Journal of Experimental Social Psychology **36,** 660–676 (2000)
doi:10.1006/jesp.2000.1430, available online at http://www.idealibrary.com on IDEAL®

Electrophysiological Evidence of Implicit and Explicit Categorization Processes

Tiffany A. Ito

University of Colorado

and

John T. Cacioppo

University of Chicago

Received November 1, 1999; revised February 29, 2000; accepted February 29, 2000

Event-related brain potentials (ERPs), which provide on-line measures of categorization processes, were used to assess the implicit and explicit categorization of stimuli along evaluative (positive and negative) and nonevaluative (people and no-people) dimensions. Participants were exposed to stimuli that simultaneously varied along both dimensions, but were explicitly instructed to categorize along only one of them. Consistent with prior research, the late positive potential (LPP) of the ERP was sensitive to participants' explicit categorization task. However, the LPP also revealed implicit categorization along the non-task-relevant dimension. Additionally, there was evidence of an implicit negativity bias in which negative stimuli spontaneously received greater processing than positive stimuli and a people/no-people analog in which stimuli containing people received greater processing than stimuli without people. Results suggest the operation of adaptively beneficial implicit categorization processes, triggered by significant proximal stimuli, serving broad, cross-situational goals. © 2000 Academic Press

We thank Bernadette Park and Irene Blair for comments on a draft of this article and Dave Lozano and Daniel Litvack for technical assistance. Preparation of this article was supported by National Science Foundation Grant SBR-9512459 and National Institute of Mental Health Grant P50MH52384-01A1. Preliminary reports of this research were presented at the annual meeting of the Society for Psychophysiological Research, Denver, CO, September 1998.

Address correspondence and reprint requests to Tiffany A. Ito, Department of Psychology, University of Colorado, Boulder, CO 80309-0345, or John T. Cacioppo, Department of Psychology, University of Chicago, 5848 S. University Ave., Chicago, IL 60637. E-mail: tito@psych.colorado.edu or cacioppo@uchicago.edu.

0022-1031/00 $35.00

The many demands of everyday life preclude the intentional processing of every potentially important stimulus we encounter. Instead, we appear to possess the ability to monitor the environment and make note of potentially relevant stimuli, even when not explicitly intending to do so. A variety of such implicit categorization processes that presumably ease social processing have been demonstrated.[1] Among these are the encoding of social behaviors in terms of trait concepts (Gilbert, 1989; Gilbert, Pelham, & Krull, 1988; Winter & Uleman, 1984) and monitoring of the environment for self-relevant information (Bargh, 1982).

Research in the field of lie detection also indicates that personally significant information may be implicitly processed. In one such demonstration, participants enacted mock crime scenarios and were later given the task of distinguishing between two arbitrary, experimenter-defined classes of stimuli (Farwell & Donchin, 1991). Embedded within the stimuli were items related to participants' "crime." Event-related brain potential (ERP) responses differentiated crime-related and crime-irrelevant material, suggesting that even though participants were performing an unrelated classification task, they were nevertheless implicitly processing stimuli in terms of its relation to their guilt (see also Rosenfeld, Agnell, Johnson, & Qian, 1991).

The implicit processes described above are likely to be beneficial because they allow us to preattentively monitor the environment for stimuli with high informational value. Perhaps one of the most important classes of stimuli to monitor are threat-related and reward- or safety-related cues. Not surprisingly, extant research contains a range of evidence suggesting that we do implicitly evaluate items in our environment (Murphy & Zajonc, 1993; Niedenthal, 1990; Pratto & John, 1991). Much of this research has employed a sequential priming paradigm in which a valenced word is presented at short stimulus onset asynchronies before a target word. Participants' primary task is to make speeded evaluative judgments about the target word (Bargh, Chaiken, Govender, & Pratto, 1992; Bargh, Chaiken, Raymond, & Hymes, 1996; Fazio, Sanbonmatsu, Powell, & Kardes, 1986). Priming with a positive attitude object tends to facilitate responses to positive targets and inhibit responses to negative targets, whereas priming with a negative attitude object facilitates responses to negative targets and inhibits responses to positive targets, suggesting that participants implicitly evaluate the attitude primes, even though they have not been explicitly instructed to do so.

Although creative, the methodologies employed to capture implicit evaluations tend to rely on relatively indirect measures. The logic of these designs is that information implicitly processed in one context (e.g., a priming stimulus) will persist long enough to have an effect in a subsequent context (e.g., some reaction

[1] Following Wittenbrink, Judd, and Park (1997), we define an *explicit* process as requiring intentional, conscious control, whereas an *implicit* process occurs spontaneously in response to environmental cues.

time task; Bargh, 1989). These paradigms can be insensitive to implicit categorization if the effects fail to persist long enough to influence subsequent task performance or if extraneous sources of information intervene and affect performance on the subsequent task. Some priming paradigms also provide participants with an active evaluative set (e.g., Bargh et al., 1992; Fazio et al., 1986; Murphy & Zajonc, 1993; Niedenthal, 1990). Said differently, participants are often explicitly instructed to evaluate stimuli (but see Bargh et al., 1996). To be maximally adaptive, the ability to process the environment in terms of its evaluative implications should occur even in instances when we lack an explicit evaluative goal. Finally, implicit and explicit processes are typically assessed with different measures, leaving open the issue of the extent to which observed differences are due to process or measurement differences.

In the present experiment, event-related brain potentials (ERPs) were used to assess both implicit and explicit processing. The major advantage of ERP measures in this context is that they provide on-line assessments of information processing activities and offer the potential to use the same measure to assess both implicit and explicit processing. ERPs reflect the electrical activity of synchronously activated neural generators involved in information processing operations. One potential of the ERP in particular, the late positive potential (LPP), is sensitive to both nonevaluative and evaluative categorization processes, with the amplitude of the LPP reflecting the extent to which the particular categorization process involves context updating. In a frequently used nonevaluative categorization paradigm, participants are presented with a series of tones of two different pitches, with one of the tones presented much more frequently than the other. When participants attend to and categorize the tones, a larger amplitude LPP is elicited by the low probability stimulus, presumably because the processing of categorically inconsistent stimuli requires greater processing resources and results in an updating of the current representation of the environment (Donchin, 1981; Donchin & Coles, 1988). Conceptually similar results are obtained when the stimuli differ in hedonic tone and participants perform an evaluative categorization task. We have presented participants with sequences of valenced stimuli, with a particular valence being shown on the majority of the trials (Cacioppo, Crites, Berntson, & Coles, 1993; Crites, Cacioppo, Gardner, & Berntson, 1995; Ito, Larsen, Smith, & Cacioppo, 1998). Evaluatively inconsistent stimuli (e.g., a negative item embedded within a series of predominantly positive items) elicit a larger amplitude LPP relative to evaluatively consistent stimuli (e.g., a negative item embedded within a series of other negative items).

An important feature of the LPP is its differential sensitivity to the categorization as opposed to response selection or output processes (e.g., Allen, Iacono, & Danielson, 1992; Crites et al., 1995; Magliero, Bashore, Coles, & Donchin, 1984; Renault, Signoret, Debruille, Breton, & Bolgert, 1989). Using an evaluative task, Crites et al. (1995) recorded LPPs to positive, negative, and neutral stimuli embedded within sequences of positive context stimuli. One-half of the participants were asked to accurately report their evaluations, whereas others were instructed to misreport their evaluations to either negative or neutral stimuli. The misreport instructions had the intended effect on overt responses. The LPPs, however, were enhanced to evaluatively inconsistent as compared to consistent stimuli regardless of the accuracy of the overt report (see also, Gardner & Cacioppo, in press).

The sensitivity of ERPs to underlying categorization processes suggests that they may be especially useful in assessing implicit evaluative categorization processes. It is not clear, however, whether the LPP is sensitive to categorization processes that are not explicitly intended. In lie-detection paradigms, the LPP has been shown to differentiate between crime-relevant and crime-irrelevant stimuli (Farwell & Donchin, 1991; Rosenfeld et al., 1991). In these paradigms, participants typically perform a crime-irrelevant task; sensitivity to the crime-related material therefore suggests implicit processing. By contrast, Crites and Cacioppo (1996) found that the LPP varied as a function of explicit categorization processes only. Participants in that study were presented with stimuli that varied simultaneously along both a nonevaluative (vegetable and nonvegetable) and an evaluative dimension (positive and negative). One-half of the participants explicitly classified the stimuli along the nonevaluative dimension (i.e., is it a vegetable or nonvegetable), whereas the others did so along the evaluative dimension (i.e., is it liked or disliked). Analyses revealed an effect of the *explicit* categorization task only. The LPPs of participants performing the nonevaluative task were sensitive to variations in the semantic inconsistency of a given stimulus relative to the semantic category of the preceding stimuli, but did not vary as a function of evaluative inconsistency. Similarly, the LPPs of participants performing the evaluative task were sensitive to variations in the evaluative inconsistency of a given stimulus, but did not vary as a function of its semantic inconsistency.

One possible explanation for the lack of sensitivity to implicit processing in Crites and Cacioppo (1996) may have been the relatively low level of stimulus potency. This was a function of both the nature of the categorical dimension employed (i.e., vegetables/nonvegetables) and the presentation of the items in written form as opposed to, for example, pictures or actual instances of the items. If implicit categorization does occur spontaneously in response to external cues, it is reasonable to assume that this process is restricted to stimuli that exceed some threshold of importance or relevance, which stimuli in Crites and Cacioppo possibly failed to reach. To provide a more sensitive test of implicit categorization effects in the present experiment, we again exposed participants to stimuli that simultaneously varied along a nonevaluative and evaluative dimension. In this present experiment, however, we used emotionally evocative pictorial stimuli instead of mildly evocative words. We also changed the nonevaluative dimension to whether people are or are not present, which likely represents a more fundamental categorization (Kurbat & Farah, 1998). Thus, the purpose of this experiment was to investigate the extent to which implicit and explicit evaluative and nonevaluative processing occurs using socially meaningful dimensions and a measure that (a) provides an on-line processing assessment, (b) does not require an active set, and (c) can be used to measure both implicit and explicit processing.

METHOD

Participants

Twenty-three right-handed Ohio State University undergraduates (12 males and 11 females) participated for class credit.

Materials

Sixty-four pictures that simultaneously varied along an evaluative (positive and negative) and a nonevaluative (people and no-people) dimension were selected from the International Affective Picture System (Center for the Study of Emotion and Attention, 1995). There were 16 pictures each in the categories of (a) positive-people (e.g., a couple embracing), (b) positive-no-people (e.g., chocolate bar), (c) negative-people (e.g., mourners at a graveside), and (d) negative-no-people (e.g., littered beach). In addition to the valence and people/no-people status of the pictures, we choose stimuli so as to maximize evaluative extremity while also equating for arousal. Mean arousal ratings within all four categories were equal and the positive and negative pictures were equated for evaluative extremity based on normative data obtained from the same population in a previous term (Ito, Cacioppo, & Lang, 1998).[2]

Procedure

Stimuli were presented in sequences of six on a computer monitor. Although all participants viewed the same stimuli, 12 explicitly categorized the stimuli along the evaluative dimension by either silently counting the number of positive or negative pictures within each sequence and 11 explicitly categorized along the nonevaluative dimension by either silently counting the number of pictures with people or without people.

Electroencephalographic activity was recorded during the presentation of a single *target* stimulus within each six-stimulus sequence. We refer to the non-target stimuli as the *context.* Targets were of 1 of the 16 types described in Table 1. As can be seen, target stimuli were either (a) evaluatively and nonevaluatively consistent with the context (types 1, 6, 11, and 16); (b) evaluatively consistent but nonevaluatively inconsistent (types 2, 5, 12, and 15); (c) evaluatively inconsistent

[2] For the four categories of pictures, mean normative arousal ratings were (in the order listed in the method section) 5.36, 5.35, 5.60, and 5.64 and mean normative valence ratings were 7.55, 7.49, 2.39, and 2.46. All ratings were made on 9-point scales, with higher ratings indicating greater arousal and greater positivity. Separate 2 (Picture Valence; positive, negative) $\times$ 2 (Picture Content: people, no-people) analyses of variance on normative arousal and valence ratings revealed no significant differences between the categories of pictures in either normative arousal or evaluative extremity, with valence ratings for negative pictures reverse scored so they could be compared to positive ones for evaluative extremity.

TABLE 1
Types of Six-Picture Sequences Presented to Participants

Type of sequence	Context	Target stimulus
	Block A	
1	Positive people	Positive people
2	Positive people	Positive nonpeople
3	Positive people	Negative people
4	Positive people	Negative nonpeople
	Block B	
5	Positive nonpeople	Positive people
6	Positive nonpeople	Positive nonpeople
7	Positive nonpeople	Negative people
8	Positive nonpeople	Negative nonpeople
	Block C	
9	Negative people	Positive people
10	Negative people	Positive nonpeople
11	Negative people	Negative people
12	Negative people	Negative nonpeople
	Block D	
13	Negative nonpeople	Positive people
14	Negative nonpeople	Positive nonpeople
15	Negative nonpeople	Negative people
16	Negative nonpeople	Negative nonpeople

Note. Data were recorded from a single target position within each sequence. On the one-half of trials for which a picture from the same category as the target pictures was shown later in the sequence, data were recorded during the first presentation from the target category.

but nonevaluatively consistent (types 3, 8, 9, and 14); and (d) evaluatively and nonevaluatively inconsistent (types 4, 7, 10, and 13). Across the experiment, each picture appeared an equal number of times in the target and context positions.

To prevent participants from anticipating when an evaluatively or nonevaluatively inconsistent picture might be shown, the target stimulus randomly appeared in either the third, fourth, or fifth position in a sequence. To further reduce predictability, a picture drawn from the same category as the target picture was shown later in the sequence during one-half of the trials. Sequences were arranged in the four separate blocks shown in Table 1 (A, B, C, and D) such that the majority of the pictures within a block were drawn from the same category (e.g., positive people for Block A). Twenty instances of each sequence type were shown in a block. To decrease predictability, half of these sequences within a block included just the one picture from the target category and the other half included a second picture drawn from the same category as the target picture shown later in the sequence. Order of the four blocks was counterbalanced across participants. Each picture was shown for 1000 ms, followed by a 1200-ms intertrial interval. Participants were prompted to write down their count after each sequence of six stimuli and then pressed a button to initiate presentation of the next sequence of pictures.

For each participant, we averaged across all instances of a particular sequence type. This procedure was performed separately on data recorded from midline frontal (Fz), central (Cz), and parietal (Pz) locations. Prior research and visual inspection of data from the present experiment indicate that the LPP peaks between 400 and 900 ms after stimulus onset. As a result, the LPP was quantified within each averaged ERP waveform as the largest positive deflection in that latency window. Data acquisition and reduction closely followed those described in greater detail in Ito et al. (1998) and Crites et al. (1995), respectively.

RESULTS

Within each task condition, there were no significant effects of the specific counting task performed (i.e., whether participants in the evaluative categorization condition were counting positive or negative stimuli and whether those in the nonevaluative categorization condition were counting people or no-people stimuli). This factor was therefore dropped from subsequent analyses. Consistent with prior research (e.g., Crites et al., 1995; Donchin, 1981), preliminary analyses involving sagittal scalp site revealed a parietal maximum distribution for the LPP, with larger LPPs at the Pz site. Effects of the experimental variables were also biggest at Pz. For ease of presentation, we therefore report only analyses on data from the Pz electrode.

Prior research showing the LPP's sensitivity to explicit evaluative and nonevaluative categorization tasks led us to expect LPP amplitude to vary as a function of a target picture's inconsistency with context stimuli along the task-relevant dimension. In addition, if the LPP is also sensitive to implicit processing effects, it should vary as a function of inconsistency along the non-task-relevant dimension. The effects of evaluative and nonevaluative inconsistency are evaluated by looking within a given context and comparing responses to contextually inconsistent and consistent pictures. For example, evaluative inconsistency effects within the positive context would be assessed by comparing responses to inconsistent negative to consistent positive pictures.

We performed a 2 (Target Category: people/no-people) $\times$ 2 (Context Category: people/no-people) $\times$ 2 (Target Valence: positive/negative) $\times$ 2 (Context Valence: positive/negative) $\times$ 2 (Task: evaluative/nonevaluative) multivariate analysis of variance on the LPP amplitudes, with all but the last factor within-subject. Main effects of Target Category, $F(1, 21) = 9.39$, $p < .01$, and of Target Valence, $F(1, 21) = 28.14$, $p < .001$, were obtained. These revealed larger LPPs to pictures of people ($M = 9.57\ \mu V$) than of nonpeople ($M = 8.26\ \mu V$) and to negative pictures ($M = 9.95\ \mu V$) than to positive ($M = 7.88\ \mu V$)

pictures. The greater sensitivity to negative stimuli replicates prior research on explicit evaluative processes demonstrating a processing bias in favor of negative stimuli (Cacioppo & Berntson, 1994; Ito et al., 1998). The greater sensitivity to stimuli containing people suggests an analogous people-based processing bias.

These processing biases in favor of people and negative stimuli also combined in a Target Category × Target Valence interaction, $F(1, 21) = 6.68$, $p < .05$. To probe the nature of this interaction, three linear contrasts were performed. The first compared LPPs to stimuli that possessed a single biasing feature (i.e., either people or negative valence), revealing equally large LPPs to positive people ($M = 9.79\ \mu$V) and to negative nonpeople ($M = 9.03\ \mu$V), $F(1, 22) = 1.60$, $p > .20$. The second contrast compared responses to these same two classes of stimuli to stimuli that contained a conjunction of both people and negative valence. This comparison revealed equally large LPPs to stimuli containing a single feature ($M = 9.41\ \mu$V) as to stimuli containing both features (i.e., negative people, $M = 10.12\ \mu$V), $F(1, 22) = 2.90$, $p > .10$. The final contrast revealed that stimuli containing *either* a single *or* both features (i.e., negative people, positive nonpeople, and negative people, $M = 9.65\ \mu$V) elicited significantly larger LPPs than stimuli containing neither feature (i.e., positive nonpeople, $M = 6.72\ \mu$V), $F(1, 22) = 41.21$, $p < .001$. Together, this pattern of results suggests that processing biases in favor of people and negatively valenced items operate in a conjunctive manner, with equally large responses to stimuli possessing one or both features.

In addition to the general sensitivity to people and negative stimuli, the LPP was sensitive to *changes* along the nonevaluative dimension. Consistent with prior research showing the LPP's sensitivity to nonevaluative inconsistency effects, a significant Target Category × Context Category interaction was obtained, $F(1, 21) = 36.28$, $p < .001$. Pairwise comparisons revealed significantly larger LPPs to nonevaluatively inconsistent pictures within both contexts. Within the nonpeople context, people pictures elicited larger LPPs than nonpeople pictures [$M = 11.18$ and 6.76 μV, respectively, $F(1, 21) = 47.53$, $p < .001$]. Similarly, within the people context, nonpeople pictures elicited larger LPPs than people pictures [$M = 9.75$ and 7.96 μV, respectively, $F(1, 21) = 6.54$, $p < .05$].

The effects of nonevaluative inconsistency were also moderated by task condition: Target Category × Context Category × Task interaction, $F(1, 21) = 9.48$, $p < .01$. Figure 1 shows the grand average waveforms as a function of presence or absence of people in the target and context stimuli, with *explicit nonevaluative* categorization effects in the top panel (i.e., participants performing the nonevaluative task) and *implicit nonevaluative* categorization effects in the bottom panel (i.e., participants performing the evaluative categorization task). Simple effects analyses revealed significant Target Category × Context Category interactions for both task conditions: $F(1, 10) = 36.04$, $p < .001$, for nonevaluative task participants, and $F(1, 11) = 5.00$, $p < .05$ for evaluative task participants. We first examined the LPP's sensitivity to *explicit nonevalu-*

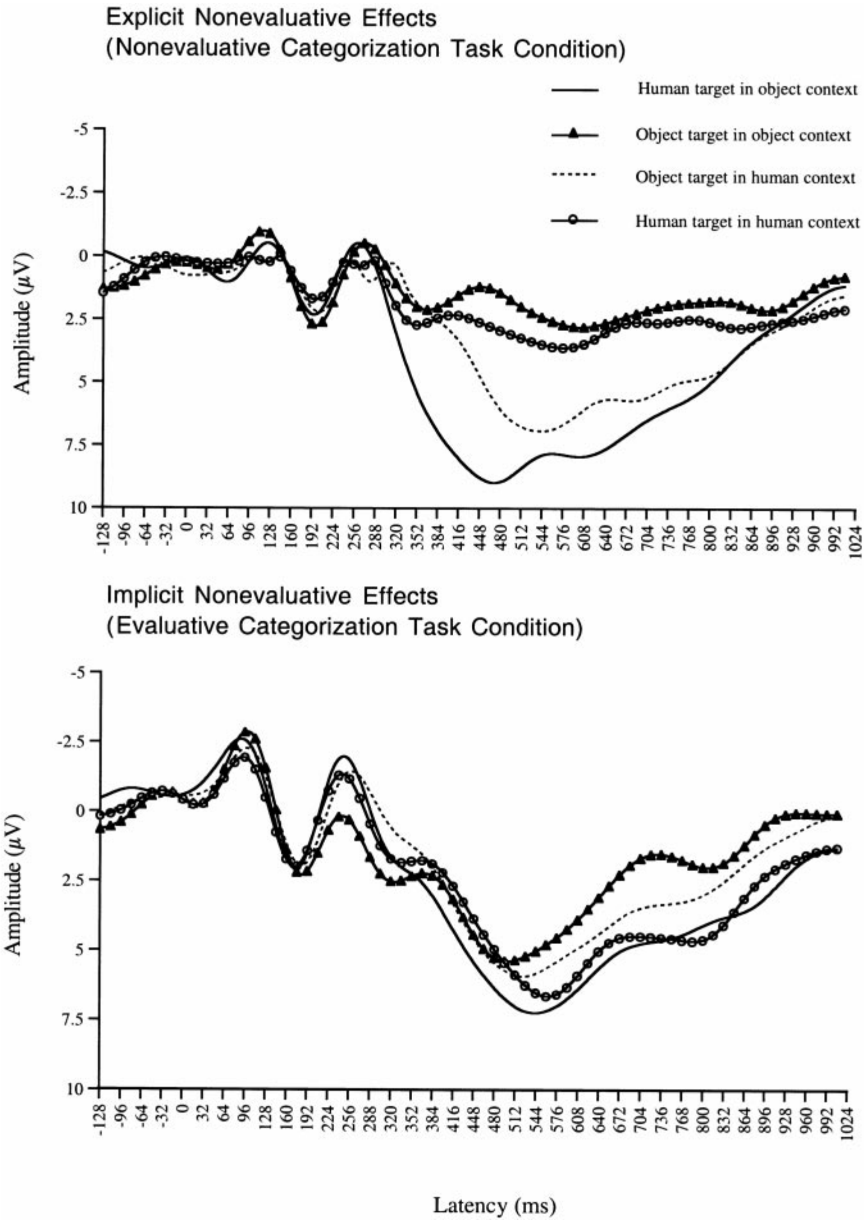

FIG. 1. Averaged event-related potential waveforms at electrode Pz as a function of the presence or absence of people in target and context stimuli. The top panel depicts explicit nonevaluative categorization effects (data from participants in the nonevaluative task condition). The bottom panel depicts implicit nonevaluative categorization effects (data from participants in the evaluative task condition). The late-positive potential is the positive (downward) deflection peaking at approximately 450–550 ms.

TABLE 2
Mean LPP Amplitude as a Function of the Presence or Absence of People in Target and Context Stimuli

Context category	Target category: People	Target category: No-people
Explicit nonevaluative effects (nonevaluative categorization task condition)		
People	6.34_{a} (1.02)	9.26_{b} (1.37)
No-people	11.73_{c} (1.73)	5.24_{a} (0.87)
Implicit nonevaluative effects (evaluative categorization task condition)		
People	9.61_{ab} (1.31)	10.21_{ab} (1.30)
No-people	10.73_{b} (1.21)	8.29_{a} (1.37)

Note. Values are in microvolts. Standard errors of the mean are in parentheses. Within each categorization task condition, means with different subscripts differ at $p < .05$.

ative categorization effects by performing follow-up pairwise comparisons among the nonevaluative task participants. Within both contexts, nonevaluatively inconsistent targets elicited larger LPPs than consistent pictures (Table 2, top panel).[3] Furthermore, comparing responses to people and nonpeople when each was inconsistent with its context revealed larger LPPs to people targets within the context of nonpeople than to equally rare nonpeople targets within the context of people. This effect further demonstrates a processing bias in favor of people. *Implicit nonevaluative* categorization effects were also observed among participants performing the evaluative task within the nonpeople context (Table 2, bottom panel), with inconsistent people pictures associated with larger LPPs than consistent nonpeople pictures. LPPs were equally large to inconsistent and consistent pictures within the people context.

The preceding effects show that all participants were sensitive to some degree to variations along the nonevaluative dimension. The LPP was also sensitive to the evaluative dimension, as shown in a Target Valence × Context Valence interaction, $F(1, 22) = 26.90$, $p < .001$. When presented with a context of positive pictures, evaluatively inconsistent negative pictures were associated with larger LPPs than consistent positive pictures [$M = 11.19$ and 6.78 μV, respec-

[3] Peak amplitudes in the grand average waveforms in Figs. 1 and 2 can appear smaller than the mean amplitudes reported in Tables 2 and 3 due to latency jitter, which occurs when the peak amplitude of the LPP differs among participants. The mean amplitudes in Tables 2 and 3 are not affected by latency jitter because they are measured at the peak response for each individual. The grand average waveforms, however, are aligned at picture onset, which can smear the LPP peak when latency jitter exists. Both representations of the data provide important information, with peak amplitudes providing the best characterization of the size of the LPP response, and the grand average waveforms providing a summary characterization of the shape of the overall waveform across all participants.

tively, $F(1, 21) = 40.38$, $p < .001$]. In the negative context, evaluatively inconsistent positive and consistent negative pictures produced equally large LPPs [$M = 8.97$ and 8.72 μV, respectively, $F(1, 21) < 1$]. The absence of evaluative inconsistency effects in the negative context may be a manifestation of the negativity bias, directing processing to negative pictures even when they are evaluatively consistent with the context.

The evaluative inconsistency effects were also moderated by task condition, Target Valence × Context Valence × Task interaction, $F(1, 22) = 4.43$, $p < .05$. The grand average waveforms as a function of target and context valence are shown in Fig. 2, with *explicit evaluative* effects in the top panel (i.e., participants performing the evaluative task) and *implicit evaluative* effects in the bottom panel (i.e., participants performing the nonevaluative task). Simple effects analyses revealed significant Target Valence × Context Valence interactions for both task conditions: $F(1, 11) = 17.23$, $p < .005$, for evaluative task participants, and $F(1, 10) = 13.95$, $p < .005$, for nonevaluative task participants. Pairwise comparisons revealed similar effects at the explicit and implicit level (Table 3). First, in the positive context, evaluatively inconsistent negative targets elicited larger LPPs than did consistent positive pictures. Second, within the negative context, LPPs were equally large to evaluatively inconsistent positive and consistent negative pictures. Third, comparing responses to negative and positive pictures when each was inconsistent with its context revealed larger LPPs to negative targets within the context of positive pictures than to equally rare positive targets within the context of negative pictures. These latter two effects are consistent with a negativity bias that increases responses to negative stimuli both when they are evaluatively consistent and inconsistent with their surrounding context.

DISCUSSION

Capitalizing on the LPP's demonstrated sensitivity to categorization processes, we recorded ERPs as participants were exposed to stimuli that simultaneously varied strongly along both a nonevaluative and an evaluative dimension. Doing so allowed us to eliminate the need to infer implicit and explicit processing from two different measures or to provide participants with an active set that was consistent with the implicit categorization process of interest. Consistent with prior research, LPP amplitude varied as a function of the dimension on which participants were explicitly categorizing the stimuli. Among those performing the nonevaluative categorization task, LPP amplitude was larger to categorically inconsistent as compared to consistent stimuli in both the people and nonpeople context. Among those performing the evaluative task, LPP amplitude was larger to evaluatively inconsistent as compared to consistent stimuli in the positive context.

These data also replicate prior research demonstrating a negativity bias (Ito et al., 1998). This was evidenced in a main effect of Target Valence, with negative stimuli eliciting larger LPPs than positive ones, regardless of task condition or

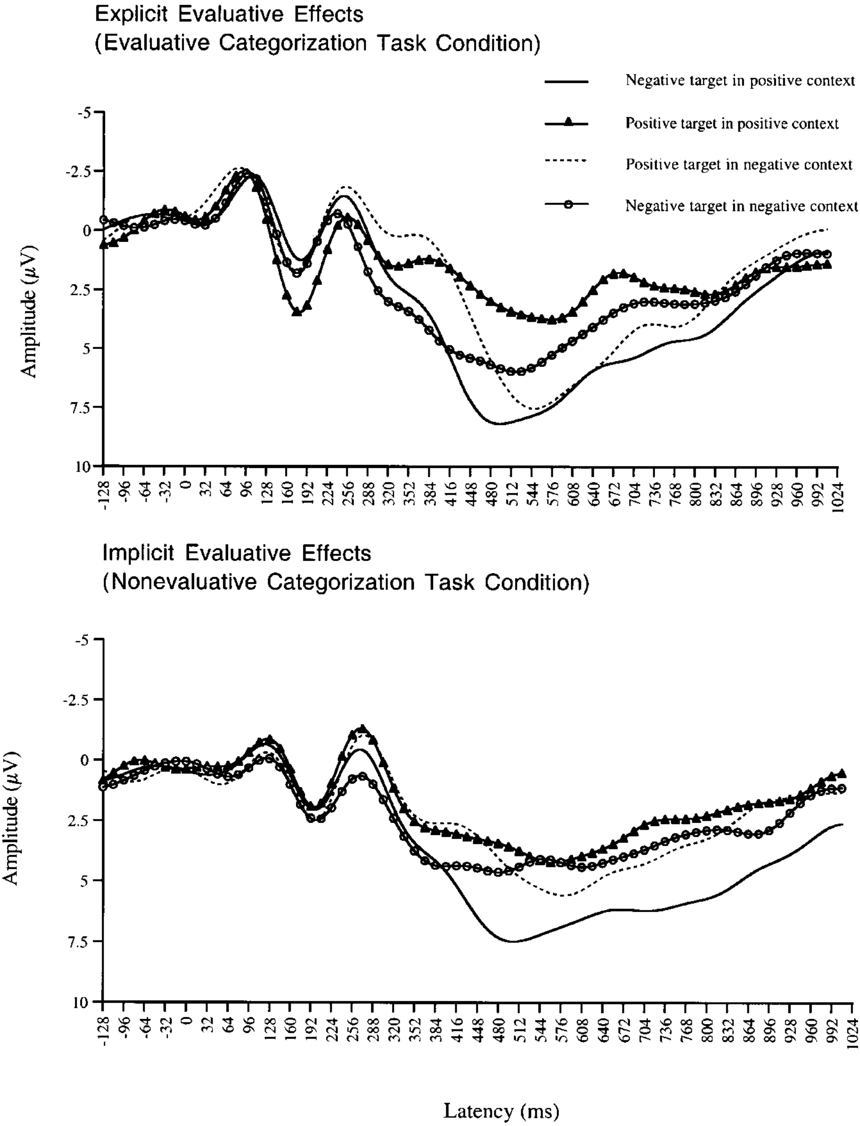

FIG. 2. Averaged event-related potential waveforms at electrode Pz as a function of target and context valence. The top panel depicts explicit evaluative categorization effects (data from participants in the evaluative task condition). The bottom panel depicts implicit evaluative categorization effects (data from participants in the nonevaluative task condition). The late-positive potential is the positive (downward) deflection peaking at approximately 450–550 ms.

evaluative consistency. In addition, in both task conditions, LPPs were larger to negative than positive targets when each was embedded within an evaluatively inconsistent context. A steeper gradient for the currency function underlying

TABLE 3

Mean LPP Amplitude as a Function of the Valence of Target and Context Stimuli

	Target valence	
Context valence	Negative	Positive
Explicit evaluative effects (evaluative categorization task condition)		
Negative	9.42_a (1.19)	10.14_a (1.23)
Positive	12.54_b (1.50)	6.74_c (1.30)
Implicit evaluative effects (nonevaluative categorization task condition)		
Negative	8.05_a (1.31)	7.78_a (1.33)
Positive	9.89_b (1.20)	6.84_a (1.00)

Note. Values are in microvolts. Standard errors of the mean are in parentheses. Within each categorization task condition, means with different subscripts differ at $p < .05$.

negative as compared to positive evaluative responses is thought to produce this effect (Cacioppo, Berntson, & Gardner, 1997). Interestingly, the main effect of target category revealed a processing bias in favor of stimuli containing people. In addition, within the nonevaluative task condition, pictures of people elicited larger LPPs than nonpeople pictures when each was inconsistent with its context. Together, these results are consistent with the notion that people in our environment receive greater processing resources (cf. Farah, Wilson, Drain, & Tanaka, 1998). The presence of negative and people features also appeared to operate conjunctively such that the presence of either or both of these features attracted equally high processing resources.

Whereas categorization processes clearly can arise out of the conscious goal of classifying stimuli along some particular dimension, the present results suggest that they can also occur more spontaneously, triggered by relevant stimuli in the environment. Participants for whom only the evaluative dimension was task-relevant were nevertheless sensitive to nonevaluative differences in the stimuli. In the nonpeople context, inconsistent people pictures elicited larger LPPs than consistent nonpeople pictures. At the explicit level, however, nonevaluative inconsistency effects were significant in the people context as well as in the nonpeople context. Examination of the relevant means reveals that the absence of this effect at the implicit level may be due to a people processing bias. When nonevaluative categorization was implicit and the context was people, relatively large LPPs were observed to both inconsistent nonpeople pictures ($M = 10.21$ μV) and consistent people pictures ($M = 9.61$ μV). The processing bias in favor of people may have increased responses to the nonevaluatively consistent people pictures, thereby obscuring nonevaluatively inconsistency effects in this condition. An examination of Table 2 reveals a similar trend when categorization was explicit; the biggest effects occurred when nonevaluative inconsistency and "people-ness" covaried. Overall, the results suggest similar forms of processing at both the implicit and explicit levels. At both levels, participants displayed a people bias and were sensitive to changes along the nonevaluative dimension to varying degrees.

For evaluative processing, the similarity between explicit and implicit effects was striking. In both task conditions, (a) within the positive context, evaluatively inconsistent negative targets elicited larger LPPs than did consistent positive pictures; (b) within the negative context, LPPs were equally large to evaluatively inconsistent positive and consistent negative pictures; and (c) negative pictures were associated with larger LPPs than were positive pictures when each was inconsistent with its context.[4] As we have discussed, these latter two effects are consistent with a negativity bias that increases responses to negative stimuli both when they are evaluatively consistent and inconsistent with their surrounding context, and their presence at the implicit level suggests that the negativity bias can operate at the level of automatic processing.

Equally large LPPs to positive and negative pictures in the negative context have not always been found. In other studies, evaluative inconsistency effects have been obtained with positive targets embedded in a negative context (Cacioppo et al., 1995). These prior demonstrations have varied the valence of the context between subjects, whereas context valence varied within subject in the present experiment. The difference in results suggest that the negativity bias may habituate with longer periods of negative context but that the strength of this response had not yet attenuated in the shorter blocks of negative context necessitated by the within-subject nature of the present design. An additional difference is the use of relatively impactful color photos in the present experiment as opposed to word stimuli in Cacioppo et al. (1995). The potency of the stimuli may also affect the rate of habituation for the negativity bias.

The strength of the negativity and people biases at the implicit level suggest that these classes of stimuli attract greater processing resources, both when we intend to evaluate objects and events and to identify the presence of people and when such explicit goals are absent. The negativity bias has been described as a tendency for the negative evaluative system to respond more intensely than the positive evaluative system to equally activating evaluative input, serving adaptive purposes by ensuring appropriately strong responses to proximate negative cues (Cacioppo & Berntson, 1994). The ability to do so spontaneously given only

[4] We also analyzed data from six homologous scalp sites over the left and right hemispheres. Consistent with prior research showing a right-hemisphere specialization for evaluative processing (Cacioppo, Crites, & Gardner, 1996), the LPP was larger over right as compared to left scale sites. Importantly, this scalp distribution did not interact with task condition, indicating that the set of neural generators underlying implicit and explicit evaluative processing was similar. This further supports the conclusion that nonevaluative task participants performed implicit evaluative categorizations by showing engagement of right-lateralized neural sources associated with affect and evaluative judgments among these participants. Greater detail on the laterality analyses is available from the authors.

relevant environmental cues likely increases the adaptive utility of the negativity bias by allowing organisms to avoid harm even when they are not explicitly sensitized to do so (Pratto & John, 1991). The greater processing resources dedicated to processing people suggests a similar beneficial mechanism.

The sensitivity of the LPP to implicit variations in the nonevaluative dimension is consistent with prior results, such as those obtained in lie-detection paradigms (Farwell & Donchin, 1991; Renault et al., 1991). The present results extend the findings of implicit nonevaluative classification to a more universal social dimension, that of distinguishing people from nonpeople. The implicit evaluative categorization effects observed here are consistent with Bargh et al. (1996) in showing that an active evaluative set is not necessary in order for implicit evaluative processing to occur. The present results are also the first, to our knowledge, to demonstrate the LPP's sensitivity to implicit evaluative processing. In a prior experiment assessing implicit evaluative categorization, participants classified familiar words (e.g., lettuce) as vegetable or nonvegetable (Crites & Cacioppo, 1996). Although the stimuli differed in valence (i.e., some of the food items were liked by the participants and others were disliked), these variations were modest. The LPP was sensitive to the explicit, nonevaluative categorization task, but did not vary as a function of item valence. The divergence between the Crites and Cacioppo results and the present results suggests that the implicit evaluative categorization effects we assessed are somewhat selective, occurring with the relatively more potent visual stimuli used in the present experiment. Other factors such as attitude accessibility (Fazio et al., 1986) and consistency (Bargh et al., 1992) may also influence the likelihood of implicit categorization.

Prior research has shown that the LPP is sensitive to the underlying categorization and not to the overt response. By demonstrating sensitivity to spontaneously activated evaluations and judgments, the present results suggest that the LPP may be useful not only in marking evaluative and nonevaluative processes that an individual may not want to report (Crites et al., 1995; Farwell & Donchin, 1991; Rosenfeld et al., 1991), but also in exploring responses that we might not even intend. Our results also reveal two interesting biases such that negative and people stimuli tended to receive greater processing resources. Thus, not only do we appear to spontaneously attend to certain classification dimensions, but within these dimensions, there are certain classes of stimuli that are especially important.

REFERENCES

Allen, J. J., Iacono, W. G., & Danielson, K. D. (1992). The identification of concealed memories using event-related potentials and implicit behavioral measures: A methodology for prediction in the face of individual differences. *Psychophysiology,* **29,** 504–522.

Bargh, J. A. (1982). Attention and automaticity in the processing of self-relevant information. *Journal of Personality and Social Psychology,* **43,** 425–436.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J.A. Bargh (Eds.), *Unintended thought* (pp. 3–51). New York: Guilford.

Bargh, J. A., Chaiken, S., Govender, T., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology,* **62,** 893–912.

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

Cacioppo, J. T., & Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. *Psychological Bulletin,* **115,** 401–423.

Cacioppo, J. T., Crites, S. L., Jr., Berntson, G. G., & Coles, M. G. H. (1993). If attitudes affect how stimuli are processed, should they not affect the event-related brain potential? *Psychological Science,* **4,** 108–112.

Center for the Study of Emotion and Attention [CSEA-NIMH] (1995). *The international affective picture system: Digitized photographs.* Gainesville, FL: The Center for Research in Psychophysiology, University of Florida.

Crites, S. L., Jr., & Cacioppo, J. T. (1996). Electrocortical differentiation of evaluative and nonevaluative categorizations. *Psychological Science,* **7,** 318–321.

Crites, S. L., Jr., Cacioppo, J. T., Gardner, W. L., & Berntson, G. G. (1995). Bioelectrical echoes from evaluative categorizations: II. A late positive brain potential that varies as a function of attitude registration rather than attitude report. *Journal of Personality and Social Psychology,* **68,** 997–1013.

Donchin, E. (1981). Surprise! · · · Surprise? *Psychophysiology,* **18,** 493–513.

Donchin, E., & Coles, M. G. H. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences,* **11,** 354–356.

Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is "special" about face perception. *Psychological Bulletin,* **105,** 482–498.

Farwell, L. A., & Donchin, E. (1991). The truth will out: Interrogative polygraphy ("lie detection") with event-related brain potentials. *Psychophysiology,* **28,** 531–547.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology,* **50,** 229–238.

Gardner, W. L., & Cacioppo, J. T. (in press). A brain-based index of evaluation. *Social Cognition.*

Gilbert, D. T. (1989). Thinking lightly about others: Automatic components of the social inference process. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 189–211). New York: Guilford.

Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busyness: When persons perceivers meet persons perceived. *Journal of Personality and Social Psychology,* **54,** 733–740.

Ito, T. A., Cacioppo, J. T., & Lang, P. J. (1998). Eliciting affect using the International Affective Picture System: Bivariate evaluation and ambivalence. *Personality and Social Psychology Bulletin,* **24,** 856–879.

Ito, T. A., Larsen, J. T., Smith, N. K., & Cacioppo, J. T. (1998). Negative information weighs more heavily on the brain: The negativity bias in evaluative categorizations. *Journal of Personality and Social Psychology,* **75,** 887–900.

Kurbat, M. A., & Farah, M. J. (1998). Is the category-specific deficit for living things spurious? *Journal of Cognitive Neuroscience,* **10,** 355–361.

Magliero, A., Bashore, T. R., Coles, M. G. H., & Donchin, E. (1984). On the dependence of P300 latency on stimulus evaluation processes. *Psychophysiology,* **21,** 171–186.

Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology,* **64,** 723–739.

Niedenthal, P. M. (1990). Implicit perception of affective information. *Journal of Experimental Social Psychology,* **26,** 505–527.

Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology,* **61,** 380–391.

Renault, B., Signoret, J.-L., Debruille, B., Breton, F., & Bolgert, F. (1989). Brain potentials reveal covert facial recognition in prosopagnosia. *Psychophysiology,* **27,** 905–912.

Rosenfeld, J. P., Angell, A., Johnson, M., & Qian, J. (1991). An ERP-based, control-question lie detector analog: Algorithms for discrimination effects within individuals' average waveforms. *Psychophysiology,* **28,** 319–335.

Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of trait inference. *Journal of Personality and Social Psychology,* **47,** 237–252.

Wittenbrink, B., Judd, C. M., & Park, B. (1997) Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology,* **72,** 262–274.