Bidirectional associations

MIRKO GEROLIN and HELENA MATUTE
Universidad de Deusto, Bilbao, Spain

Associations are usually assumed to have only a predictive directionality, from Event 1 (E1) to Event 2 (E2). However, this unidirectionality, if it really exists, could be specific to the use of conditioned responses as the index of learning (i.e., conditioned responses are ordinarily elicited only when a conditioned stimulus [CS] predicts an unconditioned stimulus [US]). The present research used neutral stimuli, devoid of CS-US directionality, in order to test whether the underlying associative mechanism was unidirectional in nature. We used colors and figures as E1s and E2s. In Experiment 1, human subjects saw the events in the E1→E2 direction during training. In the *consistent* group, the same E1 was always followed by the same E2; in the *inconsistent* group, E1s and E2s were paired randomly. In a subsequent test phase, we presented an E2 and subjects had to judge whether a particular E1 was associated with it. The judgments of the consistent group were higher than those of the inconsistent group, thereby suggesting that the association that the consistent subjects had learned was bidirectional. Experiments 2 and 3 confirmed the results of Experiment 1 while controlling for some alternative interpretations based on the representation-mediated formation of associations.

In contrast to the predictions of the older contiguity theories, current theories of associative learning generally assume that the establishment of an association between two events requires not only that the two events occur contiguously in time and space, but also that the event presented first (E1) predicts the event presented second (E2). This is consistent with prevailing explanations of several well-established findings such as the degraded contingency effect (Rescorla, 1968) and cue-competition effects (Kamin, 1968), as well as with the weak responding that is generally observed after simultaneous (E1-E2) and backward (E2→E1) pairings as opposed to forward pairings (E1→E2). Consequently, most current theories of associative learning (e.g., Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972) implicitly assume that associations are unidirectional (from E1 to E2).

The alternative possibility, that associations are not predictive, is an old idea, but it has received little attention during recent decades. For example, Beritov (1924) suggested that a bidirectional association was formed between the conditioned stimulus (CS) and the unconditioned stimulus (US) when a typical forward training procedure (i.e., CS→US) was used. Similarly, researchers working in the human paired associate tradition concluded that a bidirectional association between two events (E1 and E2) was possible. For example, in a very influential paper,

Asch and Ebenholtz (1962) presented evidence for the principle of associative symmetry, which they stated as follows: "When an association is formed between two distinct terms, a and b, it is established simultaneously and with equal strength between b and a" (p. 136). The topic of backward associations was a very popular one at that time, as reflected by the many experimental and review articles on bidirectional associations that were published in the sixties (e.g., Ekstrand, 1966; Horowitz, Brown, & Weissbluth, 1964; Murdock, 1962). Indeed, the discussion back then was not about whether backward associations existed, but about whether that they were as strong as forward associations. Moreover, researchers in Pavlovian conditioning long ago demonstrated the existence of backward excitatory conditioning in experiments in which the CS was seen to produce a strong conditioned response after one US→CS training trial (see Spetch, Wilkie, & Pinel, 1981, for a review). Those data on backward excitatory conditioning could be interpreted in terms of a bidirectional association between the CS and the US. If a bidirectional association was formed between the US and the CS during the US→CS training trials, this could explain why subjects were able to use this information in the opposite direction when the CS was presented at test. However, because backward excitatory responding normally vanishes after extended training, explanations in terms of backward associations have been questioned (see Spetch et al., 1981, for a review). Although recent developments in backward conditioning have provided some clearer data supporting the existence of bidirectional associations (e.g., Matzel, Held, & Miller, 1988), most of those experiments have been subject to alternative interpretations in terms of forward associations formed between the CS and the US (see our General Discussion, as well as Ward-Robinson & Hall, 1996, for more extensive discussion on these alternative views).

Support for this research was provided by Grant PB95-0440 from Dirección General de Enseñanza Superior (Spain) and Grant PI96-006 from Departamento de Educación, Universidades e Investigación (Basque Government) to H.M. We thank Geoffrey Hall and Jasper Ward-Robinson for illuminating discussions and Ralph Miller, Nuria Ortega, Oskar Pineño, and Sonia Vegas for their helpful comments on an earlier version of this manuscript. Correspondence concerning this article should be addressed to H. Matute, Departamento de Psicología, Universidad de Deusto, Apartado 1, 48080 Bilbao, Spain (e-mail: univem05@sarenet.es).

It should be noted, however, that most current studies of associative learning use conditioned stimuli (CSs) as E1s, unconditioned stimuli (USs) as E2s, and the development of conditioned responses as the index of learning. Even though a predictive CS-US relation seems to be necessary ordinarily for the CS to produce a conditioned response (and thus to obtain evidence of learning in those studies), this does not necessarily imply that the underlying associative structure is unidirectional. For example, according to Miller and Barnet (1993), associations are bidirectional, but they can be expressed as conditioned responses only when the CS predicts the US. In a similar vein, Rescorla (1980) noted that when neutral events were used as E1 and E2 (i.e., sensory preconditioning), simultaneous pairings produced stronger associations than did forward pairings. This is consistent with contiguity theory and suggests that predictiveness is not necessary for the acquisition of associations.

Thus, a relevant question is still whether the generally assumed E1→E2 forward directionality of associations is actually something inherent to the underlying associative process, or, by contrast, something specific to the particular stimuli and assessment techniques traditionally chosen for examination (i.e., CSs, USs, and conditioned responses). A functional analysis of associations suggests that associations can sometimes be useful in either direction. Consider, for example, the association between the two sides of a coin, or between Paris and the Eiffel tower, or between a chair and a table. Do these associations go only from one event to the other one? If so, which is E1 and which is E2? Or can these associations act in either direction?

Thus, our purpose was to develop a more basic, pure associative preparation, devoid of the CS-US specifications, in order to allow for a more general understanding of the directionality of the, presumably common, underlying associative process. For this purpose, the experiments that we report here used only neutral, initially unrelated stimuli (colors and figures). During training, human subjects were able to see each pair of events in the color→figure direction. At test, we asked subjects to report the correct color upon being presented with a figure. Under these conditions, we expected that the subjects would be able to use the acquired color→figure associations in the direction opposite to that of training, thereby showing evidence of having learned a bidirectional association.

EXPERIMENT 1

Method

Subjects

Eighteen undergraduate students from Deusto University volunteered for the study and were randomly assigned to Group Consistent ($n = 10$) and Group Inconsistent ($n = 8$). All subjects were run simultaneously. The distance between subjects was approximately 1.5 m, and adjacent subjects were assigned to different groups and counterbalanced stimuli.

Procedure

The experiment was presented to the subjects as a computerized card game. Within the game, the subjects had to bet on different cards. Each card had a color on one side and a figure on the other side. During training, the colored side of each card was always shown before the side containing the figure. Thus, each color represented an E1 and each figure represented an E2. The colors were blue, yellow, and red, which, counterbalanced, served as E1a, E1b, and E1c. The figures on the other side of the card were ASCII 01 (a smiling face), which served as E2a, ASCII 30 (similar to a house), which served as E2b, and ASCII 127 (similar to a pyramid), which served as E2c.

Training. There were 54 trials in this phase, 18 each of E1a, E1b, and E1c. In each trial, we presented on the screen one of three different colored cards (E1a, E1b, or E1c). The subjects had to bet on each trial, but they won the points that they bet only if they found E2a on the other side of the card. Otherwise, they lost the points that they had bet. The subjects made their bets by increasing or decreasing the height of a bar within a vertical scale provided on the right side of the screen. The scale was anchored at 0, 25, 50, 75, and 100 points. We used bets as our dependent variable (and, in Experiments 2 and 3, certainty ratings) rather than a more conventional measure such as, for example, percent of E2 correctly anticipated, because they provided greater sensitivity. Presumably, the more certain the subjects were about their responses, the greater the bet (or their certainty rating in Experiments 2 and 3). Once the subjects had selected the level of their bet, they had to press <ENTER>. Then they saw the reverse side of the card (i.e., E2a, E2b, or E2c). The number of points that they had won or lost in that trial was shown at the upper-left corner of the screen until the subjects pressed any key to start the next trial. Additionally, the cumulative score was shown at the upper-right corner of the screen. After each trial, a black screen showed "NEXT BET" in the middle of the screen. The subjects controlled the duration of each screen presentation by pressing a key in order to go on to the next presentation. A translation of the instructions for this phase reads as follows:

Welcome to Las Vegas casino. Here you can win as much money as you could ever imagine, but look out!, you can also lose everything unless you learn how the game works. In each play, you will see one colored card at a time, and you will win points if you correctly guess when the winning figure is on the other side of the card.

Here is the card that will allow you to win:

[A card showing the E2a figure was presented at this time]

You have to remember that the other side of the card is what will be visible when you bet. You will be able to bet on the card before you turn it over. If this particular figure is on the back of the card, you will win the number of points that you bet. If any other figure is on the back of the card, your bet will be subtracted from the total number of points you will have accumulated to that moment.

Table 1 summarizes the design for this experiment. For Group Consistent, each E1 was always paired with a particular E2 (e.g., E1a was always paired with E2a). For Group Inconsistent, stimuli E1a, E1b, and E1c were randomly paired with E2a, E2b, and E2c. Each of the E1s and E2s was presented on a total of 18 trials to both groups, but only the 18 trials with E2a presentation were reinforced with a gain of the points bet on that trial. Thus, Group Consistent received 18 E1a→E2a pairings, 18 E1b→E2b pairings, and 18 E1c→E2c pairings, in pseudorandom order; for Group Inconsistent, each of the 18 presentations of each E1 (e.g., E1a) was paired six times with E2a, six times with E2b, and six times with E2c, also in pseudorandom order.

Test. In the test phase, our purpose was to assess potential backward associations acquired during training. The difference with respect to the training phase was that the test phase had only one trial and in this trial we tested the association between E1a and E2a, in

Table 1
Design Summary of Experiment 1

Group	Phase 1	Test
Consistent	E1a→E2a E1b→E2b E1c→E2c	E2a
Inconsistent	E1a→E2a,b,c E1b→E2a,b,c E1c→E2a,b,c	E2a

Note—E1a, E1b, and E1c were three cards of different colors. E2a, E2b, and E2c were three different figures (ASCII codes) which represented the reverse side of the cards. E2a,b,c in Group Inconsistent means that each E1 was followed by E2a, by E2b, and by E2c the same number of times. During training, only E2a was positively reinforced. At test, we presented E2a and instructed subjects that they would win the points they bet if E1a was on the other side of the card.

the E2a→E1a direction. Specifically, in this phase, we presented stimulus E2a and asked subjects to guess whether E1a was on the other side. The instructions for this phase appeared immediately after the last training trial. A translation of these instructions reads as follows:

An eccentric millionaire has observed your ability with this game. He proposes ONE LAST BET in which you are going to be able to show how good you are. The proposal is the following: THE DECK IS GOING TO REMAIN THE SAME, but the cards will be presented up-side-down. You will have to guess the color that is at the other side of a figure. All you have to do is to remember everything you have learned to this point, and use it in this last bet, BECAUSE THE DECK IS STILL THE SAME.

This is the card that will allow you to win the bet:

[A card showing the E1a color was presented here]

If you think that this card is on the other side of the card that you will now see, bet on it and you will win.

Then the E2a figure was presented and subjects made their bets using the same scale as they had used during training.

Results and Discussion

The main finding of this study was that the bets of Group Consistent during the test trial were higher than

those of Group Inconsistent. This suggests that subjects who had learned the E1a→E2a association during training were able to use this association in the backward direction at test. These results are depicted in Figure 1. A Mann–Whitney U test confirmed that Group Consistent bet more on an expectancy of E1a given an E2a presentation at test than did Group Inconsistent ($U = 1.0, p < .001$). Moreover, as can be seen in Figure 1, the backward responding to E2a→E1a shown by Group Consistent in the test trial was as strong as the forward responding to E1a→E2a that this group showed in the last training trial (Wilcoxon $T = 4.5, p > .1$); however, a ceiling effect may have contributed to this near equality.

Because E2a had been consistently reinforced, an E2a–reinforcement association could potentially provide a simpler explanation than the backward-association account for the strong responding of Group Consistent to E2a at test. Note, however, that E2a was also consistently reinforced in Group Inconsistent and that responding to E2a at test was weaker in this group. Nevertheless, the amount of reinforcement occurring each time that E2a was presented was smaller in Group Inconsistent because subjects in this group bet fewer points during training (see Figure 1). Thus, an E2a–reinforcement association cannot categorically be rejected as an explanation of the differential outcome observed in these two groups. Nevertheless, the results suggest that subjects in Group Consistent were able to use the acquired association in the direction opposite to that of training.

EXPERIMENT 2

In addition to our interpretation of the results of Experiment 1 as supporting the existence of bidirectional associations or reflecting differential reinforcement during training, there is also an alternative explanation in terms of forward associations. This interpretation hinges on the instructions of the test phase, which told the sub-

Figure 1. The left side of the figure shows the expectation of E2a during the presentation of E1a, E1b, and E1c during training for Group Consistent and Group Inconsistent in Experiment 1. The right side of the figure shows the expectation of E1a during the presentation of E2a at test for both groups.

jects that they should look for E1a. Our mentioning of E1a in the test instructions may have caused the activation of the E1a→E2a association in Group Consistent but not (as much) in Group Inconsistent, in which case it would not be surprising that Consistent subjects bet more on E2a at test than did Inconsistent subjects. Consequently, the results of Experiment 1 are open to an alternative interpretation in terms of a more traditional forward association.

Therefore, in Experiment 2, we modified the test instructions and procedure: To avoid telling subjects that they had to look for E1a, we simply presented E2a and asked them, without any additional guide, to tell us which color was on the other side. This, we hoped, would provide an unambiguous backward response. If subjects knew, without any instruction, which color was on the other side, they would be demonstrating that they could use the association in the backward direction, thereby showing the bidirectionality of the acquired associations. Moreover, in Experiment 2, we completely eliminated reinforcement and bets throughout all phases and instructions. This change eliminated the difference between groups in reinforcement during training. The training phase consisted of an observational phase in which the subjects were simply told to observe the colored side of each card and to use the computer to turn it over in order to observe the figure on the reverse side. This procedure avoided the potential problem of different colors' or figures' being differentially associated with reinforcement.

Method

Subjects

The subjects were 9 undergraduate students from Deusto University who volunteered for the study. The subjects were run individually using a personal computer. The experimenter was in the room with the subjects while the instructions were presented. After that, the experimenter stayed in an adjacent room.

Procedure

This experiment used a within-subjects design. During training, E1a color was always paired with the E2a figure. Colors E1b and E1c were randomly assigned to the E2b and E2c figures. Thus, the subjects received the following pairings of stimuli, in pseudorandom order: eight E1a→E2a, four E1b→E2b, four E1b→E2c, four E1c→E2b, and four E1c→E2c. This design is summarized in Table 2.

Table 2
Design Summary of Experiment 2

Training	Tests
E1a→E2a	E2a/E2b/E2c
E1b→E2b,c	
E1c→E2b,c	

Note—In Experiment 2, a within-subjects design was used. E1a, E1b, and E1c were three cards of different colors. E2a, E2b, and E2c were the three figures shown in Figure 2, which represented the reverse side of each card. E2b,c means that E2b and E2c were presented the same number of times in that condition. No bets or reinforcement were used in this experiment. At test, each of the three E2s was presented in counterbalanced order, and the subjects were asked to indicate which color was on the other side of each of them, as well as to indicate their certainty rating for that response.

The E1 side of the cards was represented by one of three different colors, red, blue, and white, and the E2 side of the cards was represented by one of the three different figures shown in Figure 2. Because E1b played the same role as E1c and E2b played the same role as E2c, these stimuli were not fully counterbalanced. Instead, E1b and E1c as well as E2b and E2c were partially collapsed, and our main effort was to counterbalance them against E1a and E2a. That is, each of the three possible colors of E1a was paired with each of the three possible figures in Figure 2 (e.g., when E1a was red, E2a could be Figure 2A, Figure 2B, or Figure 2C, with the other two stimuli, b and c, being designated E2b and E2c). The reason for using the figures in Figure 2 instead of the ASCII codes used in Experiment 1 was to make the task more attractive to potential volunteers and to provide for an easier discrimination between the three E2s, given that no overt reinforcement was used in this experiment.

Training. This phase was similar to the training phase in Experiment 1, but no reinforcement was used and there were only 24 trials (8 for each stimulus). The subjects did not have to bet on any card and were instructed to simply observe the cards presented on the screen. When a colored card was presented, the subjects had to press any key in order to see the other side of the card on the screen. At this point, if a subject pressed any key, the following sentence was presented on a white screen: *Press any key in order to see the next card.*

The instructions for this phase were as follows:

This experiment consists of a card game, but before you begin the game, I am going to give you a little assistance: I am going to show you both sides of each of the cards, one by one.

All you have to do is to pay attention and make use of this information in order to learn to distinguish between the cards.

Test. To avoid alternative forward interpretations of this experiment, we did not present any stimuli that could activate the E1a→E2a association in a forward direction within the test instructions or during the test phase. Before the test phase, the following sentence was presented alone:

From now on, please, type THE COLOR on the other side of each card as it is presented. Press 'C' to continue.

[After that, Figure 2A was displayed and the following questions were presented.]

What color is on the other side of this card?

How certain are you about it? (Please, type a number between 0 [no idea], and 9 [absolutely sure].)

The same procedure was repeated for Figure 2B and then for Figure 2C. Because the figures in Figure 2 had been counterbalanced in their serving as E2a, E2b, and E2c, presenting them in that order counterbalanced the order of the testing of E2a, E2b, and E2c.

Results and Discussion

The main finding of this experiment was that all of the subjects demonstrated that they had learned an unambiguous bidirectional association between E1a and E2a but not between the other pairs of events (i.e., not between E1b and E2b or E2c or between E1c and E2b or E2c). All of the subjects in this experiment typed the color for E1a when they were asked which color was on the other side of E2a, and all of them gave this response with the maximum degree of certainty allowed by our rating scale. In contrast, they responded inconsistently, either E1b or E1c, and with low certainty when asked which color was on the other side of E2b and E2c. The certainty scores are shown in Figure 3. Because E2b and

Figure 2. Panels A, B, and C, counterbalanced, were used as E2a, E2b, and E2c in Experiment 2. Panels B and C, counterbalanced, were used as E2a and E2b in Experiment 3.

E2c played the same role as control stimuli in this experiment and because their certainty ratings were identical ($M = 4.0$, $SE = 1.58$ in both cases; see Figure 3), their certainty ratings were collapsed for purposes of analysis. The certainty with which subjects responded to the test on E2a was significantly higher than the mean certainty for E2b and E2c (Wilcoxon $T = 0.00$, $p < .01$).

EXPERIMENT 3

Experiment 2 provided additional support for the bidirectional hypothesis and discounted explanations of Experiment 1 that depended on differential reinforcement or unintentional memory activation by the test instructions. However, another possible explanation purely in terms of forward association could explain the data obtained in Experiments 1 and 2. Holland (1981) has sug-

gested a representation-mediated view of associative learning, according to which mental representations of stimuli that are not physically present may become associated with actual events. In this framework, once the first training trial of $E1 \rightarrow E2$ has occurred in our experiments, a mental representation of E2 could be activated whenever E1 was presented. Thus, this mental representation of E2 could become associated with E1 in the forward $E2 \rightarrow E1$ direction, because E1 was retained on the screen for some time after the mental representation of E2 had been activated by the presentation of E1 in Trial 2 and in subsequent training trials. In this framework, a forward $E1a \rightarrow E2a$ association as well as a forward $E2a \rightarrow E1a$ association could be formed starting on the second trial of training in the consistent conditions but less so in the inconsistent ones. The only trial in which a forward $E2a \rightarrow E1a$ association cannot be formed because the representation of E2a cannot possibly be activated by E1a is the first trial. For this reason, we used only one training trial in Experiment 3.

Another potential problem of Experiments 1 and 2 was that the experiments had been superimposed on a card game in which stimulus E1 was the colored side of a card and E2 was the figure on the other side of the card. This cover story allows for a configural explanation of our data. In this configural view, E1 and E2 can be regarded as two parts of the same stimulus: the card. If so, it would be arguable as to whether there was any meaningful directionality from one event to the other, because subjects could have considered E1 and E2 as different aspects of the same stimulus. In Experiment 3, the plausibility of this explanation was reduced by changing the cover story used to instruct the subjects about the experiment. In the present experiment, cards were substituted for folders that contained different pieces of papers. The folders were E1s, and the papers were E2s.

Method

Subjects

The subjects were 12 undergraduate students from Deusto University who volunteered for the study. The subjects were run indi-

Figure 3. Mean certainty for E2a, E2b, and E2c in the test phase of Experiment 2, which used a within-subjects design and observational training. Error bars indicate standard errors of the means. The lack of an error bar for E2a is due to $SE = 0$ for that stimulus.

Table 3
Design Summary of Experiment 3

Training	Tests
E1a→E2a	E2a/E2b
E1b→E2b	
E1c→E2b	

Note—In Experiment 3, a within-subjects design was used. E1a, E1b, and E1c were three folders of different colors. E2a and E2b were two different pieces of paper (shown in Panels B and C of Figure 2) which could be found inside the folders. Only one training trial for each pair of events (in counterbalanced order) was given in this experiment. No bets or reinforcement were used in this experiment. At test, E2a and E2b were presented in counterbalanced order and subjects were asked to indicate which folder they belonged to, as well as to indicate their certainty rating for that response.

vidually using a personal computer. The experimenter was in the room with the subjects while the instructions were presented. After that, the experimenter stayed in an adjacent room.

Procedure

This experiment used a within-subjects design that can be seen in Table 3. The E1 folder was one of three different colors, red, blue, and white. Because E1b and E1c played the same role, the color for E1b was kept constant (red) whereas the color for E1a and E1c was counterbalanced (blue and white). The papers contained in the folders are shown in Figures 2B and 2C. These papers, counterbalanced, were used as E2a and E2b.

Training. This phase was identical to Phase 1 in Experiment 2, but only one trial for each association was presented. When a colored folder was presented, the subjects had to press any key in order to see the paper contained in the folder. After that, the paper was presented on the screen. At this point, if a subject pressed any key, the following sentence was presented on a white screen: *Press any key in order to see the next folder.*

During training, the E1a color was paired with the E2a paper, whereas the E2b paper was assigned to both E1b and E1c. Thus, subjects received one trial of E1a→E2a, one trial of E1b→E2b, and one trial of E1c→E2b, with trial order counterbalanced between subjects (see the Appendix).

The instructions for this phase were as follows:

In this experiment you will see three folders: a blue, a red, and a white one. The only thing you have to do is to open the folders and to pay attention to the pictures contained in each of them.

Test. The order of testing for stimuli E2a and E2b was counterbalanced as shown in the Appendix. When each of these stimuli was presented, the following question was shown on the screen: *What is the color of the folder that contains this picture?* After the subjects answered this question, the following question was formulated: *How certain are you about it?* The subjects were required to answer this question using a response bar similar to that used in Experiment 1 to indicate their certainty on a 0–100 scale.

Results and Discussion

In a total of 12 subjects, 10 of them indicated the E1a color correctly when they were asked which folder contained E2a. Not surprisingly, these subjects responded inconsistently when they were asked which folder contained E2b (4 of them indicated E1b, and 6 of them indicated E1c). The certainty with which those 10 subjects responded to E2a ($M = 99$, $SE = 1$) was higher than that with which they responded to E2b ($M = 70.60$, $SE = 10.61$; Wilcoxon $T = 0.00$, $p < .05$). This gives further support to the hypothesis that the acquired association

was bidirectional and that mediated forward conditioning cannot account for these results.

GENERAL DISCUSSION

The three experiments taken together provide evidence that subjects can acquire bidirectional associations and can use the associative information in either the forward or the backward direction. Although the idea of bidirectional associations is an old one (e.g., Beritov, 1924), most contemporary theories of learning have assumed that associations are unidirectional. The reason for this is probably that the studies in the conditioning literature have not provided unambiguous results on bidirectionality (see, e.g., Spetch et al., 1981, for a review), and that studies in the human paired-associate literature were generally involved with language-related issues (see, e.g., Ekstrand, 1966, for a review), which may have prompted researchers interested in basic associations to ignore those findings. The present results, however, suggest that the issue of directionality needs to be reconsidered.

There are also some experiments in the current associative learning literature which could be interpreted as supporting the idea of bidirectional associations, but most of them are open to alternative explanations in terms of forward associations. For example, Holland (1981) and Ward-Robinson and Hall (1996) paired two events in Phase 1 (E1→E2) and then paired E1 to a US in Phase 2 (E1→US). At test, they presented E2 and observed a conditioned response. This could be interpreted as supporting the existence of a bidirectional association between E2 and E1. However, Holland and Ward-Robinson and Hall provided an explanation in terms of forward associations. In Phase 2, E1 activated a mental representation of E2. This mental representation of E2 acquired a forward association to the US. Therefore, as the authors noted, the response to E2 observed during testing in those experiments may reflect the forward (representation-mediated) E2→US association rather than the backward association between E2 and E1.

Matzel et al. (1988) provided some clearer experimental support for the bidirectional hypothesis. They paired two events (E1→E2) in Phase 1 and then paired a shock to E2 in Phase 2 (i.e., shock→E2). Although presentations of E2 did not yield a (backward) conditioned response, the presentation of E1 in a subsequent test phase did produce a conditioned response. In this experiment, E1 had never been paired with shock and E2 was associated with shock only in the backward direction (i.e., shock→E2). Thus, Matzel et al. interpreted the conditioned response that they observed at test in terms of bidirectional associations. That is, according to Matzel et al., E1 produced a conditioned response because the presentation of E1 at test activated a forward representation of E2, which in turn activated a backward representation of the shock. This, of course, was an indirect demonstration of bidirectionality, but, as noted by Matzel et al., bidirectional associations are ordinarily not directly expressed when conditioned responses are used as

the index of learning (i.e., conditioned responses usually occur only when a CS predicts the US).

The purpose of the present research, therefore, was to provide a more direct demonstration of bidirectional associations and to show that bidirectionality could be directly expressed, assuming that no directionality was implied in the content of the events or type of response that the experimenter chose to investigate. To study the directionality of associations in a more general framework, we used neutral stimuli, colors and figures, as the associates. Each color was associated with one figure. Colors and figures were arbitrarily defined as E1s and E2s, respectively. Our central hypothesis was that the associations that subjects could learn between those two stimuli could be readily tested in either direction, because, unlike the inherent functional directionality that exists between CSs and USs, no inherent directionality was implied by the present color-figure pairings. That is, subjects who learned which color (E1) was followed by which figure (E2) should have been able to tell us what figure was preceded by which color if they had learned a bidirectional association. The present experiments showed that this was the case. Moreover, Experiments 2 and 3 discarded the possibilities that differential reinforcement or forward, representation-mediated, associations were responsible for the observed results. In Experiment 3, only one training trial was used, which makes representation-mediated conditioning an implausible explanation for the results. We know of no other alternative explanation for our results except that a bidirectional association, or perhaps two unidirectional associations, one forward and one backward, were formed when a forward pairing was presented. Moreover, there was no difference between the certainty with which subjects responded to the last forward trial of training and the backward test trial (see Figure 1). As suggested by Asch and Ebenholtz (1962), the backward association seems to be equal in strength to the forward association.

In looking for further alternative interpretations in terms of forward associations, it could be argued, for example, that even with a single trial (Experiment 3), during the presentation of E2a, the context might have activated a representation of E1a, and hence a representation-mediated forward association might have been acquired between E2a and E1a. By the same reasoning, however, the context should have activated not only a representation of E1a when E2a was presented, but also a representation of each of the E1s and E2s when each of the E2s was presented. Thus, this view does not provide a compelling explanation for the present data.

An alternative possibility could be that, even in a single trial (Experiment 3), presentation of E1a might have activated a representation of E1a that remained active until and into the presentation of E2a. Hence, the presentation of E2a could in some sense precede an active representation of E1a, and, hence, a forward association could be formed from E2a to E1a. By the same argument, however, the representation of E2a should also stay active until pre-

sensation of the following E1 (whether E1b or E1c, as a function of the counterbalancing sequence). Thus, in this framework, it is not clear why E2a should acquire a forward association with the representation of the stimulus that had preceded it (i.e., E1a) rather than with the stimulus that followed it (E1b or E1c). A bidirectional association between E1a and E2a seems to provide a far more parsimonious interpretation of the present findings.

In summary, our results suggest that the generally assumed predictive directionality of the associations is probably a consequence of the dependent variables and experimental stimuli that are typically used in associative research rather than being an inherent quality of the presumably underlying associative structure. Indeed, our data suggest that, when a dependent variable different from those which are commonly monitored is used to assess the directionality of the associations, responses indicative of bidirectional associations can be readily observed.

REFERENCES

- ASCH, S., & EBENHOLTZ, S. M. (1962). The principle of associative symmetry. *Proceedings of the American Philosophical Society*, **106**, 135-163.
- BERITOV, I. S. (1924). On the fundamental nervous processes in the cortex of the cerebral hemispheres. *Brain*, **47**, 109-148.
- EKSTRAND, B. R. (1966). Backward associations. *Psychological Bulletin*, **65**, 50-64.
- HOLLAND, P. C. (1981). Acquisition of representation-mediated conditioned food aversion. *Learning & Motivation*, **12**, 1-12.
- HOROWITZ, L. M., BROWN, Z. M., & WEISSBLUTH, S. (1964). Availability and the direction of associations. *Journal of Experimental Psychology*, **68**, 541-549.
- KAMIN, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the prediction of behavior: Aversive stimulation* (pp. 9-31). Miami: University of Miami Press.
- MACKINTOSH, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, **82**, 276-298.
- MATZEL, L. D., HELD, F. P., & MILLER, R. R. (1988). Information and expression of simultaneous and backward associations: Implications for contiguity theory. *Learning & Motivation*, **19**, 317-344.
- MILLER, R. R., & BARNET, R. C. (1993). The role of time in elementary associations. *Current Directions in Psychological Science*, **2**, 106-111.
- MURDOCK, B. B. (1962). Direction of recall in short-term memory. *Journal of Verbal Learning & Verbal Behavior*, **1**, 119-124.
- PEARCE, J. M., & HALL, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, **87**, 532-552.
- RESCORLA, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative & Physiological Psychology*, **66**, 1-5.
- RESCORLA, R. A. (1980). Simultaneous and successive associations in sensory preconditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, **6**, 207-216.
- RESCORLA, R. A., & WAGNER, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton.
- SPETCH, M. L., WILKIE, D. M., & PINEL, J. P. J. (1981). Backward conditioning: A reevaluation of the empirical evidence. *Psychological Bulletin*, **89**, 163-175.
- WARD-ROBINSON, J., & HALL, G. (1996). Backward sensory preconditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, **22**, 395-404.

APPENDIX
Stimulus Assignment As Well As
Training and Testing Order in Experiment 3

Subject	E1a	E1b	E1c	Trial Sequence	Test 1	Test 2	E2a	E2b
1	Blue	Red	White	E1a→E2a E1c→E2b E1b→E2b	E2a	E2b	Figure 2C	Figure 2B
2	Blue	Red	White	E1b→E2b E1a→E2a E1c→E2b	E2a	E2b	Figure 2C	Figure 2B
3	Blue	Red	White	E1c→E2b E1b→E2b E1a→E2a	E2a	E2b	Figure 2C	Figure 2B
4	White	Red	Blue	E1a→E2a E1c→E2b E1b→E2b	E2b	E2a	Figure 2B	Figure 2C
5	White	Red	Blue	E1b→E2b E1a→E2a E1c→E2b	E2b	E2a	Figure 2B	Figure 2C
6	White	Red	Blue	E1c→E2b E1b→E2b E1a→E2a	E2b	E2a	Figure 2B	Figure 2C
7	Blue	Red	White	E1a→E2a E1c→E2b E1b→E2b	E2a	E2b	Figure 2B	Figure 2C
8	Blue	Red	White	E1b→E2b E1a→E2a E1c→E2b	E2a	E2b	Figure 2B	Figure 2C
9	Blue	Red	White	E1c→E2b E1b→E2b E1a→E2a	E2a	E2b	Figure 2B	Figure 2C
10	White	Red	Blue	E1a→E2a E1c→E2b E1b→E2b	E2b	E2a	Figure 2C	Figure 2B
11	White	Red	Blue	E1b→E2b E1a→E2a E1c→E2b	E2b	E2a	Figure 2C	Figure 2B
12	White	Red	Blue	E1c→E2b E1b→E2b E1a→E2a	E2b	E2a	Figure 2C	Figure 2B

(Manuscript received March 26, 1998;
revision accepted for publication September 6, 1998.)