

Running Head: Selective generalization of likes and dislikes

Feature-specific attention allocation modulates
the generalization of recently acquired likes and dislikes

Adriaan Spruyt¹, Karl Christoph Klauer², Anne Gast¹, Maarten De Schryver¹, & Jan De Houwer¹

Ghent University¹

Albert-Ludwigs-Universität Freiburg²

IN PRESS: EXPERIMENTAL PSYCHOLOGY

Adriaan Spruyt

Department of Psychology

Ghent University

Henri Dunantlaan 2

B-9000 Ghent, Belgium

TEL: ++32-9-264-86-18

FAX: ++32-9-264-64-89

E-mail: Adriaan.Spruyt@UGent.be

Abstract

We examined whether the generalization of recently acquired likes and dislikes depends on feature-specific attention allocation. Likes and dislikes were established by means of an evaluative-conditioning procedure in which participants were presented with several exemplars of two subordinate categories (e.g., young men vs. old women). Whereas exemplars of one category were consistently paired with negative stimuli, exemplars of the second category were consistently paired with positive stimuli. In addition, we manipulated feature-specific attention allocation for specific stimulus dimensions (e.g., gender vs. age), either during (Experiments 1 and 2) or before the acquisition phase of the experiment (Experiment 3). Both direct and indirect attitude measures revealed a clear impact of this manipulation on attitude generalization. More specifically, only generalization stimuli that were similar to the CSs in terms of the stimulus dimension that was selectively attended to were evaluated in a manner that was congruent with the acquired liking of those CSs.

Feature-specific attention allocation modulates the generalization of recently acquired likes and dislikes

Likes and dislikes govern human behavior (Allport, 1935). The activities that people engage in, the products they buy, their interpersonal behavior, etc. are all determined, at least to some degree, by personal preferences. One way to establish new likes and dislikes is to pair a neutral stimulus with another stimulus that has a clear evaluative meaning. Typically, such a procedure causes the valence of the initially neutral stimulus (hereafter referred as the *Conditioned Stimulus* or *CS*) to shift towards the affective meaning of the positive or negative stimulus with which it was paired (hereafter referred to as the *Unconditioned Stimulus* or *US*). This phenomenon is commonly referred to as the evaluative conditioning (EC) effect, and has now been replicated in an overwhelming number of experiments across a wide range of study domains (for reviews, see De Houwer, Thomas, & Baeyens, 2001; Hofmann, De Houwer, Perugini, Baeyens, & Crombez, 2010).

The aim of the present research was to examine the conditions under which recently acquired likes and dislikes generalize to novel, untrained stimuli (hereafter referred to as *attitude generalization*). Intuitively, attitude generalization must play a pervasive role in everyday life. How else would one be able to interact with novel attitude objects in a meaningful manner? Nevertheless, only a handful of EC studies have been performed to examine this important issue (see Hofmann et al., 2010). Perhaps the most compelling study showing that (recently acquired) attitudes do generalize to novel stimuli was published by Olson and Fazio (2006, Experiment 2). To examine whether the EC paradigm can be exploited as a means to reduce racial prejudice, they presented White participants with two types of EC trials: trials on which pictures of Black individuals were paired with positive USs and trials on which pictures of White individuals were

paired with negative USs. As a result of this procedure, participants later exhibited less negative racial attitudes towards Black individuals as compared to a control condition. Crucially, the stimuli used during the test phase of the experiment were different from those that were used during the EC procedure. That is, the EC effect clearly generalized to other exemplars of the categories “Blacks” and “Whites”. Similar generalization effects, were also reported by a (small) number of other researchers (e.g., Bierley, McSweeney, & Vannieuwkerk, 1985; Till & Priluck, 2000; but see Baeyens, Kaes, Eelen, & Silverans, 1996; Unkelbach, Stahl, & Förderer, 2012).

The focus of the present research is not so much attitude generalization per se, but rather the extent to which attitude generalization depends on feature-specific attention allocation (hereafter referred to as FSAA). We hypothesized that attitude generalization must be confined to generalization stimuli that are similar to the originally trained stimuli in terms of stimulus features that are selectively attended to. Two independent lines of research led us to postulate this hypothesis. A first line of research concerns the vast amount of fear-conditioning studies showing that the generalization of conditioned fear is highly dependent upon the similarity between a novel stimulus and a known entity. Consider, for example, the findings of Lissek et al. (2008). They presented participants with 10 rings of gradually increasing size with one of the extremes serving as a predictor (CS) of a highly uncomfortable electric shock (US). Both fear-potentiated startle data and online self-report ratings revealed a continuous decrease in fear generalization as the presented stimuli became less similar to the CS. Similar findings were also reported by Vervliet and colleagues (e.g., Vervliet, Iberico, Vervoort, & Baeyens, 2011; Vervliet, Vansteenwegen, & Eelen, 2004, 2006; see also Kalish, 1969; Razran, 1949). Although none of these studies included attitude measures, they do suggest that a novel stimulus is likely to be evaluated in the same way as liked/disliked stimuli showing a high degree of visual overlap with

that novel stimulus. Direct empirical support for this idea has also been reported by Fazio, Eiser, and Shook (2004). They first presented participants either with positive or negative outcomes upon approaching a series of target objects and then examined the extent to which newly formed attitudes generalized to new targets. Although negative attitudes generalized more strongly than positive attitudes, Fazio et al. (2004) found attitude generalization in general to depend heavily on similarity. The more the novel targets visually resembled the known targets, the more likely the novel targets were assumed to share the same valence as the known targets (see also Shook, Fazio, & Eiser, 2007).

A second line of research that inspired the present work concerns the impact of FSAA on the *perceived* similarity of two stimuli. In line with the Generalized Context Model (GCM) of classification (Nosofsky, 1984, 1986, 1986; Nosofsky & Palmeri, 1997; see also Medin & Schaffer, 1978), numerous studies have shown that stimuli varying on stimulus dimensions that are selectively attended to are perceived as more dissimilar to each other than stimuli that vary on stimulus dimensions that are not selectively attended to (e.g., Goldstone & Steyvers, 2001; Lamberts, 2000; Nosofsky, 1986). To the extent that attitude generalization depends on the perceived similarity between a trained CS and a generalization stimulus, the GCM of classification thus predicts that attitude generalization must be dependent upon FSAA too.

Interestingly, such a finding would coincide with recent findings obtained by Vervliet, Kindt, Vansteenwegen, and Hermans (2010). In a study on fear generalization, these authors presented participants with a yellow triangle that was predictive of an aversive electric shock (CS+) and a black cross that was never followed by an electric shock (CS-). Crucially, whereas one group of participants was instructed to use the color difference between the CS+ and CS- to predict the occurrence of the electric shock (i.e., hereafter referred to as the Color Group), a

second group of participants was told that the occurrence of the electric shock would depend on the shape difference between the CS+ and the CS- (i.e., hereafter referred to as the Shape Group). Next, in a subsequent measurement phase, fear generalization was examined using two generalization stimuli that were similar to the CS+ in terms of one specific stimulus feature: a yellow rectangle and a blue triangle. In line with the authors' predictions, both online shock-expectancy ratings and skin-conductance responses revealed a selective generalization of conditioned fear: Whereas participants in the Color Group showed more generalized fear towards the yellow rectangle than the blue triangle, the opposite was true for participants in the Shape Group. It should again be noted, however, that the experiment of Vervliet et al. (2010) was designed to study fear generalization and did not include attitude measures. We are therefore the first to examine whether attitude generalization depends on FSAA.

To evaluate this hypothesis, we ran a series of EC studies in which the acquisition phase involved the presentation of a large number of CSs that varied on two stimulus dimensions simultaneously. For example, in Experiment 1, participants were presented either with pictures of young men and old women or with pictures of young women and old men. That is, age and gender were systematically confounded during the acquisition phase. Whereas one CS category (e.g., young men) was always paired with negative USs, the other CS category (e.g., old women) was always paired with positive USs. In addition, we manipulated the extent to which participants assigned attention to specific stimulus dimensions (e.g., gender vs. age), either during (Experiments 1 and 2) or before the acquisition phase of the experiment (Experiment 3). Next, we collected both implicit and explicit (attitude) measures for all CSs and an equal number of generalization stimuli. These generalization stimuli varied on the same two stimulus dimensions as did the CSs, but the correlation between the two dimensions was now reversed. In

Experiment 1, for instance, if pictures of young men and old women were used as CSs for a particular participant, pictures of old men and young women were used as generalization stimuli. Assuming that attitude generalization depends on FSAA, we expected generalization stimuli that were similar to the CSs in terms of the stimulus dimension that was selectively attended to be evaluated in a manner that was congruent with the acquired liking of those CSs. Consider again Experiment 1 as an example. In the attention-to-age condition, we expected participants to evaluate pictures of young women in a negative manner and pictures of old men in a positive manner if, during the acquisition phase of the experiment, pictures of young men were paired with negative USs and pictures of old women were paired with positive USs. In contrast, participants assigned to the attention-to-gender condition were expected to evaluate pictures of young women in a positive manner and pictures of old men in a negative manner, despite being exposed to the same CS-US pairings during the training phase of the experiment. In sum, depending on FSAA, we expected identical learning experiences to lead to completely different generalization effects. In all experiments, EC effects and generalization effects were assessed by means of two different attitude measures. First, classic evaluative ratings were used to capture deliberate and slow evaluations. In addition, we also administered the Affect Misattribution Paradigm (AMP) developed by Payne, Cheng, Govorun, and Stewart (2005) to capture relatively fast and spontaneous evaluations (for more information about the AMP, see Bar-Anan & Nosek, 2012; Imhoff, Schmidt, Bernhardt, Dierksmeier, & Banse, 2011; Payne, Govorun, & Arbuckle, 2008; Payne, Hall, Cameron, & Bishara, 2010; Payne et al., 2013).

Experiment 1

Method

Participants. Participants were 88 students at Ghent University (14 men, 74 women) who received course credit for their participation or were paid €8 for their help in this experiment and another unrelated experiment. All participants were Dutch speakers and had normal or corrected-to-normal vision. All participants gave informed consent before participation. One participant admitted at the end of the experiment that she did not read any of the instructions. Another participant pressed the same key throughout the AMP. The data of these participants were excluded from the analyses. Two additional participants were run to replace the excluded data sets.

Materials. On the basis of norm data collected by of Spruyt, Hermans, De Houwer, and Eelen (2002), we selected 15 positive and 15 negative color pictures to be used as USs (all 512 pixels wide and 384 pixels high). Some of these pictures originated from the International Affective Picture System (IAPS; Lang, Bradley, & Cuthbert, 2001). On a scale ranging from -5 (“very negative”) to + 5 (“very positive”), the mean valence rating of the positive USs was significantly larger than zero, $M = 2.36$, $SE = 0.18$, $t(14) = 12.78$, $p < .001$. Likewise, the mean valence rating of the negative USs was significantly smaller than zero, $M = -3.08$, $SE = 0.20$, $t(14) = -15.34$, $p < .001$. Black-and-white face pictures of old men (8), young men (8), old women (8), and young women (8) were used as CSs (all 384 pixels wide and 512 pixels high). For the AMP, 200 different Chinese pictographs were used as targets. All Chinese pictographs were presented in white and were 256 pixels wide and 256 pixels high.

All stimuli were presented against the black background of a 21-inch computer monitor (100 Hz, 24 bits per pixel, screen resolution 1024×768). An Affect 4.0 program (Spruyt,

Clarysse, Vansteenwegen, Baeyens, & Hermans, 2010) controlled the presentation of the stimuli as well as the registration of the responses. The experiment was run on a Dell Optiplex GX520 computer.

Procedure. Across participants, all stimuli of the four different CS categories (young men, old men, young women, and old women) were presented equally often during the acquisition phase of the experiment. For every individual participant, however, stimuli stemming from just two different CS categories were used as CSs. Either participants were presented with pictures of old men and young women or participants were presented with pictures of young men and old women. Each picture from these categories was presented exactly once, leading to a total of 16 EC trials. The USs that were paired with the CSs were randomly drawn from the complete list of USs, with the restriction that (a) no US could be presented more than once and (b) all stimuli of one CS category would be followed by a US of the same valence category. The assignment of different CS categories to either positive or negative USs was counterbalanced across participants.

CSs were presented for 3000 ms and were then followed by a US that was presented for 2000 ms. The inter-trial interval varied randomly between 1500 ms and 2500 ms. Crucially, to manipulate the extent to which participants assigned attention to specific stimulus dimensions of the CSs, a question concerning the CSs was presented immediately after the offset of each US. In one group of participants (hereafter referred to as the ‘Age Group’), participants were asked to indicate the age of the person shown on the black-and-white face picture by pressing either a left key (old) or a right key (young). In a second group of participants (hereafter referred to as the ‘Gender Group’), participants were asked to indicate the gender of the person shown on the black-and-white face picture, also by pressing either a left key (male) or right key (female). Half

of the participants ($n = 44$) were assigned to the Age Group; the other half of the participants were assigned to the Gender Group. Within the Age Group, 11 participants were randomly assigned to each of the 2 (CSs: young men and old women vs. old men and young women) \times 2 (USs: positive vs. negative) balancing conditions. Participants in the Gender Group were also randomly assigned to each of the 4 balancing conditions. Due to an experimenter error, however, only 10 participants were assigned to the balancing condition in which pictures of young men were presented together with positive USs and 12 participants were assigned to the balancing condition in which pictures of old men were presented together with negative USs. In the two other 2 balancing conditions, the number of participants was 11.

During the measurement phase of the experiment, participants were first asked to provide valence ratings for all pictures shown during the acquisition phase of the experiment (i.e., the CSs) as well as the CS pictures that were not shown during the acquisition phase of the experiment (i.e., the generalization stimuli). All stimuli (16 CSs and 16 generalization stimuli) were presented exactly once in a random order (32 trials in total). Participants were asked to indicate how much they liked each person shown on the pictures. To indicate their evaluation, participants moved a slider on a 21-point rating scale ranging from minus 10 to plus 10 by means of the arrow keys of the computer keyboard.

Next, participants completed a series of AMP trials, modeled after the recommendations of Payne et al. (2005). Each AMP trial started with a 500-ms presentation of a fixation cross. Next, 500 ms after the offset of the fixation cross, either a CS or a generalization stimulus was presented for 75 ms, followed by a blank screen for 125 ms, and then the presentation of a Chinese pictograph for 100 ms. Following the Chinese pictograph, a black-and-white masking

stimulus was presented until a response was registered. The inter-trial interval varied randomly between 500 ms and 1500 ms.

Participants were told that they would see pairs of pictures flashed one after the other, the first one being a black-and-white face picture and the second being a Chinese character. Similar to Payne et al. (2005), the black-and-white face pictures were described as warning signals for the Chinese characters that required no response at all. Instead, participants were asked to focus on the Chinese pictographs and to indicate their visual pleasantness by pressing either a left key (negative) or a right key (positive) of the computer key board. In total, participants completed 192 AMP trials. Each CS and each generalization stimulus was presented exactly six times. For each participant separately, each CS and each generalization stimulus was combined with a unique Chinese pictographs (sampled randomly from the complete list without replacement).

Finally, at the very end of the experiment, all the CSs and generalization stimuli were again presented in an intermixed random order, and participants were asked to indicate, for each picture separately, whether they thought it had been paired with a negative US (left arrow key, coded as -1), a positive US (right arrow key, coded as +1), or not at all (enter key, coded as 0). In line with Gast, De Houwer, and De Schryver (2012), we will refer to this dependent variable as “valence awareness”.

Results

Acquisition effects. In a first step, we subjected each dependent measure to a 2 (Group: Gender vs. Age) \times 2 (CS type: paired with positive USs vs. paired with negative USs) repeated measures ANOVA. The analysis of the valence awareness ratings revealed that participants were able to indicate, at least on average, whether a particular CS had been paired with positive CSs ($M = 0.49$) or negative CSs ($M = -0.46$), $F(1, 86) = 67.22$, $p < .001$, $\eta_p^2 = .44$. This effect was

unaffected by the group factor, $F < 1$. A similar analysis of the valence ratings revealed a significant EC effect: On average, although the effect was numerically small, participants rated CSs that had been paired with positive USs ($M = 2.95$) as more positive than CSs that had been paired with negative USs ($M = 2.21$), $F(1, 86) = 4.27$, $p < .05$, $\eta_p^2 = .05$. Again, this effect was unaffected by the group factor, $F < 1$. The AMP, however, failed to reveal a significant overall EC effect, $F < 1$.

In a second step, we examined whether EC effects were dependent upon valence awareness. Linear mixed-effect analyses were performed to allow for an assessment of valence awareness effects at the item level (see Pleyers, Corneille, Luminet, & Yzerbyt, 2007; Pleyers, Corneille, Yzerbyt, & Luminet, 2009). Items and participants were defined as crossed random-effects terms. Fixed effects were the effect-coded factors CS type (paired with a positive US vs. paired with a negative US), valence awareness (correct valence awareness vs. incorrect valence awareness), as well as their interaction. The mixed-model F tests were computed using the Kenward-Roger's adjusted degrees of freedom solution (see Kenward and Roger, 1997).

Both the valence ratings, $F(1, 1301.90) = 45.56$, $p < .001$, and the AMP data, $F(1, 1321.81) = 48.56$, $p < .001$, revealed a significant interaction between CS type and valence awareness. Follow-up analyses showed that the EC effect in the valence ratings was significantly positive when considering only those CSs that were classified correctly during the valence awareness task (hereafter referred to as *aware CSs*), $M_{CSpos} = 3.36$, $M_{CSneg} = 1.96$, $F(1, 851.53) = 65.49$, $p < .001$. CSs that were classified incorrectly during the valence awareness task (hereafter referred to as *unaware CSs*) also produced a significant EC effect, but in the opposite direction, $M_{CSpos} = 2.05$, $M_{CSneg} = 2.81$, $F(1, 394.37) = 9.33$, $p < .005$. The results obtained with the AMP mimic this data pattern: aware CSs produced a positive EC effect, $M_{CSpos} = 60.47\%$,

$M_{CSneg} = 50.52 \%$, $F(1, 863.73) = 32.93$, $p < .001$, whereas unaware CSs produced a negative EC effect, $M_{CSpos} = 48.87 \%$, $M_{CSneg} = 60.19 \%$, $F(1, 407.43) = 21.69$, $p < .001$.

Generalization effects. To examine whether FSAA during acquisition modulates the generalization of recently acquired affective stimulus information, we first calculated a single generalization index for each of the three dependent measures for each participant. The generalization index is defined as

$$G = G_{\text{gender}} - G_{\text{age}},$$

where G_{gender} is the extent to which participants generalized recently acquired stimulus information along the gender dimension, and G_{age} is the extent to which participants generalized recently acquired stimulus information along the age dimension. Positive numbers thus indicate that generalization along the gender dimension is stronger than generalization along the age dimension. Negative numbers indicate that generalization along the age dimension is stronger than generalization along the gender dimension. The generalization scores for the two stimulus dimensions themselves were defined as

$$G_d = (T_{\text{pos}})_d - (T_{\text{neg}})_d,$$

where d is either the gender or the age dimension, and $(T_{\text{pos}})_d$ and $(T_{\text{neg}})_d$ are the mean scores of all generalization stimuli that were similar to the CSs in terms of stimulus dimension d . As an example, consider the case where, during the acquisition phase of the experiment, pictures of young men were paired with positive USs and pictures of old women were paired with negative USs. To obtain the gender generalization score for the valence ratings, G_{gender} , we subtracted the mean valence rating of all pictures showing young women from the mean valence rating of all pictures showing old men. In contrast, for the calculation of the age generalization score in this example, G_{age} , we subtracted the mean valence rating of all pictures showing old men from the

mean valence rating of all pictures showing young women. In sum, G_{gender} and G_{age} differ in sign only. By subtracting both generalization scores, however, we obtained a single generalization index, G , that is intuitively interpretable. More specifically, when $G > 0$, generalization along the gender dimension was stronger than generalization along the age dimension. Conversely, when $G < 0$, generalization along the age dimension was stronger than generalization along the gender dimension. We thus expected, for each dependent variable (valence ratings, valence awareness ratings, and AMP), G to be positive in the Gender Group and negative in the Age Group.

As can be seen in Table 1, the results confirm our predictions. For each dependent measure (valence awareness ratings, valence ratings, and AMP), the generalization index, G , was positive in the Gender Group and negative in the Age Group. Despite the fact that none of the generalization stimuli was ever shown during the acquisition phase of the experiment, participants were inclined to (incorrectly) point out that (at least some of) the generalization stimuli had been presented together with either a positive or a negative US. More specifically, (incorrect) valence awareness ratings for the generalization stimuli shifted towards the valence of USs that had been paired with CSs that were similar in terms of the stimulus dimension that was selectively attended to. Likewise, participants were inclined to evaluate (novel) generalization stimuli in the same manner as CSs that were similar in terms of the stimulus dimension that was selectively attended to, both at the explicit level (valence ratings) and the implicit level (AMP). Both for the valence awareness ratings and the valence ratings, the difference between both conditions was statistically reliable, $F(1, 86) = 21.73, p < .001, \eta_p^2 = .20$, and, $F(1, 86) = 4.45, p < .05, \eta_p^2 = .05$, respectively. For the AMP, however, the contrast between the Gender Group and the Age Group just missed conventional significance levels, $F(1, 86) = 3.64, p = .06, \eta_p^2 = .04$.

We also examined whether mean generalization indices were different from zero within each condition. As can be seen in Table 1, this was the case for the valence awareness ratings, both in the Age Group, $t(43) = 2.69, p < .01, d = .41$ and the Gender Group, $t(43) = 3.85, p < .001, d = .60$. For the valence ratings, the generalization index approached significance in the Age Group, $t(43) = 1.79, p = .07, d = .27$, but was statistically unreliable in the Gender Group, $t(43) = 1.16, p = .25, d = .17$. The generalization index for the AMP scores approached marginal significance in the Gender Group, $t(43) = 1.59, p = .12, d = .24$, but was far from significant in the Age Group, $t(43) = 1.17, p = .25, d = .18$.

Finally, we examined whether generalization effects were contingent upon valence awareness. By definition, generalization stimuli are never presented together with a USs during the acquisition phase. Valence awareness was therefore treated as a between-subjects factor for this analysis (accurate vs. inaccurate valence awareness). Remember that we used 16 different CSs for each participant. Eight of these CSs were paired with a positive US whereas the other eight were paired with a negative US. Valence awareness ratings ranged from -1 (paired with a negative US) over 0 (not shown during the EC phase) to +1 (paired with a positive US). Participants were classified as having accurate valence awareness when the average valence awareness rating across all CSs paired with a positive USs exceeded the average valence awareness rating across all CSs paired with a negative USs. On the basis of this criterion, 18 participants were identified as being unable to correctly point out, at least on average, whether a particular CS category had been paired with positive or negative USs. Valence awareness had no impact whatsoever on the selective generalization effects captured by the valence ratings and the AMP, $F < 1$. The selective generalization effect in the valence awareness ratings, however, did depend upon this factor, $F(1, 84) = 13.29, p < .005, \eta_p^2 = .14$. Participants who had accurate

valence awareness exhibited a significant selective generalization effect, $F(1, 66) = 38.80$, $p < .005$, $\eta_p^2 = .37$. In the remaining subset of participants, the selective generalization effect in the valence awareness ratings was absent, $F < 1$.

Discussion

The results of Experiment 1 can be summarized as follows. First, despite the fact that each CS was presented only once during the acquisition phase of the experiment, participants acquired accurate knowledge about the contingency between the CSs and the valence of the USs. In addition, our conditioning procedure resulted in a significant overall EC effect in the explicit valence ratings. The AMP, in contrast, failed to reveal a significant overall EC effect, but the anticipated EC effect did show up when considering only those CSs that were classified correctly during the valence awareness task.

Second, each of our dependent measures revealed selective generalization effects in the predicted direction, albeit the effect just missed conventional significance levels in the AMP. More specifically, attitude generalization along the gender dimension was more pronounced in participants who were encouraged to focus their attention on the gender dimension during the acquisition phase of the experiment. Conversely, attitude generalization along the age dimension was more pronounced in participants who were encouraged to focus their attention on the age dimension. While this data pattern clearly demonstrates that attitude generalization depends on FSAA, one might object that not all generalization indices were statistically different from zero within each condition. It should be emphasized, however, that the generalization indices reflect the extent to which generalization was more pronounced along a particular dimension *relative to* another dimension. Absolute G values thus provide little, if any, information concerning the absolute rate of generalization along each dimension separately.

Third, the EC effects captured by the evaluative ratings and the AMP were dependent upon valence awareness: aware CSs produced positive EC effects whereas negative EC effects were obtained with unaware CSs. Similar effects, albeit captured by valence ratings only, were reported by Stahl, Unkelbach, and Corneille (2009). As discussed by these authors, one way to account for this data pattern is to assume that participants simply tried to behave conform the expectations of the experimenter. For two reasons, however, we consider such a demand account rather unlikely. First, the mere fact that valence awareness effects emerged both at the explicit and the implicit level is difficult to reconcile with such an explanation (Förderer & Unkelbach, 2012; Stahl et al., 2009). Second, a more fine-grained analysis of the AMP data shows that the interaction between the EC effect and valence awareness was more pronounced when participants responded relatively rapidly (i.e., faster than their own median response latency), $F(1, 1286.40) = 39.11, p < .001$, as compared to when they responded relatively slowly, $F(1, 1326.46) = 20.57, p < .001$. Even if it is assumed that the AMP is susceptible to demand effects (see Bar-Anan & Nosek, 2012; Payne et al., 2013), one would expect the exact opposite data pattern if demand effects were indeed responsible for the EC effects obtained in the present study. So, how can we account for the modulation of EC effects by valence awareness? Based on recent work by Huetter, Sweldens, Stahl, Unkelbach, and Klauer (2012) as well as findings obtained by Stahl et al. (2009), one might argue that participants based their valence awareness ratings on their liking/disliking of a particular CS. The fact that the selective generalization effect also emerged in the valence awareness ratings is consistent with this viewpoint. Irrespectively, it should be noted that neither the valence ratings nor the AMP data revealed negative EC effects for unaware CSs in Experiments 2 and 3. On the contrary, in Experiment 3, unaware CSs even

produced a significant positive EC effect. We are therefore reluctant to put too much weight on this finding.

Before discussing the broader theoretical implications of our findings, we would like to present the results of two follow-up studies. In a first study, we sought to replicate Experiment 1 using a different set of CSs and generalization stimuli. The motivation for this approach was straightforward. Face pictures are seldom, if ever, affectively neutral. Pictures of old men, for example, are typically evaluated in a (relatively) negative manner. To be sure, the counterbalancing conditions implemented in Experiment 1 guaranteed that our findings were not a by-product of pre-existing, overlearned attitudes. We were unable to prevent, however, that pre-existing attitudes introduced a considerable amount of error variance. We therefore decided to replicate Experiment 1 using CSs and generalization stimuli that were truly neutral. Instead of face pictures, we used artificial, gray-scale figures in Experiment 2 (i.e., Gabor patches, see below).

A second important modification concerns the nature of the CS-categorization task that was used to manipulate FSAA. In Experiment 1, participants received no error feedback at all and the sequence of events on any given trial was completely unaffected by participants' performance. To ensure that participants were motivated to assign attention to the relevant stimulus dimension of the Gabor patches in Experiment 2, we decided to make the presentation of the USs contingent upon a correct CS classification.

A final modification concerns the number of EC trials. In Experiment 1, each CS was presented only once. There is some evidence, though, showing that EC effects tend to increase with increasing numbers of pairings (Baeyens, Eelen, Crombez, & Van den Bergh, 1992; De Houwer et al., 2001; Sachs, 1975; Staats & Staats, 1959). Therefore, to maximize the likelihood

of obtaining solid EC effects, we decided to present all CSs several times instead of just once during the acquisition phase of Experiment 2.

Experiment 2

Method

Participants. Participants were 41 students at Ghent University (11 men, 30 women) who received course credit for their participation. All participants were Dutch speakers and had normal or corrected-to-normal vision. They gave informed consent before participation. One participant made 42 % errors on the CS-categorization task during the acquisition phase of the experiment and was thus exposed to just 58 % of the CS – US pairings (see below). Because this participant was clearly an outlier in comparison with the complete sample ($M = 8.49\%$, $SD = 7.56\%$), the data of this person were excluded from the analyses.

Materials. All materials were identical to those used in Experiment 1, with the exception of the stimuli that were used as CSs. Instead of black-and-white face pictures of young and old men and women, we now used 20 grayscale Gabor patches (384 x 384 pixels). These Gabor patches varied on two, perceptually separable dimensions: spatial frequency and spatial orientation (for examples, see Figure 1). Each quadrant of the stimulus space comprised 5 stimuli.¹ Values used for the spatial frequency dimension were: 4.25, 5.5, 6.75, 9.25, 10.5, and 11.75 cycles. Values used for the orientation dimension were: 11.25, 22.5, 33.75, 56.25, 67.5, and, 78.75 degrees. Exact coordinates for the construction of all the Gabor patches used in the present experiment can be obtained from the first author.

Procedure. Similar to Experiment 1, only a subset of Gabor patches was used for a given participant during the acquisition phase of the experiment. Within each attention condition, half of the participants were presented with 5 Gabor patches high in spatial frequency and an

orientation above 45° and 5 stimuli low in spatial frequency and an orientation below 45° . The remaining participants were presented with 5 Gabor patches high in spatial frequency and an orientation below 45° and 5 stimuli low in spatial frequency and an orientation above 45° . The Gabor patches that were not used during the acquisition phase were later used as generalization stimuli during the measurement phase of the experiment.

For each participant separately, the computer program selected 5 positive and 5 negative USs from the complete list of available USs (random sampling without replacement). All stimuli within a particular CS category were then paired either with all the positive USs or all the negative USs, leading to a total of 50 EC trials ($2 \text{ categories} \times 5 \text{ CSs} \times 5 \text{ USs}$). The assignment of different CS categories/stimuli to either positive or negative USs was again counterbalanced across participants.

Similar to Experiment 1, the acquisition phase involved a CS-classification task aimed at directing attention towards a specific stimulus dimension (i.e., spatial frequency or orientation). The nature of the CS-classification task, however, was quite different from the one implemented in Experiment 1. Participants were now asked to classify the CSs in two arbitrary categories, i.e., ‘Category A’ and ‘Category B’. In one condition (hereafter referred to as the Frequency Condition), participants were informed that assigning attention to ‘the thickness of the lines’ would help them discriminate between the two CS categories ($n = 21$). Conversely, in the other condition (hereafter referred to as the Orientation Condition), participants were informed that assigning attention to ‘the orientation of the lines’ would be an efficient strategy to optimize their performance ($n = 19$). The cutoff values for assigning a particular CS to either Category A or Category B were 45 degrees and 8 cycles, for the orientation dimension and the spatial frequency dimension respectively. The CSs were presented until a classification response was registered

and participants were asked to learn which CS belonged to which category by relying on the feedback given on the computer screen. In case of an erroneous response, a 3000-ms error message (i.e., ‘FOUT!’) was displayed. In case of a correct response, the US was presented for 3000 ms. Participants were thus required to guess on the first trial but quickly learned to classify the CSs correctly. After exclusion of one participant who made an exceptionally high number of errors during the acquisition phase of the experiment (see above), the overall mean error rate was 7.65 % ($SD = 5.45$). The difference between the mean error rate on negative-US trials ($M = 8.2$ %) and positive-US trials ($M = 7.1$ %) was statistically unreliable, $F < 1$. The error rate was about twice as high during the first half of the acquisition phase as compared to the second half of the acquisition phase (i.e., 10.15 % vs. 4.78 %).

The remainder of the experiment was highly similar to Experiment 1. Participants were first asked to provide valence ratings for the CSs and the generalization stimuli. Next, participants completed an AMP in which each CS and each generalization stimulus was presented exactly once. Finally, participants were asked to indicate, for each stimulus separately, whether they thought it had been paired with a negative or a positive US (i.e., valence awareness ratings). Except for the nature of the stimuli (Gabor patches instead of face pictures) and the number of stimuli used (5 stimuli in each category instead of 8), the procedures used in the measurement phase of Experiment 2 were identical to those implemented in Experiment 1. As in Experiment 1, the inter-trial interval varied randomly between 1500 ms and 2500 ms during the acquisition phase of the experiment. During all other phases, a variable inter-trial interval between 500 ms and 1500 ms was implemented.

Results

Acquisition effects. For each dependent measure, we performed a 2 (Group: Gender vs. Age) \times 2 (CS type: paired with positive USs vs. paired with negative USs) repeated measures ANOVA. Valence awareness ratings revealed that participants were clearly able to indicate, at least on average, whether a particular CS had been paired with positive CSs ($M = 0.64$) or negative CSs ($M = -0.60$), $F(1, 38) = 50.35$, $p < .0001$, $\eta_p^2 = .57$. More importantly, Gabor patches paired with positive USs ($M = 2.29$) were rated more positively than Gabor patches paired with negative USs ($M = -2.92$), $F(1, 38) = 44.89$, $p < .0001$, $\eta_p^2 = .54$. Finally, the AMP revealed a higher proportion of positive responses after the presentation of CSs that had been paired with positive USs ($M = 59.47\%$) than after the presentation of CSs that had been paired with negative USs ($M = 48.17\%$), $F(1, 38) = 4.18$, $p < .05$, $\eta_p^2 = .10$. None of these main effects was qualified by an interaction with the group factor, all $F_s < 2.68$, all $p_s > .11$.

Similar to Experiment 1, we also examined whether EC effects were dependent upon valence awareness using linear mixed effect analyses. Whereas the valence ratings revealed a significant interaction between CS type and valence awareness, $F(1, 398.82) = 65.86$, $p < .001$, this effect was only marginally significant in the AMP data, $F(1, 392.03) = 2.73$, $p = .10$. Follow-up analyses showed that aware CSs produced a significant (positive) EC effect, both in the valence ratings, $M_{CSpos} = 3.21$, $M_{CSneg} = -3.43$), $F(1, 271.01) = 290.09$, $p < .001$, and the AMP data, $M_{CSpos} = 59.74\%$, $M_{CSneg} = 45.46\%$), $F(1, 277.23) = 6.47$, $p < .05$. In contrast, CSs that were classified incorrectly during the valence awareness task produced no effects whatsoever, $F_s < 1$.

Generalization effects. As for Experiment 1, we calculated a generalization index, G, for each participant and each dependent measure. Specifically, individual generalization scores for

the orientation dimension were subtracted from generalization scores for the spatial frequency dimension. Scores larger than zero thus reflect that generalization along the spatial frequency dimension was stronger than generalization along the orientation dimension. Scores smaller than zero indicate that generalization along the orientation dimension was stronger than generalization along the spatial frequency dimension. We thus expected positive generalization indices in the Frequency Group and negative generalization indices in the Orientation Group.

As can be seen in Table 2, the results matched our expectations. More specifically, valence awareness ratings for the generalization stimuli shifted towards the valence of USs that had been paired with CSs similar in terms of the stimulus dimension that was selectively attended to, $F(1, 38) = 93.00, p < .0001, \eta_p^2 = .71$. Likewise, participants were inclined to evaluate (novel) generalization stimuli in the same manner as CSs that were similar in terms of the stimulus dimension that was selectively attended to, both at the explicit level (valence ratings), $F(1, 38) = 23.34, p < .0001, \eta_p^2 = .38$, and the implicit level (AMP), $F(1, 38) = 10.75, p < .005, \eta_p^2 = .22$. Moreover, with the exception of the generalization index for the AMP in the Orientation Group ($t < 1$), all generalization indices reached significance within each condition, all $ts > 3.38$, all $ps < .05$.

Similar to Experiment 1, we also examined whether generalization effects were contingent upon valence awareness. For these analyses, valence awareness was again treated as a between-subjects factor (see above). Only 5 participants were identified as having inaccurate valence awareness. Both the generalization effect in the valence awareness ratings and the (explicit) valence ratings were affected by valence awareness, $F(1, 36) = 31.57, p < .005, \eta_p^2 = .47$, and $F(1, 36) = 10.03, p < .005, \eta_p^2 = .22$, respectively. The generalization effect in the AMP was not affected by valence awareness, $F < 1$. Nevertheless, for each of the three dependent

measures (including the AMP), the anticipated generalization effect was reliable only in participants with accurate valence awareness, all $F_s > 10.21$, all $p_s < .005$. In participants whose valence awareness ratings were not in line with the actual CS-US pairings, generalization effects were anything but significant, all $F_s < 1.01$. The latter analyses should be treated with caution, however, given the very low number of participants who lacked accurate valence awareness.

Discussion

The results are clear-cut. The EC procedure implemented in the present study was highly effective and impacted each of our dependent measures (i.e., valence awareness ratings, valence ratings, and AMP). Linear mixed effect analyses revealed, however, that the EC effect was reliable only for aware CSs. When the analyses were restricted to CSs that were classified incorrectly during the valence awareness task, the EC effect was absent. This data pattern is again consistent with earlier studies showing that participants are inclined, at least under certain conditions, to base their valence awareness ratings on their liking/disliking of a particular CS (see Huetter et al., 2012; Stahl et al., 2009). In addition, the absence of an EC effect for unaware CSs confirms that the use of (neutral) Gabor patches is an effective strategy to rule out the influence of pre-existing attitudes.

More importantly, Experiment 2 also revealed highly significant selective generalization effects. More specifically, attitude generalization along the spatial frequency dimension was more pronounced in participants who learned to categorize Gabor patches in terms of spatial frequency. Conversely, attitude generalization along the orientation dimension was more pronounced in participants who were encouraged to focus their attention on the spatial orientation of the Gabor patches. This data pattern was found in the explicit valence ratings, in the valence awareness ratings, and even in the AMP. Our findings therefore convincingly

demonstrate that, consistent with our predictions, attitude generalization is indeed dependent upon FSAA.

One limitation of our findings so far is that participants were actively responding to specific stimulus features of the CSs at the same time as they were presented with the CS-US pairings. According to our framework, this need not be the case in order to obtain selective generalization effects. As soon as attention is assigned to a particular stimulus dimension, attitude generalization should be confined to generalization stimuli similar to the CSs in terms of stimulus features that are selectively attended to, with or without the need to actively categorize the CSs in terms of a particular stimulus dimension. We therefore decided to run a final study in which FSAA was manipulated prior to the actual EC phase. In Experiment 3, participants first completed a category-learning task in order to manipulate FSAA. The stimuli used during this category-learning task were Gabor patches that varied along the spatial frequency as well as the spatial orientation. Similar to Experiment 2, participants were asked to classify these Gabor patches in two arbitrary categories (i.e., ‘Category A’ and ‘Category B’) after being informed that assigning attention to either the spatial frequency dimension or the orientation of the Gabor patches would help them discriminate between the two categories. None of the stimuli presented during this category-learning task were later used as a CS or generalization stimulus. Based on our framework, we expected to obtain similar selective generalization effects as in Experiments 1 and 2.

Experiment 3

Method

Participants. Participants were 32 students at Ghent University (4 men, 28 women). They received course credit for their participation or were paid €8 for their help in this experiment and

an unrelated other experiment. All participants were Dutch speakers and had normal or corrected-to-normal vision. All participants gave informed consent before participation. One participant pressed the same key throughout the AMP. The data of this participant were excluded from all analyses.

Materials. All materials were identical to those used in Experiment 2, with one exception. Next to the set of Gabor patches used in Experiment 2, we used an additional set of 20 Gabor patches during the category learning task (hereafter referred to as *induction stimuli*). Similar to the Gabor patches used in Experiment 2, this new set of Gabor patches covered the entire stimulus space, with 5 stimuli in each quadrant. Values used for the spatial frequency dimension were: 3.625, 4.875, 5.500, 6.125, 8.000, 9.875, 10.500, 11.125, and 12.375 cycles. Values used for the orientation dimension were: 5.625, 16.875, 22.500, 28.125, 45.000, 61.875, 67.500, 73.125, and 84.375 degrees. Exact coordinates for the construction of all the Gabor patches used in the present experiment can be obtained from the first author.

Procedure. The experiment started with a category learning phase in which each of the induction stimuli were presented exactly once (20 trials in total). As in Experiment 2, participants were asked to classify these stimuli in two arbitrary categories, i.e., ‘Category A’ and ‘Category B’. In one condition (hereafter referred to as the *Frequency Group*), participants were informed that assigning attention to ‘the thickness of the lines’ would help them discriminate between the two CS categories ($n = 16$). In the other condition (hereafter referred to as the *Orientation Group*), participants were informed that assigning attention to ‘the orientation of the lines’ would be an efficient strategy to optimize their performance ($n = 16$). Participants were asked to learn which stimuli belonged to which category by relying on the feedback given on the computer screen. In case of an erroneous response, a 3000-ms error message was

displayed in a red font (i.e., ‘FOUT!’, incorrect). In case of a correct response, a positive feedback message was presented for 3000 ms (i.e., ‘JUIST!’, correct). As in Experiment 2, participants were thus required to guess on the first trial but they quickly learned to classify the stimuli correctly. The overall mean error rate was 3.75 % ($SD = 4.75$).

After participants had completed the induction phase, they moved on to the acquisition phase of the experiment. Unlike Experiment 1 and 2, participants were no longer required to perform a task during this acquisition phase of the experiment. They were simply asked to watch all stimuli presented on the computer screen carefully. Each CS was presented for 2000 ms and was immediately followed by a 3000-ms US. In all other aspects, the procedures used in Experiment 3 were identical to those used in Experiment 2.

Results

Acquisition effects. For each dependent measure, we performed a 2 (Group: Gender vs. Age) \times 2 (CS type: paired with positive USs vs. paired with negative USs) repeated measures ANOVA. Valence awareness ratings demonstrated that participants were able to indicate, on average, whether a particular CS had been paired with positive CSs ($M = 0.63$) or negative CSs ($M = -0.69$), $F(1, 29) = 74.17$, $p < .0001$, $\eta_p^2 = .72$. Gabor patches paired with positive USs ($M = 3.86$) were also rated more positively than Gabor patches paired with negative USs ($M = -3.68$), $F(1, 29) = 43.10$, $p < .0001$, $\eta_p^2 = .60$. Finally, the AMP revealed a higher proportion of positive responses after the presentation of CSs that had been paired with positive USs ($M = 67\%$) than after the presentation of CSs that had been paired with negative USs ($M = 41\%$), $F(1, 29) = 14.41$, $p < .001$, $\eta_p^2 = .33$.

Similar to Experiments 1 and 2, we also examined whether EC effects were dependent upon valence awareness using linear mixed effect analyses. Whereas valence ratings revealed a

significant interaction between CS type and valence awareness, $F(1, 299.62) = 32.30, p < .01$, this effect was far from significant in the AMP data, $F(1, 303.98) = 1.61, p = .20$. Follow-up analyses showed that aware CSs produced a significant (positive) EC effect, both in the valence ratings, $M_{CSpos} = 4.68, M_{CSneg} = -4.32$), $F(1, 204.65) = 359.61, p < .01$, and the AMP data, $M_{CSpos} = 67.79\%, M_{CSneg} = 38.91\%$), $F(1, 202.75) = 23.40, p < .05$. Interestingly, the valence ratings also revealed a positive EC effect for unaware CSs, $M_{CSpos} = 1.19, M_{CSneg} = -1.44$), $F(1, 66.92) = 5.31, p < .01$. Numerically, the AMP data also revealed a positive EC effect for unaware CSs, but this effect was not significant, $M_{CSpos} = 60.62, M_{CSneg} = 46.56\%$), $F(1, 66.32) = 1.34, p = .25$.

Generalization effects. The generalization data again matched our expectations (see Table 2). Valence awareness ratings for the generalization stimuli shifted towards the valence of USs that had been paired with CSs similar in terms of the stimulus dimension that was selectively attended to, $F(1, 29) = 87.55, p < .0001, \eta_p^2 = .75$. Likewise, participants were inclined to evaluate (novel) generalization stimuli in the same manner as CSs that were similar in terms of the stimulus dimension that was selectively attended to, both at the explicit level (valence ratings), $F(1, 29) = 35.13, p < .0001, \eta_p^2 = .55$, and the implicit level (AMP), $F(1, 29) = 7.34, p < .05, \eta_p^2 = .20$. Generalization indices within each condition reached significance for the valence awareness ratings and the valence ratings, all $ts > 3.73$, all $ps < .005$, but just missed conventional significance levels for the AMP, $ts > 1.90, p < .08$. only a small number of participants were unable to report the contingency between the CSs and the valence of USs ($n = 2$). Therefore, we did not examine whether these generalization effects depended on valence awareness.

Discussion

The results were again clear-cut. Consistent with our predictions, each of our dependent measures revealed significant selective generalization effects. This set of observations adds

further weight to the idea that attitude generalization is critically dependent upon FSAA, especially considering the fact that participants were no longer required to categorize the CSs in terms of specific stimulus features during the EC procedure. The findings of the present experiment are also important for the discussion concerning the relationship between EC effects and valence awareness. Unlike Experiments 1 and 2, the valence ratings of Experiment 3 revealed a significant (positive) EC effect in the absence of accurate valence awareness. Although this effect was rather weak, it corroborates the claim that EC effects can be obtained in the absence of accurate valence awareness (Huetter et al., 2012).

General Discussion

Based on the fear-conditioning literature (e.g., Lissek et al., 2008; Vervliet et al., 2004, 2011, 2006) and models of human categorization learning (e.g., Nosofsky, 1984, 1986, 1986; Medin & Schaffer, 1978), we hypothesized that attitude generalization is critically dependent upon FSAA. To test this hypothesis, we conducted three EC experiments and examined whether EC effects generalized to novel stimuli as a function of FSAA. Our observations were clearly in line with our expectations. Irrespectively of whether FSAA was manipulated during (Experiments 1 and 2) or before (Experiment 3) the acquisition phase, (untrained) generalization stimuli were evaluated in the same manner as (trained) CSs that were similar in terms of stimulus features that were selectively attended to. Both implicit and explicit attitude measures revealed this selective generalization effect. Moreover, valence awareness ratings revealed the same data pattern. Participants were inclined to (incorrectly) point out that a particular generalization stimulus had been paired with a positive or negative US when this stimulus was similar to positive or negative CSs, respectively, in terms of the stimulus dimension that was selectively attended to.

The question now arises how the selective generalization effects obtained in our experiments can be accounted for at the mental process level. In principle, just like the EC effect itself can reflect the operation of different underlying mechanisms (see De Houwer, 2009), a number of different processes may have produced our effects. At the most abstract level, one can divide these processes into three broad categories: demand characteristics, processes operating at the category level, and processes operating at the exemplar level. We will discuss the viability of each of these (types of) explanations one by one.

Let us first consider an explanation in terms of demand characteristics. Because valence awareness was high in each of our experiments and EC effects were clearly dependent upon this factor, one anonymous reviewer suggested that demand effects may have produced our findings. According to this viewpoint, participants simply tried to behave conform the expectations of the experimenter (Stahl et al., 2009). For several reasons, however, we consider such an explanation rather implausible. First, in each of our experiments, EC effects and selective generalization effects were picked up by an implicit attitude measure as well as an explicit attitude measure.² Second, each of our experiments yielded additional findings that are hard, if not impossible, to reconcile with a demand account. Remember that participants were always presented with four stimulus categories, each of which consisted of several exemplars. If demand effects were the driving force behind our findings, one would expect participants to rate all exemplars of a particular stimulus category as consistently positive or consistently negative. Across Experiments 1 and 2, however, no more than 3.13 % of all participants responded in a consistent manner to all exemplars of each stimulus category (0.00 % and 10.00 %, in Experiment 1 and 2, respectively). In Experiment 3, the proportion of participants who responded in a consistent manner to all exemplars of each stimulus category was substantially higher (i.e., 40.00 %), but in

this experiment a reliable (positive) EC effect was found in the absence of accurate valence awareness. Taken together then, it seems rather unlikely that demand characteristics were responsible for the effects obtained in our experiments.

Having ruled out an explanation in terms of demand characteristics, let us discuss the possibility that processes operating at the category level were responsible for the selective generalization effects obtained in our studies. As a first example of such an account, consider the propositional account of EC (see De Houwer, 2009; Mitchell et al., 2009). According to this model, EC effects are mediated by the formation of propositions about the relationship between the CS and the US. Translated to our experiments, it might be argued that participants formed a proposition about the occurrence of a specific stimulus feature and the occurrence of a positive or negative US. Such a process can be seen as a form of category learning as the crucial difference between the stimulus categories in our experiments was always the absence or presence of a particular stimulus feature. Crucially, the formation of a proposition is assumed to be a non-automatic process that requires awareness and cognitive resources. A propositional model of EC can therefore explain why the EC effects obtained in our studies were dependent upon valence awareness. Moreover, as propositions can be retrieved from memory in an automatic fashion once they are formed (Bar-Anan, De Houwer, & Nosek, 2010; Zanon, De Houwer, & Gast, 2012), a propositional account can also explain why we were able to capture selective generalization effects using an implicit attitude measure.

It should be emphasized, however, that a propositional account is not the only model that can provide a category-level explanation for the selective generalization effects obtained in our studies. Association formation models can deal equally well with this phenomenon. It could be argued, for example, (a) that individual stimulus features can enter the association formation

process (see also Rescorla, 1976) and (b) that FSAA determines which stimulus features are involved in the associations, for example by increasing the salience and/or informational relevance of features that are selectively attended to (Vervliet et al., 2010). Selective generalization effects can then be attributed to the fact that FSAA maximizes conditioning to features shared by the CS and a specific generalization stimulus. Again, such a mechanism can be seen as a form of category learning because the stimulus categories used in our experiments were always defined in terms of the presence or absence of specific stimulus features.

Interestingly, this framework also coincides with recent findings obtained by Le Pelley, Reimers, Calvini, Spears, Beesley, and Murphy (2010). These authors demonstrated that cues previously experienced as predictive of neutral outcomes are more likely to acquire an evaluative connotation through a conditioning procedure than cues experienced as non-predictive (see also Le Pelley, Calvini, & Spears, 2013; Le Pelley, Suret, & Beesley, 2009). Crucially, Le Pelley et al. (2010) argued that this effect was mediated by attentional processes: Participants learned to attend to particular cues, and to ignore others, on the basis of their predictiveness (Mackintosh, 1975). In other words, the effects observed by Le Pelley et al. (2010) were most pronounced for specific cues that were selectively attended to. As a logical consequence, one can expect attitudes evoked by a generalization stimulus to depend on the presence or absence of such cues.

To summarize, both an association formation model and a propositional model can readily explain the selective generalization effects obtained in our studies. Nevertheless, it should be noted that both models have difficulty dealing with other aspects of our data. Consider, for example, the observation that a reliable positive EC effect was found in the absence of accurate valence awareness in Experiment 3. This observation is difficult to reconcile with a propositional account because the formation of propositions is assumed to require awareness. Conversely, an

association formation model is less suited to explain the observation that our effects were dependent on valence awareness in the first place.

To resolve these inconsistencies, one might argue that both propositional processes and associative processes contributed to our effects simultaneously, an idea that would be in perfect accordance with recent dual-process models (Gawronski, & Bodenhausen, 2006, 2011). As an alternative solution, however, it could also be argued that our findings are best accounted for in terms of processes operating at the exemplar level. According to exemplar-based models of categorization and memory (e.g., Hinzman, 1984, 1986, 1988; Klauer, 2009; Smith & Zárate, 1992), memory traces correspond to specific objects, persons, or experiences as interpreted by the perceiver. When memory is probed with a particular target stimulus, each of these exemplar representations is assumed to contribute to the overall memory response to some degree: The stronger the overlap between a target stimulus and a particular exemplar representation, the stronger the influence of that exemplar representation on the memory response. Crucially, it has been argued that the perceiver's attention to specific stimulus dimensions determines the weight of each stimulus dimension in computing the similarity between the exemplar representations and the target stimulus (Smith & Zárate, 1992). Within such a framework, it can be argued that FSAA determines the weight of each stimulus dimension in computing the similarity between the exemplar representations and the target stimulus at the time of memory retrieval (Smith & Zárate, 1992). One can therefore expect the overall memory response to be driven mainly by exemplar information stemming from exemplars similar in terms of stimulus features that are selectively attended to. Interestingly, an explanation in terms of processes operating at the exemplar level has no difficulty accounting for the observation that we obtained a significant EC effect in the absence of valence awareness in Experiment 3. All memory traces are assumed to

contribute to the overall memory response, irrespective of whether an individual can actively remember making an evaluation when encountering a specific exemplar or not. Moreover, given that explicit abstractions and inferences made at the time of encoding are an integral part of an exemplar representation, an exemplar model of EC can also deal with the observation that valence awareness moderated our effects. Nevertheless, we hasten to confirm, that (much) more research would be needed to fully substantiate an exemplar-based account of EC.

As another avenue for future research, it also seems interesting to examine whether selective generalization effects can be obtained by manipulating FSAA *after* evaluative learning took place. From a theoretical perspective, such an approach would be particularly important because it can potentially shed light on the interplay between processes operating at the time of learning and processes operating at the time of the retrieval. In addition, should research confirm that changes in FSAA at *time 2* can impact the generalization of attitudes acquired at *time 1*, this approach might be exploited as a new means to alter evaluative responding. That is, instead of changing people's attitudes by exposing them to a massive amount of EC-conditioning trials, an attentional (re-) training in which participants learn to focus their attention on specific stimulus dimensions in an automatic fashion may be more fruitful.

To summarize, the present research convincingly demonstrates that attitude generalization depends upon FSAA, but more research is needed to determine the nature of the processes (or combination of processes) that underlie this effect. Irrespectively, our findings have important implications for researchers who use EC procedures to reduce or alter (implicit) attitudes in applied settings (e.g., Olson & Fazio, 2006). Our work clearly demonstrates that the same training procedure can produce generalization effects along different stimulus dimensions,

depending on participants' attentional mindset. It is therefore recommended to take FSAA effects into account when turning to EC procedures in order to alter evaluative responding.

References

- Allport, G. W. (1935). Attitudes. In C. Murchison (Ed.), *Handbook of social psychology* (pp. 798-844). Worcester, Mass.: Clark University Press.
- Baeyens, F., Eelen, P., Crombez, G., & Van den Bergh, O. (1992). Human evaluative conditioning: Acquisition trials, presentation schedule, evaluative style and contingency awareness. *Behaviour Research and Therapy*, 30, 133-142. doi: 10.1016/0005-7967(92)90136-5
- Baeyens, F., Kaes, B., Eelen, P., & Silverans, P. (1996). Observational evaluative conditioning of an embedded stimulus element. *European Journal of Social Psychology*, 26, 15-28. doi: 10.1002/(SICI)1099-0992(199601)26:1<15::AID-EJSP729>3.0.CO;2-5
- Bar-Anan, Y., De Houwer, J., & Nosek, B. A. (2010). Evaluative conditioning and conscious knowledge of contingencies: A correlational investigation with large samples. *Quarterly Journal of Experimental Psychology*, 63, 2313-2335. doi: 10.1080/17470211003802442
- Bar-Anan, Y., & Nosek, B. A. (2012). Reporting intentional rating of the primes predicts priming effects in the affective misattribution procedure. *Personality and Social Psychology Bulletin*, 38, 1194-1208. doi: 10.1177/0146167212446835
- Bierley, C., McSweeney, F. K., & Vannieuwkerk, R. (1985). Classical conditioning of preferences for stimuli. *Journal of Consumer Research*, 12, 316-323. doi: 10.1086/208518
- De Houwer, J. (2009). The propositional approach to associative learning as an alternative for association formation models. *Learning & Behavior*, 37, 1-20. doi: 10.3758/LB.37.1.1

- De Houwer, J., Thomas, S., & Baeyens, F. (2001). Associative learning of likes and dislikes: A review of 25 years of research on human evaluative conditioning. *Psychological Bulletin*, 127, 853-869. doi: 10.1037//0033-2909.127.6.853
- Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology*, 87, 293-311. doi: 10.1037/0022-3514.87.3.293
- Fazio, R.H, Jackson, J. R., Dunton, B. C., & Williams, C. J (1995). Variability in automatic activation as an unobtrusive measure of racial-attitudes: A bona-fide pipeline. *Journal of Personality and Social Psychology*, 69, 1013-1027 DOI: 10.1037//0022-3514.69.6.1013
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132, Pages: 692-731. DOI: 10.1037/0033-2909.132.5.692
- Gawronski, B., & Bodenhausen, G. V. (2011). The associative-propositional evaluation model: Theory, evidence, and open questions. *Advances in Experimental Social Psychology*, 44, 59-127. DOI: 10.1016/B978-0-12-385522-0.00002-0
- Gast, A., De Houwer, J., & De Schryver, M. (2012). Evaluative Conditioning Can be Modulated by Memory of the CS-US Pairings at the Time of Testing. *Learning and Motivation*, 43, 116-126. doi: 10.1016/j.lmot.2012.06.001
- Goldstone, R. L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General*, 130, 116-139. doi: 10.1037//0096-3445.130.1.116
- Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behavior Research Methods and Instrumentation*, 16, 96-101.

- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411-428. doi: 10.1037//0033-295X.93.4.411
- Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, 95, 528–551. doi: DOI: 10.1037/0033-295X.95.4.528
- Hofmann, W., De Houwer, J., Perugini, M., Baeyens, F., & Crombez, G. (2010). Evaluative Conditioning in Humans: A Meta-Analysis. *Psychological Bulletin*, 136, 390-421. doi: 10.1037/a0018916
- Huetter, M., Sweldens, S., Stahl, C., Unkelbach, C., & Klauer, K. C. (2012). Dissociating contingency awareness and conditioned attitudes: Evidence of contingency-unaware evaluative conditioning. *Journal of Experimental Psychology: General*, 141, 539-557 doi: 10.1037/a0026477
- Imhoff, R., Schmidt, A. F., Bernhardt, J., Dierksmeier, A., & Banse, R. (2011). An inkblot for sexual preference: A semantic variant of the Affect Misattribution Procedure. *Cognition and Emotion*, 25, 676-690. doi: 10.1080/02699931.2010.508260
- Kalish, H. (1969). Stimulus generalization. In M. Marx (Ed.), *Learning: Processes* (pp. 205-297). Oxford, England: Macmillan.
- Kenward, M. G., & Roger, J. H. (1997). Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics*, 53, 983-997. DOI: 10.2307/2533558
- Klauer, K. C. (2009). Spontaneous evaluations. In F. Strack & J. Förster (Eds.), *Social cognition: The basis of human interaction* (pp. 199-217). New York: Psychology Press.
- Lamberts, K. (2000). Information-accumulation theory of speeded categorization. *Psychological Review*, 107, 227-260. doi: 10.1037//0033-295X.107.2.227

- Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (2001). *International affective picture system (IAPS): Instruction manual and affective ratings*. Technical Report A-5. Gainesville, FL: Center for Research in Physiology, University of Florida.
- Le Pelley, M. E., Calvini, G., & Spears, R. (2013). Learned predictiveness influences automatic evaluations in human contingency learning. *Quarterly Journal of Experimental Psychology*, 66, 217-228.
- Le Pelley, M. E., Reimers, S. J., Calvini, G., Spears, R., Beesley, T., & Murphy, R. A. (2010). Stereotype formation: Biased by association. *Journal of Experimental Psychology: General*, 139, 138-161. doi: 10.1037/a0018210
- Le Pelley, M. E., Suret, M. B., & Beesley, T. (2009). Learned predictiveness effects in humans: A function of learning, performance, or both? *Journal of Experimental Psychology: Animal Behavior Processes*, 35, 312-327. doi: 10.1037/a0014315
- Lissek, S., Biggs, A. L., Rabin, S. J., Cornwell, B. R., Alvarez, R. P., Pine, D. S., et al. (2008). Generalization of conditioned fear-potentiated startle in humans: Experimental validation and clinical relevance. *Behaviour Research and Therapy*, 46, 678-687. 276-298. doi: 10.1016/j.brat.2008.02.005
- Mackintosh, N. J. (1975). A theory of attention. Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207-238. doi: 10.1037//0033-295X.85.3.207
- Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, 32, 183-246. doi: 10.1017/S0140525X09000855

- Nosofsky, R. M. (1984). Choice, similarity, and the context of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 104-114. doi: 10.1037/0278-7393.10.1.104
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.
- Nosofsky, R. M. (1987). Attention and learning processes in identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 87-108. doi: 10.1037//0278-7393.13.1.87
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104, 266-300. doi: 10.1037/0033-295X.115.2.446
- Olson, M. A., & Fazio, R. H. (2006). Reducing automatically activated racial prejudice through implicit evaluative conditioning. *Personality and Social Psychology Bulletin*, 32, 421-433. doi: 10.1177/0146167205284004
- Payne, B. K., Brown-Iannuzzi, J., Burkley, M., Arbuckle, N. L., Cooley, E., Cameron, C. D., & Lundberg, K. B. (2013). Intention invention and the affect misattribution procedure: Reply to Bar-Anan and Nosek (2012). *Personality and Social Psychology Bulletin*, 39, 375-386. doi:10.1177/0146167212475225
- Payne, B. 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. doi: 10.1037/0022-3514.89.3.277
- Payne, B. K., Govorun, O., & Arbuckle, N. L. (2008). Automatic attitudes and alcohol: Does implicit liking predict drinking? *Cognition and Emotion*, 22, 238-271. doi: 10.1080/02699930701357394

- Payne, B. K., Hall, D. L., Cameron, C. D., & Bishara, A. J. (2010). A process model of affect misattribution. *Personality and Social Psychology Bulletin*, 36, 1397-1408. doi: 10.1177/0146167210383440
- Pleyers, G., Corneille, O., Luminet, O., & Yzerbyt, V. (2007). Aware and (dis)liking: Item-based analyses reveal that valence acquisition via evaluative conditioning emerges only when there is contingency awareness. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 130-144. doi: 10.1037/0278-7393.33.1.130
- Pleyers, G., Corneille, O., Yzerbyt, V., & Luminet, O. (2009). Evaluative conditioning may incur attentional costs. *Journal of Experimental Psychology: Animal Behavior Processes*, 35, 279-285. doi: 10.1037/a0013429
- Razran, G. (1949). Stimulus generalization of conditioned responses. *Psychological Bulletin*, 46, 337-365. doi: 10.1037/h0060507
- Rescorla, R. A. (1976). Stimulus generalization: Some predictions from a model of Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 2, 88-96. doi: 10.1037//0097-7403.2.1.88
- Sachs, D. H. (1975). Conditioning of affect to a neutral stimulus: Number of trials. *Perceptual and Motor Skills*, 40, 895-901.
- Shook, N. J., Fazio, R. H., & Eiser, J. R. (2007). Attitude generalization: Similarity, valence, and extremity. *Journal of Experimental Social Psychology*, 43, 641-647. doi: 10.1016/j.jesp.2006.06.005
- Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, 99, 3-21. doi: 10.1037/0033-295X.99.1.3

- Spruyt, A., Clarysse, J., Vansteenwegen, D., Baeyens, F., & Hermans, D. (2010). Affect 4.0: A free software package for implementing psychological and psychophysiological experiments. *Experimental Psychology*, 57, 36-45. doi: 10.1027/1618-3169/a000005
- Spruyt, A., Hermans, D., De Houwer, J., & Eelen, P. (2002). On the nature of the affective priming effect: Affective priming of naming responses. *Social Cognition*, 20, 227-256. doi: 10.1521/soco.20.3.227.21106
- Spruyt, A., Hermans, D., De Houwer, J., Vandekerckhove, J., & Eelen, P. (2007). On the predictive validity of indirect attitude measures: Prediction of consumer choice behavior on the basis of affective priming in the picture – picture naming task. *Journal of Experimental Social Psychology*, 43, 599-610. doi: 10.1016/j.jesp.2006.06.009
- Staats, A. W., & Staats, C. K. (1959). Effect of number of trials on the language conditioning of meaning. *The Journal of General Psychology*, 61, 211-223.
- Stahl, C., Unkelbach, C., & Corneille, O. (2009). On the respective contributions of awareness of unconditioned stimulus valence and unconditioned stimulus identity in attitude formation through evaluative conditioning. *Journal of Personality and Social Psychology*, 97, 404-420. doi: 10.1037/a0016196
- Till, B. D., & Priluck, R. L. (2000). Stimulus generalization in classical conditioning: An initial investigation and extension. *Psychology & Marketing*, 17, 55-72. doi: 10.1002/(SICI)1520-6793(200001)17:1<55::AID-MAR4>3.0.CO;2-C
- Unkelbach, C., Förderer, S., & Stahl, C. (in press). Changing CS features alters evaluative responses in evaluative conditioning. *Learning & Motivation*, 43, 127-134. doi: 10.1016/j.lmot.2012.04.003

- Vervliet, B., Iberico, C., Vervoort, E. & Baeyens, F. (2011). Generalization gradients in human predictive learning: Effects of discrimination training and within-subjects testing. doi:*Learning and Motivation*, 42, 210-220. doi: 10.1016/j.lmot.2011.03.004
- Vervliet, B., Kindt, M., Vansteenwegen, D., & Hermans, D. (2010). Fear generalization in humans: Impact of verbal instructions. *Behaviour Research and Therapy*, 48, 38-43. doi: 10.1016/j.brat.2009.09.005
- Vervliet, B., Vansteenwegen, D., & Eelen, P. (2004). Generalization of extinguished skin conductance responding in human fear conditioning. *Learning & Memory*, 11, 555-558. doi: 10.1101/lm.77404
- Vervliet, B., Vansteenwegen, D., & Eelen, P. (2006). Generalization gradients for acquisition and extinction in human contingency learning, *Experimental Psychology*, 53, 132-142. doi: 10.1027/1618-3169.53.2.132
- Zanon, R., De Houwer, J., & Gast, A. (2012). Context effects in evaluative conditioning of implicit evaluations *Learning and Motivation*, volume: 43, 155-165 doi: 10.1016/j.lmot.2012.02.003

Author Note

Correspondence concerning this article should be addressed to Adriaan Spruyt, Department of Psychology, Ghent University, H. Dunantlaan 2, B-9000 Ghent, Belgium. Electronic mail may be sent to Adriaan.Spruyt@UGent.be. Preparation of this paper was supported by Grant BOF/GOA2006/001 and Methusalem Grant BOF09/01M00209 of Ghent University. Adriaan Spruyt is Postdoctoral Fellow of the Flemish Research Foundation (FWO - Vlaanderen).

Footnotes

¹ An additional set of 9 Gabor patches was used during the assessment phase of the experiment. Similar to the stimuli used as CSs, these additional Gabor patches varied in spatial frequency and orientation. None of these stimuli, however, was ever presented during the acquisition phase of the experiment. These stimuli were included for exploratory modeling purposes. The results obtained with these stimuli will not be reported here.

² The degree to which the AMP is able to capture implicit attitudes has recently been called into question (Bar-Anan & Nosek, 2012; but see Payne et al., 2013). It might therefore be worthwhile to replicate the present experiments using a different implicit attitude measures, such as the affective priming paradigm (e.g, Fazio, Jackson, Dunton, & Williams, 1995; Spruyt, Hermans, De Houwer, Vandekerckhove, & Eelen, 2007).

TABLE 1

Experiments 1: Mean generalization indices (G) and selective generalization effects (STE) for each dependent measure as a function of attention group.

<i>Dependent Measure</i>	G		STE
	Gender Group	Age Group	
valence ratings	.95	-1.63 [†]	2.58*
valence awareness ratings	1.15**	-0.72**	1.87**
AMP scores	12.18	-10.77	22.95 [†]

[†] $p < .10$, * $p < .05$, ** $p < .01$

TABLE 2

Experiments 2 and 3: Mean generalization indices (G) and selective generalization effects (STE) for each dependent measure as a function of attention group.

<i>Dependent Measure</i>	G		STE
	Frequency Group	Orientation Group	
Experiment 2			
valence ratings	9.37**	-8.69**	18.06**
valence awareness ratings	2.67**	-2.61**	5.28**
AMP scores	55.24**	-.06	55.30**
Experiment 3			
valence ratings	14.97**	-11.17**	26.14**
valence awareness ratings	2.73**	-2.43**	5.16**
AMP scores	42.50	-37.33*	79.83*

* $p < .05$, ** $p < .005$

Figure Caption

Figure 1. Examples of Gabor patches used in Experiments 2 and 3.

Figure 1

