The Quarterly Journal of Experimental Psychology
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t716100704

Processing fluency as a predictor of salience asymmetries in the Implicit Association Test

Betty P. I. Chang [a]; Chris J. Mitchell [a]
[a] University of New South Wales, Sydney, New South Wales, Australia

First Published on: 26 February 2009

To cite this Article Chang, Betty P. I. and Mitchell, Chris J.(2009)'Processing fluency as a predictor of salience asymmetries in the Implicit Association Test',The Quarterly Journal of Experimental Psychology,62:10,2030 — 2054

To link to this Article: DOI: 10.1080/17470210802651737

URL: http://dx.doi.org/10.1080/17470210802651737

THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY
2009, 62 (10), 2030–2054

Psychology Press
Taylor & Francis Group

Processing fluency as a predictor of salience asymmetries in the Implicit Association Test

Betty P. I. Chang and Chris J. Mitchell
University of New South Wales, Sydney, New South Wales, Australia

The Implicit Association Test (IAT) is the most popular indirect measure of attitudes in social psychology. Rothermund and Wentura (2001, 2004) suggested that artifacts such as salience asymmetries are a source of compatibility effects in the IAT, and, therefore, the IAT does not necessarily measure attitude. They claim that salience asymmetries correspond with visual search asymmetries, such that the stimulus categories that are more quickly detected in a visual search task are also compatible in the IAT. We propose that processing fluency is a more reliable indicator of salience asymmetries in the IAT than are visual search asymmetries. To test this hypothesis, we set processing fluency in opposition to visual search asymmetry to see which variable better predicted IAT effects. In one pair of categories, the category that was more quickly detected in visual search was also more fluently processed in a binary classification task. In a second pair of categories, the category that was more quickly detected in visual search was the less fluently processed category. Across four experiments, we demonstrated that compatibility effects in the IAT corresponded with differences in processing fluency between categories, rather than with visual search asymmetries.

Keywords: Implicit Association Test; Salience asymmetry; Figure–ground asymmetry; Processing fluency; Visual search asymmetry.

The Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998) is the most widely used indirect measure of attitudes in psychology. The IAT gauges the attitude to a target concept by measuring how quickly target exemplars are classified with pleasant or unpleasant attributes in two speeded classification tasks. For example, in a race IAT, stimuli belonging to the target categories of white people (e.g., photos of white faces) and black people (e.g., photos of black faces) serve as target stimuli, which are categorized with pleasant words (e.g., peace, joy) and unpleasant words (cancer, death), known as attributes (Dasgupta, McGhee, Greenwald, & Banaji, 2000). In one classification task, white and pleasant attributes share one response key, and black and unpleasant attributes share another response key. In another task the reverse is true; white and unpleasant attributes share one response key, and black and pleasant attributes share another response key.

Correspondence should be addressed to Betty Chang, School of Psychology, University of New South Wales, Sydney, NSW, 2052, Australia. E-mail bchang@psy.unsw.edu.au

We thank Sachiko Kinoshita, Klaus Rothermund, and an anonymous reviewer for their helpful comments on a previous draft of this paper. We are also grateful to Klaus Rothermund for his advice in preparing Experiments 3a and 3b.

http://www.psypress.com/qjep
DOI:10.1080/17470210802651737

The conventional finding is that white people more easily classify white with pleasant attributes (responding is faster and more accurate) than with unpleasant attributes, and white people more easily classify white than black with pleasant attributes (Greenwald et al., 1998). This pattern of data represents an IAT effect for white. From this it is inferred that white people are preferred to black people.

Not all IAT effects can be easily explained by evaluative associations. In an IAT with insects and nonwords as the target categories, people find it easier to classify insects with pleasant words, and nonwords with unpleasant words, than vice versa (Brendl, Markman, & Messner, 2001; Experiment 2). This result contradicts the associative account of the IAT. Because nonwords are novel, they should have no prior associations. Therefore, the associative account of the IAT would predict that nonwords should be equally compatible with both pleasant and unpleasant words. Morever, because insects were rated more negatively than nonwords on a self-report measure, insects should have been more compatible with unpleasant words than with pleasant words.

Rothermund and Wentura (2001, 2004) explained this effect in terms of salience asymmetries, or figure–ground asymmetries. They suggested that in many circumstances, the unpleasant attribute category is the more salient "figure" category, making it the focus of attention in comparison to the less salient "ground" category of pleasant attributes. This salience asymmetry between the attribute categories makes unpleasant attributes compatible with whichever target category is more salient. In the case of the insect/nonword IAT effect, they supposed nonwords to be more salient than words (insect names), as nonwords are less familiar than words. Thus, in the insect/nonword IAT, nonwords and unpleasant attributes are easily categorized together because they are the more salient of the target and attribute categories, respectively. Rothermund and Wentura (2001, 2004) based their assumptions about the salience of negative and unfamiliar items on findings from the visual search literature. A variety of visual search studies show that negative items (e.g., angry faces) are detected more quickly among an array of positive items (e.g., happy faces) than vice versa (Fox et al., 2000; Ohman, Flykt, & Esteves, 2001a). Similarly, unfamiliar targets are more readily detected among familiar distractors than vice versa (e.g., Strayer & Johnston, 2000; Wang, Cavanagh, & Green, 1994). This search asymmetry has been taken as evidence that negative and unfamiliar stimuli capture and/or hold attention to a greater extent than positive and familiar items and are therefore more salient.

One problem with the visual search asymmetry account of salience asymmetries in the IAT is that search asymmetries involving affective stimuli do not fall along a 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 were embedded among neutral distractors and when they were embedded among other emotional faces. This led Ohman et al. (2001b) to argue that angry advantage displayed in other visual search experiments (e.g., Fox et al., 2000) 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 on the specific emotion that is shown. Because most IATs use positive and negative attributes that are not specific to a particular emotion, this implies that there is probably no consistent mapping between affective search asymmetries and affective salience asymmetries in the IAT.

Despite the fact that visual search asymmetries do not fall along a simple positive/negative dichotomy, Rothermund and Wentura (2004) managed to demonstrate a consistent correspondence between visual search asymmetries and IAT effects. To investigate salience asymmetries in the IAT, Rothermund and Wentura (2004; Experiments 1b, 1a, 1c, and 1e) adopted visual search measures to cross-validate their salience manipulations. Using a variety of stimuli (e.g., young vs. old names, single- vs. multicoloured strings), they found that stimuli from certain categories (i.e., old names and multicoloured strings) were detected more quickly from among stimuli of the other category (young names and single-coloured strings) than the reverse. These search asymmetries were taken by Rothermund and Wentura to indicate salience asymmetries between the two categories. When the old and young names were then classified in an IAT with multi- and single-coloured strings, responses were faster when old names shared a key with multicoloured strings, and young names shared a key with single-coloured strings, than with the reverse assignment. From this, Rothermund and Wentura concluded that the category that is detected more quickly in visual search (old names and multicoloured strings) is also more salient in the IAT, and that the more salient categories are compatible 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.

Kinoshita and Peek-O'Leary (2006) have also suggested that IAT effects can be caused by salience asymmetries between target categories. However, they argue that it is the familiar and positive categories that are the figure, making those categories compatible in the IAT (Kinoshita & Peek-O'Leary, 2005, 2006). They consider that familiar and positive categories are the figure in the IAT because these are also the figure categories in psycholinguistics research (e.g., Clark, 1973; Greenberg, 1966). Although Kinoshita and Peek-O'Leary (2005, 2006) disagree with Rothermund and Wentura (2001, 2004) as to exactly which category pairs are the figure (positive and familiar are the figure in the former case, and negative and unfamiliar are the figure in the latter case), their account is consistent with Rothermund and Wentura's position that familiar and positive stimuli are compatible, and unfamiliar and negative stimuli are compatible in terms of salience in the IAT.

Kinoshita and Peek-O'Leary (2006) point out that the categories they claim to be the figure—that is, familiar and positive categories—are generally more fluently processed. This is supported by research showing that familiar words are processed more quickly than unfamiliar words on a lexical decision task (Balota & Chumbley, 1984; Whaley, 1978) and that 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, 2008; see Unkelbach, Fiedler, Bayer, Stegmüller, & Danner, 2008, for a similar result involving German words). Building upon this notion, we suggest that processing fluency may be a better indicator of nonassociative compatibility effects in the IAT than are visual search asymmetries.

Thus it appears that visual search asymmetry and processing fluency are two potential predictors of salience asymmetry effects in the IAT. These two factors can vary independently between category pairs; in some cases, stimuli that are detected more quickly in a visual search task are less fluently processed, but in other cases stimuli that are detected more quickly in a visual search task are more fluently processed. Examples of these conditions are described in the experimental design below. Because search asymmetries and processing fluency can vary independently, this suggests that only one of these factors can be diagnostic of salience asymmetries in the IAT. Therefore it is important to clarify whether it is processing fluency or visual search asymmetries that are associated with salience asymmetry in the IAT. This will allow us to better identify circumstances in which salience asymmetries may occur and to more accurately interpret IAT effects.

To determine which is a better indicator of salience asymmetries in the IAT—visual search asymmetries or processing fluency—we manipulated visual search asymmetries and processing fluency independently. This led us to create two different category pairs that varied orthogonally in these dimensions. To anticipate the results of Experiments 1b–1d, in one pair of categories (upright elephants vs. inverted elephants), the category that was more quickly detected on a visual search task (inverted elephants) was processed more slowly on a separate binary classification task than was the other category (upright elephants). In another pair of categories (big cows vs. small cows) the category that was more easily detected in the visual search task (big cows) was classified more quickly on a separate binary classification task than was 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 similiarly 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 (i.e., upright elephants and big cows) to produce similar effects in the IAT.

EXPERIMENTS 1A AND 1B

Experiments 1a and 1b were conducted to verify that the stimuli used in our experiments exhibited the expected visual search asymmetries. To do this, we used a same/different visual search task in which participants are required to judge whether all the stimuli in a visual search array are the same, or whether one of the stimuli is different. The same/different visual search task was deployed instead of the more conventional directed search task (in which participants are required to respond whether a prespecified target is present or absent) for two reasons. First, this type of task was used by Rothermund and Wentura (2004) to assess search asymmetries between target categories. Secondly, prespecifying a target means that search performance is influenced by participants actively searching for a target, which may obscure naturally occurring search asymmetries. For example, directed search for a familiar upright elephant may speed up detection for the target more so than directed search for an unfamiliar inverted elephant, because participants are likely to have a stronger mental representation of what to look for in an upright elephant than in an inverted elephant. In this way, directed search may artificially create salience asymmetries.

Treisman and Gormican (1988, Experiment 1) demonstrated that a big item (an 8-mm line) also was detected more quickly among smaller items (6.5-mm or 5-mm lines), than vice versa. Based on this principle, we created another two categories that differed only in size: big cows versus 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. We also predicted that search for an inverted elephant silhouette among upright elephant silhouettes would be faster than the reverse, as previously demonstrated by Wolfe, Alvarez, Wong, and Klempen, 2000 (cited in Wolfe, 2001).

Method

Participants

First-year psychology students from the University of New South Wales volunteered to participate in the experiment in exchange for course credit. A total of 8 students participated in Experiment 1a (cow visual search task), and 12 students participated in Experiment 1b (elephant visual search task).

Stimuli and apparatus

The experiment was conducted on an IBM-compatible PC, which was used in all the following experiments. The stimuli in Experient 1a were four coloured line drawings of cows on a background of grass. The cows had identical features, except that two of the cows measured 27.8 mm × 20.1 mm, and the other two measured 44.6 mm × 33.6 mm. All cows were presented

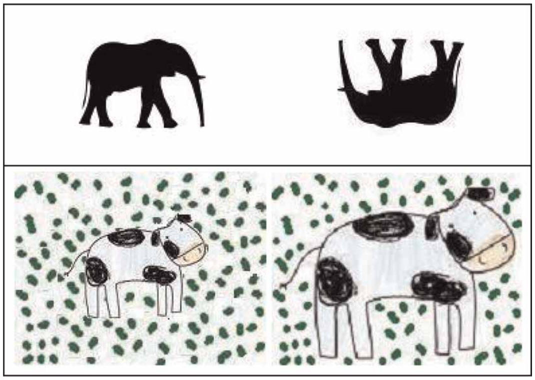

Figure 1. *Elephant stimuli (above) and cow stimuli (below). The stimuli from the elephant categories differ in their upright/inverted orientation, whereas stimuli from the the cow categories differ in size. [To view a colour version of this figure, please see the online issue of the Journal.]*

on a background of green dots that was 49.5 mm × 38.1 mm. The elephant and cow stimuli are presented in Figure 1. For half the participants, the two smaller cows belonged to the "Zif" category (denoting one breed of cow), and the two larger cows belonged to the "Wug" category (denoting another 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 Rothermund and Wentura (2001, 2004) and Kinoshita and Peek-O'Leary (2005, 2006), 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 produce, or their processing fluency). Although linguistic markedness as a source of salience asymmetries is more of an issue in later experiments, these labelling measures were adopted in Experiment 1a to maintain consistency of the stimuli across experiments. The stimuli in Experiment 1b consisted of four elephant sihouette drawings that measured 24.6 mm × 17.5 mm. Two elephants were upright; they belonged to the category of "live" elephants. The other two were inverted; they belonged to the category of "dead" elephants. In each category, one elephant faced left, and the other faced right.

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 the Zif or Wug category (Experiment 1a), or the live or dead category (Experiment 1b). 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 on screen, 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. 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 on screen until a response was made. Incorrect responses received the feedback of "WRONG RESPONSE" presented in red font at the centre of the screen for 200 ms. In Experiment 1a, this feedback allowed participants to learn which size of cow belonged to which breed.

Following this, participants completed the visual search trials, consisting of 12 practice trials and 192 test trials. A fixation cross appeared for 500 ms 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 (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., big/small cows) by two response types (same/different). There were 24 randomized test trials in each of the four conditions. 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.

Results

One participant was replaced in Experiment 1a for committing 16 errors in at least one of the conditions (33.3%). Erroneous responses were omitted from the analysis (5.5% in Experiment 1a, 5.6% in Experiment 1b), as were those that were 3.5 standard deviations above the mean in each condition (0.7% in Experiment 1a, 1.0% in Experiment 1b).

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

In Experiment 1b, when there was an odd elephant out, participants were faster to detect an inverted elephant among upright elephants (1,087 ms) than vice versa (1,147 ms), $t(11) = 2.72$, $p < .05$. The search asymmetry for inverted elephants was also characterized by fewer errors ($M_{\text{error}} = 2.33$) than in the condition where the upright elephant was the odd one out ($M_{\text{error}} = 3.50$), $t(11) = 2.65$, $p < .05$. Participants were equally quick to respond to an array that consisted only of upright elephants (1,127 ms) as they were to an array that consisted only of inverted elephants (1,131 ms), $t < 1$. However, more errors were committed in the latter condition ($M_{\text{error}} = 1.83$ for upright elephants, $M_{\text{error}} = 3.08$ 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.

EXPERIMENTS 1C AND 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, we 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. For Experiment 1d, we 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 preexisiting mental representation and thus would require more processing to reach the threshold of stimulus identification.

Method

Participants

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

Procedure

The stimuli were 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, big cows facing left, big cows facing right, small cows facing left, small cows facing right), the first 16 of which were practice trials (4 presentations of each stimulus). The intertrial interval was 300 ms. On each trial, a stimulus was presented on screen, 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 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. Those that completed the elephant binary classification task categorized elephants as being either live (upright) or dead (inverted). The category labels "Zif" and "Wug" or "live" and "dead" 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 200 ms.

Results

One participant was replaced in Experiment 2d for having an error rate of 47.9%, suggesting that she or he was not able to distinguish between the two categories in the classification task. Erroneous responses were excluded from the analysis (4.5% in Experiment 1c, 4.0% in Experiment 1d), as were those that were below 200 ms (2.1% in Experiment 1c, none in Experiment 1d) and those that were 3.5 standard deviations above the mean in each condition (0.1 % in Experiment 1c, 1.3% in Experiment 1d).

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

Experiments 1c–1d demonstrated that the elephant and cow category pairs exhibited different patterns of search asymmetries and fluency asymmetries. In the cow pairing, the category that was more quickly detected in the visual search task (big cow) was more fluently responded to on a classification task. In contrast, 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. Therefore, these stimuli can be used to investigate the effects of search asymmetries and fluency asymmetries in the IAT.

EXPERIMENTS 2A AND 2B

Experiments 2a and 2b examined whether stimuli that are detected more quickly in visual search (big cows and inverted elephants) 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, and are detected more quickly among words in a visual search task than vice versa (Rothermund & Wentura, 2004, Experiment 1c). 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 categories that were harder to detect in visual search when there was a search asymmetry (i.e., small cows and upright elephants) to be compatible with words and the animal category that was easier to distinguish among stimuli of the other category (i.e., big cows and inverted elephants) to be compatible with nonwords. It may be, however, that processing fluency is the basis of salience asymmetries in the IAT. On this view, it is predicted that the more fluently processed animal categories (big cows and upright elephants) will be more easily classified with words than with nonwords, as one would expect words to be more fluently processed than nonwords due to their familiarity. This was tested in Experiments 2a and 2b by placing cow and elephant targets, respectively, in an IAT with words and nonwords. Note that Experiment 2a produces divergent predictions for the fluency and visual search hypotheses of salience asymmetries. If IAT effects align with fluency asymmetries, then big cows should be compatible with words. However, if IAT effects align with visual search asymmetries, then big cows should be compatible with nonwords. In contrast, Experiment 2b produces convergent predictions for the fluency and visual search hypotheses. Both accounts would predict that upright elephants and words should be compatible in the IAT. Therefore, only Experiment 2a can distinguish between the two accounts of salience asymmetries in the IAT, whereas Experiment 2b functions as a control condition.

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 16 participants in Experiment 2a (cow/word IAT) and 8 participants in Experiment 2b (elephant/word IAT).

Stimuli and apparatus

Experiments 2a and 2b used the same stimuli as those from Experiments 1c and 1d, respectively. In addition, Experiments 2a–2b included eight word stimuli (*angry*, *bad*, *cold*, *crude*, *mean*, *rude, cruel*, *nasty*) and eight nonword stimuli (*clure*, *cren*, *dolab*, *druc*, *meed*, *nady*, *staun*, *yarg*). We selected word stimuli with an unpleasant meaning to minimize the possibility that a picture category would be compatible with the word category on an affective basis, because words might be preferred to nonwords due to their greater familiarity (Zajonc, 1968). Participants classified the word stimuli as belonging to the category "English" and the nonword stimuli as belonging to the category "foreign". We also examined whether there were any processing fluency differences between words and nonwords using a separate binary classification task. Participants classified words (603 ms) more quickly than nonwords (624 ms), $t(23) = 2.54$, $p < .05$. Response accuracy was the same in both conditions ($M_{error} = 5.21$ vs. 4.54, respectively), $t < 1$, suggesting that there was no speed–accuracy trade-off in the corresponding reaction times.

Procedure

Experiments 2a and 2b. Participants performed five classification tasks in the IAT: (a) an animal classification task (Zif vs. Wug cows, or live vs. dead elephants), (b) a word classification task (English vs. foreign words), (c) an animal and word/combined classification task, (d) an animal classification task with reversed response assignment, and (e) an animal and word combined classification task with reversed pairings.

The animal classification task (Task 1) was the same as the classification trials of Experiments 1c and 1d, except that participants only received 24 trials. In addition, participants in Experiment 2a (cow/word IAT) were told that the two categories of cows that 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. The word classification task (Task 2) followed the same procedure as the animal classification task (24 trials), except that participants classified words as belonging to the "English" category and nonwords as belonging to the "foreign" category.

In the combined classification task (Task 3) participants categorized the animal categories with the word categories using the same keys as those that were assigned to them in Tasks 1 and 2. Thus, 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"). In Task 1, half the participants in Experiment 2a classified big cows using the "a" key and small cows using the "5" key, and half the participant in Experiment 2b classified upright elephants using the "a" key and inverted elephants using the "5" key. For the other half of participants in Experiments 2a and 2b, the response assignment for the animal categories was reversed. The category labels "English"/"foreign", and "Zif"/"Wug" or "live"/ "dead" were presented on the side of the screen corresponding 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 (e.g., big or small cows facing left/right) presented eight times and each of the eight word/nonword stimuli presented twice. 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.

The animal classification task (Task 4) with reversed response assignment was similar to Task 1, except for two changes. First, the response assignment to the animal categories was reversed. Secondly, the number of trials was doubled from 24 to 48, to reduce the carry-over order effects from the response assignment of the first animal classification task (Nosek, Greenwald, & Banaji, 2005). Participants then completed the animal and word combined classification task (Task 5) with reversed pairings. This task was the same as the first combined classification task, except with the new response assignment to the animal categories from Task 4.

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

Results

A total of 5 participants in Experiment 2a were replaced as they made 10 or more errors in at least one of the combined classification conditions. No participants fitted this replacement criterion from Experiment 2b. Incorrect responses were discarded from the analysis (6.1% in Experiment 2a, 11.0% in Experiment 2b), as were those that were less than 300 ms[1] (0.0% in Experiment 2a, 0.5% in Experiment 2b) or 3.5 standard deviations above the mean for each participant for the condition (1.1% in Experiment 2a, 0.9% in Experiment 2b). Mean reaction time scores were calculated for each of the two combined classification tasks in Experiments 2a and 2b. Because some IAT experiments required the analysis of an interaction effect, all IAT analyses were peformed using an analysis of variance (ANOVA) to maintain consistency.

Experiment 2a: Cow/word IAT

The analysis for Experiment 2a was conducted with the breed of cow (Zif/Wug) assigned to the large cow category as a between-groups factor. This factor did not interact with the contrasts of interest, $F(1, 8) = 2.56$, $MSE = 21{,}406.22$, $p > .13$ for reaction times, and $F < 1$ for errors. Figure 2 shows that the reaction time for the task in which big cows were classified with words, and small cows were classified with nonwords, was 109 ms faster than the reaction time for the task in which the reverse was true, $F(1, 8) = 8.92$, $MSE = 21{,}406.22$, $p < .05$. There was no difference between the number of errors made in the two conditions ($M_{\text{error}} = 1.63$

[1] The IAT experiments used a 300-ms 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 (200 ms), which is in keeping with the relative difficulty of the two tasks.

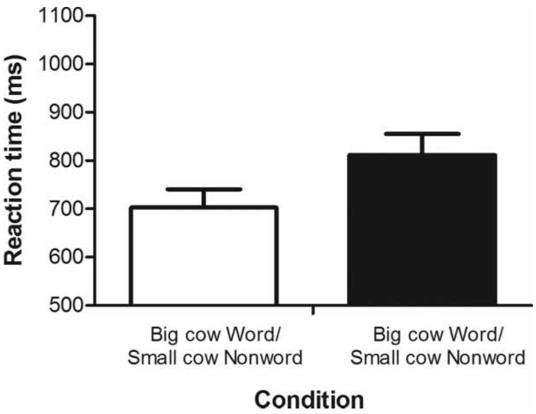

Figure 2. *Mean reaction times for the cow/word Implicit Association Test (IAT) of Experiment 2a. The open bar represents the condition in which big cows were classified with words, and small cows were classified with nonwords. The filled bar represents the condition where big cows were classified with nonwords, and small cows were classified with words. Error bars represent the standard error of the mean.*

vs. 2.25, respectively); $F(1, 8) = 1.67$, $MSE = 3.75$, $p = .23$, which suggests that the IAT effect seen in reaction times was not due to a speed–accuracy trade-off.

Experiment 2b: Elephant/word IAT

Figure 3 shows that participants responded 188 ms faster when upright elephants shared a key with words, and inverted elephants shared a key with nonwords, than vice versa $F(1, 7) = 10.58$, $MSE = 26{,}782.71$, $p < .05$. There was also a trend toward responses being more accurate in this condition ($M_{error} = 3.31$ for upright words/inverted nonwords vs. $M_{error} = 3.75$ for upright nonwords/inverted words), $F(1, 7) = 3.94$, $MSE = 0.39$, $p = .09$.

Discussion

This result of Experiment 2b, in which upright elephants were more easily classified with words, and inverted elephants were more easily classified with nonwords, 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

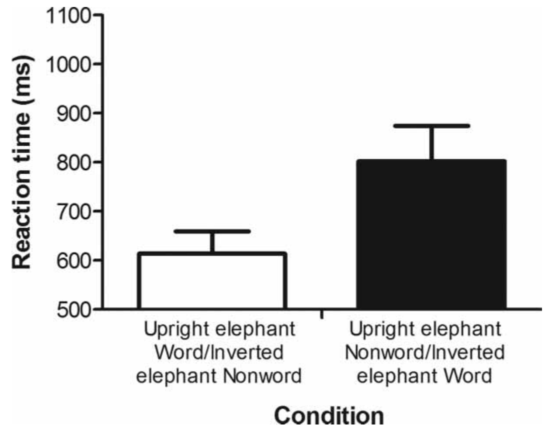

Figure 3. *Mean reaction times for the elephant/word Implicit Association Test (IAT) of Experiment 2b. The open bar represents the condition in which upright elephants were classified with words, and inverted elephants were classified with nonwords. The filled bar represents the condition where upright elephants were classified with nonwords, and inverted elephants were classified with words. Error bars represent the standard error of the mean.*

of Experiment 2a 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 similiarly in a visual search task (inverted elephants and big cows) do not, as Rothermund and Wentura (2001, 2004) claim, behave similarly in the IAT. Instead, we see that the more fluently processed (as indexed by the binary classifcation task) animal categories (big cows and upright elephants) are compatible with words, and the less fluently processed animal categories (small cows and inverted elephants) are compatible with nonwords.

The finding that word stimuli were more fluently processed than the nonword stimuli 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 and 2b suggest that processing fluency is a more reliable predictor of nonassociative compatibility effects in the IAT than are visual search asymmetries.

However, it might be argued that the compatibility effects found are due to affective associations instead of salience asymmetries. Although the category English consisted of unpleasant words, participants may actually find English words preferable to foreign words if they process the items with respect to their category-level features. That is, English stimuli may be preferred to foreign stimuli because the category "English" is preferred to the category of "foreign". Indeed, De Houwer (2001) has demonstrated that a positive target category is still compatible with positive attributes even when half of the target exemplars are negative (but see Govan & Williams, 2004). With reference to Experiment 2b, "live" elephants, which are presumably more positive than inverted "dead" elephants, may be compatible with "English" words because "English" is more positive than "foreign". Similarly, with respect to Experiment 2a, big cows may be compatible with "English" words because these two categories are the preferred categories in their respective pairings. The next series of experiments was conducted to address this possibility.

EXPERIMENTS 3A AND 3B

In Experiments 2a and 2b, the two animal categories that were more fluently processed in a binary classification task (upright elephants and big cows) were compatible with the more fluently processed words over nonwords in the IAT. These results suggest that differences in processing fluency may correspond with compatibility effects in the IAT. However, this interpretation rests on the assumption that the results of Experiments 2a and 2b represent salience asymmetry effects and are not due to evaluative associations between the animal and word categories. In Experiments 3a and 3b, we tested whether valence confounds between the animal and word categories may be responsible for the IAT effects observed in Experiments 2a and 2b. This was addressed using two methods. First, we replicated the previous experiments, but replaced the word category labels of "English" and "foreign" with the more neutral labels of "word" and "nonword", respectively. Secondly, we had participants rate the pleasantness of each stimulus and category presented in the IAT. If the animal and word categories that were found to be compatible were also the two preferred categories in their respective pairings, this would provide evidence that the IAT effects seen were due to affective associations.

To further investigate the contribution of salience asymmetries and/or evaluative associations to the IAT effects of Experiments 2a and 2b, we conducted a correlational study to examine whether individual differences in visual search asymmetries, processing fluency asymmetries, or evaluative ratings were predictive of individual differences in IAT effects. The variables assessed were: (a) IAT performance, (b) processing fluency as measured by the animal and word classification trials (Tasks 1 and 2) in the IAT, (c) processing fluency as measured by a separate binary classification task, (d) visual search asymmetries, (e) self-rated stimulus pleasantness, and (f) self-rated category pleasantness. We used two processing fluency measures to compare whether a more localized measure of processing fluency in the context of the IAT would be more predictive of IAT performance than processing fluency outside the context of the IAT (in the separate binary classification task trials).

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 32 participants in Experiment 3a (cow condition) and 28 participants in Experiment 3b (elephant condition).

Stimuli and apparatus

Experiment 3a used the same stimuli as those from Experiment 1a, and Experiment 3b used the same stimuli as those from Experiment 1b.

Procedure

All participants performed six tests in each experiment, presented in a randomized order. Each participant performed the following tests:

(a) an animal/word IAT, as performed in Experiments 2a–2b; (b) an animal binary classification task, as performed in Experiments 1c–1d; (c) a word binary classification task corresponding to the same task as that described in the stimulus section of Experiments 2a–2b; (d) an animal visual search task, as performed in Experiments 1a–1b; (e) a word visual search task; and (f) an animal/word pleasantness rating task. Participants in Experiment 3a (cow condition) classified large and small cow stimuli as belonging to the categories "Zif" and "Wug". Participants in Experiment 3b (elephant condition) classified the inverted and upright elephants as belonging to the categories "live" and "dead". Tests 1–4 were the same as those presented in the previous experiments, with a few minor exceptions. In the animal/word IAT and the word binary classification task, in Tests 1 and 3, instead of classifying words and nonwords as being either "English" or "foreign", participants classifed the word and nonword stimuli into the categories "word" and "nonword", respectively. Because all participants received extended practice with the stimuli in their six tests, there was also a concern that these practice effects may dilute differences in familiarity between the familiar catgories (upright elephant, word) and unfamiliar categories (inverted elephant, nonword). In an effort to minimize this effect, the initial classification trials in the visual search task (Tests 4 and 5) were reduced from 48 to 8 (two presentations of each animal stimulus).

The word visual search task (Test 5) was similar to the animal visual search tasks of Experiments 1a and 1b. After first classifying words and nonwords for 16 trials (one presentation of each stimulus), partcipants were presented with 104 visual search trials. Each visual search array consisted of four word/nonword stimuli presented in a rectangular configuration. Participants had to judge whether all the stimuli belonged to the same category (all words, or all nonwords), or whether one of the stimuli belonged to a different category (one word among three nonwords, or one nonword among three words).

In Test 6, the animal/word pleasantness rating task, participants were presented with all of the words, nonwords, and relevant animal stimuli, and they were asked to rate how pleasant they found each item to be on a 7-point scale (1 = extremely unpleasant, 7 = extremely pleasant). As well as rating the invidual pleasantness of the stimuli, participants were asked to rate how pleasant they found each of the stimulus categories (e.g., "The category of words", "The category of Zif cows"). Each of the items (eight words, eight nonwords, two animal stimuli from each animal category, and four categories) was presented in a randomized order.

Results

A total of 7 participants were replaced in Experiment 3a, and 9 participants were replaced in Experiment 3b because they exceeded the maximum error criterion (10 errors in the IAT and 16 errors in the visual search task) in at least one of the conditions for at least one of the six tests. Another participant was replaced in Experiment 3a for having mean reaction times that were greater than 3.5 standard deviations from the group mean on the cow visual search task. Erroneous responses were omitted from the analysis (across all six tests, there were 5.2% in Experiment 3a, and 5.8% in Experiment 3b), as were those that were 3.5 standard deviations above the mean in each condition (1.1% in both Experiments 3a and 3b), or which were below the minimum outlier cut-off for the task (none in both cases). For each experiment, we first present the analysis for the IAT and pleasantness ratings task. We then present the reaction times and errors for the animal and word classification tasks of the IAT (Tasks 1 and 2), the separate binary classification tasks, and the visual search tasks. These data are followed by a multivariate correlation and a multiple regression analysis.

Experiment 3a

Cow/word IAT. An analysis was conducted with the breed of cow (Zif/Wug) assigned to the large cow category as a between-groups factor. This factor did not interact with the contrasts of interest ($F > 1$ for both the reaction time and error data).

The mean reaction time data for the cow/word IAT are presented in Figure 4. Reaction times were 55 ms faster when big cows were classified with words, and small cows were classified with nonwords, than vice versa, $F(1, 24) = 6.89$, $MSE = 13{,}753.81$, $p < .05$. Responses were equally accurate in both conditions ($F < 1$), suggesting that the IAT effect in the reaction time data was not due to a speed–accuracy tradeoff. This result replicates the cow/word IAT effect demonstrated in Experiment 2a.

Cow/word pleasantness rating task. Big cows and small cows were rated as equally pleasant at both the stimulus level (5.13 and 5.17, respectively), $t < 1$, and the category level (4.13 and 4.28, respectively), $t < 1$. Words were rated as being less pleasant than nonwords at the stimulus level (2.52 and 3.52, respectively), $t(31) = 6.07$, $p < .001$. However, at the category level, words (4.28) were considered to be more pleasant than nonwords (3.66), $t(31) = 2.70$, $p < .05$.

Cow and word binary classification and visual search tasks. The mean reaction times and errors for the binary classification and visual search tasks are presented in Table 1. As can be seen from Table 1, the direction of the means in the binary classification tasks (Tasks 1 and 2 of the IAT and the separate binary classification task) and visual search task involving cows is consistent with the findings of Experiments 1a and 1c and the fluency difference between words and nonwords outlined in Experiment 2a. The variance for each task is quite high, which may be explained by the fact that each participant performed multiple tasks featuring the same stimuli. Because these tasks were presented in a randomized order, it is likely that practice effects were introduced at different stages for each participant, producing a wide variety of response patterns.

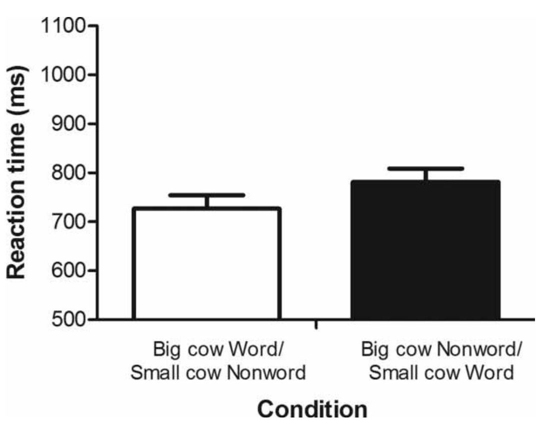

Figure 4. *Mean reaction times for the cow/word Implicit Association Test (IAT) of Experiment 3a. The open bar represents the condition in which big cows were classified with words, and small cows were classified with nonwords. The filled bar represents the condition where big cows were classified with nonwords, and small cows were classified with words. Error bars represent the standard error of the mean.*

Correlational analysis. For each participant, we calculated six variables to be entered into the correlational analysis: (a) an IAT effect score; (b) a processing fluency score from the IAT animal and word classification trials (Tasks 1 and 2 of the IAT); (c) a processing fluency score from a separate binary classification task; (d) a visual search asymmetry score; (e) a stimulus pleasantness score; and (f) a category pleasantness score.

The IAT effect score (Factor 1) was calculated by deducting the reaction times for the condition in which big cows were classified with words, and small cows were classified with nonwords, from the reaction times for the condition in which big cows were classified with nonwords, and small cows were classified with words. Thus, a positive score indicated that a participant was faster to classify big cows with words and small cows with nonwords than vice versa.

The score for Factor 2 (processing fluency) was calculated from the reaction times in the animal and word classification trials in Tasks 1 and 2 of the IAT. This score represented the fluency compatibility between the cow and word categories. The equation for this score is presented below:

$$\pm \frac{\left|\left(\left(\overline{M}_{\text{smallcowRT}} - \overline{M}_{\text{bigcowRT}}\right) - \mu\right)/\sigma\right| + \left|\left(\left(\overline{M}_{\text{nonwordRT}} - \overline{M}_{\text{wordRT}}\right) - \mu\right)/\sigma\right|}{2}$$

Table 1. *Reaction times and errors for responses in the classification and visual search tasks of Experiment 3a involving cows and words/nonwords*

	Task											
	Binary classification (IAT Tasks 1 & 2)				*Separate binary classification*				*Visual search*			
	RT (ms)		*Errors*		*RT (ms)*		*Errors*		*RT (ms)*		*Errors*	
Stimulus	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Big cows	666	383.86	1.41	1.95	492	58.26	2.47	2.29	804	147.29	2.13	1.93
Small cows	737	565.38	1.56	2.20	505	69.89	3.03	2.39	864	176.77	3.41	2.49
Words	649	145.8	0.41	0.50	652	142.47	3.38	3.60	2029	607.90	6.41	3.09
Nonwords	715	165.99	0.69	0.82	669	140.78	4.0	3.55	2036	619.20	4.22	2.93

Note: For the visual search task, the stimulus (e.g., "Big cows") refers to the stimulus that was the the odd one out. IAT = Implicit Association Test. RT = reaction time.

In this equation the sign of the score is the product of the signs of the corresponding unstandardized difference scores.

To obtain this score, we first calculated the median reaction times for responses to big cows, small cows, words, and nonwords. The median reaction time was taken instead of the mean reaction time because there were only 12 trials per condition, and the first few trials usually consisted of very long reaction times that were considered to be outliers. We then calculated a difference score for the cow classification task by subtracting the median reaction time for big cows from the median reaction time for small cows for each participant. A positive score indicates that big cows were responded to more quickly than small cows. We also calculated a difference score for the word/nonword classification task by subtracting the median reaction time for words from the median reaction time for nonwords. A positive score indicates that words were classified more quickly than nonwords. Because the variances of the two difference scores were quite discrepant, all difference scores for the two classification tasks were then standardized to z scores to ensure that the variances for each measure were comparable.

For each participant, we then took the absolute value (i.e., the numerical value without the sign) of the standardized cow difference score and the standardized word difference score and averaged these two scores to obtain an overall measure of processing fluency. The resulting score represents the magnitude of the compatibility between the cow and word categories, with a high score indicating a high degree of compatibility. To illustrate, if the fluency difference between the cow categories is large (as represented by a large standardized difference score), and the fluency difference between the word categories is also large, then the average of these two scores will also be large. The high magnitude of this score indicates that the two fluently processed categories should be highly compatible. Conversely, if the fluency difference between the two cow categories and between the two word categories is small, then the average of these two scores will also be small, reflecting low compatibility between the more fluently processed cow and word categories.

The direction of compatibility (i.e., positive/negative value) of the overall processing fluency score was determined with reference to the original (unstandardized) difference scores for cows and words. If both the cow and word (unstandardized) difference scores were in the same direction (i.e., both big cows and words were classified more quickly, yielding two positive difference scores, or both small cows and nonwords were classified more quickly, yielding two negative scores), then the sign of the corresponding overall fluency score was also made positive. However, if one (unstandardized) difference score was positive, and the other difference score was negative, then the corresponding overall fluency score was made negative. In effect, the positive/negative sign for the overall fluency score was determined by taking the product of the positive/negative signs of the coresponding difference scores. Overall then, a higher fluency compatibility score indicated greater compatibility between big cows and words (or between small cows and nonwords), and a lower score indicted greater compatibility between big cows and nonwords (or between small cows and words).

The processing fluency score for the separate binary classification task (Factor 3) was calculated in the same manner as that for Factor 2 (which used the animal and word classification trials of Tasks 1 and 2 in the IAT) except that the difference scores were calculated from the mean reaction times for each condition of the separate binary classification task, instead of the median reaction time. The mean (rather than median) reaction time was taken in this case because there was a large number of trials in the separate binary classification task (160). The large number of trials allows the mean to provide a more realiable measure of reaction times.

The scores for the other factors were calculated based on the same logic as that used to calculate the processing fluency scores. For the visual search asymmetry score (Factor 4), a difference score was obtained on the cow visual search task by subtracting each participant's mean reaction time to detect the big cow as the "odd one out" from their mean reaction time to detect small cows as the "odd one out". This would result in a positive score if the big cows were detected more quickly and a negative score if the small cows were detected more quickly. We used the same method to calculate the difference score in the word visual search task, so that a positive score would indicate that a word was more quickly detected as the "odd one out", and a negative score would indicate that a nonword was more quickly detected as the "odd one out". The visual search asymmetry score was derived in the same way as the processing fluency score (Factor 2), except that the visual search asymmetry score was based on the visual search difference scores instead of classification difference scores. If big cows and words were more quickly detected as the "odd one out" in their respective visual search asymmetry tasks, then the corresponding visual search asymmetry score would have a positive sign. In contrast, if big cows and nonwords were more quickly detected as the "odd one out", the corresponding visual search asymmetry score would have a negative sign.

To calculate the scores of stimulus pleasantness (Factor 5), we obtained difference scores on the cow/word pleasantness ratings task using a method similar to that used to calculated the processing fluency and visual search asymmetry scores. For each participant, their mean pleasantness rating for small cow stimuli was subtracted from their mean pleasantness rating of big cow stimuli to create a cow stimuli difference score. A positive score indicated that big cows were preferred to small cows. Also, participants' mean pleasantness ratings for nonword stimuli were subtracted from their mean pleasantness ratings of word stimuli to create a word stimulus difference score. Thus, a positive score indicated that word stimuli were preferred to nonword stimuli. As with the scores on the previous factors, the overall score for stimulus pleasantness was derived in the same way as the processing fluency score (Factor 2), except that the stimulus pleasantness score was based on the stimulus pleasantness difference scores instead of classification difference scores. If big cow stimuli and word stimuli were preferred, then this would produce a positive difference score. However, if big cow stimuli and nonword stimuli were preferred, then this would produce a negative difference score. A category pleasantness score (Factor 6) was calculated in the same way as the stimulus pleasantness score (Factor 5), except that instead of deriving scores from the pleasantness ratings of each stimulus, we derived the scores from each participant's ratings of category pleasantness.

The scores for the six factors just described give a measure of the compatibility between the animal and word categories on the relevant task dimension. The sign of the score indicates which categories should be compatible on the relevant dimension, and the magnitude of the score indicates the degree of compatibility that the categories share. For all factors, a positive score indicates that big cows are compatible with words, and small cows are compatible with nonwords, and a negative score indicates that big cows are compatible with nonwords, and small cows are compatible with words. It should be noted that the scores for these factors do not indicate which categories are more fluently processed (in the case of the classification tasks), more easily detected (in the case of the visual search tasks), or preferred (in the case of the pleasantness ratings). A high score may indicate that big cows and words have a higher score on the relevant dimension, or it may indicate that small cows and nonwords have a higher score on the relevant dimension. Similarly, a low score may reflect that big cows and nonwords scored more highly on the relevant dimension, or it may reflect that small cows and words scored more highly on the relevant dimension.

If we assume that a usual IAT effect implies that big cows are compatible with words, and small cows are compatible with nonwords (as found in Experiments 2a and 3a), then the scores on the other five factors that indicate the same category compatibilities will correlate positively with the IAT score. For example, a high score for processing fluency (Factors 2 and 3) would indicate that big cows are compatible with words, and small cows are compatible with nonwords in terms of fluency asymmetries. According to the fluency hypothesis, the categories that are compatible in terms of fluency asymmetries should also be compatible in the IAT, and thus a high score for processing fluency should correspond with a high score on the IAT, resulting in a positive correlation between the two factors. Similar relationships also exist between the other factor scores and the IAT effect score. If the categories that are compatible in the IAT are also compatible in terms of visual search asymmetries (Factor 4) or pleasantness (Factors 5 and 6), then these will result in a positive correlation with the IAT effect score. However, if the categories that are compatible in the IAT are incompatible in terms of visual search asymmetries or pleasantness, then these factors will correlate negatively with the IAT effect score.

Table 2 presents the correlation matrix for the six factors outlined above. None of the independent variables were significantly related to the IAT effect. However, there was a trend toward a negative relationship between the IAT effect score and the visual search asymmetry score ($r = -.34$, $p = .06$). This suggests that the more similarly two categories behave in the visual search task (a high score on Factor 4), if anything, the less compatible they are on the IAT, leading to a reduction in the IAT effect. Among the other factors, there was one significant correlation. Category pleasantness (Factor 6) was negatively related to processing fluency in the IAT animal and word classification trials (Tasks 1 and 2; $r = -.49$, $p < .01$). Thus, the animal and word categories that were compatible in terms of category pleasantness were incompatible in terms of relative fluency in the animal and word classification trials of the IAT. For example, if the category of small cows and and words were preferred, then big cows and words were processed more quickly in the animal and word classification trials of the IAT.

Multiple regression analysis. In the multiple regression analysis, scores for the IAT effect (Factor 1) were entered as the dependent variable, and scores for Factors 2–6 were entered as the independent variables. All independent variables were entered simultanously into the regression equation to predict the IAT effect. This method was adopted as there was no theoretical justification for some variables to be given greater priority than others in capturing the variance in the dependent variable. The results of the analysis are presented in Table 3, including beta weights (ß) and significance tests for each predictor. There was only one factor that was significantly related to the IAT effect score. Specifically, visual search asymmetry was the only predictor of the IAT effect, ß = −.67, $t(26) = 3.04$, $p < .01$. Therefore as two categories behaved more similarly in the visual search task, they became less compatible in the IAT. One factor that was very close to significance as a predictor of the IAT effect was stimulus pleasantness, ß = −.39, $t(26) = 2.06$, $p = .050$. This indicates that

Table 2. *Correlation matrix among the six factors of Experiment 3a involving cows and words/nonwords*

Factor		*IAT effect*	*Processing fluency (IAT)*	*Processing fluency (binary classification)*	*Visual search asymmetry*	*Stimuli pleasantness*	*Category pleasantness*
Factor 1	IAT effect	1.000					
Factor 2	Processing fluency (IAT)	−.104	1.000				
Factor 3	Processing fluency (binary classification)	.019	.132	1.000			
Factor 4	Visual search asymmetry	−.337	−.249	.298	1.000		
Factor 5	Stimuli pleasantness	−.223	.093	.097	−.217	1.000	
Factor 6	Category pleasantness	.034	−.487	−.068	−.054	−.038	1.000

Note: IAT = Implicit Association Test. Processing fluency (IAT): processing fluency in the separate target and classification trials of the IAT. Processing fluency (binary classification): processing fluency in the separate binary classification task.

Table 3. *Experiment 3a: Summary of the multiple regression analysis on IAT effect with processing fluency, visual search asymmetry, stimuli pleasantness, and category pleasantness as predictors*

	Variable	*ß*	*t* value	*Sig.*
Factor 2	Processing fluency (IAT)	− .404	− 1.593	.123
Factor 3	Processing fluency (binary classification)	.288	1.439	.162
Factor 4	Visual search asymmetry	− .673	− 3.041	.005
Factor 5	Stimuli pleasantness	− .388	− 2.055	.050
Factor 6	Category pleasantness	− .171	− .782	.441

Note. IAT = Implicit Association Test. The beta weights represent the coefficients of each variable in the regression equation.

stronger IAT effects were obtained when the categories compatible in the IAT were incongruent in terms of their stimulus pleasantness.

Experiment 3b

Elephant/word IAT. The mean reaction times for the elephant/word IAT are shown in Figure 5. Responses were 285 ms faster when upright elephants were classified with words, and inverted elephants were classified with nonwords, than vice versa, $F(1, 27) = 69.97$, $MSE = 32{,}470.86$, $p < .001$. Slightly more errors were committed when upright elephants shared a key with nonwords, and inverted elephants shared a key with words ($M_{\text{error}} = 1.89$) than when the animal and word categories were reversed ($M_{\text{error}} = 1.36$); however, this effect did not reach significance, $F(1, 27) = 3.78$, $MSE = 2.128$, $p = .06$. Nonetheless, this result does suggest that there was no speed–accuracy trade-off in the pattern of reaction times.

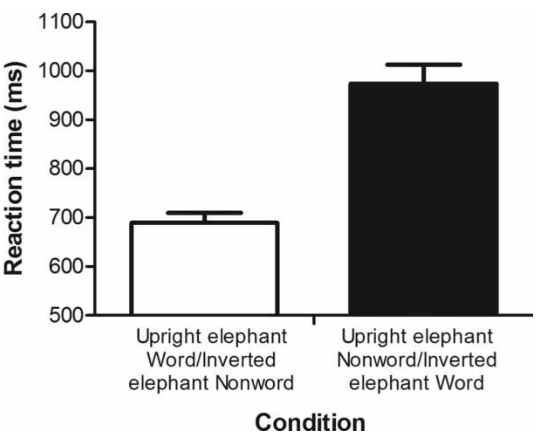

Figure 5. *Mean reaction times for the elephant/word IAT of Experiment 2b. The open bar represents the condition in which upright elephants were classified with words, and inverted elephants were classified with nonwords. The filled bar represents the condition where upright elephants were classified with nonwords, and inverted elephants were classified with words. Error bars represent the standard error of the mean.*

Elephant/word pleasantness rating task. Live elephants were rated as being more pleasant than dead elephants at both the stimulus level (5.36 and 3.09, respectively), $t(27) = 6.39$, $p < .001$, and the category level (5.36 and 2.50, respectively), $t(27) = 6.61$, $p < .001$. Word stimuli were rated to be less pleasant than nonword stimuli (2.61 and 3.51, respectively), $t(27) = 3.55$, $p < .01$. However, the category of words was considered to be more pleasant than the category of nonwords (4.64 and 3.43, respectively), $t(27) = 4.36$, $p < .001$.

Elephant and word binary classification and visual search tasks. The mean reaction times and errors for the binary classification tasks (Task 1 of the IAT and the separate binary classification task) and visual search tasks are in Table 4. It is surprising to note that, in contrast to the results of the separate binary classification task in Experiment 1d, inverted elephants were classified more quickly than upright elephants in the binary classification task of the IAT (Task 1), $t(27) = 2.96$, $p < .01$, and upright and inverted elephants were classified equally quickly in the separate binary classification task, $t < 1$. These results may have occurred because most participants were preexposed to the elephant stimuli before completing the binary classification task of the IAT (Task 1) and the separate binary classification task. Compared to Experiment 1d, preexposure to the elephant stimuli in Experiment 3d is likely to decrease the relative difference in familiarity between upright and inverted elephants.

Table 4. *Reaction times and errors for responses in the classification and visual search tasks of Experiment 3b involving elephants and words/nonwords*

		Task											
		Binary classification (IAT Tasks 1 & 2)				*Separate binary classification*				*Visual search*			
		RT (ms)		*Errors*		*RT (ms)*		*Errors*		*RT (ms)*		*Errors*	
Stimulus		*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Elephants	Upright	576	164.75	0.89	1.45	524	80.24	3.82	2.72	980	149.88	3.32	2.80
	Inverted	514	118.90	0.68	1.49	524	79.40	3.14	2.10	950	173.20	2.65	1.75
Words		633	106.87	0.39	0.88	618	87.59	2.64	1.95	1,860	447.22	6.61	3.08
Nonwords		712	173.03	0.57	0.88	640	77.21	3.18	2.18	1,846	350.54	3.93	1.86

Note: For the visual search task, the stimulus (e.g., "Upright elephants") refers to the stimulus that was the the odd one out. IAT = Implicit Association Test. RT = reaction time.

Table 5. *Correlation matrix among the six factors of Experiment 3b involving elephants and words/nonwords*

Factor		*IAT effect*	*Processing fluency (IAT)*	*Processing fluency (binary classification)*	*Visual search asymmetry*	*Stimuli pleasantness*	*Category pleasantness*
Factor 1	IAT effect	1.000					
Factor 2	Processing fluency (IAT)	.509	1.000				
Factor 3	Processing fluency (binary classification)	.138	.368	1.000			
Factor 4	Visual search asymmetry	.087	− .071	− .239	1.000		
Factor 5	Stimuli pleasantness	− .249	− .088	− .039	− .147	1.000	
Factor 6	Category pleasantness	− .108	− .258	.005	− .150	.154	1.000

Note: IAT = Implicit Association Test. Processing fluency (IAT): processing fluency in the separate target and classification trials of the IAT. Processing fluency (binary classification): processing fluency in the separate binary classification task.

These conditions are likely to make the binary classification tasks less sensitive to actual differences in processing fluency between upright and inverted elephants.

Correlational analysis. Six variables were entered into the correlational analysis: (a) an IAT effect score, (b) a processing fluency score for the separate target and attribute classification trials in the IAT, (c) a processing fluency score for the separate binary classification task, (d) a visual search asymmetry score, (e) a stimulus pleasantness score, and (f) a category pleasantness score. These scores were derived using the same methods as those used to calculate the same scores in Experiment 3a, except that upright and inverted elephants replaced big and small cows, respectively. The correlation matrix for the six factors are presented in Table 5. The only factor that showed a significant correlation with the IAT effect score was processing fluency on the separate target and attribute classification trials of the IAT ($r = +.51$, $p < .01$). This indicates that as upright elephants and words (or inverted elephants and nonwords) increased in processing fluency, the magnitude of the IAT effect also increased. There were no significant relations among the other factors, but the correlation between processing fluency on the IAT separate target and attribute classification tasks and processing fluency on the binary classification task closely approached significance ($r = +.37$, $p = .054$). This suggests that as people became faster at processing upright elephants and words (or inverted elephants and nonwords) in the IAT target and classification task, they were also faster to process the same categories in the binary classification task.

Multiple regression analysis. We conducted a simultanous multiple regression analysis (see the same analysis for Experiment 3a). Scores for the IAT effect were entered as the dependent variable, and scores for the remaining five factors were entered as the independent variables. Table 6 shows the results of the analysis, including beta weights (ß) and significance tests for each predictor. Only one independent variable was significantly related to the IAT effect. Processing fluency in the separate target and

Table 6. *Experiment 3b: Summary of the multiple regression analysis on IAT effect with processing fluency, visual search asymmetry, stimuli pleasantness, and category pleasantness as predictors*

	Variable	*ß*	*t* value	*Sig.*
Factor 2	Processing fluency (IAT)	.590	2.697	.013
Factor 3	Processing fluency (binary classification)	− .055	− .221	.827
Factor 4	Visual search asymmetry	.103	.519	.609
Factor 5	Stimuli pleasantness	− .243	− 1.114	.277
Factor 6	Category pleasantness	.105	.399	.694

Note: IAT = Implicit Association Test. The beta weights represent the coefficients of each variable in the regression equation.

attribute classification trials of the IAT predicted the IAT effect, $t(22) = 2.70$, $p < .05$. This indicates that the faster people were to classify one of the animal and one of the word categories in the separate animal and word classification tasks of the IAT, the more easily these two categories were classified together in the combined classification tasks of the IAT. Therefore, compatibility in fluency predicted compatibility in the IAT.

Discussion

In Experiments 3a and 3b, we replicated the animal/word IAT effects using "word" and "nonword" labels. This result supports the robustness of the animal/word IAT effects by showing that these effects generalize to different attribute category labels. The cow/word IAT effect in Experiment 3a is unlikely to be caused by affective associations between the animal and word categories, because big and small cows were rated as equally pleasant at both the stimulus and category levels. The cow/word IAT effect also cannot be attributed to visual search asymmetries, as it appears that visual search asymmetries negatively predict the direction of IAT effects. That is, the big cow and nonword categories that are more quickly detected in a visual search task are not compatible in the IAT. This effect is consistent with the results of Experiments 2a. Therefore, the results of Experiment 3a provide some support for the fluency account of salience asymmetries in the cow/word IAT effect, because they rule out alternative explanations for the effect.

In Experiment 3b with elephant stimuli, the pleasantess ratings reveal that, at the category level, upright elephants and words were both preferred to inverted elephants and nonwords, respectively. This finding suggests that affective associations are responsible for the elephant/word IAT effect, because the two preferred categories (upright elephants and words) were compatible in the IAT. However, there was no evidence that individual variations in evaluative associations corresponded with individual variations in IAT effects. Nor was there a reliable relationship between visual search asymmetry and the elephant/word IAT effect.

There was evidence that processing fluency (as measured by Tasks 1 and 2 in the IAT) predicts the elephant/word IAT effect. This provides some evidence for the hypothesis that fluency asymmetries on a classification task correspond with IAT effects. However, other measures of processing fluency did not correlate with IAT effects in Experiments 3a and 3b. The discrepancy in correlations between the two measures of processing fluency in Experiment 3b may have occurred because the animal and word classification tasks in the IAT (Tasks 1 and 2) provide a more localized index of processing fluency that is more pertinent to the IAT than is that of the separate binary classification task. For example, a participant who performs the IAT as the first task is likely to show larger fluency asymmetries between the familiar and unfamiliar categories than a participant who performs the IAT as the sixth task. Thus, fluency differences measured close to the time of IAT performance are likely be more closely associated with compatibility effects in the IAT than are fluency differences measured on a separate binary classification task that is performed at a more distal point in time.

In contrast to Experiment 1d, inverted elephants were classified more quickly than upright elephants in the animal classification task (Task 1) of the IAT. This reversal in fluency asymmetry may have resulted because, in Experiment 3b, most participants had already processed the upright and inverted elephant stimuli on previous tasks before they performed the animal classification task of the IAT. Therefore, the participants in Experiment 3b would be more familiar with the inverted elephant stimuli than would the participants in Experiment 1d. Once the participants were familiarized with the elephant stimuli, some aspect of the stimuli other than familiarity may become more salient and produce faster responding to the dead elephants. For example, if participants were to focus on discriminating between live and dead elephants based simply on their physical features, it may be that the upward extension of a dead elephant's legs is easier to discern than the downward extension of a live elephant's legs.

The results of Experiment 3b involving elephant stimuli are less critical in distinguishing between the different accounts of IAT effects than are those of Experiment 3a involving cow stimuli. This is because the compatibility of upright elephants and words is predicted by all three accounts of IAT effects (the fluency hypothesis of salience asymmetries; Rothermund and Wentura's (2001, 2004) account of salience asymmetries that correspond with visual search asymmetries; and an evaluative association account). In contrast, the compatibility of big cows and words found in Experiment 3a is predicted by only the fluency account of salience asymmetries in the IAT.

GENERAL DISCUSSION

Experiments 1a–1d established that the stimuli in our two category pairings varied independently along the dimensions of visual search asymmetry and processing fluency. In Experiment 1a, big cows were more quickly detected among small cows than small cows were among big cows, and participants were faster to categorize big cows than small cows on a binary classification task (Experiment 1c). In Experiment 1b, search was quicker 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 (Experiment 1d). Thus in one pair of categories (big/small cows), the category that was detected more quickly in visual search was the more fluently processed category, whereas in the other pair of categories (upright/inverted elephants), the category that was more quickly detected in visual search was the less fluently processed category in the binary classification task.

Experiments 2a and 2b used the same cow and elephant stimuli to reveal that the category that was more fluently processed in the binary classification task (big cows and upright elephants) was also more easily classified with words over nonwords in the IAT. These data support the idea that compatibility effects in the IAT may be due to salience asymmetries based on differences in processing fluency.

Experiments 3a and 3b replicated the results of Experiments 2a and 2b in showing that the more fluently processed animal category was compatible with words over nonwords in the IAT. These studies also examined whether processing fluency, visual search asymmetries, or affective associations could predict individual differences in the the IAT effects seen in Experiments 2a and 2b. There was some support for the fluency asymmetry hypothesis in Experiment 3b, with a significant positive correlation between IAT effects and processing fluency in the target and attribute classification tasks of the IAT. There was also evidence against the visual search asymmetry account in Experiment 3a, with a negative relationship between visual search asymmetries and IAT effects. In general, however, individual variation in processing fluency, visual search asymmetry, and evaluative associations did not reliably predict individual variations in IAT effects.

Taken together, Experiments 3a and 3b lend support to the idea that compatibility effects in the IAT can sometimes be due to salience asymmetries rather than associations. As mentioned previously, it would be very difficult to argue that the cow/word IAT effect is the consequence of any attitude held by the participants towards the target stimuli in these experiments, because big and small cows were equally preferred on an explicit measure of attitudes, and there is no reason for participants to conceal their attitudes towards these stimuli. However, the cow/word IAT effect is contrary to Rothermund and Wentura's (2001, 2004) assertion that search asymmetries between categories correspond with salience asymmetries in the IAT, as the two categories that were more easily detected in visual search (big cows and nonwords) were not compatible in the IAT. Instead, fluency assymetries between category pairs may be a better predictor of salience asymmetry effects, because the two categories that were compatible in the IAT (big cows and words), were generally the more fluently processed categories of their respective pairings.

Our results clarify the salience asymmetry account of Rothermund and Wentura (2001, 2004) by demonstrating that processing fluency, rather than visual search asymmetries, is diagnostic of salience asymmetries in the IAT. In fact, it can be seen that differences in category processing fluency can also account for Rothermund and Wentura's (2004) own findings that items that are detected faster in visual search tend to be compatible in the IAT. In their Experiment 1c using a visual search task, participants were faster to detect the name of an old person among the names of young people, a nonword among words, and a multicoloured string among single-coloured strings than they were when the category assignments were reversed. In other experiments (Rothermund & Wentura, 2004; Experiments 1a and 1e), participants classified old and young names with each of the other category pairs in an IAT. It was found that old names were compatible with other categories that were also more quickly detected as being "different" in the visual search task (i.e., nonwords and multicoloured strings). Rothermund and Wentura (2004) 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. Old names and nonwords are less familiar than young names and words, thus, 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, multicoloured strings are likely to take longer to process than single-coloured strings, because multicoloured strings are more complex and do not form a unified perceptual configuration. Thus, multicoloured 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.

When Rothermund and Wentura (2001, 2004) claimed that visual search asymmetries coincide with salience asymmetries in the IAT, they conceived of search asymmetries as being due to differences in attentional disengagement between the two categories. However, there appear to be multiple mechanisms responsible for visual search asymmetries, which may explain why search asymmetries produce inconsistent compatibility effects in the IAT. First, search asymmetries can be due to the ease with which distractors are rejected (e.g., Strayer & Johnston, 2000; Treisman & Souther, 1985; Wang et al., 1994). Evidence shows that targets are easier to detect when distractors are familiar (Rauschenberger & Yantis, 2006; Shen & Reingold, 2001) or prototypical (Levin & Angelone, 2001). This situation is illustrated by Experiments 1b and 1d involving the elephant stimuli. In the visual search task of Experiment 1b, 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 1d).

Secondly, however, the characteristics of the target are a determinant in other types of search asymmetry. Treisman and Gelade (1980) point out that when the target can be distinguished from the distractor by the presence of a feature, participants engage in "feature search", which allows the target to "pop out" from the background. In these circumstances, efficient search is caused by the target being more easily identifiable, and thus it would also be more fluently processed in a binary classification task. Feature search is likely to have characterized the visual search experiment involving cow stimuli (Experiment 1a), as 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. In support of this, the big cows that were detected more quickly than small cows in visual search (Experiment 1a) were also identified more rapidly in the classification task (Experiment 1c).

The idea that the feature-present category is more fluently processed in a binary classification task has implications for the race IAT. Levin (1996) demonstrated that white people are faster to classify faces from the unfamiliar black category than the familiar white category (see Valentine & Endo, 1992, for a similar effect involving Asian faces). Moreover, visual search tasks reveal that white people are quicker to detect a black face among white faces than vice versa (Levin, 2000). Levin (1996, 2000) explained that these effects occur 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. This means that in the race IAT, "blackness" is more salient than "whiteness", because racial category membership is more apparent with black faces. In this case, the unfamiliar category of black is more fluently classified than the familiar category of white.

This difference in processing fluency between black and white faces suggests that in the race IAT, black faces may be compatible with the more fluently processed positive attribute category due to salience asymmetries. Of course, in terms of associations, white faces are thought to be compatible with positive attributes (and black faces with negative attributes) when white people are tested in the IAT (Greenwald et al., 1998). Thus, if black faces and pleasant attributes are indeed compatible on a nonassociative (processing fluency) basis, this would serve to attenuate the IAT effect for white. That is, the race IAT effect would be an underestimate of the true attitude. Given the controversy surrounding the race IAT, it would be helpful to investigate whether salience asymmetries between black and white target categories have the potential to increase or decrease the race IAT effect. Our laboratory is currently investigating this issue.

The experiments reported here provide support for Rothermund and Wentura's (2001, 2004) proposal that the IAT may be influenced by salience asymmetries. However, contrary to their theory, it appears that salience asymmetries in the IAT correspond with processing fluency, not visual search asymmetries. This has some important implications for social psychologists using the IAT who wish to control for nonassociative factors. By measuring the speed and accuracy with which target categories are classified, it would be possible to obtain at least a rough measure of the degree to which salience asymmetries contribute to particular IAT effects.

Original manuscript received 7 October 2008
Accepted revision received 24 November 2008
First published online 26 February 2009

REFERENCES

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.

Brendl, M., Markman, A., & Messner, C. (2001). How do indirect measures of salience asymmetries in the IAT evaluation work? Evaluating the inference of prejudice in the Implicit Association Test. *Journal of Personality and Social Psychology*, *81*, 760–773.

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.

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.

De Houwer, J. (2001). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology*, *37*, 443–451.

Estes, Z., & Adelman, J. S. (2008). Automatic vigilance for negative words in lexical decision and naming: Comment on Larsen, Mercer, and Balota (2006). *Emotion, 8*, 441–444.

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.

Govan, C. L., & Williams, K. D. (2004). Changing the affective valence of the stimulus items influences the IAT by re-defining the category labels. *Journal of Experimental Social Psychology*, *40*, 357–365.

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

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.

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.

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.

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.

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 revisted: A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology*, *8*, 381–396.

Rauschenberger, R., & Yantis, S. (2006). Perceptual encoding efficiency in visual search. *Journal of Experimental Psychology: General, 135*, 116–131.

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.

Shen, J., & Reingold, E. M. (2001). Visual search asymmetry: The influence of stimulus familiarity and low-level features. *Perception & Psychophysics*, *63*, 464–475.

Strayer, D. L., & Johnston, W. A. (2000). Novel popout is an attention-based phenomenon: An ERP analysis. *Perception & Psychophysics*, *62*, 459–470.

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

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. (2008). Why positive information is processed faster: The density hypothesis. *Journal of Personality & Social Psychology, 95*, 36–49.

Valentine, T., & Endo, M. (1992). Toward an exemplar model of face processing: The effects of race and distinctiveness. *Quarterly Journal of 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 Behaviour, 17*, 143–154.

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.

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