 2.33$) than in the condition in which the upright elephant was the odd-one-out ($M_{error} = 3.50$), t(11) = 2.65, p < .05. Participants were equally quick to respond to an array that consisted only of upright elephants (1131ms), t<1. However, more errors were committed in the latter condition (1.83 for upright elephants vs. 3.08 for inverted elephants), t(11) = 2.53, p < .05, implying that it was more difficult for participants to identify all the inverted elephants as being the same.

In Experiment 1b, a search asymmetry between big and small cows indicated that big cows were more quickly detected among small cows (866ms), than the reverse (987ms), t(7) = 6.11, p < .05. Responses were also more accurate when the odd-cow-out was big (M_{error} = 2.00), than when it was small (M_{error} = 4.25), t(7) = 2.55, p < .05. There was no difference in reaction time between identifying small cows as all the same (898ms), and identifying big cows as all the same (968ms), t(7) = 1.56, p = .16, and there were an equal number of errors in each condition (3.13 vs. 1.25 respectively), t(7) = 1.74, p = .13.

The results of Experiments 1a-1b demonstrate the expected search asymmetries between the elephant and cow categories. Participants were faster to find inverted elephants among upright elephants than vice versa, and they were also faster to find big cows among small cows than vice versa.

2.2. Experiments 1c-1d

Experiments 1c and 1d were conducted to establish that the two pairs of categories that exhibited visual search asymmetries in the previous experiments varied orthogonally in terms of processing fluency. In Experiment 1c, it was expected that inverted elephants that were easier to detect in the visual search task would be classified more slowly than upright elephants. This is because inverted elephants, due to their novelty, would not have a pre-existing mental representation, and thus would require more processing to reach the threshold of stimulus identification. For Experiment 1d, it was predicted that the big cows that were more quickly detected in the visual search task would be classified more quickly than the small cows. This is because Treisman and Gormican (1988) consider that larger values on a quantifiable dimension (such as size) mark the presence of a feature that allows larger stimuli to be discriminated from smaller stimuli more easily than vice versa.

2.2.1. Method

Participants

First-year psychology students from the University of New South Wales volunteered to participate in the experiment in exchange for course credit. There were 14 participants in Experiment 1c (elephant binary classification task), and 12 participants in Experiment 1d (cow binary classification task).

Procedure

The stimuli were the individual elephant and cow exemplars taken from Experiments 1a and 1b. Each binary classification task consisted of 160 randomized trials (40 presentations each of, for example, live elephants facing left, live elephants facing right, dead elephants facing left, dead elephants facing right), the first sixteen of which were practice trials (four presentations of each stimulus). The intertrial interval was 300ms. On each trial, a stimulus was presented onscreen, and participants were required to assign it to one of two categories using either the left 'a' key or the right '5' key. Key assignment was counterbalanced between participants. Participants in the elephant binary classification task categorized elephants as either being live (upright) or dead (inverted). Those that completed the cow binary classification task categorized cows as belonging to either the 'Zif' breed of cow, or the 'Wug' breed of cow. Participants were told that the two breeds differed in size, but were not informed which breed of cow was bigger or smaller. For half the participants, the big cows were of the Zif breed, and for the remaining participants the big cows were of the Wug breed. The category labels 'live' and 'dead' or 'Zif and 'Wug' were presented on the side of the screen that corresponded with the assigned response for that category. Participants were informed that they would receive feedback for incorrect responses, which allowed them to learn which size of cow belonged to which breed. Incorrect responses received the feedback of 'WRONG RESPONSE' which was presented in red at the centre of the screen for 200ms.

2.2.2. Results and Discussion

Data reduction

Participants in this and all subsequent experiments reported in this thesis were replaced if they made 24 or more errors (16.7%) in the separate binary classification task (none in Experiment 1c, one in Experiment 1d). A lower error rate was adopted as the criterion in the classification task than in the visual search task, because the former is easier to perform than the latter. Erroneous responses were excluded from the analysis (4.0% in Experiment 1c, 4.5% in Experiment 1d), as were those that were below 200ms (none in Experiment 1c, 2.1% in Experiment 1d), and those that were 3.5 standard deviations above the mean in each condition (1.3% in Experiment 1c, 0.1% in Experiment 1d).

Classification task analysis

T-tests were conducted to compare the mean reaction times and errors for the two categories in each classification task. In Experiment 1c, participants were quicker to classify upright elephants (482ms), than inverted elephants (499ms), t(13) = 2.75, p < .05. Equal accuracy in both cases (M_{error} = 3.00 for upright elephants vs. 2.79 for inverted elephants, t<1) suggests that there was no speed-accuracy trade-off in the reaction time data. In Experiment 1d, responses to big cows (481ms) were faster than to small cows (502ms), t(11) = 2.86, p < .05, and the number of errors in each condition was the same (3.08 vs. 3.42 respectively, t<1), again indicating that there was no speed-accuracy trade-off.

Experiments 1a-1d demonstrated that the elephant and cow category pairs exhibit different patterns of search asymmetries and fluency asymmetries. In the elephant pairing, the category that was more quickly detected in the visual search task (dead elephant) was less fluently responded to on a binary classification task. In contrast, in the cow pairing, the category that was more quickly detected in the visual search task (big cow) was also more fluently responded to on a binary classification task. Therefore, these stimuli were next used to investigate the effects of search asymmetries and fluency asymmetries in the IAT.

2.3. Experiments 2a-2c

Experiments 2a and 2b examined whether stimuli that are detected more quickly in visual search (inverted elephants and big cows) also behave similarly when classified with words and nonwords in the IAT. Rothermund and Wentura (2001, 2004) consider that nonwords are more salient than words, because nonwords are unfamiliar. They predict that items that are more easily detected in visual search are more salient, and thus should be compatible with nonwords in the IAT. If this principle is correct, then we would expect the animal category that was easier to distinguish among stimuli of the other category (i.e. inverted elephants and big cows) to be compatible with nonwords, and the animal categories that were harder to detect in visual search when there was a search asymmetry (i.e. upright elephants and small cows) to be compatible with words. It may be, however, that processing fluency underlies salience asymmetries in the IAT. On this view, it is predicted that the more fluently processed animal categories (upright elephants and big cows) will be more easily classified with words over nonwords, because words are expected to be more fluently processed than nonwords due to their familiarity. This was tested in Experiments 2a and 2b by placing elephants and cow targets respectively in an IAT with words and nonwords. To verify that words are more fluently processed than nonwords, I also tested the word and nonword stimuli in a separate binary classification task in Experiment 2c.

2.3.1. Method

Participants

First-year psychology students from the University of New South Wales volunteered in exchange for course credit. There were 8 participants in the elephant/word

IAT (Experiment 2a), 16 participants in the cow/word IAT (Experiment 2b), and 24 participants in the word/nonword binary classification task (Experiment 2c).

Stimuli and Apparatus

The stimuli and apparatus in Experiments 2a and 2b were the same as those in Experiments 1c and 1d respectively. In addition, Experiments 2a-2c included eight word stimuli (*angry*, *bad*, *cold*, *crude*, *mean*, *rude*, *cruel*, *nasty*), and eight nonword stimuli (*clure*, *cren*, *dolab*, *druc*, *meed*, *nady*, *staun*, *yarg*). Word stimuli with an unpleasant meaning were chosen to minimize the possibility that an animal category would be compatible with the word category on an affective basis, because words might be preferred to nonwords due to their greater familiarity (e.g. Zajonc, 1968). The category label 'nonword' was avoided, because its negative linguistic structure may be interpreted as having unpleasant connotations. Instead, participants classified words as belonging to the category 'English', and nonwords as belonging to the category 'Foreign'.

Procedure

Experiments 2a and 2b. Participants performed five classification tasks in the IAT: 1) an animal classification task (live vs. dead elephants, or Zif vs. Wug cows), 2) a word classification task (English vs. Foreign words), 3) an animal and word combined classification task, 4) an animal classification task with reversed response assignment, and 5) an animal and word combined classification task with reversed response assignment for animals.

Task 1: The animal classification task (Task 1) was the same as the classification trials of Experiments 1a and 1b, except that participants only received 24 trials (12 presentations of each category). In addition, participants in Experiment 2b (cow/word IAT) were told that the two categories of cows they were to classify differed in size, but were not told which size of cow belonged to which breed. They were informed that they would receive feedback only when they made an incorrect response, and this allowed them to work out which cow belonged to which breed.

Task 2: The word classification task followed the same procedure as the animal classification task, except that participants classified words as belonging to the 'English' category, and nonwords as belonging to the 'Foreign' category.

Task 3: In the combined classification task participants categorized the animal categories with the word categories using the same keys that were assigned to them in Tasks 1 and 2. One animal category was assigned the same key response as words (e.g. 'a'), and the other animal category was assigned the same key response as nonwords (e.g. '5'). For the other half of participants, the response assignment for the animal categories was reversed. The category labels 'English', 'Foreign' and 'live'/'dead' or 'Zif'/'Wug' were presented on the side of the screen that corresponded to the assigned responses.

There were 80 trials in the combined classification task, the first 16 of which were practice trials (eight animals and eight words/nonwords randomly selected) that were excluded from the analysis. The remaining 64 trials consisted of each of the four animal stimuli (live/dead elephants or big/small cows facing left or right) presented eight times, and each of the eight word/nonword stimuli presented four times. Animal classification trials alternated with word classification trials. Participants were instructed to respond as quickly as possible in categorizing all items, but not so quickly that they made many errors. *Task 4:* The animal classification task with reversed response assignment was similar to Task 1, but included two changes. Firstly, the response assignment to the animal categories was reversed. Secondly, the number of trials was doubled from 24 to 48, following the recommendation of Nosek, Greenwald and Banaji (2005) that this adjustment reduces the carry-over effects of the target response assignment from Task 2 and Task 3.

Task 5: Participants then completed the animal and word combined classification task with reversed response assignment for animals. This task was identical to Task 3, except that the response assignment to the target categories was reversed.

For all the trials, stimuli were presented onscreen until a response was made. Incorrect responses received the feedback of 'WRONG RESPONSE' presented in red font at the centre of the screen for 200ms. The intertrial interval was 300ms.

Experiment 2c. The binary classification task was identical to that used in Experiment 1c, except that participants classified words and nonwords as belonging to the categories 'English' and 'Foreign'. Each stimulus was presented 20 times.

2.3.2. Results

Data reduction

Participants in this and all subsequent experiments reported in this thesis were replaced if they made 10 or more errors (31.25%) in any one condition on the IAT combined classification tasks, and (there was no such cases in Experiment 2a, and 5 in Experiment 2b). Incorrect responses were discarded from the analysis (11.0% in Experiment 2a, 6.1% in Experiment 2b, 6.8% in Experiment 2c). Reaction time outliers were also excluded. Outliers at the lower end of the spectrum were those below 300ms¹ in Experiments 2a-2b (0.5% in Experiment 2a, and none in Experiment 2b) and less than 200ms in Experiment 2c (none in this case). Outliers at the higher end of the spectrum were set at 3.5 standard deviations above the mean for each participant for the condition (0.9% in Experiment 2a, 1.1% in Experiment 2b; 1.0% in Experiment 2c). Mean reaction time scores were calculated for each of the two combined classification tasks in Experiments 2a and 2b, and for each condition in Experiment 2c. Because some IAT experiments required the analysis of an interaction effect, all IAT analyses were peformed using an ANOVA to maintain consistency. A set of orthogonal contrasts was tested using a multivariate, repeated measures model (O'Brien & Kaiser, 1985).

Experiment 2a: elephant/word IAT

Mean reaction time scores were calculated for each of the two combined classification tasks. Figure 2.2 shows that participants responded 188ms faster when upright elephants shared a key with words, and inverted elephants shared a key with nonwords, than vice versa F(1,7) = 10.58, p < .05. There was also a trend toward responses being more accurate when upright elephants were classified with words, and inverted elephants words, and inverted elephants were classified with words, and inverted elephants were classified with words, and inverted elephants were classified with nonwords (M_{error} = 3.31) than in the other condition (M_{error} = 3.75), F(1,7) = 3.94, p = .09.

¹ The IAT experiments used a 300ms criterion consistent with standard IAT experiments (e.g. Greenwald et al. (1998). This minimum outlier value is higher than that adopted in the binary classification tasks (200ms), which is in keeping with the relative difficulty of the two tasks.



Figure 2.2. Mean reaction times for the elephant/word IAT of Experiment 2a. The open bar represents the condition in which upright elephants were classified with words, and inverted elephants were classified with nonwords. The filled bar represents the condition in which upright elephants were classified with nonwords, and inverted elephants were classified with nonwords.

Experiment 2b: cow/word IAT

In Experiment 2b, the breed of the two cows (Zif/Wug) did not interact with the contrast of interest (F(1,14) = 2.54, p = .13 for reaction times, and F<1 for errors). Thus, the following analysis was averaged across both of the cow breed conditions. Figure 2.3 shows that the reaction time for the task in which big cows were classified with words and small cows were classified with nonwords was 109ms faster than the reaction time for the task in which the reaction time for the task in which be compared to the task in which the reaction time for the task in which the reaction time for the task in which be compared to the task in which the reaction time for the task in which the reverse was true, F(1,15) = 8.04, p < .05. There was no difference

between the number of errors made in the two conditions (1.63 for Big Words/Small Nonwords, vs. 2.25 for Big Nonwords/Small Words; F(1,15) = 1.94, p = .18). This suggests that the IAT effect seen in reaction times was not due to a speed-accuracy trade-off.



Figure 2.3. Mean reaction times for the cow/word IAT of Experiment 2b. The open bar represents the condition in which big cows were classified with words, and small cows were classified with nonwords. The filled bar represents the condition in which big cows were classified with nonwords, and small cows were classified with words. Error bars represent the standard error of the mean.

Experiment 2c: word/nonword binary classification task

In Experiment 2c, participants classified words (603ms) more quickly than nonwords (624ms), t(23) = 2.54, p < .05. Response accuracy was the same in both conditions ($M_{error} = 5.21$ for words vs. $M_{error} = 4.54$ for nonwords), *t*<1, suggesting that there was no speed-accuracy trade-off in the corresponding reaction times.

2.3.3. Discussion

In Experiment 2a, upright elephants were more easily classified with words, and inverted elephants were more easily classified with nonwords. This result is consistent with Rothermund and Wentura's (2001, 2004) account that stimuli that are more easily detected in visual search (inverted elephants) are compatible with the unfamiliar category of nonwords. However, the result of Experiment 2b shows the reverse of this prediction. In this case, the category that was more easily detected in visual search (big cows) was compatible with words, not nonwords in the IAT. Thus, stimuli that behave similarly in a visual search task (inverted elephants and big cows), do not, as Rothermund and Wentura would predict, behave similarly in the IAT. Instead, the animal categories that were more fluently processed on the binary classification task (upright elephants and big cows) are compatible with words, and the less fluently processed animal categories (inverted elephants and small cows) are compatible with nonwords.

The finding that words were more fluently processed than nonwords supports our prediction that the more fluently processed animal categories are compatible with the more fluently processed word category in the IAT. Taken together, Experiments 2a-2c suggest that processing fluency is a more reliable predictor of non-associative compatibility effects in the IAT than visual search asymmetries.

However, the compatibility effects observed between the animal and word categories may also have been based on valence differences. For example, in Experiment 2a, participants may have found it easier to classify upright ('live') elephants with ('English') words, and inverted ('dead') elephants with ('Foreign') nonwords, because 'live' and 'English' may have been preferred to 'dead' and 'Foreign' due to their familiarity (the mere exposure effect, e.g. Zajonc, 1968). Alternatively, this effect may have occurred because participants considered 'live' and 'English' to have positive connotations, and/or 'dead' and 'Foreign' to have negative connotations. It is possible that similar associations also exist between the word categories and the cow categories that differ in size. If this is the case, then these factors may have obscured salience asymmetries between the word and animal categories as defined by Rothermund and Wentura (2001, 2004). That is, Rothermund and Wentura may have been correct in their prediction that items that are easily detected in visual search are more salient, and hence should be more easily classified with nonwords than words, but this effect was overshadowed by valence/associative factors between the animal and word categories.

The results of Experiments 2a and 2b appear to provide evidence against the idea that salience asymmetries in the IAT correspond with visual search asymmetries. However, a more direct test of this notion should follow the design of Rothermund and Wentura (2004, Experiments 1b, 1d, and 1e), who placed stimuli that exhibited visual search asymmetries into an IAT to be classified together. Therefore, Experiment 3 was conducted to verify whether the elephant and cow categories that were more quickly detected in a visual search task in Experiments 1a and 1b would also be compatible in the IAT.

2.4. Experiment 3

The results of Experiments 2a and 2b suggest that salience asymmetries in the IAT do not consistently correspond with visual search asymmetries. Experiment 3 examined this issue more directly by having participants classify upright/inverted elephants with big/small

cows in the IAT. If salience asymmetries correspond with visual search asymmetries, then the two categories that are detected more quickly in a visual search task (inverted elephants and big cows) should be compatible in the IAT. However, if salience asymmetries correspond with differences in processing fluency, then the two categories that are more quickly classified in a binary classification task (upright elephants and big cows) should be compatible in the IAT.

2.4.1. Method

Participants

Thirty-two first-year psychology students from the University of New South Wales volunteered in the experiment in exchange for course credit.

Procedure

The procedure was the same as that of Experiment 2a, except that instead of words and nonwords, participants classified Zif and Wug cow categories. The cow stimuli were those used in Experiment 2b. As in the previous experiments involving the cow categories, the assignment of cows breed (Zif/Wug) to the two cow sizes was counterbalanced.

2.4.2. Results

Data reduction

Five participants were replaced because they made 10 or more errors in at least one of the conditions. Incorrect responses were discarded from the analysis (7.0%), as were those which were less than 300ms (none in this case), or 3.5 standard deviations above the mean for each participant for the condition (0.8%).

IAT analysis

The mean reaction time scores for each of the two combined classification tasks are illustrated in Figure 2.4. The between groups factor of cow breed (Zif/Wug) did not interact with the contrasts of interest F<1 (for both reaction times and errors), allowing the data to be collapsed across both cow breed conditions. Responses were 92ms faster when upright elephants shared a key with big cows, and inverted elephants shared a key with small cows, compared to when the pairings were reversed, F(1,31) = 6.54, p < .05.



Figure 2.4. Mean reaction times for the elephant/cow IAT of Experiment 3. The open bar represents the condition in which upright elephants were classified with big cows, and inverted elephants were classified with small cows. The filled bar represents the condition in which upright elephants were classified with small cows, and inverted elephants were classified with big cows. Error bars represent the standard error of the mean.

There was equal accuracy in both conditions, ($M_{error} = 2.25$ for Upright Big/Inverted Small vs. $M_{error} = 2.23$ for Upright Small/Inverted Big) *F*<1, suggesting that there was no speed-accuracy trade-off in the corresponding reaction time data. These results provide evidence that salience asymmetry effects in the IAT are related to fluency asymmetries between categories and not to visual search asymmetries.

2.5. General Discussion

Rothermund and Wentura (2004) propose that salience asymmetries in the IAT follow visual search asymmetries, and have reported several experiments to support this claim. However, another variable that may be involved in these effects is processing fluency. The experiments in this chapter provide evidence that processing fluency, rather than search asymmetries, is diagnostic of salience asymmetries in the IAT.

To compare search asymmetries against processing fluency, stimuli were used that varied independently along these dimensions (Experiments 1a-d). In Experiments 1a and 1c, search was faster for inverted elephants among upright elephants than vice versa, but inverted elephants were responded to more slowly than upright elephants on a binary classification task. In Experiments 1b and 1d, big cows were more quickly detected as the odd-one-out among small cows than small cows were among big cows, and participants were also faster to categorize big cows than small cows on a binary classification task. Thus in one pair of categories (upright/inverted elephants), the category that was more quickly detected in visual search was the *less* fluently processed category, whereas in another pair of categories (big/small cows), the category that was more quickly detected was the *more* fluently processed category.

Experiments 2a and 2b used the same elephant and cow stimuli to reveal that the category that was more fluently processed in the binary classification task (upright elephants and big cows) was also more easily classified with words over nonwords in the IAT. These effects occurred regardless of whether or not that category was quickly detected in a visual search task. Experiment 2c established that words were more fluently processed than nonwords, confirming that the IAT effects of Experiments 2a and 2b were in line with fluency asymmetries between the category pairs.

It is also possible that the results of Experiments 2a and 2b are due to affective/semantic similiarity between the animal and word categories. In an attempt to rule out these factors as an explanation of the animal/word IAT effects, the results of Experiments 2a and 2b were replicated in another sets of experiments² using the more affectively neutral labels of 'word' and 'nonword' as the attribute labels, instead of 'English' and 'Foreign'. In the same study, participants also rated the pleasantness of the target and attribute stimuli and categories. Nonwords were rated as more positive than words at the stimulus level, but words were rated as more positive than nonwords at the category level. Upright elephants were rated as more pleasant than dead elephants at the stimulus and category level. Therefore, the elephant/word IAT effect of Experiment 2a may have been due to category level affective associations. In contrast, big and small cows were rated as being equally pleasant at the stimulus and category level. This suggests that the cow/word IAT effect is unlikely to reflect affective associations, as there were no evaluative differences between big and small cows.

 $^{^{2}}$ These experiments were conducted for the purposes of publication just prior to the submission of this thesis. Therefore, there was insufficient time for these experiments to be included in this chapter.

In Experiment 3, compatibility effects between categories displaying visual search asymmetries and fluency asymmetries were tested by conducting an IAT in which participants classified upright/inverted elephants with big/small cows. It was found that the more fluently processed target categories (upright elephants and big cows) were compatible on the IAT, rather than the categories that were detected more quickly in visual search (inverted elephants and big cows). This provides further evidence that salience asymmetries are driven by differences in processing fluency, rather than visual search asymmetries.

These results clarify the salience asymmetry account of Rothermund and Wentura (2001, 2004) by demonstrating that processing fluency, rather than visual search asymmetry, predicts salience asymmetries in the IAT. The processing fluency account may also explain other salience asymmetry effects in the IAT, such as Rothermund and Wentura's (2004) findings that items that are detected faster in visual search tend to be compatible in the IAT. In their Experiment 1d, participants performed a visual search task in which they were required to judge whether all the stimuli in the visual search array were the same, or if one was different. Participants were faster to detect the name of an old person among the names of young people, a nonword among words, and a multi-colored string among single-colored strings, than they were when the category assignments were reversed. In other experiments (Experiments 1b and 1e), participants classified old and young names with each of the other category pairs (words/nonwords, single-/multi-colored strings) in an IAT. It was found that old names were compatible with other categories that were also more quickly detected (as the odd-one-out) in the visual search task (i.e. nonwords and multi-colored strings). Rothermund and Wentura took these results as evidence that salience asymmetries in the IAT correspond with visual search asymmetries. However, it is possible that these search asymmetries were confounded with processing
fluency. As old names and nonwords are less familiar than young names and words, they are likely to be processed more slowly. These fluency differences may have accounted for the compatibility effects between old names and nonwords, and between young names and words in an IAT. Similarly, multi-colored strings are likely to take longer to process than single colored strings, because multi-colored strings are more complex and do not form a unified perceptual configuration. Thus, multi-colored strings may have been compatible with old names in the IAT because those two categories were the less fluently processed of their respective category pairs.

One reason why visual search asymmetries are difficult to reconcile with salience asymmetries in the IAT is that there are multiple mechanisms responsible for visual search asymmetries. Firstly, search asymmetries can be due to the ease with which distractors are rejected (e.g. Treisman & Souther, 1985; Strayer & Johnston, 2000; Wang et al., 1994). Evidence shows that targets are easier to detect when distractors are familiar (Shen & Reingold, 2001; Rauschenberger & Yantis, 2006) or prototypical (Levin & Angelone, 2001). This situation is illustrated by Experiments 1a and 1c involving the elephant stimuli. In the visual search task of Experiment 1a, the odd-elephant-out was more quickly detected when the inverted elephant was the target, and the upright elephants were the distractors. This result may have been due to the fluency of distractor processing, as upright elephants were processed more quickly than inverted elephants in the binary classification task (Experiment 1c).

A second determining factor in other types of search asymmetry is the nature of the target. This can be seen in Experiments 1b and 1d involving cow stimuli, as the big cows which were detected more easily in visual search (Experiment 1b) were also identified more quickly in the classification task (Experiment 1d). This pattern of results suggests that

participants were detecting a salient feature in big cows, rather than rejecting small cows more quickly. Treisman and Gormican (1988) consider that larger values on a quantifiable dimension (such as size) mark the presence of a feature that distinguishes targets from distractors. Because participants focus on this feature as a means of discriminating between the two types of cow categories on the classification task, the category which possesses this feature is, by implication, the more salient category. As big cows are also more fluently processed than small cows, this suggests that the more fluently processed category is more salient in the IAT.

The experiments in this chapter support the idea that processing fluency, instead than visual search asymmetries, is a predictor of salience asymmetry effects. The aim of the next chapter will be to further investigate the effects of processing fluency in the context of a standard IAT with valence attribute categories.

Chapter 3. Distinguishing between the effects of valence and salience in the IAT: the development of a new procedure

The previous chapter evaluated visual search asymmetries and processing fluency as potential predictors of salience asymmetries in the IAT. The evidence strongly favored processing fluency as a predictor of salience asymmetries. When the processing fluency and visual search asymmetry of category pairs were set in opposition, it was fluent processing, rather than visual search asymmetries, which determined compatibility effects in the IAT. This provides preliminary evidence that salience asymmetries in the IAT are accompanied by differences in processing fluency.

This chapter further explores the effects of processing fluency by investigating whether patterns of classification fluency correspond with salience asymmetry effects in the standard IAT which features positive and negative attributes (rather than words and nonwords, as in Chapter 2). To examine this issue, it is necessary to distinguish between IAT effects based on shared valence, and IAT effects based on similar salience. In a standard IAT, valence and salience are confounded, such that a particular attribute category (positive according to Kinoshita and Peek-O'Leary (2005, 2006), and negative according to Rothermund and Wentura (2001, 2004)) is always more salient than the other attribute category. This makes it difficult to distinguish whether an IAT effect is due to valence differences, or to salience asymmetries. One method used by Rothermund and Wentura (2001, 2004) and Kinoshita and Peek-O'Leary (2006) to detect salience asymmetries in an IAT is to have participants classify target categories with words and nonwords instead of positive and negative attributes, as was used in Experiments 2a and 2b of Chapter 2. The assumption is that words and nonwords are asymmetrically salient because they vary in familiarity, but do not differ in salience. Thus, if the target categories show an IAT effect with the word and nonword attributes, then one can infer that there are salience asymmetries between the target categories. However, there may also be valence differences between words and nonwords in the IAT. If there are valence differences between words and nonwords, then any IAT effect using these attributes may influenced by evaluative factors. Thus, we could not be certain whether an IAT effect involving these attributes was due to salience or to valence. Even if there are no valence differences between words and nonwords, although this method may be able to detect salience asymmetries, it will not be able to measure valence differences between the target categories. This would pose a problem in IATs where there are both salience and valence differences between the target categories. Once a salience difference had been detected using word/nonword attributes, the result of the standard IAT would then be rejected on the basis of salience asymmetries. This allows for the possibility that a true difference in attitude would be rejected on the basis of the screening process for salience asymmetries, constituting a Type II error. To overcome these problems, a modified version of the IAT was adopted in the current studies to distinguish between valence differences and salience asymmetries in the IAT.

3.1. Overview of research methods

Each of the following experiments in this thesis follows a similar format. A standard IAT was first conducted. Thus, two target categories were to be classified with two attribute categories, usually pleasant and unpleasant attributes. Because I wanted to examine the effects of attribute valence independently of processing fluency, the pleasant and unpleasant attribute categories were matched on processing fluency. As in Experiments 1c, 1d and 2c of the previous chapter, an independent measure of target processing fluency

was taken. This consisted of the binary classification task used in Experiments 1c, 1d and 2c, which measured the speed and accuracy with which participants responded to the two target categories.

Following this, a modified version of the IAT was used to discriminate between effects of valence and salience in the IAT. The original IAT, which consists of two target categories and two attribute categories (pleasant and unpleasant), was divided to create two different IATs. In one IAT (the Pleasant IAT), the target categories were classified with pleasant and neutral attributes, and in another IAT (the Unpleasant IAT), the target categories were classified with unpleasant and (the same) neutral attributes. A critical aspect of this design is that positive and negative words are generally assumed to be more salient than neutral words, as demonstrated by two findings. Firstly, people show superior retention of both positive and negative valence words compared with neutral words (Rubin & Friendly, 1986), indicating that valence words receive more processing. Secondly, valence words disrupt recall performance to a greater extent than neutral distractors (Buchner, Rothermund, Wentura, & Mehl, 2004), suggesting that they command greater attentional resources than neutral words. With the design of the modified IAT, valence differences and salience asymmetries should produce two different outcomes in the Pleasant and Unpleasant IATs. If a target category is more positive than another target category, then that category should be compatible with the more positive attributes, that is, pleasant attributes (over neutral attributes) in the Pleasant IAT, and neutral attributes (over unpleasant attributes) in the Unpleasant IAT. This result will be referred to as a valence effect for the category that is compatible with the more positive attributes. This valence effect may be due to affective similarity with positive attributes, and/or semantic similarity with positive attributes. However, if the IAT effect is caused by salience asymmetries, then that category should be compatible with the more salient attributes, that is, pleasant attributes in the Pleasant IAT, and unpleasant attributes in the Unpleasant IAT. This result will be referred to as a salience effect for the category that is compatible with the more salient attributes. The Pleasant and Unpleasant IATs together shall hereafter be referred to as the split IAT.

The aim of this chapter was to examine whether the split IAT can discriminate between the effects of valence and salience in the IAT. To do this, I first applied the split IAT to target categories that have 'near-universal evaluative differences' (Greenwald et al., 1998), namely, flowers and insects. Secondly, I tested whether the split IAT can also detect salience asymmetries between Go/Nogo target categories, because these categories have been shown by Rothermund and Wentura (2001, 2004) to differ in salience. Because processing fluency was shown to be a predictor of salience asymmetry effects in the previous chapter, I also examined whether differences in processing fluency between the target categories also corresponded with salience effects on the split IAT.

3.2. Experiment 4a

The target categories of flowers and insects were used to test whether the split IAT can measure valence differences. These categories were chosen because people have stronger positive associations with flowers than with insects on both IAT and self-report measures (Greenwald et al., 1998, Experiment 1). In preparation for conducting a split IAT with flowers and insects, I first conducted a standard IAT using the same target categories to ensure that I could replicate the original flower/insect IAT effect with the current stimuli.

3.2.1. Method

Participants

First-year psychology students from the University of New South Wales volunteered to participate in the experiment in exchange for course credit. Eight students participated in the flower/insect IAT, and 16 other students performed a flower/insect binary classification task.

Stimuli

The target stimuli consisted of eight flower exemplars (*lily*, *tulip*, *violet*, *daffodil*, dandelion, marigold, iris), and eight insect exemplars (wasp, moth, locust, mosquito, cockroach, beetle, termite, dragonfly). The attribute stimuli were eight pleasant words (charming, cheerful, ethical, generous, lovely, loyal, wise, witty) and eight unpleasant words (angry, bad, cold, crude, cruel, mean, nasty, rude). To avoid confounding processing fluency with attribute valence, I also controlled for the processing fluency of the attribute categories. Attributes from the two categories were matched on word frequency based on the Kucera and Frances (1967) count ($M_{positive} = 24.5$, $M_{negative} = 74.8$ per million; t(14) =1.72, p = .11). The word length of the attributes was not controlled, because matching the word length of positive and negative words results in faster responding of positive words than negative words (Estes & Adelman, in press; Unkelbach et al., in press). Instead, the word length of pleasant attributes (M = 6.38) was increased relative to unpleasant attributes (M = 4.68; t(14) = 3.21 p < .01), to offset the fluency superiority of positive words. This ensured that the pleasant and unpleasant attributes were classified equally quickly (616ms vs. 622ms respectively, t < 1) and accurately (M_{error} = 3.69 vs. 2.81 respectively, t(15) =1.16, p = .27) on a separate classification task that consisted of 160 trials (10 presentations

of each stimulus). All stimuli were presented in 19-point, uppercase font in the centre of the computer screen.

Procedure

The procedure was the same as in Experiments 2a and 2b (animal/word IATs), except for two changes; the target categories were flowers and insects, and the attributes were pleasant and unpleasant words. This format follows that of the standard form of the IAT. The binary classification task was similar to the task used in Experiments 1c-1d, except that participants classified flowers and insects as the target categories.

3.2.2. Results

Data reduction

Three participants met the replacement criterion for the separate binary classification task. Incorrect responses were discarded from the analysis (9.1% in the IAT, 6.6% in the binary classification task). Reaction times were discarded as outliers if they were less than 300ms in the IAT and less than 200ms in the binary classification task (none in both cases), or were 3.5 standard deviations above the mean for each participant for the condition (1.6% in both tasks).

IAT analysis

The mean reaction time scores are shown in Figure 3.1. Responses were faster when flowers shared the same response key as pleasant attributes, and insects shared the same key as unpleasant attributes, than when the target and attribute assignment was reversed, F(1,7) = 43.60, p < .001. This indicated an IAT effect for flowers (317ms). The mean number of errors was also calculated for each of the two conditions. The same number of errors was made in the condition in which flowers were classified with pleasant attributes and insects with unpleasant attributes (2.06), as in the condition with the reverse target and attribute assignment (3.75), F(1,7) = 2.85, p = .14. A trend towards greater accuracy in the former condition suggests that the corresponding pattern in reaction times was not due to a speed-accuracy trade-off.



Figure 3.1. Mean reaction times for the flower/insect IAT of Experiment 4a. The open bar represents the condition in which flowers were classified with pleasant attributes, and insects were classified with unpleasant attributes. The filled bar represents the condition in which flowers were classified with unpleasant attributes, and insects were classified with unpleasant attributes, and insects were classified with pleasant attributes. The filled bar represents the condition in which flowers were classified with unpleasant attributes, and insects were classified with pleasant attributes. Error bars represent the standard error of the mean.

Processing fluency of the target categories

When flowers and insects were classified in a separate classification task, reaction times did not differ between flowers (639ms) and insects (632ms), t<1. Also, an equal number of errors was committed in both conditions (4.38 for flowers vs. 5.16 for insects), t<1, indicating that the two categories were processed equally fluently.

3.3. Experiment 4b

Having demonstrated the standard IAT effect with flowers and insects in Experiment 4a, the objective of Experiment 4b was to test whether the same effect could be replicated in a split IAT procedure. If the split IAT is sensitive to valence differences, flowers should be compatible with the more positive attributes. This means in that in the Pleasant IAT, flowers should be compatible with pleasant attributes and insects should be compatible with neutral attributes. In the Unpleasant IAT, however, flowers should be compatible with neutral attributes and insects should be compatible with unpleasant attributes. Such a pattern of results would suggest that the split IAT can detect valence differences between target categories.

3.3.1. Method

Participants

Forty-eight first-year psychology students from the University of New South Wales volunteered for the experiment in return for course credit.

Stimuli

The target and attribute stimuli were those used in Experiment 4a. In addition, the neutral attribute stimuli used in the Pleasant IAT and the Unpleasant IAT consisted of the words: *solid, concave, round, textured, curved, near, slippery, typical.* Ten participants rated the pleasant, unpleasant and neutral stimuli on a seven-point scale of valence (1 = extremely unpleasant, 7 = extremely pleasant). The ratings of the pleasant (6.34) and unpleasant (1.64) attributes were equidistant from the ratings of neutral attributes (3.75, *t* <1), indicating that the pleasant and unpleasant attributes were matched on valence extremity.

Design and Procedure

Participants were randomly assigned to perform either the Pleasant IAT, or the Unpleasant IAT. The design and procedure of the two IATs was the same as the IAT of Experiment 4a, except for one modification. Instead of classifying pleasant and unpleasant attributes, participants in the Pleasant IAT condition classified pleasant and neutral attributes, and participants in the Unpleasant IAT condition classified unpleasant and neutral attributes.

3.3.2. Results

Data reduction

Thirteen participants were replaced for committing 10 or more errors in any one condition. One other participant was replaced for having mean reaction times that were 3.5 standard deviations higher than the group mean. The analysis excluded reaction times to erroneous responses (8.85%), outliers below the minimum cutoff (0.4%), or above the

maximum cutoff (0.9%). Separate analysis was conducted for the Pleasant IAT and Unpleasant IAT, then the data from both IATs were combined.

Processing fluency of the attribute categories

To test the idea that pleasant and unpleasant attributes are more salient than neutral attributes, I first compared the reaction times of the attribute categories in the attribute classification task (Task 2) performed prior to the IAT. In the Pleasant IAT, pleasant attributes (788ms) were responded to more quickly than neutral attributes (1035ms), t(23) = 4.78, p < .001. The number of errors was equal in both conditions (1.08 for pleasant vs. 0.8 for neutral), t(23) = 1.05, p = .30. In the Unpleasant IAT, unpleasant attributes (806ms) were classified more quickly than neutral attributes (938ms), t(23) = 3.09, p < .01, and there was equal accuracy in both conditions, (M_{error} = 1.54 vs. 1.25 respectively), t<1. This indicates that valence stimuli are indeed more fluently processed than neutral stimuli. This difference in processing fluency did not interact with the valence of the salient attributes; the effect size was the same regardless of whether the attributes were positive or negative (*F*<1). Therefore, the two valence categories were matched on fluency relative to the neutral attribute categories.

Pleasant IAT

The left panel of Figure 3.2 shows the mean reaction time scores for each of the conditions in the Pleasant IAT. The open bar indicates the reaction times for when flowers were classified with pleasant attributes and insects were classified with neutral attributes, and the filled bar show the reverse combinations. The previous flower/insect IAT effect shown in Experiment 4a was replicated; responses were faster when flowers shared a key

with pleasant attributes and insects shared a key with neutral attributes, than vice versa, F(1,23) = 20.85, p < .001. The mean number of errors was also calculated for each condition. There was no difference in the number of errors committed in the two IAT conditions (2.52 for Flower Pleasant/Insect Neutral vs. 2.54 for Flower Neutral/Insect Pleasant), F<1. This result supports the idea that the Pleasant IAT effect in the reaction time data was not due to a speed-accuracy trade-off.



Figure 3.2. Mean reaction times for the flower/insect split IAT of Experiment 4b. The open bars represent the conditions in which flowers were classified with comparatively positive attributes, and insects were classified with comparatively negative attributes. The filled bars represent the conditions in which the target and attribute assignments were reversed. Error bars represent the standard error of the mean.

Unpleasant IAT

The mean reaction time data for the two Unpleasant IAT conditions are depicted in the right panel of Figure 3.2. The open bar indicates the reaction times for when flowers were classified with neutral attributes and insects were classified with unpleasant attributes, and the filled bar shows the reverse target and attribute assignment. Participants were quicker to classify flowers with neutral attributes, and insects with unpleasant attributes, than vice versa, F(1,23) = 26.44, p < .001. Again, this result replicates the standard flower/insect IAT effect, suggesting that flowers are more positive than insects. Responses were equally accurate regardless of whether flowers were classified with neutral attributes and insects with unpleasant attributes (M_{error} = 3.10), or whether the reverse was true (M_{error} = 3.17), F<1. This suggests that the IAT effect in response times was not due to a speedaccuracy trade-off.

Combined analysis of the Pleasant and Unpleasant IATs

The contributing effects of valence and salience in the split IAT were examined by combining the mean reaction times of the Pleasant and Unpleasant IAT. In this and all subsequent split IAT experiments, an overall valence effect was tested by comparing the reaction times from two pairs of conditions. In one pair of conditions, one target category (e.g. flowers) was classified with the comparatively positive attributes (pleasant in the Pleasant IAT, and neutral in the Unpleasant IAT), and another target category (e.g. insects) was classified with the comparatively negative attributes (neutral in the Pleasant IAT, and unpleasant IAT). These conditions were compared to the other pair of conditions in which the target and attribute assignments were reversed (in this case, flowers were paired with the comparatively negative attributes and insects were paired with

comparatively the positive attributes). With respect to the two panels of Figure 3.2, this analysis compares the two open bars to the two filled bars. In the current experiment, the mean reaction time difference between these two pairs of conditions was 134ms. This was a reliable difference, F(1,46) = 47.01, p < .001, indicating that flowers were more easily classified with the more positive attributes, and insects were more easily classified with the more negative attributes when both IATs of the split IAT were combined. An equal number of errors were made in both pairs of conditions, F<1.

An IAT effect based on salience asymmetries was also examined in this and all subsequent split IAT experiments. To do this, I compared the mean reaction time of the conditions in which flowers were paired with (the more salient) pleasant and unpleasant attributes and nonwords were paired with (the less salient) neutral attributes, against the conditions in which the combinations were reversed. In the two panels of Figure 3.2, the outer bars were compared to the inner bars. There was no difference in reaction times between these two pairs of conditions (6ms, F<1), indicating that there was no overall salience in the flower/insect split IAT. Nor was there a salience effect in the error data, as responses were equally accurate in both pairs of conditions, F<1. This result is consistent with the results of the binary classification task, in which flowers and insects were classified equally quickly, suggesting that there were no salience differences between the categories (as indexed by processing fluency). The pattern of results across the Pleasant IAT and the Unpleasant IAT suggests that the split IAT can detect valence differences between flower and insect categories.

3.4. Experiment 5a

The results of Experiment 4b suggest that the split IAT is sensitive to valence differences between target categories. In the present experiments, to investigate whether the split IAT can also assess salience asymmetries between target categories in the IAT, I employed the same manipulation used by Rothermund and Wentura (2001, 2004) to vary salience asymmetries – the Go/Nogo task. In this task, participants were instructed to respond to Go items and to withhold responses to Nogo items in an initial target classification task. Treating the target categories in a Go/Nogo fashion focuses attention on the Go items, which should in turn make those items more fluently processed. Therefore, we would expect the Go category to be compatible with the more salient attribute category.

Rothermund and Wentura (2001, 2004) claimed that of the two attribute categories, unpleasant is more salient and thus should be compatible with the Go target category in the IAT. In contrast, Kinoshita and Peek-O'Leary (2005, 2006) would predict that the Go target category should be compatible with pleasant attributes, as they suggest that pleasant attributes are more salient. I also measured the processing fluency of the Go/Nogo categories using a separate binary classification task, to examine whether patterns of processing fluency correspond with salience asymmetries in the IAT.

3.4.1. Method

The method was the same as in Experiment 4a, with some exceptions. Sixteen participants volunteered for the IAT, and another 16 participated in the separate binary classification task. The target (Go/Nogo) stimuli consisted of 16 words which belonged to the two categories Flight or Teeth. Flight stimuli were objects that fly but which do not

have teeth (*airplane*, *parrot*, *kite*, *bullet*, *arrow*, *rocket*, *missile*, *balloon*). Teeth stimuli were objects that have teeth but do not fly (*koala*, *bear*, *saw*, *dog*, *deer*, *rabbit*, *kitten*, *cat*).

Design and Procedure

The design and procedure were identical to the IAT in Experiment 4a for the attribute classification task (Task 2), and the combined classification tasks (Tasks 3 and 5). Participants performed a Go/Nogo task in place of the two the target classification tasks (Tasks 1 and 4). In these tasks, participants were told that they were to categorize sets of words as belonging to either the Flight category, or the Teeth category. For half the participants Flight was the Go category and Teeth was the Nogo category. The reverse was true for the remaining participants. All participants were required to press the spacebar if they saw a stimulus belonging to the Go category (e.g. FLIGHT), and to make no response if the stimulus belonged to the Nogo category (e.g. TEETH). Presentations of the 12 Go trials and the 12 Nogo trials were given in a random order, with an intertrial interval of 300ms. On Go trials, stimuli were presented onscreen until the participant made the correct response. On Nogo trials, stimuli were presented for 2000ms each. Any responses made on Nogo trials had no effect on the exposure duration of the Nogo stimuli. Above each stimulus was a reminder to press the spacebar for Go items – for example, "FLIGHT = SPACEBAR". This Go/Nogo manipulation was not used in the combined classification tasks of the IAT, where, just as in the standard IAT (see Experiment 4a), participants responded to both the Flight and Teeth categories.

A separate binary classification task was used to assess the processing fluency of the target categories. Prior to performing the binary classification task, participants first performed a Go/Nogo task with the Flight/Teeth categories. This task was similar to that of the Go/Nogo task in the IAT of the present experiment, except that it consisted of 64 trials. The binary classification task followed the same procedure as the binary classification in Experiment 4a, with participants classifying Flight/Teeth stimuli into their respective categories using the left and right key responses.

3.4.2. Results

Data reduction

No participants met the replacement criterion for either the IAT or the binary classification task in the present experiment. Incorrect responses were discarded from the analysis (6.3% in the IAT, 5.4% in the classification task), as were those which were below the minimum cutoff (0.3% in the IAT, none in the classification task), or above the maximum cutoff (0.2% in the IAT, 1.4% in the classification task).

IAT analysis

Because half the participants were assigned Flight as the Go category, and half the participants were assigned Teeth as the Go category, an analysis was conducted with the category assigned to the Go condition (Flight/Teeth) as a between groups factor. This variable did not interact with the contrast of interest (F<1 for the reaction time data, and F(1,14) = 1.72, p = .21 for the error data), and so the data were combined across groups. Figure 3.3 shows that the reaction time for the task in which Go was classified with pleasant attributes and Nogo was classified with unpleasant attributes was faster than the reaction time for the task in which the combinations were reversed. There was an IAT effect of 162ms for the Go category, F(1,15) = 6.70, p < .05. There was no difference between the number of errors made in each condition (1.97 for Go Pleasant/Nogo

Unpleasant vs. 2.09 for Go Unpleasant/Nogo Pleasant; F < 1), which suggests that the IAT effect seen in reaction times was not due to a speed-accuracy trade-off.



Figure 3.3. Mean reaction times for the Go/Nogo IAT of Experiment 5a. The open bar represents the condition in which Go items were classified with pleasant attributes, and Nogo items were classified with unpleasant attributes. The filled bar represent the condition in which Go items were classified with unpleasant attributes, and Nogo items were classified with pleasant attributes. Error bars represent the standard error of the mean.

Processing fluency of the target categories

The following analysis was conducted with the category assigned to the Go condition (Flight/Teeth) as a between groups factor, because this factor interacted with the comparison of interest for the error data F(1,14) = 21.81, p < .001. The same interaction with the reaction time data did not reach significance, F(1,14) = 1.21, p = .29. Go items

were responded to more quickly than Nogo items in the binary classification task (634ms vs. 662ms respectively), F(1,14) = 5.08, p < .05. Because there was an interaction effect on the error data, t-tests were used to analyze the simple effects that comprised this interaction. There were more errors in the Nogo category compared to the Go category when Teeth was the Go category (t(7) = 5.61, p = .001), but no difference in errors between the Go and Nogo categories when Flight was the Go category (t(7) = 1.87, p = .10). Overall, an equal number of errors were committed in response to Go items (3.50) and Nogo items (4.25), F(1,14) = 2.42, p = .14. This indicates that the reaction time data was not due to a speed-accuracy trade-off.

3.4.3. Discussion

The Go IAT effect indicates that manipulating category salience using a Go/Nogo task allows the more salient Go category to be more easily classified with pleasant attributes (and/or the less salient Nogo category to be more easily classified with unpleasant attributes). In this case, it was the more fluently processed target category (Go) that was compatible with pleasant attributes in the IAT.

An IAT effect for Go items may result from two possible influences. Firstly, in accordance with the salience asymmetry account, it may be the case that Go items and pleasant words are compatible because Go items are more salient than Nogo items, and, as Kinoshita and Peek-O'Leary (2005, 2006) claim, pleasant attributes are more salient than unpleasant attributes. However, it may also be the case that the data reflect that Go items are more positive than Nogo items. This may occur because Go items were more fluently processed than Nogo items on the binary classification task, and people prefer items that are fluently processed. For example, there is evidence that stimuli that are identified more

readily are judged to be prettier or more pleasant (Reber, Winkielman & Schwarz 1998; Whittlesea, 1993, Experiment 5), and can also induce positive affect (Winkielman & Cacioppo, 2001). To tease apart the salience and valence theories of the Go/Nogo IAT effect, we conducted a split IAT in Experiment 5b using the same Go and Nogo categories.

3.5. Experiment 5b

In Experiment 5b, a split IAT was conducted to test whether the Go IAT effect found in Experiment 5a was due to a salience asymmetry between Go and Nogo items, or whether it was due to Go items being more positive than Nogo items. If salience asymmetries are responsible for the Go/Nogo IAT effect shown in Experiment 5a, then the Go category should be more easily classified with the more salient pleasant and unpleasant attributes, than with the less salient neutral attributes in a split IAT. In contrast, if the Go IAT effect is due to the Go items being more positive than Nogo items, then participants will classify Go items more easily with pleasant than neutral attributes in the Pleasant IAT, and more easily with neutral than unpleasant attributes in the Unpleasant IAT.

3.5.1. Method

The method was the same as in Experiment 4b, with two exceptions. Firstly, instead of classifying flowers and insects, 32 participants classified the Flight and Teeth stimuli from Experiment 5a. Also, to bring this procedure in line with that of Rothermund and Wentura's (2001, 2004) Go/Nogo IAT, participants performed two Go/Nogo tasks instead of the one Go/Nogo task used in Experiment 5a, and they used either a left-/right- hand response key instead of a spacebar. The first Go/Nogo task occurred at the outset of the experiment. In this task (Task 1), half of the participants were required to press the left

response key ('a'), and half were required to press the right response key ('5') on Go trials. This key assignment was compatible with the key assignment in the combined classification trials of Task 3. The second Go/Nogo task (Task 4) occurred just before the second combined classification task, and this time the response assignment ('a' or '5') for Go trials was reversed to be compatible with the key assignment in the combined classification task of Task 5. The number of trials was doubled from 24 to 48 in the second Go/Nogo task (Task 4) to counteract the carry-over effects of the target response assignment from Task 2 and Task 3 (Nosek et al., 2005).

3.5.2. Results and Discussion

Data reduction

Five participants met the criterion for replacement. For the IAT analysis, reaction times associated with incorrect responses (8.2%), outliers below the minimum cutoff (0.2%), and above the maximum cutoff (0.7%) were rejected from the analyses.

Pleasant IAT

The between groups factor of the category assigned to the Go condition (Flight/Teeth) did not interact with the comparison of interest (F(1,14) = 1.91, p = .19 for the reaction time data, F < 1 for the error data), and thus the data were averaged across both groups. The mean reaction times are presented in the left panel of Figure 3.4. The open bar shows the condition in which Go items shared a key with pleasant attributes and Nogo items shared a key with neutral attributes, and the filled bar shows the condition in which the target and attribute assignment was reversed. Reaction times were faster in the former condition than the latter condition, F(1,15) = 11.52, p < .01. This result replicates the previous IAT effect for Go items observed in Experiment 5a. There were an equal number of errors in both conditions (2.13 for Go Pleasant/Nogo Neutral vs. 2.81 for Go Neutral/Nogo Pleasant), F(1,15) = 1.91, p = .19. This suggests that the IAT effect observed in the reaction time data was not due to a speed-accuracy trade-off.



Figure 3.4. Mean reaction times for the Go/Nogo split IAT of Experiment 5b. The open bars represent the conditions in which Go items were classified with comparatively positive attributes, and Nogo items were classified with comparatively negative attributes. The filled bars represent the conditions in which Go items were classified with comparatively negative attributes, and Nogo items were classified with comparatively positive attributes. Error bars represent the standard error of the mean.

Unpleasant IAT

The between groups factor of the category assigned to the Go condition (Flight/Teeth) did not interact with the comparison of interest (F < 1 for the reaction time

data, F(1,14) = 2.83 for the error data, p = .12), so the data were collapsed across groups for the following analysis. The mean reaction time scores for each combined classification task in the Unpleasant IAT are presented in the right panel of Figure 3.4. The open bar shows when Go was classified with neutral attributes and Nogo with unpleasant attributes, and the filled bar shows the reverse pairings. Responses were faster when Go shared a key with unpleasant attributes and Nogo shared a key with neutral attributes, F(1,15) = 6.76, p < .05. In contrast to the findings of the previous IATs (both Experiment 5a, and the Pleasant IAT in the present experiment), these data show that the Go category is more easily classified with the less pleasant (but more salient) attributes. Responses were also more accurate in the condition in which Go was classified with unpleasant attributes and Nogo was classified with neutral attributes (M_{error} = 2.00), than vice versa (M_{error} = 3.59), F(1,15) = 6.03, p <.05. This pattern of results supports the IAT effect found on the reaction time measure.

Combined analysis of the Pleasant and Unpleasant IATs

The following analysis was conducted with the category assigned to the Go condition (Flight/Teeth) as a between groups factor, because this factor interacted with the salience effect for the error data, F(1,28) = 6.21, p < .05. This factor did not interact with any of the other contrasts of interest on either dependent measure (largest F(1,28) = 1.08, p = .31)

A valence effect in the split IAT was tested by comparing the conditions in which Go was classified with comparatively positive attributes and Nogo with comparatively negative attributes, to the conditions in which the target and attribute combinations were reversed. In Figure 3.4, the open bars were compared to the filled bars. The difference between these two pairs of conditions was 3ms, which was not a significant difference (*F*<1). Thus, there was no overall valence effect for Go items on the split IAT. Nor was there a valence effect in the error data, as an equal number of errors were made in both conditions, F(1,28) = 1.40, p = .25.

A salience effect was examined by comparing the mean reaction time of the conditions in which Go was paired with pleasant and unpleasant attributes and Nogo was paired with neutral attributes, against the remaining two conditions in the split IAT. In Figure 3.4, the outer bars were compared to the inner bars. There was a difference of 170ms between the two pairs of conditions. This was a reliable salience effect, revealing that the Go category was more compatible with the more salient categories (pleasant or unpleasant) than the less salient (neutral) category, F(1,28) = 17.53, p < .001. A similar pattern of results was reflected in the error data, with fewer errors when Go shared a key with the more salient attributes, and Nogo shared a key with neutral attributes, F(1,28) = 8.89, p < .01. However, as noted above, this salience effect interacted with the between groups factor of the category assigned to the Go condition (Flight/Teeth). Analysis of this interaction as a function of the Go category revealed that this salience effect only occurred when Flight was the Go category (F(1,15) = 6.58, p < .05), but that both sets of conditions were equally accurate when Teeth was the Go category (F(1,15) = 1.67, p = .22).

Experiment 5b revealed that Go items are compatible with the more salient attribute categories of pleasant and unpleasant over neutral in the split IAT. The split IAT data also suggest that the effect in Experiment 5a, in which Go was more easily classified with pleasant attributes in a standard IAT, is due to salience asymmetries rather than valence differences. This salience asymmetry is characterized by the more fluently processed target category (Go) being compatible with pleasant attributes in the standard IAT.

3.6. General Discussion

The purpose of this chapter was to investigate whether the split IAT is able to discriminate between IAT effects due to valence differences, and those due to salience asymmetries. In Experiment 4a, I first established the standard IAT effect for flowers over insects. A split IAT with flower and insect target categories in Experiment 4b showed that flowers were compatible with the more positive attributes (pleasant in the Pleasant IAT, and neutral in the Unpleasant IAT), regardless of the salience of the attributes. In demonstrating that the split IAT is able to replicate the classic flower/insect IAT effect, I present evidence that this measure is sensitive to valence differences in the IAT.

In Experiment 5a, salience asymmetries were induced between the target categories using a Go/Nogo manipulation. On a standard IAT, Go items were more easily classified with pleasant attributes, and Nogo items with unpleasant attributes. Experiment 5b used a split IAT to show that the Go category was compatible with the more salient attributes of pleasant and unpleasant over neutral, when compared with the Nogo category. This effect was not influenced by the valence of the attributes, suggesting that the Go/Nogo IAT effect is caused by salience asymmetries and not valence differences. Thus we see that that the split IAT is able to assess salience asymmetries in the IAT.

One trend that emerged was that the patterns of processing fluency between the target categories corresponded with the pattern of salience effects in the split IAT. The target category that was more fluently processed in the binary classification task (Go in Experiment 5a) was also compatible with the more salient attribute categories (pleasant and unpleasant) in the corresponding split IAT (Experiment 5b). This suggests that the more fluently processed target is the more salient category in the IAT. Consistent with this,

94

flowers and insects were processed equally fluently in the binary classification task of Experiment 4a, and produced no salience effect on the corresponding split IAT (Experiment 4b). These results extend the findings of the previous chapter in two ways. Firstly, they show that processing fluency is diagnostic of salience asymmetry effects in a standard IAT with valence attributes. Secondly, they reveal that the more fluently classified category is the more salient of the two target categories.

The salience asymmetry effect in Experiment 5b was characterized by the more fluently processed target category (Go) being compatible with the more fluently processed attribute categories (pleasant and unpleasant) in the split IAT. This result is similar to the results of the IAT experiments in Chapter 2, in which the more fluently processed categories were again compatible with one another (Experiments 2a, 2b and 3). In the standard IAT, the pleasant and unpleasant attribute categories were processed equally quickly, and there was no fluency asymmetry between them. Nevertheless, the more fluently processed Go category was compatible with pleasant attributes in the standard IAT (Experiment 5a). There are two classes of explanation for this effect. One account is based on salience asymmetries, and another based on affective causes. One category of stimuli that is almost certainly salient is the Go category following a Go/Nogo task, as the Go/Nogo manipulation directs participants to focus on that category. Because the Go category was compatible with pleasant attributes over unpleasant attributes in Experiment 5a, this suggests that positive valence is also likely to be the focus of attention. One explanation for why positive valence might attract attention comes from the concept of linguistic markedness, described in Chapter 1 (Section 1.6.2). As mentioned previously, positivity is considered to be the unmarked, basic form, because people will ask, "How good is it?" when evaluating whether a stimulus is good or bad, rather than "How bad is

it?" which implies that a stimulus is already judged to be bad (e.g. Lyons, 1977). This suggests that stimuli are classified using positive as the frame of reference, which is consistent with the idea that positivity is the focal category in the IAT (Kinoshita & Peek-O'Leary, 2005, 2006).

There may also be an affect-based explanation for the compatibility between fluently processed items and positive attributes. Processing fluency may be compatible with positive attributes because people prefer target items that are easily processed. However, there is no direct evidence of this relationship in our results. Although the more fluently processed Go category was compatible with positive attributes in the standard IAT (Experiment 5a), there was no valence effect in the corresponding split IAT (Experiment 5b).

This chapter explored the relationship between processing fluency and salience asymmetry effects in the IAT. It was shown that the target category that was more fluently processed was compatible with pleasant attributes in the standard IAT, due to salience asymmetries. Another source of fluency that has been proposed to mediate IAT effects is familiarity. Some researchers claim that familiarity is compatible with positive attributes in the IAT on the basis of salience asymmetries (Rothermund & Wentura, 2001, 2004; Kinoshtia & Peek-O'Leary, 2005, 2006). Others suggest that this compatibility is caused by a preference for familiarity (Brendl et al., 2001). In this chapter, it was shown that the split IAT can distinguish between the effects of salience and valence in the IAT. Therefore, in the next chapter, I use the split IAT to investigate whether IAT effects involving familiar/unfamiliar target categories are driven by salience asymmetries or valence differences.

Chapter 4. Familiarity in the Implicit Association Test

Chapter 3 had two aims. The first aim was to develop a version of the IAT – the split IAT – that allowed the two contributions of salience and valence to be distinguished from one another. The evidence suggested that the split IAT was sensitive to both valence differences and salience asymmetries between target categories. The second aim was to further test the conclusion of Chapter 2, that fluency is a non-associative contaminant of the IAT. It was found that the more fluently processed target category was compatible with salient attributes in the split IAT, and with pleasant attributes in the standard IAT. One source of processing fluency that may influence IAT effects is familiarity. Familiarity has been hypothesized by Rothermund and Wentura (2001, 2004) and Kinoshita and Peek-O'Leary (2005, 2006) to play an important role in salience asymmetry effects in the IAT. Therefore, the purpose of this chapter is to more directly investigate whether familiarity can mediate IAT effects.

It is critical to clarify the role of familiarity in the IAT, because familiarity has the potential to influence both salience asymmetries and valence. In the past, researchers have examined the role of familiarity in the IAT by using pre-existing target categories that differed in familiarity (e.g. Rothermund & Wentura, 2001, 2004; Kinoshita & Peek-O'Leary, 2006). One pair of categories which are thought to exhibit salience asymmetries due to familiarity are even numbers and odd numbers (Kinoshita & Peek-O'Leary, 2006). Kinoshita and Peek-O'Leary (2006, Experiment 1) demonstrated that in an IAT with even and odd numbers as the target categories, participants were faster to classify even numbers with pleasant attributes, and odd numbers with unpleasant attributes. In the same study, another group of participants performed the flower/insect IAT. Both groups of participants

then rated how favorably they felt toward their respective target categories on a self-report measure. The size of the even/odd number IAT effect was equal to that of the flower/insect IAT, but the results of the self-report measure differed between the two groups. Participants who completed the even/odd number IAT rated the even and odd number stimuli to be equally favorable, but the category of even numbers to be more favorable than the category of odd numbers. Those who completed the flower/insect IAT rated the flower stimuli to be more favorable than the insect stimuli, and the category of flowers to be more favorable than the category of insects. The inconsistent pattern between the IAT effects and the selfreport results led Kinoshita and Peek-O'Leary to claim that the flower/insect IAT effect was driven by evaluative differences, but the even/odd number IAT effect was driven by salience asymmetries. More specifically, they suggested that even numbers are compatible with pleasant attributes because they are the more familiar and/or linguistically unmarked number category.

The evidence cited by Kinoshita and Peek-O'Leary (2006) that even numbers are more familiar than odd numbers comes from a study by Lochy, Seron, Delazer and Butterworth (2000), in which participants were slower to reject a false answer to multiplication equations when the answer was even compared to when it was odd. Lochy et al. (2000) theorized that this effect occurred because even numbers occur three times more often than odd numbers as products in the multiplication tables. In terms of linguistic markedness, even numbers are thought to be unmarked relative to odd numbers because people are faster to classify even numbers using the right-hand response, and odd numbers using the left-hand response, than vice versa; this is known as the linguistic Markedness of Response Codes (MARC) effect (Nuerk, Iversen, & Willmes, 2004). Because the right end of the horizontal dimension is considered to be unmarked (Cho & Proctor, 2005), Nuerk et al. (2004) interpreted the compatibility of even numbers with the right-hand side as indicating that even numbers are similarly unmarked.

To test the hypothesis that salience asymmetries are responsible for the number IAT effect, Kinoshita and Peek-O'Leary (2006, Experiments 3 and 4) created a variant of the IAT in which pleasant and unpleasant attributes were replaced with words and nonwords. In one of their experiments (Experiment 3), the word stimul were valence-neutral words (e.g. item), whereas in another experiment (Experiment 4), the word stimuli were unpleasant words (e.g. evil). In both experiments, responses were faster when even numbers shared a response key with words, and odd numbers shared a response key with nonwords. That is, the familiar target category of even numbers was compatible with the familiar category of words, even when the category consisted of unpleasant words. On the basis of these results in combination with the self-report data showing equal preference for even and odd number stimuli, Kinoshita and Peek-O'Leary (2006) proposed that salience asymmetries are the dominant factor underlying the even/odd number IAT effect. They did not consider that the IAT effect with word/nonword attributes could have been due to preference for familiar items, as they assumed that the classification of words and nonwords is based primarily on the dimension of familiarity, and not valence. Furthermore, even numbers were compatible with the negatively valenced words in their Experiment 4. This result lends some support to the idea that even numbers were not preferred to odd numbers, and thus the IAT effect for even numbers on the word/nonword IAT was not due to evaluative differences.

Based on the assumption that categories differing in familiarity also differ in salience, Rothermund and Wentura (2001, 2004) also used words and nonwords in the same way as Kinoshita and Peek-O'Leary (2006) to detect salience asymmetries in the IAT.

For example, when university undergraduates classified old and young names with words and nonwords in an IAT, they were faster to respond when old names shared a key with unfamiliar nonwords, and young names shared a key with familiar words. Again, their justification for the use of words and nonwords is that these categories differ in salience (based on familiarity) but not in valence. Rothermund and Wentura (2001, 2004) claimed that words and nonwords were equal in valence because they supposed nonwords to be neutral. However, this does not imply that there are no valence differences between the category of words and the category of nonwords. In fact, there is reason to suppose that words may be preferred to nonwords. The mere exposure effect shows that pre-exposed stimuli are rated as being more pleasant than novel stimuli (e.g. Zajonc, 1968). This may be because familiar words are easier to process than unfamiliar words (Balota & Chumbley, 1984; Whaley, 1978), and research has shown that fluently processed items are considered to be more positive than less fluently processed items (Reber et al., 1998; Whittlesea, 1993; Experiment 5; Winkielman & Cacioppo, 2001). These findings imply that the word and nonword categories used to detect salience asymmetries may be confounded by differences in valence or meaningfulness. Rothermund and Wentura (2001, 2004) did attempt to equate word and nonword exemplars for valence in their studies. However, this does not resolve the problem that words may be preferred to nonwords at the category level. For example, people's performance on the IAT may be influenced by whether they prefer words to nonwords in general, rather than by how they evaluate the individual exemplars of each category. This distinction is similar to the argument made by Rothermund and Wentura (2004), and Kinoshita and Peek-O'Leary (2005), that equating the familiarity of exemplars between categories does not mean that the familiarity of the categories themselves is equated (see Chapter 1, Section 1.5.1.).

Thus, it is possible that previous IAT effects that are claimed to be caused by salience asymmetries based on familiarity may actually reflect affective similarity between target categories and words and nonwords. That is, young names may be more easily classified with words, and old names with nonwords, because the young and word categories may be preferred to the old and nonword categories. A similar principle may also apply to IAT results with even/odd numbers when classified with words/nonwords, particularly as participants in Kinoshita and Peek-O'Leary's (2006) study preferred even to odd numbers at the category level.

Therefore, the aim of the current experiments was to clarify the role of familiarity in the IAT. Given that familiarity has been suggested to artificially inflate IAT effects, it is surprising that no studies have experimentally manipulated the familiarity of target categories to examine their effect on the IAT. Instead, researchers have used pre-existing categories whose effects may be obscured by confounds such as pleasantness. To investigate whether familiarity produces salience asymmetry effects or preference effects in the IAT, it would be helpful to have a measure that could distinguish between IAT effects based on valence, and those based on salience. One such measure may be the split IAT. If the split IAT can discriminate between valence differences between flowers and insects (Experiment 4b) and salience asymmetries between Go/Nogo categories (Experiment 5b), then it should be able to reveal the underlying source of compatibility in IATs involving familiar/unfamiliar target categories.

The starting point for an examination of familiarity in the IAT is Brendl et al.'s (2001) insect/nonword IAT effect. The insect/nonword IAT effect was selected because it has been claimed to be caused by a preference for familiarity (Brendl et al.), or salience asymmetries based on familiarity (Rothermund & Wentura, 2001, 2004). A replication of

the standard insect/nonword IAT effect was conducted in Experiment 6a, and an insect/nonword split IAT was conducted in Experiment 6b. However, the results of Experiments 6a-6b may also be influenced by other differences between insects and nonwords, such as valence and/or meaningfulness. Therefore, I also investigated familiarity in the IAT more directly by experimentally manipulating the familiarity of novel target categories in a standard IAT (Experiment 7a), and the split IAT (Experiment 7b). To test whether the findings of Experiment 7a and 7b can be extended to another set of categories considered to exhibit salience asymmetries due to familiarity/linguistic markedness, I conducted an IAT and split IAT with even and odd number target categories (Experiments 8a and 8b respectively). An independent measure of processing fluency was included for all the target category pairings to assess whether processing fluency corresponds with patterns of salience and/or valence effects in the split IAT.

4.1. Experiment 6a

The current study examined whether the insect/nonword IAT effect is due to a preference for insects over nonwords (as hypothesized by Brendl et al., 2001) or to salience asymmetries between the target categories (as proposed by Rothermund & Wentura, 2001, 2004). Because the examination of the insect/nonword IAT effect involved using the split IAT, the basic insect/nonword IAT effect was first replicated using the standard IAT. It was expected that responses would be faster when insects were classified with pleasant words, and nonwords were classified with unpleasant words. Insects and nonwords were also placed in a separate binary classification task to assess whether there were any differences in processing fluency between the two categories. This data was then used for comparison

with split IAT effects in Experiment 6b to assess the relationship between processing fluency and salience asymmetries in the IAT.

4.1.1. Method

Eight participants volunteered for the IAT, and a further 10 participants completed the separate binary classification task. The method was the same as in Experiment 4a (standard flower/insect IAT), except that pronounceable nonwords replaced flowers as one of the target categories. The eight nonword exemplars were *krad*, *rish*, *hocart*, *peshuto*, *krostoak*, *telber*, *mittear*, *nordaloge*. Nonwords were referred to as 'foreign' words instead of nonwords, to avoid nonwords being associated with neutral attributes on the basis of the category labels both beginning with the letter 'n'.

4.1.2. Results

Data reduction

One participant met the replacement criterion. Incorrect responses were discarded from the analysis (5.2% in the IAT, 5.1% in the classification task), as were outliers below the minimum cutoff (none in both cases), and above the maximum cutoff (1.4% in the IAT, 1.5% in the classification task).

IAT analysis

Figure 4.1 shows that the mean reaction time for the task in which insects were classified with pleasant attributes and nonwords were classified with unpleasant attributes was 235ms faster than the reaction time for the task in which the target and attribute assignment was reversed, F(1,7) = 17.67, p < .01. This replicates Brendl et al.'s (2001)

insect/nonword IAT effect. There was no difference between the number of errors made in each condition (1.31 for Insect Pleasant/Nonword Unpleasant vs. 2.00 for Insect Unpleasant/Nonword Pleasant), F(1,7) = 1.36, p = .28. This suggests that the IAT effect seen in reaction times was not due to a speed-accuracy trade-off.



Figure 4.1. Mean reaction times for the insect/nonword IAT of Experiment 4a. The open bar represents the condition in which insects were classified with positive attributes, and nonwords were classified with negative attributes. The filled bar represents the condition in which insects were classified with negative attributes, and nonwords were classified with negative attributes, and nonwords were classified with negative attributes, and nonwords were classified with positive attributes. Error bars represent the standard error of the mean.

Processing fluency of the target categories

In the separate binary classification task, insects were responded to more quickly than nonwords (650ms vs. 706 ms respectively), t(9) = 2.68, p < .05. Equal accuracy in
both conditions suggests that there was no speed-accuracy trade-off ($M_{error} = 3.70$ in both cases), t<1. This shows that words were processed more fluently than nonwords.

4.2. Experiment 6b

Experiment 6b was conducted to test whether the IAT effect for insects found in Experiment 6a was due to a preference for insects over nonwords, or to a salience asymmetry between the two categories. According to Rothermund and Wentura (2001, 2004), the unfamiliar target category of nonwords is more salient than the category of words. If this is the case, then nonwords should be more easily classified with the salient categories of pleasant and unpleasant attributes than with neutral attributes in the split IAT. In contrast, Kinoshita and Peek-O'Leary (2005, 2006) consider that familiar stimuli are more salient in the IAT. They would predict that insects should be more easily classified with pleasant and unpleasant attributes than with neutral attributes. The third possibility, as suggested by Brendl et al. (2001), is that the insect/nonword IAT effect is due to a preference for insects. If this is the case, then participants should find it easier to classify insects with pleasant attributes when compared with neutral attributes in the Pleasant IAT, and with neutral attributes when compared with unpleasant attributes in the Unpleasant IAT.

4.2.1. Method

The method was the same as in Experiment 4b, except that participants classified insects and nonwords as the target categories. Thirty-two participants were tested, and the target stimuli were the same insect and nonwords exemplars as those used in Experiment

6a.

4.2.2. Results and Discussion

Data reduction

Ten participants were replaced for exceeding the maximum error criterion. Reaction times associated with incorrect responses (7.6%), outliers below the minimum cutoff (none in this case), and above the maximum cutoff (0.1%) were rejected from the analyses.

Pleasant IAT

The reaction times in the Pleasant IAT are depicted in the left panel of Figure 4.2. The open bar shows the reaction times for when insects were classified with pleasant attributes and nonwords with neutral attributes, and the filled bar shows the reverse pairings. Responses were faster in the former condition than in the latter condition, F(1,15) = 35.29, p < .001. This result replicates the previous IAT effect for insects observed in Experiment 6a. In addition, fewer errors were committed when insects were classified with pleasant attributes and nonwords were classified with neutral attributes (1.69) than in the other condition (2.53), F(1,15) = 6.36, p < .05. This supports the reaction time data in showing that responding was easier when insects were classified with pleasant attributes, and nonwords were classified with neutral attributes.

Unpleasant IAT

The mean reaction times in the Unpleasant IAT are presented in the right panel of Figure 4.2. Responses were faster when insects shared a key with unpleasant attributes and nonwords shared a key with neutral attributes (the open bar), than in the condition when the pairings were reversed (the filled bar), F(1,15) = 44.25, p < .001. Accuracy was equal in the two IAT classification tasks (M_{error} = 3.34 for Insect Unpleasant/Nonword Neutral vs. 2.28

for Insect Neutral/Nonword Unpleasant), F(1,15) = 3.17, p = .10, suggesting that there was no speed-accuracy trade-off in the reaction time data. In contrast to the findings of the previous IATs (both Experiment 6a, and the Pleasant IAT in the present experiment), these data show that insects are more easily classified with the less pleasant (but more salient) attributes.



Figure 4.2. Mean reaction times for the insect/nonword split IAT of Experiment 6b. The open bars represent the conditions in which insects were classified with comparatively positive attributes, and nonwords were classified with comparatively negative attributes. The filled bars represent the conditions in which insects were classified with comparatively negative attributes, and nonwords were classified with comparatively positive attributes. Error bars represent the standard error of the mean.

Combined analysis of the Pleasant and Unpleasant IATs

A valence effect was tested by comparing the conditions in which insects were classified with comparatively positive attributes and nonwords with comparatively negative attributes, against the conditions in which the combinations were reversed. In Figure 4.2, the open bars were compared to the filled bars. The difference between these two pairs of conditions was 10ms, which was not significant (F<1). This analysis suggests that there is no overall valence effect for either insects or nonwords. Thus, when the Pleasant and the Unpleasant IATs were combined, insects and nonwords were classified equally quickly with the more positive and the more negative attributes. A similar pattern of results occurred in the error data, with the same number of errors made in each pair of conditions (F<1).

A salience effect was tested by comparing the conditions in which insects were classified with pleasant or unpleasant attributes and nonwords with neutral attributes (outer bars), against the conditions in which these combinations were reversed (inner bars). This resulted in a difference of 275ms between the two pairs of conditions. This was a reliable salience effect, revealing that insects were more compatible with the more salient attribute category (pleasant or unpleasant) than the less salient (neutral) category, F(1,30) = 77.82, p < .001. This salience effect was also seen in the error data, with fewer errors occurring when insects shared a key with the more salient attributes, and nonwords shared a key with neutral attributes, F(1,30) = 7.77, p < .01.

In Experiment 6b, insects were more easily classified with the more salient attribute categories (pleasant and unpleasant) in the split IAT than were nonwords. This result supports Rothermund and Wentura's (2001, 2004) suggestion that the insect/nonword IAT effect is due to salience asymmetries, not to a preference for insects over nonwords.

However, contrary to Rothermund and Wentura's claim that the unfamiliar category is more salient, it appears that the familiar insect category is more salient than the nonword category, as insects were compatible with the more salient attribute categories of pleasant and unpleasant in the split IAT. Moreover, in the standard IAT of Experiment 6a, insects were compatible with positive attributes over negative attributes, suggesting that positive attributes may also be more salient. Thus, these results support Kinoshita and Peek-O'Leary's (2005, 2006) proposal that familiar and positive items are more salient in the IAT.

4.3. Experiment 7a

Experiment 6a investigated the effect of familiarity in the IAT with insect and nonword target categories. However, Experiment 6a and previous experiments that have examined this issue used pre-existing target categories that may have confounded familiarity with other factors such as valence and/or meaningfulness. The simplest and cleanest way to test for the effects of familiarity would be to use novel stimuli and categories, and to manipulate familiarity through pre-exposure. This was done in the present experiment using abstract paintings as the novel stimuli, and hypothetical artists as the novel categories. From what is known about the mere exposure effect, it is predicted that the (pre-exposed) familiar items will be compatible with pleasant attributes in the IAT, and unfamiliar items will be compatible with unpleasant attributes. This effect is also predicted by the salience asymmetry account, which claims that familiar items are compatible with pleasant attributes because they are similarly salient. Assuming that familiar items are processed more fluently than unfamiliar items, this prediction is also consistent with the results of the previous chapter, in which the more fluently processed target category was compatible with pleasant attributes, and the less fluently processed target category was compatible with unpleasant attributes. Indeed, in the insect/nonword IAT of Experiment 6a and corresponding split IAT of Experiment 6b, the familiar and more fluently processed category of insect was compatible with pleasant attributes on the grounds of salience asymmetries and not valence, according to the split IAT data. To further test the relationship between processing fluency and IAT effects, a separate binary classification task involving familiar and unfamiliar paintings was included as the independent measure of processing fluency.

4.3.1. Method

The method was similar to that of Experiment 4a (standard flower/insect IAT), with some changes. Sixteen participants volunteered for the IAT, and another 16 participated in the separate binary classification task. The target categories were named after two hypothetical artists, 'Xanthie' and 'Quanto'. These names were based on the artist exemplars 'Xanthie' and 'Quan' used in the IAT studies of Ashburn-Nardo, Voils, & Monteith, (2001). The target stimuli were sixteen abstract paintings, half of which were ostensibly painted by the artist Xanthie, and the other half were ostensibly painted by Quanto. The Xanthie and Quanto paintings were randomly selected for each participant. The familiarity of the target stimuli was manipulated by pre-exposing participants to eight of the target paintings for 128 trials (16 trials of each painting). At the beginning of the preexposure phase, half the participants were informed that they would view paintings by the artist Xanthie, and half were told that they would view paintings by the artist Quanto. Each painting was then presented for 1500ms on each trial. The pre-exposed (familiar) and new (unfamiliar) paintings were then classified in the IAT as Xanthie vs. Quanto paintings. In the IAT instructions, participants were told that paintings belonging to the 'Xanthie' category were painted by the artist Xanthie, and paintings belonging to the 'Quanto', category were painted by the artist Quanto. Because the paintings from one of the categories were novel to participants, participants could only correctly classify the stimuli by discriminating between the paintings they had seen before (e.g. Xanthie's paintings), and the paintings they had not seen before (e.g. Quanto's paintings). To minimize the possibility that repeated exposure to both Xanthie and Quanto paintings in the target classification tasks would weaken the familiarity manipulation, each painting was presented only once in both target classification tasks (Task 1 and Task 4), constituting 16 randomized trials.

The separate binary classification task used to assess the processing fluency of the target categories was similar to that used in Experiment 4a, with some minor exceptions. Prior to classifying the Xanthie and Quanto paintings, participants were pre-exposed to eight paintings 16 times each following the same procedure used in the pre-exposure phase prior to the IAT. To reduce exposure to the unfamiliar paintings so that the pre-exposure effects would not be diluted in the binary classification task, the number of classification trials was halved from 160 to 80, the first 16 of which were practice trials. Each stimulus was presented 10 times.

4.3.2. Results and Discussion

Data reduction

Two participants met the replacement criterion, one in the IAT, and one in the binary classification task. The reaction times of erroneous responses were excluded (4.8 % in the IAT, 3.0% in the classification task), as were those that were less than the minimum

outlier values (none in both cases), or above the maximum outlier values (0.8% in the IAT, 0.2% in the binary classification task).

IAT analysis

Figure 4.3 illustrates the mean reaction times for the conditions in the IAT. There was an IAT effect for familiarity (156ms), with responses being faster when familiar paintings shared a key with pleasant attributes, and unfamiliar paintings shared a key with unpleasant attributes, than when the target and attribute pairings were reversed, F(1,15) = 6.38, p < .05. A similar pattern occurred in the number of errors for the two conditions, with fewer errors when familiar paintings were classified with pleasant attributes, and unfamiliar paintings were classified with unpleasant attributes, and unfamiliar paintings were classified with pleasant attributes, and unfamiliar paintings were classified with unpleasant attributes (2.06), than in the condition in which the target and attribute pairings were reversed (1.03), F(1,15) = 11.74, p < .01. This pattern of errors confirms that the IAT effect in reaction times was not due to a speed-accuracy trade-off.

Processing fluency of the target categories

Responses were faster to familiar paintings than unfamiliar paintings (558ms vs. 583ms respectively), t(15) = 2.42, p < .05. There was no speed-accuracy trade-off, as an equal number of errors was made in both categories (1.19 for familiar paintings vs. 0.75 for unfamiliar paintings), t(15) = 1.39, p = .19. These data reveal that familiar (pre-exposed) paintings were responded to more quickly than unfamiliar paintings.

The IAT result for familiar items may be due to a preference for familiarity, or it may be due to salience asymmetries. Therefore, Experiment 7b was conducted to test which of these causes was the source of the effect.



Figure 4.3. Mean reaction times for the familiar/unfamiliar paintings IAT of Experiment 7a. The open bar represents the condition in which familiar paintings were classified with pleasant attributes, and unfamiliar paintings were classified with unpleasant attributes. The filled bar represents the condition in which familiar paintings were classified with unpleasant attributes, and unfamiliar paintings were classified with pleasant attributes. Error bars represent the standard error of the mean.

4.4. Experiment 7b

Experiment 7b used a split IAT to investigate whether the IAT effect for familiar items observed in Experiment 7a was due to valence differences or salience asymmetries between the target categories. If the familiarity IAT effect shown in Experiment 7a reflects a preference for the more familiar category, then familiar paintings should be compatible with the more positive attributes in a split IAT (pleasant in the Pleasant IAT, and neutral in the Unpleasant IAT). However, if the familiarity IAT effect is due to salience asymmetries,

then the more salient target category (familiar paintings) should be compatible with the more salient attribute categories (pleasant in the Pleasant IAT, and unpleasant in the Unpleasant IAT).

4.4.1. Method

The method was similar to that of Experiment 4b (flower/insect split IAT), except that it used the same Xanthie and Quanto stimuli, and the same pre-exposure method as that used in Experiment 7a. Thirty-two participants took part in the experiment.

4.4.2. Results and Discussion

Data reduction

Seven participants were replaced according to the rejection criterion. Reaction times were eliminated if they were associated with errors (6.3%), or were below the minimum cutoff (none in this case), or above the maximum cutoff (1.1%).

Pleasant IAT

The mean reaction times for each of the two conditions in the Pleasant IAT are presented in the left panel of Figure 4.4. Responses were faster when familiar paintings were classified with pleasant attributes, and unfamiliar paintings were classified with neutral attributes (the open bar), than vice versa (the filled bar) F(1,15) = 18.40, p = .001. Fewer errors were made when familiar paintings and pleasant attributes shared the same key, and unfamiliar paintings and neutral attributes shared the same key (1.25), than vice versa (3.03), F(1,15) = 23.81, p < .001. This replicates the IAT effect for familiar paintings found on the standard IAT.



Figure 4.4. Mean reaction times for the familiar/unfamiliar paintings split IAT of Experiment 7b. The open bars represent the conditions in which familiar paintings were classified with comparatively positively attributes, and unfamiliar paintings were classified with comparatively negative attributes. The filled bars represent the conditions in which familiar paintings were classified with comparatively negative attributes, and unfamiliar paintings and unfamiliar paintings were classified with comparatively negative attributes. The filled bars represent the conditions in which familiar paintings were classified with comparatively negative attributes, and unfamiliar paintings were classified with comparatively positive attributes. Error bars represent the standard error of the mean.

Unpleasant IAT

The right panel of Figure 4.4 presents the mean reaction time data for the two IAT conditions. The open bar shows when familiar paintings were classified with neutral attributes and unfamiliar paintings with unpleasant attributes, and the filled bar shows when the target and attribute combinations were reversed. Participants did not differ in their responses to the two conditions, either in terms of reaction times (F<1), or number of errors

(2.13 for Familiar Neutral/Unfamiliar Unpleasant vs. 1.75 for Familiar

Unpleasant/Unfamiliar Neutral; F < 1). Thus, familiar and unfamiliar paintings were classified equally easily with unpleasant and neutral attributes. This is quite different from the results of the previous IATs presented here that used familiar and unfamiliar paintings as targets (Experiment 7a and the Pleasant IAT in the present experiment).

Combined analysis of the Pleasant and Unpleasant IATs

A valence effect was tested by comparing the reaction times from two pairs of conditions. In one pair of conditions, familiar paintings were classified with comparatively positive attributes and unfamiliar paintings were classified with comparatively negative attributes (the open bars in Figure 4.4). These conditions were compared to the other pair of conditions in which the target and attribute combinations were reversed (the filled bars in Figure 4.4). The difference between these pairs of conditions was 165 ms (F(1,30) = 10.62, p < .01), demonstrating a valence effect for familiar paintings over unfamiliar paintings. A valence effect in the error data indicated greater accuracy when familiar paintings where paired with the more positive attributes, and unfamiliar paintings were paired with the more negative attributes, than vice versa, F(1,30) = 6.15, p < .05. This shows that across the two IATs, familiar paintings were compatible with the more negative attributes, and unfamiliar paintings are more positive than unfamiliar paintings.

A salience effect was examined by comparing the mean reaction time of the conditions in which familiar paintings were paired with (the more salient) pleasant and unpleasant attributes and unfamiliar paintings were paired with (the less salient) neutral attributes, against the conditions in which these pairings were reversed. In Figure 4.4, the

outer bars were compared to the inner bars. The difference between these pairs of conditions was 166ms. This difference was significant, F(1,30) = 10.74, p < .01, indicating that familiar paintings were more easily classified with the more salient attributes overall. There was also a salience effect for familiar paintings in the error data, with greater accuracy when familiar paintings where paired with salient attributes, and unfamiliar paintings were paired with neutral attributes, than vice versa, F(1,30) = 14.45, p < .01. These results suggest that familiar paintings are compatible with the more salient attributes, and, therefore, familiar paintings are more salient than unfamiliar paintings.

The results of Experiment 7b provide evidence that the more fluently processed familiar items are the more salient target category in the IAT. However, familiar items also appeared to be preferred to unfamiliar items in the split IAT. Because familiarity is associated with pleasantness, previous research which has used differentially familiar categories to detect salience asymmetries may instead have produced effects based on affective compatibility. Thus, even numbers and words may be compatible with pleasant attributes (and with each other) because they are preferred, not because there is a salience asymmetry in the IAT. Alternatively, even numbers and words may be both more salient and more preferred to odd numbers and nonwords.

4.5. Experiment 8a

The following experiments examined whether categories that that have been used to investigate salience asymmetries in past IATs actually do exhibit differences in salience, and whether they may also have valence differences. Kinoshita and Peek-O'Leary (2006) claim that even and odd numbers produce an IAT effect that is not evaluative, but based on salience asymmetries due to linguistic markedness and/or familiarity. To test whether

Kinoshita and Peek-O'Leary are correct in these assumptions, even and odd numbers were tested for valence effects and salience asymmetries in the split IAT. Prior to placing these categories in a split IAT, I first replicated the basic even/odd number IAT effect of Kinoshita and Peek-O'Leary (2006, Experiment 1). Even and odd numbers were also placed in a separate binary classification task to examine whether there were any fluency asymmetries between the target categories.

4.5.1. Method

The method was similar to Experiment 4a (flower/insect IAT), except for two minor changes. Twelve participants volunteered for the IAT, and a further 18 participants performed a separate binary classification task. In both tasks, participants classified even and odd numbers as the target categories. The even numbers (*four, eight, twelve, eighteen, twenty-two*) and odd numbers (*five, seven, eleven, thirteen, twenty-one*) were the same stimuli as those used by Kinoshita and Peek-O'Leary (2006), and were matched as closely as possible to each other on length and word frequency. In the binary classification task, participants classified even numbers as the target categories.

4.5.2. Results and Discussion

Data reduction

One participant met the replacement criterion for the IAT. The reaction times of erroneous responses (8.1% in the IAT, 6.1% in the classification task), and those below the minimum cutoff (none in both tasks), or above the maximum cutoff (1.4% in the IAT, 1.2% in the classification task) were discarded from the analysis.

IAT analysis

The mean reaction time scores for the conditions in the combined classification task are presented in Figure 4.5. The open bar shows the reaction times for when even numbers were classified with pleasant attributes and odd numbers with unpleasant attributes, and the filled bar shows reaction times for the reverse target and attribute assignment. There was an IAT effect for even numbers (458ms), with responses being faster when even numbers shared a key with pleasant attributes, and odd numbers shared a key with unpleasant attributes, than vice versa (F(1,11) = 12.05, p < .01). Responses were also more accurate in the former condition than the latter condition ($M_{error} = 2.00$ vs. 3.17 respectively, F(1,11) =5.96, p < .01). This replicates the number IAT effect for even numbers found by Kinoshita and Peek-O'Leary (2006).

Processing fluency of the target categories

In the separate classification task, reaction times to even numbers (636ms) were slightly faster than to odd numbers (658ms), but this difference was not reliable t(17) = 1.45, p = .17. Slightly more errors were made in response to odd numbers (5.00) than to even numbers (3.89), but this was not a reliable result, t(17) = 1.76, p = .10. Thus there were no significant processing fluency differences between even and odd numbers.



Figure 4.5. Mean reaction times for the even/odd number IAT of Experiment 8a. The open bar represents the condition in which numbers were classified with pleasant attributes and odd numbers were classified with unpleasant attributes. The filled bar represents the condition in which even numbers were classified with unpleasant attributes, and odd numbers were classified with pleasant attributes. Error bars represent the standard error of the mean.

4.6. Experiment 8b

Experiment 8a replicated the IAT effect for even numbers first demonstrated by Kinoshita and Peek-O'Leary (2006). This result may indicate that even numbers are more positive than odd numbers, and/or that even numbers are more salient than odd numbers. According to Kinoshita and Peek-O'Leary, even numbers are more salient than odd numbers, and thus should be compatible with the more salient attribute category in the split IAT (pleasant in the Pleasant IAT, and unpleasant in the Unpleasant IAT). In contrast, the fluency account would predict that there are no salience asymmetries between even and odd numbers, because even and odd numbers were processed equally fluently on a binary classification task (Experiment 8a). Based on this reasoning, we would infer that the IAT effect for even numbers indicates that even numbers are more positive than odd numbers. If this is the case, then even numbers should be compatible with the more positive attribute categories in the split IAT.

4.6.1. Method

The method was the same as in Experiment 4b (flower/insect split IAT), except that 32 participants classified the even and odd number stimuli used in Experiment 8a as the target categories.

4.6.2. Results and Discussion

Data reduction

Ten participants met the replacement criterion in this experiment. The reaction times of erroneous responses were discarded (5.8%), as were those that were below the minimum cutoff (none in this experiment), or above the maximum cutoff (1.4%).

Pleasant IAT

The left panel of Figure 4.6 shows the mean reaction time scores for each of the Pleasant IAT conditions. The previous IAT effect for even numbers in Experiment 8a was replicated; responses were faster when even numbers shared a key with pleasant attributes and odd numbers shared a key with neutral attributes (the open bar), than when the target and attribute combinations were reversed (the filled bar), F(1,15) = 11.36, p < .01. There

was no difference in the number of errors committed in the two IAT conditions (1.56. for Even Pleasant/ Odd Neutral vs. 1.81. for Even Neutral/Odd Pleasant), F<1. This result supports the idea that the Pleasant IAT effect in the reaction time data was not due to a speed-accuracy trade-off.



Figure 4.6. Mean reaction times for the even/odd number split IAT of Experiment 8b. The open bars represent the conditions in which even numbers were classified with comparatively positively attributes, and odd numbers were classified with comparatively negative attributes. The filled bars represent the conditions in which even numbers were classified with comparatively negative attributes, and odd numbers were classified with comparatively positive attributes. Error bars represent the standard error of the mean.

Unpleasant IAT

The mean reaction time data for the two Unpleasant IAT conditions are presented in the right panel of Figure 4.6. Participants were quicker to classify even numbers with neutral attributes, and odd numbers with unpleasant attributes (the open bar), than vice versa (the filled bar), F(1,15) = 9.10, p < .01. Again, this result replicates the even/odd number IAT effect of Experiment 8a, suggesting that even numbers are more positive than odd numbers. There was a trend toward fewer errors being committed when even numbers shared a key with neutral attributes, and when odd numbers shared a key with unpleasant attributes (2.28 vs. 1.81 respectively), F(1,15) = 2.74, p = 0.12, indicating that the corresponding reaction time data was not due to a speed-accuracy trade-off.

Combined analysis of the Pleasant and Unpleasant IATs

A valence effect was tested by comparing the conditions in which even numbers were classified with comparatively positive attributes and odd numbers with comparatively negative attributes, against the conditions in which the target and attribute assignments were reversed (in Figure 4.6, the open bars were compared to the filled bars). The difference between these two pairs of conditions revealed a valence effect for even numbers of 180ms (F(1,30) = 19.71, p < .001). This indicates that even numbers were more easily classified with the more positives attributes, and odd numbers were more easily classified with the more negative attributes, when both IATs were combined in the split IAT. There was a trend toward the same effect in the error data, with greater accuracy when even numbers were classified with the more positive attributes and odd numbers were classified with the more negative attributes, than in the other condition (F(1,30) = 3.42, p = .07).

A salience effect was tested by comparing the conditions in which even numbers were classified with the more salient (pleasant and unpleasant) attributes and odd numbers with the less salient (neutral) attributes, against the conditions in which these combinations were reversed. In Figure 4.6, the outer bars were compared with the inner bars. Although responses were 55ms faster when even numbers shared a key with the more salient attributes, this was not a reliable effect, (F(1,30) = 1.83, p = .19). Nor was there an effect in the error data; an equal number of errors was made in both pairs of conditions (F<1).

There is a slight trend toward a salience effect for even numbers, however there appears to be a much larger valence effect for even numbers. The lack of a significant salience effect in the split IAT is consistent with the finding that there was also no fluency asymmetry between the target categories in Experiment 8a. The valence effect for even numbers is contrary to Kinoshita and Peek-O'Leary's (2006) claim that evaluative differences between even and odd numbers are not responsible for the IAT effect. The possible reasons for this discrepancy will be discussed in the General Discussion. However, this valence effect is consistent with Kinoshita and Peek-O'Leary's (2006, Experiment 1) finding that the category of even numbers is rated to be more favorable than the category of odd numbers.

4.7. General Discussion

The purpose of this chapter was to investigate the role of familiarity in the IAT. This was done by first examining whether the insect/nonword IAT effect (Brendl et al., 2001) was due to salience asymmetries caused by familiarity, or to preference for the familiar insect category. Experiment 6a demonstrated that insects were compatible with pleasant attributes in the IAT, compared with nonwords. In the split IAT of Experiment 6b, insects were compatible with the more salient attribute categories in the split IAT, regardless of the valence of the attributes. This suggests that the insect/nonword IAT effect is due to salience asymmetries, rather than to preference for the more familiar target category. However, the results of Experiments 6a and 6b may have been due to pre-existing differences in valence and/or meaningfulness between insects and nonwords.

To control for pre-existing differences between target categories, Experiments 7a and 7b explored familiarity in the IAT by experimentally manipulating the familiarity of novel target categories (familiar/unfamiliar paintings). In Experiment 7a, familiar paintings were more easily classified with pleasant attributes, and unfamiliar paintings were more easily classified with unpleasant attributes in the IAT. Familiar paintings were also responded to more quickly than unfamiliar paintings in a binary classification task, indicating that the pre-exposure manipulation served to increase the processing fluency of the familiar paintings. When attribute valence and salience were manipulated independently in a split IAT (Experiment 7b), familiar paintings were compatible with the more salient attributes (pleasant and unpleasant), and the more positive attributes (pleasant in the Pleasant IAT, and neutral in the Unpleasant IAT). Thus, it would appear that the standard IAT effect seen in Experiment 7a was caused by familiar paintings being more salient, and more positive than unfamiliar paintings. These data suggest that familiarity can contribute to IAT effects through both salience asymmetries and valence.

Experiment 8a replicated the number IAT effect of Kinoshita and Peek-O'Leary (2006), in which even numbers are more easily classified with pleasant attributes, and odd numbers with unpleasant attributes. Even and odd numbers were responded to equally quickly on the binary classification task. In a subsequent split IAT (Experiment 8b), even numbers were compatible with the more positive categories (pleasant in the Pleasant IAT, and neutral in the Unpleasant IAT), and odd numbers were compatible with the more numbers were compatible with the more numbers were compatible with the more positive categories (pleasant in the Pleasant IAT).

This suggests that the previous number IAT effect was largely driven by valence differences between even and odd numbers.

The finding that even numbers are more positive than odd numbers on the split IAT is contrary to Kinoshita and Peek-O'Leary's (2006) assumptions that the even and odd numbers IAT effect is driven by salience asymmetries rather than valence. The experimental evidence provided by Kinoshita and Peek-O'Leary (2006, Experiment 4) to support the assumption that even and odd numbers claim comes from an IAT in which even and odd numbers were classified with (unpleasant) words and nonwords. This study is very similar to the Unpleasant IAT of Experiment 8b, in which even and odd numbers were classified with unpleasant words and neutral words. In Kinoshita and Peek-O'Leary's study, even numbers were compatible with (unpleasant) words over nonwords. By contrast, in the Unpleasant IAT of Experiment 8b, even numbers were compatible with neutral words over unpleasant words. The discrepancy between the results of Kinoshita and Peek-O'Leary's study and the Unpleasant IAT may be due to the different task demands involved. In Kinoshita and Peek-O'Leary's study, when unpleasant words and nonwords are classified as 'words' and 'nonwords', attention may be drawn to the salience rather than to the valence of the categories, because word status is the relevant feature. However, in the split IAT, when 'unpleasant' and 'neutral' serve as the category labels, participants are required to encode the valence of the stimuli. Therefore, valence is likely to exert a stronger influence in the split IAT than in an IAT with word and nonword attributes. Because the attribute categories in the standard IAT are also defined by valence, any valence differences between even and odd numbers may also be highlighted by the task demands of the standard IAT. Thus, although Kinoshita and Peek-O'Leary's data suggest that there are salience asymmetries between even and odd numbers, the results of the even/odd number

split IAT imply that valence differences between even and odd numbers may play a larger role than salience asymmetries in the even/odd number IAT effect.

It is also of interest to note the discrepancy between the results of the familiar/unfamiliar painting split and the insect/nonword split IAT. In particular, the former split IAT showed the familiar painting category to be more positive than the unfamiliar painting category, whereas the latter split IAT did not reveal any valence differences between the familiar insect and the unfamiliar nonword categories. The result of the insect/nonword split IAT may be caused by conflicting valence differences between insects and nonwords. On the one hand, insects may be preferred to nonwords because they are more familiar. On the other hand, insects may also be more negative than nonwords on an associative basis, as suggested by the self-report ratings in Brendl et al.'s (2001) study. These two opposing effects of valence may have cancelled out each other, resulting in the absence of a valence effect on the insect/nonword split IAT. Taken together, the results of the insects/nonword split IAT and the familiar/unfamiliar paintings split IAT suggest that familiar items are both more salient and positive than unfamiliar items in the IAT. However, if the familiar category is also more negative than the unfamiliar category (as in the case of the insect/nonword split IAT), this negativity serves to offset the positivity caused by familiarity, minimizing any valence effects in the IAT.

A similar explanation is given by Kinoshita and Peek-O'Leary (2006) to explain the results of their experiments in which flower and insect targets were classified with 'word' and 'nonword' attributes. In one modified flower/insect IAT, the word category consisted of neutral words, and in another modified flower/insect IAT, the word category consisted of unpleasant words. When flowers and insects were classified with (neutral) words and nonwords, flowers were compatible with words, and insects were compatible with words, and insects were compatible with words.

nonwords. However, when the same target categories were classified with (unpleasant) words and nonwords, the IAT effect with flowers/insects disappeared, as flowers and insects were equally compatible with (unpleasant) words and nonwords. According to Kinoshita and Peek-O'Leary, these effects may have occurred because the positive category of flowers was compatible with the familiar category of words on the basis of salience asymmetries, because the positive and familiar are both salient in the IAT. However, when the familiar attribute category consisted of unpleasant stimuli, this may have introduced an additional valence effect that counteracted the effect of salience asymmetries. That is, in the combined classification condition in which flowers shared a key with (unpleasant) words, and insects shared a key with nonwords, flowers were still compatible with words on the basis of salience asymmetries. However, in the combined classification condition in which flowers shared a key with nonwords, and insects shared a key with (unpleasant) words, insects were compatible with (unpleasant) words on the basis of shared valence. Kinoshita and Peek-O'Leary suggest that these two conflicting compatibility effects cancelled out each other, resulting in the target categories being equally compatible with the unpleasant (word) and nonword attribute categories in the IAT.

Although Kinoshita and Peek-O'Leary (2006) maintain that there are no valence differences between words and nonwords in the IAT, their explanation of the modified flower/insect IAT effect does not rule out the possibility that categories differing in familiarity also differ in valence. That is, it is not certain whether the compatibility between flowers and (neutral) words, and between insects and nonwords, is due to salience asymmetries or valence differences. The valence effect in the familiar/unfamiliar paintings split IAT (Experiment 7b) implies that categories differing in familiarity also differ in valence. Therefore, using categories that differ in familiarity to measure salience asymmetries may introduce valence-based confounds into an IAT. However, the results of the insect/nonword split IAT (Experiment 6b), and Kinoshita and Peek-O'Leary's modified flower/insect IAT, suggest that it may be possible to counteract these confounds by using unpleasant exemplars in the familiar category.

What is most interesting about the present data is that in all split IAT experiments reported in this chapter, categories that differed in processing fluency produced salience asymmetries in the split IAT. Specifically, the more fluently processed category is compatible with the more salient attributes in the split IAT (Experiments 6a and 7a). Where there were no fluency asymmetries between the target categories (even/odd numbers in Experiment 8a), there was also no effect of salience asymmetries in the corresponding split IAT (Experiment 8b). These data are consistent with the results of the flower/insect and Go/Nogo experiments in Chapter 2 (Experiments 4a-6b), and lend further support to the idea that processing fluency is a reliable predictor of salience asymmetry effects in the IAT.

The split IAT with familiar/unfamiliar paintings (Experiment 7b) demonstrated that target categories differing in familiarity can produce valence and salience effects in the IAT. Taken together with the results of the insect/nonword split IAT (Experiment 6b), it appears that the effects of familiarity can interact with pre-existing differences in valence to influence IAT effects. However, the results of the even/odd number split IAT (Experiment 8b) suggest that any effects of salience between categories differing in familiarity/linguistic markedness may also be overshadowed by valence differences between the categories. Therefore, categories differing in familiarity/linguistic markedness may not always produce salience asymmetry effects in the IAT. Because familiarity has been suggested to contaminate IATs evaluating social categories, the next chapter examines whether the effects of familiarity seen in this chapter may also apply to IATs involving social targets.

Chapter 5. Social categories in the Implicit Association Test

In the previous chapter, familiarity between target categories (familiar/unfamiliar paintings in Experiments 7a-7b) was experimentally manipulated to investigate familiarity in the IAT. In these experiments, it was demonstrated that familiar targets are compatible with salient attributes in the split IAT (perhaps through differences in processing fluency), and they are also compatible with the more positive attributes in the split IAT. However, pre-existing target categories differing in familiarity produced inconsistent results on the split IAT. In Experiment 6b, insects and nonwords produced only a salience effect on the split IAT, whereas even and odd numbers produced only a valence effect on the split IAT (Experiment 8b). The present chapter examines whether the effects of salience and/or valence also applies to IATs that assess social categories that differ in familiarity. For instance, Kinoshita and Peek-O'Leary (2005) contend that IAT effects for white people over black people are caused entirely by salience asymmetries, which are driven by greater familiarity with the white ingroup. Thus, they claim that there is no attitude underlying the race IAT. Similarly, Rothermund and Wentura (2001, 2004) claim that salience asymmetries may inflate the age IAT. They argue that although an IAT effect for the familiar young category may be based on shared valence, this effect may be further increased by compatibility between the young category and positive attributes based on salience asymmetries. These arguments, combined with the findings from Experiments 6a-7b, suggest that the potential for familiarity to create or mediate IAT effects is the main problem facing the use of the IAT. This problem threatens to undermine the conclusions of many of the papers published using the IAT.

Previous experiments have tried to rule out the contribution of familiarity in the IAT by controlling for the familiarity of the target exemplars (Rudman et al., 1999; Dasgupta et al., 2000; Ottaway et al., 2001). For example, as discussed in Chapter 1 (Section 1.5.1), Dasgupta et al. controlled for familiarity in one race IAT by using novel photographs as white and black exemplars, and in another race IAT, they matched the familiarity of black and white exemplar names. Under both circumstances, participants still demonstrated an IAT effect for white. Nonetheless, as argued by Rothermund and Wentura (2004), and Kinoshita and Peek-O'Leary (2005), differences in category familiarity may still have influenced these IAT effects.

To investigate the role of familiarity in IATs involving social categories, Rothermund and Wentura (2001, 2004) examined whether social targets produce salience asymmetry effects in the IAT. As discussed in Chapter 4, Rothermund and Wentura (2001, 2004) examined this issue using an age IAT with word and nonword attributes. They found that young names were compatible with words and old names were compatible with nonwords. Based on the reasoning that words and nonwords differ in salience but not valence, this IAT effect led Rothermund and Wentura (2001, 2004) to conclude that salience asymmetries exist between the young and old categories due to familiarity. The present experiments extended upon this approach by assessing the contributions of both salience and valence to IATs involving social targets. This was examined by conducting standard IATs and split IATs with white/black and young/old target categories.

5.1. Experiments 9a-9b

The current series of experiments was conducted to establish standard race IAT effects. To begin with, two standard IATs were conducted using white and black target

categories, one with face stimuli, and another with name stimuli. The processing fluency of these target stimuli was then assessed using a separate binary classification task. In keeping with the race IAT effects of Greenwald et al. (1998) and Dasgupta et al. (2000), it was predicted that there would be an IAT effect for the white category for both face and name stimuli. Previous cross-race binary classification tasks have demonstrated that white participants are quicker to classify faces belonging to other races than faces belonging to their own race (e.g. Levin, 1996; Valentine & Endo, 1992). Levin (1996, 2000) explained that this effect occurs because people from the majority race encode the race of minority members as a feature more so than they encode the race of majority members. That is, people emphasize information specifying race in minority members to a greater extent than they do in majority members. Thus, when participants classify stimuli according to race, the necessary information for the task is more readily available when processing minority members. Based on these findings, it was expected that black faces would be classified more quickly than white faces in the binary classification task. Because previous studies have not directly compared the classification of white names against black names, no specific predictions were made as to whether there would be any processing difference between these two categories on the binary classification task. Although familiar words are generally responded to more quickly than unfamiliar words (Balota & Chumbley, 1984; Whaley, 1978), it may be that like black faces, black names possess a racial feature that makes them easier to categorize than white names.

5.1.1. Method

The method was the same as in Experiment 4a, with some exceptions. There were 8 participants in the race IAT with faces (Experiment 9a), and 12 participants in the race IAT

with names (Experiment 9b). The target categories were white and black, and the participants were all Caucasian. For the race IAT with faces, each target category consisted of 4 female and 4 male grayscale photographs of people, taken from a face database compiled by Minear and D.C. Park (2004). All faces exhibited neutral expressions. The stimuli measured 54mmx74mm, and were presented on a white background. White exemplars were Caucasian faces, and black exemplars were African American faces. All these faces were of people between 18-22 years of age, featuring neutral expressions. For the race IAT with names, the target stimuli were a subset of white and black names taken from Greenwald et al. (1998). They were eight names typical of European Americans (*Courtney, Emily, Stephanie, Megan, Ryan, Justin, Matthew, Adam*) and eight names typical of African Americans (*Jamal, Theo, Leroy, Jerome, Latisha, Shereen, Yvette, Latoya*).

Another set of participants performed a separate binary classification task to assess the processing fluency of the target categories (18 in the race classification with faces, 22 in the race classification with names). These tasks were identical to that used in Experiment 4a, except that participants classified white and black as the target categories.

5.1.2. Results and Discussion

Data reduction

Two participants met the criterion for replacement (one in each of the IATs of Experiments 9a-9b, none in the classification tasks). Incorrect responses were discarded from the analysis (4.9% in the race IAT with faces, 3.5% in classification task with faces, 8.9% in the race IAT with names, 5.3% in the classification task with names). Reaction times were also omitted if they were below the minimum cutoff (none in each of the IATs

and binary classification tasks of Experiments 9a-9b), or above the maximum cutoff (1.4% in the race IAT with faces, 0.9% in the classification task with faces, 0.9% in the race IAT with names, 1.6% in the classification task with names).

Experiment 9a: race IAT with faces

IAT analysis. The left panel of Figure 5.1 depicts the mean reaction time scores for each IAT condition. The open bar shows the condition in which white faces were classified with pleasant attributes and black faces with unpleasant attributes, and the filled bar shows the condition in which the target and attribute assignment was reversed. Responses were faster when white faces shared the same key as pleasant attributes, and black faces shared the same key as unpleasant attributes, F(1,7) = 12.50, p < .01. This indicated an IAT effect for white (191ms). There was no difference in the number of errors between the two conditions (1.25 for White Pleasant/Black Unpleasant vs. 1.88 for White Unpleasant/Black Pleasant), F(1,7) = 1.02, p = .35, suggesting that the corresponding pattern in reaction times was not due to a speed-accuracy trade-off.

Processing fluency of the target categories. Black faces elicited faster responses than white faces in the separate binary classification task (536ms vs. 565ms respectively), t(17) = 3.55, p < .01. Equal accuracy to both types of faces indicated that there was no speed-accuracy trade-off (M_{error} = 2.78 for white faces vs. 2.44 for black faces), t<1. This result replicates Levin's (1996) finding that white participants are faster to classify black faces than white faces.



Figure 5.1. Mean reaction times for the race IATs of Experiment 9a with faces (left panel) and Experiment 9b with names (right panel). The open bars represent the conditions in which white stimuli were classified with pleasant attributes, and black stimuli were classified with unpleasant attributes. The filled bars represent the conditions in which white stimuli were classified with unpleasant attributes, and black stimuli were classified with pleasant attributes. The filled bars represent the conditions in which white stimuli were classified with unpleasant attributes, and black stimuli were classified with pleasant attributes. The filled bars represent the conditions in which white stimuli were classified with unpleasant attributes.

The experiments so far have shown that the fluently processed category is compatible with pleasant attributes in the IAT based on salience asymmetries. Thus, in order for salience asymmetries to increase IAT effect sizes, the category that is compatible with pleasant attributes should also be the more fluently processed target category. Based on this principle, these data suggest that the race IAT effect with faces was not inflated by differences in processing fluency between the target categories. This is because the white category that was compatible with pleasant attributes was less fluently processed than the black category. On the basis of salience asymmetries, therefore, white is compatible with unpleasant attributes.

Experiment 9b: race IAT with names

IAT analysis. The right panel of Figure 5.1 shows the reaction time scores for the IAT conditions in which white names were classified with pleasant attributes and black names were classified with unpleasant attributes (the open bar), and vice versa (the filled bar). Responses were 294ms faster in the former condition, representing an IAT effect for white names, F(1,11) = 16.63, p < .01. There was no difference in the number of errors committed in the two conditions (2.58 for White Pleasant/Black Unpleasant vs. 3.13 for White Unpleasant/Black Pleasant), F<1, which suggests that there was no speed-accuracy trade-off in the reaction time data.

Processing fluency of the target categories. In the separate classification task, black names (665ms) were classified more quickly than white names (689ms), t(21) = 2.13, p < .05. There were an equal number of errors in the two conditions (3.41 for black names vs. 4.27 for white names), t(21) = 1.26, p = .22, which indicates that there was no speed-accuracy trade-off. This reaction time difference suggests that the race IAT effect with names was not magnified by the more fluent processing of the white target category.

Comparison of the race IAT with faces and the race IAT with names. A crossexperimental comparison was conducted to compare the effect sizes between the race IAT with faces (Experiment 9a) and the race IAT with names (Experiment 9b). Although the name stimuli produced a numerically larger race IAT effect than the face stimuli (294ms vs 191ms respectively), this difference did not reach significance F(1,18) = 1.09, p = .31.

Familiar items are often more fluently processed than unfamiliar items. However, the binary classification tasks with white and black categories (Experiments 9a-9b) demonstrate that this relationship is not always reliable. In these experiments, the unfamiliar black category was more fluently processed than the familiar white category. If salience asymmetries are based on processing fluency, then the more fluently processed target category (in the context of a classification task, this would be the black category) would be compatible with the more salient attribute category of pleasant, and the less fluently processed target category of white would be compatible with the less salient attribute category of unpleasant. In this case, salience asymmetries could lead the race IAT effect to underestimate prowhite bias, as the white category would be compatible with pleasant attributes on the basis of valence, but compatible with unpleasant attributes on the basis of salience. This prediction is opposite to that made by Rothermund and Wentura (2001, 2004), and Kinoshita and Peek-O'Leary (2005), who suggest that white should be compatible with pleasant attributes on the basis of salience asymmetries. To tease apart these two interpretations of the race IAT effect, the next experiments placed the same race categories in a split IAT to examine whether the more salient category was the one that was more fluently processed (black) or more familiar (white).

5.2. Experiments 9c-9d

The split IAT was used in Experiment 9c-9d to examine the contribution of valence and salience to the race IAT effects obtained in Experiments 9a-9b. In the binary classification tasks of Experiments 9a-9b, black faces and names were processed more quickly that white faces and names respectively. If familiarity is a source of salience asymmetries (as predicted by Rothermund and Wentura (2001, 2004), and Kinoshita and Peek-O'Leary (2005, 2006)), then the familiar target category of white should be compatible with the more salient attributes (pleasant and unpleasant) in the split IAT. However, if processing fluency is responsible for salience asymmetry effects, then the more fluently processed black category should be compatible with the more salient attributes in the split IAT. The Go/Nogo IAT and Go/Nogo split IAT (Experiments 5a-5b respectively) showed that fluency is predictive of salience asymmetries in the IAT independently of familiarity. Based on this finding, I predicted that the more fluently processed categories of black faces and names should be compatible with the more salient attributes in the split IAT.

5.2.1. Method

The method was the same as in Experiment 4b, except in the following respects. There were 32 participants in race split IAT with faces (Experiment 9c), and 32 participants in the race split IAT with names (Experiment 9d). All participants were Caucasian. Participants classified the same target stimuli from Experiments 9a-9b as belonging to the white and black categories.

5.2.2. Results

Data reduction

Four participants were replaced for exceeding the maximum error criterion (all from the race split IAT with names). Erroneous reaction times were omitted (5.6% in the race split IAT with faces, 7.7% in the race split IAT with names), as were those below the minimum cutoff (none in both race split IATs), and above the maximum cutoff (1.1% in the race split IAT with faces, 1.0% in the race split IAT with names).

Experiment 9c: race split IAT with faces

Pleasant IAT. The upper left panel of Figure 5.2 shows the mean reaction times in the Pleasant IAT, when white faces were classified with pleasant attributes and black faces with neutral attributes (the open bar), and when the assignment of targets and attributes was reversed (the filled bar). Contrary to the race IAT effect obtained in Experiment 9a, participants classified black and white faces equally quickly with pleasant and neutral attributes, F<1. There was also an equal number of errors in both conditions of the Pleasant IAT (1.69 for White Pleasant/Black Neutral, vs. 1.78 for White Neutral/Black Pleasant, F<1). This effect is consistent with the idea that the more fluently processed black faces are compatible with pleasant attributes based on salience asymmetries, and white names are compatible with pleasant attributes on affective grounds. Thus, these two competing effects may have cancelled each other out.

Unpleasant IAT. The mean reaction times in the Unpleasant IAT are shown in the upper right panel of Figure 5.2. In one condition, white faces were paired with neutral attributes and black faces with unpleasant attributes (the open bar), in the other condition the reverse was true (the filled bar). Responses were faster when white faces shared a key with neutral attributes, and black faces shared a key with unpleasant attributes, than vice versa, F(1,15) = 8.14, p < .05. Accuracy was equal in both conditions (M_{error} = 1.59 for White Neutral/Black Unpleasant vs. 2.06 for White Unpleasant/Black Neutral) F(1,15) = 1.07, p = .32, indicating that the pattern of reaction time data was not due to a speed-accuracy trade-off. This result replicates the standard race IAT effect for faces found in Experiment 9a.




Combined analysis of the Pleasant and Unpleasant IATs. A valence effect across both the Pleasant and Unpleasant IATs was tested by comparing the conditions in which white faces were paired with comparatively positive attributes and black faces were paired with comparatively negative attributes, to the conditions in which the target and attribute pairings were reversed. In the upper panels of Figure 5.2, the open bars were compared to the filled bars. There was a difference of 99ms, F(1,30) = 7.89, p < .01, indicating that white faces were compatible with the more positive attributes, and black faces were compatible with the more negative attributes in the split IAT. This demonstrates a valence effect for white faces. An equal number of errors were committed in both pairs of conditions, F(1,30) = 1.09, p = .31.

A salience effect was examined by comparing the conditions in which white faces were classified with the more salient attributes, and black faces were classified with the less salient attributes, to the conditions in which the target and attribute assignments were reversed. In the upper panels of Figure 5.2, the outer bars were compared to the inner bars. Across both the Pleasant and Unpleasant IATs, black faces were classified 88ms faster with the more salient attribute category, that is, pleasant and unpleasant attributes, than with neutral attributes, F(1,30) = 6.19, p < .05. This result demonstrates a salience effect for black faces. There was no difference in the number of errors in both conditions, F<1.

Experiment 9d: race split IAT with names

Pleasant IAT. The mean reaction times in the Pleasant IAT are presented in the lower left panel of Figure 5.2. The open bar shows when white names were classified with pleasant attributes and black names with neutral attributes, and the filled bar shows when the target and attribute assignment was reversed. There was no difference in the reaction

times between these two conditions, F < 1. Nor was there a difference in errors between the conditions (2.44 for White Pleasant/Black Neutral vs. 2.72 for White Neutral/ Black Pleasant), F < 1. This result stands in contrast to the IAT effect for white names obtained in Experiment 9b. However, it replicates the race Pleasant IAT effect for faces obtained in Experiment 9c.

Unpleasant IAT. The lower right panel of Figure 5.2 shows the mean reaction times in the Unpleasant IAT. The open bar represents the condition in which white names shared a key with neutral attributes and black names shared a key with unpleasant attributes, and the filled bar represents the condition in which the combinations were reversed. Responses were faster when white names were paired with neutral attributes, and black names with unpleasant attributes, than vice versa, F(1,15) = 31.19, p < .001. Participants were also more accurate in the former condition than the latter condition (M_{error} = 1.72 vs. 3.03 respectively), F(1,15) = 10.10, p < .01. This pattern of results is in line with the standard race IAT effect found in Experiment 9b.

Combined analysis of the Pleasant and Unpleasant IAT. A valence effect was tested by comparing the conditions in which white names shared a key with comparatively positive attributes and black names shared a key with comparatively negative attributes, against the conditions in which the target and attribute assignments were reversed. In the lower panels of Figure 5.2, the open bars were compared to the filled bars. A difference of 197ms between the two pairs of conditions indicated a significant valence effect for white names, F(1,30) = 16.72, p < .001. This valence effect was also found in the error data; accuracy was greater when white names were classified with the more positive attributes, and black names were classified with the more negative attributes, than vice versa, F(1,30)= 9.61, p < .01. A salience effect was examined by comparing the conditions in which white names were classified with the more salient attributes and black names with the less salient attributes (the outer bars of the lower panels of Figure 5.2), against the reverse conditions (the inner bars of the same panels). Responses were 191ms faster when black names were paired with the more salient attribute category (pleasant/unpleasant), F(1,30) = 15.83, p <.001, revealing a salience effect for black names. There was also a trend towards greater accuracy when black names were classified with the more salient attributes, and white names were classified with neutral attributes, F(1,30) = 4.02, p = .054.

5.2.3. Discussion

In both race split IATs there was a valence effect for the white category, and a salience effect for the black category. Because black stimuli were more fluently processed than white stimuli on the binary classification task, and were more compatible with salient attributes in the split IAT, this provides further evidence that processing fluency is a source of salience asymmetries in the IAT.

What we know from previous work presented here is that salient categories are compatible with pleasant attributes in the standard IAT. This can be seen in Experiments 5a-5b with Go/Nogo categories, and Experiments 6a-6b with insect/nonword target categories. In these experiments, the target that was compatible with salient attributes in the split IAT (Go/insect), was also compatible with pleasant attributes in the corresponding standard IAT. The salient target was also the more fluently processed category on a binary classification task. If we assume that, in general, fluent categories are compatible with positive attributes on the basis of salience asymmetries, then it could be inferred that the more fluently processed black category is, to some degree, also compatible with pleasant attributes over unpleasant attributes. Thus, we can conceive of the standard race IAT effect as reflecting a valence effect for white that is countered by a salience effect for black. However, of course, the compatibility between black and pleasant attributes must be weaker than the compatibility between white and pleasant attributes, because pleasant attributes are compatible with white on the standard IAT. This interpretation contradicts Kinoshita and Peek-O'Leary's (2005) hypothesis that the race IAT effect is at least partly, if not entirely, due to salience asymmetries. The next set of experiments examined whether the effects found with white and black targets could be generalized to IATs involving other social categories. In this case, young and old target stimuli were used, as Rothermund and Wentura (2001, 2004) have claimed that the age IAT effect is inflated by salience asymmetries.

5.3. Experiments 10a-10b

Once again, the aim of these experiments was to examine the effects of salience and valence in an IAT with social categories. The target categories were young and old people previously used by Rothermund and Wentura (2001, 2004) to demonstrate salience asymmetries in the IAT. To begin with, two standard IATs were conducted, one with young and old faces (Experiment 10a), and another with young and old names (Experiment 10b). The processing fluency of the target categories was assessed using a separate binary classification task.

5.3.1. Method

The method followed that of Experiment 4a, with some changes. There were 8 participants in the age IAT with faces, and 8 other participants in the age IAT with names.

Participants ranged between 17-26 years of age. The target categories were young and old people. The stimuli for the age IAT with faces were sixteen 54mmx74mm grayscale photographs taken from the face database compiled by Minear and D.C. Park (2004). Young exemplars ranged between 18-22 years of age, and old exemplars ranged between 71-82 years of age. There were four males and four females in each target category. All the faces were Caucasian in appearance, with neutral facial expressions. The stimuli were presented on a white background. For the age IAT with names, the young stimuli consisted of eight names that were among the 40 most popular female and male names for babies born in 1990 (*Tiffany, Ashley, Danielle, Megan, Ryan, Kevin, Jordan, Jason*) and the old stimuli consisted of eight names that were among the 40 most popular female and male name for babies born in 1930 (*Beverly, Martha, Florence, Dorothy, Arthur, Harold, Walter, Clarence*). Names were taken from http://www.babynames.com.au/search-categories-popular.htm.

The separate binary classification task was the same as that in Experiment 4a, except that participants classified white and black stimuli as the target categories. Ten participants completed the age classification task with faces, and 18 other participants completed the age classification task with names.

5.3.2. Results

Data reduction

Two participants met the criterion for replacement (both in the age IAT with names). Incorrect responses were discarded from the analysis (5.6% in the age IAT with faces, 4.4% in the classification task with faces, 5.0% in the age IAT with names, 5.2% in the classification task with names). I also omitted reactions times that were below the

minimum cutoff (0.1% in the age IAT with faces, 0.6% in the classification task with faces, 0.4% in the age IAT with names, 0.2% in the classification task with names) and above the maximum cutoff (0.8% in the age IAT with faces, 1.2% in the classification task with faces, 1.7% in the age IAT with names, 1.4% in the classification task with names).

Experiment 10a: age IAT with faces

IAT analysis. The mean reaction times appear in the left hand panel of Figure 5.3. In one condition, young faces shared a key with pleasant attributes and old faces shared a key with unpleasant attributes (the open bar), these combinations were reversed in the other condition (the filled bar). There was a difference of 160ms between these conditions, revealing an IAT effect in which young faces were compatible with pleasant attributes and old faces were compatible with unpleasant attributes, F(1,7) = 28.94, p < .001. An equal number of errors in both conditions suggest that there was no speed-accuracy trade-off in either condition (4.38 for Young Pleasant/Old Unpleasant vs. 5.12 for Young Unpleasant), F<1. This result replicates the standard age IAT effect.

Processing fluency of the target categories. In the separate classification task, participants responded to old faces (494ms) more quickly than to young faces (525ms), t(9) = 4.01, p < .01. Responses were equally accurate to both types of stimuli (M_{error} = 3.20 for old faces vs. 3.40 for young faces), t<1, indicating that there was no speed-accuracy tradeoff in the corresponding reaction time data.



Figure 5.3. Mean reaction times for the age IATs of Experiment 10a with faces and Experiment 10b with names. The open bars represent the conditions in which young stimuli were classified with pleasant attributes, and old stimuli were classified with unpleasant attributes. The filled bars represent the conditions in which young stimuli were classified with unpleasant attributes, and old stimuli were classified with pleasant attributes. Error bars represent the standard error of the mean.

Experiment 10b: age IAT with names

IAT analysis. The right panel of Figure 5.3 shows the mean reaction times for the conditions in the combined classification task. In one condition, young names were classified with pleasant attributes and old names with unpleasant attributes (the open bar), and in another condition the target and attribute pairings were reversed (the filled bar). Responses were faster when young names shared a key with pleasant attributes, and old names shared a key with unpleasant attributes, than vice versa, F(1,7) = 19.24, p < .01. This

represents an IAT effect for young names (378ms). Fewer errors were made in the condition in which young names were classified with pleasant attributes and old names with unpleasant attributes (2.25), than vice versa (4.25), F(1,7) 13.18, p < .01. This result shows that the corresponding pattern in reaction times was not due to a speed-accuracy trade-off.

Processing fluency of the target categories. In a separate binary classification task, participants classified young and old names equally quickly (731ms vs. 742ms respectively), t<1, and accurately (M_{error} = 3.56 vs. 3.89 respectively), t<1. This indicates that there are no processing fluency differences between the two categories.

Comparison of the age IAT with faces and the age IAT with names. A comparison of the two age IATs revealed that the age IAT with names (Experiment 10b) yielded an IAT effect of 378ms, which is larger than the IAT effect of 160ms produced by the age IAT with faces (Experiment 10a), F(1,14) = 5.17, p < .05. There was also a trend toward a greater IAT effect in errors for the age IAT with names than the age IAT with faces, F(1,14) = 4.51, p = .052, indicating that the difference in IAT sizes in the reaction time data was not due to a speed-accuracy trade-off.

5.3.3. Discussion

In the age IAT with faces, there was an IAT effect for young, and greater processing fluency of old. The previous findings suggest that the more fluently processed target category is compatible with pleasant attributes on the grounds of salience asymmetries (Experiment 5a-5b with Go/Nogo categories, Experiment 6a-6b with insects/nonwords). If this finding can be generalized, we can infer that the age IAT effect with faces is not inflated by the fluency of old. Thus, it appears that the age IAT effect with faces indicates that young faces are considered to be more positive than old faces. In fact, the IAT effect for young may actually be reduced by the greater processing fluency of old faces. This is because old faces may be compatible with pleasant attributes in terms of salience asymmetries, which may serve to counteract the valence effect for young faces.

There was also an IAT effect for young names over old names. In the binary classification task, young and old names were responded to equally fluently. To the extent that fluency asymmetries correspond with salience asymmetries, this suggests that young and old names are equally salient. This makes it likely that the age IAT effect obtained with names will not be influenced by salience asymmetries. In contrast, the age IAT effect with faces appears to be reduced by the salience of old faces. In support of this hypothesis, the age IAT effect with names (378ms) was significantly larger than the age IAT effect with faces (160ms). These results imply that when the preferred target category (e.g. young) is the one that is classified more slowly (as was the case for young faces), then the size of the IAT effect is decreased.

The aim of the next experiments was to further confirm the relationship between fluency and salience in the age IAT. This was done by conducting split IATs using the same age face and name stimuli, and examining whether fluency asymmetries between target categories correspond with salience effects. This also allowed me to investigate whether the difference in effect size between the age IAT for faces and the age IAT for names may be due to the influence of salience asymmetries in the former IAT but not in the latter IAT.

5.4. Experiments 10c-10d

The age IATs of Experiments 10a-10b both demonstrated an IAT effect for young. Based on this result, in the split IAT we would expect young stimuli to be compatible with pleasant attributes in the Pleasant IAT and neutral attributes in the Unpleasant IAT. In addition, because old faces were more fluently processed than young faces, an age split IAT with faces should produce an overall salience effect for the old category, with old faces being compatible with pleasant attributes in the Pleasant IAT, and unpleasant attributes in the Unpleasant IAT. In contrast, because young and old names were classified equally quickly on the binary classification task, there should no salience effect for either of these categories.

5.4.1. Method

The method was similar to Experiment 4b with flowers and insects, except in the following respects. There were 32 participants in the age split IAT with faces (Experiment 10c), and 32 participants in the age split IAT with names (Experiment 10d). Participants ranged from 17-24 years of age. The same young and old target stimuli from Experiments 10a-10b were used.

5.4.2. Results and Discussion

Data reduction

Nine participants met the replacement criterion (all in the age split IAT with names). Incorrect responses were excluded (5.5% in the age split IAT with faces, 9% in the age split IAT with names), as were those which were below the minimum cutoff (none in

the age split IAT with faces, 0.4% in the age split IAT with names), or above the maximum cutoff (1.4% in the age split IAT with faces, 1.0% in the age split IAT with names).

Experiment 10c: age split IAT with faces

Pleasant IAT. The mean reaction times for the Pleasant IAT are presented in the upper left panel of Figure 5.4. The open bar shows the condition in which young faces were paired with pleasant attributes and old faces were paired with neutral attributes, and the filled bar shows the condition in which the target and attribute pairings were reversed. Young and old faces were classified equally quickly (F(1,15) = 1.05, p = .32) and accurately ($M_{error} = 1.88$ vs. 1.94 respectively, F < 1) with pleasant and neutral attributes. This result differs from the standard age IAT effect, and may have been caused by old being compatible with pleasant attributes on the grounds of salience (because old was more fluently processed), and young being compatible with pleasant attributes on the grounds of valence. In this way, the two effects may have counteracted each other.

Unpleasant IAT. The upper right panel of Figure 5.4 shows the mean reaction times for the Unpleasant IAT. The open bar shows the condition in which young faces were classified with neutral attributes and old faces with unpleasant attributes, and the filled bar shows the condition in which the target and attribute assignment was reversed. Responses were faster when young faces were classified with neutral attributes and old faces were classified with unpleasant attributes, than vice versa, F(1,15) = 47.47, p < .001. This result is consistent with the standard age IAT effect. There was also a trend toward fewer errors when young faces were paired with neutral attributes and old faces were paired with unpleasant attributes (1.34) than when the category pairings were reversed (1.91), F(1,15) =



Figure 5.4. Mean reaction times for the age split IATs of Experiment 10c with faces (upper panels) and Experiment 10d with names (lower panels). The open bars represent the conditions in which young was paired with comparatively positive attributes, and old was paired with comparatively negative attributes. The filled bars represent the conditions in which young was paired with comparatively negative attributes, and old was classified with paired positive attributes. Error bars represent the standard error of the mean.

2.08, p = .17. This suggests that the corresponding reaction time data were not caused by a speed-accuracy trade-off.

Combined analysis of the Pleasant and Unpleasant IATs. An overall valence effect was examined by comparing the conditions in which young faces were paired with the comparatively positive attributes, and old faces were paired with the comparatively negative attributes, against the conditions in which the combinations were reversed. In the upper panels of Figure 5.4, the open bars were compared to the filled bars. Responses were 144ms faster when young faces were classified with the more positive attributes, and old faces were classified with the more positive attributes, and old faces were classified with the more positive attributes, and old faces were classified with the more negative attributes in the split IAT, F(1,30) = 25.17, p < .001. There was no difference in errors between the two pairs of conditions, F(1,30) = 1.25, p = .27. Taken together, these results demonstrate a valence effect for young faces.

A salience effect was tested by comparing the conditions in which young faces were classified with the more salient attributes, and old faces with the less salient attributes, to the conditions in which the target and attribute assignments were reversed. In the upper panels of Figure 5.4, the outer bars were compared to the inner bars. Reaction times were 97ms faster when old faces were classified with the more salient attribute category (pleasant in the Pleasant IAT and unpleasant in the Unpleasant IAT), F(1,30) = 11.52, p < .01. An equal number of errors were made in both pairs of conditions, F < 1. This salience effect for old faces in the reaction time data is consistent with the finding that old faces were responded to more quickly than young faces on the binary classification task of Experiment 10a.

Experiment 5d: age split IAT with names

Pleasant IAT. The lower left panel of Figure 5.4 shows the mean reaction times for each of the Pleasant IAT conditions. The open bar shows the condition in which young names were classified with pleasant attributes and old names were classified with neutral attributes, and the filled bar shows the condition in which these target and attribute pairings were reversed. In keeping with the standard age IAT effect, responses were faster when young names shared a key with pleasant attributes, and old names shared a key with neutral attributes, than vice versa, F(1,15) = 10.61, p < .01. There was no difference in the number of errors committed in the two conditions (2.94 for Young Pleasant/Old Neutral vs. 3.06 for Young Neutral/Old Pleasant, F<1). This result supports the idea that the Pleasant IAT effect in the reaction time data was not due to a speed-accuracy trade-off

Unpleasant IAT. The mean reaction time data for the two Unpleasant IAT conditions are presented in the lower right panel of Figure 5.4. The open bar shows the condition in which young names were classified with neutral attributes and old names with unpleasant attributes, and the filled bar shows the condition in which the target and attribute pairings were reversed. Unlike the results of the standard age IAT, participants were equally quick to classify young and old names with unpleasant and neutral attributes, F(1,15) = 1.02, p = .33. Responses in the two conditions were also equally accurate (M_{error} = 2.50 for Young Neutral/Old Unpleasant vs. 3.25 for Young Unpleasant/Old Neutral, F<(1), F(1,15) = 1.11, p = .31.

Combined analysis of the Pleasant and Unpleasant IAT. A valence effect was tested by comparing the conditions in which young names were paired with comparatively positive attributes, and old names were paired with comparatively negative attributes, to the conditions in which the combinations were reversed. In the lower panels of Figure 5.4, the filled bars were compared to the open bars. There was a difference of 112 ms between the two pairs of conditions, F(1,30) = 8.43, p < .01, showing that young names were more easily classified with the more positive attributes, and old names were more easily classified with the more negative attributes in the split IAT. This indicates a valence effect for young names. An equal number of errors was made in both pairs of conditions, F(1,30) = 1.04, p = .32.

To test for a salience effect, the conditions in which young names were paired with the more salient attributes, and old names were paired with the less salient attributes were compared to the conditions in which the target and attribute assignments were reversed. In the lower panels of Figure 5.4, the outer bars were compared to the inner bars. Participants were, on average, 53ms faster to respond when young names shared a key with the more salient pleasant or unpleasant attributes, compared to when young names shared a key with neutral attributes. However, this was not a reliable difference, F(1,30) = 1.90, p = .18, indicating that there was no overall salience effect in the age split IAT with names. Nor was there any difference in the number of errors between the two pairs of conditions, F<1.

The two age split IATs (faces and names) both produced a valence effect for young, which is consistent with the IAT effect for young obtained on the standard IAT in Experiments 10a-10b. Old faces were responded to more quickly than young faces on a binary classification task, and this resulted in a corresponding salience effect for old faces in the split IAT. In contrast, young and old names were classified equally quickly on a binary classification task, and there was no salience effect in the split IAT with names. These results further support the idea that fluency in a binary classification task predicts salience asymmetries in the split IAT.

5.5. General Discussion

This chapter examined the contribution of valence and salience to IAT effects involving social targets that differed in race or age. In Experiments 9a -9d, it was found that white stimuli were preferred to black stimuli on the standard and split IAT, but that black stimuli were processed more fluently in the separate binary classification task. Consistent with experiments from previous chapters, the more fluently processed category of black was compatible with the more salient attribute categories in the split IAT (Experiments 9c-9d). When the target categories of young and old people were tested using the same procedure, young stimuli were preferred to old stimuli on both the standard IAT and the split IAT (Experiments 10a-10d). However, compatibility between the targets and attributes also depended on the fluency with which the target categories were classified. Old faces were more fluently processed than young faces in the binary classification task of Experiment 10a, producing a salience effect for old faces in Experiment 10c. In contrast, old and young names were processed equally fluently on the binary classification task in Experiment 10b, and there was no salience effect in Experiment 10d. A comparison of the age IAT with faces (Experiment 10a) and the age IAT with names (Experiment 10b) suggested that when the more fluently processed target category (old faces) is not the category compatible with positive attributes in the IAT (young faces), this results in a reduction in the magnitude of the IAT effect; the effects of salience and valence appear to work against one another. Thus the same target categories (young/old) can produce difference sized IAT effects depending on the relative fluency with which the stimuli (faces/names) are processed. More specifically, when the target category that is compatible with pleasant attributes is also the less fluently processed target category in a binary

classification task (young faces), a smaller IAT effect is seen compared to when there are no fluency differences between the target categories.

Rothermund and Wentura (2001, 2004), and Kinoshita and Peek-O'Leary (2005, 2006) proposed that salience asymmetries may artificially inflate IAT effects in which there is a preference for the more familiar category. Contrary to this, the results of the two race IATs (Experiments 9a-9b), and the age IAT with faces (Experiment 10a) suggest that often when the familiar target category (white and young) is positive, salience asymmetries may actually *reduce* IAT effects. However, our findings can be reconciled with the reasoning behind Kinoshita and Peek-O'Leary's claims. When Kinoshita and Peek-O'Leary argue that familiarity is a source of salience asymmetries in the race IAT, they assume that familiar items are more fluently processed. In support of this, they cite evidence of the "other race effect", in which people have poorer recognition for faces belonging to other races compared to faces belonging to their own race (Bothwell, Brigham, & Malpass, 1989; Chiroro & Valentine, 1995). In contrast, our experiments show that white participants actually processed *black* faces more fluently, in that they were quicker to classify black faces than white faces according to their racial label. This result is consistent with research showing that white participants classify black faces faster than white faces (Levin, 1996). Thus Kinoshita and Peek-O'Leary were correct to assume that the more fluent target category is salient in the IAT, they were simply incorrect about which categories were likely to be most fluently processed in the IAT. Although our participants were faster to classify the particular old faces, black faces, and black names used in our experiments, these effects may not be universally applicable to other populations and with other stimulus sets of the same categories. Therefore, to establish whether a particular stimulus set is likely to show a salience asymmetry effect in an IAT, the stimuli should be tested in a split IAT.

In all the experiments that demonstrated salience effects in the split IAT (Experiments 9c-9d and 10c), fluency and familiarity were set against each other, so that the target category that was familiar (white and young) was actually less fluently processed on a binary classification task than the target category that was unfamiliar (black and old). The category that was unfamiliar, but more fluently processed, was compatible with the more salient attributes in the split IAT. This suggests that it is processing fluency, rather than the familiarity of the target category, that determines the contribution of salience to the IAT.

The experiments in this chapter support the previous findings in showing that fluency asymmetries between target categories parallel salience asymmetries, with the more fluently processed target category being compatible with the more salient pleasant and unpleasant attributes in the split IAT. A comparison of the standard IATs involving young and old target categories (Experiments 10a-10b) also suggests that the fluently processed target category is compatible with positive attributes in the standard IAT on the basis of salience asymmetries. This interpretation is consistent with previous experiments showing a similar relationship among processing fluency, standard IAT effects, and split IAT effects (Experiments 5a-5b with Go/Nogo categories, Experiments 6a-6b with insect/nonwords). However, there is one effect in the IAT literature that challenges this assumption, known as the mere acceptance effect. This IAT effect is examined in the next chapter because it is the only experiment that undermines the idea that salience in the IAT is determined by processing fluency.

Chapter 6. Mere acceptance and processing fluency in the Implicit Association Test

The experiments reported thus far have shown a consistent pattern between the processing fluency of the target categories (as assessed by the separate binary classification task), and salience asymmetry effects in the split IAT. The data allow us to say three things about the possible role of salience asymmetries in the IAT. Firstly, if there are no processing fluency differences between the target categories (e.g. flowers and insects in Experiment 4a), an IAT effect is less likely to be contaminated by salience asymmetries. Secondly, if the category that is compatible with *pleasant* attributes in the IAT (e.g. familiar paintings in Experiment 7a) is also the more fluently processed target category, then an IAT effect is likely to be artificially inflated. Thirdly, if the category that is compatible with *unpleasant* attributes in the IAT is the more fluently processed target category (e.g. black faces and names in Experiments 10a and 10b respectively), then this is likely to decrease the size of the resulting IAT effect.

However, there is one IAT effect in the existing literature that appears to contradict these principles. This effect stems from the mere acceptance IAT of C.J. Mitchell (2004). In the mere acceptance IAT, target stimuli are defined according to whether they conform to a specified rule. In one mere acceptance experiment (C.J. Mitchell, 2004, Experiment 1), people were presented with the names of objects that belonged to one of two target categories: Flight or Teeth. Flight objects could fly, but did not have teeth (e.g. airplane); Teeth objects had teeth, but could not fly (e.g. kitten). With a few minor exceptions, the stimuli used were identical to the Flight and Teeth stimuli of the Go/Nogo categories in Experiments 5a-5b. Half the people classified target items as conforming to the Flight/No Flight rule (i.e. whether it is an object that flies or not), and half the people classified the same items as conforming to the Teeth/No Teeth rule (i.e. whether it is an object that has teeth or not). The category that conforms to the rule will be hereafter referred to as the 'Accept' category, and the category that does not conform to the rule will be referred to as the 'Reject' category. It was found that the Accept category was more easily classified with pleasant than unpleasant attributes in the IAT. The mere acceptance effect cannot be due to the pleasantness of the target stimuli themselves, because the Flight/Teeth items are counterbalanced between the Accept and Reject categories, so that the Accept/Reject conditions feature both Flight and Teeth stimuli. Thus, mere acceptance can create apparent preference on the IAT between categories that are comparable in pleasantness.

C.J. Mitchell (2004) proposed a range of possible mechanisms for the mere acceptance effect, including affective or semantic associations between the target and and attribute categories, or non-associative factors such as salience asymmetries. He suggested that participants may have preferred Accept items to Reject items because they found the act of accepting a category to be more pleasant than rejecting it. There may also be a semantic association between 'Accept' and 'pleasant', and between 'Reject' and 'unpleasant', with participants finding it easier to classify the two categories that have a similar meaning together.

A non-associative explanation of the mere acceptance effect is based on the salience asymmetry account (Rothermund & Wentura, 2001, 2004). In the mere acceptance IAT, participants were required to classify target stimuli based on whether they obeyed a particular rule, i.e. as belonging to the categories Flight/No Flight, or Teeth/No Teeth. This strategy would have focused their attention on the rule-conforming 'Accept' category, creating a salience asymmetry between the Accept and Reject categories. Thus, the Accept category may have been compatible with positive attributes because they were similarly salient compared to the Reject category and negative attributes. Under these circumstances, Accept items are likely to be more fluently processed than Reject items because they receive greater attention. If this is indeed the case, then the mere acceptance effect is consistent with a fluency account of salience asymmetries.

The particular mere acceptance effect that is problematic for a processing fluency account of salience asymmetries is one in which C.J. Mitchell (2004, Experiment 2) manipulated the relative processing fluency of the Accept and Reject categories. C.J. Mitchell created two target categories using stimuli that consisted of numbers arranged in three rows of two digits (see Figure 6.1). Half the stimuli had a pair of identical two-digit numbers, such that two rows of the numbers matched (Row Same stimuli, presented on the left-hand side of Figure 6.1). The remaining stimuli had three different digits repeated in each column, so that there were matching numbers in each column (Column Same stimuli, presented on the right-hand side of Figure 6.1). The two categories of stimuli were mutually exclusive, so that the Row Same stimuli did not have matching numbers in each column, and the Column Same stimuli did not have any matching rows. Participants in the Row condition classified the stimuli according to whether they had matching rows (Row Same) or not (Row Different), making Row Same the Accept category, and Row Different the Reject category. Participants in the Column condition classified the stimuli according to whether they had matching columns (Column Same) or not (Column Different), making Column Same the Accept category, and Row Different the Reject category. Therefore, the design of the Row/Column mere acceptance IAT was similar to the mere acceptance IAT with Flight/Teeth categories.

Example Row Same stimulus:	Example Column Same stimulus:
32	11
32	42
12	24

Figure 6.1. Examples of the stimuli used in the mere acceptance experiment of C.J. Mitchell (2004, Experiment 2).

Same response in the Row condition, participants would only have to detect a single match. With reference to the Row Same stimulus in Figure 6.1., participants need only detect a match between the first two rows. However, to make a Column Same response in the Column condition, participants would have to match all aspects of the stimulus, that is, all three numbers in each column must match. With reference to the Column Same stimulus in Figure 6.1., participants must compare all the numbers in the first column to all the numbers in the second column. A Column Different decision is simpler, it only requires a single mismatch to be detected. In support of this principle, participants responded to Row Same stimuli more quickly than to Row Different stimuli in the IAT, but were slower at responding to Column Same than to Column Different stimuli. That is, the Accept response was more fluent in the Row condition, but less fluent in the Column condition. Despite these fluency differences, there was a mere acceptance effect in both the Row and Column conditions. In fact, the mere accept effect was larger in the Column condition than the Row condition. This pattern of results is quite at odds with findings from all the previous studies presented here, showing that the more fluently processed category is compatible with pleasant attributes on the basis of salience asymmetries. Instead we see that in the Column condition, the fluently classified category of Column Different was actually compatible with unpleasant attributes in the standard IAT.

If it is assumed that the mere acceptance effect is due to salience asymmetries, then the mere acceptance effect with Row/Column stimuli challenges the hypothesis that the more fluently processed target category is compatible with pleasant attributes on the grounds of salience asymmetries. Therefore, further research is needed to investigate the causes of this apparent discrepancy. To address this issue, the first step was to examine whether the basic mere acceptance effect is caused by valence differences and/or salience asymmetries. This was done by conducting a mere acceptance IAT, split IAT, and separate binary classification task using Flight/Teeth categories. The same tasks were then applied to the mere acceptance IAT with Row Same and Column Same categories, in which the fluency of the Accept and Reject categories was manipulated.

6.1. Experiment 11a

In preparation for investigating whether the mere acceptance effect is due to valence or salience, I first replicated the mere acceptance effect of C.J. Mitchell (2004, Experiment 1). A standard IAT was conducted in which the categories of Flight (vs. No Flight), and Teeth (vs. No Teeth) were classified with pleasant and unpleasant attributes. Following the results of C.J. Mitchell (2004), I predicted that there would be a mere acceptance effect in which the Accept category would be compatible with pleasant attributes, and the Reject category would be compatible with unpleasant attributes.

6.1.1. Method

The method was similar to that of Experiment 4a, with some modifications. Sixteen participants took part in the mere acceptance IAT, and a further 16 completed a separate binary classification task. Participants classified the same Flight and Teeth stimuli used in

Experiment 5a. They were randomly assigned to either the Flight or Teeth condition. In the Flight condition, target stimuli were classified as belonging to the Flight/No Flight categories, making Flight the Accept category and Teeth the Reject category. The remaining participants in the Teeth condition classified the same stimuli as belonging to the Teeth/No Teeth categories, thereby making Teeth the Accept category and Flight the Reject category.

6.1.2. Results and Discussion

Data reduction

Three participants were replaced in the in the IAT for exceeding the maximum number of errors allowed; there were none in the binary classification task who met this criterion. Incorrect responses were discarded from the analysis (7.1% in the IAT, 5.6% in the classification task), as were outliers that were below the minimum cutoff (none in the IAT nor the classification task), or above the maximum cutoff (0.9% in the IAT, 1.1% in the classification task).

IAT analysis

The between groups factor of the category (Flight/Teeth) assigned to the Accept condition did not interact with the comparison of interest, F < 1, and so the data were averaged across both conditions. The mean reaction times are shown in Figure 6.2. Responses were faster when Accept stimuli shared a key with pleasant attributes and Reject stimuli shared a key with unpleasant attributes, compared to the condition in which the target and attribute pairings were reversed. There was a difference of 487ms between these conditions, indicating a mere acceptance effect, F(1,15) = 15.25, p = .001. This

replicates the mere acceptance effect obtained by C.J. Mitchell (2004). Equal error rates in both conditions ($M_{error} = 1.85$ for Accept Pleasant/Reject Unpleasant, vs. 2.69 for Accept Unpleasant/Reject Pleasant) suggests that there was no speed-accuracy trade-off, F(1,15) = 2.48, p = .14.



Figure 6.2. Mean reaction times for the mere acceptance IAT of Experiment 11a. The open bar represents the condition in which Accept (Flight/Teeth) was classified with pleasant attributes, and Reject (No Flight/No Teeth) was classified with unpleasant attributes. The filled bar represents the condition in which Accept was classified with unpleasant attributes, and Reject was classified with pleasant attributes. Error bars represent the standard error of the mean.

Processing fluency of the target categories

The between groups factor of the category (Flight/Teeth) assigned to the Accept condition did not interact with the comparison of interest (F<1 for the reaction time data,

and F(1,10) = 1.04, p = .33 for the error data), and so the following analyses were averaged across both conditions. When the reaction times to Accept and Reject stimuli in the classification task were compared, participants responded to Accept stimuli (680ms) more quickly than to Reject stimuli (728ms), t(11) = 4.27, p = .001. Responses to both categories were equally accurate ($M_{error} = 4.67$ vs. 3.33 respectively for Accept and Reject), t<1, implying that there was no speed-accuracy trade-off in the corresponding reaction time data. These data show that Accept stimuli were more fluently processed than Reject stimuli. In light of the previous experiments, we would expect there to be a salience effect for the more fluently processed Accept category in a mere acceptance split IAT.

6.2. Experiment 11b

Experiment 11b was conducted to test whether the mere acceptance effect obtained in Experiment 11a was due to the Accept category being more positive than the Reject category, or due to the Accept category being more salient than the Reject category. If the mere acceptance effect reflects an affective/semantic association with the Accept category, then Accept should be compatible with the more positive attributes in the split IAT (i.e. pleasant attributes over neutral attributes in the Pleasant IAT, and neutral attributes over unpleasant attributes in the Unpleasant IAT). However, if the mere acceptance effect is due to salience asymmetries, then the Accept category should be compatible with the more salient attributes, that is, pleasant attributes in the Pleasant IAT, and unpleasant attributes in the Unpleasant IAT.

6.2.1. Method

The method followed that of Experiment 4b, except that 32 participants classified the Flight and Teeth stimuli from Experiment 5a into either Flight/No Flight categories (Flight condition), or Teeth/No Teeth categories (Teeth condition). Participants were randomly allocated to either the Flight or Teeth condition.

6.2.2. Results

Data reduction

No participants met the criterion for replacement in the split IAT. I omitted erroneous responses (9.0%), and outliers that were below the minimum cutoff (0.2%) or those that were above the maximum cutoff (1.3 %).

Pleasant IAT

The mean reaction times for each Pleasant IAT condition are illustrated in the left panel of Figure 6.3. The open bar indicates the condition in which Accept stimuli shared a key with pleasant attributes and Reject stimuli shared a key with neutral attributes, and the filled bar shows the condition with the reverse target and attribute assignment. Because the category (Flight/Teeth) assigned to the Accept condition interacted with the contrast of interest for the error data (F(1,14) = 5.83, p < .05), the data (both reaction times and errors) were analyzed with the Accept category assignment (Flight/Teeth) as a between groups factor.

Overall, responses were faster when Accept stimuli were classified with pleasant attributes, and Reject stimuli were classified with neutral attributes, F(1,14) = 34.32, p < .001, than vice versa. There was a trend toward this effect being greater when Teeth was assigned as the Accept category, than when Flight was assigned as the Accept category, (F(1,14) = 4.28, p = .06). An analysis of simple effects showed a mere acceptance effect

for both the Flight (F(1,7) = 5.67, p < .05) and Teeth (F(1,7) = 42.87, p < .001) conditions.



Figure 6.3. Mean reaction times for the mere acceptance split IAT of Experiment 11b. The open bars represent the conditions in which Accept (Flight/Teeth) was classified with comparatively positive attributes, and Reject (No Flight/ No Teeth) was classified with comparatively negative attributes. The filled bars represent the conditions in which Accept was classified with comparatively negative attributes. Error bars represent the standard error of the mean.

There were also fewer errors when Accept stimuli were paired with pleasant attributes and Reject stimuli were paired with unpleasant attributes (1.56), compared to the other condition (4.00), F(1,14) = 30.68, p < .001. This suggests that the corresponding pattern of reaction times was not due to a speed-accuracy trade-off. When the interaction between the error rate and the Accept category was analyzed in terms of simple effects, the

analysis revealed a significant IAT effect when Teeth was the Accept category (F(1,7) = 31.18, p < .05), but not when Flight was the Accept category (F(1,7) = 4.95, p = .06). Thus, the between-groups results in the error data are in the same direction as the between-groups results in the reaction time data. Overall, although the effect was stronger in one condition (when Teeth was Accept), the Pleasant IAT effect for the Accept category replicates the mere acceptance effect on the standard IAT.

Unpleasant IAT

The right panel of Figure 6.3 shows the mean reaction times for the two Unpleasant IAT conditions. The open bar shows when Accept stimuli were classified with neutral attributes and Reject stimuli with unpleasant attributes, the filled bar shows when the target and attribute pairings were reversed. The between groups factor of the category (Flight/Teeth) assigned to the Accept condition did not interact with the comparison of interest (F < 1 for both the reaction time and error data), allowing the data to be combined across both conditions. Responses were faster when Accept stimuli were classified with unpleasant attributes and Reject stimuli were classified with neutral attributes, than vice versa, F(1,15) = 17.08, p = .001. Fewer errors were made when Accept stimuli were paired with unpleasant attributes and Reject stimuli were paired with neutral attributes (2.63), compared to the other condition (3.38), but this was not a reliable difference, F(1,15) = 2.46, p = .14. Nonetheless, this trend in accuracy does suggest that the corresponding reaction time data was not caused by a speed-accuracy trade-off. Unlike the results of the standard mere acceptance IAT, and the mere acceptance Pleasant IAT, these findings show that the Accept category is compatible with the less positive, but more salient, attribute category.

Combined analysis of the Pleasant and Unpleasant IATs

An analysis with the category (Flight/Teeth) assigned to the Accept condition as a between groups factor did not interact with any the contrasts of interest. The interaction with the valence effect was F(1,28) = 3.80, p = .06 for the reaction time data, and F(1,28) =2.57, p = .12 for the error data. The interaction with the salience effect was F(1,28) = 1.69, p = .20 for the reaction time data, and F(1,28) = 2.57, p = .12 for the error data. Therefore, the data were combined across groups for the following analysis. A valence effect was tested by comparing the conditions in which Accept was paired with the relatively positive attributes, and Reject with the relatively negative attributes (the open bars in Figure 6.3), to the conditions in which the target and attribute combinations were reversed (the filled bars in Figure 6.3). The difference between these conditions was 139ms. Although Accept was more quickly classified with the more positive attributes, and Reject with the more negative attributes, this difference did not quite reach significance, F(1,30) = 3.95, p = .06. However, there was a significant valence effect in the error data; fewer errors were committed when Accept was classified with the more positive attributes, and Reject was classified with the more negative attributes across the Pleasant and Unpleasant IATs, F(1,30) = 5.87, p < .05.

A salience effect was tested by comparing the conditions in which Accept was classified with the more salient attributes (pleasant and unpleasant), and Reject with the less salient attributes (neutral), against the remaining conditions in the split IAT. In Figure 6.3, the outer bars were compared with the inner bars. Reaction times were 470ms faster when Reject stimuli were classified with the more salient attribute category (pleasant in the Pleasant IAT and unpleasant in the Unpleasant IAT). This difference indicates a salience effect for Accept stimuli, F(1,30) = 45.11, p < .0001. The same analysis of the error data

revealed a similar salience effect, with fewer errors when Accept was classified with pleasant/unpleasant attributes and Reject with neutral attributes, than vice versa, F(1,30) = 20.94, p < .001.

6.2.3. Discussion

The results of Experiment 11b suggest that the mere acceptance effect is caused by salience asymmetries between the target categories. The Accept category is more salient than the Reject category, because it was compatible with the more salient attribute category (pleasant and unpleasant) in the split IAT, regardless of the valence of that attribute. In the error data, Accept items were more accurately classified with the more positive attributes over the more negative attributes, implying that participants also found Accept to be more positive than Reject.

The Accept category was also more fluently processed than the Reject category in Experiment 11a. This result is likely to have been caused by participants focusing on the Accept category. However, it is not always the case that Accept stimuli are responded to more quickly than Reject stimuli. As discussed in the introduction of the present chapter, C.J. Mitchell (2004, Experiment 2) used the number stimuli presented in Figure 6.1 to demonstrated that mere acceptance effects occur even when the Accept response is more difficult to make than the Reject response. In fact, there was a larger mere acceptance effect for the Accept category that was processed *less* quickly, which contradicts the fluency account of salience asymmetries. Therefore, to clarify the reasons for this discrepancy, the following experiments were conducted to examine the contribution of valence and salience to this mere acceptance effect.

6.3. Experiment 12a

The present experiments investigated how processing fluency influences the mere acceptance effect. This was done by adopting the categories used by C.J. Mitchell (2004, Experiment 2) to manipulate the processing fluency of the Accept and Reject categories: Row Same/Row Different, and Column Same/Column Different. C.J. Mitchell (2004) also showed that Accept responses were faster (i.e. more fluent) than Reject responses when participants classified stimuli as belonging to the Row Same/Row Different category, but were slower when stimuli were classified as belonging to the Column Same/Column Different category. This effect was demonstrated in the IAT combined classification trials of C.J. Mitchell's study. Responses were faster to Row Same than to Row Different when they were classified with pleasant and unpleasant attributes, and faster to Column Different than Column Same when they were classified with the same attributes. However, C.J. Mitchell did not use a separate measure to evaluate processing speed. Therefore, to verify the fluency asymmetries between the Row and Column categories, I used a separate binary classification task as an independent measure of processing fluency.

6.3.1. Method

The method was similar to Experiment 4a, except that 16 participants classified the number stimuli in Figure 6.1 as belonging to either the Row Same/Row Different categories (Row condition), or Column Same/Column Different categories (Column condition) in the IAT. There were 8 participants in each condition. The number stimuli are presented in Appendix A. In the Row condition, participants were informed that each stimulus consisted of three rows of 2-digit numbers. If two of the 2-digit numbers were the same, then the stimulus belonged to the Row Same (Accept) category. However, if none of

the 2-digit numbers matched, then the stimulus belonged to the Row Different (Reject) category. The Column condition consisted of the same number stimuli, except that participants were informed that each stimulus was made up of 2 columns of 3 single-digit numbers. A stimulus belonged to the Column Same category if each of the three numbers in the first column also appeared in the second column, but presented in a different order. If the numbers in the first column did not match the numbers in the second column, then the stimulus belonged to the Column Different category.

There were 20 participants who performed binary classification task. Ten participants classified the same target stimuli as belonging to the Row Same/Row Different categories. The remaining participants classified the same target stimuli as belonging to the Column Same/Column Different categories.

6.3.2. Results

Data reduction

Two participants were replaced in the IAT due to programming error. Two other participants were replaced according to the maximum error criterion (one in the IAT, and one in the binary classification task). The reaction times of erroneous responses (5.5% in the IAT, and 5.3% in the classification task), and outliers that were below the minimum cutoff (none in the IAT and classification task), or that were above the maximum cutoff (0.01% in the both the IAT and classification task) were excluded from the analysis.

IAT analysis

Row condition. The left panel of Figure 6.4 shows the mean reaction time scores for the combined classification conditions. The open bar indicates the condition in which Row

Same was classified with pleasant attributes and Row Different with unpleasant attributes, and the filled bar shows the condition in which the target and attribute assignment was reversed. An IAT effect for Row Same (204ms) was characterized by responses being faster when Row Same shared a response key with pleasant attributes, and Row Different shared a key with unpleasant attributes, F(1,7) = 10.18, p < .05, than vice versa. The same number of errors was made in both conditions (1.69 for Row Same Pleasant/Row Different Unpleasant vs. 2.25 for Row Same Unpleasant/Row Different Pleasant), F<1, indicating that the corresponding pattern in reaction times was not due to a speed-accuracy trade-off.

Column condition. The mean reaction time scores for the combined classification conditions are presented in the right panel of Figure 6.4. The open bar represents the condition in which Column Same was classified with pleasant attributes and Column Different with unpleasant attributes, and the filled bar represents the condition in which the target and attribute pairings were reversed. Responses were 573ms faster in the former condition than the latter condition, indicating an IAT effect for Column Same, F(1,7) = 14.79, p < .01. The same number of errors was made in both conditions (1.88 for Column Same Pleasant/Column Different Unpleasant vs. 1.19 for Column Same Unpleasant/Column Different Pleasant), F(1,7) = 1.43, p = .27.

Comparison of the Row and Column conditions. A comparison was made of the IAT effect sizes for the two different mere acceptance conditions (Row vs. Column). An interaction effect showed that the mere acceptance effect in the Column condition was significantly larger than in the Row condition, F(1,14) = 5.18, p < .05. The same number of errors occurred in both mere acceptance conditions, F(1,14) = 2.03, p = .18, indicating that there was no speed-accuracy trade-off.



Figure 6.4. Mean reaction times for the Row/Column mere acceptance IAT of Experiment 12a. The left panel shows the condition in which target stimuli were classified according to the Row Same rule, and the right panel shows the condition in which target stimuli were classified according to the Column Same rule. The open bars represent the conditions in which Accept (Row Same/Column Same) was classified with pleasant attributes, and Reject (Row Different/Column Different) was classified with unpleasant attributes. The filled bars represent the conditions in which the target and attribute combinations were reversed. Error bars represent the standard error of the mean.

Processing fluency of the target categories

Row condition. Row Same stimuli (675ms) were responded to more quickly than Row Different stimuli (893ms), t(9) = 5.54, p < .001. Equal accuracy in both conditions ($M_{error} = 2.00$ for Row Same vs. 2.10 for Row Different) suggests that there was no speedaccuracy trade-off in the reaction time data (t < 1). *Column condition.* Column Same stimuli (1367ms) were responded to more slowly than Column Different stimuli (1151ms), t(9) = 4.96, p < .001. Equal error rates in both categories again shows that there was no speed-accuracy trade-off ($M_{error} = 6.10$ for Column Same vs. 5.00 for Row Same, t < 1). This means that across both the Row and Column conditions, the same stimulus set (i.e. Row Same/Column Different) was responded to more quickly than the other stimulus set (i.e. Row Different/Column Same).

6.3.3. Discussion

The mere acceptance effects with Row Same and Column Same stimuli replicate the results obtained by C.J. Mitchell (2004, Experiment 2). The binary classification task showed that when participants classified the target stimuli in the Row condition, responses were faster to the Accept category (Row Same) than to the Reject category (Row Different). However, when the target stimuli were classified in the Column condition, the Reject category (Column Different) was processed more quickly than the Accept category (Column Same). When the two different category pairs were placed in an IAT, there was a larger mere acceptance effect when participants performed the task according to the Column Same rule than the Row Same rule. In this instance, the less fluently processed category (Column Same), showed greater compatibility with pleasant attributes, and the more fluently processed category (Column Different) showed greater compatibility with unpleasant attributes, compared to the Row condition. This data pattern appears to contradict previous results showing that the more fluently processed target category is compatible with pleasant attributes over unpleasant attributes in the IAT. However, there may be valence-related processes underlying the two mere acceptance effects. Therefore,
the next experiment examined the contribution of salience and valence to the Row and Column mere acceptance effects.

6.4. Experiment 12b

To investigate the role of salience and valence in the Row/Column mere acceptance effect, a split IAT was conducted that followed the Row Same rule, and another split IAT was conducted that followed the Column Same rule. Experiment 12a demonstrated that the Column mere acceptance effect is larger than the Row mere acceptance effect. If this difference is caused by salience asymmetries, then we would expect the Column split IAT to show a larger salience effect than the Row split IAT. However, if valence-based factors are responsible for the difference between the two IAT effects, then the Column split IAT should show a larger valence effect than the Row split IAT.

6.4.1. Method

The method was the same as the IAT in Experiment 4b, except in the following respects. Forty-eight participant volunteered for the experiment. Twenty-four participants were assigned to the Row condition, where they classified Row Same/Different target categories. The remaining participants were assigned to the Column condition, where they classified Column Same/Different target categories. Within each of these conditions, half the participants performed a Pleasant IAT, and the other half performed an Unpleasant IAT. The target stimuli were the same as those used in Experiment 12a.

6.4.2. Results

Data reduction

Five participants met the replacement criterion for errors. Reaction times to erroneous responses were excluded (6.20%), as were outliers that were below the minimum cutoff (none in this case), or above the maximum cutoff (0.8%).

Row split IAT

Pleasant IAT. The upper left panel of Figure 6.5 shows the mean reaction times for the Pleasant IAT. The open bar indicates the condition in which Row Same was classified with pleasant attributes and Row Different was classified with neutral attributes, and the filled bar shows the condition with the reverse target and attribute pairings. Responses were faster when Row Same shared a key with pleasant attributes and Row Different shared a key with neutral attributes, F(1,11) = 22.22, p < .01. An IAT effect was also evident in the error data; responses were more accurate when Row Same was classified with pleasant attributes, and Row Different was classified with neutral attributes (1.67), than vice versa (2.79), F(1,11) = 12.79, p < .01. This result replicates the mere acceptance effect in Experiment 11a.

Unpleasant IAT. The mean reaction times are presented in the upper right panel of Figure 6.5. Participants were quicker to classify Row Same with unpleasant attributes, and Row Different with neutral attributes (the filled bar), than vice versa (the open bar), F(1,11) = 5.23, p < .05. There was a trend towards greater accuracy when Row Same was paired with unpleasant attributes and Row Different with neutral attributes ($M_{error} = 1.33$), than in the condition in which the target and attribute combinations were reversed ($M_{error} = 2.21$), F(1,11) = 3.60, p = .08. This suggests that there was no speed-accuracy trade-off in the

reaction time data. Contrary to the standard mere acceptance effect in Experiment 11a, this demonstrates that Row same is compatible with the more salient, but less positive, attribute category.

Combined analysis of the Pleasant and Unpleasant IATs. An overall valence effect was tested by comparing the conditions in which Row Same was classified with the more positive attributes, and Row Different with the more negative attributes, the conditions in which the target and attribute pairings were reversed (in the upper panels of Figure 6.5, the open bars were compared to the filled bars). There was some suggestion of a valence effect for the Row Same category. Participants responded 72ms more quickly when Row Same was paired with the more positive attributes and Row Different was paired with the more negative attributes, but this effect was not significant, F(1,22) = 2.30, p = .14. An equal number of errors was made in the two pairs of conditions, F < 1.

A salience effect was examined by comparing the conditions in which Row Same was classified with the more salient attributes (pleasant and unpleasant) and Row Different with the less salient attributes (neutral) against the conditions in which the target and attribute combinations were reversed. In the upper panels of Figure 6.5, the outer bars were compared to the inner bars. Responses were 465ms faster when Row Same shared a key with the more salient attributes and Row Different shared a key with the less salient attributes, indicating a salience effect for the Row Same category, F(1,22) = 23.79, p < .001. A similar effect was seen in the error data, with fewer errors made when Row Same was classified with more salient attributes, and Row Different was classified with neutral attributes, than vice versa, F(1,22) = 12.84, p < .01. The results of the Row split IAT suggest that Row Same is more salient than Row Different, but that there are no valence differences between the two categories.





Column split IAT

Pleasant IAT. The mean reaction times for each of the combined classification conditions of the Pleasant IAT are shown in the lower left panel of Figure 6.5. Responses were faster when Column Same and pleasant attributes shared a key, and Column Different and neutral attributes shared a key (the open bar), than when the target and attribute assignments were reversed (the filled bar), F(1,11) = 20.57, p < .01. This IAT effect for Column Same was mirrored in the error data; fewer errors were committed when Column Same was classified with pleasant attributes, and Column Different was classified with neutral attributes (1.79), than vice versa (3.13), F(1,11) = 6.07, p < .05. Thus, the Pleasant IAT data replicates the standard mere acceptance effect.

Unpleasant IAT. The lower right panel of Figure 6.5 shows the mean reaction times for each Unpleasant IAT condition. Contrary to the mere acceptance effect of Experiment 12a, participants were just as quick to classify Column Same with neutral attributes, and Column Different with unpleasant attributes (the open bar), as vice versa (the filled bar), F<1. The two conditions also exhibited an equal number of errors (1.71 for Column Same Neutral/Column Different Unpleasant vs. 1.25 for Column Same Unpleasant/Column Same Neutral), F<1.

Combined analysis of the Pleasant and Unpleasant IATs. A valence effect was tested by comparing the conditions in which Column Same was classified with the comparatively positive attributes and Column Different with the comparatively negative attributes, against the conditions in which the combinations were reversed. In the lower panels of Figure 6.5, the open bars were compared to the filled bars. The difference between these two pairs of conditions revealed a valence effect for Column Same of 309ms. This was a reliable difference, F(1,22) = 15.64, p < .01, indicating that Column

Same was more easily classified with the more positive attributes, and Column Different was more easily classified with the more negative attributes in the split IAT. There was no difference in the number of errors made in the two pairs of conditions, F(1,22) = 1.36, p = .26.

A salience effect was tested by comparing the conditions in which Column Same was classified with the more salient (pleasant and unpleasant) attributes and Column Different with the less salient (neutral) attributes, against the conditions in which these combinations were reversed. Responses were 227ms faster when Column Same was classified with the more salient attributes, and Column Different was classified with the less salient attributes, and Column Different was classified with the less salient attributes (the outer bars in Figure 6.5), than vice versa (the inner bars in Figure 6.5), F(1,22) = 8.46, p < .01. This indicates a salience effect for Column Same. There was also a salience effect for Column Same in the error data, with fewer errors being committed when Column Same was paired with more the salient attributes, and Column Different was paired with the neutral attributes (F(1,22) = 5.70, p < .05).

6.4.3. Discussion

In the Row split IAT, participants exhibited a salience effect for the Row Same category, but did not show a valence effect. Considering that Row Same was compatible with pleasant attributes in the standard IAT (Experiment 12a), this suggests that the Row mere acceptance effect was due to the salience of the Row Same category, rather than to Row Same being more positive than Row Different. In contrast, the Column split IAT revealed both a valence effect and a salience effect for the Column Same category. These two sources of compatibility may account for why there was a larger mere acceptance effect for Column Same than for Row Same in Experiment 12a. Specifically, the mere

acceptance effect in the Row condition was due only to the salience of the Row Same category, whereas the mere acceptance effect in the Column condition was driven by both the salience and positivity of the Column Same category.

6.5. General Discussion

The experiments in this chapter examined the influence of processing fluency, salience asymmetries, and valence differences in the mere acceptance effect. Experiment 11a replicated C.J. Mitchell's (2004, Experiment 1) mere acceptance effect showing that the Accept category (Flight/Teeth) was compatible with pleasant attributes and the Reject category was compatible with unpleasant attributes in an IAT. Experiment 11b revealed that the Accept category was compatible with the more salient attribute category (pleasant and unpleasant) in the split IAT, independently of the valence of that attribute category. There was also a marginal valence effect, with Accept being compatible with the more positive attributes in the split IAT on the error data. Accept was also more fluently processed than Reject on a separate binary classification task.

In Experiment 12a, I replicated C.J. Mitchell's (2004, Experiment 2) mere acceptance effect which featured number stimuli that differed in the relative processing fluency of their Accept and Reject categories. In a separate binary classification task, Accept was more fluently processed than Reject in the Row condition, but less fluently processed than Reject in the Column condition. Also, although there was a mere acceptance effect for both the Row and Column conditions, this effect was larger in the Column condition. When both conditions were subject to a split IAT in Experiment 12b, the Accept category was compatible with the more salient attribute categories, demonstrating a salience effect for Accept in both conditions. In addition, the Column condition (but not the Row condition) exhibited a valence effect for the Accept condition, indicating that Column Same is more positive than Column Different. Thus, the larger mere acceptance effect for Column Same in the standard IAT (Experiment 12a) may be due to the Accept category in the Column condition being both more salient *and* positive than the Reject category.

The valence effect for Column Same may be due semantic or evaluative similarity between Column Same and positive attributes (and/or between Column Different and negative attributes). However, it is not obvious why Column Same and positive attributes should be semantically similar, especially given that there was no valence effect in the Row condition. This discrepancy between the Column and Row conditions may have been due to differences in the way the target categories were processed under their respective mere acceptance rules. As mentioned previously, accepting a Column Same stimulus is more difficult than accepting a Row Same stimulus. The Row Same rule requires only that a simple perceptual match be identified. However, in a Column Same stimulus, the different order of the three digits in each column prevents participants from using a strategy based on simple perceptual matching, and thus deeper processing is required. It may be that the more elaborative processing of Column Same stimuli (in comparison to the shallow encoding of Row Same stimuli) increased participants preference for Column Same over Column Different.

The idea that elaborative processing increases pleasantness is illustrated in a study by Phillips (2000). In this study, participants were shown complex image advertisements accompanied by explanatory headlines that differed in how explicitly they explained the visual metaphor presented in the advertisement. In a moderate verbal anchoring condition, the headline provided only a clue to the message of the advertisement, and thus further processing was required on the part of the participant to understand the advertisement. In a complete verbal anchoring condition, the headline completely explained the message of the advertisement, and thus the participant did not have to process the advertisement as deeply to reach comprehension. Liking of the advertisement was greater in the moderate verbal anchoring condition than in the complete verbal anchoring condition. This is suggestive evidence that deeper processing increases the pleasantness of the advertisement. According to Phillips (2000), this effect may have occurred because people enjoy solving puzzles. Thus, it may not be elaborative processing per se that increases pleasantness, but elaborative processing that results in obtaining a solution to a problem. Applying this principle to the Row/Column mere acceptance IAT (Experiment 12a), participants were required to work out whether a stimulus matched the accept rule given. Because Accept stimuli can be considered to be the correct answer to the accept rule, participants may have obtained greater satisfaction from correctly identifying the harder-to-process Column Same stimuli than from identifying the easier-to-process Row Same stimuli.

The notion that processing difficulty can increase the valence discrepancy between Accept and Reject is also consistent with the pattern of valence effects found across the mere acceptance category conditions. Looking at the mean reaction times across the mere acceptance split IATs (Figures 6.3 and 6.5), it appears that, overall, longer reaction times were related to larger valence effects in the split IAT. The longest reaction times occurred in the Column condition of the split IAT (Experiment 12b), which also produced a significant valence effect. The next longest reaction time occurred in the split IAT with Flight/Teeth categories (Experiment 11b), which yielded a marginal valence effect. The shortest reaction times occurred in the Row condition of the split IAT (Experiment 12b), which did not show any valence effect. Thus, it may be that elaborative processing of the stimuli (as indexed by response speed) produced preference for Accept over Reject, resulting in a greater valence effect. In this case, processing difficulty increases preference. Alternatively, it may be that preference increases processing difficulty. That is, preference for the Accept category in addition to salience asymmetries may have increased the difficulty of performing the split IAT, thereby slowing down responses.

Given that ease of processing is also associated with liking (Reber et al., 1998; Whittlesea, 1993; Winkielman & Cacioppo, 2001), it seems contradictory to claim that processing difficulty can also engender preference for Accept. However, previous studies showing that processing ease produces liking have only done so by varying how people process the physical properties of the stimuli (e.g. through priming, manipulating figure/ground contrast). In contrast, it may be that the difficulty of conceptual processing can increase liking when it is associated with a correct response, as demonstrated by Phillips (2000). In the mere acceptance IATs presented in this chapter, it is possible that classifying stimuli according to a rule involves some form of conceptual processing, which may explain why processing difficulty may have increased liking of the Accept category.

In the standard IATs of Experiments 11a and 12a, the salient Accept category was compatible with pleasant attributes over unpleasant attributes. This effect is unlikely to be entirely due to the positivity of the Accept category, because there was an unreliable relationship between the Accept category and valence in the corresponding split IATs. Accept was strongly compatible with the more positive attributes in the Column condition, weakly compatible with the more positive attributes in the Flight/Teeth condition, and minimally compatible with more the positive attributes in the Row condition. These effects suggest that any contribution of valence to the mere acceptance effect is variable and slight. The influence of salience asymmetries, in contrast, is consistent across three of these stimulus types. Therefore it appears that the mere acceptance effect is largely driven by salience asymmetries.

Accept may be more salient than Reject because, from the outset, Accept is assigned as the referent target category in the IAT. This makes Accept the target category that participants are likely to focus on throughout the task. Accept was also the more fluently processed category in the mere acceptance binary classification task with Flight/Teeth categories (Experiment 11a), which suggests that attention can lead to processing fluency. However, the Accept category is not always more fluently processed than the Reject category. In the Column condition of the mere acceptance IAT (Experiment 12a), the Reject category (Column Different) was more fluently processed than the Accept category (Column Same). In the same condition of the corresponding split IAT (Experiment 12b), there was a salience effect for the less fluently processed Accept (Column Same) category. This result contradicts previous findings showing that the more fluently processed target category is salient in the split IAT. The possible reasons for this will be discussed in the following chapter.

This chapter investigated the underlying mechanisms of the mere acceptance effect. Salience asymmetries were shown to be the primary factor responsible for the mere acceptance effect, with the Accept category being more salient than the Reject category. Of particular interest was the Row/Column mere acceptance effect, because it produces results which are discrepant with a fluency account of salience asymmetries. In the Row/Column mere acceptance IAT (Experiment 12a), there was a larger mere acceptance effect in the Column condition than in the Row condition, even though Accept was the less fluently processed target category in the Column condition, but the more fluently processed target category in the Row condition. This difference in IAT effect size may have been caused by the Accept category being the more positive target category in the Column condition, but not the Row condition, as suggested by the results of the corresponding split IAT (Experiment 12b) The findings of the Row/Column experiments also imply that processing fluency (as measured on a binary classification task) is not a perfect predictor of salience in the IAT. This requires a revision of the fluency account of salience asymmetries, which will be discussed in the next chapter.

Chapter 7. General Discussion

The experiments in this thesis were conducted to investigate non-associative contaminants of the IAT. More specifically, the aim was to determine the predictors of salience asymmetries in the IAT, to clarify the mechanism/s behind them, and to ascertain the impact of these factors on IAT effects. The present chapter will discuss these issues, and then evaluate the implications of salience asymmetries for the wider IAT literature.

7.1. Processing fluency as a predictor of salience asymmetry

The three broad predictors of salience asymmetries in the IAT that have been proposed are search asymmetries (Rothermund & Wentura, 2001, 2004), linguistic markedness (Kinoshita & Peek-O'Leary, 2005, 2006), and processing fluency. The results of the experiments reported here provide evidence that processing fluency, rather than search asymmetries or linguistic markedness, is a predictor of salience asymmetry effects in the IAT.

In Chapter 2 'Determinants of salience asymmetry: visual search asymmetries versus processing fluency', I explored whether visual search asymmetries or differences in classification fluency better predicted salience asymmetry effects in the IAT. Two sets of categories were created to test this issue. In one set of categories (upright/ inverted elephants), the category that was more fluently classified on a binary classification task was detected less quickly on a visual search task than the other category (Experiments 1a and 1c). In contrast, in another set of categories (big/small cows), the more fluently classified category (big cows) was detected more quickly on a visual search task than the less fluently

processed category (small cows; Experiments 1b and 1d). When both of these sets of categories were classified in an IAT with words and nonwords (Experiments 2a and 2b), it was found that the animal categories that were more fluently classified on a binary classification task (upright elephants and big cows) behaved similarly in the IAT. That is, the more fluently processed categories were compatible with words over nonwords in the IAT. This result is inconsistent with an account of salience asymmetries in terms of visual search asymmetries, because that account claims that categories that exhibit similar visual search asymmetries (e.g. inverted elephants and big cows) should be similarly salient in the IAT. In addition, when the two elephant categories and the two cow categories were classified together in an IAT (Experiment 3), compatibility effects were along the lines of fluency asymmetries, rather than visual search asymmetries. That is, the more fluently classified categories (upright elephants and big cows) were compatible, and the less fluently classified categories (inverted elephants and small cows) were compatible in the IAT. If it is assumed that these IAT effects are salience-based, these data suggest that salience asymmetry effects correspond with fluency asymmetries (as indexed on a binary classification task), rather than with visual search asymmetries. The assumption of salience asymmetries is only critical in the experiments where fluency and search asymmetries were set in opposition, which were those involving the cow categories. The IAT effects involving the cow categories were unlikely to have an affective basis, because there was no evaluative difference between big and small cows on a self-report measure.

In Chapter 3 'Distinguishing between the effects of salience and valence in the IAT: the development of a new procedure', I investigated the relationship between processing fluency and salience asymmetries in a standard IAT with positive and negative attributes. The split IAT was trialed as a method of discriminating between the contribution of valence and salience asymmetries in the IAT. This consisted of a measure in which the valence and salience of the attribute categories were independently manipulated. Thus, if there are valence differences between the target categories, the more positive target would be compatible with the more positive attributes in the split IAT. By contrast, if there are salience asymmetries between the target categories, then the more salient target would be compatible with the more salient attributes in the split IAT. It was established that the split IAT was able to detect valence differences between flower and insect categories (Experiment 4b), and salience asymmetries between Go and Nogo categories (Experiment 5b). In Experiments 5a-5b, the category that was more fluently processed on a binary classification task (Go), was compatible with the more salient attributes in the corresponding split IAT. This suggests that the fluently processed category is the more salient target in the IAT. The more fluently processed category was also compatible with pleasant attributes in the standard IAT, on the basis of salience asymmetries. This implies that pleasant attributes are the more salient attribute category in the IAT. Moreover, when there were no fluency differences between the target categories (as with flowers and insects in Experiment 4a), there was no salience effect in the corresponding split IAT. These studies corroborate those from Chapter 2 in suggesting that processing fluency is a predictor of salience asymmetries in the IAT.

In Chapter 4 'Familiarity in the Implicit Association Test', I examined whether familiarity produced valence and/or salience effects in the IAT. This issue was investigated in Experiments 6a-6b by examining whether the insect/nonword IAT effect was due to valence or salience differences between the familiar insect category and the unfamiliar nonword category. The results of the split IAT suggested that the insect/nonword IAT effect was due to salience asymmetries rather than to valence, with familiar insects being more salient than unfamiliar nonwords (Experiment 6b). However, this result may have been influenced by pre-existing differences in valence or meaningfulness between insects and nonwords. Therefore, Experiment 7a-7b controlled for pre-existing differences between target categories by experimentally manipulating the familiarity of the target categories (familiar/unfamiliar paintings). In the familiar/unfamiliar paintings split IAT, there was both a salience effect and a valence effect for the familiar category, suggesting that the familiar category was more salient and positive than the unfamiliar category. Experiments 8a-8b explored the same issues with even and odd number target categories claimed by Kinoshita and Peek-O'Leary (2006) to exhibit salience asymmetries due to familiarity. There was a valence effect for even numbers on the split IAT, but no salience effect for either category. Thus, it appears that the familiar category is sometimes more positive (familiar paintings in Experiment 7b, and even numbers in Experiment 8b), and sometimes more salient (insects in Experiment 6b, and familiar paintings in Experiment 7b) than the unfamiliar category. Once again, a reliable predictor of salience asymmetries was processing fluency. The target category that was more fluently processed in a binary classification task (insects in Experiment 6a, and familiar paintings in Experiment 7a) was also more compatible with the more salient attributes in the corresponding split IAT (Experiments 6b and 7b respectively). When both the target categories were processed equally quickly (even and odd numbers in Experiment 8a), there was no salience effect on the corresponding split IAT (Experiment 8b).

The relationship between processing fluency and salience asymmetries was also supported in Chapter 5 'Social categories in the Implicit Association Test'. Experiments 9a-9d tested white and black target categories using face and name stimuli. For both faces and names, the unfamiliar black category was actually processed more fluently than the familiar white category on a binary classification task (Experiment 9a-9b). In the corresponding split IATs (Experiments 9c-9d), there was a valence effect for white, but a salience effect for black. Experiments 10a-10d tested young and old categories using face and name stimuli. Old faces were processed more fluently than young faces on a binary classification task (Experiment 10a). The results of the corresponding split IAT (Experiment 10c) showed a valence effect for young faces, and a salience effect for old faces. For name stimuli, both young and old were processed equally quickly on a binary classification task (Experiment 10b). On the split IAT with young and old names (Experiment 10d), there was a valence effect for young, but no salience effect for either category. The experiments in Chapter 5 are notable because they show that when the familiarity and fluency of the target categories are set in opposition (e.g. white is familiar, but black is fluent), processing fluency on a binary classification task is a better predictor of salience asymmetries than is familiarity.

Taken together, the results of the experiments in Chapters 4-5 provide evidence against the linguistic markedness account of salience asymmetries. According to Kinoshita and Peek-O'Leary (2005, 2006), the unmarked category (such as even numbers compared to odd numbers, and familiar items compared to unfamiliar items) is the more salient category in the IAT. However, there was no salience effect for even numbers in the split IAT of Experiment 8b, and there was an unreliable relationship between familiarity and salience in other experiments. On the one hand, there was a salience effect for the *familiar* category in the insect/nonword split IAT (Experiment 6b) and the familiar/unfamiliar painting split IAT (Experiment 7b). On the other hand, there was a salience effect for the *unfamiliar* category in the race split IATs (Experiments 10c-10d) and the split age IAT with faces (Experiment 10c). Lastly, there was no salience effect in the age split IAT with names (Experiment 10d). Thus, it appears that salience asymmetry effects in the IAT do not correspond with the linguistic markedness of the target categories.

So far, the pattern of data across the experiments has shown that the more fluently processed category in a binary classification task is also the more salient target category in the IAT. There was, however, one exception to this rule in Chapter 6 'Mere acceptance and processing fluency in the Implicit Association Test'. This chapter investigated the contribution of valence and salience to the mere acceptance effect. This was first examined by using Flight/Teeth as the Accept and Reject categories in Experiments 11a-11b. Experiment 11a showed a mere acceptance effect, and showed that the Accept category was processed more quickly than the Reject category in a binary classification task. In the mere acceptance split IAT (Experiment 11b), there was a salience effect for the Accept category, and a marginal valence effect for the Accept category (as shown by the error data). This result is consistent with the fluency account of salience asymmetries. The relationship between processing fluency and the mere acceptance effect was further explored in Experiments 12a-12b. In one pair of target categories (the Row condition), the Accept category was processed *more* fluently than the Reject category. In another pair of target categories (the Column condition), the Accept category was processed less fluently than the Reject category. In the split IATs of Experiment 12b, there was a salience effect for the Accept category in both the Row and Column conditions, regardless of the relative processing fluency of those categories. Thus, we see that processing fluency does not always predict salience asymmetries in the IAT. The possible reasons for this will be discussed later in this chapter.

The results of the studies presented here suggest that, in general, processing fluency predicts salience in the IAT. The more fluently processed category appears to be the more

salient target, as it is compatible with the more salient attributes in the split IAT. This pattern is consistent across manipulations of visual search asymmetries, familiarity/linguistic markedness, and valence asymmetries between the target categories. It also seems that the salient target category is compatible with pleasant attributes in the standard IAT. Before we can accept this analysis however, we must first evaluate the assumptions on which these conclusions are based.

7.2. Evaluating the assumptions of the split IAT

An interpretation of the split IAT in terms of valence and salience effects is based on three assumptions that deserve further examination. The first assumption is that the salience asymmetry between the attribute categories in the Pleasant IAT is equal in magnitude to the salience asymmetry between the attribute categories in the Unpleasant IAT. This would mean that pleasant and unpleasant attributes are equally salient when compared with neutral attributes. The second assumption is that the valence difference between the attribute categories in the Pleasant IAT is the same as the valence difference between the attribute categories in the Unpleasant IAT. This would mean that neutral falls halfway between positive and negative on the valence dimension. Addressing these issues will allow us to evaluative whether it is valid to conduct a combined analysis of the split IAT to obtain separate measures of valence and salience. The third assumption is that the split IAT measures the effects of valence and salience as they would occur in the standard IAT. Evaluating this assumption is critical in determining whether the split IAT is an appropriate measure of valence and salience in the standard IAT.

7.2.1. The assumption of equal salience

The combined analysis of the split IAT rests on the assumption that the salience asymmetry between pleasant and neutral attributes is the same size as the salience asymmetry between unpleasant and neutral attributes. However, this assumption may be incorrect. Because pleasant attributes appear to be more salient than unpleasant attributes in the IAT, it may also be that pleasant attributes are more salient than unpleasant attributes when both are compared to neutral attributes in the split IAT. It is possible to reconcile the assumptions of unequal salience in the standard IAT and equal salience asymmetries in the split IAT if we consider that the way salience asymmetries are determined between attributes in the standard IAT differs from how they are determined between attributes in the split IAT. This is because pleasant and unpleasant attributes in the standard IAT may differ primarily along one dimension (e.g. linguistic markedness), whereas valence and neutral attributes in the split IAT may differ primarily along another dimension (e.g. valence extremity). Therefore, it is possible for pleasant attributes to be more salient than unpleasant attributes, but for the two attributes to share the same degree of salience relative to neutral attributes.

In evaluating the assumption of equal salience asymmetries in the split IAT, it would be useful to examine how alternative accounts (of unequal salience asymmetries) would interpret split IAT effects. To consider an extreme case, what would happen if pleasant attributes are more salient than neutral attributes, but unpleasant and neutral attributes are equally salient? This would mean that pleasant attributes were more salient than unpleasant attributes in the split IAT. Let us apply this interpretation to the Go/Nogo split IAT in Experiment 5b, in which the Go category was equally compatible with both the

pleasant and unpleasant attributes over neutral attributes. Based on the assumption that unpleasant and neutral attributes are equally salient, this would make Go more negative than Nogo (as suggested by the results of the Unpleasant IAT), and Go more salient than Nogo (as suggested by the results of the Pleasant IAT). On the corresponding standard IAT, Go was compatible with pleasant rather than unpleasant attributes, which would suggest that the standard IAT effect is driven by salience asymmetries, rather than valence. This conclusion is also consistent with the current interpretation of the Go/Nogo IAT effect, in which Go is considered to be more salient than Nogo, but there are no valence differences between them.

Another possibility is that unpleasant attributes are more salient than neutral attributes in the split IAT, but pleasant and neutral attributes are equally salient. Thus, consistent with Rothermund and Wentura's (2001, 2004) salience account, unpleasant attributes would be more salient than pleasant attributes. Based on the assumption that pleasant and neutral attributes are equally salient, this would make Go more positive than Nogo (as suggested by the results of the Pleasant IAT), and Go more salient than Nogo (as suggested by the results of the Unpleasant IAT). On the corresponding standard IAT, Go was compatible with pleasant attributes, and Nogo was compatible with unpleasant attributes. This would imply that the standard IAT effect reflects valence differences between Go and Nogo. This interpretation contradicts the current interpretation that the Go/Nogo IAT effect, based on the idea that pleasant and unpleasant attributes are equally salient when compared to neutral attributes.

Although it is possible that the salience asymmetry between attributes in the Pleasant IAT differs from that of the Unpleasant IAT, the assumption of equal salience corresponds with the results of an independent measure of salience in the IAT, the binary classification task. The target category that was more fluently processed on the binary classification task was also compatible with both positive and negative attributes in the split IAT. Thus, the category that was the more salient target on the binary classification task was also the more salient target in the split IAT, based on the assumption of equal salience asymmetries in the Pleasant and Unpleasant IAT. In the binary attribute classification trials of the split IAT, pleasant and unpleasant attributes were classified more quickly than neutral attributes, and this reaction time difference was the same for both valence categories. These results support the idea that both pleasant and unpleasant attributes are, to a similar extent, more salient than neutral attributes in the split IAT.

7.2.2. The assumption of equal valence

The current analysis of the split IAT also assumes that the valence difference between the attribute categories is equal in the Pleasant and Unpleasant IATs. However, it may be that the valence difference between the attribute categories is unequal between the Pleasant and Unpleasant IAT. To consider the extreme case, the attributes may differ in valence in one of the IATs in the split IAT, but not in the other IAT. If there are no valence differences between the attributes in either the Pleasant or Unpleasant IATs, then this would undermine the calculation of a valence effect in the split IAT. This is because any valence differences could only be represented by one half of the split IAT. For example, if pleasant is more positive than neutral in the Pleasant IAT, but neutral and unpleasant attributes are equal in valence in the Unpleasant IAT, then we could only detect valence differences using the Pleasant IAT, but not the Unpleasant IAT. This has implications for the conclusions we can draw in the split IAT with regards to valence.

Let us see how the assumption of equal valence differences between the attributes in the Pleasant and Unpleasant IATs compares to the assumption that there are valence differences between the attributes of the Pleasant IAT, but not the Unpleasant IAT. To do this, these two assumptions will be applied to interpret the results of the Go/Nogo split IAT effect (Experiment 5b, Chapter 3, see Figure 3.4 on p.91). In the Pleasant IAT, Go is compatible with pleasant over neutral attributes. This could mean either that Go is more salient than Nogo, or that Go is more positive than Nogo. In the Unpleasant IAT, Go is compatible with unpleasant over neutral attributes. This is also consistent with the idea that Go is salient more salient than Nogo. If unpleasant attributes are more negative than pleasant attributes, then we could infer that because Go is compatible with unpleasant attributes in the Unpleasant IAT, it cannot be more positive than Nogo. This would allow us to conclude that the Go/Nogo split IAT effect is due to salience asymmetries and not valence differences between Go and Nogo. However, if there are no valence differences between the attributes in the Unpleasant IAT, we cannot infer anything about the valence of Go from the Unpleasant IAT. Although we can conclude that Go is more salient than Nogo, we could not use the result of the Unpleasant IAT to rule out the possibility that Go is also more positive than Nogo. This would make it difficult to assess whether there are valence differences between the target categories.

Therefore, if the assumption of equal valence differences between attributes in the Pleasant and Unpleasant IATs is violated, this would limit the conclusions that could be drawn from the results of the split IAT. Nevertheless, there is independent evidence to support the assumption of equal valence differences in the split IAT. Participants' pleasantness ratings of pleasant and unpleasant attributes were equidistant from their pleasantness ratings of neutral attributes, which suggests that they found the valence

difference in the Pleasant IAT to be equal to that in the Unpleasant IAT. Of course, it is possible that these ratings may differ at the category level, so that the valence difference between the categories of pleasant and neutral words is greater than the valence difference between the categories of neutral and unpleasant words. The magnitude of category valence can only be directly assessed by having participants rate the valence of the category of 'pleasant words', 'neutral words' and 'unpleasant words'. The nature of this task may encourage participants to assume that their rating of 'neutral words' must fall halfway along the rating scale, or halfway between 'pleasant words' and 'unpleasant' words. Thus, such as task would be unlikely to yield illuminating results. Given these circumstances, it would be very difficult to rule out the possibility of unequal valence differences between attributes at the category level.

7.2.3. The assumption that the split IAT is a valid measure of IAT effects

Another assumption of the split IAT is that it can detect salience asymmetries as they would occur in the standard IAT. This assumption is also made of the modified IAT with word/nonword attributes used by Rothermund and Wentura (2001, 2004) and Kinoshita and Peek-O'Leary (2006). It may be argued that these modified forms of the IAT create rather than reveal salience asymmetry effects. That is, because salience asymmetries between attribute categories are more apparent in the modified IATs than in the standard IAT, people may be more likely to categorize stimuli along the lines of salience asymmetries in the former than in the latter case. In the split IAT, salience asymmetry between the attribute categories may be influenced by differences in valence extremity, with the valence attributes being more salient than the neutral attributes because they are more emotionally vivid. In an IAT with word/nonword attributes, salience asymmetry between the attribute categories is thought to be due to differences in familiarity. Differences in either valence extremity or familiarity are less pronounced between pleasant and unpleasant attributes in the standard IAT than they are between the attributes in the modified forms of the IAT. This is likely to produce a smaller salience asymmetry between the attributes of the standard IAT, relative to either the split IAT or the IAT with word/nonword attributes.

Therefore, it could be argued that the standard IAT is, by and large, resistant to salience asymmetries, and salience asymmetries are revealed only when a modified IAT makes them salient. However, evidence from the insect/nonword IAT effect of Brendl et al. (2001) suggests that salience asymmetries can still override the effect of valence in the standard IAT. In Brendl et al.'s study, insects were rated more negatively than nonwords on a self-report measure. In contrast, insects were actually compatible with positive attributes on the IAT, and nonwords were compatible with negative attributes. This implies that the standard IAT does not always reflect actual attitude, but may in fact be influenced by salience asymmetries. Thus, a measure is needed to distinguish between IAT effects that are based on valence, and those that are based on salience asymmetries.

Although modified forms of the IAT may overestimate the contribution of salience asymmetries to the standard IAT effect, the split IAT appears to suffer from this problem less so than the IAT with word/nonword attributes. This is because attributes in the split IAT are defined by their valence, and so are more likely to draw attention to valence differences between the target categories than are word/nonword attributes. This distinction is illustrated by a comparison of the two modified IATs with even and odd number target categories. As mentioned in Chapter 4, when Kinoshita and Peek-O'Leary (2006, Experiment 4) used an IAT with (unpleasant) word and nonword attributes to detect salience asymmetries between even and odd numbers, they found that even numbers were compatible with words and odd numbers were compatible with nonwords. This result led them to conclude that the even/odd number IAT effect was driven by salience asymmetries, but not valence differences. In contrast, when even and odd numbers were classified in a split IAT in Experiment 8b, there was a valence effect for even numbers, but no salience effect. This suggests that the split IAT may be less sensitive to salience differences (and more sensitive to valence differences) between even and odd numbers than are words and nonwords. Like the split IAT, the attributes of the standard IAT are also defined by valence, which serves to highlight any valence differences between the target categories. Therefore, relative to an IAT with word/nonword attributes, the split IAT may be more sensitive to the processes that occur in the standard IAT.

The split IAT is a complicated measure designed to assess the contributions of valence and salience in the standard IAT. There is no definitive evidence to support the validity of combining the Pleasant and Unpleasant IATs to obtain separate salience and valence measures in the split IAT. However, the results of the independent measures of processing fluency and valence provide support for the assumption that the salience and valence asymmetries in the Pleasant IAT are equal to those in the Unpleasant IAT. Furthermore, the split IAT appears to be a closer analogue of the standard IAT than is the IAT using word/nonword attributes. In light of these results, and for the purposes of the remaining discussion, I will assume that the split IAT is a valid measure of valence and salience effects in the IAT. In the following section, I explore the relationship between salience asymmetries and processing fluency in the IAT.

7.3. The relationship between salience asymmetries and processing fluency in the IAT

If salience asymmetries are indeed, as much of the present data suggest, related to the fluency of classification, then what exactly is the nature of this relationship? Rothermund and Wentura (2001) claim that salience asymmetries in the IAT are driven by selective attention, whereby people focus predominately on one of the two target categories, and on one of the two attribute categories. To understand the mechanism underlying this phenomenon in the IAT, one must understand which types of categories people tend to focus on when classifying stimuli, and why these categories are attended to. This requires a consideration of the cognitive processes underlying categorization. In the following sections, I first discuss which factors determine whether a category will be salient in binary classification. I then discuss how salience asymmetries can lead to fluency asymmetries in classification. Lastly, I discuss circumstances under which salience asymmetries do not correspond with fluency asymmetries.

7.3.1. How is salience determined in classification?

Traditionally, categorization theorists have assumed that classification is based on the similarity of stimuli to stored representations of category exemplars (e.g. Medin & Schaffer, 1978), or prototypes (e.g. Homa, Cross, Cornell, Goldman & Schwartz, 1973). In a straightforward binary classification task (e.g. classifying flowers and insects) decisions are thought by some accounts to be to be based on some comparison to representations of both categories (e.g. Nosofsky, 1986). The stimulus is then classified as belonging to the category whose representations it most resembles. According to exemplar models of categorization, under these circumstances, representations from both categories are similarly activated, entailing that both categories are equally attended to.

The assumption that each category is equally activated is likely to be violated if the categories differ in the strength of their representation. For example, it should be easier to compare a stimulus to a category representation when the stimulus belongs to a familiar (e.g. words) rather than an unfamiliar category (e.g. nonwords), as representations of an unfamiliar category are non-existent or poorly formed. Therefore, in instances where participants must classify stimuli as belonging to either a familiar or unfamiliar category, it would be reasonable to assume that they compare stimuli to representations of the familiar category only. For example, in lexical decision, if the stimulus matches an existing word representation, then the stimulus is classified as a word; if the stimulus does not match a representation (i.e. there is a mismatch with the closest representation) then it is classified as a nonword. In this case, the familiar category serves as the focal reference category against which all stimuli are compared. This process is consistent with the idea that category learning may involve the extraction of simple, logical rules (Nosofsky, Palmeri, & McKinley, 1994). In the previous example, participants use the rule that a stimulus is a nonword if it does not match a representation from the word category. In the current studies, such processes are likely to have operated when classifying live/dead elephants (Experiments 1c, 2a and 3), insects/nonwords (Experiments 6a-6b), and familiar/unfamiliar paintings (Experiments 7a-7b). Thus, the strength of a category's representations may determine whether it will be the focal category. In this way, the ease with which the exemplars from that category are matched to that category representation can lead to that category being more salient.

Under some circumstances, instead of comparing all aspects of a stimulus to stored category exemplars, classification can be simplified by shifting attention to the dimension/s that most effectively distinguish one category from another (Getty, Swets, Swets, & Green, 1979; Nosofsky, 1984, 1986; Reed, 1972; Shepard, Hovland, & Jenkins, 1961). This method may have been used to discriminate between black and white faces in Experiments 9a and 9c. According to Levin (2000), white participants classify black faces more quickly than they do white faces because they focus on a racial marker that is more apparent in cross-race faces than same-race faces. This allows stimuli to be classified by comparing them to exemplars of the critical feature/s, rather than to exemplars of the whole category. Although a familiar category (e.g. white faces) has a stronger global representation, it may be represented less strongly in terms of the critical diagnostic dimension (e.g. race). The category with the weaker global representation (e.g. black faces), but which is better represented on the critical dimension, will be more salient. Therefore, attention to the critical dimension, combined with how well a category is represented on that dimension, can determine salience asymmetries in the IAT.

In the experiments presented here, participants may also have shifted their attention to the critical dimension when classifying Accept/Reject (Flight/Teeth) items in Experiments 11a-11b, big/small cows in Experiments 1d, 2b and 3, and the attribute categories in the split IAT. In Experiments 11a-11b, Accept items (e.g. Flight) can be differentiated from Reject items (No Flight) by the presence of the 'flight' feature that defines the Accept category. Similarly, the two cow breeds in Experiments 1d, 2b and 3 are classified by focusing on the dimension of size. Big cows have a higher value of on this dimension, and this acts as the presence of a feature in big cows. It follows then, that the category which possesses this feature becomes the referent category. When classifying attributes in the split IAT, valence is the critical dimension that participants focus on. Positive and negative attributes in the split IAT are better represented on this dimension, because valence is defined by the presence of positive/negative associations, whereas neutrality is defined by the absence of such associations. This makes pleasant/unpleasant the salient focal category compared to neutral.

Therefore, we see that the categories that are focused upon in the IAT are the ones that are better represented as a whole, or better represented on the critical dimension. In some cases, representation strength depends on the features of the stimuli (e.g. the size of the stimuli in the cow categories), whereas in other cases, representation strength is determined by the features of the category (e.g. the valence extremity of attributes in the split IAT). In the following section, I explain how selective attention caused by salience asymmetries can produce fluency asymmetries in classification.

7.3.2. How do salience asymmetries influence processing fluency?

In the previous section, I described that when representation strength differs between categories, people may compare stimuli to representations of the exemplars of the more salient category (e.g. words over nonwords), or to representations of the critical feature of the more salient category (e.g. the race feature in black vs. white faces). In this way, stimuli from the more salient category produce a match with the representations they are compared against, whereas stimuli from the less salient category produce a mismatch. Matches are generally processed more quickly than mismatches, as illustrated by findings from sentence-picture verification tasks. In these tasks, participants are presented with a concept (e.g. "green square"), and are asked whether a particular object (e.g. a picture of a green square) is or is not an exemplar of that concept (e.g. "Is this a green square?"). Their reaction time to make this decision is recorded. Sentence-picture verification tasks demonstrate that it is faster to identify a match (e.g. green square) than a mismatch (e.g. red circle; Gough, 1965; Just & Carpenter, 1971; Trabasso, Rollins, & Shaughnessy, 1971). Therefore, in the context of salience asymmetries in the IAT, responses to stimuli from the more salient category (indicating a match with the reference category) should be quicker than to stimuli from the less salient category (indicating a mismatch with the reference) category can produce differences in processing fluency. This principle can be seen most clearly in the mere acceptance binary classification task with Flight/Teeth categories (Experiment 11a). Participants in this task were required to classify stimuli according to the mere acceptance rule 'Does it fly?' (Flight condition), or 'Does it have teeth?' (Teeth condition). Accept items indicated a match with the reference category.

The relationship between processing fluency and salience asymmetries in classification appears to be bi-directional. The fluency with which a category is classified (e.g. because it is more familiar) can lead to that category being more salient in the IAT. In addition, selective attention toward the salient category (e.g. the Go category) can increase the fluency with which that category is classified. These two processes may operate in a positive-feedback loop, reinforcing the relationship between fluency and salience asymmetries. This association between fluency and salience is a consistent pattern seen in most of the experiments presented here. However, it is not always the case that the salient category is the more fluently processed category. The reason for this discrepancy will be explained in the following section.

7.3.3. When salience and fluency do not correspond

Although the relationship between salience asymmetries and processing fluency has been widely replicated in this thesis, it must be noted that not all salient categories are more fluently processed. One counterexample is the Column Same stimuli of Experiments 12a-12b (Chapter 6). These stimuli appear to be more salient than Column Different stimuli (Experiment 12b). However it is Column Different stimuli which are more fluently processed in a binary classification task (Experiment 12a). This may have occurred because, although the mere acceptance instruction focused participants' attention on the Column Same category, the Column Different stimuli were actually easier to process. As can be seen by the Row/Column stimuli in Figure 6.1, (Chapter 6, p.162), Column Different (Row Same) stimuli may be easier to process because they possess more adjacent matching numbers than Column Same stimuli, which allows the numbers of Column Same to be more readily grouped together for greater processing efficiency. Thus, in addition to salience asymmetries affecting fluency asymmetries, the features of the target stimuli can also contribute to processing fluency.

The influence of stimulus features on processing fluency can also be seen with the pleasant and unpleasant attributes in the current experiments. In these experiments, pleasant attributes appear to be more salient than unpleasant attributes, because they were compatible with target categories that were shown to be salient in the split IAT (Go in Experiments 5a-5b, insects in Experiment 6a-6b, Accept in Experiments 11a-11b, and Row Same in Experiment 12a-12b). As mentioned in Chapter 3, the pleasant category may be more salient than the unpleasant category because positive is linguistically unmarked, and thereby serves as the reference category on the evaluative dimension. If so, then

participants are likely to focus on the pleasant category from the outset of the IAT, independently of the stimulus features of the attributes. However, the pleasant and unpleasant attributes were classified equally quickly on a separate binary classification task (Experiment 4a). This may have occurred because the pleasant attribute stimuli in these experiments were processed less fluently than unpleasant attribute stimuli. This is because the pleasant exemplars used were, on average, longer than the negative exemplars, and longer words are processed more slowly (Balota, Cortese, Sergent-Marshall, Spieler, & Yap, 2004). Indeed, when pleasant and unpleasant attributes are matched on word length and frequency, pleasant attributes tend to be processed more quickly than unpleasant attributes (Estes & Adelman, in press; Unkelbach et al., in press).

Therefore, it appears that salience asymmetries and fluency do not always correspond when salience asymmetries are determined by factors that are independent of stimulus features, such as task instruction (in the mere acceptance IAT), and linguistic markedness (between pleasant and unpleasant attributes). This means that controlling for the processing fluency of stimuli in the IAT does not necessarily eliminate salience asymmetries. This argument is similar in principle to the idea that familiarity in the IAT cannot be controlled for by equating the familiarity of category stimuli (Rothermund & Wentura, 2004; Kinoshita & Peek-O'Leary, 2005), because the categories themselves may still differ in familiarity. Given that salience asymmetries cannot be reliably predicted or controlled for, it is necessary to use the split IAT (or a similar measure) to distinguish between the effects of valence and salience asymmetries in the IAT. In the following section, I examine how valence and salience asymmetries may combine to influence IAT effects.

7.4. How do salience asymmetries combine with valence in the IAT?

In this section, I first evaluate how the evidence presented in this thesis addresses the contribution of salience asymmetries to valance-based IAT effects. After identifying issues that require further clarification, I propose a study that more directly examines how salience asymmetries combine with valence in the IAT.

Salience asymmetries appear to increase the compatibility between the more salient target category and positive attributes in the IAT. Because of this, salience asymmetries also have the potential to distort valence-based IAT effects. The results of the current experiments suggest two clear patterns of how salience asymmetries may influence IAT effects. Firstly, IAT effects are likely to be inflated when the more salient target is also the more positive target category in the IAT. This is because the IAT effect in this case is a combination of two congruent compatibility effects, one of which is valence-based, the other of which is salience-based. This can be seen in the comparison of the Row and Column mere acceptance conditions of Experiments 12a-12b. There was a salience effect for the Accept category in both conditions, with an additional valence effect for Accept in the Column condition. On the standard IAT, the Column condition produced a larger IAT effect than the Row Same condition. From this, it would appear that the salience and positivity of the Column Same category combined to inflate the mere acceptance effect.

Secondly, IAT effects are likely to be reduced when the more salient target is also the more negative target category in the IAT. This occurs because, although the less salient target category is compatible with positive attributes on the grounds of valence, the more salient target category is compatible with positive attributes on the grounds of salience asymmetries. This effect can be seen most clearly in the comparison of the age split IAT with faces (Experiment 10c) and the age split IAT with names (Experiment 10d). In both split IATs, there was a valence effect for the young category, with the addition of a salience effect for old faces in Experiment 10c, but no salience effect for old names in Experiment 10d. On the corresponding standard IATs, the age IAT for faces produced a smaller effect than the age IAT for names. This is consistent with the idea that the salience compatibility between old faces and pleasant attributes to some extent counteracted the valence compatibility between young faces and pleasant attributes.

The above analyses assume that the effects of valence and salience asymmetries combine in an additive manner to influence IAT effects. This assumption is based on comparing the size of IAT effects from similar target categories that vary salience (e.g. young and old faces vs. young and old names) or valence (e.g. Row Same/Different vs. Column Same/Different). However, because the target categories/stimuli being compared are not identical, there may be other differences between them that account for variations in standard IAT effect sizes. Therefore, the question of how salience and valence are combined in the IAT can only be conclusively answered by comparing the results of an IAT that is purely valence based, such as the flower/insect IAT, against the same IAT in which the salience of the target categories is manipulated. Such a study could be conducted by varying salience by means of a Go/Nogo manipulation, because the results from the Go/Nogo split IAT (Experiment 5b) suggest that such a manipulation does not affect the valence of the target categories. If the effects of salience and valence combine in an additive manner, then we would expect that, compared to the standard flower/insect IAT effect, IAT effects should be larger when flower is the Go category and insect is the Nogo category (Go-flower/insect IAT), and smaller when insect is the Go category and flower is the Nogo category (Go-insect/flower IAT).

211

Alternatively, valence and salience asymmetries may not have additive effects in the IAT, but instead function as mutually exclusive bases of compatibility in the IAT. This principle may operate in the Go-insect/flower IAT. In the valence-compatible condition of the Go-insect/flower IAT, the affectively congruent categories share a key (e.g. Nogo flower and pleasant attributes share a key, and Go insect and unpleasant attributes share a key). In this case, participants may use a valence-based strategy to perform the IAT. In the valence-incompatible condition, the affectively incongruent categories share the same key (e.g. Go insect and pleasant attributes share a key, and Nogo flower and unpleasant attributes share a key). However, in this case the categories sharing the same key are similarly salient, and thus a salience-based strategy may be used. This means that the target and attribute categories sharing the same key are compatible in both the valence-compatible and valence-incompatible conditions, but for different reasons. These conditions may encourage participants to shift from focusing on the valence dimension in the former condition, to focusing on the salience dimension in the latter condition.

If salience asymmetries function as an alternative, rather than as an additional source of compatibility in the IAT, then we would expect salience asymmetries to operate only when a valence-based strategy cannot be used, and then only when the categories sharing the same key are similarly salient. As can be seen in the previous analysis of the Go-insect/flower IAT, this means that salience asymmetries are only likely to have an impact when the negative target category is the more salient target, and only in the valence-incompatible condition of the IAT. In this case, salience asymmetries would enhance the compatibility between affectively incongruent target and attribute categories. This would make performance faster (and more accurate) in the valence-incompatible condition of the Go-insect/flower IAT, compared to same condition in the standard flower/insect IAT. In
this way, a salience-based strategy can be used to counteract any valence-based IAT effect effects, thereby reducing the size of the overall IAT effect. However, we would not expect to see any differences between the two IATs in the valence-compatible condition, because the same valence-based strategy should be used in both cases.

If salience asymmetries provide an alternative source of compatibility to valence in the IAT, then they will not effect IATs in which the positive target category is the more salient target, such as the Go-flower/insect IAT. In this case, valence and salience asymmetries are aligned, so when the target and attribute categories sharing the same key are compatible in terms of salience (e.g. Go-flower shares a key with pleasant attributes, and Nogo-insect shares a key with unpleasant attributes), they are also compatible in terms of valence. Therefore, any source of compatibility on the grounds of salience is redundant in the valence-compatible condition, and non-existent in the valence-incompatible condition. From this, we would expect the resulting Go-flower/insect IAT effect to be equal to that of the standard flower/insect IAT effect. Such a result would imply that IAT effects are not inflated when the positive target category is the more salient target.

In summary, the data from the current experiments involving the Row/Column mere acceptance categories and the young/old categories suggest that valence and salience asymmetries have additive effects in the IAT. However, these data only provide indirect evidence supporting this claim, because the results may be due to category differences between the stimuli being compared (in the case of Row/Column categories), or stimulus differences between the categories being compared (in the case of young/old faces and names). The question of how salience asymmetries combine with the effects of valence can only be definitively answered by a direct comparison of a valence-based IAT with and without a salience manipulation of the target categories. One study that used a similar method was conducted by Sargent, Kahan and Mitchell (2007), which will be discussed in the next section. If valence and salience asymmetries have additive effects in the IAT, then salience asymmetries have the potential to both increase and decrease IAT effects. However, if the operation of valence factors and salience asymmetries is mutually exclusive, this implies that salience asymmetries can only ever reduce valence-based IAT effects, not increase them. If so, then IAT effects will only be reduced when the negative target category is more salient, as this allows salience asymmetries to increase the compatibility between target and attribute categories in the valence-incompatible condition.

7.5. Implications of salience asymmetries for other IAT effects

In this section, I examine the implications that salience asymmetries have for other IAT effects in the literature, and suggest ways in which these implications may be tested.

7.5.1. The race IAT effect

Sargent et al. (2007) have claimed that the race IAT effect may be caused by participants adopting a mere acceptance strategy when performing the race IAT. That is, the white target category may be compatible with pleasant attributes because participants treat white as the 'Accept' category in the IAT, and black as the 'Reject' category. Sargent et al. investigated this in a mere acceptance race IAT in which participants classified white and black names according to a mere acceptance rule. In the white Accept condition, the target categories were classified as white/not white (making white the Accept category); in the black Accept condition, the target categories were classified as black/not black (making black the Accept category). The race IAT effect was larger when white was made the Accept category than when black was made the Accept category. When both the white Accept and black Accept conditions were compared to a standard IAT (in which the two target categories were classified as white/black), the white Accept condition produced an IAT effect size that was comparable to that of the standard race IAT effect, but the black Accept condition produced an IAT effect size which was smaller than that obtained with the standard race IAT. The smaller IAT effect in the black Accept condition is consistent with the idea that increasing the salience of black will increase the compatibility of black with pleasant attributes on the grounds of salience asymmetries, thereby reducing the IAT effect. However, increasing the salience of white did not produce a corresponding change in the white Accept condition. Sargent et al. proposed that this null effect was caused by participants adopting a similar strategy in the standard race IAT as they do in the white Accept condition, that of focusing on the white category. That is, they claimed that white is more salient than black in the standard race IAT. This position contradicts the argument presented in this thesis, that black is the more salient target in the race IAT.

However, Sargent et al.'s (2007) finding is also consistent with the idea that valence and salience do not combine in an additive manner in the IAT. As discussed in the previous section, participants may choose between a valence-based strategy and a salience-based strategy when performing each condition of the IAT. If so, then increasing the salience of the positive white category will not increase the magnitude of the IAT effect. Of course, this interpretation of Sargent et al.'s results is predicated on the assumption that there are no salience differences between the black and white target categories. If black is actually more salient than white (as was found in the race split IATs of Experiments 9c-9d, Chapter 5), then increasing the salience of the white category should actually increase the resulting IAT effect. This is because increasing the salience of white would increase the compatibility between white and positive attributes, and decrease the compatibility between black and positive attributes, on the grounds of salience asymmetries.

Although the race split IAT result of Experiment 9c suggest that black names to be more salient than white names, the participants in Sargent et al.'s (2007) study may not show the same salience asymmetry, due to population differences between the two studies. Participants in the present experiments were recruited from an Australian university, and, generally speaking, would have had little experience with black names. In contrast, the participants in Sargent et al.'s study were recruited from an American university, and would be more familiar than the Australian participants with black names. It may be that the American participants in Sargent et al.'s study found the two name categories to be equally salient (as Australian participants appeared to have with young and old names in Experiment 10d).

To investigate whether mere acceptance is responsible for the race IAT effect, American participants could perform a split IAT with name stimuli. One group of participants could be instructed to classify the target categories as white/black (the standard condition), another group of participants could classify the target categories as white/not white (the white Accept condition), and a third group could classify the target categories as black/not black (the black Accept condition). In Sargent et al.'s (2007) study, there was a smaller IAT effect in the black Accept condition than in the standard race IAT. If this difference is due to salience asymmetries, there should be a larger salience effect for black in the black Accept condition of the split IAT than in the standard condition of the split IAT. Sargent et al. also claim that white is more salient than black in both the standard IAT and the white Accept condition. If so, then there should be a salience effect for the white category in both the standard and white Accept conditions of the race split IAT. Equal salience effect sizes in these two conditions would provide evidence that participants adopt the strategy of focusing on the white category in the standard race IAT. However, if there is no salience effect for white in the standard condition of the split IAT, this suggests that the race split IAT effect is largely valence-based. This result, in combination with Sargent et al.'s data, would also imply that increasing the salience of the white category has no effect on the standard IAT, and thus salience asymmetries cannot inflate the race IAT effect.

To summarize, the race mere acceptance effect is consistent with two interpretations of salience asymmetries in the race IAT. According to Sargent et al. (2007), this effect suggests that white names are more salient than black names in the IAT. Alternatively, the race mere acceptance effect may also occur if there are no salience asymmetries in the standard race IAT with names. This is based on the assumption that the effects of valence and salience do not combine in an additive manner in the IAT. Either way, a comparison of the race split IAT with names (Experiment 9c) and Sargent et al.'s study implies that Australian and American participants differ with respect to how they process the same target categories in the IAT. In particular, black names are more salient than white names for Australian participants, but this salience asymmetry may be reduced or reversed for American participants. This suggests that, although two different populations may exhibit the same race IAT effect, this IAT effect may be valence-based for one population, but salience-based for another population. Any such cross-cultural differences may be discerned by the split IAT, but not the standard IAT.

7.5.2. Category-level IAT effects

The IAT has been claimed to assess the valence of target categories rather than of individual exemplars (De Houwer, 2001; Fazio & Olson, 2003; Gawronski et al., 2007). This principle is illustrated by De Houwer's (2001) British/Foreign IAT effect. As mentioned in Chapter 1 (Section 1.5.2), British participants in this experiment classified the target categories 'British' and 'Foreign' with positive and negative attributes. Each target category consisted of the names of three positively evaluated public figures, and three negatively evaluated public figures. There was an IAT effect for British that did not appear to interact with the valence of the stimuli, as both positive and negative target exemplars were responded to equally quickly in the IAT. From this result, De Houwer (2001) concluded that evaluations in the IAT occur at the category level rather than at the level of individual exemplars. Support for this claim also comes from a study by J.P. Mitchell et al. (2003, Experiment 1) in which participants completed an IAT in which the target exemplars were three liked African American athletes and three disliked Caucasian politicians. The standard race IAT effect was obtained when the targets were classified as belonging to the categories 'black' and 'white'. However, this effect was reversed when participants classified the same stimuli as either 'athlete' or 'politician'. This was taken to indicate that the same stimulus can be evaluated differently depending on which category it represents. Other researchers have interpreted this finding as evidence that evaluations in the IAT are based on the valence of the category labels rather than on the valence of individual exemplars (e.g. Rothermund & Wentura, 2004; Gawronski et al., 2007).

As mentioned in Chapter 1 (Section 1.5.2), Rothermund and Wentura (2001, 2004) claim that De Houwer's (2001) British/Foreign IAT effect may have been due to salience

asymmetries between the categories. They suggest that salience asymmetries may have arisen because the British participants in the study found the British category to be more familiar than the Foreign category. However, the evidence presented in this thesis demonstrates that the familiar ingroup is not always compatible with pleasant attributes on the grounds of salience asymmetries. For example, in the race IATs and split IATs of Experiments 9a-9d (Chapter 5), the unfamiliar category of black was the more salient target category in the split IAT, but it was compatible with unpleasant attributes on the standard IAT. Nonetheless, in De Houwer's study, the British category may have been more salient than the Foreign category due to mere acceptance. That is, instead of comparing the target exemplars equally to category representations of British and Foreign personalities, participants adopted a strategy in which they classified targets as being either 'British' or 'Not British', rather than 'British' or 'Foreign'. This may have occurred because the British category was much more homogenous than the Foreign category (whose exemplars were sourced from five nations), and thus easier to classify. This would allow participants to simplify the task by focusing on the more specific category of British. Thus, an IAT effect for British over Foreign may be caused by salience asymmetry between the target categories, rather than the positivity of British at the category level.

To examine whether category-level valence or salience asymmetries underlie the British/Foreign IAT effect, a split IAT could be conducted with the same target categories and exemplars. If British is compatible with the more positive attributes in the split IAT, then this would support De Houwer's (2001) hypothesis that category valence rather than exemplar valence is the relevant feature in the IAT. However, if British is compatible with the more salient attributes in the split IAT, then this suggests that the British/Foreign IAT effect is caused by salience asymmetries. A split IAT could also be used to assess J.P. Mitchell et al.'s (2003) athlete/politician IAT effect. Kinoshita and Peek-O'Leary (2005) claim that this effect may be due to salience asymmetries between the target categories, with athlete being compatible with positive attributes because it is the more salient target category, and politician being compatible with negative attributes because it is the less salient target category.

In summary, the category-level IAT effects of De Houwer (2001) and J.P. Mitchell et al. (2003) have been assumed to reflect differences in category valence, but they may actually be due to salience asymmetries between the target categories. If the category-level IAT effects are due to salience asymmetries, this would undermine the idea that the IAT assesses the valence of target categories rather than of individual exemplars. The split IAT can be used to distinguish between the valence and salience accounts of category-level IAT effects.

7.6. Concluding remarks

The IAT was originally designed to measure evaluative associations, but it has since been discovered to be susceptible to contamination from salience asymmetries. The experiments reported here clarify several issues regarding the nature of salience asymmetries and their influence on IAT effects. Firstly, salience asymmetries between target categories are generally characterized by differences in processing fluency of the categories, as indexed by reaction times on a binary classification task. Secondly, salience asymmetries do not consistently correspond with of visual search asymmetries, familiarity, and linguistic markedness. The split IAT was introduced as a method to distinguish between salience asymmetries and valence in the IAT. The category is compatible with positive attributes in the split IAT is also compatible with pleasant attributes in the standard IAT. The target category that is revealed to be more salient in a split IAT is, all things being equal, also compatible with positive attributes on the standard IAT. However, it is not yet certain how the effects of valence and salience are combined in the IAT. It is possible that when the target categories differ in valence, salience asymmetries can increase IAT effects when the positive category is the more salient target, and decrease IAT effects when the negative category is more salient. It is also possible that salience asymmetries cannot inflate IAT effects, but only reduce them.

The data described here present a consistent pattern between the results of the binary classification task and the split IAT, and between the split IAT and the standard IAT. However, there are some caveats to keep in mind. Firstly, the IAT is a new measure that rests on certain assumptions about the salience and valence of the categories used. Secondly, although the present data do not reveal the split IAT to be flawed, it is possible that future data will. Nonetheless, the research presented contributes much to an understanding of how salience asymmetries operate, and how they can be manipulated, measured, and distinguished from the effects of valence. Based on this foundation, future research in this area should provide many fruitful results.

References

- Adolphs, R. A. (1998). The human amygdala in social judgment. Nature, 393, 470-474.
- Allport, G.W. (1935). Attitudes. In C. Murchison (Ed.), Handbook of social psychology (pp. 798–844). Worcester, MA: Clark University Press.
- Arkes, H., Tetlock, P. E. (2004). Attributions of Implicit Prejudice, or "Would Jesse Jackson 'Fail' the Implicit Association Test?" *Psychological Inquiry*, *15*, 257-278.
- Ashburn-Nardo, L., Voils, C. I., & Monteith, M. J. (2001). Implicit associations as the seeds of intergroup bias: How easily do they take root? *Journal of Personality and Social Psychology*, 81, 789–799.
- Balota, D. A., & Chumbley, J. I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 10, 340–357.
- Balota, D. A., Cortese, M. J., Sergent-Marshall, S. D., Spieler, D. H., & Yap, M. J. (2004).
 Visual word recognition of single-syllable words. *Journal of Experimental Psychology: General, 133,* 283–316.
- Baron, A. S., & Banaji, M. R. (2006). The development of implicit attitudes: Evidence of race evaluations from ages 6, 10, and adulthood. *Psychological Science*, 17, 53–58.
- Blair, I.V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, 6, 242–261.
- Blair, I.V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology*, 81, 828–841.

- Blanton, H., Jaccard, J., Gonzales, P., & Christie, C. (2006). Decoding the implicit association test: Perspectives on criterion prediction. *Journal of Experimental Social Psychology*, 42, 192–212.
- Bornstein, R.F., & D'Agostino, P.R. (1992). Stimulus recognition and the mere exposure effect. *Journal of Personality and Social Psychology*, 63, 545-552.
- Bothwell, R.K., Brigham, J.C., & Malpass, R.S. (1989). Cross-racial identification. Personality & Social Psychology Bulletin, 15, 19-25.
- Breiter, H. C., Etcoff, N. L., Whalen, P. J., Kennedy, W. A., Rauch, S. L., Buckner, R. L.,
 Strauss, M. M., Hyman, S., & Rosen, B. (1996). Response and habituation of the
 human amygdala during visual processing of facial expression. *Neuron*, 17, 875–887.
- Brendl, M., Markman, A., & Messner, C. (2001). How do indirect measures of evaluation work? Evaluating the inference of prejudice in the Implicit Association Test. *Journal of Personality and Social Psychology*, 81, 760-773.
- Buchner, A., Rothermund, K., Wentura, D., & Mehl, B. (2004). Valence of distractor words increases the effects of irrelevant speech on serial recall. *Memory & Cognition*, 32, 722-731.
- Chiroro, P. & Valentine, T. (1995). An investigation of the contact hypothesis of the ownrace bias in face recognition. *Quarterly Journal of Experimental Psychology*, 48A, 879-894.
- Cho, Y. S., & Proctor, R. W. (2005). Representing response position relative to display location: Influence on orthogonal stimulus-response compatibility. *Quarterly Journal of Experimental Psychology*, 58A, 839–864.

- Clark, H. H. (1973). Space, time, semantics, and the child. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 27-63). New York: Academic Press.
- Crowne, D.P., & Marlowe, D. (1960). A new scale of social desirability independent of psychopathology. *Journal of Consulting Psychology*, *24*, 349–354.
- Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes:
 Combating automatic prejudice with images of admired and disliked individuals.
 Journal of Personality and Social Psychology, 81, 800–814.
- Dasgupta, N., McGhee, D. E., Greenwald, A. G., & Banaji, M. R. (2000). Automatic preference for White Americans: Eliminating the familiarity explanation. *Journal of Experimental Social Psychology*, 36, 316–328.
- Davis, M. (1997). Neurobiology of fear responses: The role of the amygdala. *Journal of Neuropsychiatry and Clinical Neurology*, *9*, 382–402.
- De Houwer, J. (2001). A structural and process analysis of the Implicit Association Test. Journal of Experimental Social Psychology, 37, 443-451.
- De Houwer, J. (2003). The Extrinsic Affective Simon task. *Experimental Psychology*, 50, 77-85.
- De Houwer, J. (2006). What are implicit measures and why are we using them? In R.W. Wiers & A.W. Stacy (Eds.), *The handbook of implicit cognition and addiction* (pp. 11–28). Thousand Oaks, CA: Sage.
- De Houwer, J., Beckers, T., & Moors, A. (2007). Novel attitudes can be faked on the Implicit Association Test. *Journal of Experimental Social Psychology*, 43, 972-978.

- De Houwer, J., Geldof, T., & De Bruycker, E. (2005). The Implicit Association Test as a general measure of similarity. *Canadian Journal of Experimental Psychology*, 59, 228-239.
- Estes, Z. & Adelman, J. S. (in press). Automatic vigilance for negative words in lexical decision and naming: Comment on Larsen, Mercer, and Balota (2006). *Emotion*.
- Fazio, R. H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 247–282). Mahwah, NJ: Erlbaum.
- 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, 1013–1027.
- Fazio, R.H., & Olson, M.A. (2003). Implicit measures in social cognition research: Their meaning and use. Annual Review of Psychology, 54, 297–327.
- 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*, 229–238.
- 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 and Emotion*, 14, 61-92.
- Fox, E., Russo, R., & Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition and Emotion*, *16*, 355-379.
- Freud, S. (1958). The question of lay analysis. In J. Strachey (Ed.), *The standard edition of the complete psychological works of Sigmund Freud* (Vol. 20, pp. 183-250).
 London: Hogarth Press. (Original work published 1926).

- Garcia-Marques, T., Mackie, D. M., Claypool, H. M., & Garcia- Marques, L. (2004).
 Positivity can cue familiarity. *Personality and Social Psychology Bulletin, 30*, 585–593.
- Gawronski, B., & Bodenhausen, G.V. (2005). Accessibility effects on implicit social cognition: The role of knowledge activation versus retrieval experiences. *Journal of Personality and Social Psychology*, 89, 672–685.
- Gawronski, B., LeBel, E. P., & Peters, K. R. (2007). What do implicit measures tell us?
 Scrutinizing the validity of three common assumptions. *Perspectives on Psychological Science*, 2, 181-193.
- Getty, D. J., Swets, J. B., Swets, J. A., & Green, D. M. (1979). On the prediction of confusion matrices from similarity judgments. *Perception & Psychophysics*, 26, 1-19.
- Gough, P. B. (1965). Grammatical transformations and speed of understanding. *Journal of Verbal Learning and Verbal Behavior*, *4*, 107-111.

Greenberg, J.H. (1966). Language universals. The Hague: Mouton Publishers.

Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., & Mellott,
D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and
self-concept. *Psychological Review*, 109, 3–25.

Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., & Rosier, M. (2000). Prologue to a unified theory of attitudes, stereotypes, and self-concept.
In Forgas, J. P. (Ed.), *Feeling and thinking: The role of affect in social cognition. Studies in emotion and social interaction, second series.* (pp. 308-330). New York:
Cambridge University Press.

- Greenwald, A. G., & Farnham, S. D. (2000). Using the Implicit Association Test to measure self-esteem and self-concept. *Journal of Personality and Social Psychology*, 79, 1022–1038.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. K. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, 74, 1464-1480.
- Greenwald, A. G., Nosek, B. A., Banaji, M. R., & Klauer, K. C. (2005). Validity of the Salience Asymmetry Interpretation of the IAT: Comment on Rothermund and Wentura (2004). *Journal of Experimental Psychology: General, 134*, 420-425.
- Heit, E., & Bott, L. (2000). Knowledge selection in category learning. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 39, pp. 163-199). San Diego: Academic Press.
- Homa, D., Cross, J., Cornell, D., Goldman, D., & Schwartz, S. (1973). Prototype abstraction and classification of new instances as a function of number of instances defining the prototype. *Journal of Experimental Psychology*, 101, 116-122.
- Hugenberg, K., & Bodenhausen, G.V. (2003). Facing prejudice: Implicit prejudice and the perception of facial threat. *Psychological Science*, *14*, 640–643.
- Hugenberg, K., & Bodenhausen, G. V. (2004). Ambiguity in social categorization: The role of prejudice and facial affect in race. *Psychological Science*, *15*, 342–345.
- Just, M. A., & Carpenter, P. A. (1971). Comprehension of negation with qualification. Journal of Verbal Learning and Verbal Behavior, 10, 244-253.
- Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. Journal of Personality and Social Psychology, 81, 774–788.

- Kester, J. D. (2001). A revolution in social psychology. *The American Psychological* Society Observer, 14, 9-26.
- Kinoshita, S., & Peek-O'Leary, M. (2005). Does the compatibility effect in the race
 Implicit Association Test (IAT) reflect familiarity or affect? *Psychonomic Bulletin*& *Review*, 12, 442-452.
- Kinoshita, S., & Peek-O'Leary, M. (2006). Two bases of the compatibility effect in the Implicit Association Test (IAT). *The Quarterly Journal of Experimental Psychology*, 59, 2102-2120.
- Klauer, K. C., & Mierke, J. (2005). Task-set inertia, attitude accessibility, and compatibility- order effects: New evidence for a task-set switching account of the Implicit Association Test effect. *Personality and Social Psychology Bulletin, 31*, 208–217.
- Kucera, H., & Francis, W. N. (1967). Computational analysis of present-day American English. Providence, RI: Brown University Press.
- Kunst-Wilson, W. R., & Zajonc, R. B. (1980). Affective discrimination of stimuli that cannot be recognized. *Science*, 207, 557-558.
- Lane, K. A., Banaji, M. R., Nosek, B. A., & Greenwald, A. G. (2007). Understanding and using the Implicit Association Test: IV: Procedures and validity. In B. Wittenbrink & N. Schwarz (Eds.), *Implicit measures of attitudes: Procedures and controversies* (pp. 59-102). New York: Guilford Press.
- LeDoux, J. E. (1996). The emotional brain. New York: Simon & Schuster.
- Levin, D. T. (1996). Classifying faces by race: The structure of face categories. Journal of Experimental Psychology: Learning, Memory and Cognition, 22, 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*, 559-574.
- Levin, D. T., & Angelone, B. L. (2001). Visual search for a socially defined feature: What causes the search asymmetry favoring cross-race faces? *Perception & Psychophysics*, 63, 423-435.
- Lochy, A., Seron, X., Delazer, M., & Butterworth, B. (2000). The odd–even effect in multiplication: Parity rule or familiarity with even numbers? *Memory & Cognition*, 28, 358–365.
- Lowery, B. S., Hardin, C. D., & Sinclair, S. (2001). Social influence effects on automatic racial prejudice. *Journal of Personality and Social Psychology*, *81*, 842–855.

Lyons, J. (1977). Semantics I. London: Cambridge University Press.

- Maison, D., Greenwald, A. G., & Bruin, R. (2001). The Implicit Association Test as a measure of implicit consumer attitudes. *Polish Psychological Bulletin, 32*, 61–69.
- Malinowski, P., & Hübner, R. (2001). The effect of familiarity on visual-search
 performance: Evidence for learned basic features. *Perception & Psychophysics*, 63, 458-463.
- Mandler, G., Nakamura, Y., & Van Zandt, B.J.S. (1987). Nonspecific effects of exposure on stimuli that cannot be recognized. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 646-648.
- McConnell, A. R., & Leibold, J. M. (2001). Relations among the Implicit Association Test, discriminatory behavior, and explicit measures of racial attitudes. *Journal of Experimental Social Psychology*, 37, 435–442.

- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207-238.
- Mierke, J., & Klauer, K. C. (2001). Implicit association measurement with the IAT: Evidence for effects of executive control processes. *Zeitschrift für Experimentelle Psychologie, 48*, 107–122.
- Mierke, J., & Klauer, K.C. (2003). Method-specific variance in the Implicit Association Test. Journal of Personality and Social Psychology, 85, 1180–1192.
- Minear, M. & Park, D. C. (2004). A lifespan database of adult facial stimuli. Behavior Research Methods, Instruments, and Computers, 36, 630-633.
- Mitchell, C. J. (2004). Mere acceptance produces apparent attitude in the Implicit Association Test. *Journal of Experimental Social Psychology*, 40, 366–373.
- Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2003). Contextual variations in implicit evaluation. *Journal of Experimental Psychology: General, 132*, 455–469.
- Nisbett, R.E., & Wilson, T.D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*, 231–259.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting intergroup implicit attitudes and beliefs from a demonstration Web site. *Group Dynamics*, *6*, 101-115.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin*, *31*, 166–180.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10, 104-114.

- Nosofsky, R. M. (1986). Attention, similarity, and the identification categorization relationship. *Journal of Experimental Psychology: General, 115,* 39-57.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101(1), 53-79.
- Nuerk, H.C., Iversen, W., & Willmes, K. (2004). Notational modulation of the SNARC and the MARC (linguistic markedness of response codes) effect. *Quarterly Journal of Experimental Psychology*, 57A, 853–863.
- O'Brien, R. G., & Kaiser, M. K. (1985). MANOVA method for analyzing repeated measures designs: An extensive primer. *Psychological Bulletin*, 97, 316–333.
- Ohman, A., Flykt, A., & Esteves, F. (2001a). Emotion drives attention: Detecting the snake in the grass. *Journal of Experimental Psychology: General*, 130, 466-478.
- Ohman, A., Flykt, A., & Esteves, F. (2001b). The face in the crowd revisited: A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology*, 8, 381-396.
- Ottaway, S. A., Hayden, D. C., & Oakes, M. A. (2001). Implicit attitudes and racism: Effects of word familiarity and frequency on the Implicit Association Test. *Social Cognition, 19*, 97–144.
- Paulhus, D.L., (1984). Two component models of social desirable responding. *Journal of Personality and Social Psychology*, 46, 598–609.
- Phelps, E.A., Cannaistraci, C.J., & Cunningham, W.A. (2003). Intact performance on an indirect measure of race bias following amygdala damage. *Neuropsychologia*, 41, 203-208.

- Phelps, E.A., O'Connor, K.J., Cunningham, W.A., Funayma, E.S., Gatenby, J.C., Gore, J.C., Banaji, M.R. (2000). Performance on indirect measures of race evaluation predicts amygdala activity. *Journal of Cognitive Neuroscience*, 12, 1-10.
- Phillips, Barbara J. (2000). The Impact of Verbal Anchoring on Consumer Response to Image Ads. *Journal of Advertising*, 29, 15-24.
- Rauschenberger, R., & Yantis, S. (2006). Perceptual encoding efficiency in visual search. Journal of Experimental Psycology: General, 135, 116-131.
- Reber, R., Winkielman, P. & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, *9*, 45-48.
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, *3*, 382-407.
- Rothermund, K., & Wentura, D. (2001). Figure-ground asymmetries in the Implicit Association Test (IAT). *Zeitschrift fur Experimentelle Psychologie*, 48, 94-106.
- Rothermund, K., & Wentura, D. (2004). Underlying processes in the Implicit Association Test (IAT): Dissociating salience from associations. *Journal of Experimental Psychology: General*, 133, 139-165.
- Rubin, D.C., & Friendly, M. (1986). Predicting which words get recalled: Measures of free recall availability, goodness, emotionality, and pronouncability for 925 nouns.
 Memory & Cognition, 14, 79 94.
- Rudman, L. A., Greenwald, A. G., Mellott, D. S., & Schwartz, J. L. K. (1999). Measuring the automatic components of prejudice: Flexibility and generality of the Implicit Association Test. *Social Cognition*, 17, 437–465.

- Sargent, M. J., Kahan, T. A., & Mitchell, C. J. (2007). The mere acceptance effect: Can it influence responses on racial Implicit Association Tests? *Journal of Experimental Social Psychology*, 43, 787-793.
- Shen, J., & Reingold, E. M. (2001). Visual search asymmetry: The influence of stimulus familiarity and low-level features. *Perception & Psychophysics*, 63, 464-475.
- Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied*, 75, 1-42.
- Strayer, D. L., & Johnston, W. A. (2000). Novel popout is an attention-based phenomenon: An ERP analysis. *Perception & Psychophysics*, 62, 459–470.
- Tipples, J., Young, A. W., Quinlan, P., Broks, P., & Ellis, A. W. (2002). Searching for threat. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 55A, 1007-1026.
- Teachman, B. A., Gregg, A. P., & Woody, S. R. (2001). Implicit associations for fearrelevant stimuli among individuals with snake and spider fears. *Journal of Abnormal Psychology*, 110, 226–235.
- Trabasso, T., Rollins, H., & Shaughnessy, E. (1971). Storage and verification stages in processing concepts. *Cognitive Psychology*, *2*, 239-289.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95, 15-48.
- Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, 114, 285-310.

- Unkelbach, C., Fiedler, K., Bayer, M., Stegmüller, M., & Danner, D. (in press). Why positive information is processed faster: The density hypothesis. *Journal of Personality & Social Psychology*.
- Valentine, T., & Endo, M. (1992). Toward an exemplar model of face processing: The effects of race and distinctiveness. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 44A, 161-204.
- Wang, Q., Cavanagh, P., & Green, M. (1994). Familiarity and pop-out in visual search. Perception and Psychophysics, 56, 495-500.
- Whaley, C. P. (1978). Word–nonword classification time. *Journal of Verbal Learning and Verbal Behavior*, 17, 143–154.
- Whittlesea, B.W.A. (1993). Illusions of familiarity. Journal of Experimental Psychology: Learning, Memory, & Cognition, 19, 1235-1253.
- Williams, M.A., Moss, S.A., Bradshaw, J.L., & Mattingley, J.B. (2005). Look at me, I'm smiling: Searching for threatening and non-threatening facial expressions. *Visual Cognition*, 12, 29-50.
- Wilson, T.D., & Dunn, E.W. (2004). Self-knowledge: Its limits, value, and potential for improvement. Annual Review of Psychology, 55, 493–518.
- Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face:
 Psychophysiological evidence that processing facilitation increases positive affect.
 Journal of Personality and Social Psychology, 81, 989–1000.
- Wittenbrink, B., Judd, C.M., & Park, B. (2001). Spontaneous prejudice in context:
 Variability in automatically activated attitudes. *Journal of Personality and Social Psychology*, 81, 815–827.

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

- Zajonc, R.B. (1968). Attitudinal effects of mere exposure. Journal of Personality and Social Psychology, 9, 1-27.
- Zajonc, R. B. (2001). Mere exposure: A gateway to the subliminal. *Current Directions in Psychological Science*, 10, 224–228.

Appendix A: Target stimuli for Experiments 12a-12b (Row/Column mere acceptance

IAT and split IAT)

Row Same stimuli:

32	98	21	77
32	45	63	77
12	45	63	19
23	57	25	67
23	14	32	67
48	14	32	83

Column Same stimuli

11	73	98	35
42	47	82	56
24	34	29	63
72	41	31	59
27	66	16	92
88	14	63	25

Appendix B: Summary tables of statistical analysis

Experiment 1a: Elephant visual search task

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
SameLive-SameDead (RT)	-4.69	200.22	57.80	11	122.52, 131.90	-0.08	.94
SameLive-SameDead (Errors)	-1.25	1.71	0.49	11	016, -2.34	-2.53	.03
DifferentLive-DifferentDead (RT)	-60.49	76.94	22.21	11	-11.61, 109.38	-2.72	.02
DifferentLive-DifferentDead (Errors)	-1.17	1.53	0.44	11	-0.20, -2.14	-2.65	.02

Experiment 1b: Cow visual search task

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
SameBig-SameSmall (RT)	69.95	127.07	44.93	7	-176.18, 36.28	1.56	.16
SameBig-SameSmall (Errors)	-1.88	3.04	1.08	7	-0.67, 4.42	1.74	.13
DifferentBig-DifferentSmall	120.99	55.97	19.79	7	-167.78, -	6.11	.000
(RT)					74.19		
DifferentBig-DifferentSmall	2.25	2.49	0.88	7	-4.33, -0.17	2.55	.04
(Errors)							

Experiment 1c: Elephant classification task

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Live-Dead (RT)	-17.64	24.00	6.41	13	-3.78, -31.50	-2.75	.02
Live-Dead (Errors)	0.21	2.64	0.71	13	1.74, -1.31	0.34	.77

Experiment 1d: Cow classification task

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Big-Small (RT)	-21.02	25.44	7.34	11	-4.86, -37.19	-2.86	.02
Big-Small (Errors)	-0.33	2.74	0.79	11	1.41, -2.08	42	.68

Experiment 2a: Elephant/word IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	521369.43	7	74481.34		
	Within (IAT)	283426.24	1		10.58	.01
	Error	187478.97	7	26782.71		
Errors	Between	59.22	7	8.46		
	Within (IAT)	1.53	1		3.94	.09
	Error	2.72	7	0.39		

Experiment 2b: Cow/word IAT

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between (Cow breed)	9019.63	1		0.110	.75
	Error	1.14e+6	14	81633.38		
	Within (IAT)	191013.34	1	191013.34	8.86	.01
	BetweenWithin	54663.23	1		2.54	.13
	Error	301820.06	14	21558.58		
Errors	Between (Cow breed)	3.06	1		0.48	.50
	Error	89.19	14	6.37		
	Within (IAT)	6.25	1		1.84	.20
	BetweenWithin	0.56	1	0.56	0.17	.69
	Error	47.69	14			

ANOVA summary table: averaged across groups

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.15e+6	15	76792.46		
	Within (IAT)	191013.34	1		8.04	.01
	Error	356483.29	15	23765.55		
Errors	Between	92.25	15	6.15		
	Within (IAT)	6.25	1		1.94	.18
	Error	48.25	15	3.22		

Experiment 2c: Word/nonword classification task

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Word-Nonword (RT)	-20.76	40.06	8.18	23	-3.84, -37.67	-2.54	.02
Word-Nonword (Errors)	0.67	4.08	0.83	23	2.39, -1.06	0.80	.43

Experiment 3: Elephant/cow IAT

ANOVA	summary	table:	between	groups	analysis
	~			() · · · (· ·)	

Variable	Source	SS	DF	MS	F	Р
RT	Between (Cow breed)	284385.86	1		2.59	.12
	Error	3.29e+6	30	109823.87		
	Within (IAT)	272957.29	1		6.49	.02
	BetweenWithin	31438.96	1		0.75	.39
	Error	1.26e+6	30	42062.86		
Errors	Between (Cow breed)	0.01	1		0.001	.98
	Error	220.73	30	7.36		
	Within (IAT)	0.01	1		0.004	.95
	BetweenWithin	1.32	1		0.62	.44
	Error	63.92	30	2.13		

ANOVA summary table: averaged across groups

Variable	Source	SS	DF	MS	F	Р
RT	Between	3.578e+6	31	115454.90		
	Within (IAT)	272957.29	1		6.54	.02
	Error	1.29e+6	31	41720.16		
Errors	Between	220.74	31	7.12		
	Within (IAT)	0.008	1		0.004	.84
	Error	65.24	31	2.11		

Experiment 4a: Flower/insect IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	139709.99	7	19958.57		
	Within (IAT)	804879.50	1		43.60	.000
	Error	129227.91	7	18461.13		
Errors	Between	34.47	7	4.92		
	Within (IAT)	22.78	1		2.85	.14
	Error	55.97	7	7.80		

Flower/insect classification task

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Flower-Insect (RT)	6.58	52.48	13.12	15	34.54, -21.39	0.50	.62
Flower-Insect (Errors)	-0.69	4.50	1.12	15	1.71, -3.08	-0.61	.55

Experiment 4b: Flower/insect split IAT

Attribute classification trials

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Pleasant-Neutral (RT)	246.53	252.56	51.5 5	23	-139.89, -353.18	-4.78	.000
Pleasant-Neutral (Errors)	0.29	1.37	0.28	23	.89,29	1.05	.30
Neutral-Unpleasant (RT)	132.75	210.47	42.96	23	43.88, 221.62	3.09	.005
Neutral-Unpleasant	0.29	1.60	0.33	23	-0.38, 0.97	0.89	.38
(Errors)				_			

Pleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	4.30e+6	23	186866.11		-
	Within (IAT)	398232.96	1		20.85	.000
	Error	398232.96	23	398232.96		
Errors	Between	194.66	23	194.65		
	Within (IAT)	0.01	1		0.002	.97
	Error	105.74	23	4.60		

Unpleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	991236.55	23	43097.24		
	Within (IAT)	471429.35	1		26.44	.000
	Error	410161.86	23	17833.12		
Errors	Between	208.49	23	9.07		
	Within (IAT)	0.09	1		0.03	.86
	Error	84.16	23	3.66		

Combined analysis

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	Р
RT	Between	298941.87	1	298941.87	2.60	.11
	Error	5.29e+6	46	114981.67		
	Within	868119.40	1		47.01	1.50e-8
	(Preference Effect)					
	BetweenWithin (Salience Effect)	1542.91	1		0.08	.78
	Error	849474.28	46	18466.83		
Errors	Between	17.52	1			
	Error	403.14	46	8.76		
	Within	0.08	1		.02	.89
	(Preference Effect)					
	BetweenWithin (Salience Effect)	0.02	1		.01	.94
	Error	189.90	46	4.13		

Experiment 5a:

Go/Nogo IAT (Flight/Teeth)

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between (Flight/Teeth)	224954.63	1		0.63	.44
	Error	4.95e+6	14	353537.83		
	Within (IAT)	420405.01	1		6.58	.02
	BetweenWithin	48038.02	1		0.75	.40
	Error	893876.79	14	63848.34		
Errors	Between (Flight/Teeth)	0.25	1		0.03	.87
	Error	110.19	14	7.87		
	Within (IAT)	0.25	1		0.09	.77
	BetweenWithin	5.06	1		1.72	.21
	Error	41.19	14	2.94		

ANOVA summary table: averaged across groups

Variable	Source	SS	DF	MS	F	Р
RT	Between	5.17e+6	15	344965.62		
	Within (IAT)	420405.01	1		6.70	.02
	Error	941914.81	15	62794.32		
Errors	Between	110.44	15	7.36		
	Within (IAT)	0.25	1		0.08	.78
	Error	46.25	15	3.08		

Go/Nogo classification task (Flight/Teeth)

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between (Flight/Teeth)	112.34	1		0.005	.95
	Error	345967.63	14	24711.97		
	Within (IAT)	6211.74	1		5.08	.04
	BetweenWithin	1481.55	1		1.21	.29
	Error	17107.45	14	1221.96		
Errors	Between (Flight/Teeth)	0.50	1		0.07	.80
	Error	106.00	14			
	Within (IAT)	4.50	1		2.42	.14
	BetweenWithin	40.50	1		21.81	.000
	Error	26.00	14	1.86		

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
FlightGo – FlightNogo	1.50	2.27	0.80	7	3.40,40	1.87	.10
(Errors)							
TeethGo – TeethNogo	3.00	1.51	0.53	7	-1.74, -	-5.6	.001
(Errors)					4.26		

Experiment 5b: Go/Nogo split IAT (Flight/Teeth)

Pleasant IAT

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between (Flight/Teeth)	70698.89	1		0.40	.54
	Error	2.48e+6	14	176945.37		
	Within (IAT)	548090.71	1		12.21	.004
	BetweenWithin	85555.92	1		1.91	.19
	Error	628301.48	14	44878.67		
Errors	Between (Flight/Teeth)	4.00	1		0.51	.49
	Error	108.94	14	7.78		
	Within (IAT)	7.56	1		1.79	.20
	BetweenWithin	0.25	1		0.06	.81
	Error	59.19	14	4.23		

ANOVA summary table: averaged across groups

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.55e+6	15	169862.27		
	Within (IAT)	548090.71	1		11.52	.004
	Error	713857.41	15	47590.49		
Errors	Between	112.94	15	7.53		
	Within (IAT)	7.56	1		1.91	.19
	Error	59.44	15	3.96		

Unpleasant IAT

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between (Flight/Teeth)	264.24	1		0.001	.98
	Error	2.64e+6	14	188592.53		
	Within (IAT)	379579.51	1		6.32	.03
	BetweenWithin	1885.80	1		0.03	.87
	Error	840732.23	14	60052.30		
Errors	Between (Flight/Teeth)	2.64	1			
	Error	68.47	14	4.89		
	Within (IAT)	40.64	1		6.77	.02
	BetweenWithin	17.02	1		2.83	.12
	Error	84.09	14	6.00		

ANOVA summary table: averaged across groups

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.64e+6	15	176037.83		
	Within (IAT)	379582.29	1		6.76	.02
	Error	842615.24	15	56174.35		
Errors	Between	71.11	15	4.74		
	Within (IAT)	40.64	1		6.03	.03
	Error	101.11	15	6.74		

Experiment 5b: Go/Nogo split IAT - continued

Combined analysis

ANOVA summary table:

Flight/Teeth as the between groups factor (B1)

Pleasant IAT/Unpleasant IAT as the between groups factor (B2)

Variable	Source	SS	DF	MS	F	Р
RT	Between (B1)	39803.78	1		0.22	.64
	Between (B2)	56.80	1		0.00	.99
	Error	5.11753e+6	28	182768.95		
	Within	7716.70	1		0.15	.70
	(Preference Effect)					
	Between(B1)Within	56422.87	1		1.08	.31
	(Flight/Teeth x Preference)					
	Between(B2)Within	919953.52	1		17.53	.000
	(Salience Effect)					
	Between(B1)Within(Salience Effect)	31018.85	1		0.59	.45
	(Flight/Teeth x Salience)					
	Error	1.47e+6	28	52465.49		
Errors	Between (B1)	6.57	1		1.04	.32
	Between (B2)	3.45	1		0.54	.47
	Error	177.47	28	6.34		
	Within	6.57	1		1.40	.25
	(Preference Effect)					
	Between(B1)Within	0.38	1		0.08	.78
	(Flight/Teeth x Preference)					
	Between(B2)Within	41.63	1		8.89	.006
	(Salience Effect)					
	Between(B1)Within(Salience Effect)	29.07	1		6.21	.02
	(Flight/Teeth x Salience)					
	Error	131.0	28	4.68		

Experiment 6a:

Insect/nonword IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.04598e+6	7	292282.93		
	Within (IAT)	443249.65	1		17.67	.004
	Error	175560.82	7	25080.12		
Errors	Between	37.47	7	5.35		
	Within (IAT)	3.78	1		1.36	.28
	Error	19.47	7	2.78		

Experiment 6a: - continued

Insect/nonword classification task

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Insect-Nonword (RT)	-56.01	66.15	20.92	9	-8.69, 103.33	-2.68	.03
Insect -Nonword (Errors)	0.00	3.80	1.20	9	2.72, -2.72	0.00	1.00

Experiment 6b: Insect/nonword split IAT

Pleasant IAT

ANOVA summary table

Source	SS	DF	MS	F	Р
Between	966726.44	15	64448.43	35.29	.000
Within (IAT)	1.29e+6	1			
Error	550058.78	15	36670.59		
Between	210.48	15	14.03		
Within (IAT)	11.39	1		6.36	.02
Error	26.86	15	1.79		
	Source Between Within (IAT) Error Between Within (IAT) Error	Source SS Between 966726.44 Within (IAT) 1.29e+6 Error 550058.78 Between 210.48 Within (IAT) 11.39 Error 26.86	SourceSSDFBetween966726.4415Within (IAT)1.29e+61Error550058.7815Between210.4815Within (IAT)11.391Error26.8615	SourceSSDFMSBetween966726.441564448.43Within (IAT)1.29e+61Error550058.781536670.59Between210.481514.03Within (IAT)11.391Error26.86151.79	SourceSSDFMSFBetween966726.441564448.4335.29Within (IAT)1.29e+611Error550058.781536670.59Between210.481514.03Within (IAT)11.3916.36Error26.86151.79

Unpleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.018e+6	15	67899.55		
	Within (IAT)	1.12e+6	1		44.25	7.71e-6
	Error	380926.47	15	25395.10		
Errors	Between	112.25	15	7.48	3.17	.10
	Within (IAT)	18.06	1			
	Error	85.44	15	5.70		

Combined analysis

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	Р
RT	Between	626421.32	1		9.47	.004
	Error	1.99e+6	30	66173.99		
	Within (Preference Effect)	3011.57	1		0.10	.75
	BetweenWithin (Salience Effect)	2.41e+6	1		77.82	7.79e-10
	Error	930985.26	30	31032.84		
Errors	Between	15.82	1		1.47	.24
	Error	322.73	30	10.76		
	Within (Preference Effect)	0.38	1		0.10	.75
	BetweenWithin (Salience Effect)	29.07	1		7.77	.01
	Error	112.30	30	3.74		

Experiment 7a:

Familiar/unfamiliar paintings IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	3.18e+6	15	212035.65		
	Within (IAT)	389279.48	1		6.38	.023
	Error	915432.04	15	61028.80		
Errors	Between	95.61	15	6.37		
	Within (IAT)	17.02	1		11.74	.004
	Error	21.73	15	1.45		

Familiar/unfamiliar paintings classification task

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Familiar-Unfamiliar (RT)	-24.17	39.91	9.98	15	-2.91, -45.44	-2.42	.03
Familiar-Unfamiliar (Errors)	0.44	1.26	0.32	15	1.11, -0.24	1.39	.19

Experiment 7b: Familiar/unfamiliar paintings split IAT

Pleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	4.21e+6	15	280828.65		
	Within (IAT)	1.75e+6	1		18.40	.001
	Error	1.43e+6	15	95376.93		
Errors	Between	154.48	15	10.30		
	Within (IAT)	50.77	1		23.81	. 000
	Error	31.98	15	2.13		

Unpleasant IAT

A	V	0	VA	summary	tabl	le
---	---	---	----	---------	------	----

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.53024e+6	15	102016.54		
	Within (IAT)	12.70	1		0.000	.99
	Error	1.03e+6	15	68932.02		
Errors	Between	104.25	15	6.95		
	Within (IAT)	2.25	1		0.75	.40
	Error	45.25	15	3.02		

Experiment 7b: Familiar/unfamiliar paintings split IAT - continued

Combined analysis

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	Р
RT	Between	638038.30	1		3.33	.08
	Error	5.74e+6	30	191422.60		
	Within (Preference Effect)	872655.75	1		10.62	.003
	BetweenWithin (Salience Effect)	882096.82	1		10.74	.003
	Error	2.46e+6	30	82154.47		
Errors	Between	1.32	1		0.15	.70
	Error	258.73	30	8.62		
	Within (Preference Effect)	15.82	1		6.15	.02
	BetweenWithin (Salience Effect)	37.20	1		14.45	.001
	Error	77.23	30	2.57		

Experiment 8a:

Even/odd number IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.81e+6	11	255622.08		
	Within (IAT)	2.52e+6	1		12.05	.005
	Error	2.30e+6	11	209176.66		
Errors	Between	87.17	11	7.92		
	Within (IAT)	16.33	1		5.96	.03
	Error	30.17	11	2.74		

Even/odd number classification task

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Even-Odd (RT)	-21.78	63.72	15.02	17	9.91, -53.46	-1.45	.17
Even-Odd (Errors)	-1.171	2.81	0.66	17	0.23, -2.57	-1.76	.10

Experiment 8b: Even/odd number split IAT

Pleasant IAT

	ANOVA	summary	table
--	-------	---------	-------

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.63e+6	15	108694.09		
	Within (IAT)	883313.98	1		11.36	.004
	Error	1.17e+6	15	77737.70		
Errors	Between	84.75	15	5.65		
	Within (IAT)	1.00	1		0.88	.36
	Error	17.00	15	1.13		

Experiment 8b: Even/odd number split IAT - continued

Unpleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.82e+6	15	188220.48		
	Within (IAT)	250963.07	1		9.10	.009
	Error	413537.54	15	27569.17		
Errors	Between	63.61	15	4.24		
	Within (IAT)	3.52	1		2.74	0.12
	Error	19.23	15	1.28		

Combined analysis

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	Р
RT	Between	81257.26	1		0.55	
	Error	4.45e+6	30	148457.29		
	Within (Preference Effect)	1.04e+6	1		19.71	.000
	BetweenWithin (Salience Effect)	96310.33	1		1.83	.19
	Error	1.58e+6	30	52653.43		
Errors	Between	4.13	1			
	Error	148.36	30	4.95	0.84	.37
	Within (Preference Effect)	4.13	1		3.42	.07
	BetweenWithin (Salience Effect)	0.38	1		0.32	.58
	Error	36.23	30	1.21		

Experiment 9a:

Race IAT with faces

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.17e+6	7	166769.19		
	Within (IAT)	290592.90	1		12.50	.01
	Error	162746.76	7	23249.54		
Errors	Between	22.38	7	3.20		
	Within (IAT)	3.13	1		1.02	.35
	Error	21.38	7	3.05		

Race classification task with faces

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
White-Black (RT)	29.55	35.35	8.33	17	47.13, 11.97	3.55	.002
White-Black (Errors)	-0.33	1.65	0.39	17	0.48, -1.15	-0.86	.40

Experiment 9b:

Race IAT with names

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	3.72e+6	11	337988.02		
	Within (IAT)	1.04e+6	1		16.63	.002
	Error	6875600.00	11	62509.09		
Errors	Between	127.73	11	11.61		
	Within (IAT)	3.52	1		0.77	.40
	Error	50.23	11	4.57		

Race classification task with names

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
White-Black (RT)	24.11	53.06	11.31	21	47.63, 0.59	2.13	.05
White-Black (Errors)	-0.86	3.22	0.69	21	0.57, -2.29	-1.26	.22

Race IAT with faces vs. Race IAT with names

ANOUA		1. 1	1		
ANUVA	summarı	' tanie:	netween	groups	analysis
	<i>y</i>			$S' \circ np \circ$	

Variable	Source	SS	DF	MS	F	Р
RT	Between (Faces/names)	1.51e+6	1		5.55	.03
	Error	4.89e+6	18	271403.09		
	Within (IAT)	1.13e+6	1		23.89	.000
	BetweenWithin	51608.17	1		1.09	.31
	Error	850350.21	18			
Errors	Between (Faces/names)	32.03	1		3.84	.07
	Error	150.10	18	8.34		
	Within (IAT)	6.53	1		1.64	.22
	BetweenWithin	0.03	1		0.01	.92
	Error	71.60	18	3.98		

Experiment 9c: Race split IAT with faces

Pleasant IAT

ANOVA	summary	table
-------	---------	-------

Variable	Source	SS	DF	MS	F	P
RT	Between	1.65e+6	15	109717.60		
	Within (IAT)	2058.13	1		0.19	.67
	Error	165799.64	15	11053.31		
Errors	Between	100.23	15	6.68		
	Within (IAT)	0.14	1		0.10	.76
	Error	20.61	15	1.37		
Experiment 9c: Race split IAT with faces - continued

Unpleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.90e+6	15	126354.87		
	Within (IAT)	560532.63	1	560532.63	8.14	.012
	Error	1.03e+6	15	68851.49		
Errors	Between	59.86	15	3.99		
	Within (IAT)	3.52	1		1.07	.32
	Error	49.23	15	3.28		

Combined analysis

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	Р
RT	Between	12814.62	1		0.11	.74
	Error	3.54e+6	30	118036.24		
	Within (Preference Effect)	315260.80	1		7.89	.009
	BetweenWithin (Salience Effect)	247329.97	1		6.19	.02
	Error	1.199e+6	30	39952.40		
Errors	Between	0.28	1			
	Error	160.09	30	5.34	.05	.83
	Within (Preference Effect)	2.53	1		1.09	.31
	BetweenWithin (Salience Effect)	1.13	1		0.48	.49
	Error	69.84	30	2.33		

Experiment 9d: Race split IAT with names

Pleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.34e+6	15	155821.89		
	Within (IAT)	449.79	1		0.01	.92
	Error	1.06212e+6	15	70808.56		
Errors	Between	78.36	15	5.22		
	Within (IAT)	1.27	1		0.84	.37
	Error	22.48	15	1.50		

Unpleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	5.49e+6	15	366071.22		
	Within (IAT)	2.41	1		31.19	.000
	Error	1.15771e+6	15	77180.68		
Errors	Between	107.50	15	7.18		
	Within (IAT)	27.56	1		10.10	.006
	Error	40.94	15	2.73		

Experiment 9d: Race split IAT with names - continued

Combined analysis

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	P
RT	Between	432101.83	1			
	Error	7.83e+6	30	260946.56	1.66	.21
	Within (Preference Effect)	1.24e+6	1		16.72	.000
	BetweenWithin (Salience Effect)	1.17e+6	1		15.83	.000
	Error	2.22	30	73994.62		
Errors	Between	1.32	1			
	Error	185.86	30	6.20	0.21	.65
	Within (Preference Effect)	20.32	1		9.61	.004
	BetweenWithin (Salience Effect)	8.51	1		4.02	.054
	Error	63.42	30	2.11		

Experiment 10a:

Age IAT with faces

ANOVA summary table

Variable	Source	SS	DF	MS	F	P
RT	Between	605815.87	11	55074.17		
	Within (IAT)	679246.01	1		28.94	.000
	Error	258140.92	11	23467.36		
Errors	Between	83.42	11	7.58		
	Within (IAT)	0.75	1		0.40	.54
	Error	20.75	11	1.89		

Age classification task with faces

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Young-Old (RT)	30.98	24.42	7.72	9	48.45, 13.5	4.01	.003
Young-Old (Errors)	-0.20	2.10	0.66	9	1.30, -1.70	-0.30	.77

Experiment 10b:

Age IAT with names

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.57e+6	7	366611.29		
	Within (IAT)	1.15e+6	1		19.24	.003
	Error	416793.50	7	59541.93		
Errors	Between	81.00	7	11.57		
	Within (IAT)	32.00	1		13.18	.008
	Error	17.00	7	2.43		

Experiment 10b: - continued

Age classification task with names

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Young-Old (RT)	-11.72	85.67	20.20	17	30.88, -54.32	-0.58	.57
Young-Old (Errors)	-0.33	3.50	0.82	17	1.41, -2.07	-0.40	.69

Age IAT with faces vs. Age IAT with names

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between (Faces/names)	1.48e+6	1		6.73	.02
	Error	3.08e+6	14	220191.27		
	Within (IAT)	1.16097e+6	1		31.58	.000
	BetweenWithin	190140.53	1		5.17	.04
	Error	514754.87	14	36768.21		
Errors	Between (Faces/names)	60.06	1		7.98	.01
	Error	105.38	14	7.52		
	Within (IAT)	20.25	1		7.46	.02
	BetweenWithin	12.25	1		4.51	.052
	Error	38.00	14	2.71		

Experiment 10c: Age split IAT with faces

Pleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.83e+6	15	122152.98		
	Within (IAT)	34699.13	1		1.05	.32
	Error	496169.17	15	33077.95		
Errors	Between	156.94	15	156.93		
	Within (IAT)	0.06	1		0.02	.89
	Error	38.44	15	2.56		

Unpleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	755591.63	15	50372.78		
	Within (IAT)	932672.00	1		47.47	5.26e-6
	Error	294717.33	15	19647.82		
Errors	Between	99.50	15	6.63		
	Within (IAT)	5.06	1		2.08	.17
	Error	36.44	15	2.43		

Experiment 10c: Age split IAT with faces - continued

Combined analysis

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	Р
RT	Between	47144.28	1	· · · · · · · · · · · · · · · · · · ·	0.55	.46
	Error	2.59e+6	30	86262.88		
	Within (Preference Effect)	663582.48	1		25.17	.000
	BetweenWithin (Salience Effect)	303788.64	1	26362.88	11.52	.002
	Error	790886.51	30			
Errors	Between	2.53	1		0.30	.59
	Error	256.44	30	8.55		
	Within (Preference Effect)	3.13	1	3.13	1.25	.27
	BetweenWithin (Salience Effect)	2.00	1	2.00	0.80	.38
	Error	74.88	30	2.50		

Experiment 10d: Age split IAT with names

Pleasant IAT

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.57e+6	15	171056.31		
	Within (IAT)	436220.62	1		10.61	.005
	Error	616771.50	15			
Errors	Between	94.50	15	6.30		
	Within (IAT)	0.25	1		0.07	.80
	Error	54.25	15	3.62		

ANOVA s	ummary table	?
---------	--------------	---

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.77e+6	15	117959.44		
	Within (IAT)	55477.14	1		1.02	.33
	Error	812649.75	15	54176.65		
Errors	Between	140.00	15	9.33		
	Within (IAT)	9.00	1		1.11	.31
	Error	122.00	15	8.13		

Experiment 10d: Age split IAT with names - continued

Combined analysis

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	Р
RT	Between	228543.30	1			
	Error	4.34+6	30	144507.88	1.58	.22
	Within (Preference Effect)	401413.25	1		8.43	.007
	BetweenWithin (Salience Effect)	90284.51	1		1.90	.18
	Error	1.43+6	30	47647.38		
Errors	Between	0.50	1		0.06	.81
	Error	234.50	30	7.82		
	Within (Preference Effect)	6.13	1		1.04	.32
	BetweenWithin (Salience Effect)	3.13	1		0.53	.47
	Error	176.25	30	5.88		

Experiment 11a:

Mere acceptance IAT (Flight/Teeth)

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between	128858.49	1		0.21	.66
	Error	8.78e+6	14	626875.19		
	Within (IAT)	3.79e+6	1		14.33	.002
	BetweenWithin	24652.32	1	264715.96	0.09	.77
	Error	3.71e+6	14			
Errors	Between	9.77	1		1.80	.20
	Error	75.97	14	5.43		
	Within (IAT)	11.40	1		2.33	.15
	BetweenWithin	0.39	1		0.08	.78
	Error	68.47	14	4.89		

ANOVA summary table: averaged across groups

Variable	Source	SS	DF	MS	F	Р
RT	Between	8.90511e+6	15	593674.08		
	Within (IAT)	3.79296e+6	1		15.25	.001
	Error	3.73067e+6	15	248711.71		
Errors	Between	85.73	15	5.72		
	Within (IAT)	11.39	1		2.48	.14
	Error	68.86	15	4.59		

Mere acceptance classification task (Flight/Teeth)

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
Accept-Reject (RT)	-48.14	39.05	11.27	11	-23.33, -72.95	-4.27	.001
Accept-Reject (Errors)	1.33	5.10	1.47	11	4.58, -1.91	0.91	.39

Experiment 11b: Mere acceptance split IAT (Flight/Teeth)

Pleasant IAT

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between (Flight/Teeth)	652754.45	1		2.34	.15
	Error	3.91e+6	14	279022.35		
	Within (IAT)	5.93e+6	1		34.32	.000
	BetweenWithin	740528.56	1		4.28	.06
	Error	2.42e+6	14	172932.57		
Errors	Between (Flight/Teeth)	27.56	1		4.34	.06
	Error	88.88	14	6.35		
	Within (IAT)	95.06	1		30.68	.000
	BetweenWithin	18.06	1		5.83	.03
	Error	43.38	14	3.10		

ANOVA summary table: simple effects analysis

Variable	Source	SS	DF	MS	F	Р
RT (Flight)	Between	937115.89	7	133873.70		
	Within (Flight IAT)	1.24e+6	1		5.67	.049
	Error	1.53e+6	7	219112.83		
RT (Teeth)	Between	2.97e+6	7	424170.76		
	Within (Teeth IAT)	5.43e+6	1		42.87	.000
	Error	887261.34	7	126751.62		
Errors (Flight)	Between	19.00	7	2.71		
	Within (Flight IAT)	15.13	1		4.95	.06
	Error	21.38	7	3.05		
Errors (Teeth)	Between	69.88	7			
	Within (Teeth IAT)	98.00	1		31.18	.001
	Error	22.00	7	3.14		

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between (Flight/Teeth)	66590.26	1			
	Error	5.08e+6	14	362934.92	0.18	.68
	Within (IAT)	1.75e+6	1	16.25		
	BetweenWithin	29762.75	1		0.28	.61
	Error	1.51e+6	14	107802.10		
Errors	Between (Flight/Teeth)	0.25	1		0.02	.89
	Error	175.75	14	12.55		
	Within (IAT)	9.00	1		2.29	.15
	BetweenWithin	.00	1		.00	.99
	Error	55.00	14	3.93		

Experiment 11b: Mere acceptance split IAT (Flight/Teeth) - Unpleasant IAT - continued

	ANOVA	summary	table:	averaged	across	groups
--	-------	---------	--------	----------	--------	--------

Variable	Source	SS	DF	MS	F	Р
RT	Between	5.15e+6	15	343178.61		
	Within (IAT)	1.75e+6	1		17.08	.001
	Error	1.53899e+6	15	102599.48		
Errors	Between	176.00	15	11.73		
	Within (IAT)	9.00	1		2.46	.14
	Error	55.00	15	3.67		

Combined analysis

Flight/Teeth as the between groups factor (B1)

Pleasant IAT/Unpleasant IAT as the between groups factor (B2)

Variable	Source	SS	DF	MS	F	Р
RT	Between (B1)	568158.01	1		1.77	.19
	Between (B2)		1			
	Error	8.99e+6	28	320978.40		
	Within (Preference Effect)	618771.11	1		4.41	.045
	Between(B1)Within	533605.19	1		3.80	.06
	(Flight/Teeth x Preference)					
	Between(B2)Within	7.068e+6	1		50.36	1.01e-
	(Salience Effect)					7
	Between(B1)Within(Salience Effect)	236688.34	1		1.69	.20
	(Flight/Teeth x Salience)					
	Error	3.93027e+6	28	140367.06		
Errors	Between (B1)	16.53	1		1.75	.20
	Between (B2)	1.53	1		0.16	.69
	Error	264.63	28	9.45		
	Within (Preference Effect)	22.78	1		6.48	.02
	Between(B1)Within	9.03	1		2.57	.12
	(Flight/Teeth x Preference)					
	Between(B2)Within	81.28	1		23.14	.000
	(Salience Effect)					
	Between(B1)Within(Salience Effect)	9.03	1		2.57	.12
	(Flight/Teeth x Salience)					
	Error	98.38	28	3.51		

ANOVA summary table: Pleasant IAT/Unpleasant IAT as the between groups factor

Variable	Source	SS	DF	MS	F	Р
RT	Between	40212.69	1		.124	
	Error	9.71e+6	30	323558.22		
	Within (Preference Effect)	618766.41	1		3.95	.06
	BetweenWithin (Salience Effect)	7.07e+6	1		45.11	.000
	Error	4.70e+6	30	156685.89		
Errors	Between	1.53	1		.16	.69
	Error	292.44	30	9.75		
	Within (Preference Effect)	22.78	1		5.87	.02
	BetweenWithin (Salience Effect)	81.28	1	3.88	20.94	.000
	Error	116.44	30			

Experiment 12a:

Mere acceptance IAT (Row condition)

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	689772.96	7	98539.00		
	Within (IAT)	332504.44	1		10.18	.02
	Error	228714.65	7	32673.52		
Errors	Between	36.22	7	5.17		
	Within (IAT)	2.53	1		0.72	.42
	Error	24.72	7	3.53	=	

Mere acceptance IAT (Column condition)

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.52e+6	7	217562.48		
	Within (IAT)	2.62e+6	1		14.79	.006
	Error	1.24e+6	7	177399.03		
Errors	Between	41.22	7	5.89		
	Within (IAT)	3.78	1		1.43	.27
	Error	18.47	7	2.64		

Mere acceptance Row and Column IAT effect size interaction

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.26046e+6	1			
	Error	2.21271e+6	14	158050.74	7.98	
	Within (IAT)	2.41208e+6	1		22.96	.000
	BetweenWithin	544065.95	1		5.18	.04
	Error	1.47050e+6	14	105036.28		
Errors	Between	3.06	1		0.55	
	Error	77.44	14	5.53		
	Within (IAT)	6.25	1		2.03	.18
	BetweenWithin	0.06	1		0.02	
	Error	43.19	14	3.09		

Mere acceptance classification task (Row condition)

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
RowSame-RowDifferent (RT)	-218.36	124.64	39.41	9	129.20, 307.52	-5.54	.000
RowSame-RowDifferent (Errors)	-1.00	2.18	0.69	9	1.46, - 1.66	-0.15	.89

Experiment 12a: - continued

Mere acceptance classification task (Column condition)

Paired samples t-test summary table

Variables	Mean Difference	SD	SEM	DF	95% CI	t	Р
ColumnSame- ColumnDifferent (RT)	216.00	137.59	43.5	9	314.42, 117.57	4.96	.001
ColumnSame- ColumnDifferent (Errors)	1.100	5.20	1.64	9	4.82, -2.62	0.67	.52

Experiment 12b: Mere acceptance split IAT

Pleasant IAT (Row condition)

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.69283e+6	11	153894.13		
	Within (IAT)	1.11667e+6	1		22.22	.001
	Error	552908.54	11	50264.41		
Errors	Between	104.73	11	9.52		
	Within (IAT)	15.19	1		12.79	.004
	Error	13.06	11	1.19		

Unpleasant IAT (Row condition)

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.79547e+6	11	254133.83		
	Within (IAT)	308541.84	1		5.23	.04
	Error	648916.78	11	58992.44		
Errors	Between	59.73	11	5.43		
	Within (IAT)	9.19	1		3.60	.08
	Error	28.06	11	2.55		

Combined analysis (Row condition)

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	Р
RT	Between	230.11	1		0.00	.99
	Error	4.488e+6	22	204013.98		
	Within (Preference Effect)	125633.04	1		2.30	.14
	BetweenWithin (Salience Effect)	1.30e+6	1		23.79	.00
	Error	1.20e+6	22	54628.42		
Errors	Between	5.04	1		0.67	.42
	Error	164.46	22	7.48		
	Within (Preference Effect)	0.38	1		0.20	.66
	BetweenWithin (Salience Effect)	24.00	1		12.84	.002
	Error	41.13	22	1.87		

Experiment 12b: Mere acceptance split IAT - continued

Pleasant IAT (Column condition)

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	2.04e+6	11	185353.40		
	Within (IAT)	3.45e+6	1		20.57	.001
	Error	1.84e+6	11	167488.09		
Errors	Between	58.92	11	5.36		
	Within (IAT)	21.33	1		6.07	.03
	Error	38.66	11	3.52		

Unpleasant IAT (Column condition)

ANOVA summary table

Variable	Source	SS	DF	MS	F	Р
RT	Between	1.91e+6	11	173789.33		
	Within (IAT)	80229.77	1		0.64	.44
	Error	1.38e+6	11	125137.33		
Errors	Between	46.23	11	4.20		
	Within (IAT)	2.52	1		0.78	.40
	Error	35.73	11	3.25		

Combined analysis (Column condition)

ANOVA summary table: between groups analysis

Variable	Source	SS	DF	MS	F	<u> </u>
RT	Between	93202.24	1		0.52	.48
	Error	3.95e+6	22	179571.36		
	Within (Preference Effect)	2.29e+6	1		15.64	.001
	BetweenWithin (Salience Effect)	1.24e+6	1		8.46	.008
	Error	3.22e+6	22	146312.71		
Errors	Between	23.01	1		4.82	.04
	Error	105.15	22	4.78		
	Within (Preference Effect)	4.59	1		1.36	.26
	Between Within (Salience Effect)	19.26	1		5.70	.03
	Error	74.40	22	3.38		

Appendix C: Raw reaction time data

		Mean read	tion time (m	s)		Number of e	errors		
Condition		8 Live (Same)	8 Dead (Same)	7 Live, 1 Dead (Different)	7 Dead, 1 Live (Different)	8 Live (Same)	8 Dead (Same)	7 Live, 1 Dead (Different)	7 Dead, 1 Live (Different)
Subject	Counter- balancing condition								
1	1	927.44	932.58	805.27	919.79	3	4	2	4
2	2	1386.30	1508.03	1300.38	1303.10	7	11	7	8
3	3	1712.23	1388.91	1456.82	1581.77	1	1	3	4
4	4	782.11	739.98	669.02	645.09	3	2	1	2
5	1	1008.70	1211.07	1026.15	1108.83	0	3	1	2
6	2	1297.19	1079.62	1051.23	1138.39	1	0	1	4
7	3	1088.81	1355.53	1302.07	1311.59	0	2	1	2
8	4	1575.91	1931.34	1496.15	1679.32	1	1	0	1
9	1	901.47	798.56	1065.78	1110.29	1	2	3	6
10	2	956.51	874.95	994.87	1125.67	2	5	3	2
11	3	820.78	823.76	735.33	798.06	2	2	3	1
12	4	1061.24	930.65	1137.60	1044.69	1	4	3	6
	mean	1126.56	1131.25	1086.72	1147.22	1.83	3.08	2.33	3.50

Experiment 1a: Elephant visual search task

Experiment 1b: Cow visual search task

		Mean read	Mean reaction time (ms)				Number of errors			
Condition		8 Small (Same)	8 Big (Same)	7 Small, 1 Big (Different)	7 Big, 1 Small (Different)	8 Small (Same)	8 Big (Same)	7 Small, 1 Big (Different)	7 Big, 1 Small (Different)	
Subject	Counter- balancing condition									
1	1	814.13	892.72	876.51	1004.61	3	1	1	4	
2	2	766.30	744.28	819.36	979.52	1	0	0	4	
3	3	749.74	842.87	689.91	824.27	1	2	2	4	
4	4	851.54	836.47	762.98	888.07	0	0	0	7	
5	5	757.00	712.52	680.24	793.89	4	1	3	2	
6	6	1477.26	1835.24	1405.66	1538.17	0	2	0	0	
7	7	793.69	860.85	766.83	759.90	9	2	6	8	
8	8	975.54	1019.87	929.23	1110.19	7	2	4	5	
	mean	898.15	968.10	866.34	987.33	3.13	1.25	2.00	4.25	

Experiment 1c: Elephant classification task

		Mean reaction t	ime (ms)	Number of erro	ors
Condition		Live	Dead	Live	Dead
Subject	Counter- balancing condition				
1	1	401.97	468.72	4	3
2	2	433.30	493.10	2	2
3	1	738.92	732.09	5	3
4	2	607.74	639.49	2	1
5	1	360.39	366.78	4	3
6	2	445.67	464.55	2	8
7	1	594.18	589.50	9	3
8	2	433.77	439.89	1	1
9	1	438.07	429.59	2	3
10	2	502.34	519.56	1	3
11	1	414.65	428.82	3	4
12	2	416.41	437.41	4	2
13	1	506.01	496.30	1	2
14	2	451.93	486.54	2	1
	mean	481.81	499.45	3.00	2.79

		Mean reaction t	ime (ms)	Number of erro	ors
Condition		Big	Smail	Big	Small
Subject	Counter- balancing condition			· · · · · · · · · · · · · · · · · · ·	
1	1	391.33	428.64	4	6
2	2	553.16	565.49	1	4
3	1	445.76	486.86	1	2
4	2	485.46	526.53	0	4
5	1	463.93	471.37	3	1
6	2	529.61	560.65	3	2
7	1	393.97	450.19	2	3
8	2	480.31	443.80	3	1
9	1	523.85	554.15	12	6
10	2	490.91	494.97	3	5
11	1	460.30	455.97	1	1
12	2	549.24	581.48	4	6
	mean	480.65	501.67	3.08	3.42

Experiment 1d: Cow classification task

Experiment 2a: Elephant/word IAT

		Mean read	lean reaction time (ms)				Number of errors			
Condition		Live Word	Live Nonword	Dead Word	Dead Nonword	Live Word	Live Nonword	Dead Word	Dead Nonword	
Subject	Counter- balancing condition									
1	1	661.16	1076.96	1072.57	771.67	1	4	1	5	
2	2	455.38	685.30	675.33	496.28	2	2	2	2	
3	3	608.45	679.57	651.30	619.66	2	4	2	2	
4	4	673.79	829.83	819.00	563.48	4	3	4	2	
5	5	536.93	518.65	589.87	441.87	4	5	9	9	
6	6	556.56	721.78	667.17	595.37	4	5	2	2	
7	7	556.56	1017.34	1108.64	595.37	4	2	7	2	
8	8	838.76	940.63	780.90	852.00	2	5	3	6	
	mean	610.95	808.76	795.60	616.96	2.88	3.75	3.75	3.75	

Experiment 2b: Cow/word IAT

		Mean read	ction time (ms)		Number of e	errors		
Condition		Big Word	Big Nonword	Small Word	Small Nonword	Big Word	Big Nonword	Small Word	Small Nonword
Subject	Counter- balancing condition								
1	1	821.31	831.63	686.68	801.64	3	0	0	3
2	2	755.87	1094.97	892.26	710.13	0	3	4	1
3	3	614.91	807.23	637.60	630.93	0	1	1	2
4	4	662.94	750.43	735.71	604.23	0	2	4	0
5	5	607.71	610.00	537.61	656.90	1	0	1	1
6	6	924.88	1062.35	961.55	824.88	0	0	0	0
7	7	679.21	691.57	599.24	655.07	2	2	7	3
8	8	638.93	806.87	686.62	765.62	4	1	3	3
9	1	497.69	669.79	637.08	474.37	3	3	7	2
10	2	579.23	641.68	694.06	566.57	1	1	1	1
11	3	594.86	1120.50	1035.52	654.43	2	4	1	3
12	4	790.19	1050.77	980.04	823.24	5	5	4	2
13	5	808.86	682.34	766.97	920.60	3	0	2	1
14	6	1138.81	1003.22	1096.69	983.09	1	0	0	0
15	7	553.27	629.81	747.07	569.00	2	6	3	2
16	8	573.06	818.57	1010.14	597.74	1	3	3	0
	mean	702.61	829.48	794.05	702.40	1.75	1.94	2.56	1.50

		Mean reaction t	ime (ms)	Number of erro	ors
Condition		Word	Nonword	Word	Nonword
Subject	Counter- balancing condition				
1	1	568.14	530.37	8	3
2	2	664.73	727.17	5	3
3	1	733.95	684.66	10	4
4	2	531.08	577.45	18	15
5	1	521.33	530.94	6	4
6	2	516.45	546.16	5	3
7	1	494.54	571.51	3	12
8	2	547.77	556.05	8	14
9	1	676.83	644.64	1	2
10	2	507.52	598.35	6	6
11	1	515.93	548.62	5	4
12	2	566.58	567.81	7	4
13	1	679.27	705.67	9	2
14	2	520.23	543.51	4	1
15	1	585.87	574.75	2	1
16	2	586.69	567.89	1	0
17	1	750.16	775.49	2	0
18	2	512.32	541.63	6	15
19	1	650.13	613.08	4	4
20	2	593.72	644.69	2	3
21	1	681.32	740.74	1	1
22	2	580.38	559.94	7	1
23	1	890.55	947.54	2	3
24	2	601.22	676.27	3	4
	mean	603.20	623.95	5.21	4.54

Experiment 2c: Word/nonword classification task

		Mean read	ction time (ms)		Number of e	errors		
Condition		Big	Big	Small Word	Small	Big Word	Big	Small	Small
Subject	Counter-	VVOIG	Noriword		Nortword		Nonworu		Nonwora
0	balancing								
	condition								
1	1	775.19	1185.00	1046.26	783.63	1	2	1	0
2	2	634.47	744.67	546.07	584.52	1	2	1	2
3	3	701.44	695.63	665.29	583.59	0	2	1	2
4	4	925.77	1318.50	1177.23	726.22	6	2	1	4
5	5	927.74	1021.21	877.58	1081.29	8	3	7	3
6	6	589.10	642.60	575.66	586.45	1	1	2	0
7	7	598.28	1208.13	1046.13	575.48	0	0	1	1
8	8	832.00	763.57	644.27	773.70	3	4	2	1
9	1	658.32	865.56	723.89	641.63	1	5	4	4
10	2	550.96	729.30	608.40	644.23	4	2	2	2
11	3	592.10	665.69	563.03	648.78	1	0	1	0
12	4	832.64	558.00	565.52	845.97	4	0	7	1
13	5	623.29	615.63	722.53	676.41	2	0	0	2
14	6	771.03	1100.83	982.97	952.81	1	2	2	1
15	7	784.07	1264.44	1078.24	1138.16	2	7	3	7
16	8	773.77	896.77	814.34	863.13	1	0	2	1
17	1	541.25	615.27	609.62	572.66	4	2	3	2
18	2	825.48	595.39	671.90	715.00	1	1	3	1
19	3	561.68	610.74	646.17	670.61	1	0	3	1
20	4	1197.03	809.03	744.15	1339.78	3	1	5	9
21	5	760.23	993.00	980.00	859.07	0	0	6	1
22	6	530.55	651.03	560.93	613.65	1	2	3	1
23	7	1132.75	1358.92	1300.52	1159.46	4	6	3	5
24	8	728.64	873.03	732.48	679.62	4	1	2	3
25	1	736.34	879.97	1037.81	663.53	2	2	1	2
26	2	677.96	991.91	770.93	703.50	4	0	5	5
27	3	661.29	913.97	811.90	669.26	1	2	1	1
28	4	456.87	568.97	576.37	437.96	1	1	5	4
29	5	753.46	615.61	595.34	685.96	4	4	2	6
30	6	721.65	929.71	789.22	724.87	0	1	0	1
31	7	718.90	786.16	792.47	624.73	2	1	2	2
32	8	629.94	1090.93	1227.54	703.16	1	2	4	0
	mean	725.13	861.22	796.40	747.78	2.16	1.81	2.66	2.34

Experiment 3: Elephant/cow IAT

Experiment 4a:

Flower/insect IAT

		Mean reaction	on time (ms)			Number of errors			
Condition		Flower Pleas	Flower Neutral	Insect Pleas	Insect Neutral	Flower	Flower	Insect Pleas	Insect Neutral
Subject	Counter- balancing condition	11000					Hourd	1 1003	Houtar
1	1	583.42	830.20	974.72	566.97	0	6	7	1
2	2	541.03	1010.07	1054.29	575.73	2	1	1	5
3	3	634.90	763.46	770.62	606.91	1	3	2	0
4	4	629.23	956.28	1128.84	712.72	2	5	7	2
5	1	648.13	1123.68	1149.72	595.46	2	6	0	4
6	2	744.17	917.00	1148.80	736.70	2	5	7	2
7	3	625.92	930.59	844.75	613.00	5	4	4	2
8	4	729.97	988.81	814.30	786.81	2	1	1	1
	mean	642.10	940.01	985.75	649.29	2.00	3.88	3.63	2.13

Experiment 4a: - continued

Flower/insect classification task

		Mean reaction t	ime (ms)	Number of errors		
Condition		Flower Insect		Flower	Insect	
Subject	Counter- balancing condition					
1	1	598.76	543.03	5	1	
2	2	663.80	617.73	10	2	
3	1	537.49	513.46	3	3	
4	2	591.63	602.40	10	13	
5	1	755.66	650.08	1	13	
6	2	669.86	717.14	7	5	
7	1	618.98	573.23	5	3	
8	2	595.24	646.49	4	1	
9	1	728.91	702.34	1	4	
10	2	647.82	690.46	2	4	
11	1	544.91	601.98	3	6	
12	2	550.95	522.97	5	4	
13	1	813.42	811.57	4	5	
14	2	736.81	657.05	4	5	
15	1	587.41	628.65	3	9	
16	2	576.04	633.88	3	3	
	mean	638.61	632.03	4.38	5.06	

Experiment 4b: Flower/insect split IAT

Pleasant IAT

		Mean reaction	on time (ms)			Number of errors			
Condition		Flower Pleas	Flower Neutral	Insect Pleas	Insect Neutral	Flower Pleas	Flower Neutral	Insect Pleas	insect Neutral
Subject	Counter- balancing condition								
1	1	860.60	937.71	995.63	771.45	1	1	2	2
2	2	1174.85	986.00	841.90	1043.22	5	2	3	5
3	3	553.16	709.23	742.90	618.97	7	6	3	3
4	4	865.65	675.80	644.71	829.71	0	1	0	1
5	1	634.32	778.81	771.80	579.38	6	4	6	5
6	2	1416.57	1333.72	1702.67	1295.74	9	3	1	5
7	3	697.75	789.30	836.57	713.66	0	2	2	3
8	4	1078.32	1375.00	1192.81	1280.84	1	1	1	1
9	1	654.74	910.04	924.44	634.23	0	7	0	0
10	2	610.65	829.71	921.19	575.41	1	4	1	5
11	3	916.39	1324.34	1119.24	778.07	0	0	3	2
12	4	1165.63	1366.69	1070.70	971.69	0	0	2	0
13	1	734.52	803.79	1111.81	792.77	1	2	1	1
14	2	603.45	752.00	876.81	614.47	0	3	4	1
15	3	607.73	997.89	911.38	708.97	2	5	3	3
16	4	551.60	775.97	626.97	647.72	1	1	1	0
17	1	687.77	906.43	768.30	813.32	1	3	5	1
18	2	934.63	1398.64	1048.10	1056.27	5	4	1	6
19	3	603.50	729.32	631.80	708.89	0	4	1	4
20	4	729.94	836.55	846.89	746.14	1	1	4	3
21	1	819.16	1020.90	986.50	1017.08	7	1	4	7
22	2	713.70	870.38	853.96	640.10	1	3	5	1
23	3	1456.07	1470.15	910.48	1143.71	5	6	2	4
24	4	561.35	845.13	724.55	688.72	1	2	1	3
	mean	818.00	975.98	919.25	819.60	2.29	2.75	2.33	2.75

Experiment 4b: Flower/insect split IAT - continued

Pleasant IAT attribute classification task

		Mean reaction t	ime (ms)	Number of erro	ors
Condition		Pleasant	Neutral	Pleasant	Neutral
Subject	Counter- balancing condition				
1	1	965.75	2105.455	0	1
2	2	838	1102.667	1	0
3	1	593.4167	783.0833	0	0
4	2	726.8333	846.7	0	2
5	1	865	838.1	1	2
6	2	994.9	1122.818	2	1
7	1	578.4167	698.8182	0	1
8	2	1124.625	1590.727	4	1
9	1	674.7778	709.0313	4	0
10	2	685.3333	710.5833	0	0
11	1	844.5455	1095.75	1	0
12	2	818	1298.364	1	1
13	1	930.2	1198.333	2	0
14	2	767.5455	757.6364	1	1
15	1	730.9	903.5833	2	0
16	2	847.8182	1035	1	2
17	1	699	1085.25	0	0
18	2	773.9167	993.25	0	0
19	1	926.7273	1049.909	1	1
20	2	683.3636	750.8182	1	1
21	1	683.3636	750.8182	1	1
22	2	655.5	962	0	0
23	1	835	1492.182	1	1
24	2	684.8	963.6667	2	3
	mean	788.6555	1035.189	1.083333	0.791667

	Mean reaction time (ms)					Number of errors			
Condition		Flower	Flower	Insect	Insect	Flower	Flower	Insect	Insect
		Neutral	Unpleas	Neutral	Unpleas	Neutral	Unpleas	Neutral	Unpleas
Subject	Counter- balancing condition								
1	1	748.32	740.10	747.79	856.14	1	2	2	3
2	2	678.19	809.40	910.83	599.20	6	7	8	7
3	3	717.59	1005.69	804.72	703.59	5	6	7	3
4	4	634.53	580.84	652.10	686.38	1	6	2	3
5	1	679.87	973.86	847.40	649.61	0	3	1	0
6	2	941,16	867.33	960.26	886.00	0	4	0	2
7	3	641.07	678.17	792.87	579.74	3	3	1	5
8	4	619.64	980.39	1056.35	653.85	6	4	1	5
9	1	870.56	944.90	990.11	1022.52	6	2	4	5
10	2	704.72	851.18	826.34	684.52	3	4	3	3
11	3	610.63	743.03	683.33	602.14	8	2	4	3
12	4	565.83	670.67	606.07	582.12	3	2	3	6
13	1	961.87	820.36	929.00	915.52	2	4	0	4
14	2	949.07	951.48	1134.29	881.84	1	1	1	0
15	3	816.62	642.53	681.61	693.03	6	0	1	2
16	4	812.34	931.45	1021.11	793.79	0	1	4	3
17	1	855.22	1102.08	1170.96	716.20	5	7	3	1
18	2	711.41	929.74	1138.90	785.47	3	4	3	2
19	3	702.63	862.46	920.18	591.22	1	6	7	5
20	4	772.66	1011.41	911.90	654.03	2	3	2	2
21	1	645.17	996.47	1032.21	792.36	6	6	3	6
22	2	772.66	896.79	934.87	675.37	2	2	2	2
23	3	700.37	777.77	826.70	648.70	4	1	5	2
24	4	752.38	774.32	847.35	724.53	0	4	1	1
	mean	744.35	855.94	892.80	724.08	3.08	3.50	2.83	3.13

Experiment 4b: Flower/insect split IAT - continued

onpreasant IAT attribute classification tas	Unpleasant IAT	attribute	classification	task
---	----------------	-----------	----------------	------

		Mean reaction t	ime (ms)	Number of erro	Number of errors		
Condition		Neutral	Unpleasant	Neutral	Unpleasant		
Subject	Counter- balancing condition		<u></u>		· · · · · · · · · · · · · · · · · · ·		
1	1	923.375	947.3333	4	0		
2	2	834.2857	769.6	5	2		
3	1	838.7	679	2	2		
4	2	1130.917	1008.182	0	1		
5	. 1	672.0833	585.8182	1	0		
6	2	1389.727	1410.4	1	2		
7	1	629.7273	547.6667	0	0		
8	2	843.5556	879.625	3	4		
9	1	1121.364	711.0909	1	1		
10	2	770.5833	760	0	1		
11	1	700.6667	589.8182	0	1		
12	2	911.1818	714.1	1	2		
13	1	1285.833	924.8182	0	1		
14	2	856	1055	3	1		
15	1	855	531.6	0	2		
16	2	1581.778	747	3	0		
17	1	694.4444	901.5455	3	1		
18	2	1141.1	1053.909	2	1		
19	1	651.3636	553.9	1	2		
20	2	904.0909	723.4167	1	0		
21	1	849.6	812.8889	2	3		
22	2	979.7	883.2727	2	1		
23	1	952.7	716.2727	2	1		
24	2	1005.833	831.3636	0	1		
	mean	938.4837	805.7342	1.541667	1.25		

Experiment 5a:

Go/Nogo IAT (Flight/Teeth)

		Mean reaction	on time (ms)			Number of errors			
Condition		Go Pleas	Go Unpleas	Nogo Pleas	Nogo Unpleas	Go Pleas	Go Unpleas	Nogo Pleas	Nogo Unpleas
Subject	Counter- balancing condition								
1	1	780.47	1120.61	1035.04	670.21	0	4	9	4
2	2	1531.20	1527.28	1326.21	1308.59	2	3	3	5
3	3	965.10	1356.93	1183.13	973.40	2	2	2	2
4	4	1226.58	1178.50	1410.45	1492.44	0	0	0	1
5	1	936.90	1316.20	1346.66	1142.66	1	2	0	0
6	2	658.17	924.54	881.00	797.04	1	3	2	8
7	3	868.97	1152.73	957.67	875.23	0	2	2	1
8	4	1567.45	1499.87	1475.86	1732.46	2	4	3	6
9	1	1304.72	1264.84	1312.75	1196.61	0	0	0	1
10	2	1123.91	1196.61	1304.72	1329.59	0	1	0	0
11	3	688.96	694.17	731.77	732.12	4	2	2	5
12	4	1493.32	930.71	1386.77	1529.24	4	0	2	3
13	1	1360.13	1562.16	2052.11	1163.48	1	2	4	1
14	2	575.74	712.28	1052.15	585.46	1	2	4	3
15	3	1149.70	1676.06	1745.03	1300.81	1	1	0	0
16	4	1115.10	1852.48	1723.48	1527.90	1	2	4	3
	mean	1084.15	1247.87	1307.80	1147.33	1.25	1.88	2.31	2.69

Experiment 5a: - continued

Go/Nogo classification task (Flight/Teeth)

		Mean reaction t	ime (ms)	Number of errors		
Condition		Go	Nogo	Go	Nogo	
Subject	Counter- balancing condition					
1	1	890.22	863.93	4	4	
2	2	501.02	519.43	6	2	
3	3	720.59	719.00	5	9	
4	4	516.79	586.90	5	8	
5	1	507.53	528.04	2	1	
6	2	869.73	795.71	5	1	
7	3	635.74	596.85	2	4	
8	4	603.61	618.21	1	3	
9	1	499.13	519.07	3	3	
10	2	650.51	728.44	4	0	
11	3	636.55	728.73	4	6	
12	4	615.58	690.26	1	6	
13	1	548.94	572.16	8	7	
14	2	644.67	699.02	4	6	
15	3	654.64	666.48	1	2	
16	4	649.36	758.22	1	6	
	mean	634.04	661.90	3.50	4.25	

Experiment 5b: Go/Nogo split IAT (Flight/Teeth)

Pleasant IAT

		Mean reaction	tion time (ms) Number of errors						
Condition		Go Pleas	Go Neutral	Nogo Pleas	Nogo Neutral	Go Pleas	Go Neutral	Nogo Pleas	Nogo Neutral
Subject	Counter- balancing condition								
1	1	658.93	1213.60	922.53	636.00	3	7	2	2
2	2	573.13	838.03	722.10	611.90	0	3	0	1
3	3	891.50	1109.09	1275.88	1196.54	1	9	8	4
4	4	737,90	882.00	883.57	957.85	0	5	2	5
5	1	587.59	923.17	893.38	741.04	2	3	8	5
6	2	617.26	1291.43	983.45	720.48	1	2	0	0
7	3	854.34	964.41	1107.97	1092.45	0	0	2	3
8	4	724.65	1396.73	1055.14	1009.97	1	2	4	1
9	1	616.58	872.39	709.23	677.93	1	4	1	5
10	2	806.23	1512.47	1064.00	867.28	1	2	1	3
11	3	836.39	810.23	863.31	1015.93	3	0	0	3
12	4	607.10	660.20	707.76	581.14	1	1	1	2
13	1	1565.28	1563.65	1151.37	1437.21	3	0	2	2
14	2	737.00	1198.40	852.07	726.13	2	1	4	0
15	3	1171.97	1200.12	888.00	1297.15	2	6	1	5
16	4	847.04	1303.79	1496.65	991.61	4	3	6	2
	mean	802.05	1108.73	973.52	910.04	1.56	3.00	2.63	2.69

Experiment 5b: Go/Nogo split IAT (Flight/Teeth) - continued

Unpleasant IAT

		Mean reaction	n time (ms)			Number of errors			
Condition		Go Neutral	Go Unpleas	Nogo Neutral	Nogo Unpleas	Go Neutral	Go Unpleas	Nogo Neutral	Nogo Unpleas
Subject	Counter- balancing condition								
1	1	961.09	591.37	586.50	870.76	9	2	0	8
2	2	1063.86	1105.90	1023.47	912.83	4	2	0	3
3	3	759.37	568.00	628.44	766.12	5	2	5	7
4	4	928.45	596.50	741.21	843.64	3	0	4	3
5	1	859.57	978.34	1205.73	741.81	4	3	2	1
6	2	790.81	686.26	781.17	821.35	5	0	2	6
7	3	1240.97	1066.97	1144.17	1486.10	3	2	2	3
8	4	1658.13	1009.48	1307.33	1606.69	1	1	2	3
9	1	1183.41	706.23	760.20	1032.79	0	2	2	4
10	2	800.66	687.50	671.04	692.10	2	3	4	2
11	3	952.88	830.71	827.90	854.55	5	1	1	3
12	4	842.23	1128.58	955.87	931.19	0	0	1	0
13	1	1764.48	1049.71	1081.28	1321.59	7	1	3	3
14	2	1124.81	691.97	620.33	1083.92	5	0	1	6
15	3	962.00	898.24	1118.59	772.18	1	6	3	4
16	4	1142.81	853.78	1030.24	1088.68	1	4	3	4
	mean	1064.72	840.60	905.22	989.14	3.44	1.81	2.19	3.75

Experiment 6a:

Insect/nonword IAT

		Mean reaction	on time (ms)			Number of errors			
Condition		Insect Pleas	Insect Unpleas	Nonword Pleas	Nonword Unpleas	Insect Pleas	Insect Unpleas	Nonword Pleast	Nonword Unpleas
Subject	Counter- balancing condition								
1	1	653.40	920.03	1009.16	673.50	1	1	1	1
2	2	818.45	842.58	843.03	788.72	0	1	1	3
3	3	695.21	884.13	781.75	670.17	2	8	3	2
4	4	1310.65	1351.10	1966.13	964.93	0	0	0	1
5	1	689.63	914.71	816.00	733.69	2	3	4	0
6	2	1167.56	1389.56	1553.93	1106.07	0	0	2	5
7	3	664.31	737.59	753.48	661.00	3	3	3	1
8	4	572.13	909.29	918.83	655.71	0	0	2	0
	mean	821.42	993.62	1080.29	781.72	1.00	2.00	2.00	1.63

Insect/nonword classification task

		Mean reaction t	ime (ms)	Number of errors		
Condition		Insect	Nonword	Insect	Nonword	
Subject	Counter- balancing condition					
1	1	915.23	921.59	2	2	
2	2	555.90	571.03	4	1	
3	1	646.33	735.87	1	0	
4	2	584.64	632.73	11	10	
5	1	817.82	1047.37	3	3	
6	2	570.25	605.51	3	0	
7	1	576.45	618.18	5	15	
8	2	594.23	612.25	5	2	
9	1	685.20	698.56	3	3	
10	2	550.17	613.22	0	1	
	mean	649.62	705.63	3.70	3.70	

Experiment 6b: Insect/nonword split IAT

Pleasant IAT

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Insect Pleas	Insect Neutral	Nonword Pleas	Nonword Neutral	Insect Pleas	Insect Neutral	Nonword Pleas	Nonword Neutral
Subject	Counter- balancing condition								
1	1	626.57	980.55	968.72	704.00	2	1	3	1
2	2	739.59	1018.67	929.13	755.66	0	1	1	0
3	3	625.04	1002.82	919.14	735.67	4	3	9	2
4	4	719.93	867.72	736.43	863.68	3	3	3	4
5	1	1017.83	1133.58	958.06	1167.37	2	1	1	2
6	2	1062.93	1318.93	1420.13	1260.55	1	0	2	1
7	3	636.34	1287.97	1343.57	695.97	0	0	2	2
8	4	646.06	1169.78	1226.18	655.65	1	5	4	1
9	1	718.28	898.56	1026.70	748.90	0	0	2	2
10	2	746.13	911.56	983.06	911.35	1	0	1	0
11	3	711.77	1064.63	936.50	718.66	1	1	0	0
12	4	756.59	962.44	988.74	933.22	2	6	9	8
13		788.67	1361.53	1386.54	853.29	1	2	4	0
14	2	798.19	1087.72	1113.06	778.44	0	0	0	0
15	3	580.93	915.35	984.36	733.54	4	6	6	6
16	4	705.35	1141.07	1228.63	774.66	1	4	1	2
	mean	742.51	1070.18	1071.81	830.66	1.44	2.06	3.00	1.94

		Mean reaction	on time (ms)			Number of e	errors		
Condition	1	Insect Neutral	Insect Unpleas	Nonword Neutral	Nonword Unpleas	Insect Neutral	Insect Unpleas	Nonword Neutral	Nonword Unpleas
Subject	Counter- balancing condition								•
1	1	1087.04	754.63	740.72	901.17	6	1	0	9
2	2	930.76	680.83	707.80	841.11	2	2	6	3
3	3	890.63	559.54	642.69	787.41	5	4	5	9
4	4	938.23	539.24	637.43	638.32	1	2	4	6
5	1	1182.62	788.60	918.90	1453.19	3	2	2	0
6	2	935.00	528.06	595.14	965.19	1	1	4	6
7	3	827.17	722.62	688.72	1037.23	2	2	2	6
8	4	661.50	671.61	719.65	701.38	0	1	0	2
9	1	1391.88	665.87	786.97	1269.21	7	2	1	3
10	2	643.47	533.88	604.50	565.63	1	0	4	2
11	3	770.65	466.03	547.52	695.65	1	1	1	6
12	4	949.90	563.22	665.46	916.83	1	4	6	3
13	1	847.81	612.44	667.85	981.13	1	4	6	8
14	2	1126.10	622.50	633.03	1039.35	2	1	0	6
15	3	800.65	662.73	662.39	1094.28	1	1	3	3
16	4	785.91	715.59	697.83	827.84	0	0	1	1
	mean	923.08	630.46	682.29	919.68	2.13	1.75	2.81	4.56

Experiment 7a:

Familiar/unfamiliar paintings IAT

		Mean reacti	on time (ms)			Number of e	errors		
Condition		Familiar Pleas	Familiar Unpleas	Unfamiliar Pleas	Unfamiliar Unpleas	Familiar Pleas	Familiar Unpleas	Unfamiliar Pleas	Unfamiliar Unpleas
Subject	Counter- balancing condition								
1	1	1281.48	2094.53	1752.90	1140.97	2	2	2	0
2	2	677.33	1420.67	1460.64	701.38	4	4	4	3
3	3	763.46	796.65	870.57	764.80	3	6	4	2
4	4	812.56	746.55	796.30	846.58	0	0	1	0
5	1	621.23	711.87	692.83	644.66	1	1	2	0
6	2	1023.68	777.84	711.42	852.60	3	1	1	0
7	3	505.58	684.00	688.71	535.00	1	5	4	1
8	4	753.10	821.83	737.32	765.09	3	3	7	0
9	1	529.06	596.29	618.42	523.56	0	1	0	0
10	2	603.55	860.23	798.33	673.32	3	2	2	0
11	3	642.78	835.29	816.63	662.22	0	1	0	0
12	4	735.34	864.22	922.28	760.91	0	0	0	0
13	1	811.22	959.63	1000.07	843.42	0	0	2	0
14	2	658.97	906.16	786.55	655.31	3	1	3	2
15	3	819.40	801.24	977.77	898.90	1	2	1	0
16	4	789.84	728.13	703.50	650.69	1	1	3	0
	mean	751.79	912.82	895.89	744.96	1.56	1.88	2.25	0.50

Familiar/unfamiliar paintings classification task

		Mean read	ction time (ms)	Number of e	errors
Condition		Familiar	Unfamiliar	Familiar	Unfamiliar
Subject	Counter- balancing condition				
1	1	560.06	548.10	0	1
2	2	487.14	534.00	2	1
3	1	493.07	553.87	1	1
4	2	602.37	643.00	1	0
5	1	601.26	633.97	0	1
6	2	529.16	564.71	1	0
7	1	606.45	602.23	2	0
8	2	574.90	567.79	2	1
9	1	517.47	604.73	0	1
10	2	604.03	574.10	0	0
11	1	571.90	521.47	2	0
12	2	549.06	620.97	0	0
13	1	526.10	521.13	2	0
14	2	501.44	548.59	0	0
15	1	557.96	562.21	4	2
16	2	651.57	719.86	2	4
	mean	558.37	582.54	1.19	0.75

Experiment 7b: Familiar/unfamiliar paintings split IAT

Pleasant IAT

		Mean reaction	n time (ms)			Number of e	errors		
Condition		Familiar Pleas	Familiar Neutral	Unfamiliar Pleas	Unfamiliar Neutral	Familiar Pleas	Familiar Neutral	Unfamiliar Pleas	Unfamiliar Neutral
Subject	Counter- balancing condition								
1	1	678.34	1224.35	1453.37	733.90	0	1	2	1
2	2	655.59	862.72	1005.74	760.34	0	2	4	0
3	3	675.81	1613.56	1283.07	1167.58	0	4	4	1
4	4	1053.03	1811.68	1585.65	880.32	2	7	9	9
5	1	789.43	1633.58	1266.13	796.06	2	1	1	0
6	2	503.26	803.96	828.33	533.00	1	5	1	6
7	3	568.81	748.97	756.78	621.72	0	3	0	0
8	4	556.86	778.77	699.89	677.14	2	2	4	2
9	1	1059.59	1989.92	1431.61	712.70	0	6	8	4
10	2	1109.47	1451.25	1077.79	1575.03	1	4	3	1
11	3	552.84	875.87	780.93	713.32	1	1	2	0
12	4	1501.09	1156.11	1350.17	1557.28	0	3	2	0
13	1	543.00	1165.96	1254.93	687.37	0	5	3	2
14	2	762.26	935.93	1007.53	695.66	1	2	2	3
15	3	550.42	696.82	778.56	796.34	0	3	0	0
16	4	629.45	949.32	1072.30	637.22	1	1	2	0
	mean	761.83	1168.67	1102.05	846.56	0.69	3.13	2.94	1.81

ł		Mean reaction	on time (ms)			Number of e	errors		
Condition		Familiar Pleas	Familiar Unpleas	Unfamiliar Pleas	Unfamiliar Unpleas	Familiar Pleas	Familiar Unpleas	Unfamiliar Pleas	Unfamiliar Unpleas
Subject	Counter- balancing condition								
1	1	951.91	1230.53	1059.29	983.35	0	0	0	1
2	2	858.97	683.34	825.84	764.27	1	0	1	1
3	3	903.65	768.11	817.28	824.21	1	3	2	3
4	4	1215.24	665.32	813.37	1240.26	7	4	5	5
5	1	765.61	1239.32	1073.50	730.63	0	1	1	0
6	2	687.24	652.25	628.85	746.81	2	4	4	1
7	3	821.54	1189.10	1178.88	785.23	3	2	0	0
8	4	959.40	682.93	713.31	1123.10	2	1	0	1
9	1	726.69	816.62	1058.37	873.97	3	3	2	2
10	2	860.32	613.63	764.90	639.19	1	1	0	1
11	3	831.30	642.19	915.16	776.28	2	0	1	3
12	4	713.64	510.29	559.84	686.93	3	1	0	3
13	1	885.00	1120.30	1490.80	940.39	1	1	2	1
14	2	991.33	633.48	598.70	903.19	5	2	1	6
15	3	551.47	571.54	634.00	520.63	1	4	5	2
16	4	665.54	684.00	665.07	601.37	4	2	3	2
	mean	836.80	793.94	862.32	821.24	2.25	1.81	1.69	2.00

Experiment 8a:

Even/odd number IAT

		Mean reaction	on time (ms)			Number of errors			
Condition		Even	Even	Odd	Odd	Even	Even	Odd	Odd
		Pleas	Unpleas	Pleas	Unpleas	Pleas	Unpleas	Pleas	Unpleas
Subject	Counter-								
	balancing								
	condition								
1	1	641.00	991.45	1209.97	677.63	1	0	3	0
2	2	1146.68	695.00	688.39	1147.23	7	5	1	5
3	3	814.32	1045.93	1107.78	731.97	3	5	4	1
4	4	787.83	1653.15	1999.68	813.33	2	5	4	1
5	1	793.33	2114.17	2292.28	852.22	2	2	3	0
6	2	865.69	1373.00	1379.66	734.87	2	2	0	1
7	3	624.75	903.77	1402.04	560.92	3	5	8	6
8	4	639.23	711.31	803.69	596.91	1	2	2	0
9	1	705.26	1438.65	1374.30	902.89	0	6	2	5
10	2	663.10	932.79	1112.74	736.39	2	2	1	0
11	3	715.80	1046.17	1192.37	585.47	2	2	4	2
12	4	546.80	679.03	655.48	518.10	2	2	6	0
	mean	745.32	1132.04	1268.20	738.16	2.25	3.17	3.17	1.75

Even/odd number classification task

		Mean reaction t	ime (ms)	Number of errors		
Condition		Even	Odd	Even	Odd	
Subject	Counter- balancing condition					
1	1	1054.69	991.39	0	0	
2	2	604.13	637.90	3	3	
3	1	693.79	695.09	8	5	
4	2	605.53	604.36	4	5	
5	1	448.46	462.98	7	7	
6	2	605.53	604.36	4	5	
7	1	644.36	662.80	0	2	
8	2	631.38	698.20	2	5	
9	1	565.90	590.69	9	7	
10	2	712.09	666.64	4	0	
11	1	654.03	637.77	1	7	
12	2	672.93	609.91	4	4	
13	1	575.26	586.27	9	9	
14	2	701.22	914.09	0	2	
15	1	544.57	554.61	2	5	
16	2	563.33	670.25	4	5	
17	1	537.67	587.39	7	14	
18	2	631.67	663.82	2	5	
	mean	635.92	657.70	3.89	5.00	

Experiment 8b: Even/odd number split IAT

Pleasant IAT

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Even Pleas	Even Neutral	Odd Pleas	Odd Neutral	Even Pleas	Even Neutral	Odd Pleas	Odd Neutral
Subject	Counter- balancing condition								
1	1	754.90	860.00	1034.91	820.80	1	2	0	1
2	2	643.00	745.47	694.34	733.38	2	0	2	3
3	3	755.13	1837.59	1627.61	974.27	0	5	1	2
4	4	825.30	905.48	1090.63	1084.06	1	0	2	1
5	1	835.97	1193.27	1264.84	718.17	0	1	0	1
6	2	572.03	1102.17	1232.71	605.46	1	2	4	3
7	3	764.14	888.61	1149.96	702.70	2	0	5	1
8	4	833.48	1005.70	817.45	854.83	4	2	3	3
9	1	869.26	971.34	1074.66	842.23	1	2	3	0
10	2	1072.28	711.71	774.69	1000.70	0	1	0	1
11	3	643.67	1002.38	821.07	666.93	2	3	4	3
12	4	1007.67	1428.48	1558.55	1114.53	1	0	1	0
13	1	570.72	816.13	814.72	575.61	3	1	0	1
14	2	649.38	1203.47	932.63	830.13	0	0	1	0
15	3	653.81	1054.43	1021.30	836.63	0	3	1	2
16	4	948.80	844.40	852.39	1054.38	7	6	3	3
	mean	774.97	1035.67	1047.65	838.43	1.56	1.75	1.88	1.56

		Mean reaction	on time (ms)			Number of e	errors		
Condition	1	Even Neutral	Even Unpleas	Odd Neutral	Odd Unpleas	Even Neutral	Even Unpleas	Odd Neutral	Odd Unpleas
Subject	Counter- balancing condition								
1	1	655.40	794.12	1143.41	636.39	1	5	0	0
2	2	600.56	517.50	530.13	688.82	4	3	1	3
3	3	653.63	764.79	716.38	582.79	1	3	2	3
4	4	731.93	869.52	914.41	808.43	2	4	2	1
5	1	659.10	838.82	830.33	806.96	2	3	8	4
6	2	694.84	755.77	878.69	669.90	0	2	2	2
7	3	841.07	734.92	947.32	805.61	5	5	3	1
8	4	1509.32	1504.58	1547.16	1185.10	1	1	1	1
9	1	649.37	721.31	804.17	565.29	2	2	2	0
10	2	767.10	763.28	796.23	787.45	1	0	1	1
11	3	1300.32	1188.69	900.73	1149.76	3	3	1	3
12	4	710.55	803.16	904.45	728.55	0	0	3	1
13	1	836.16	868.62	1131.90	748.90	1	2	1	1
14	2	705.63	1022.04	1218.03	561.25	0	3	3	4
15	3	936.17	890.45	924.22	875.03	2	1	0	1
16	4	1219.37	1596.63	1143.29	886.60	5	5	1	2
	mean	841.91	914.64	958.18	780.43	1.88	2.63	1.94	1.75

Experiment 9a:

Race IAT with faces

		Mean reaction	on time (ms)			Number of e	errors		
Condition		White Pleas	White Unpleas	Black Pleas	Black Unpleas	White Pleas	White Unpleas	Black Pleas	Black Unpleas
Subject	Counter- balancing condition								
1	1	575.80	790.45	784.41	517.74	1	3	5	1
2	2	664.90	790.93	773.75	597.10	1	1	3	1
3	3	602.53	783.31	657.48	549.31	1	2	3	0
4	4	577.07	650.52	718.03	557.40	3	3	2	2
5	1	806.14	1097.60	1537.97	1234.59	2	2	1	0
6	2	738.42	1241.34	1095.13	609.09	0	0	1	0
7	3	947.71	999.32	922.55	834.25	0	1	1	0
8	4	562.71	645.91	523.50	588.00	4	0	2	4
	mean	684.41	874.92	876.60	685.94	1.50	1.50	2.25	1.00

Race classification task with faces

		Mean reaction t	ime (ms)	Number of erro	ors
Condition		White	Black	White	Black
Subject	Counter- balancing condition				
1	1	748.59	761.79	2	1
2	2	649.20	685.52	1	2
3	1	595.36	591.30	2	1
4	2	435.72	397.74	4	6
5	1	541.96	496.63	1	1
6	2	571.31	585.80	0	1
7	1	564.80	487.94	3	4
8	2	555.65	501.40	9	6
9	1	485.60	494.72	0	2
10	2	573.62	569.55	2	5
11	1	423.58	388.33	7	9
12	2	547.93	468.79	0	1
13	1	546.14	494.27	2	4
14	2	625.41	597.58	3	1
15	1	682.55	607.81	2	0
16	2	508.23	491.44	2	2
17	1	553.76	530.54	0	0
18	2	555.64	482.04	4	4
	mean	564.73	535.71	2.44	2.78

Experiment 9b:

Race IAT with names

		Mean reaction	on time (ms)			Number of e	errors		
Condition		White	White	Black	Black	White	White	Black	Black
		Pleas	Unpleas	Pleas	Unpleas	Pleas	Unpleas	Pleas	Unpleas
Subject	Counter-								
	balancing								
	condition								
1	1	868.64	932.33	723.58	818.50	4	2	1	5
2	2	921.16	1403.00	1288.00	863.00	7	4	8	3
3	3	753.29	1186.90	1059.33	921.10	0	2	2	1
4	4	796.00	1299.93	1647.76	1415.42	1	3	3	1
5	1	870.23	1621.29	1200.12	668.57	1	4	5	3
6	2	636.76	748.54	763.96	655.96	6	4	6	9
7	3	1066.44	874.28	807.81	939.22	0	0	1	0
8	4	794.50	1589.54	1229.72	821.14	1	6	2	4
9	1	1608.65	1985.96	2140.81	1486.72	1	5	1	2
10	2	654.40	821.60	836.97	714.14	1	7	3	3
11	3	773.59	1243.20	1107.89	735.31	0	1	5	3
12	4	953.66	1360.28	1113.47	1187.14	3	0	0	3
	mean	891.44	1255.57	1159.95	935.52	2.08	3.17	3.08	3.08

Race classification task with names

		Mean reaction t	ime (ms)	Number of erro	ors
Condition		White	Black	White	Black
Subject	Counter- balancing condition				
1	1	561.36	543.19	3	9
2	2	513.98	530.48	9	7
3	1	593.73	678.16	4	9
4	2	937.48	869.01	1	1
5	1	702.60	710.98	6	6
6	2	726.54	699.76	4	0
7	1	795.63	794.73	1	0
8	2	827.62	728.74	5	5
9	1	626.93	610.76	2	3
10	2	677.67	570.61	1	4
11	1	642.68	617.99	9	1
12	2	643.60	709.82	1	3
13	1	676.88	632.26	4	4
14	2	654.37	597.26	3	4
15	1	857.69	790.28	2	2
16	2	568.45	589.44	3	8
17	1	792.03	702.73	1	4
18	2	542.52	565.78	5	7
19	1	594.57	559.91	6	3
20	2	925.04	815.35	2	5
21	1	593.39	607.94	0	3
22	2	707.53	706.68	3	6
	mean	689.20	665.09	3.41	4.27

Experiment 9c: Race split IAT with faces

Pleasant IAT

		Mean reaction	on time (ms)			Number of e	errors		
Condition		White Pleas	White Neutral	Black Pleas	Black Neutral	White Pleas	White Neutral	Black Pleas	Black Neutral
Subject	Counter- balancing condition								
1	1	848.70	877.58	804.32	858.06	1	1	0	1
2	2	821.19	724.39	610.74	722.52	5	4	5	7
3	3	823.86	1167.10	934.73	925.21	2	1	1	3
4	4	642.72	597.27	554.00	608.03	3	2	1	2
5	1	722.87	790.54	925.24	701.53	2	4	3	2
6	2	848.52	975.23	778.07	652.54	3	1	5	3
7	3	650.29	776.07	708.10	773.60	1	1	0	1
8	4	802.56	610.19	578.37	556.65	5	0	2	0
9	1	868.59	986.81	1162.22	1030.81	0	1	0	0
10	2	1177.55	1046.55	773.53	1042.59	1	3	0	2
11	3	641.57	669.94	524.30	632.74	1	1	1	1
12	4	864.52	750.23	747.68	762.28	0	1	1	0
13	1	536.76	658.33	654.86	643.80	2	2	3	2
14	2	1058.93	965.75	1174.73	886.10	3	3	5	0
15	3	1123.69	1216.23	944.83	1000.50	0	1	3	1
16	4	624.39	689.94	529.23	690.50	0	0	1	0
	mean	816.04	843.88	775.31	780.47	1.81	1.63	1.94	1.56

		Mean reaction	on time (ms)			Number of e	errors		
Condition		White Neutral	White Unpleas	Black Neutral	Biack Unpleas	White Neutral	White Unpleas	Black Neutral	Black Unpleas
Subject	Counter- balancing condition								
1	1	875.69	798.00	780.79	767.91	0	0	2	0
2	2	599.10	798.00	780.79	568.90	3	0	2	1
3	3	911.38	937.71	902.31	890.65	0	1	0	0
4	4	506.61	587.70	617.07	514.13	1	1	2	2
5	1	1183.73	1372.93	1617.90	805.71	1	2	3	3
6	2	899.47	597.17	572.55	854.48	0	2	2	3
7	3	672.40	1356.07	1063.61	631.44	1	3	0	0
8	4	934.96	1159.40	976.65	733.43	4	1	5	1
9	1	617.94	1193.27	1366.90	599.61	0	2	1	0
10	2	567.86	620.17	720.90	513.38	4	2	1	2
11	3	667.13	909.72	836.00	508.72	1	3	4	3
12	4	894.85	805.63	717.10	728.24	5	1	2	7
13	1	643.97	1088.66	910.11	633.20	0	3	4	1
14	2	1102.44	807.10	795.45	837.00	0	1	1	1
15	3	831.22	1054.89	1272.88	851.53	0	3	6	2
16	4	544.14	633.04	710.24	480.00	4	3	3	1
	mean	778.31	919.97	915.08	682.40	1.50	1.75	2.38	1.69

Experiment 9d: Race split IAT with names

Pleasant IAT

		Mean reaction	on time (ms)			Number of e	errors		
Condition		White Pleas	White Neutral	Black Pleas	Black Neutral	White Pleas	White Neutral	Black Pleas	Black Neutral
Subject	Counter- balancing condition						<u></u>		
1	1	880.81	666.79	642.10	856.11	1	3	1	3
2	2	1324.90	1520.23	1936.29	822.32	3	1	4	3
3	3	1189.63	827.81	934.35	1217.20	2	1	1	2
4	4	947.76	1248.85	1043.59	1008.79	2	5	3	4
5	1	929.62	963.86	896.34	938.78	3	3	2	5
6	2	642.69	664.53	672.22	750.22	3	0	5	0
7	3	804.61	1033.20	1428.84	880.18	8	6	7	3
8	4	765.31	813.03	780.10	981.69	3	1	3	3
9	1	918.07	1226.37	737.32	969.79	2	2	3	4
10	2	1267.07	985.21	672.90	1029.10	2	3	1	2
11	3	816.13	988.32	854.75	715.50	0	0	3	0
12	4	1035.79	717.20	663.63	1045.23	3	1	1	2
13	1	703.87	596.97	582.39	736.21	2	2	3	2
14	2	847.23	1024.03	983.92	816.76	1	2	6	2
15	3	571.00	600.37	538.41	579.30	1	4	3	2
16	4	1035.79	1036.28	962.11	1045.23	3	3	4	2
	mean	917.52	932.07	895.58	899.52	2.44	2.31	3.13	2.44

		Mean reaction	on time (ms)			Number of e	errors		
Condition		White Neutral	White Unpleas	Black Neutral	Black Unpleas	White Neutral	White Unpleas	Black Neutral	Black Unpleas
Subject	Counter- balancing condition								
1	1	699.23	1281.00	1124.90	698.90	1	3	1	3
2	2	522.75	619.94	614.81	472.97	0	0	0	1
3	3	960.84	1891.18	1104.35	1177.38	0	4	1	3
4	4	782.72	969.57	1022.84	585.50	3	2	1	1
5	1	640.50	649.44	897.46	572.52	4	7	4	2
6	2	1202.30	1789.48	1464.04	756.25	2	3	8	0
7	3	923.90	1333.94	1452.97	960.34	1	1	2	1
8	4	1208.59	1143.43	1394.93	787.87	2	1	4	1
9	1	584.83	758.71	997.40	560.14	1	4	7	4
10	2	924.77	1314.48	1562.86	847.97	0	1	3	1
11	3	1665.31	1636.03	1438.20	930.61	0	0	2	3
12	4	1208.59	1148.03	1394.93	787.87	2	1	4	1
13	1	1313.03	2598.16	1982.09	727.78	2	7	8	4
14	2	624.40	902.24	961.19	732.03	1	3	1	2
15	3	786.07	1108.66	843.27	756.13	3	3	1	1
16	4	624.16	937.58	744.26	643.36	1	1	9	4
	mean	917.00	1255.12	1187.53	749.85	1.44	2.56	3.50	2.00

Experiment 10a:

Age IAT with faces

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Young Pleas	Young Unpleas	Old Pleas	Old Unpleas	Young Pleas	Young Unpleas	Old Pleas	Old Unpleas
Subject	Counter- balancing condition						• • • • • • • • • • • •	, <u>, , , , , , , , , , , , , , , , , , </u>	
1	1	564.90	652.31	669.70	561.17	1	2	2	3
2	2	602.13	815.40	744.25	607.23	1	2	0	1
3	3	748.39	698.88	749.59	663.39	1	0	0	1
4	4	518.87	560.47	523.03	494.81	1	0	0	0
5	1	586.09	807.61	726.26	541.23	0	1	4	1
6	2	855.28	1011.43	1019.69	808.68	0	1	0	1
7	3	656.35	1052.61	1148.72	753.63	0	4	3	0
8	4	573.19	900.10	755.90	734.89	4	2	2	4
	mean	638.15	812.35	792.14	645.63	1.00	1.50	1.38	1.38

Age classification task with faces

		Mean reacti	on time (ms)	Number	oferrors
Condition		Young	Old	Young	Old
Subject	Counterbalancing condition			· · · · ·	• • • • • • • • • • • • • • • • • • •
1	1	560.46	562.51	1	1
2	2	499.20	507.07	1	2
3	1	455.83	395.60	11	9
4	2	631.70	609.72	1	0
5	1	486.97	442.79	6	10
6	2	516.51	476.03	1	4
7	1	510.81	448.06	3	1
8	2	493.29	450.94	2	2
9	1	574.88	538.21	2	3
10	2	520.75	509.64	4	2
	mean	525.04	494.06	3.20	3.40

Experiment 10b:

Age IAT with names

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Young Pleas	Young Unpleas	Old Pleas	Old Unpleas	Young Pleas	Young Unpleas	Old Pleas	Old Unpleas
Subject	Counter- balancing condition								
1	1	639.52	698.86	753.52	603.10	2	2	3	1
2	2	1444.17	2290.77	1448.66	1134.35	2	5	2	0
3	3	792.04	832.97	812.33	814.28	8	2	8	3
4	4	857.75	1692.44	1426.00	810.23	3	7	7	2
5	1	648.83	869.92	1018.26	622.36	2	7	5	6
6	2	737.74	1065,17	1370.36	755.89	1	7	4	3
7	3	938.25	1536.82	1560.96	1070.84	3	3	4	0
8	4	778.71	1019.58	1052.00	746.42	0	1	1	0
	mean	854.63	1250.82	1180.26	819.69	2.63	4.25	4.25	1.88

Experiment 10b: - continued

Age classification task with names

		Mean reaction t	ime (ms)	Number of erro	ors
Condition		Young	Old	Young	Old
Subject	Counter- balancing condition				
1	1	660.90	652.06	0	1
2	2	454.58	487.88	11	8
3	1	687.15	643.62	4	2
4	2	1107.24	1007.43	4	10
5	1	765.44	746.14	6	1
6	2	666.92	700.53	6	9
7	1	569.32	597.34	2	1
8	2	599.58	756.08	5	9
9	1	527.01	526.32	1	2
10	2	1044.13	993.59	5	1
11	1	794.83	851.45	1	3
12	2	762.15	1030.86	2	6
13	1	819.16	765.74	2	4
14	2	576.61	574.97	0	1
15	1	512.81	543.89	1	6
16	2	959.04	943.81	0	1
17	1	601.21	547.93	9	4
18	2	1044.13	993.59	5	1
	mean	730.84	742.40	3.56	3.89

Experiment 10c: Age split IAT with faces

Pleasant IAT

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Young Pleas	Young Neutral	Old Pleas	Old Neutral	Young Pleas	Young Neutral	Old Pleas	Old Neutral
Subject	Counter- balancing condition								
1	1	582.32	654.58	680.10	583.37	1	1	2	2
2	2	644.71	679.33	613.77	596.03	0	1	0	0
3	3	691.50	883.07	698.74	696.39	1	1	0	0
4	4	753.74	889.87	648.96	903.20	4	1	4	7
5	1	792.73	984.31	1186.70	713.18	2	3	5	3
6	2	752.10	792.74	809.84	1017.04	0	1	0	4
7	3	1335.80	807.93	932.57	1213.63	2	1	1	0
8	4	584.73	894.94	749.97	661.71	1	1	1	1
9	1	604.65	716.76	647.63	650.97	1	2	0	1
10	2	940.94	925.56	1008.63	882.87	1	5	5	2
11	3	1016.60	1234.73	1124.57	872.80	2	6	9	7
12	4	692.06	718.06	685.94	641.47	2	1	0	1
13	1	1033.14	1085.57	943.56	1182.00	3	2	5	4
14	2	839.53	1221.81	1031.23	737.43	0	0	1	1
15	3	597.86	627.63	523.75	678.09	3	0	0	0
16	4	560.36	626.44	525.39	611.52	3	0	3	1
	mean	776.42	858.96	800.71	790.10	1.63	1.63	2.25	2.13

Experiment 10c: Age split IAT with faces - continued

Unpleasant IAT

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Young Neutral	Young Unpleas	Old Neutral	Old Unpleas	Young Neutral	Young Unpleas	Old Neutral	Old Unpleas
Subject	Counter- balancing condition								
1	1	664.15	916.44	1193.21	522.57	6	4	8	4
2	2	956.38	1068.34	847.50	841.67	3	0	3	1
3	3	733.29	892.87	893.97	693.10	0	2	2	0
4	4	595.70	809.37	884.60	724.67	1	2	2	2
5	1	561.97	578.21	571.71	562.53	0	3	0	0
6	2	989.29	1118.41	1119.84	714.15	4	2	0	4
7	3	802.19	1246.06	892.69	640.97	0	1	0	2
8	4	664.63	734.03	775.19	579.14	0	1	1	4
9	1	569.07	1100.93	967.90	559.58	1	2	2	0
10	2	876.97	927.87	975.26	842.74	0	1	1	0
11	3	713.00	985.76	1089.72	609.30	4	3	3	2
12	4	819.50	1180.84	960.30	784.52	1	1	2	1
13	1	784.93	912.38	1089.79	705.87	1	7	2	0
14	2	874.44	1034.61	1336.63	714.30	0	1	0	1
15	3	683.58	969.00	831.70	643.23	0	0	1	1
16	4	1009.22	1049.81	945.86	738.19	0	0	4	0
	mean	768.64	970.31	960.99	679.78	1.31	1.88	1.94	1.38

Experiment 10d: Age split IAT with names

Pleasant IAT

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Young Pleas	Young Neutral	Old Pleas	Old Neutral	Young Pleas	Young Neutral	Old Pleas	Old Neutral
Subject	Counter- balancing condition								
1	1	868.50	961.31	924.73	1098.28	4	0	1	6
2	2	518.83	663.04	716.22	554.57	8	5	8	2
3	3	718.50	864.63	1031.00	844.59	5	5	4	3
4	4	1102.09	1748.72	1832.88	1215.86	0	0	8	3
5	1	767.81	962.48	766.15	837.14	6	5	4	2
6	2	896.71	851.21	884.54	814.96	4	3	6	4
7	3	638.85	830.81	763.50	681.69	5	5	6	0
8	4	589.91	1058.68	762.50	653.92	0	4	4	5
9	1	643.59	681.93	696.53	667.19	0	3	0	4
10	2	915.14	1142.00	1080.35	894.34	2	2	0	3
11	3	849.74	968.93	1062.83	851.55	1	3	2	1
12	4	776.94	628.20	655.03	717.76	0	2	0	3
13	1	802.68	1160.52	937.41	863.66	6	1	4	3
14	2	591.52	1120.16	1283.78	613.33	1	1	5	2
15	3	704.31	826.86	787.75	800.55	2	3	0	2
16	4	543.66	659.00	628.28	620.04	3	2	2	4
	mean	745.55	945.53	925.84	795.59	2.94	2.75	3.38	2.94

Experiment 10d: Age split IAT with names - continued

Unpleasant IAT

		Mean reaction	n time (ms)			Number of e	errors		
Condition		Young Neutral	Young Unpleas	Old Neutral	Old Unpleas	Young Neutral	Young Unpleas	Old Neutral	Old Unpleas
Subject	Counter- balancing condition						· · · · · ·		
1	1	679.38	581.84	707.13	683.74	2	6	9	5
2	2	1294.07	1033.53	851.31	1266.35	1	0	3	1
3	3	691.16	811.90	801.32	711.03	6	2	1	2
4	4	942.55	805.36	1029.67	1044.86	3	3	8	3
5	1	730.94	942.48	1033.04	649.20	1	9	8	2
6	2	1174.09	1179.33	1264.47	1052.40	0	2	2	2
7	3	826.23	1088.13	1139.83	1010.96	1	8	7	5
8	4	989.73	990.00	832.28	965.18	2	4	2	0
9	1	1040.07	1078.14	1264.03	1086.68	1	4	1	6
10	2	673.85	628.79	717.68	684.87	5	3	3	2
11	3	895.28	957.69	885.27	997.80	2	0	0	6
12	4	882.34	1587.40	1486.55	843.48	3	2	1	1
13	1	814.47	674.17	794.44	718.84	1	0	0	0
14	2	860.41	1186.14	899.56	683.90	2	6	2	2
15	3	877.34	737.13	780.45	1094.89	2	0	1	4
16	4	1090.90	968.93	1208.57	1105.26	2	3	4	5
	mean	903.93	953.18	980.98	912.47	2.13	3.25	3.25	2.88

Experiment 11a:

Mere acceptance IAT (Flight/Teeth)

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Accept Pleas	Accept Unpleas	Reject Pleas	Reject Unpleas	Accept Pleas	Accept Unpleas	Reject Pleas	Reject Unpleas
Subject	Counter- balancing condition								
1	1	888.84	1573.96	1631.43	610.96	0	4	2	3
2	2	696.28	839.07	743.03	639.86	0	2	0	4
3	3	797.00	1336.06	1408.37	643.41	0	0	1	2
4	4	1613.38	1081.38	1136.34	1460.30	2	2	3	5
5	5	985.86	711.64	696.23	1066.38	3	6	1	3
6	6	664.52	1516.87	1254.90	668.19	0	9	3	0
7	7	1461.09	2538.66	3404.73	1740.13	0	2	2	0
8	8	1042.81	1724.90	1656.45	1112.84	0	1	0	0
9	1	623.16	1171.64	1346.82	638.20	1	4	4	1
10	2	747.81	1071.54	1166.93	694.20	5	5	2	2
11	3	599.06	777.16	916.61	621.37	1	1	1	2
12	4	709.70	1127.79	1802.08	803.82	1	3	7	4
13	5	1178.39	1689.14	2087.03	1508.93	1	4	3	5
14	6	840.87	1403.77	2098.77	822.65	2	1	2	8
15	7	930.03	1536.86	2291.00	904.68	0	3	6	4
16	8	1164.91	1014.30	944.73	1240.16	0	1	1	0
	mean	933.98	1319.67	1536.59	948.50	1.00	3.00	2.38	2.69

Experiment 11a: - continued

Mere acceptance classification task (Flight/Teeth)

		Mean reaction t	ime (ms)	Number of erro	ors
Condition		Accept	Reject	Accept	Reject
Subject	Counter- balancing condition				
1	1	588.99	686.86	4	2
2	2	486.90	520.92	4	4
3	3	603.56	663.58	12	4
4	4	658.62	714.22	3	4
5	5	688.26	774.06	2	0
6	6	683.09	661.71	13	0
7	7	609.60	628.09	2	1
8	8	1243.36	1298.25	2	6
9	1	627.22	669.11	1	5
10	2	572.11	565.10	6	2
11	3	654.63	703.64	1	5
12	4	747.38	855.91	6	7
	mean	680.31	728.45	4.67	3.33

Experiment 11b: Mere acceptance split IAT (Flight/Teeth)

Pleasant IAT

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Accept Pleas	Accept Neutral	Reject Pleas	Reject Neutral	Accept Pleas	Accept Neutral	Reject Pleas	Reject Neutral
Subject	Counter- balancing condition								
1	1	717.93	1156.92	953.14	835.83	3	6	3	1
2	2	1229.34	578.03	628.83	1171.70	0	1	2	5
3	3	770.84	1012.93	1009.52	862.66	0	1	1	2
4	4	554.14	948.81	1109.48	563.25	3	0	5	0
5	1	672.10	1675.24	1423.37	626.10	0	3	2	1
6	2	572.13	1290.33	1161.54	672.23	1	1	4	0
7	3	582.28	1527.42	1305.52	601.59	2	6	5	0
8	4	966.30	1859.89	1397.94	1337.96	2	4	1	3
9	1	786.11	1008.10	977.57	626.42	3	2	8	1
10	2	716.03	1657.28	1208.17	813.35	0	3	8	1
11	3	637.91	1322.56	1493.73	753.73	0	7	6	2
12	4	562.14	1408.08	1669.54	694.11	2	7	6	4
13	1	743.80	1494.63	1593.32	863.37	6	8	7	4
14	2	743.40	3005.00	2215.90	1503.23	1	0	2	1
15	3	625.97	1301.44	1301.43	630.29	1	7	4	0
16	4	861.09	2044.94	2011.12	965.00	0	1	7	1
	mean	733.84	1455.73	1341.26	845.05	1.50	3.56	4.44	1.63

Experiment 11b: Mere acceptance split IAT (Flight/Teeth) - continued

Unpleasant IAT

		Mean reactic	on time (ms)			Number of e	errors		
Condition		Accept Neutral	Accept Unpleas	Reject Neutral	Reject Unpleas	Accept Neutral	Accept Unpleas	Reject Neutral	Reject Unpleas
Subject	Counter- balancing condition								
1	1	1306.20	1018.17	1152.58	1114.31	2	1	0	2
2	2	1414.79	768.60	744.00	1890.41	3	1	2	3
3	3	1631.43	847.18	981.81	1653.65	4	3	0	6
4	4	629.72	510.15	554.92	958.74	6	5	6	5
5	1	1348.52	1784.47	1203.66	1624.23	0	0	2	1
6	2	708.36	597.35	621.04	654.26	7	9	7	5
7	3	808.80	613.03	594.28	1102.12	1	0	3	6
8	4	968.43	657.83	778.67	1598.74	1	1	1	1
9	1	1612.61	1449.93	1232.15	1684.13	0	2	4	0
10	2	931.15	583.52	949.86	797.41	6	2	4	4
11	3	1215.92	995.89	1307.10	1275.90	6	4	3	2
12	4	1755.16	899.70	1005.40	1639.82	0	4	1	9
13	1	912.94	576.20	637.23	858.48	1	2	1	4
14	2	687.03	1082.39	1073.78	868.96	2	4	5	6
15	3	1070.46	555.62	625.03	1070.83	4	3	0	8
16	4	1670.13	923.97	1252.23	1703.90	2	2	2	1
	mean	1166.98	866.50	919.61	1280.99	2.81	2.69	2.56	3.94

Experiment 12a:

Mere acceptance IAT (Row condition)

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Row Same Pleas	Row Same Unpleas	Row Different Pleas	Row Different Unpleas	Row Same Pleas	Row Same Unpleas	Row Different Pleas	Row Different Unpleas
Subject	Counter- balancing condition								
1	1	553.68	846.34	881.11	594.17	4	3	5	2
2	2	651.21	856.80	1063.75	663.93	2	6	4	3
3	3	654.91	823.69	875.43	811.60	0	5	2	2
4	4	881.65	906.56	874.45	944.00	1	0	0	1
5	1	717.35	1107.61	1338.07	719.03	0	3	2	0
6	2	616.26	824.00	949.65	624.70	1	2	1	2
7	3	781.25	1167.32	1147.59	1608.90	3	1	0	3
8	4	596.45	754.13	893.25	628.76	1	2	0	2
	mean	681.59	910.81	1002.91	824.39	1.50	2.75	1.75	1.88

Mere acceptance IAT (Column condition)

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Column Same Pleas	Column Same Unpleas	Column Different Pleas	Column Different Unpleas	Column Same Pleas	Column Same Unpleas	Column Different Pleas	Column Different Unpleas
Subject	Counter- balancing condition								
1	1	739.10	1700.32	2022.40	741.03	2	4	7	2
2	2	849.35	1547.77	1674.47	810.52	1	1	1	1
3	3	581.97	1195.54	1164.20	669.11	1	4	2	4
4	4	1233.39	1046.53	1144.58	991.26	4	2	1	0
5	1	1153.13	1469.97	1529.52	1141.19	1	0	0	0
6	2	990.13	1918.73	1336.88	807.13	1	2	0	1
7	3	687.94	1963.59	1645.90	889.39	0	0	0	1
8	4	618.63	721.59	669.04	684.97	0	2	4	0
	mean	856.70	1445.50	1398.37	841.82	1.25	1.88	1.88	1.13

Experiment 12a: - continued

Mere acceptance classification task (Row condition)

		Mean reaction ti	me (ms)	Number of errors		
Condition		Row Same	Row Different	Row Same	Row Different	
Subject	Counterbalancing condition					
1	1	550.82	801.69	0	4	
2	2	645.32	947.31	0	0	
3	1	788.99	793.51	1	4	
4	2	702.42	927.62	2	0	
5	1	767.74	847.26	2	1	
6	2	514.92	730.88	5	7	
7	1	871.85	1286.93	6	5	
8	2	586.23	681.66	2	0	
9	1	737.10	1048.01	2	0	
10	2	582.76	866.92	0	0	
	mean	674.81	893.18	2.00	2.10	

Mere acceptance classification task (Column condition)

		Mean reaction time	e (ms)	Number of errors	
Condition		Column Same	Column Different	Column Same	Column Different
Subject	Counterbalancing condition				
1	1	863.14	653.56	22	15
2	2	1741.45	1526.18	6	17
3	1	1109.64	825.90	0	0
4	2	670.20	569.41	0	3
5	1	2180.59	1785.69	0	0
6	2	2246.78	1851.88	11	5
7	1	1681.97	1453.11	4	0
8	2	750.08	809.76	6	4
9	1	1663.44	1391.86	5	1
10	2	760.47	640.41	7	5
	mean	1366.78	1150.78	6.10	5.00

Experiment 12b: Mere acceptance split IAT

Pleasant IAT (Row condition)

		Mean reaction	on time (ms)			Number of e	errors		
Condition		Row Same Pleas	Row Same Neutral	Row Different Pleas	Row Different Neutral	Row Same Pleas	Row Same Neutral	Row Different Pleas	Row Different Neutral
Subject	Counter- balancing condition								
1	1	859.77	902.03	1050.96	1039.00	2	2	3	3
2	2	511.35	637.14	736.84	567.20	1	2	1	2
3	3	793.71	1462.41	1335.38	909.97	0	0	2	1
4	4	761.04	982.00	1117.04	843.83	3	2	4	1
5	1	559.37	693.48	670.24	596.22	2	5	6	4
6	2	659.19	770.83	747.03	724.13	0	1	2	0
7	3	702.83	1290.38	1898.68	1208.21	1	6	4	2
8	4	616.19	1402.17	1028.21	658.00	5	3	3	0
9	1	678.93	911.71	1548.92	746.77	3	7	6	6
10	2	625.47	996.78	1155.10	747.78	1	0	1	0
11	3	738.61	799.78	888.23	801.63	1	2	1	0
12	4	529.73	817.16	953.00	595.35	2	0	4	0
	mean	669.68	972.16	1094.14	786.51	1.75	2.50	3.08	1.58

Experiment 12b: Mere acceptance split IAT - continued

Unpleasant IAT (Row condition)

		Mean reaction time (ms)				Number of errors			
Condition		Row Same Neutral	Row Same Unpleas	Row Different Neutral	Row Different Unpleas	Row Same Neutral	Row Same Unpleas	Row Different Neutral	Row Different Unpleas
Subject	Counter- balancing condition								
1	1	604.57	593.93	641.75	635.55	2	2	0	2
2	2	1297.38	653.66	727.13	991.66	3	2	0	2
3	3	1222.48	692.00	950.16	1434.90	2	0	0	1
4	4	752.61	789.43	839.39	791.17	1	3	1	2
5	1	718.03	545.19	574.80	722.19	3	0	2	1
6	2	939.67	660.50	742.66	1178.15	8	3	2	5
7	3	1355.38	1139.66	1096.50	1747.35	0	0	0	1
8	4	606.13	552.34	623.39	579.67	2	0	4	2
9	1	1179.97	1407.21	1552.76	1170.23	1	4	3	2
10	2	845.40	653.32	986.52	824.50	1	0	0	1
11	3	794.14	590.57	685.00	779.56	4	2	1	5
12	4	942.33	628.67	809.87	871.74	1	1	2	1
	mean	938.17	742.21	852.49	977.22	2.33	1.42	1.25	2.08

Pleasant IAT (Column condition)

		Mean reaction time (ms)				Number of errors				
Condition		Column Same Pleas	Column Same Neutral	Column Different Pleas	Column Different Neutral	Column Same Pleas	Column Same Neutral	Column Different Pleas	Column Different Neutral	
Subject	Counter- balancing condition									
1	1	1186.10	1477.41	1745.21	1100.19	2	2	3	1	
2	2	1072.83	2278.59	1934.00	993.77	2	3	1	6	
3	3	640.70	1580.61	1201.19	696.75	3	4	5	4	
4	4	1238.38	1690.15	1562.67	880.76	2	6	5	2	
5	1	678.42	1413.85	1407.32	827.47	0	5	1	0	
6	2	783.31	1104.37	1179.13	747.94	0	2	1	0	
7	3	727.04	1254.07	1219.96	827.55	5	3	6	1	
8	4	1498.59	1069.38	832.04	1297.74	3	2	4	1	
9	1	975.52	1045.66	1000.06	855.24	1	0	1	2	
10	2	1003.45	1928.53	1418.93	898.60	1	0	3	1	
11	3	961.33	2123.00	1908.72	973.84	1	5	7	1	
12	4	642.76	1496.00	1239.59	741.39	3	1	5	1	
	mean	950.70	1538.47	1387.40	903.44	1.92	2.75	3.50	1.67	

Unpleasant IAT (Column condition)

		Mean reaction time (ms)				Number of errors			
Condition		Column Same Neutral	Column Same Unpleas	Column Different Neutral	Column Different Unpleas	Column Same Neutral	Column Same Unpleas	Column Different Neutral	Column Different Unpleas
Subject	Counter- balancing condition								
1	1	1205.50	810.87	865.13	1012.24	2	1	1	3
2	2	1401.24	1538.97	1371.31	1025.00	3	1	0	0
3	3	1586.78	1485.13	1415.48	1863.27	4	2	1	6
4	4	987.10	1848.89	1820.33	935.69	1	4	5	0
5	1	1435.38	1137.73	1399.03	1180.25	3	2	1	0
6	2	1484.81	1915.65	1741.50	1086.80	0	1	0	1
7	3	1305.28	1334.61	1366.28	1032.74	3	1	2	4
8	4	941.33	938.87	967.50	1034.59	2	2	1	0
9	1	1259.19	1097.72	1110.71	990.28	0	0	0	0
10	2	1389.55	1207.87	884.71	1290.60	2	1	0	2
11	3	1538.97	1700.31	1277.03	996.66	0	0	0	0
12	4	1118.83	1103.03	818.19	1092.38	2	3	1	3
	mean	1304.50	1343.30	1253.10	1128.37	1.83	1.50	1.00	1.58
FIN LIB.

.