Zeitschrift für Experimentelle Psychologie, 48, 2, 107–122 © Hogrefe-Verlag, Göttingen 2001

Implicit Association Measurement with the IAT: Evidence for Effects of Executive Control Processes

Jan Mierke and Karl Christoph Klauer

Rheinische Friedrich-Wilhelms-Universität Bonn

Abstract. It is argued that a model of goal-independent spreading activation in a social or semantic knowledge structure is insufficient to explain implicit association effects in the IAT (Greenwald, McGhee, & Schwartz, 1998). An alternative account is proposed, which attributes IAT effects to differential costs for switching between task sets. Two experiments were conduced to test this account. In Experiment 1, specific task-set switching cost was a function of IAT condition: switching between tasks was associated with significantly more cost in the incompatible IAT phase. In a second experiment the magnitude of the IAT effect was reduced when task-set reconfiguration was possible in advance of or simultaneously with the upcoming stimulus. The results are discussed with respect to recently suggested accounts of the effect.
Key words: Implicit measurement, IAT, task switching, executive control

Implizite Assoziationsmessung mit dem IAT: Evidenz für Effekte exekutiver Kontrollprozesse

Zusammenfassung: Es wird argumentiert, dass ein Modell zielunabhängiger Aktivierungsausbreitung in einem sozialen oder semantischen Netzwerk nicht ausreicht, um Assoziationseffekte im IAT (Greenwald, McGhee, & Schwartz, 1998) zu erklären. Ein alternativer Ansatz wird vorgeschlagen, der IAT Effekte auf bedingungsanhängige Aufgabenwechselkosten zurückführt. Zwei Experimente prüfen diesen Ansatz. Im ersten Experiment zeigte sich, dass spezifische Aufgabenwechselkosten eine Funktion der IAT-Bedingung sind: Aufgabenwechsel waren in der inkompatiblen IAT-Bedingung mit bedeutsam höheren Kosten verbunden. In einem zweiten Experiment war das Ausmaß des IAT-Effekts reduziert, wenn die Möglichkeit zu vorbereitender oder simultaner Einstellung auf den Aufgabenwechsel gegeben wurde. Die Ergebnisse werden im Hinblick auf aktuelle Erklärungsmodelle des Effekts diskutiert.
Schlüsselwörter: Implizites Messen, IAT, Aufgabenwechsel, exekutive Kontrolle

Since its introduction in 1998, the Implicit Association Test (Greenwald et al., 1998) has been discussed controversially, mainly due to a lack of a theoretical fundament. The purpose of the present study is to gain further insight into the mechanisms that underlie this effect. We aim at a closer understanding of *how* conceptual relations between concepts affect IAT performance. It is argued that an account based on conceptual spreading activation alone is not sufficient to explain the effect. We will furthermore discuss alternative accounts that have been proposed recently and finally suggest a model that attributes the IAT effect to differences in cost related to task-set switching (Allport, Styles, & Hsieh, 1994; Rogers & Monsell, 1995).

The authors wish to thank Katja Ehrenberg, Rainer Banse, Klaus Rothermund, Jan De Houwer, and two independent reviewers for valuable comments on an earlier version of this article. We thank Tom Swinehart for his assistance in improving the English. Portions of the data appearing here were presented at the 1999 meeting of the Division of Social Psychology, Kassel.

Basic Design of the IAT

The IAT is typically composed of two independent binary choice tasks that have to be performed in an alternating or random sequence. Each task demands the discrimination of two conceptually defined sets of stimuli. One of the tasks requires the discrimination of two target categories or concepts, the other judgments about an attribute. Each of the four response categories is instantiated by a distinct set of typical stimuli.

Two types of blocks are constructed, differing in the response assignment for the target discrimination task. A block is termed *compatible*, if the "instructions oblige highly associated categories (e.g., flower + pleasant) to share a response key", and *incompatible* "when less associated categories (e.g., insect + pleasant) share a key" (Greenwald et al., 1998, p. 1464). The typical finding is that performance is better in compatible than in incompatible blocks. These experimental conditions are preceded by a training phase in which the attribute discrimination task is practiced, and two training phases for the target discrimination task, in which the response mapping for the upcoming alternation blocks is trained (see Greenwald et al. 1998, for details).

Can Spreading Activation Account for the IAT Effect?

Spreading activation in semantic (Collins & Loftus, 1975; Collins & Quillian, 1969) or social network structures (Greenwald et al., 2000) is a powerful theoretical concept that is capable of explaining a host of different phenomena such as semantic and associative priming effects in a variety of paradigms. The essential mechanism is a distribution or spreading of activation between interconnected nodes that represent conceptual knowledge. The spreading activation metaphor is capable of representing a large variety of mental structures or processes in the form of nodes and interconnections.

Spreading activation accounts have been tested by means of semantic and associative priming paradigms (see Neely, 1991, for a review) as well as by appropriate variants of the affective priming paradigm (Bargh, Chaiken, Raymond, & Hymes, 1996; Klauer & Musch, in press). A compelling difference between the IAT and these paradigms has to be noted: The compatible and incompatible IAT conditions do not differ with respect to stimulus composition. This fact, although it may seem trivial, amounts to the conclusion that exclusively stimulus-triggered spread of activation should lead to identical activation patterns in both conditions, and thus cannot by itself explain the IAT effect. After considering some general issues concerning the interpretation of IAT effects in the following section, we will portray some alternative accounts that have been proposed recently.

Category vs. Stimulus Feature Basis of the IAT

De Houwer (in press) and Neumann, Totzke, Popp, & Fernandez (2000) provide empirical evidence for the generally held claim that the IAT effect reflects relations between properties of the *target categories* and the attribute, but not between properties of *individual stimuli* and the attribute.

Neumann et al. (2000) report an attitude-IAT effect with pictures of Caucasians vs. black persons, if and only if the labels used as response categories were meaningful to the participants. The effect occurred, if the pictures had to be classified as "Germans" vs. "Foreigners", but was absent if the pictures had to be categorized as "white persons" vs. "black persons", the latter being a less meaningful dichotomy to German participants. This indicates that the IAT effect is based on properties of the target categories.

De Houwer (in press) disentangled category evaluation and stimulus evaluation by using target stimuli that are evaluatively incongruent to the category they instantiate in half of the cases and evaluatively congruent in the other half. If the IAT effect is indeed based on target-category evaluation, a compatibility effect should emerge with respect to category valence, but not stimulus valence. In fact, contrasting famous (positive and negative) British citizens with famous (positive and negative) citizens from foreign countries produced an IAT effect based on category evaluation, but no effect of stimulus evaluation. Steffens & Plewe (2001), however, report effects at stimulus level. In their study, the valence of stimuli exemplifying the attribute dimension was confounded with the target category. The manipulated direction of stimulus-level association was clearly reflected in the IAT measure.

A plausible explanation for these findings might be that stimuli are perceived in terms of their category membership, as soon as they are processed with respect to a strong social category. Research on person categorization in the "Who said what?" paradigm, for example, demonstrates that memory for the behavior of persons in a group discussion is frequently encoded with respect to group membership, almost ignoring individual stimulus features

(Klauer & Wegener, 1998). Furthermore, there is reliable evidence from cross-categorization research, indicating that some social categories are applied more readily than others (Ehrenberg, Klauer, & Wegener, 2000; Stangor, Lynch, Duan, & Glass, 1992). Similar categorization processes may occur in the IAT, if the target categories are instantiated by stereotypical names (e.g. Greenwald et al., 1998; Neumann et al., 1998).

Figure-Ground Asymmetry Model by Rothermund and Wentura (2001)

The model proposed by Rothermund and Wentura (2001) is based on the assumption that, for both discrimination tasks, participants perceive one response category as a figure on the ground of the opposing response category. The performance difference between differently mapped IAT conditions is accounted for by the assumption that participants can base their responses on figure-ground information alone in the compatible condition, where "figure stimuli" are consistently mapped to one response key and "ground stimuli" to the other. See Rothermund & Wentura (2001) for details.

Acquired Meaning of Response Keys (De Houwer, in press)

The model proposed by De Houwer (in press) is based on the assumption that response keys become temporarily associated with evaluative meaning based on the task instructions[1]. It is argued that the meaning of responses in the compatible condition of the IAT is less ambiguous with respect to valence than in the incompatible condition, because stimulus-response pairs are more homogenous in this respect. The IAT effect itself is elegantly explained by a stimulus-response compatibility mechanism, similar to those proposed to underlie the Simon effect (Simon, 1990; see Kornblum, 1992, for a review): Higher similarity of stimulus features and response features are assumed to underlie faster responses in the compatible condition.

Random Walk Model by Brendl, Markman, and Messner (2000)

Brendl et al. (2000) account for the IAT effect by combining a simple random walk model of information processing with the assumption of a criterion shift in the incompatible IAT condition. The authors assume that incoming attribute and identity information, i.e. information on the target category of the presented stimulus, is accumulated in a random walk process on a response-related decision dimension. The net accumulation rate of evidence for stimuli from the target categories should be lower in the incompatible condition, as information concerning the identity of a stimulus and information concerning the attribute push the counter towards different responses here.

Brendl et al. (2000) further assume that because of the slower net accumulation-rate the actual and perceived difficulty of the incompatible IAT condition is higher, leading participants to employ a more conservative response criterion in incompatible IAT blocks. This should result in a longer accumulation interval for both target and attribute stimuli in the incompatible condition[2].

While the properties of the random walk model were not explicitly tested, a hypothesis derived from the assumption of a shift in response criterion was confirmed in three experiments, i.e. that responses to instances of the attribute concepts should be slower in the incompatible than in the compatible IAT condition. A random-walk model without criterion shift does not predict this, as the lower net accumulation rate in the incompatible IAT condition is restricted to target-concept stimuli.

A Task-Set Switching Account of the IAT Effect

Imagine a Stroop experiment (Stroop, 1935; see MacLeod, 1991, for a review) in which you have to judge a stimulus composed of a letter and a digit. On some of the trials you are required to judge whether the digit is odd or even, while on other trials, you have to decide whether the letter is a vowel or a consonant. Now imagine there were blocks in which digit and letter consistently trigger the same response (the co-occurring stimulus aspects are mapped to the same response key) and other blocks, in which digit and letter always require different responses (co-oc-

[1] See also Neumann et al. (1998) who assume that response-keys acquire a meaning based on a learning process.

[2] Although not explicitly mentioned by the authors, this should also lead to a higher contrast between attribute and target discrimination in this condition.

curring stimulus aspects are mapped to different response keys). Presumably, performance would be better when the co-occurring stimulus aspects trigger the same response, i.e. when performing accurately does not depend on switching between the tasks as instructed.

The analogy between the experiment described above and the IAT is obvious: The irrelevant attribute information of instances of the target categories triggers the same response as does their categorical identity in the compatible condition. In the incompatible condition, however, different responses are mapped to these co-occurring stimulus aspects.

To explain IAT effects it is necessary to analyze how the described design properties influence performance. The model we propose assumes that switching between different task sets might provide this causal link. A *task set* is assumed to be a complex of numerous settings, required for performing a given task. These settings include "which attribute of the stimulus to attend to, which response mode and value to get ready, what classification of the relevant stimulus attribute to perform, how to map those classes to response values, with what degree of caution to set one's criterion for response etc." (Monsell, Yeung, & Azuma, 2000, p. 252). Processes of task-set switching are generally attributed to a combination of endogenously initiated and exogenously triggered processes (Meiran, 1996; Rogers & Monsell, 1995; Monsell et al. 2000), and are typically associated with performance cost.

The model we propose may be called a probabilistic task-switch-neglect model. The model states that participants neglect the instruction to switch between task sets in the compatible IAT condition on a substantial proportion of experimental trials. Neglecting task switches is possible in this condition, because basing the responses on attribute-related information alone allows fast and accurate responding. As there is no necessity to switch between tasks, switching costs can be evaded. The IAT effect, according to this model, reflects general costs for switching between two different task sets that specifically affect the incompatible IAT condition.

An ongoing debate in the literature on task switching concerns the question whether the performance cost typically found in task switching experiments directly reflect the operation of a stage-like executive control function that is responsible for implementing an appropriate task set on task-switch trials. Allport et al. (1994; cf. Wylie & Allport, 2000) argue that task-switching costs might be due to proactive interference from the preceding trial. In this model, task-switching costs do not reflect the operation of an extra process specific to task-switch trials, but rather differential time costs for the same set of processes for switch vs. non-switch trials. This claim does neither challenge nor replace the assumption that control processes take place in this type of experiments (Monsell et al., 2000). It questions whether task-switching costs reflect the *duration* of these control processes.

Although the model we proposed above is compatible with both accounts of task-switching costs, we briefly discuss how a proactive task-set interference model might explain IAT effects. The major difference between such a model and the probabilistic task-switch-neglect model is that task switches are assumed to occur in both IAT conditions. Task-set switching, however, is associated with less cost in the compatible condition. The rationale for this is as follows: In order to perform a given task appropriately, response-incongruent task-sets have to be inhibited. Besides the cost directly associated with response incongruence on the current trial, this should result in an increased difficulty to retrieve the inhibited task set on an upcoming trial, i.e. a negative priming of task set (Wylie & Allport, 2000), adding to performance costs in the incompatible IAT phase. Note that task-switch trials in the incompatible condition are preceded either by neutral stimuli without an irrelevant feature (attribute-set stimuli) or by stimuli with a response-incongruent irrelevant feature (target-set stimuli). In the compatible IAT condition, on the contrary, all trials are response congruent or neutral and thus are preceded by response-congruent or neutral trials, i.e. the antecedent for negative priming is not given. We will discuss these issues in more detail in the general discussion.

Some non-trivial predictions may be derived from both mechanisms. First, besides general performance cost for alternating between two tasks on a blocked scale, an additional component of switching task set should emerge on a sequential level, i.e. a performance difference between trials that are preceded by the same vs. a different task (Meiran, 1996; Rogers & Monsell, 1995). This sequential effect will be referred to as *specific* task-set switching cost (Kray & Lindenberger, 2000). If the task-switching account is correct, specific cost should be more pronounced in the incompatible IAT phase.

Distinguishing between specific and general task-switching costs empirically requires the incorporation of task-repetition trials in the design. Task-repetition trials emerge if the order of stimuli in a block is completely randomized. Some authors have used this completely randomized procedure (e.g. Kühnen et al., 2001; Rothermund & Wentura, 2001), while others used strictly alternating task sequences (e.g. Greenwald et al, 1998). Banse, Seise, & Zerbes

(2001) compared both types of stimulus ordering and found little difference in the IAT effects.

A second prediction concerns effects of response repetition. Rogers and Monsell (1995) repeatedly found that a facilitative effect of response repetition was disrupted if the presently required task switched from the previous trial, while it occurred when the task was repeated. Although the explanation for this effect is still under debate, the effect itself seems to be a reliable property of task-set switching (Meiran, 1996; Rogers & Monsell, 1995). Based on the account described above, we expect effects of response repetition for all trials in the compatible condition, but for non-switch trials only in the incompatible condition. Note that neither of these predictions is expected on the basis of the random-walk model with or without criterion shift (Brendl, et al., in press), which predicts response-repetition effects on all trials.

Experiment 1

Experiment 1 was designed to test some of the predictions derived above. The Hypotheses for Experiment 1 were

H1) The IAT effect is replicable with German material selected on the basis of word norms.

H2) Specific costs for task-set switching are more pronounced in the incompatible condition of the IAT (performance on trials preceded by the same vs. other task interacts with IAT condition).

H3) Effects of response repetition are disrupted on task-switch trials in the incompatible IAT condition, but occur irrespective of task switches in the compatible condition (interaction of IAT condition, task switching and response repetition).

Trial sequences were coded with respect to the repetition of features from the directly preceding trial (Pashler & Baylis, 1991). One repetition factor concerned the repetition of responses, i.e. whether the current and the preceding trial required the same response. The second sequential factor concerned the discrimination task to be performed. If moving from the preceding to the actual trial required a switch in the discrimination task (an evaluative decision followed by a flower-insect discrimination or vice versa) the trial was coded as a task-switch trial. If both trials of a sequence required the same type of decision, the trial was coded as a task-repetition trial. The relations between sequence factors are depicted in Table 1.

The coding of trial sequences results in two orthogonal within-participants factors. These factors are orthogonal to other factors manipulated in the IAT. The response-stimulus interval (RSI) was varied between participants (100 ms vs. 1000 ms) to test whether the specific sequence effects we expected would generalize over different time constraints. An interaction of specific task-switching cost with the RSI would indicate a dissipation of carryover. It cannot, however, be attributed to an intentional switch of task set, which is initiated before stimulus presentation, as the task-set appropriate for the upcoming trial is unpredictable.

Method

Following Greenwald et al. (1998, Experiment 1), we used insect and flower names in conjunction with positive and negative words. The experiment is based on a 2 × 2 × 2 design with the within-participants factor mapping compatibility, and the between-participants factors response-stimulus interval (100 ms vs. 1000 ms), and order of mapping conditions (compatible vs. incompatible mapping-condition first). A between-participants counterbalancing-factor permuted response assignments to control for any unwanted systematic variance caused by differences in the treatment of response keys. Furthermore, two within-participants factors emerged as a product of the random sampling of stimuli. Unless noted otherwise, design and procedure follow Greenwald et al. (1998).

Table 1. Relations between mapping compatibility, response-repetition, and task-switching

	Valence repetition	Valence switch
Compatible		
Task repetition	Response repetition	Response switch
Task switch	Response repetition	Response switch
Incompatible		
Task repetition	Response repetition	Response switch
Task switch	Response switch	Response repetition

Note: Response repetition and valence repetition are confounded in the compatible condition, which follows from the contingency of valence and responses. Task switching moderates this relation in the incompatible condition.

Participants

Participants were 32 students (22 females, 10 males) recruited from different faculties of the University of Bonn. They received either partial course credit or a monetary gratification of DM10 for their participation.

Material

The material consisted of 96 words referring to insects, flowers, positive objects, and negative objects. To minimize material effects, words were matched in quadruples that were maximally similar on three criteria, i.e. the number of characters, an estimation of the word's frequency of use based on the CELEX lexical database (Celex, 1995), and a rating of the word's valence. In particular, the stimulus words were thereby matched for frequency of occurrence to rule out a familiarity-based explanation of the effect (Dasgupta, McGhee, Greenwald, & Banaji, 2000).

800 positive and negative words were selected from a rating study of German substantives by Schwibbe, Räder, Schwibbe, Borchhardt, and Geiken-Pophanken (1994). Based on the 200 most positive and the 200 most negative words, three independent raters excluded those words that were either judged outdated or were unknown to them. The raters also eliminated all words describing plants or animals, to prevent semantic overlap between the categories. 132 positive and 143 negative stimuli remained in the sample.

The same raters also rated the initial 186 words referring to insects or flowers that were collected from different sources with respect to the same criteria. Based on the judgments, 46 flower words and 62 insect words were chosen. Ten participants then rated the valence of these 108 words in a pilot study, using the same question and response format as in the Schwibbe et al. (1994) studies. An estimation of the word's frequency of use was taken from the CELEX lexical database (Celex, 1995).

All words were represented as points in a 3-dimensional space defined by the z-transformed valence rating, frequency estimate, and word length. The matching process was based on an algorithm that sequentially extracted word quadruples including a word from each of the four sets such that the mean Euclidean distance between them was minimal. The 24 quadruples with the smallest distance were selected for the main experiments and are listed in Appendix A.

Procedure

All blocks consisted of the sequential presentation of 48 single words. The words were sampled without replacement from the stimulus lists for each block such that words from each list appeared equally frequently. The ordering of words within each block was randomized. The words were presented in a 20 mm × 120 mm rectangle on the computer screen, and written in black on light gray background. The presentation of a new word began either 100 ms or 1000 ms after the participants' response to a previous item, depending on the between-participants RSI-condition. Responding was allowed as soon as the stimulus was visible. In total, 14 blocks were presented, thus each participant underwent 672 separate trials. Performance data were recorded for every trial.

Participants were explicitly instructed to make an evaluative decision, if the presented stimulus was a positive or negative word, and a flower-insect discrimination, if it was a flower or insect name. Participants were told to respond to each word as rapidly as possible while avoiding errors. They started the upcoming block by sequentially pressing the two response keys. After a short countdown, the block was initiated.

The experiment started with two training phases consisting of two blocks each. In the first of these phases, only insect and flower names were presented, and participants practiced the insect-flower discrimination task. In the second phase, the evaluative decision task was practiced. The four training blocks were followed by four combined blocks, in which both tasks were mixed and were mapped either compatibly or incompatibly, depending on the order-balancing condition. The remaining six blocks consisted of two simple and four compound blocks, for which the compatibility of response mapping was switched. At the beginning of each phase, participants were informed about the word categories that were to appear in the upcoming block and their assignment to the response keys.

Results

All trials with latencies below 300 ms (0.5%) and above 3000 ms (0.5%) were excluded from the analyses. Another 8.7% of the trials were excluded from analyses of latency data, since responses were incorrect on these trials. Mean latencies and error proportions were calculated for each participant in each of the 2 × 2 × 2 within-participants conditions.

Mean latencies and error proportions were examined by repeated-measures ANOVAS with order

Table 2. Mean latencies in milliseconds as a function of mapping compatibility, response-stimulus interval, and the trial-sequence factors in Experiment 1

	Task Switch		No Task Switch	
	Response Switch	Response Repetition	Response Switch	Response Repetition
RSI 100 ms				
Compatible	798	772	775	713
Incompatible	1030	1037	919	801
RSI 1000 ms				
Compatible	675	675	690	633
Incompatible	819	798	758	715

of mapping conditions (compatible vs. incompatible condition first) and response-stimulus interval (100 ms vs. 1000 ms) as between-participants factors. Within-participants factors were mapping condition (compatible vs. incompatible), response repetition (repetition of response vs. switch of response between trial n and trial n-1), and task repetition (repetition of task vs. switch of task between trial n and trial n-1). The mean untransformed latencies are depicted in Table 2.

As predicted by Hypothesis 1, mean response latencies were significantly longer in the incompatible than in the compatible IAT condition, $F(1, 28) = 49.39$, $p < .01$, indicating a relative preference for flowers compared to insects. The mean aggregated latencies for the compatible and incompatible conditions were 716 ms (*SD* 26 ms) and 859 ms (*SD* 26 ms) respectively, resulting in an IAT effect of 143 ms.

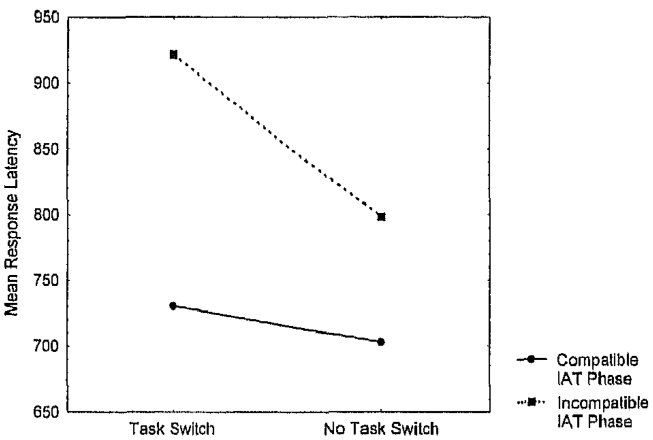

Figure 1. Average response latencies in Experiment 1 as a function of compatibility condition and task-switching. The implicit association may be defined by the differences between compatible and incompatible phases in these conditions.

The performance cost associated with switch and non-switch trials differed significantly between the two IAT conditions, $F(1, 28) = 42.53$, $p < .01$ as predicted by Hypothesis 2: Task switching had a stronger impact on performance in the incompatible than in the compatible IAT condition. This effect is depicted in Figure 1. Additionally, a main effect of task switching, $F(1, 28) = 62.52$, $p < .01$, reached significance. Planned contrasts revealed, that specific task-switch cost was present in the incompatible condition, $F(1, 28) = 81.88$, $p < .01$, as well as in the compatible condition, $F(1, 28) = 7.24$, $p < .05$.

The three-way interaction between compatibility, task switching and response repetition predicted by Hypothesis 3, did not reach significance, $F(1, 28) = 1.38$, n.s. The two-way interaction of task switching and response repetition was highly significant, $F(1, 28) = 17.23$, $p < .01$, and is displayed in Figure 2. Furthermore, a significant main effect of response repetition emerged, $F(1, 28) = 36.69$, $p < .01$. A planned contrast revealed, however, that performance on response-repetition trials was not increased when the task had to be switched, $F(1, 28) = 1.09$, n.s., in line with Hypothesis 3.

A main effect of response-stimulus interval, $F(1, 28) = 6.03$, $p < .05$, reveals that responses were faster with long than with short RSI. The two-way interaction between task switching and response-stimulus interval, $F(1, 28) = 11.59$, $p < .01$, shows that this effect was more pronounced for the task-switch trials. In other words, specific task-switch costs were reduced under long RSI.

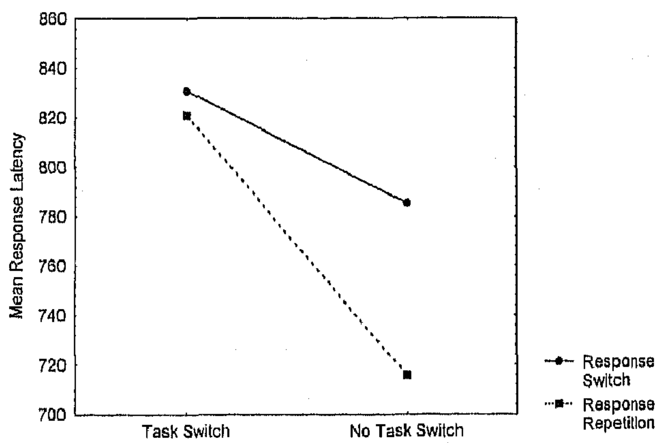

Figure 2. Mean response latencies in Experiment 1 as a function of task-switching and the repetition vs. switch of response. Note that no effect of response-repetition occurs on task-switch trials.

Discussion

Testing the hypotheses derived from the task-set switching account of the IAT produced a number of informative results. First and most notably, the prediction of higher cost for specific task switching in the incompatible IAT condition than in the compatible condition was confirmed. This results is not only in line with the idea that the IAT effect represents general task-set switching costs, but directly confirms that a specific component of these costs is more pronounced in the incompatible IAT condition. It is also noteworthy that this interaction is neither predicted by nor compatible with any of the alternative accounts presented earlier, with the exception of the figure-ground asymmetry model in its recent form (Rothermund & Wentura, 2001) that also postulates task-switching costs. Furthermore, it challenges a model that is exclusively based on goal-independent spreading activation between conceptual nodes, as task-related effects cannot be explained in such a network.

The finding that facilitative effects of response repetition disappear if the task has to be switched between two trials supports our assumption that performance in the IAT paradigm is affected by task-set reconfiguration processes. The effect, however, did not interact with mapping compatibility. In combination with the finding of generally higher specific task-switch cost in the incompatible mapping condition, this result is somewhat surprising. It seems to indicate that reconfiguration processes take place in the compatible condition, but that these processes are not accompanied by any specific performance cost. This, together with the observation that specific task-switching costs were smaller when the RSI was increased, favors the task-set interference account we outlined earlier, rather than the task-switch neglect model: The longer RSI might have allowed a passive dissipation of task-set activation and inhibition, thereby differentially reducing carryover effects. This interpretation is not limited by the absence of generally accepted standards of how activation and inhibition should dissipate in time (e.g. Los, 1996), since this is critical only if such effects do not emerge.

Note that the disruption of response-repetition effects on task-switch trials is incompatible with both variants of the random-walk model proposed by Brendl et al. (2000). According to this model, evidence for one response alternative is accumulated without respect to the source of this evidence. The only moderator for response-repetition effects should be the state of the hypothetic evidence counter.

Experiment 2

The results of Experiment 1 imply that the IAT effect reflects *higher costs* of switching between discrimination tasks in the incompatible IAT condition than in the compatible condition. They are, however, not conclusive with respect to the causal direction of this effect. Experiment 2 was designed to address this question directly, by manipulating the cost associated with task switching by experimental means. The advantage of this strategy is that it enables us to demonstrate a reduced IAT effect *caused* by smaller cost for task-set switching. The hypothesis for this experiment was therefore that, besides a replication of the effects found in Experiment 1,

H4) the IAT effect is reduced, if the costs for task-set switching are reduced.

Task-switch costs should be generally reduced, if participants are allowed (or even encouraged) to reconfigure processing mode in advance of the stimulus. Several experimental manipulations can be expected to achieve this, e.g. explicitly announced predictable runs of trials (Rogers & Monsell, 1995; Wylie & Allport, 2000) or instructed switches in small stimulus lists (Gopher, 1996; Kramer, Hahn, & Gopher, 1999). Applying these methods to the IAT would require more or less dramatic changes in procedural structure, which we wanted to avoid. Instead, we attempted to make sure that the compatibility effects we investigate are indeed the same type of effects found in IAT studies with a more applied focus.

We therefore preferred a paradigm that makes use of task cues (Meiran, 1996), which are presented in advance of the stimulus, but do not require special instructions or changes in the randomized nature of the stimulus sequences presented. Several other properties of this procedure are noteworthy. First, effects of advance reconfiguration are not confounded with effects of response-stimulus interval, as the cueing paradigm allows the comparison of a cued vs. uncued condition with *constant RSI.* Second, the cues may be presented without any further instruction to use them. Instructions could result in different processing strategies and thus reduce the generalizability of the cueing effect. Third, the information contained in task cues is not confounded with any property of the upcoming task that would be expected to have an effect on the basis of other IAT models. That is, neither the valence of the presented stimulus, nor the required response can be predicted on the basis of the cue. Fourth, the cues can be administered in the compatible and incompatible IAT condition without further changes in the design.

Method

The design of Experiment 2 is a 2 × 4 × 2-design with mapping condition (compatible vs. incompatible) as a within-participants factor, and cue condition, and order of mapping conditions (compatible vs. incompatible mapping condition as first compound block) as between-participants factors. The complete response-key assignment was again switched for one half of the participants. Two additional sequence factors emerge from the coding of trial sequences.

Participants

Sixty students (42 females, 18 males) recruited from different faculties of the University of Bonn participated in the experiment. They received either partial course credit or a monetary gratification of DM10 for their participation.

Material

The material for Experiment 2 was based on the selection of words for Experiment 1, except for minor changes. The changes were motivated by poor discrimination performance on some of the German words, i.e. "Beteuerungen", "Eigenmächtigkeit", and "Eigensinn", which probably result from a change in affective connotation for these words. In order to uphold similarity between the stimulus lists, the quadruples containing these words were excluded and replaced by the most homogenous unused quadruples. The additional stimuli can be found in Appendix B.

In addition to the stimulus words, task cues were presented in two conditions of the experiment. The task cues were realized by using simple symbols that represent the word categories used in the task. These symbols can be found in Appendix C.

Procedure

The presentation of task cues was varied in four conditions between participants. Task cues consisted of symbols that represent the concepts used in the task, i.e. a flower and an insect symbol for the flower-insect discrimination task, and two pictures of symbolized thumbs pointing upwards and downwards for the evaluative decision task. The symbols were presented simultaneously to the left and right side of the stimulus word such that the position of each symbol also specified the appropriate response key. Although the cues specified the (constant) stimulus-response mapping, they were uninformative with respect to the correct response or the valence or category of the stimulus word.

Task cues were either presented 600 ms in advance of the imperative stimulus or simultaneously. Two control conditions were realized, one without task cues (Control 1), and the other one with all symbols presented simultaneously (Control 2). The latter symbol constellation contained information on the stimulus-response mappings, but did not predict the appropriate task set. In these two conditions, an unmitigated implicit association effect was expected, as temporal information and information on stimulus-response mappings was not assumed to have an impact on this effect. In contrast, a task-cueing effect on the IAT measure and the more specific indicator of task-switch cost was expected for both cueing conditions. The effect of advance reconfiguration is expected to be higher for the cue presented 600 ms before the stimulus.

The response-stimulus interval was fixed at 800 ms for all trials. Moreover, an asterisk ("*") was presented 600 ms before the stimulus in all conditions. This provided all participants with the same temporal information to predict stimulus onset. A mere temporal warning function of the advance cues is thereby ruled out. Experiment 2 was identical to Experiment 1 in all other respects.

Results

Trials with latencies below 300 ms (0.9 %) and above 3000 ms (1 %) were excluded from the analyses. On 7.2 % of the remaining trials, participants made response errors. These trials were also excluded from the analyses of response latencies.

Performance data were examined by repeated measures analysis of variance with order of mapping conditions (compatible vs. incompatible condition first) and cueing condition (advance task-cue vs. simultaneous task-cue vs. Control 1 vs. Control 2) as between-participants factors. Within-participant factors included mapping compatibility and the two repetition factors. Mean untransformed latencies are given in Table 3.

The analysis of response latencies revealed a main effect of mapping compatibility, $F(1, 52) = 75.54$, $p < .01$, i.e. a large IAT effect in the expected direction. The mean aggregated latencies for compatible and incompatible condition were 685 ms (*SD* 18 ms) and 806 ms (*SD* 21 ms) respectively, resulting in an IAT effect of 121 ms.

Table 3. Means latencies in milliseconds as a function of mapping compatibility, cueing condition, and the trial-sequence factors in Experiment 2

	Task Switch		No Task Switch	
	Response Switch	Response Repetition	Response Switch	Response Repetition
Advance Task-Cue				
Compatible	700	717	668	664
Incompatible	795	823	725	708
Simultaneous Task-Cue				
Compatible	652	645	645	619
Incompatible	773	792	728	661
Control 1				
Compatible	700	707	703	662
Incompatible	941	932	867	797
Control 2				
Compatible	750	747	727	682
Incompatible	876	890	807	763

As predicted by Hypothesis 4, there was a significant interaction of task cueing and IAT condition, $F(3, 52) = 3.63$, $p < .05$. This indicates that the size of the implicit association effect was reduced by presenting task cues and is depicted in Figure 3. The effect, however, did not interact with sequential-level task-switching, $F(3, 52) = 0.64$, n.s., which would have demonstrated that the effect of the cueing manipulation was mediated by *specific* task-switch cost. A planned contrast reveals that the effect of cueing condition was due to a smaller IAT effect in the two conditions with a task cue, compared to those without task cueing, $F(1, 52) = 5.34$, $p < .05$. Testing the magnitude of the implicit association effect with advance task-cueing against the other three conditions does not reveal an effect.

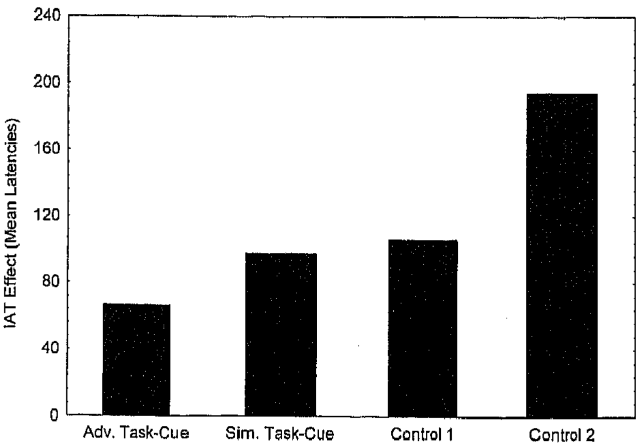

Figure 3. The IAT effect as a function of task-cueing condition. Bars display the difference between incompatible minus compatible IAT condition in terms of mean response latencies.

The interaction of task switching and IAT condition predicted in Hypothesis 2 was replicated, $F(1, 52) = 41.69, p < .01$, i.e. a larger IAT effect for task-switch trials. An overall performance difference between switch and no-switch trials also reached significance, $F(1, 52) = 128.25, p < .01$. Thus, the interesting pattern of results from Experiment 1 could be replicated: The implicit association effect was moderated by differences in specific task switching. In addition, the moderating role of task switching seems to be so dominant that it could not be eliminated by our cueing manipulation.

As predicted by Hypothesis 3, the three-way interaction of response repetition, task switching, and IAT condition was significant, $F(1, 52) = 4.16$, $p < .05$. Response-repetition effects were more strongly disrupted by task switches in the incompatible than in the compatible condition. Like in Experiment 1, overall performance was better, if the required response to the actual stimulus was the same as on the previous trial, $F(1, 52) = 8.59$, $p < .01$. This effect was also moderated by the task-switch factor, independently of IAT condition, $F(1, 52) = 29.31$, $p < .01$, i.e. there were generally smaller facilitative effects, if the task had to be switched.

Discussion

The results of Experiment 2 demonstrated a direct effect of task cues on an IAT measure: The performance difference between the compatible and incompatible phase was smaller when task cues were presented than when no cues were presented. As this manipulation was neither confounded with foreperiod, nor attributable to the priming of stimulus or response features for the upcoming trial, this effect is most probably due to a facilitation of task-appropriate advance-reconfiguration of processing mode, but not to passive dissipation of activation. This effect provides support for the assumption that the IAT effect reflects differential costs associated with task-set switching.

Further support for this account arises from the fact that the specific prediction on effects of response repetition was confirmed in Experiment 2:

Response-repetition effects *were* disrupted more strongly by task switches in the incompatible IAT condition than in the compatible condition. Moreover, the interaction of IAT condition and specific task-switching costs, already found in the first experiment, could be replicated. Task switching had highly specific differential costs and effects on the two IAT conditions. These findings are incompatible with any other current account of the IAT effect.

The question remains why task-switching costs were not fully eliminated by the presentation of task cues. The two-way interaction of task-switch cost and compatibility condition reflects a large component of variance that was independent of the cueing condition. It has to be noted that the task-cueing effect we found is an incidental effect, i.e. participants were neither encouraged nor instructed to make use of the task cues. Indeed, residual cost for task switches was found in various studies even when advance reconfiguration would have been possible (e.g. Rogers and Monsell, 1995). Rogers and Monsell suggest that "completion of the reconfiguration is triggered only by, and must wait upon, the presentation of a task-associated stimulus" (Rogers & Monsell, 1995, p. 224). Because there was no significant difference between advance and simultaneous cues, it seems possible that the task cues used in the present experiment were most helpful in this latter, stimulus-triggered reconfiguration phase.

Furthermore, task cues in the present experiments did not contain information that was *necessary* for task execution, since the appropriate decision could always be inferred from the current stimulus alone. It was therefore possible be to direct attention exclusively to the imperative stimuli, and to ignore the task cues. If so, a more dramatic reduction – or even elimination – of the IAT effect might be achieved by incorporating other means of reducing task-switch cost. Doing so, however, would require a more fundamental change in the IAT structure.

General Discussion

As discussed in the introduction, there is some evidence suggesting that IAT effects are based on properties of the target categories to a larger degree than on properties of the specific stimuli used to instantiate these categories. Steffens and Plewe (2001) on the other hand have provided challenging evidence for an effect of stimulus features. Providing a model of the mechanisms and structures underlying the perceptual processes that lead to the activation of irrelevant attribute information is beyond the scope of the current paper. Based on the evidence by De Houwer (in press) and Neumann et al. (2000), it seems plausible to assume that these processes are not completely stimulus-driven.

The results presented here do not challenge the internal validity of the IAT as such, but the assumption that the magnitude of the IAT effect is monotonically related to associative strength. Both mechanisms we have proposed are based on the assumption that attribute-related information becomes available upon processing a stimulus from the target categories. Performance costs in the incompatible IAT condition, however, may not directly depend on the amount of activation of irrelevant attribute information, but rather on the duration of executive control processes responsible for task-set organization or the passive dissipation of residual task-set activation.

Distinguishing empirically between specific and general costs of task-set switching requires trials in which the task is repeated. The relative contribution of specific and general cost factors cannot be disentangled when tasks always alternate from trial to trial (e.g. Banse et al., 2001; Greenwald et al., 1998). Using alternating tasks may have two distinct consequences. First, it allows the participants to prepare for the upcoming task, if the response-stimulus interval is sufficiently long. This might lead to a reduction of task-switching cost, and thereby a reduction of the IAT effect. This is not very probable, however, as the response-stimulus interval used in IAT studies is typically much lower than the 600 ms Rogers and Monsell (1995) found to be optimal for advance reconfiguration. Second, the use of strictly alternating tasks maximizes the number of task-switch trials in a block. Banse et al. (2001) report a significant correlation of the randomized, but not the alternating variant of their Homosexuality-IAT with an explicit measure. The mean IAT-score was, however, smaller for the randomized version, which points to the possibility that the IAT-score variance caused by differential task-switching costs is not content related.

Some wider implications of the present findings concern the discriminant validity of the IAT, as costs for task-set switching are known to correlate with a number of other variables. Kray and Lindenberger (2000) report substantial correlations of general task-set switching costs with fluid, mechanical aspects of intelligence. Further, a substantial age-related decrement in performance with alternating tasks was observed (Kramer et al., 1999; Kray & Lindenberger, 2000), even though Kramer et al. (1999) also found a significant reduction of specific task-set switching cost for older adults after a substantial amount of practice. With respect to the IAT, these findings indicate possible contaminations of the measure, mediated by the influence of executive control functions on performance. One practical implication of such findings is that the IAT should be interpreted with special care in quasi-experimental designs, if the contribution of the factors mentioned above cannot be controlled. A recent experiment from our own labs has shown a substantial correlation of IAT effects and a pure measure of task-set switching, if the incompatible IAT phase was performed first. Whether the validity of IAT results is improved by using such a measure as a covariate, however, is an open question at present.

Based on the model we presented, we may deduce that associations in social or conceptual knowledge structures may be a sufficient, but by no means necessary condition for IAT effects. This is underlined by Rothermund and Wentura's (2001) results demonstrating IAT-like effects when attribute and target categories are related indirectly by a third variable. The task-switching model presented here may be extended to account for such related effects as well. Effects mediated by a specific third variable, however, may not reflect a typical IAT-effect.

In sum, the findings we presented make a step toward a psychological model of the processes and structures underlying the IAT effect. In line with the above-mentioned theoretical arguments against a straightforward explanation in terms of spreading activation, our data demonstrate several effects that are inconsistent with such an explanation. The expected effect of higher specific task-switching costs for the incompatible IAT condition was reliably demonstrated in two experiments: Task switching contributed to the implicit association measure's magnitude. Experiment 2 directly tested the hypothesis that the IAT effect can be reduced by reducing task-switching costs. The effect of task cues cannot be attributed to a mere lengthening of the foreperiod, as the foreperiod was fixed to 800 ms in all experimental conditions.

The effect of response-stimulus interval on specific task-switch cost might be interpreted as evidence for a passive dissipation of carryover from one trial to another (Allport et al., 1994). Contrasting the RSI related findings (indicating passive dissipation of carryover) with the effect of task cues there is evidence that both passive dissipation as well as endogenously controlled processes affect task-switching costs. It seems plausible that the function of an endogenous control process could be to override such carryover from previous action (Monsell et al., 2000). In addition, some other results indicate that task switching occurred in both IAT conditions, but is associated with higher cost in the incompatible condition. These findings are in line with the task-set interference model we sketched in the introduction. We will outline how such a model can account for the various effects that we reported, and discuss its assumptions, in the following sections.

The task-interference model is based on a goal-oriented mechanism that prevents the system from engaging in two actions at the same time if they are physically incompatible. This may be a key feature of attentional processes, as Neumann points out in his elegant functional analysis of selective attention (e.g. Neumann, 1996), but quite plausibly is a major function of executive control processes, too. The operative repertoire of this mechanism may consist of two basic operations, i.e. activation of now-appropriate task-sets or task-set components and inhibition of now-inappropriate task-sets or task-set components (cf. Norman & Shallice, 1986). How may IAT effects be explained on this basis?

As already pointed out, the only stimuli that may induce processing conflicts are target-set stimuli in the incompatible condition. This is due to the fact that attribute-set stimuli are neutral with respect to the target-discrimination task (at least if positive stimuli are not assumed to be more "flower-like") in both IAT conditions, while the relevant and irrelevant features of target-set stimuli require physically compatible responses in the compatible IAT condition. To respond correctly to target-set stimuli in the incompatible condition, task-sets that specify a response based on the irrelevant attribute information need to be inhibited, in order to avoid physically incompatible action.

Based on the above assumptions, three consequences of such inhibitory processes may be distinguished:

1) When switching to a target-set stimulus in the incompatible condition, response-selection is slowed down, until the now-inappropriate attribute task-set is inhibited sufficiently. No such effect is expected in the compatible IAT condition, as the now-inappropriate attribute task-set does not interfere with response-selection.

2) When switching to an attribute-set stimulus in the incompatible condition, response are slowed down, because the now-appropriate task-set has been inhibited on the preceding trial, i.e. a negative task-set priming (Wylie & Allport, 2000; cf. Monsell et al., 2000). No such effect is expected in the compatible IAT condition, as the now-appropriate task-set has not been inhibited on the preceding trial. Whether release from inhibition indeed operates on the global level of task-sets, or on more specific component processes (cf. Tipper, 1985; Tipper, MacQueen, & Brehaut, 1988), has to be addressed in future research.

3) Performance is relatively faster when the processing system has already been configured optimally on the preceding trial, i.e. on task-repetition trials.

Taken together, 1) and 2) can explain the IAT effect, as they predict additional cost that are specific to the incompatible IAT condition. The basic idea is that both, the need to inhibit the attribute task-set when it is irrelevant and the need to release this inhibition when it is relevant are the source of the IAT effect. Considering Prediction 3), it becomes clear that this effect should be specific to task-switching trials, which was confirmed empirically in both reported experiments. Note that based on prediction 2) slower performance would be expected for attribute-set stimuli in the incompatible compared to the compatible condition. This was indeed reported by Brendl et al. (in press) and was explained by a criterion shift in the incompatible condition. From our point of view, this is due to a carryover of inhibition from the previous trials. Moreover, the finding of disrupted response-repetition effects when the task had to be switched in the incompatible IAT phase matches well with the discontinuity of processing that is expected under these specific conditions.

The model furthermore predicts the finding that cost for specific task-switching are more pronounced when the RSI is short. The RSI quite plausibly moderates the degree of carryover from one trial to another, i.e. carryover should be smaller, when there is more time for a passive decay of activation and release from inhibition. The finding of a reduced IAT effect when task cues were presented points to the possibility that task-set activation can not only be affected by passive decay, but also by control processes that can operate in advance of the stimulus.

Activation of the now-inappropriate task-sets should produce more interference on a given trial, if the residually activated task set is applicable to the processing of the current stimulus. As we pointed out, this should be the case for switches to the target-discrimination task, but not for switches to the attribute-discrimination task. Inhibition of the now-inappropriate task set on the preceding trial, on the contrary, should affect performance for the attribute-discrimination task more strongly. The relative contribution of these distinct factors in IAT effects seems to be an interesting field for future research.

References

Allport, D. A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umilta & M. Moscovitch (Eds.), *Attention and performance XV* (pp. 421–452). Cambridge, MA: MIT Press.

Banse, R., Seise, J., & Zerbes, N. (2001). Implicit attitudes towards homosexuality: Reliability, validity, and controllability of the IAT. *Zeitschrift für Experimentelle Psychologie, 48* (this issue), 145–160.

Bargh, J. A., Chaiken, S., Raymond, P., & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology, 32,* 104–128.

Brendl, C. M., Markman, A. B., & Messner, C. (in press). How do indirect measures of evaluation work? Evaluating the interference of prejudice in the Implicit Association Test. *Journal of Personality and Social Psychology.*

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic priming processing. *Psychological Review, 82,* 407–428.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8,* 240–248.

Celex (1995). *The Celex Lexical Database, Release 2.* Nijmegen: Center for Lexical Information.

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. (in press). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology.*

Ehrenberg, K., Klauer, K. C., & Wegener, I. (2000). *Hierarchische Verarbeitung multipler sozialer Kategorieinformation: Die Dominanz von Geschlechtskategorien* [Hierarchical processing of multiple social category information: The dominance of gender categories]. Paper presented at the 42. Tagung experimentell arbeitender Psychologen, Braunschweig.

Gopher, D. (1996). Attention control: explorations of the work of an executive controller. *Cognitive Brain Research, 5,* 23–38.

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

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology, 74,* 1464–1480.

Klauer, K. C., & Musch, J. (in press). Does sunshine prime loyal? Affective priming in the naming task. *Quarterly Journal of Experimental Psychology.*

Klauer, K. C., & Wegener, I. (1998). Unraveling social categorization in the "Who Said What?" paradigm. *Journal of Personality and Social Psychology, 75,* 1155–1178.

Kornblum, S. (1992). Dimensional overlap and dimensional relevance in stimulus-response and stimulus-stimulus compatibility. In G. E. Stelmach & J. Requin (Eds.), *Tutorials in motor behavior, Volume 2* (pp. 743–777). Amsterdam: North Holland.

Kramer, A. F., Hahn, S., & Gopher, D. (1999). Task coordination and aging: explorations of executive control processes in the task-switching paradigm. *Acta Psychologica, 101,* 339–378.

Kray, J., & Lindenberger, U. (2000). Adult age differences in task switching. *Psychology and Aging, 15,* 126–147.

Kühnen, U., Schießl, U., Bauer, N., Paulig, N., Pöhlmann, N., & Schmidthals, N. (2001). How robust is the IAT? Measuring and manipulating implicit attitudes of East- and West-Germans. *Zeitschrift für Experimentelle Psychologie, 48* (this issue), 135–144.

Los, S. A. (1996). On the origin of mixing costs: Exploring information processing in pure and mixed blocks of trials. *Acta Psychologica, 94,* 145–188.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin, 109,* 163–203.

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 1423–1442.

Monsell, S., Yeung, N., & Azuma, R. (2000). Reconfiguration of task-set: Is it easier to switch to the weaker task? *Psychological Research, 63,* 250–264.

Neely, J. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–337). Hillsdale, NJ: Erlbaum.

Neumann, O. (1996). Theories of attention. In O. Neumann & A. F. Sanders (Eds.), *Handbook of perception and action. Volume 3: Attention* (pp. 389–446). San Diego: Academic Press.

Neumann, R., Ebert, M., Gabel, B., Guelsdorff, J., Krannich, H., Lauterbach, C., & Wiedl, K. (1998). Vorurteile zwischen Bayern und Norddeutschen: Die Anwendung einer neuen Methode zur Erfassung evaluativer Assoziationen in Vorurteilen [Prejudice prevailing between Bavarians and North Germans: Application of a new method for the assessment of evaluative associations in prejudices]. *Zeitschrift für Experimentelle Psychologie, 45,* 99–108.

Neumann, R., Totzke, I., Popp, C., & Fernandez, O. (2000). *Distinct attitudes despite similar targets: Comparing German students' attitudes towards "blacks" and foreigners.* Unpublished manuscript.

Norman, D., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self regulation, Volume 4* (pp. 1–18). New York: Plenum Press.

Pashler, H., & Baylis, G. (1991). Procedural learning: 2. Intertrial repetition effects in speeded-choice tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 33–48.

Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General, 124,* 207–231.

Rothermund, K., & Wentura, D. (2001). Figure-ground-asymmetries in the Implicit Association Test (IAT). *Zeitschrift für Experimentelle Psychologie, 48* (this issue), 94–106.

Schwibbe, M., Räder, K., Schwibbe, G., Borchardt, M., & Geiken-Pophanken, G. (1994). Zum emotionalen Gehalt von Substantiven, Adjektiven und Verben (On the emotional content of nouns, adjectives, and verbs). In W. Hager & M. Hasselhorn (Eds.), *Handbuch deutschsprachiger Wortnormen* (pp. 272–284). Göttingen: Hogrefe.

Simon, J. R. (1990). The effect of an irrelevant directional cue on human information processing. In R. W. Proctor & T. G. Reeve (Eds.), *Stimulus-response compatibility: An integrated perspective. Advances in psychology, 65* (pp. 31–86). Amsterdam: North Holland.

Stangor, C., Lynch, L., Duan, C., & Glass, B. (1992). Categorization of individuals on the basis of multiple social features. *Journal of Personality and Social Psychology, 62,* 207–218.

Steffens, M. C., & Plewe, I. (2001). Items' cross-category associataions as a confounding factor in the Implicit Association Test. *Zeitschrift für Experimentelle Psychologie, 48* (this issue), 123–134.

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

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

Tipper, S. P., MacQueen, G. M., & Brehaut, J. C. (1988). Negative priming between response modalities: Evidence for the central locus of inhibition in selective attention. *Perception and Psychophysics, 43,* 45–52.

Wylie, G., & Allport, A. (2000). Task switching and the measurement of "switch costs". *Psychological Research, 63,* 212–233.

Jan Mierke

Psychologisches Institut
Universität Bonn
Römerstr. 164
D-53117 Bonn
E-Mail: Jan.Mierke@uni-bonn.de

Appendix A – Stimuli for Experiment 1

Positive words	*Negative words*	*Flower*	*Insects*
Leckerbissen	Mißbilligung	Schwertlilie	Springspinne
Echtheit	Torheit	Flieder	Moskito
Erlösung	Nässe	Lilie	Wespe
Freundlichkeit	Verheimlichung	Maiglöckchen	Wespenspinne
Lebhaftigkeit	Mutmaßungen	Butterblume	Borkenkäfer
Nektar	Einbuße	Primel	Mücke
Gemüt	Stoß	Nelke	Fliege
Diamant	Ermüdung	Edelweiß	Termite
Weise	Eigensinn	Feuerlilie	Mistkäfer
Säugling	Notfall	Orchidee	Spinne
Photographie	Beteuerungen	Chrysantheme	Fruchtfliege
Humor	Phrase	Rose	Laus
Duft	Makel	Iris	Made
Leichtigkeit	Hirngespinst	Gänseblümchen	Stubenfliege
Musikinstrument	Untauglichkeit	Vergißmeinnicht	Kartoffelkäfer
Eingebung	Bettler	Veilchen	Schabe
Bücherei	Zwielicht	Krokus	Hornisse
Fröhlichkeit	Rückschritt	Sonnenblume	Küchenschabe
Busen	Meineid	Tulpe	Wanze
Aufnahmefähigkeit	Eigenmächtigkeit	Stiefmütterchen	Blatthornkäfer
Flexibilität	Trugschluß	Pusteblume	Blattlaus
Charme	Orkan	Mohn	Floh
Redlichkeit	Banalität	Klatschmohn	Blattwanze
Herzlichkeit	Dumpfheit	Seerose	Kakerlake

English Translation

Positive words	*Negative words*	*Flowers*	*Insects*
delicacy	disapproval	yellow iris	jumping spider
authenticity	foolishness	lilac	mosquito
redemption	moistness	lily	wasp
friendliness	concealment	lily of the valley	wasp-spider
liveliness	speculation	buttercup	bark-beetle
nectar	forfeit	primrose	gnat
mind	stroke	clove	fly
diamond	fatigue	edelweiss	termite
sage	stubbornness	orange lily	dung beetle
baby	emergency	orchid	spider
photography	reaffirmation	chrysanthemum	drosophilae
humor	phrase	rose	louse
scent	flaw	iris	maggot
easiness	phantasm	daisy	housefly
musical instrument	incompetency	forget-me-not	potato beetle
intuition	beggar	violet	scraper
library	twilight	crocus	hornet
gladness	regress	sunflower	housemartin
bosom	perjury	tulip	bug
receptivity	arbitrary act	viola	(Blatthornkäfer)

Positive words	*Negative words*	*Flowers*	*Insects*
flexibility	fallacy	blowball	plant louse
charm	hurricane	poppy seed	flea
fidelity	banality	corn poppy	(Blattwanze)
cordiality	dullness	water lily	cockroach

Note: The translations provided here may only approximately reflect evaluation, meaning, and frequency of occurrence of the German words used in the original studies.

Appendix B – Additional Stimuli for Experiment 2

Positive words	*Negative words*	*Flowers*	*Insects*
Grundrecht	Abrechnung	Narzisse	Ameise
Gelassenheit	Einbildung	Pfingstrose	Rüsselkäfer
Barmherzigkeit	Fragwürdigkeit	Alpenveilchen	Nachtfalter

English Translation

Positive words	*Negative words*	*Flowers*	*Insects*
basic right	reckoning	narcissus	ant
calmness	conceit	peony	weevil
charity	dubiousness	cyclamen	moth

Note: The translations provided here may only approximately reflect evaluation, meaning, and frequency of occurrence of the German words used in the original studies.

Appendix C – Symbols Used as Task-Cues in Experiment 2

Positive Words	*Negative Words*	*Flower Names*	*Insect Names*