



Implicit Association Test as an Analogical Learning Task

Ian Hussey and Jan De Houwer

Department of Experimental Clinical and Health Psychology, Ghent University, Belgium

Abstract: The Implicit Association Test (IAT) is a popular tool for measuring attitudes. We suggest that performing an IAT could, however, also change attitudes via analogical learning. For instance, when performing an IAT in which participants categorize (previously unknown) Chinese characters, flowers, positive words, and negative words, participants could infer that Chinese characters relate to flowers as negative words relate to positive words. This analogy would imply that Chinese characters are opposite to flowers in terms of valence and thus that they are negative. Results from three studies ($N = 602$) confirmed that evaluative learning can occur when completing an IAT, and suggest that this effect can be described as analogical. We discuss the implications of our findings for research on analogy and research on the IAT as a measure of attitudes.

Keywords: Implicit Association Test, analogy, learning

It seems safe to assume that psychological assessment likely adheres to Heisenberg's observer effect in physics: by measuring, we perturb the system. This effect implies that the act of completing a psychological testing task simultaneously provides a training context for the organism. During the past two decades, the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998) has become one of the most popular measurement tasks in psychology. Across three experiments, we examined a novel way in which the IAT might function as a training task. More specifically, we tested whether the IAT is an analogical evaluative learning task.

The IAT has seen wide use as an implicit measure of attitudes in many domains of psychological research, including clinical, social, and experimental psychology (Greenwald, Poehlman, Uhlmann, & Banaji, 2009). The task requires individuals to quickly categorize stimuli as belonging to four superordinate categories that are presented on-screen (e.g., pictures of flowers, pictures of insects, positive words, and negative words). Two categories are mapped to each response option (e.g., left response = flowers or positive, right response = insects or negative). Importantly, this mapping changes across blocks (e.g., left response = flowers or negative, right response = insects or positive). Importantly, participants are often observed to exhibit an IAT effect, whereby they show faster responding in one block of categorizations relative to the other. For example, faster performance is typically observed when pictures of flowers and positive words are assigned to one key and pictures of insects and negative words to the second key compared to when flower and negative are assigned to one key and

insects and positive to the second key. Such differences in performance are assumed to reflect preexisting differences in attitudes.

Whereas the IAT has most frequently been employed as a testing task, variants of the task have also been employed as a training task. For example, Ebert, Steffens, Von Stülpmagel, and Jelenec (2009) demonstrated that completing an IAT-like categorization task induced evaluative learning (see Prestwich, Perugini, Hurling, & Richetin, 2010, for a related task). In one study, participants were asked to categorize candy, chocolate, and valenced words. When candy was mapped to the same response key as positive words and chocolate was mapped to the same response key as negative words, participants later evaluated candy as more positive than chocolate, and vice versa. It has been argued that this form of learning is one instance of learning via intersecting regularities (Hughes, De Houwer, & Perugini, 2016). Within the above example, the change in evaluation of candy versus chocolate was arguably driven by the intersection between the stimulus-response mappings that involve candy and positive (e.g., both required pressing a left key press) and those that involve chocolate and negative (e.g., both required a right key press). Such intersections have been shown to allow for a transfer of valence between the other members of the intersecting response mappings (Hughes et al., 2016). Importantly, however, research to date on the IAT as a training task has focused only on the effects of a single block of categorizations rather than the IAT as a whole.

We suggest that the IAT as a whole can function as a training task by specifying two relations and allowing the

individual to use one relation to inform the nature of the other. That is, we believe that it allows for learning via analogy. In order to understand what we are proposing here, it is important to realize that the IAT always includes two pairs of concepts. By definition, analogies are defined by the relating of relations between two pairs of concepts (Holyoak & Koh, 1987). Specifically, analogical inferences can take place when there is one source relation and one target relation, and one of the concepts of the recipient relation is similar to one of the elements of the source relation. When these conditions are fulfilled, it becomes possible to form a link between the second element of the source relation and the second element of the recipient relation (Gentner & Smith, 2013). We argue that the IAT fulfills these requirements by including two pairs of categories. Notionally, the relations between the target stimuli pair (e.g., positive and negative, which are opposites) could be taken as being indicative of the relation between the category stimuli pair (e.g., a positive category such as flowers, and an unknown novel category, such as Chinese characters), allowing the individual to construct an analogy (e.g., Chinese characters:flowers::negative:positive). If so, it therefore should be possible to use the IAT to induce learning via analogy. This would have implications for the traditional use of the measure as a testing task within many clinical and social domains of psychology. Specifically, knowing whether, and to what degree, the IAT serves to establish or change the very attitudes it seeks to assess would seem to be an important caveat to the use of the IAT. It would also have implications for learning psychology research, by potentially providing a relatively subtle way to induce (evaluative) learning.

In three experiments, we assessed changes in liking due to the IAT. Participants evaluated neutral, unknown Chinese characters, completed a (training) IAT that differed between conditions, and then evaluated the characters a second time. The training IATs were based on the original flower–insects evaluation IATs (Greenwald et al., 1998). Each contained a novel stimulus target category (Chinese characters) as well as a valenced stimulus target category (either pictures of flowers or pictures of insects). Both tasks also employed positive and negative words as attribute stimuli. The flowers IAT therefore required participants to categorize Chinese characters, flowers, negative words, and positive words with response mappings that varied across the two blocks (e.g., press left for Chinese characters and negative and right for flowers and positive in one block versus press left for flowers and negative and right for Chinese characters and positive words in a second block). In contrast, the insects IAT required participants to categorize images of Chinese characters, images of insects, positive words, and negative words. Such IATs therefore share some similarities to those employed by Brendl, Markman, and

Messner (2001), who required participants to complete an IAT containing the categories insects, nonwords, positive words and negative words, and subsequently rating the nonwords. However, in the absence of pre-IAT ratings or a comparison condition (e.g., a flowers IAT), it is not possible to know from Brendl et al.'s (2001) results whether ratings of the nonwords were affected by the completion of the IAT.

Importantly, previous research on learning via IAT-like categorization tasks has always involved only one of the IAT's block types (e.g., Ebert et al., 2009; Prestwich et al., 2010) and would therefore likely not generalize to the full task. In contrast, our paradigm involves the full IAT including both block types. As such, this is the first study to our knowledge to examine the full IAT as a training task. Furthermore, by involving both block types, our task did not directly train participants on the relation between Chinese characters and one type of valence stimuli. Because of this, it is difficult to account for any observed effects in terms of learning via intersecting regularities, or indeed in terms of other well-known learning effects such as classical or operant conditioning, given that the Chinese characters contained an equal number of intersections with both positive and negative stimuli (see De Houwer, 2007; De Houwer, Barnes-Holmes, & Moors, 2013).

Learning could, however, result on the basis of an analogy between pairs of categories. For instance, in the flowers IAT, participants could infer that Chinese characters relate to flowers as negative words relate to positive words. This analogy would imply that Chinese characters are opposite to flowers in terms of valence and thus that they are negative. Participants in the insects IAT, on the other hand, could infer that the Chinese characters are positive in valence. In sum, if the IAT can function as an analogical learning task, we expected more positive evaluations of the category “Chinese characters” when participants complete the insects IAT than when they complete the flowers IAT. For reasons that we will discuss at the end of this paper, observing these analogical learning effects would have interesting implications for both research on (learning via) analogy and research on the IAT as a measure of attitudes.

Experiment 1

All tasks were programmed using PsychoPy (Peirce, 2007) or Inquisit (Inquisit 4, 2015) and presented on-screen. All materials for both experiments including measures and R code for data processing and analyses and the preregistration for Experiment 3 are available on the Open Science Framework (<http://osf.io/t89fu>). We report how we

determined our sample size, all data exclusions (if any), all manipulations, and all measures in the study (Simmons, Nelson, & Simonsohn, 2012).

Method

Participants

Given the exploratory nature of the study, we recruited as many participants as possible within a pre-allocated data collection period of 3 weeks. Fifty-two students (37 women, 15 men, $M_{\text{age}} = 22.06$, $SD = 3.46$) at Ghent University were recruited and participated in the experiment in exchange for €5. All participants were fluent Dutch speakers and provided written informed consent prior to participation. All instructions and tasks were provided in Dutch, although the English translations are reported here.

Measures

Participants rated how much they liked each of the five Chinese characters and each of the five valenced images (flowers or insects, depending on condition). Each item employed a 1 (= *not at all*) to 5 (= *very much*) Likert scale.

For both IATs, the task parameters employed were identical to the archetype described by Nosek, Greenwald, and Banaji (2007). Nonetheless, we provide a brief description of the task parameters here. Before each block, participants were instructed that they would be presented with words and pictures in the middle of the screen and that they were to categorize these using the labels presented at the top left and top right of the screen. Participants were furthermore instructed that these labels were mapped onto the left and right response keys (E and I, respectively). Prior to each block, participants were also alerted to which categories would be presented on the next block. On each trial, participants were required to emit a correct response in order to advance to the next trial; accuracy feedback was delivered via a red X.

Five images of each of flowers, insects, and different Chinese characters served as target stimuli. Five positive and five negative words served as attribute stimuli (attractive, enjoy, favorable, likeable, and lovely; awful, despise, dirty, disgust, and nasty). Two IATs were created: the flowers IAT (flowers, Chinese characters, positive, negative) and the insects IAT (insects, Chinese characters, positive, negative).

The length and content of each block were as follows (Nosek et al., 2007): Block 1: 20 attribute trials (i.e., positive and negative words); Block 2: 20 target trials (i.e., Chinese characters and either flowers or insects), Block 3: 20 target and attribute trials; Block 4: 40 target and attribute trials; Block 5: 40 attribute trials; Block 6: 20 target and attribute

trials; and Block 7: 40 target and attribute trials. The location of the two target category labels reversed on Block 5, providing two contrasting categorization patterns across blocks. Comparisons were made between reaction times between stimulus presentation and correct response on Blocks 3 and 4 versus those on Blocks 6 and 7. The order of presentation of the blocks was also counterbalanced between participants so that half of participants first encountered a block in which Chinese characters and positive words were categorized using the same response key, and half first encountered a Chinese characters–negative words block.

Procedure

Participants were tested in individual experimental cubicles and were randomly assigned to the flowers condition or insects condition. All instructions were presented on the computer screen. The experimental sequence was as follows. First, participants rated the five Chinese characters (time point “pre”), followed by condition-appropriate valenced images (flowers or insects). Second, they completed either the flowers IAT or insects IAT. Finally, they rated the Chinese characters and the valenced images a second time (time point “post”).

Results

Data Processing and Analysis

Differences in reaction times for each participant between the IATs’ two response patterns (i.e., Blocks 3 vs. 6 and 4 vs. 7) were quantified using the *D1* scoring algorithm (Greenwald, Nosek, & Banaji, 2003). These were coded so that positive scores refer to more positive automatic evaluations of the Chinese characters. No exclusions were made based on overall accuracy within the test blocks ($M = 92.8\%$, $SD = 4.6$) or latency performances ($M = 659$ ms, $SD = 92$) in the test blocks (i.e., Blocks 3, 4, 6 and 7). The ratings of the Chinese characters, flowers, and insects were each reduced to a mean score for each time point, and change scores were then calculated.

Given the directional nature of the manipulation checks and hypothesis, we elected, *a priori*, to employ one-tailed comparisons in all *t*-tests while retaining an α value of .05. Two manipulation checks were performed. First, our analysis relies on the assumption that the flowers stimuli would be rated as more positive than the insect stimuli. A Welch’s independent *t*-test indicated that participants rated the flowers stimuli as more positive than the insects stimuli, $t(50) = 14.52$, $p < .001$, $d = 4.02$, 95% CI [3.03, 5.03]. Second, our analysis relies on the assumption that that *D1* scores on the flowers and insects IATs would differ as a function of their valenced images. An independent *t*-test

Table 1. Evaluations of the Chinese characters in Experiments 1 and 2

	Mean (SD)		<i>p</i>	<i>d</i>	95% CI	
	Flowers condition	Insects condition			Low	High
Experiment 1						
Ratings	-0.29 (0.53)	0.12 (0.40)	.001	0.87	0.27	1.46
Experiment 2						
Ratings	-0.15 (0.72)	0.15 (0.74)	.02	0.42	0.01	0.82
SC-IATs	-0.10 (0.35)	0.11 (0.33)	.001	0.62	0.20	1.03

Note. All analyses performed on difference scores before and after the IAT.

suggested that participants demonstrated larger *D1* scores on the insects IAT than the flowers IAT, $t(49.82) = -2.24$, $p = .015$, $d = 0.62$, 95% CI [0.03, 1.20].

Most crucially, the change in liking of the Chinese characters due to the completion of the IAT was compared between groups using an independent *t*-test on the rating change scores. This demonstrated significant differences of large effect size, $t(46.26) = -3.14$, $p = .001$, $d = 0.87$, 95% CI [0.27, 1.46]. Self-report ratings of the Chinese characters became more negative after the flowers IAT ($M_{\text{change}} = -0.29$, $SD = 0.53$), whereas ratings became more positive after the insects IAT ($M_{\text{change}} = 0.12$, $SD = 0.40$).¹ See Table 1 for summaries of the results of Experiments 1 and 2. See Figure 1 for plots of raw data, distributions, and inferential information. In sum, we observed an analogical learning effect.

Prompted by reviewers' comments, we used post hoc exploratory tests to examine whether the effect was influenced by the order in which participants completed the IAT blocks (i.e., IAT block order). An analysis of variance (ANOVA) with rating change scores entered as DV and condition and block order entered as IVs did not reveal a significant interaction between condition and block order, $F(1, 48) = 0.03$, $p = .95$, $\eta^2 < 0.01$.

Experiment 2

In a second experiment, we sought to replicate and extend the effects observed in Experiment 1 under more stringent conditions. Several changes were made to the design and methods in order to attempt to increase the strength of the conclusions that could be drawn. First, evaluations of the Chinese characters were also assessed using the Single Category Implicit Association Test (SC-IAT: Karpinski & Steinman, 2006). The use of an implicit measure was intended to limit demand characteristics in the sample

and to demonstrate that the evaluative learning effects have features of automaticity (e.g., can also be observed in a task that does not require the intention to evaluate the Chinese characters). The SC-IAT is a derivative of the IAT that includes only one target category (i.e., Chinese characters) and the two attribute categories (i.e., positive and negative). Second, the self-report ratings of the valenced images were taken at the end of the experiment rather than both before and after the IAT. In this manner, the only task in which the participants encountered the valenced images was the IAT, thus limiting the potential for unintended learning contexts within the procedure (e.g., tacit contrasting of the Chinese characters and valenced images within the rating scales). Third, in order to establish that the symbols were of unknown meaning, participants were asked if they understood any of the Chinese symbols at the start of the experiment. Finally, the rating scales were changed from a 1-5 scale to a 1-7 scale to allow for greater variance. Similar to Experiment 1, it was hypothesized that the Chinese characters would acquire the opposite valence to the target category that they were contrasted with in the IAT, on both the self-report and implicit measures.

Method

Participants

In contrast to Experiment 1, the sample size of $N = 100$ was determined by an *a priori* power analysis. This sample could be expected to provide sufficient power to detect change score groups differences of medium effect size (given $\alpha = .05$, $1 - \beta = .80$, $d = 0.5$; $N \geq 98$). One hundred and four students (79 women and 25 men, $M_{\text{age}} = 21.51$, $SD = 3.24$) took part in exchange for €5. All participants were fluent Dutch speakers and provided written informed consent prior to participation. Four participants provided only partial data and were excluded.

¹ All results were robust under the alternative strategy of employing an analysis of covariance (ANCOVA) with post scores as DV and condition as IV with pre scores as a covariate. Experiment 1 ratings: $F(1, 49) = 11.19$, $p = .002$, $\eta^2 = 0.12$, 90% CI [0.05, 0.33]; Experiment 2 ratings: $F(1, 97) = 9.17$, $p = .003$, $\eta^2 = 0.08$, 90% CI [0.02, 0.18]; Experiment 2 SC-IATs, $F(1, 97) = 6.18$, $p = .01$, $\eta^2 = 0.06$, 90% CI [0.01, 0.15]; combination analysis, $F(1, 149) = 18.29$, $p = .00003$, $\eta^2 = 0.09$, 90% CI [0.04, 0.19].




Figure 1. Change in evaluations of the Chinese characters. Pre–post change scores between the flowers IAT and insects IAT conditions across Experiments 1 and 2. Points represent individual participants' change scores, curved lines represent Gaussian-smoothed kernel density estimations, black lines represent group means, and gray bars represent 95% confidence intervals.

Measures

Understanding of the Chinese characters was assessed using a single item that asked the participant whether they understood any or all characters using a yes/no response format. Similar to Experiment 1, participants rated the Chinese characters and the valenced images, this time using a 1–7-point scale. The IATs were identical to Experiment 1. The SC-IAT was employed to assess automatic evaluations of the Chinese characters in the absence of a valenced image contrast category. Similar to the IAT, the SC-IAT presented stimuli in the middle of the screen and required participants to categorize them in line with category labels that were presented at the top of the screen. These labels were mapped onto the left and right response keys (E and I). The SC-IATs consisted of 3 blocks. Participants first completed a practice block of 20 trials that presented the two attribute category stimuli only (i.e., positive and negative words), followed by two test blocks that presented both target (Chinese characters) and attribute stimuli (valenced words). Whereas the IAT typically presents equal numbers of trials for each stimulus categories on each test block, the SC-IAT presented an unequal number in order to roughly to balance the number of trials requiring left and right key responses while employing only three stimulus categories. Specifically, one test block presented 20 Chinese characters trials, 20 positive words trials, and 30 negative words trials; the other test block presented 20 Chinese characters trials, 30 positive words trials, and 20 negative words trials. Progression to the next trial was contingent on providing a

correct response, and accuracy feedback was presented via a red X on incorrect trials. Brief explanatory instructions identical to the IAT were presented to the participant before each block. Block order presentation of both the IAT and SC-IATs was counterbalanced between participants, as was the congruence between the IAT and SC-IAT block orders.

Procedure

Participants were tested individually in experimental cubicles and were randomly assigned to the flowers condition or insects condition. All instructions were presented on the computer screen. The experimental sequence was as follows: participants completed the Chinese characters recognition scales, the SC-IAT, and rated the Chinese characters (but not the valenced images). Time point for this SC-IAT and rating scales was thus “pre” the IAT. Participants then completed either the flowers IAT or insects IAT, and then the rating scales and SC-IAT a second time. Time point for this SC-IAT and rating scales was thus “post” the IAT. Finally, each participant was asked to rate both the flowers and insects stimuli.

Results

All participants responded on the recognition test that they did not understand any of the Chinese characters. As such, no exclusions were made on the basis of this criterion.

Identical to Experiment 1, ratings of the 5 exemplars of the characters, the images of flowers and the images of insects, respectively, were reduced to one mean score for each. Reaction times and accuracies on the IATs' test blocks (i.e., Blocks 3 vs. 6 and 4 vs. 7) were each transformed to a single *D1* score for each participant. Responses on the SC-IATs' test blocks (i.e., Blocks 2 vs. 3) were also used to calculate a *D1* score for each SC-IAT. Finally, change scores were calculated. Both the IAT and the SC-IATs were coded so that positive scores refer to more positive evaluations of the Chinese characters (i.e., IATs: Chinese characters–positive/flowers–negative or Chinese characters–positive/insects–negative effects, respectively; SC-IATs: Chinese characters–positive/negative effects). No exclusions were made based on accuracy ($M_{SC-IATpre} = 94.6\%$, $SD = 3.6$; $M_{IAT} = 93.5\%$, $SD = 4.9$; $M_{SC-IATpost} = 92.9\%$, $SD = 4.7$) or latency performances ($M_{SC-IATpre} = 642$ ms, $SD = 108$; $M_{IAT} = 669$ ms, $SD = 98$; $M_{SC-IATpost} = 622$ ms, $SD = 79$) in the test blocks of either the IAT or SC-IATs.

As in Experiment 1, given the directional nature of the manipulation checks and hypotheses, we elected, *a priori*, to employ one-tailed comparisons in all tests (while retaining an α value of .05). Similar to Experiment 1, two manipulation checks were tested. A dependent *t*-test suggested that participants rated the flowers stimuli as more positive than the insects stimuli, $t(99) = 26.91$, $p < .001$, $d = 4.08$, 95% CI [3.59, 4.58]. An independent *t*-test demonstrated larger IAT *D1* scores on the insects IAT than the flowers IAT, $t(97.53) = -4.11$, $p < .001$, $d = 0.82$, 95% CI [0.40, 1.24].

Changes in self-reported liking of the Chinese characters due to the completion of the IAT were compared between groups using an independent *t*-test on the rating change scores. This demonstrated significant differences of medium effect size, $t(97.70) = -2.07$, $p = .02$, $d = 0.42$, 95% CI [0.01, 0.82]. Self-report ratings of the Chinese characters became more negative after the flowers IAT ($M_{change} = -0.15$, $SD = 0.72$), whereas ratings became more positive after the insects IAT ($M_{change} = 0.15$, $SD = 0.74$).

Differential changes in automatic evaluations of the Chinese characters due to the completion of the IAT between the two groups were assessed using an independent *t*-test on the SC-IAT change scores. This demonstrated significant differences of medium effect size, $t(96.77) = -3.07$, $p = .001$, $d = 0.61$, 95% CI [0.20, 1.03]. Automatic evaluations of the Chinese characters on the SC-IATs became more negative after the flowers IAT ($M_{change} = -0.10$, $SD = 0.35$), whereas automatic evaluations became more positive after the insects IAT ($M_{change} = 0.11$, $SD = 0.33$).

Post hoc exploratory tests examined whether the effect was influenced by the order in which participants completed the IAT blocks (i.e., IAT block order). First, an

ANOVA with rating change scores entered as DV and condition and IAT block order entered as IVs did not reveal a significant interaction between condition and block order, $F(1, 96) = 0.14$, $p = .71$, $\eta^2 < .01$. A second ANOVA was conducted with SCI-IAT change scores entered as DV and condition and IAT block order entered as IVs. Again, this did not demonstrate a significant interaction between condition and block order, $F(1, 96) = 1.85$, $p = .18$, $\eta^2 = .02$.

Experiment 3

In a third experiment, we sought to explore the analogical nature of the IAT's learning effect. If the effect were analogical, it would necessarily involve not only the Chinese characters and the valenced contrast category (flowers or insects) but also a second pair of stimuli, such as the attribute categories (e.g., positive and negative words). This experiment therefore manipulated the relation between the attribute categories in order to observe whether this undermined the IAT's learning effect. Participants in a first condition completed an "opposite attributes" IAT that used positive and negative words as attribute stimuli, as in the previous experiments. This was intended to invite a relatively stronger analogy among the four categories: "Chinese characters:insects::positive:negative." Participants in a second condition completed a "non-opposite attributes" IAT that employed nonwords as attribute stimuli (e.g., Niffites vs. Luupites). This was intended to invite a relatively weaker analogy among the categories: "Chinese characters:insects::Niffites:Luupites." Of course, any two categories of stimuli can be related as distinct from one another. In one sense, the IAT still required the participants to relate Niffites and Luupites as opposites by responding to them using opposite keys. However, it seems fair to say that the opposition between positive and negative is more extreme than that between Niffites and Luupites because the former is based not only on current task requirements (i.e., responding to them with opposite keys) but also on a long learning history. Hence, if our effects are analogical in nature, the learning effect should be stronger in the former than in the latter condition. In order to increase the evidential weight of the data, the design, sample size, inclusion/exclusion criteria, and analytic strategy for this study were preregistered.

Method

Participants

Sample size was determined *a priori*. Due to the difficulty of estimating power for linear mixed-effects models, we employed a power analysis for a traditional fixed-effects

ANOVA using G*power. Although not mutually substitutable, this provided a broad heuristic and informed our decision to set sample size to ≥ 440 after exclusions.² As we expected a relatively small effect between the opposite attributes and non-opposite attributes conditions, we employed a relatively large sample size. In contrast to the previous studies, we collected data online from the prolific.ac platform: 449 individuals (223 women, 222 men, and 4 who identified as a third category, $M_{\text{age}} = 34.90$, $SD = 11.45$) took part in exchange for £0.90. All participants provided informed consent prior to participation. More stringent inclusion criteria were used relative to the previous experiments given the use of online data collection. Inclusion criteria were 90% approval rate in previous studies on the platform, age 18–65 years, English as a first language, and no participation in previous studies by our research group. Exclusion criteria were incomplete data, > 10% of reaction times < 300 ms on the IAT's test blocks (Blocks 3, 4, 6, and 7). Twenty-two individuals were excluded on this basis. Whereas the previous experiments also employed an accuracy exclusion criterion, none were applied in the current study on the basis that the non-opposite attributes IAT, which included nonwords, was of unknown difficulty relative to more typical IATs.

Measures

Measures were highly similar to the previous experiment but were presented in English rather than Dutch. Two additional versions of the training IAT were created. These replaced the positive and negative attribute stimuli and their category labels with categories of nonwords (Niffites: cellanif, eskannif, lebunnif, zallanif, otrannif, borrinif; vs. Luupites: meesolup, neenolup, omellup, wenalup, turalup, loomalup). These nonwords have been employed as neutral stimuli in previous evaluative learning studies (Van Dessel, De Houwer, Gast, & Smith, 2015). Four versions of the IAT were therefore used in a 2×2 between-groups factorial design: (1) flowers IAT with opposite attributes, (2) insects IAT with opposite attributes, (3) flowers IAT with non-opposite attributes, and (4) insects IAT with non-opposite attributes.

Procedure

Participants completed one random training IAT, followed by the Chinese characters SC-IAT and rating scales. Unlike the previous experiments, participants were not asked to evaluate the stimuli before completing the IAT. This change was made on the basis of the idea that requiring

participants to evaluate stimuli at multiple time points might lead to artifactual learning effects under some circumstances (see Gawronski, Gast, & De Houwer, 2015 for more details). The order of the SC-IAT and rating scales was counterbalanced between participants, as was the SC-IAT block order. Based on the results of the previous experiment, the congruence between the IAT block order and SC-IAT block order was not manipulated between participants (i.e., participants completed tasks in the same block order).

Results

Analytic Strategy

We expected any evaluative learning effect to be small and therefore chose to employ a more powerful analytic strategy. While *D* scoring of (SC-)IAT data is common, it reduces each individual's performance on the task down to a single score. In contrast, by examining reaction time data directly, we were able to make use of all 140 reaction times provided by each participant's SC-IAT test blocks, thereby increasing power. Of course, this increase in power must be balanced with two functions usually accomplished with *D* scoring: (1) an acknowledgment of the nonindependence of the multiple reaction times produced by each participant, and (2) controlling for differences in general responding speed between participants (Greenwald et al., 2003). Both points were addressed by using linear mixed-effects models, which allow for the modeling of both experimentally manipulated fixed-effects (e.g., SC-IAT block, IAT contrast category condition, IAT attribute condition) and also sources of unknown variation due to random effects (e.g., general responding speed; Baayen, Davidson, & Bates, 2008; Bates, Mächler, Bolker, & Walker, 2015). For the sake of clarity and simplicity, only the results of effects that test specific hypotheses will be reported here. Details of other fixed main and interaction effects and random effects can be found in the supplementary materials on OSF. All steps in this analytic strategy were preregistered, and any diverges from the preregistration are noted or described as exploratory. In all models, confidence intervals and *p* values were approximated using the Wald method, which has been shown to be acceptable for larger samples such as ours (Luke, 2016).

Data from the SC-IATs' test blocks (i.e., Blocks 3 vs. 4) were first processed to remove outliers and increase normality. Reaction times < 300 ms were trimmed and then

² Parameters used in G*power: $f = 0.25$, $\alpha = .05$, power = .95, numerator $df = 1$, groups = 4: suggested $N \geq 206$. This suggested sample size was then doubled on the basis that our planned analytic strategy (i.e., multiple testing of the opposition category groups, first separately and then together) and to allow for an unknown attrition rate. Two preregistrations were made: The first incorrectly specified the sample size. This was corrected in a second preregistration. This was done during data collection before analyses had been conducted.

reciprocally transformed to increase normality (Ratcliff, 1993; Whelan, 2008). It should be noted that reciprocal of reaction time can be referred to as “speed” (i.e., responses per second vs. seconds per response). Reciprocal transformations therefore have greater interpretability than other transformations (e.g., log). For example, participants may equally be described as having either “lower reaction times” or “higher speed” in one SC-IAT block relative to the other. Outliers were defined as $> 2.5 SD$ from the mean and excluded. On this basis 1.74% of data points were excluded. No transformations or exclusions were applied to the rating data. Data were then submitted to a series of linear mixed-effects models.

Automatic Evaluations

We began by assessing whether the evaluative learning effect observed on the SC-IAT in the previous experiments was replicated. This was done using only data from the opposite attributes conditions. In the context of our mixed model, this referred to differences in responding speed being predicted by the interaction between the IAT conditions (whether Chinese characters were contrasted with flowers vs. insects) and the SC-IAT block (differences in responding speed between the SC-IAT blocks). In order to control for differences in general responding speed and acknowledge the nonindependence of the reaction times generated by each participant, we included a random slope intercept for participant. Finally, in order to model differences in the magnitude of the SC-IAT effect even between individuals in the same condition, we allowed for a random slope for SC-IAT block. The inclusion of this random slope, which is not directly required to test the hypotheses, made this a relatively conservative approach (Barr, Levy, Scheepers, & Tily, 2013). This model can be specified more precisely using the Wilkinson notation:

$$\text{speed} \sim \text{SCIATblock} * \text{IATcontrastCondition} \\ + (\text{SCIATblock} | \text{participant})$$

We first tested whether we could replicate the learning effect when using oppositional attribute categories (positive vs. negative). The presence of a learning effect was demonstrated by a significant two-way interaction between SC-IAT block and IAT contrast condition ($B = 0.016$, 95% CI = [0.007, 0.024], $\beta = 0.034$, 95% CI = [0.016, 0.052], $p = .0003$). Inspection of the estimated marginal means demonstrated that the effect was in the expected direction: When participants had previously completed the flowers IAT, they were faster (had higher speed) when categorizing Chinese characters using the same key as negative words ($M = 1.71$, 95% CI = [1.67, 1.75]) relative to positive words ($M = 1.69$, 95% CI = [1.65, 1.72]), whereas, when

they had previously completed the insects IAT, they were faster to categorize Chinese characters using the same key as positive words ($M = 1.72$, 95% CI = [1.68, 1.77]) relative to negative words ($M = 1.69$, 95% CI = [1.64, 1.73]). Results therefore suggest that the Chinese characters acquired the opposite valence to the images they were contrasted with in the IAT, thus replicating the evaluative learning effect found in the previous experiments.

Next, we assessed whether the learning effect was present when employing nonword attribute stimuli. The same mixed model specification was employed as in the analysis of the condition with oppositional attribute stimuli. As expected, the interaction effect was nonsignificant ($B = 0.005$, 95% CI = [−0.003, 0.013], $\beta = 0.011$, 95% CI = [−0.007, 0.029], $p = .23$). Inspection of the estimated marginal means demonstrated that participants responded with comparable speed in the two conditions (flowers IAT condition: Chinese characters–negative block, $M = 1.69$, 95% CI = [1.65, 1.73], Chinese characters–positive block, $M = 1.68$, 95% CI = [1.64, 1.72]; insects IAT condition: Chinese characters–negative block, $M = 1.67$, 95% CI = [1.63, 1.7], Chinese characters–positive block, $M = 1.67$, 95% CI = [1.63, 1.71]). Results therefore provide no evidence for the evaluative learning effect when the IAT’s structural analogy was undermined.

The previous analyses suggest that there was evidence for the effect in the opposite attributes condition and no evidence for the effect in the non-opposites condition. However, in order to directly test whether the evaluative learning effect is undermined when attribute stimuli were non-opposites relative to opposites, these two conditions should be directly compared. A third model was created using data from both the IAT attribute conditions and entering this condition into the model:

$$\text{speed} \sim \text{SCIATblock} * \text{IATcontrastCondition} \\ * \text{IATattributeConditions} + (\text{SCIATblock} | \text{participant})$$

We observed a marginally significant three-way interaction between SC-IAT block, IAT contrast condition, and IAT attributes condition ($B = -0.005$, 95% CI = [−0.011, 0.001], $\beta = -0.012$, 95% CI = [−0.024, 0.001], $p = .07$). Inspection of the marginal means (above) indicated that this effect was in the expected direction (see Figure 2 for estimated marginal means). We also constructed an additional, but not preregistered, exploratory model that was less conservative in that it removed the random slope for SC-IAT block. This model therefore assessed differences in the SC-IAT effect between conditions, but did not attempt to also model differences in the SC-IAT effect at the individual level. Using this model, a significant three-way interaction effect was found ($B = -0.005$, 95%




Figure 2. Estimated marginal means of response speed on the SC-IAT blocks between conditions in Experiment 3.

Table 2. Evaluations of the Chinese characters in Experiment 3

IAT contrast category	IAT attribute stimuli	Mean (SD)	
		Ratings	SC-IAT <i>D1</i>
Flowers	Non-opposites	4.24 (0.77)	−0.04 (0.29)
Insects	Non-opposites	4.39 (0.88)	0.01 (0.27)
Flowers	Opposites	4.23 (0.88)	−0.05 (0.27)
Insects	Opposites	4.42 (0.83)	0.07 (0.26)

Note. Descriptive statistics for mean ratings and SC-IAT *D1* scores are provided for reader familiarity.

CI = [−0.009, −0.002], $\beta = -0.012$, 95% CI = [−0.019, −0.005], $p = .001$.³

Based on a reviewer suggestion, we calculated SC-IAT *D1* scores (see Table 2 for descriptive statistics) and Bayesian analysis was used to allow us to quantify the evidence for the null hypothesis. We compared a fixed-effects ANOVA on *D1* scores with a Bayesian equivalent of the above mixed-effects analysis on transformed reaction time data, both of which were conducted using the R package brms (Bürkner, 2017). A default prior was employed on scaled and centered data (normal distribution with $M = 0$, $SD = 1$; Gelman, Lee, & Guo, 2015). Results from the fixed-effects ANOVA on SC-IAT *D1* scores provided no credible evidence in favor of either the null or alternative hypotheses for the interaction between IAT contrast condition and attribute condition ($\beta = 0.25$, 95% HDI = [−0.12, 0.62], $BF_{10} = 0.47$). However, results from the mixed-effects model on reciprocally transformed reaction time

data provided moderate evidence in favor of the alternative hypothesis ($\beta = 0.10$, 95% HDI = [0.04, 0.15], $BF_{10} = 4.49$). Both results were robust to the use of a wider prior (i.e., normal distribution with $M = 0$, $SD = \sqrt{2}$; fixed-effects analysis: $\beta = 0.25$, 95% HDI = [−0.10, 0.60], $BF_{10} = 0.68$; mixed-effects analysis: $\beta = 0.10$, 95% HDI = [0.04, 0.15], $BF_{10} = 4.96$). While the results from these two analytic strategies might at first seem to be at odds with one another, it is worth recalling that Bayes Factors are contingent on the data at hand. When using *D1* scoring, given fewer data points and a small effect size, the key interaction term was of uncertain utility in characterizing the data. In contrast, when all 140 data points from each participant were modeled, given a far greater number of data points, the best characterization of the data included the key interaction term. These findings were therefore consistent with our preregistered decision to employ mixed-effects models rather than *D1* scores, given our

³ We also examined the influence of IAT block order for Experiment 3. Although we did find evidence for an effect of IAT block order, we do not report the analyses here because IAT and SC-IAT block order were not manipulated independently in Experiment 3. Hence, no strong conclusions can be made regarding whether the effect of IAT block order is actually due to IAT block order or to SC-IAT block order.

sample size choices and expectations that the size of the effect would be relatively small. On the basis that the more powerful analytic strategy suggested that the inclusion of the key interaction term provided the better characterization of the data, we concluded that the data were supportive of our hypothesis that the learning effect was analogical in nature.

Self-Reported Evaluations

A comparable set of models was used to assess the same hypotheses within the self-report ratings. Each of these included a random intercept for participant to acknowledge the nonindependence of the multiple ratings produced by each participant. In the opposite attributes IAT conditions, we did not observe a significant main effect for IAT contrast condition ($B = -0.097$, 95% CI = $[-0.215, 0.022]$, $\beta = -0.088$, 95% CI = $[-0.196, 0.020]$, $p = .11$). Participants rated the Chinese characters similarly whether they previously completed the flowers IAT ($M = 4.23$, 95% CI = $[4.07, 4.39]$) or insects IAT ($M = 4.42$, 95% CI = $[4.25, 4.59]$). Results therefore do not replicate those found in the previous experiments: No evaluative learning effect was found on the self-report measures. An analysis of data from the non-opposite attributes conditions also did not reveal evidence that ratings differed between the IAT contrast category conditions ($B = -0.075$, 95% CI = $[-0.184, 0.033]$, $\beta = -0.068$, 95% CI = $[-0.166, 0.030]$, $p = .17$). Participants rated the Chinese characters similarly regardless of whether they previously completed the flowers IAT ($M = 4.24$, 95% CI = $[4.08, 4.4]$) or insects IAT ($M = 4.39$, 95% CI = $[4.24, 4.54]$). A direct comparison of the learning effect in both attribute conditions did not provide evidence for a difference in the effect in between conditions ($B = 0.011$, 95% CI = $[-0.069, 0.091]$, $\beta = -0.010$, 95% CI = $[-0.063, 0.082]$, $p = .79$).

Similar to the analysis of SC-IAT data, based on reviewer suggestions, we calculated mean ratings for each participant (see Table 2 for descriptive statistics) and compared Bayes Factors between a fixed-effects ANOVA on mean rating data with a mixed-effects analysis on the individual ratings. Results from both provided moderate evidence in favor of the null hypotheses for the interaction between IAT contrast condition and attribute condition (fixed-effects analysis: $\beta = 0.05$, 95% HDI = $[-0.32, -0.42]$, $BF_{10} = 0.20$; mixed-effects analysis: $\beta = 0.04$, 95% HDI = $[-0.25, 0.32]$, $BF_{10} = 0.15$). Results from the fixed-effects analysis were not robust to the use of a wider prior ($\beta = 0.05$, 95% HDI = $[-0.26, 0.35]$, $BF_{10} = 0.68$), whereas the mixed-effects analysis was robust ($\beta = 0.04$, 95% HDI = $[-0.24, 0.31]$, $BF_{10} = 0.20$). Similar to the analysis of the SC-IAT data, given the greater power of the mixed model analysis, we therefore concluded that there was evidence against the presence of an evaluative learning effect in Experiment 3.

General Discussion

Results from three experiments demonstrate that the IAT functions as an evaluative learning context and does so in a previously unappreciated manner. Changes in liking of the Chinese characters as a result of being contrasted with valenced images in the IAT were observed in all three experiments. This effect was found on both self-report ratings (Experiments 1 and 2, but not Experiment 3) and an implicit measure (Experiments 2 and 3, no implicit measure included in Experiment 1). Importantly, and in line with the hypotheses, the Chinese characters acquired the opposite valence to that of the target category they were contrasted with in the IATs.

The Analogical Nature of the Effect

All three experiments provide evidence for an evaluative learning effect, with consistent results on the implicit measures and more mixed results on the self-report measures. However, the nature of the effect requires consideration. Importantly, the procedural setup of the IAT and the fact that the Chinese characters acquired the opposite rather than same valence to the stimulus they were contrasted with make the effect very difficult to explain as an instance of learning via established environmental regularities, such as stimulus pairing or intersecting regularities (De Houwer et al., 2013; Hughes et al., 2016). For example, the effect does not qualify as an instance of intersecting contingencies given that the contingencies involving Chinese characters (a) intersected an equal number of times with contingencies involving the positive and negative attribute stimuli, and (b) never intersected with the contingencies involving the valenced images (flowers or insects).

As pointed out by the reviewers, it could be the case that one block of intersections (i.e., one IAT block) had more impact on the learning effect than the other. For instance, the consistent block in which flowers were assigned to the same key as positive stimuli might have had more impact than the inconsistent block in which flowers were assigned to the same key as negative stimuli. The reviewers noted that this selective impact of the consistent block might be particularly likely if participants complete this block before the inconsistent block because block order is known to affect ease of performance (i.e., the advantage for the consistent over inconsistent block is bigger when the consistent block comes first). Exploratory analyses of our data, however, provided little evidence for effects of IAT block order. Also, note that a bigger impact of the consistent IAT blocks does not exclude an explanation in terms of analogical learning. For instance, it is possible that consistent blocks are more likely to lead to the construction of an analogy (e.g., Chinese characters:negative::flowers:

positive vs. Chinese characters:positive::flowers:negative). However, this would still imply the involvement of analogical processes. Nevertheless, more research is needed to examine the extent to which IAT blocks contribute differentially to the learning effects that we observed.

In the absence of an explanation for the learning effect in terms of established regularities, Experiment 3 sought to provide support for the analogical nature of the effect. If the effect is best described as analogical, it would necessarily involve not only the Chinese characters and the valenced images, but also (a) a second pair of stimuli and (b) the relating of relations among these two pairs (Gentner & Smith, 2013). Within the context of the IAT, the Chinese characters and valenced contrast category (flowers or insects) could notionally form the target pair, and the attribute categories (positive vs. negative words) could form a source pair. If the effect were analogical, the relation between the source pair (the attribute categories) would serve to cue the relation between the target pair (the neutral category and valenced contrast category). That is, the fact that the attribute words are a pair of opposites may signal that the Chinese characters are also opposite to the contrast category (e.g., Chinese characters:insects::positive:negative). Experiment 3 therefore manipulated the nature of the relation between the attribute categories between groups in order to test whether this influenced the learning effect. Results provide evidence for the learning effect in the condition that employed opposite attributes (positive vs. negative words), whereas no evidence was found for the learning effect in the condition that employed non-opposite attribute stimuli (nonwords: Niffites vs. Luupites). The difference between these two conditions was marginally significant in the more conservative, preregistered linear mixed-effects model, although it should be noted that it was highly significant in a less conservative, exploratory model that relaxed the assumptions about the random effect. Results from a comparable Bayesian analysis were congruent with this conclusion, as discussed in Experiment 3's Results section. Because the effect appeared to be influenced by the relation among the attribute stimuli even when the Chinese characters and contrast category were kept constant, we conclude that there is at least initial evidence that the IAT's learning effect can be characterized as being analogical. This has several implications for both learning research and the use of the IAT, which will be discussed further below.

While our results refer specifically to learning via the IAT, it is likely that other tasks might induce a comparable effect. For example, Scherer and Lambert (2009) demonstrated that relatively neutral images presented within an evaluative priming task are evaluated within the task as either positive or negative depending on whether they are presented alongside positively or negatively valenced

images. It should be pointed out that Scherer and Lambert's results are not necessarily learning effects, because no assessment was made of whether the valence changes during the task were observed to change behavior following the task. Furthermore, it is equally possible that the effects observed by Scherer and Lambert are also instances of an analogical effect like that described here. Like the IAT, the evaluative priming task includes not only the neutral and valenced stimulus categories, but also a second pair or valenced attribute categories. This hypothesis could be investigated in a similar manner as in Experiment 3 within the current article, that is, by manipulating the relation among the attribute categories. Future research might therefore wish to explore whether other tasks, such as evaluative priming, also demonstrate contain structural analogies and can be used to demonstrate analogical learning effects (see also Spellman, Holyoak, & Morrison, 2001).

Implications for Learning

Results are of interest to learning psychology, as they demonstrate that IAT may have utility as a task for inducing learning via analogy. Importantly, research to date on analogy has predominantly focused on the mechanisms of analogical reasoning (e.g., the ability to generate a correct answer given an incomplete analogy; Gentner, 1989; Holyoak & Thagard, 1989) and on how analogical abilities are involved in the development of language (Christie & Gentner, 2014; Gentner & Christie, 2008; Gentner & Namy, 2006). However, relatively less work has focused on learning via analogy, when learning is defined as a change in behavior due to regularities in the environment (De Houwer et al., 2013; although see Bliznashki & Kokinov, 2009; Gentner & Smith, 2013; Ruiz & Luciano, 2011; Stewart, Barnes-Holmes, Hayes, & Lipkens, 2001). This is surprising given that several authors have argued that analogy is central to human adaptability to our environment and our ability to learn about events we have not directly experienced (e.g., Gentner, 2003; Hofstadter, 2001; Holyoak & Thagard, 1995). We therefore suggest that there is a need to directly examine learning about novel concepts via such analogical abilities, including study of the environmental regularities that determine when analogical learning takes place. We suggest that this concept of learning via analogy may have both explanatory and predictive utility. For example, while our observed effects here are not easily interpreted as being due to established environmental regularities, such as stimulus pairing, operant learning, or intersecting regularities (Hughes et al., 2016), they are readily accommodated as instances of learning via analogy. Furthermore, the concept of learning via analogy allows for future predictions about how to arrange the environment to

produce such learning. Learning psychology and behavior change research may therefore benefit from closer consideration of learning via analogy (see also De Houwer, Hughes, & Barnes-Holmes, 2016). In a broader sense, this may contribute to efforts to bridge our understanding of (evaluative) learning and attitude change between traditionally disparate fields of study, such as evaluative conditioning, instructional control, and persuasion (De Houwer & Hughes, 2016).

It is also worth noting that analogy research to date most frequently has used explicit tasks, insofar as participants are explicitly asked to select logically correct responses in order to complete analogies. These include analogy completion tasks (e.g., “start:finish::far:[near, away, travel, farther]”; Sternberg & Nigro, 1980) and relational matching-to-sample tasks (e.g., where individuals can respond to the compound stimulus “square-square” [same] with the compound stimulus “triangle-triangle” [same] or “circle-star” [different]; Christie & Gentner, 2014; Premack, 1983). Our use of the IAT here, in contrast, is a relatively subtle way of arranging the environment so as to allow the participant themselves to construct and employ analogies by making analogical responding task-irrelevant. It may be the case that this induces less reactance in participants (i.e., resistance against the experimental manipulation; Fulcher & Hammerl, 2005) than explicit instructions to learn via analogy. This property of making analogical responding a task-irrelevant goal within the IAT may therefore allow for several new areas of research, such as research into the automaticity of learning via analogy or the automaticity of the deployment of analogical abilities. It should be noted, however, that the structure of the IAT (i.e., pairs of oppositely related categories) likely makes it more amenable to establishing learning via analogies involving equivalence and opposition. Future research might explore whether analogical learning that involves other relations can also be established using IAT-like tasks.

Implications for the IAT

In line with Heisenberg’s observer effect in physics, our results suggest that the act of completing an IAT to some extent establishes or changes the attitudes that it is intended to assess. This may have important implications for its use as a measure of implicit attitudes: Consider a race IAT in which participants categorize faces of black people, white people, positive words, and negative words (e.g., Greenwald et al., 1998, 2009). Existing research suggests that completing such an IAT may influence subsequent behavior toward racial out-groups (Vorauer, 2012), but has not provided a precise account of how the IAT induces learning, and therefore, how this effect could be enhanced or mitigated. The current results suggest that,

for example, if prior to completing a race IAT a white individual possessed (a) strong automatic positive evaluations of white people but (b) no particularly strong automatic evaluations of black people, the act of completing the IAT may serve to establish the analogy that “black people are to white people as negative is to positive” (i.e., two pairs of opposites). Such an individual might subsequently evaluate black people more negatively after completing the task. Future research should therefore examine whether analogical learning within the IAT biases its outcome, or even changes the established attitudes it intends to assess. For example, one could test whether completion of a race IAT serves to increase racial bias, and how lasting any such changes may be (cf. Lai et al., 2016).

Limitations and Suggestions for Future Research

Given the procedural similarity between the IAT and SC-IAT, the use of a SC-IAT as an outcome measure in Experiments 2 and 3 may have influenced the observed learning effects in some way. Future research may therefore wish to employ an assessment task whose procedural properties are more distinct from the IAT, such as evaluative priming (e.g., Spruyt & Tibboel, 2015) or the Affective Misattribution Procedure (Payne, Cheng, Govorun, & Stewart, 2005). Future research might also address whether consistent and inconsistent IAT blocks differentially contribute to the learning effects (i.e., the influence of IAT block order). Effects were observed on the implicit measure across studies, including those that did and did not require the Chinese characters to be rated prior to completing the IAT. As such, this provides some confidence that the learning effects are not merely an artifact of pre-rating or a demand compliance effect. Nonetheless, future work is needed to examine the possibility that the absence of the pre-ratings in Experiment 3 was responsible for the elimination of the learning effect on the self-report measures in that experiment.

Conclusion

Results from three experiments demonstrated that the IAT functions as a training task as well as a testing task and does so in a previously unrecognized manner. More specifically, results suggest that the IAT can function as an analogical learning task. Our results reveal new avenues for research on both learning via analogy and the IAT as a measure of attitudes.

Acknowledgment

Funding was provided by Ghent University Grants O1P05517 to IH and BOF16/MET_V/002 to Jan De Houwer.

References

- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390–412. <https://doi.org/10.1016/j.jml.2007.12.005>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68, 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67, 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Bliznashki, S., & Kokinov, B. (2009). Analogical transfer of emotions. In B. Kokinov, K. Holyoak, & D. Gentner (Eds.), *New frontiers in analogy research* (pp. 45–53). Sofia, Bulgaria: NBU Press.
- Brendl, C. M., Markman, A. B., & Messner, C. (2001). How do indirect measures of evaluation work? Evaluating the inference of prejudice in the Implicit Association Test. *Journal of Personality and Social Psychology*, 81, 760–773. <https://doi.org/10.1037/0022-3514.81.5.760>
- Bürkner, P. C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80, 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Christie, S., & Gentner, D. (2014). Language helps children succeed on a Classic Analogy Task. *Cognitive Science*, 38, 383–397. <https://doi.org/10.1111/cogs.12099>
- De Houwer, J. (2007). A conceptual and theoretical analysis of evaluative conditioning. *The Spanish Journal of Psychology*, 10, 230–241. <https://doi.org/10.1017/S1138741600006491>
- De Houwer, J., Barnes-Holmes, D., & Moors, A. (2013). What is learning? On the nature and merits of a functional definition of learning. *Psychonomic Bulletin & Review*, 20, 631–642. <https://doi.org/10.3758/s13423-013-0386-3>
- De Houwer, J., & Hughes, S. (2016). Evaluative conditioning as a symbolic phenomenon: On the relation between evaluative conditioning, evaluative conditioning via instructions, and persuasion. *Social Cognition*, 34, 480–494. <https://doi.org/10.1521/soco.2016.34.5.480>
- De Houwer, J., Hughes, S., & Barnes-Holmes, D. (2017). Psychological engineering: A functional-cognitive perspective on applied psychology. *Journal of Applied Research in Memory and Cognition*, 6, 1–13. <https://doi.org/10.1016/j.jarmac.2016.09.001>
- Ebert, I. D., Steffens, M. C., Von Stülpnagel, R., & Jelenec, P. (2009). How to like yourself better, or chocolate less: Changing implicit attitudes with one IAT task. *Journal of Experimental Social Psychology*, 45, 1098–1104. <https://doi.org/10.1016/j.jesp.2009.06.008>
- Fulcher, E. P., & Hammerl, M. (2005). Reactance in affective-evaluative learning: Outside of conscious control? *Cognition and Emotion*, 19, 197–216. <https://doi.org/10.1080/02699930441000283>
- Gawronski, B., Gast, A., & De Houwer, J. (2015). Is evaluative conditioning really resistant to extinction? Evidence for changes in evaluative judgements without changes in evaluative representations. *Cognition and Emotion*, 29, 816–830. <https://doi.org/10.1080/02699931.2014.947919>
- Gelman, A., Lee, D., & Guo, J. (2015). Stan: A probabilistic programming language for Bayesian inference and optimization. *Journal of Educational and Behavioral Statistics*, 40, 530–543. <https://doi.org/10.3102/1076998615606113>
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning* (pp. 199–241). New York, NY: Cambridge University Press.
- Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 195–235). Cambridge, MA: MIT Press.
- Gentner, D., & Christie, S. (2008). Relational language supports relational cognition in humans and apes. *Behavioral and Brain Sciences*, 31, 136–137. <https://doi.org/10.1017/S0140525X08003622>
- Gentner, D., & Namy, L. L. (2006). Analogical processes in language learning. *Current Directions in Psychological Science*, 15, 297–301. <https://doi.org/10.1111/j.1467-8721.2006.00456.x>
- Gentner, D., & Smith, L. A. (2013). Analogical learning and reasoning. In D. Reisberg (Ed.), *The Oxford Handbook of Cognitive Psychology* (1st ed., pp. 668–681). New York, NY: Oxford University Press.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, 74, 1464–1480. <https://doi.org/10.1037/0022-3514.74.6.1464>
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85, 197–216. <https://doi.org/10.1037/0022-3514.85.2.197>
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, 97, 17–41. <https://doi.org/10.1037/a0015575>
- Hofstadter, D. R. (2001). Analogy as the core of cognition. In D. Gentner, K. J. Holyoak, & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 499–538). Cambridge, MA: MIT Press.
- Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition*, 15, 332–340. <https://doi.org/10.3758/BF03197035>
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295–355. https://doi.org/10.1207/s15516709cog1303_1
- Holyoak, K. J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press.
- Hughes, S., De Houwer, J., & Perugini, M. (2016). Expanding the boundaries of evaluative learning research: How intersecting regularities shape our likes and dislikes. *Journal of Experimental Psychology: General*, 145, 731–754. <https://doi.org/10.1037/xge0000100>
- Inquisit 4 [Computer software]. (2015). Retrieved from <https://www.millisecond.com>
- Karpinski, A., & Steinman, R. B. (2006). The single category Implicit Association Test as a measure of implicit social cognition. *Journal of Personality and Social Psychology*, 91, 16–32. <https://doi.org/10.1037/0022-3514.91.1.16>
- Lai, C. K., Skinner, A. L., Cooley, E., Murrar, S., Brauer, M., Devos, T., ... Nosek, B. A. (2016). Reducing implicit racial preferences: II. Intervention effectiveness across time. *Journal of Experimental Psychology: General*, 145, 1001–1016. <https://doi.org/10.1037/xge0000179>
- Luke, S. G. (2016). Evaluating significance in linear mixed-effects models in R. *Behavior Research Methods*, 49, 1494–1502. <https://doi.org/10.3758/s13428-016-0809-y>
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The Implicit Association Test at age 7: A methodological and conceptual review. In J. Bargh (Ed.), *Automatic processes in social thinking and behavior* (pp. 265–292). New York, NY: Psychology Press.

- Payne, K., Cheng, C. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89*, 277–293. <https://doi.org/10.1037/0022-3514.89.3.277>
- Peirce, J. W. (2007). PsychoPy: Psychophysics software in Python. *Journal of Neuroscience Methods, 162*, 8–13. <https://doi.org/10.1016/j.jneumeth.2006.11.017>
- Premack, D. (1983). The codes of man and beasts. *Behavioral and Brain Sciences, 6*, 125–136. <https://doi.org/10.1017/S0140525X00015077>
- Prestwich, A., Perugini, M., Hurling, R., & Richetin, J. (2010). Using the self to change implicit attitudes. *European Journal of Social Psychology, 40*, 61–71. <https://doi.org/10.1002/ejsp.610>
- Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin, 114*, 510–532. <https://doi.org/10.1037/0033-2909.114.3.510>
- Ruiz, F. J., & Luciano, C. (2011). Cross-domain analogies as relating derived relations among two separate relational networks. *Journal of the Experimental Analysis of Behavior, 95*, 369–385. <https://doi.org/10.1901/jeab.2011.95-369>
- Scherer, L. D., & Lambert, A. J. (2009). Contrast effects in priming paradigms: Implications for theory and research on implicit attitudes. *Journal of Personality and Social Psychology, 97*, 383–403. <https://doi.org/10.1037/a0015844>
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2012). *A 21 word solution*. Social Science Research Network. Retrieved from <http://papers.ssrn.com/abstract=2160588>
- Spellman, B. A., Holyoak, K. J., & Morrison, R. G. (2001). Analogical priming via semantic relations. *Memory & Cognition, 29*, 383–393. <https://doi.org/10.3758/BF03196389>
- Spruyt, A., & Tibboel, H. (2015). On the automaticity of the evaluative priming effect in the valent/non-valent categorization task. *PLoS One, 10*, e0121564. <https://doi.org/10.1371/journal.pone.0121564>
- Sternberg, R. J., & Nigro, G. (1980). Developmental patterns in the solution of verbal analogies. *Child Development, 51*, 27–38. <https://doi.org/10.2307/1129586>
- Stewart, I., Barnes-Holmes, D., Hayes, S. C., & Lipkens, R. (2001). Relations among relations: Analogies, metaphors, and stories. In S. C. Hayes, D. Barnes-Holmes, & B. Roche (Eds.), *Relational frame theory: A post-Skinnerian account of human language and cognition* (pp. 73–86). New York, NY: Kluwer Academic/Plenum Press.
- Van Dessel, P., De Houwer, J., Gast, A., & Smith, C. T. (2015). Instruction-based approach-avoidance effects: Changing stimulus evaluation via the mere instruction to approach or avoid stimuli. *Experimental Psychology, 62*, 161–169. <https://doi.org/10.1027/1618-3169/a000282>
- Vorauer, J. D. (2012). Completing the Implicit Association Test reduces positive intergroup interaction behavior. *Psychological Science, 23*, 1168–1175. <https://doi.org/10.1177/0956797612440457>
- Whelan, R. (2008). Effective analysis of reaction time data. *The Psychological Record, 58*, 475–482. <https://doi.org/10.1007/BF03395630>

Received February 13, 2017

Revision received April 11, 2018

Accepted April 16, 2018

Published online November 19, 2018

Ian Hussey

Department of Experimental Clinical and Health Psychology

Ghent University

Henri Dunantlaan 2

9000 Gent

Belgium

ian.hussey@ugent.be