Journal of Personality and Social Psychology
1983, Vol. 45, No. 3, 492-500

Implicit Learning and Generalization of the "Mere Exposure" Effect

Peter C. Gordon and Keith J. Holyoak
University of Michigan

Two experiments are reported in which the generalization of the "mere exposure" effect is examined. Both experiments demonstrated that positive affect, produced by repeated viewing of a set of stimuli, generalizes to previously unseen stimuli that are similar along certain abstract dimensions to the exposed stimuli. The first experiment used letter strings constructed according to a complex rule system. Positive affect attributable to exposure generalized to novel letter strings that obeyed the rule system. Affective generalization was related to subjects' judgments of whether the novel strings obeyed the rule system. The second experiment, in which the stimuli were complex visual patterns created by distorting standard forms, yielded an orderly gradient of affective generalization to novel patterns at varying levels of distortion. These experiments indicate that the exposure effect behaves in a manner similar to "implicit" concept learning and rule induction. The generalization techniques developed here provide a novel method for studying the affective processing of stimuli.

Repeated exposure to a stimulus causes increased liking of that stimulus (Zajonc, 1968). This phenomenon, known as the "mere exposure" effect, has been found for a wide variety of stimuli and is quite robust (Harrison, 1977). In a typical mere-exposure experiment, subjects are simply asked to observe a set of stimuli presented at varying frequencies and then to rate how much they like each one. Subjects generally rate the more frequently presented stimuli more favorably.

Although the stimuli in mere-exposure studies are presented with minimal instructions, it is likely that psychological processes other than the buildup of affect go on. In addition, "mere" presentation of stimuli results in something being learned about the structure of individual stimuli and about the similarities shared by the members of a set of stimuli. This type of cognitive category learning has been dubbed "implicit learning" (Reber, 1967). It is characterized as a process that occurs naturally when a person attends to the members of a structured stimulus set and that does not require conscious strategies such as hypothesis testing. Implicit learning may play an important role in basic cognitive processes such as language acquisition and pattern recognition. Exploration of the relation between implicit learning and the exposure effect may provide a way of understanding the nature of the exposure effect itself.

Reber and his colleagues have studied implicit learning using letter strings generated by artificial finite-state grammars that specify permissible ("grammatical") orders of letters. These grammars can be represented as a network of nodes connected by labeled directional pathways (see Figure 1). Strings are generated by following a route through the network, from the entrance to one of the exits, and adding the appropriate letter to the string for each path taken. Any letter string that can be obtained in this way obeys the grammar (e.g., XXRR and XMVTTRX for the grammar in Figure 1). Strings that cannot be

This research was supported by National Science Foundation Grant BNS-7904730. P. C. Gordon was supported by a National Institute of Mental Health Traineeship (1-T32-MH16892-01), and K. J. Holyoak was supported by an NIMH Research Scientist Development Award (1-K02-MH00342-02). We thank Donna Frey for her skillful typing of the manuscript and tables, Julianne Legon and Laura Pyle for valuable assistance in testing subjects, and John Jonides for providing the materials used in Experiment 1 and for several helpful suggestions. Bob Zajonc provided generous advice for which we are very grateful.

Requests for reprints should be sent to Peter C. Gordon or Keith J. Holyoak, Human Performance Center, University of Michigan, 330 Packard Road, Ann Arbor, Michigan 48104.

obtained in this way are nongrammatical (e.g., XMRX).

Reber (1967) demonstrated that after having memorized a set of grammatical strings, subjects could accurately judge the grammaticality of letter strings they had never seen before, even though they were not informed of the existence of the grammar at the time of memorization. Reber and Lewis (1977) found that subjects learning an artificial grammar first learned permissible letter bigrams and subsequently learned the positions in which the bigrams were allowed to occur. It therefore appears that subjects do not literally learn a finite-state network representation of the grammar but rather a feature representation based largely on position-specific bigrams. Reber (1976) demonstrated that subjects were actually more successful in learning a grammar when they simply memorized sample strings than when they actively examined the same set of sample strings and tried to figure out the rule system. This occurred despite the fact that when they were informed of the grammar's existence, some subjects in the memorization condition protested that they had learned nothing. When the underlying rule system is relatively complex, it seems that implicit learning may be more effective in inducing regularities than the consciously directed strategies subjects use when they are trying to learn (Reber, Kassin, Lewis, & Cantor, 1980). Reber and Allen (1978) demonstrated that it is not even necessary for subjects to memorize exemplars to learn implicitly. They simply had subjects look at a sample group of grammatical strings without being informed about the existence of a grammar. Afterwards, when informed about the rule-governed nature of the stimuli, subjects were able to make accurate grammaticality judgments. This result reinforces the conclusion that implicit learning is a natural product of attending to structured stimuli.

Implicit learning and development of affect take place under similar and quite simple conditions—the mere exposure of stimuli. In addition, the cognitive and affective consequences of mere exposure have been described in very similar terms. On the basis of subjects' introspective reports, Reber and Allen (1978) concluded that implicit learning

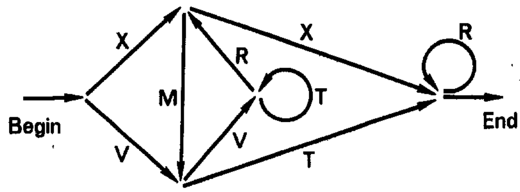

Figure 1. A representation of the finite-state grammar used to generate the grammatical letter strings used in Experiment 1.

is automatic, effortless, and relatively unconscious. This conclusion is bolstered by Reber and Lewis's (1977) finding that subjects' conscious, verbalizable knowledge about the artificial grammars always lagged behind their ability to make accurate grammaticality judgments. This characterization of implicit learning is similar to Zajonc's (1980) description of affective processes as automatic, effortless, nonverbalizable, and unconsciously based. Given these similarities, it seems possible that the mere-exposure effect and implicit learning are in some way related and that the affective consequences of stimulus exposure might generalize along the stimulus dimensions abstracted through implicit learning. If so, an investigation of affective generalization in the mere-exposure paradigm may shed some light on the relation between cognitive and affective learning processes.

Experiment 1

Experiment 1 was performed to determine whether positive affect developed through exposure to a set of letter strings generated by an artificial grammar would generalize to novel letter strings that also obey the grammar. The experiment consisted of an initial learning or exposure phase, during which subjects saw grammatical strings, followed by a test phase, during which subjects made affective and grammaticality judgments with novel grammatical and nongrammatical strings.

Method

Stimuli. Sixteen grammatical strings, ranging from three to seven letters in length, were generated from the grammar represented in Figure 1 for use in the learning

Table 1
Letter Strings Used in Experiment 1

Phase	Letter strings			
	Grammatical			
Exposure	XXRR	XMVTRX	VVRXRR	XMVRXRR
	XMVRX	VVTRXR	VVRMTR	VVTTRMT
	VTRRR	XMVRXR	XMVRMT	VVRMTRR
	XXRRR	XMTRRR	VVRMVRX	XMVTRMT
Test	XX	XMT	VVRX	VVTTTRX
	VT	VTR	VVRXR	XMVRXRRR
	XXR	XMTR	VVTTRX	VVTRMVRX
	Nongrammatical			
Test	VM	TTV	VTTX	VVRRRTX
	MT	XXM	RXTMV	RXTTVMXR
	VRT	MXXR	TTRXXM	MVTTTXVR

phase. An additional 12 grammatical strings of lengths two to eight were generated for the test phase. Twelve nongrammatical strings were also constructed for use in the test phase. These closely matched the grammatical strings with respect to letter composition and overall length. Of the 12 nongrammatical strings, 6 violated the grammar by one misplaced letter, 3 by two misplaced letters, 2 by four misplaced letters, and 1 string was randomly ordered. The letter strings employed in the two phases of the experiment are shown in Table 1.

Procedure A. Two versions of Experiment 1 were performed. In Procedure A, the 16 grammatical strings used in the learning phase appeared on four slides, with four strings on each. Subjects were presented with a slide for approximately 10 sec and were told to try to memorize the strings it contained. They were then given approximately 10 sec to try to write down the strings. This sequence was repeated five times in a row with each of the slides, giving the subjects five opportunities to try to memorize each string.

After the learning phase, subjects were informed that the strings they had just memorized all conformed to an underlying rule. They were then shown 24 slides with one letter string each and were asked to make two judgments about each. First, they were to decide whether the string followed the same rule as had the previous strings and to indicate how confident they were in their judgment, using the number 3 to express strong confidence, 2 to express moderate confidence, and 1 to express little confidence. Second, they were asked to rate how much they liked each string on a scale of 1 to 7, with 1 representing least liking and 7 representing most liking. The 24 strings on which these ratings were obtained consisted of the 12 new grammatical strings and the 12 nongrammatical strings described above. They were presented in a random order for approximately 5 sec each, and after each presentation subjects were given approximately 5 sec to write down their responses. Forty-six subjects were tested in a group as part of a course requirement.

Procedure B. In the above procedure, subjects were informed about the grammar before they made their affective judgments. Because this could have potentially influenced the assigned ratings, a second procedure was used with a different set of subjects that eliminated this problem and also did not involve memorization of the learning strings. The latter modification made the learning phase more similar to the typical exposure phase of mere-exposure studies. Procedure B was identical to Procedure A, except for the changes noted below. During the learning phase, subjects were not told to memorize the strings but were simply told to look at them. Also, the five presentations of each of the four learning slides were intermixed in a random order. After the learning phase, subjects were not immediately told about the grammar but were simply presented with the 24 test stimuli and asked to rate how much they liked each one. After these ratings were completed, subjects were informed about the existence of the grammar and were shown the test sequence again. This time they were asked to make grammaticality judgments and to give confidence ratings as in Procedure A. Thirty-six subjects were tested in a group as part of a course requirement.

Control condition. A control condition was run to assess the intrinsic likability, independent of any learning, of the grammatical and nongrammatical letter strings used in the test phase of the above procedures. Subjects were shown the 24 test strings used in Procedures A and B and were asked to rate how much they liked each string. They were given no exposure phase prior to the presentation of the test strings. Twenty-two subjects, not one of whom served in the previous conditions, were tested as a group.

Results and Discussion

The main results are summarized in Table 2. For purposes of analysis, the grammaticality judgments and confidence ratings were converted into a 6-point scale, with 6 representing a very confident "grammatical" judgment and 1 representing a very confident "nongrammatical" judgment. Separate analyses of variance (ANOVAS) were performed for each type of judgment (grammaticality and liking) in each condition (Procedures A and B and the control condition). Grammaticality of strings and subjects was a factor in these ANOVAS. In addition, a third factor of "items" was created by randomly pairing grammatical and nongrammatical letter strings of the same length. Minimum quasi-F ratios (F'_{min}) were calculated to assess the probability that the reported effects would simultaneously generalize to different samples of both subjects and letter strings (Clark, 1973).

In Procedure A, subjects gave higher grammaticality ratings to grammatical than to nongrammatical strings (M = 3.85 and 2.98, respectively), $F'_{min}(1, 15) = 11.5$, $p < .01$, indicating that they had in fact induced at

Table 2
Mean Grammaticality and Liking Ratings for Grammatical and Nongrammatical Letter Strings: Experiment 1

String	Procedure A		Procedure B		Control procedure liking
	Grammaticality	Liking	Grammaticality	Liking	
Grammatical	3.85	4.52	3.81	4.29	3.54
Nongrammatical	2.98	4.03	3.17	3.89	3.67

least part of the grammar. Subjects also gave higher affective ratings to grammatical than nongrammatical strings (M = 4.52 and 4.03, respectively), $F'_{min}(1, 27) = 7.99$, $p < .01$. In Procedure B, subjects again gave higher grammaticality ratings to grammatical than to nongrammatical strings (M = 3.81 and 3.17, respectively), $F'_{min}(1, 22) = 10.56$, $p < .01$, and they gave higher affective ratings to grammatical than to nongrammatical strings (M = 4.29 and 3.89, respectively) $F'_{min}(1, 36) = 6.13$, $p < .025$. The results for Procedures A and B were thus very similar despite their procedural differences. In the control condition the grammatical strings received slightly *lower* affective ratings than did the nongrammatical strings (M = 3.54 and 3.67, respectively); however, this reversal was not significant, $F'_{min}(1, 19) < 1$. The latter result indicates that the grammatical strings were not intrinsically more likable than the nongrammatical strings.

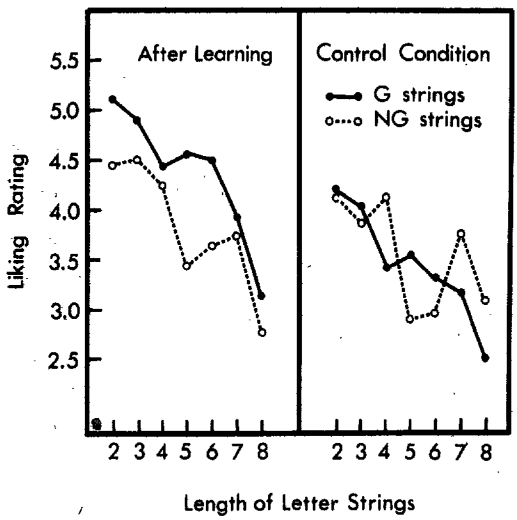

Figure 2. The results of Experiment 1 broken down by condition, string length, and grammaticality. (G = grammatical; NG = nongrammatical.)

Inspection of the data reveals an influence of length of letter string on liking ratings.[1] The combined liking ratings of Procedures A and B are shown in the left panel of Figure 2, broken down by grammaticality and length of letter string. The liking ratings from the control condition are shown in the right panel. It is apparent that amount of liking was inversely related to string length. An ANOVA with unequal numbers of observations was performed to compare the combined data of Procedures A and B (the experimental conditions) with that of the control condition. The effect of length was significant, $F(6, 630) = 14.9$, $p < .001$, and there was no Length × Condition (experimental vs. control) interaction. This result indicates that the liking ratings were influenced in an orderly fashion by a variable—string length—independent of whether an exposure phase had been given, suggesting that the exposure phase did not simply change the basis on which subjects made affective judgments about these stimuli. The effect of condition was significant, $F(1, 105) = 13.3$, $p < .001$, indicating that viewing letter strings during the exposure phase caused a general increase of liking for novel letter strings. The interaction of grammaticality and condition was also significant, $F(1, 105) = 11.5$, $p < .002$, showing that the higher ratings produced in the experimental condition were more pronounced for grammatical strings. The latter finding confirmed our earlier conclusion that prior exposure to grammatical letter strings caused a selective increase in liking for novel grammatical strings.

[1] Inspection of the grammaticality ratings did not reveal a similar pattern of influence attributable to string length.

To our knowledge, Experiment 1 provides the first demonstration of generalization of the mere-exposure effect. The results indicate that generalization of affect can take place along relatively abstract dimensions, indicating that affective responses to stimuli are based in part on complex processes that abstract structural regularities. Furthermore, the pattern of affective generalization matched subjects' cognitive discriminations regarding which stimuli obeyed the abstract rule system.

Experiment 2

Experiment 2 attempted to extend the finding of affective generalization to a very different type of stimuli than the letter strings used in Experiment 1. Each category of stimuli consisted of distortions of a single "standard" stimulus, which thus corresponded to the prototype or "central tendency" of the category (Posner & Keele, 1968). The particular stimuli and distortion method employed have been extensively studied in concept-learning experiments by Fried and Holyoak (in press). They have shown that subjects can learn these categories implicitly simply by observing a set of distortions of the standard stimulus, even when examples of two different categories are randomly intermixed. After learning, the probability of classifying a novel distortion of the standard as a category member is a decreasing function of its degree of distortion from the standard. The use of these stimuli allowed us to explore aspects of affective generalization that could not be easily examined using an artificial grammar.

In Experiment 2 we varied the degree of distortion of the stimuli used in the test phase to obtain an affective generalization gradient. In addition, the experiment included two types of exposure phase. In one condition (the abstraction condition) subjects were shown a set of distortions of a standard form, from which they could presumably abstract the commonalities through processes of implicit learning. This procedure was analogous to the type of exposure phase used in Experiment 1, in which subjects were shown a set of related letter strings. The second exposure condition (the repetition condition) involved repeated presentation of the standard itself. This condition was more analogous to the exposure conditions used in standard mere-exposure experiments, in which the same stimuli are presented repeatedly. This condition allowed us to determine if affect would generalize from a single stimulus. In addition, frequency of presentation (high vs. low) was varied. This factor, when crossed with type of exposure phase (abstraction vs. repetition), yielded four different exposure conditions.

Method

The procedure included an exposure phase, in which subjects were simply asked to view a set of stimuli, followed by a test phase, in which they were asked to rate how much they liked a new set of stimuli.

Stimuli. The stimuli were 10×10 matrices constructed of blue and orange squares. The matrices were approximately 15 by 16 cm in size and were presented on a color monitor controlled by an Apple II-Plus computer. The standards were constructed by randomly making half of the 100 squares of a matrix blue and half orange. Distortions of the standards were constructed by reversing the color (blue to orange or orange to blue) of a number of randomly selected squares.

Exposure phase. Two types of exposure phase were tested: an abstraction condition, consisting of distortions of standard stimuli, and a repetition condition, consisting of repeated presentations of the standards themselves. This variable was manipulated between subjects. For the abstraction condition, the number of squares changed in each distortion of the standard was determined by randomly sampling from the positive half of a normal distribution, with mean of zero and standard deviation of six, and then adding three. This meant that the minimum distortion of standard seen by subjects was 3%, which was also the modal value. Given the properties of a normal distribution, 68% of the stimulus presentations involved distortions of less than 9%, and 95% of the stimuli were distorted by less than 15%. For both the abstraction and the repetition conditions, two standards were used, a high-frequency one (40 presentations) and a low-frequency one (20 presentations).[2] Examples of each standard were randomly intermixed to construct the exposure phase. A given standard could be assigned to any of four possible conditions, defined by the factorial combination of abstraction versus repetition and high versus low frequency. Yoked quadruplets of subjects were therefore run with the same pair of standards, each standard occurring in each of the four conditions across the four subjects. Across quadruplets, different randomly generated standards were used. The instructions for the

[2] In a preliminary study, we used only one prototype and failed to detect any exposure effect or generalization. We switched to two prototypes because heterogeneous exposure sequences typically produce a larger exposure effect than do homogeneous sequences (Harrison & Crandall, 1972).

exposure phase told subjects simply to look at the patterns on the display screen. Each of the 60 patterns was presented for 3 sec, after which the screen went blank for 2 sec.

Test phase. The test phase was identical for all conditions. It consisted of distortions of 0, 5, 10, 25, and 50% of each of the two standards used during the exposure phase. The resulting 10 stimuli were presented in a random order. The two stimuli with 0% distortion were, of course, the standards themselves. The stimuli with 50% distortion were random with respect to the standards from which they were generated. (Distortions greater than 50% begin to approach the mirror image of the standard.) Matched quadruplets of subjects saw exactly the same distortions, in the same order, during the test phase. Across quadruplets, different presentation orders were used. Between the learning phase and the test phase, subjects were told to rate how much they liked each of the 10 test stimuli on an increasing scale from 1 to 7. Each of these stimuli was presented for 5 sec, after which the screen went blank and subjects were informed by a beep that they should respond by pressing one of seven labeled buttons, located on a response panel in front of them. The next stimulus was presented 2 sec after the subject's response.

Subjects. Forty subjects (20 in the abstraction condition and 20 in the repetition condition) were tested individually in sound-attenuating isolation chambers. They were recruited from the Human Performance Center subject pool and were paid $3 for their participation in a session that lasted approximately 15 minutes.

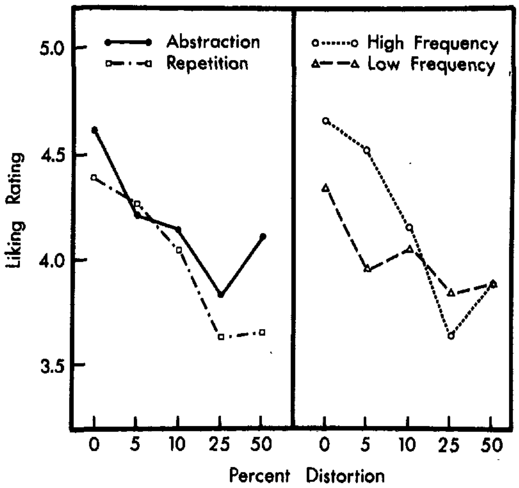

Figure 3. The affective generalization gradients obtained in Experiment 2.

Results and Discussion

A four-way ANOVA was performed on the subjects' affective ratings, treating the yoked quadruplets of subjects as matched observations. Level of distortion yielded the only significant main effect, $F(4, 36) = 4.87, p < .01$. Figure 3 depicts the effect of distortion level on liking ratings for the abstraction versus repetition conditions, collapsing over frequency level (left panel), and for high-frequency versus low-frequency stimulus categories, collapsing over training condition (right panel). As is clear from inspection of Figure 3, the overall effect of distortion level primarily consisted of a decrease in liking ratings with increasing distortion level, $t(36) = 4.39, p < .002$, by a linear-trend test. Separate linear-trend tests on the effect of distortion level were also performed for the abstraction and repetition conditions, because these constituted different ways of learning about the standard or prototype. The linear trend was significant for both the abstraction condition, $t(36) = 3.53, p < .005$, and the repetition condition, $t(36) = 5.32, p < .001$. There was no significant linear interaction of distortion level and condition, $t(36) = 1.26, p > .05$.

Separate linear-trend tests on the effect of distortion level were also performed for the high- and low-frequency conditions. (These tests did not include the 50% distortion level, because these stimuli were not statistically related to either the high- or low-frequency standards from which they were generated.) Both tests showed significant linear effects of distortion level: $t(36) = 3.42, p < .01$, for the high-frequency condition, and $t(36) = 2.52, p < .02$, for the low-frequency condition. The linear interaction of frequency and distortion level was also significant, $t(36) = 2.21, p < .05$, indicating that a stronger linear component was present in the high-frequency condition. This linear interaction did not vary significantly across the two training conditions, $t(36) = .05, p > .25$. As the data in the right panel of Figure 3 indicate, stimuli identical to the standard (0% distortion) or very similar to it (5% distortion) received higher liking ratings if they were derived from the high-frequency rather than the low-frequency standard, $t(2.48), p < .02$. This result is consistent with the positive effect of exposure frequency on liking observed in many previous mere-exposure studies (Harrison, 1977). As would be expected, the difference between the two frequency levels disappeared at higher distortion levels, at which the test stimuli were less similar to either of the two standards. It is interesting to note that the standard itself was liked as much by subjects in the abstraction condition, who saw it for the first time in the test phase, as it was by subjects in the repetition condition, who had seen it many times. This result suggests that the affective consequences of direct stimulus exposure can be equaled by indirect experience based on an implicit learning process.[3]

The results of Experiment 2 provide a second demonstration of affective generalization of the mere-exposure effect in an implicit learning paradigm, using very different stimuli than the letter strings employed in Experiment 1. Furthermore, the pattern of generalization was very orderly: Liking declined as distortion level increased. The pattern of generalization was similar regardless of whether prior experience with the standard was direct (repetition condition) or the result of implicit learning with a series of training exemplars (abstraction condition). In the latter case, the present monotonic decline in liking ratings with increasing distortion levels is similar to the effect of distortion level on classification judgments observed in many previous studies (e.g., Posner & Keele, 1968), including experiments that used stimuli very similar to the present ones (Fried & Holyoak, in press). After exposure to a set of training exemplars, classification accuracy typically declines as distortion level increases. As was the case in Experiment 1, the results of Experiment 2 suggest that mere exposure to a set of structured stimuli has similar consequences for affective and for cognitive judgments.

General Discussion

The experiments reported in this article constitute a preliminary investigation of the generalization of the affective consequences of exposure to stimuli. These experiments take as a model various paradigms developed by cognitive psychologists for studying concept learning. In particular, affective generalization was related to implicit learning, which has been described as a type of learning that takes place naturally and automatically when viewing a set of related stimuli (Reber, 1967).

In Experiment 1, the pattern of affective generalization matched subjects' cognitive discriminations regarding which stimuli obeyed an artificial grammar. This finding suggests a possible relation between affective and cognitive processes that differs from recent accounts. Zajonc (1980) argued that affect and cognition are separate processes and that affective responses occur before cognitive responses and do not depend on them. One obvious interpretation of the results of Experiment 1 is that the affective consequences of mere exposure were dependent on the cognitive processes of implicit learning. The preference for the novel grammatical letter strings would seem to require the inference that these strings obeyed the grammar. This interpretation conflicts with Zajonc's position, because it is not clear why a separate and independent affective system would rely on complex cognitive processes for a simple preference response.

That a possible dependence of affect on cognition in the present paradigm exists is not surprising when some of Zajonc's (1980) arguments are examined. Holyoak and Gordon (in press) criticized Zajonc's apparent characterization of cognitive processes as slow, conscious, and effortful, citing counterexamples from the literature on cognitive psychology. When a broader view of cognition is taken than that of Zajonc, many of his arguments for the independence of affect from cognition appear less strong. For example, the most dramatic evidence of independence cited by Zajonc (1980) is work showing that the mere-exposure effect develops under circumstances in which the subject is not aware of the stimulus and cannot later recognize it (Kunst-Wilson & Zajonc, 1980; Wilson, 1979). Zajonc argued that these studies demonstrate affective responding independent of cognitive processing. However, this argument is dependent on the identification of cognition with conscious recogni-

---

[3] Pilot data indicated that the generalization effects observed in Experiment 2 disappear entirely when the number of test trials is increased beyond the relatively small number (10) used here. Previous work on the effects of mere exposure also suggests that differences in liking ratings are attenuated with repeated testing (Harrison, 1977).

tion, a link that does not seem justified. For example, Fowler, Wolford, Slade, and Tassinary (1981) showed that very briefly presented words activate semantic memory (certainly a cognitive process), even if they are not consciously recognized. Affective responses may at times be based on more or less unconscious, automatic cognitive processes. This would explain why affective generalization in our study was closely related to implicit learning, which is in part automatic and unconscious.

Having made the above claims, we must acknowledge that it is not certain that the affective responses in our experiment were based on the cognitive processes of implicit learning. Although both cognitive and affective judgments discriminated grammatical from nongrammatical strings, we cannot be sure that they did so on the same basis. It is possible that the affective responses in our experiment were based on a different and separate system of learning that produced the same pattern of discrimination. Essentially, this would amount to a claim that there are two information processing systems—affective and cognitive—both of which induce some aspects of the grammar. Although this interpretation is unparsimonious, it cannot be ruled out on the basis of our data.

Further tests of the relation between affect and cognition are possible using the generalization techniques we have developed. Zajonc (1980) argued that the affective and cognitive systems respond to different features of the environment, cognition responding to "discriminanda" and affect to "preferenda." According to Zajonc, discriminanda are analytic features, whereas preferenda are "quite gross, vague and global" (although unfortunately he as yet "cannot be very specific about preferenda"; Zajonc, 1980, p. 159). If a more precise specification of the distinction between preferenda and discriminanda could be formulated, it might be possible to construct a set of stimuli with which one could determine whether in fact affective responses generalize along preferenda whereas cognitive judgments generalize along discriminanda. This would allow a test of whether affect and cognition respond to different aspects of the environment.

Experiment 2 allowed us to obtain affective generalization gradients. One reason these gradients (Figure 3) are of interest is that they provide a potential test of Berlyne's (1970) theory of affective responding. Berlyne (1970) argued that the arousal potential of a stimulus in part determines its ability to evoke pleasure (its "hedonic value"). He claims that hedonic value and arousal potential are related by an inverted-U function, with moderate levels of arousal having the highest hedonic value. In our experiments, repeated experience with the standards would be expected to have reduced their novelty and hence their arousal potential. The novelty of the distortions of the standard, used in the test phase, ought to have increased the farther they were from the standard. It follows that if subjects' experience with the standard during the exposure phase sufficiently reduced its arousal potential, then some nonzero distortion of the standard should have been liked more than the standard. However, as the data in Figure 3 indicate, this was not the case. Furthermore, Berlyne's theory predicts that the likelihood of finding the peak displaced from the standard ought to increase with more exposure to the prototype. This prediction was also not substantiated by our results; the right panel in Figure 3 in fact reveals a steeper slope for the high-frequency condition than for the low-frequency condition.

We do not regard Experiment 2 as a complete test of Berlyne's theory on this issue. Such a test would require the use of a greater range of stimulus complexities, exposure frequencies, and distortion levels. The present results are nonetheless suggestive and illustrate how the generalization techniques developed here could be used to test theories of affective responding, such as Berlyne's arousal-level theory or different-adaptation-level theories (see Harrison, 1977, for a discussion).

At an intuitive level, it seems that the affective responses to abstract visual stimuli of the sort studied here may be similar to the kind of emotion involved in the appreciation and enjoyment of art. It is quite likely that aesthetic appreciation involves something akin to implicit learning. Hartley and Homa (1981) studied the development of people's ability to recognize stylistic differences among painters. They concluded that this ability is influenced by the same kinds of variables as have been found to influence the recognition of artificial categories used in laboratory experiments. Meyer (1956) argued that aesthetic appreciation requires the ability to perceive structure in a stimulus. He claims that in music this structure is of two sorts. The first is defined by what the Gestalt psychologists called "good form." To the extent that a stimulus obeys the Gestalt rules of organization, one perceives structure in it and finds it pleasing. The second kind of structure develops through familiarity with a set of stimuli that share some abstract organizational principles. Once this structure is internalized, it then becomes a kind of good form to which new and old stimuli can be related. The present experiments support Meyer's (1956) conjecture that relating stimuli to a learned structure is pleasing.

References

Berlyne, D. E. Novelty, complexity and hedonic value. *Perception & Psychophysics,* 1970, *8,* 279–286.

Clark, H. H. The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior,* 1973, *12,* 335–359.

Fowler, C. A., Wolford, G., Slade, R., & Tassinary, L. Lexical access with and without awareness. *Journal of Experimental Psychology: General,* 1981, *110,* 341–362.

Fried, L. S., & Holyoak, K. J. Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* in press.

Harrison, A. A. Mere exposure. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10). New York: Academic Press, 1977.

Harrison, A. A., & Crandall, R. Heterogeneity–homogeneity of exposure sequence and the attitudinal effects of exposure. *Journal of Personality and Social Psychology,* 1972, *21,* 234–238.

Hartley, J., & Homa, D. Abstraction of stylistic concepts. *Journal of Experimental Psychology: Human Learning and Memory,* 1981, *7,* 33–46.

Holyoak, K. J., & Gordon, P. C. Information processing and social cognition. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition.* Hillsdale, N.J.: Erlbaum, in press.

Kunst-Wilson, W. R., & Zajonc, R. B. Affective discrimination of stimuli that cannot be recognized. *Science,* 1980, *207,* 557–558.

Meyer, L. B. *Emotion and meaning in music.* Chicago: University of Chicago Press, 1956.

Posner, M. I., & Keele, S. W. On the genesis of abstract ideas. *Journal of Experimental Psychology,* 1968, *77,* 353–363.

Reber, A. S. Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior,* 1967, *5,* 855–863.

Reber, A. S. Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory,* 1976, *2,* 88–94.

Reber, A. S., & Allen, R. Analogy and abstraction strategies in synthetic grammar learning: A functional interpretation. *Cognition,* 1978, *6,* 189–221.

Reber, A. S., Kassin, S. M., Lewis, S., & Cantor, G. W. On the relationship between implicit and explicit modes in the learning of a complex rule structure. *Journal of Experimental Psychology: Human Learning and Memory,* 1980, *6,* 492–502.

Reber, A. S., & Lewis, S. Implicit learning: An analysis of the form and structure of a body of tacit knowledge. *Cognition,* 1977, *5,* 333–361.

Wilson, W. R. Feeling more than we can know: Exposure effects without learning. *Journal of Personality and Social Psychology,* 1979, *37,* 811–821.

Zajonc, R. B. Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology Monograph,* 1968, *9*(2, Pt. 2).

Zajonc, R. B. Feeling and thinking: Preferences need no inferences. *American Psychologist,* 1980, *35,* 151–175.

Received September 15, 1982
Revision received December 15, 1982 ■