Relational learning with and without awareness: Transitive inference using nonverbal stimuli in humans

ANTHONY J. GREENE and BARBARA A. SPELLMAN
University of Virginia, Charlottesville, Virginia

JEFFERY A. DUSEK and HOWARD B. EICHENBAUM
Boston University, Boston, Massachusetts

and

WILLIAM B LEVY
University of Virginia, Charlottesville, Virginia

Learning complex relationships among items and representing them flexibly have been shown to be highly similar in function and structure to conscious forms of learning. However, it is unclear whether conscious learning is essential for the exhibition of flexibility in learning. Successful performance on the transitive inference task requires representational flexibility. Participants learned four overlapping premise pairs (A > B, B > C, C > D, D > E) that could be encoded separately or as a sequential hierarchy (A > B > C > D > E). Some participants (*informed*) were told prior to training that the task required an inference made from premise pairs. Other participants (*uninformed*) were told simply that they were to learn a series of pairs by trial and error. Testing consisted of unreinforced trials that included the non-adjacent pair, B versus D, to assess capacity for transitive inference. Not surprisingly, those in the informed condition outperformed those in the uninformed condition. After completion of training and testing, uninformed participants were given a postexperimental questionnaire to assess awareness of the task structure. In contrast with expectations, successful performance on the transitive inference task for uninformed participants does not depend on or correlate with postexperimental awareness. The present results suggest that relational learning tasks do not necessarily require conscious processes.

Declarative or explicit memory refers to those facts or events of which we are consciously aware and includes semantic knowledge and episodic (autobiographical) memory. Other aspects of past experience may influence behavior without conscious recollection (e.g., Jacoby & Dallas, 1981). These nondeclarative or implicit forms of learning and memory are generally characterized as acquired fluency of perceptual and motor performance (e.g., Roediger, 1990). Nondeclarative memory reveals itself in habits, procedural knowledge, perceptual priming, simple sequence learning, single-discriminant forms of classical and operant conditioning, and path-based route learning (for reviews, see Eichenbaum, 1999b; Moscovitch, 2000). One of the most convincing demonstrations that nondeclarative learning and memory are independent of declarative memories is that declarative memory may be substantially impaired without corresponding impairments of nondeclarative memory. That is, whereas damage to part or all of the *hippocampal system* (hippocampus, amygdala, parahippocampal cortices) results in anterograde amnesia, which is an inability to acquire new declarative memories, performance on nondeclarative tasks is normal or near normal (for reviews, see Milner, 1972; Squire & Knowlton, 1999).[1]

Flexible Representation

Although opinion is currently changing (for a review, see Eichenbaum, 1999a; Willingham, 1997), one putative feature of the declarative versus nondeclarative distinction is the extent to which each supports generalizable and flexible representations (e.g., Dienes & Berry, 1997). Nondeclarative forms of learning and memory have been widely characterized as relatively inflexible, automatic, and limited in generalizability (e.g., Holyoak & Spellman, 1993). Evidence supporting such views is drawn from studies showing that perceptual priming is highly specific to certain perceived surface features. For instance, study–test changes in such things as speaker's voice for auditory words (Schacter, Church, & Bolton, 1995) or object orientation for vision and touch (Srini-

vas, Greene, & Easton, 1997) diminish the magnitude of priming (for general reviews, see Richardson-Klavehn & Bjork, 1988; Roediger & Srinivas, 1993). In contrast, declarative memory can be viewed as the way in which one models the complexities of the world (Clark & Squire, 1998). Declarative memory is highly organized and represents broad associative relationships, both probabilistic and conditional, among multiple contingencies(Cohen & Eichenbaum, 1993; Hebb, 1949; Kroll, Knight, Metcalfe, Wolf, & Tulving, 1996). For example, word list recall is vastly improved if it is associatively organized (Anderson, 1981; Bower, Clark, Lesgold, & Winzenz, 1969), and well-learned arrays of objects are learned in associative hierarchies (McNamara, Hardy, & Hirtle, 1989). Similarly, association to existing knowledge, or deep encoding (Bradshaw & Anderson, 1982; Craik & Tulving, 1975), broadens the circumstances under which the information may be recalled. Such effects suggest that the more information is associated, the more accessible the information may be, even under varied recall conditions. "The full interconnectedness of such a representational system produces the ability of information to be activated regardless of the current context . . . and [such representations] can be manipulated and flexibly expressed in any number of novel situations" (Cohen & Eichenbaum, 1993, p. 62).

The notion that declarative memory may be characterized by flexible, multicontingent representation is supported by convergent neuropsychological findings. A wide array of evidence (for reviews, see Cohen & Eichenbaum, 1993; Sutherland & Rudy, 1989) shows that both conscious forms of learning and memory in humans and relational learning in humans and animals are critically dependent on the same set of structures—namely, the hippocampal system (e.g., Rickard & Grafman, 1998; Squire, Zola-Morgan, & Chen, 1988). Relational learning tasks require the concurrent representation of relationships among multiple stimuli or events. One type of relational learning involves conjunctions of stimuli that correspond to different outcomes than do any individual stimuli. For example, a rodent's response to a light and a tone presented simultaneously may result in reward, whereas a response to a light or a tone separately may result in punishment. These relational tasks include conditional discriminations and multicontingent forms of classical and operant conditioning (e.g., transverse patterning, negative patterning). For such tasks, a set of stimuli is predictive of reward only in the context of additional stimuli. Another type of relational learning requires a reorganization of learned items into a more global representation. Path integration in animals (cf. survey-based spatial navigation in humans) and transitive inference (TI) in humans and animals are two examples of this type of task. Path integration requires that individually learned routes among points form a representation that allows novel routes among those points. Such behavioral flexibility suggests that maplike representations are formed from the integration of partial experiences (e.g.,

Eichenbaum, Stuart, & Morris, 1990; O'Keefe & Nadel, 1978; Pearce, Roberts, & Good, 1998). TI requires a reorganization of individually learned discriminations in order to make an inferential judgment about a novel discrimination (Bryant & Trabasso, 1971; Piaget, 1928; Siemann & Delius, 1993, 1996).

The design of the TI task makes it well suited for exploring representational flexibility. In the prototypic TI experiment, participants are trained on a series of overlapping two-item discriminations or premise pairs (A > B, B > C, C > D, D > E, where each letter stands for a stimulus item and ">" represents the relationship "should be selected over"). For animals, the choice is between stimuli varying on some physical dimension (e.g., odor), whereas for human versions of the task, the stimuli are generally verbal (e.g., "Al is taller than Bob"). Each of the discriminations could be represented as an individual pair (AB, BC, CD, DE), or all five items could be represented as a superordinate hierarchy (A > B > C > D > E). Hierarchical organization allows correct inferences to be made about all possible pairs of items (in this case, 10) from only the subset of those possible pairs consisting of adjacent items (in this case, 4). To examine which of the representations is used, participants are tested with a pair of nonstudied, nonadjacent items, B versus D (nonadjacent pairs containing end items A or E can be solved nontransitively, because A is always correct and E is always incorrect). The capacity for inferential judgment (consistently choosing B over D) in this test is interpreted as prima facie evidence of a representation with a linearly ordered hierarchy.

The demands of a TI task are remarkably similar to the demands of many declarative memory tasks. Cohen and Eichenbaum (1993) proposed that declarative memory consists of reorganized samples of experience that make up a flexible, global representation. Such representations support the sort of general knowledge useful for inference and decisions about novel or unexpected events. Thus, it is proposed, for both TI and declarative memory, that prior samples of experience, which are reorganized for coherence and context, allow correct choices under novel circumstances.

Conscious Awareness and Relational Learning

Because relational learning and conscious forms of memory both depend on the hippocampal system, it is not surprising that some evidence (for a review, see Schacter, 1998) suggests that conscious awareness of the learned contingencies is both necessary and sufficient for relational learning. Indeed until recently, the prevailing viewpoint has been that hippocampal-system-dependent learning requires conscious awareness and that non–hippocampal-system-dependent learning does not (e.g., Clark & Squire, 1998, 1999). The suggestion that relational learning entails task awareness heightens interest in comparing declarative memory in humans with relational learning in humans and animals. Although evidence shows that good performance on numerous rela-

tional learning tasks corresponds to task awareness in humans, the opportunities for disconfirming evidence have been remarkably sparse. Nevertheless, recent human studies reveal that some forms of hippocampal-system-dependent learning may occur without awareness. Chun and Phelps (1999) showed, when searching for targets embedded in various backgrounds, that the context facilitation of the background–target configuration is disrupted in hippocampal system amnesia. Although context facilitation is observed in normal control participants, their recognition of studied backgrounds is not greater than chance. Therefore, this type of context facilitation requires the hippocampal system but does not depend on awareness of the context–target relationship. In addition, Rajaram and Coslett (2000a, 2000b) showed that priming for novel word associations, across semantically related but perceptually different sentences, is disrupted in hippocampal system amnesia but is not disrupted by basal forebrain amnesia of comparable density. Novel word associations must therefore depend on the hippocampal system, but not on conscious memory processes. Schacter and his colleagues (Schacter et al., 1995) showed that voice-specific (Type B) priming (i.e., a same-voice advantage for auditory word identification) is disrupted in hippocampal system amnesia even though voice-specific priming is not attributable to conscious memory strategies. Thus, evidence is beginning to emerge that tasks that depend on the hippocampal system need not entail conscious (declarative) processes (for a discussion, see Eichenbaum, 1999a).

These three findings tend to confuse what was previously thought to be a clear distinction between implicit and explicit tasks. Prior to such findings, a consensus was beginning to emerge that a simple dichotomy would suffice—namely, that relational learning tasks were dependent on the hippocampal system and required awareness, whereas nonrelational learning tasks were independent of the hippocampal system and did not imply declarative task awareness. Background-specific target searching, novel word association priming, and voice-specific priming challenge conventional notions of a one-to-one relationship between awareness and hippocampal system dependence, because they are all hippocampal-dependent, implicit tasks. Theories of relational learning, hippocampal-system-dependent learning, and conscious learning, would all have difficulty predicting which tasks should require (or at least correlate with) awareness. At present, it is not known what types of tasks necessarily entail awareness and which do not. One way to approach this problem is to evaluate the role of awareness in a prototypical relational task.

The Present Experiments

The multicontingent, higher order associations used to support such things as inferential judgments and other tasks typically construed to be relational are currently unexplored with respect to awareness. Here, we ask what the role of task awareness is in the more prototypical relational task, where novel relationships are inferred from broad, organized representations.

To assess this question, we examine the relationship between task awareness and performance on a nonverbal TI task. A nonverbal version of the TI task is appropriate for exploring the relationship between conscious awareness and relational learning, for several reasons. First, consistent with the theory of relational learning, Dusek and Eichenbaum (1997) found that both normal and hippocampally disabled rats could rapidly learn premise pairs but that only normal rats could successfully make an inferential judgment. Second, use of a nonverbal TI task allows a direct comparison of human and animal results (see also Siemann, Delius, & Wright, 1996; Weaver, Steirn, & Zentall, 1997).

Most important, the nonverbal TI task is well suited to assessing awareness of the superordinate hierarchy. For any TI task (verbal and nonverbal), the transitive relationship (BD) is never included in training, and the hierarchy that supports this inference is not required to learn the premise pairs. In verbal TI tasks, however, the relationships expressed in the premise pairs (e.g., longer than, better than, etc.) provide strong clues that a hierarchy can be formed among all the stimuli (Kallio, 1982; Nguyen & Revlin, 1993). Our nonverbal TI task provides no such inherent clues and thus does not bias participants toward forming overt hierarchical orderings. In addition, many TI tasks inform participants, prior to study, that the stimuli can be hierarchically organized. For this reason, we have included a comparison condition wherein participants are informed before the study session that the premise pairs do form a linear hierarchy, and an example of such a hierarchy is given (see Appendix A). Thus, for our experiments, participants who are not informed may or may not attain some degree of serendipitous awareness during the experiment, whereas those who are informed before study are aware of the hierarchy from the start

EXPERIMENT 1

In this experiment, we examined whether awareness of an overall hierarchical relationship among stimuli is necessary for inferential judgments. Because most human TI experiments begin with an explicit disclosure that there is a superordinate hierarchy to be discovered, we have included a group of participants who are informed at the outset that there is a hierarchy to be discovered. Awareness of the overall hierarchical relationships can, thus, occur in one of two ways: Either the participant may be told about it by the experimenter before training begins, or the participant may serendipitously develop awareness of it during the course of the experiment. Two questions follow directly. First, does being informed, prior to training, of the hierarchical arrangement of stimuli affect task acquisition or the capacity for TIs? Second, for those who are not informed prior to training, does the capacity for TI depend on discovery of the hierarchical organization? Two kinds of data addressed the latter

question. One is whether failure to become aware during training corresponds with chance performance on the TI task; another is whether there is a correlation between level of acquired awareness and transitive performance.

Method

Participants

The participants were 41 undergraduate and graduate students from the University of Virginia, who were unfamiliar with the Japanese Hiragana script. Graduate students participated voluntarily; undergraduates participated for course credit during the academic year or for $7.50 during the summer. The participants were randomly assigned to either an informed ($n = 19$) or an uninformed ($n = 22$) condition.

Materials

Stimulus items were five characters selected from the Japanese Hirigana script. Characters were initially selected as those judged to be highly discriminable; of those, pilot testing revealed that the five characters shown in Figure 1 were the most easily discriminable. The Hiragana characters were presented on a 15-in. color monitor in 36-point font size.

Procedure

Prior to training, the participants assigned to the informed condition were told that the task involved making a TI from premise pairs. In the uninformed condition, the participants were simply told that they were to learn the premise pairs by trial and error. See Appendix A for the full instructions.

Training. Pairs of characters were presented on the monitor. The participants were instructed to select one character from each by pressing either the "z" (left character) or the "m" (right character) on the keyboard. For each training pair, one character was always correct and the other always incorrect, and screen position (left–right) was randomized. The characters were presented for 3 sec or until a response was entered, whichever came first. Feedback consisted of the display of either "correct" in blue letters on the upper half of the screen or "incorrect" in red letters on the bottom half of the screen. The participants were trained to learn which of each pair was correct from a set of four overlapping premise pairs (A > B, B > C, C > D, D > E, where ">" indicates the correct choice is on the left) that could be encoded as distinct pairs or as a sequential hierarchy (A > B > C > D > E).

Training was divided into four blocks, each block consisting of 40 training trials. In Block 1, pairs were trained in sequence (de Boysson-Bardies & O'Regan, 1973), each pair was presented 5 times in a row, and the list was repeated once (i.e., 5 × AB, 5 × BC, 5 × CD, 5 × DE, 5 × AB, 5 × BC, 5 × CD, 5 × DE). In Block 2, pairs were trained in sequence, each pair was presented twice in a row, and the list was repeated an additional 4 times. In Block 3, pairs were presented in sequence, one presentation each, and the list was repeated an additional 9 times. In Block 4, pairs were presented in random order, 10 times each. At the end of each training block, performance on each pair was evaluated; if a participant failed to achieve at least 80% accuracy on each of the premise pairs, the training block was repeated.

Figure 1. The five Hiragana characters shown in hierarchical order. Adjacent pairs are presented at study, and the correct choice for each pair is to the left of the ">" sign.

Testing. Following completion of all four training blocks, the participants were tested, without feedback (i.e., without "correct" or "incorrect"). Test items included all four premise pairs, the transitive pair, BD, to assess capacity for TI, and the end-anchor pair, AE (e.g., Menzel, 1969) as a control. Test pairs were presented eight times each, randomly interspersed within a random presentation of the premise pairs.

Measuring awareness. After testing, all the uninformed participants were given as much time as they needed to complete a post-experimental questionnaire (shown in Appendix B). The questionnaire was designed to assess the participants' awareness of the hierarchical/logical nature of the BD task without confounding BD performance. Four raters independently evaluated postexperimental questionnaires: Awareness scores ranged from 1 to 5—1 corresponding to *no evidence of awareness of the hierarchical structure*, and 5 corresponding to *definite indications of awareness*.

Design

Informed versus uninformed was the only independent variable. Dependent measures were (1) proportion of premise pairs correct from each training block, (2) proportion of premise pairs correct at final test, (3) proportion of transitive pairs (BD) correct at final test, (4) proportion of end pairs (AE) correct at final test, and (5) for those in the uninformed condition, a posttest assessment of hierarchical awareness.

Results and Discussion

Task Awareness

The awareness ratings from the four raters resulted in an average score of 3.37, with a standard deviation of 1.02, a median of 3.50, and a mode of 4.25. Interrater reliability was high and significantly greater than zero [Pearson's multiple $R = 0.616$; $F(3,18) = 3.67$, $MS_e = 1.80$, $p < .03$].

A relatively small proportion (22.7%) of the participants received an awareness rating of less than 2. Those in this range typically reported that they guessed on the BD question (e.g., "I liked the rounder one better so I picked it") and had no clear notion that the stimuli could be ordered. Most (45.5%) of the participants scored at or better than 3, but less than 4. Such participants generally had a vague notion that the stimuli could be ordered but had little or no notion that such an ordering could support a transitive choice. A moderate proportion (31.7%) of the participants received a 4 or greater. These participants clearly explained that the stimuli could be ordered and that the ordering supported a correct transitive choice. A typical response was "If one beats out another, and that second one beats out the third, then the first one should also beat out the third."

Premise Pair Performance

As is shown in Figure 2, premise pair performance shows typical serial position effects (i.e., higher performance on pairs including an end item, AB and DE, than on pairs created entirely from inner items, BC and CD), observed ubiquitously in animal (rat, primate, pigeon) and human studies (for a review, see Wynne, 1995). In addition, both the informed and the uninformed participants learned the premise pairs at better than chance levels (all t_D tests were significant). A repeated measures
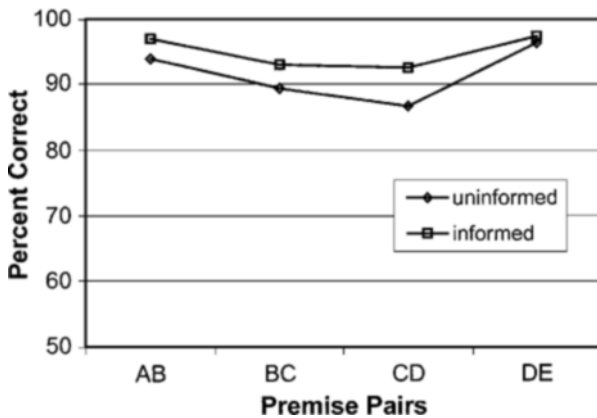
Figure 2. Serial position effects for premise pairs in Experiment 1.

analysis of variance (ANOVA) shows that those in the informed condition learned premise pairs more accurately than did those in the uninformed condition [$F(1,39) = 11.15$, $MS_e = 0.016$, $p < .01$]. For the uninformed condition, task awareness was moderate (mean awareness = 3.38, $SEM = 0.22$); a repeated measures ANOVA with awareness as a covariate shows no relationship between awareness and premise pair performance ($F < 1$).

Transitive Performance

On the critical BD probe test, both the informed and the uninformed participants scored significantly above chance [informed, .98, $t_D(18) = 39.9$, $SEM = 0.012$, $p < .01$; uninformed = .87, $t_D(21) = 5.83$, $SEM = 0.065$, $p < .01$]. The informed participants were marginally better than the uninformed participants [$t(39) = 1.58$, $SEM = 0.065$, $p < .12$]. Performance on BD in the informed condition, however, is only marginally lower than perfect performance [$t(18) = 1.60$, $SEM = 0.012$, one-tailed $p < .07$], suggesting that potential differences between informed and uninformed performance may be masked by a ceiling effect. This difficulty was addressed in Experiment 2A.

Transitive Performance Relative to Awareness

The more important result is that, for uninformed participants, performance on the transitive task does not correlate with task awareness [$r(20) = 0.10$, $R^2 = .01$ (1% explained variability), $t_r(20) = 0.64$]. This lack of relationship between postexperimental awareness and performance yields several possible interpretations. First, it could be the case that a ceiling effect is preventing a relationship between awareness and performance from being detected. However, this possibility seems unlikely, given that BD performance in the uninformed condition is significantly below ceiling [$t(21) = 1.86$, $SEM = 0.07$, one-tailed $p < .04$]. Second, it is possible that the course of learning differs for aware versus unaware participants but that performance converges by the end of training. For example, it is possible that those who

become aware would be able to correctly infer the BD relationship at an earlier point in training than those who did not become aware. Both the possibility of ceiling effects and the possibility of differential performance profiles are addressed in Experiment 2A.

EXPERIMENT 2A

In order to more closely examine the possible relationship between performance and awareness, in Experiment 2A, we assessed BD performance throughout all phases of training by testing after each of the four training blocks. This manipulation had two aims. First, it is entirely possible that task awareness may correspond to faster acquisition of the capacity to make correct inferences but that BD performance for all participants converges near asymptotic levels by the end of training. Thus, assessing BD performance throughout training may reveal a relationship between awareness and changes in BD performance. Second, because BD performance after early training blocks is likely to be much closer to chance, we can potentially alleviate the ceiling effect observed in Experiment 1 by observing performance prior to task mastery.

Method

The 58 participants were from the same population and recruited in the same manner as those in the previous experiment. The participants were randomly assigned to either an informed ($n = 17$) or an uninformed ($n = 41$) condition. The materials, design, and procedure were identical to those used in Experiment 1, with the exception that a testing procedure was conducted after each training block. As in Experiment 1, awareness was assessed only after completion of the entire experiment.

Results and Discussion
Transitive Performance

It was expected that overall BD performance would be lower in Experiment 2A than in Experiment 1 owing to the inclusion of tests from the earlier training blocks before all items were well learned. This is precisely what was observed [Experiment 1, .93; Experiment 2A, .77; $t(97) = 3.03$, $SEM = 0.05$, $p < .01$]. Unexpectedly, final BD performance (Block 4 test only) is also nominally, but not significantly, lower in Experiment 2A than in Experiment 1 [Experiment 1, .93; Experiment 2A, .85; $t(97) = 1.25$, $SEM = 0.06$, $p < .21$]. This difference may be the result of response tenacity. That is, participants tend to stick with a first choice such so that guesses after Block 1 persist [$r(56) = 0.48$ for BDs of Blocks 1 and 4]. More important, for all the blocks, BD performance is significantly lower than ceiling ($t_1 = 8.86$, $SEM = 0.04$, $p < .01$; $t_2 = 4.74$, $SEM = 0.04$, $p < .01$; $t_3 = 4.16$, $SEM = 0.04$, $p < .01$; $t_4 = 4.07$, $SEM = 0.03$, $p < .01$).

Awareness and Performance

We speculated that alleviating ceiling effects observed in Experiment 1 would reveal a significant difference between informed versus uninformed BD performance. As

can be seen in Figure 3, there was a main effect of informedness [informed, .87; uninformed, .73; $t(56) = 2.17$, $SEM = 0.06$, $p < .02$] so that being informed at the outset resulted in significantly higher BD performance.

As in Experiment 1, for the uninformed participants, we observed no correlation between final BD performance (Block 4 only) and postexperimental awareness [$r = 0.16$, $R^2 = .02$, $t_r(39) = 0.98$], indicating no direct relationship between awareness and BD performance. It was hypothesized that postexperimental awareness might correspond with different performance profiles throughout training. The awareness ratings resulted in an average score of 3.79, with a standard deviation of 1.11, a median of 4.00, and a mode of 5.00. Interrater reliability was high and significantly greater than zero [Pearson's multiple $R = 0.789$; $F(3,37) = 20.38$, $MS_e = 0.96$, $p < .01$]. Treating BD performance from Blocks 1–4 as a repeated measure and awareness as a covariate, a complete factorial and trend analysis for the repeated measure was conducted. As can be seen in Figure 3, there was a main effect of training block on BD [Block 1 = 0.64, Block 2 = 0.73, Block 3 = 0.81, Block 4 = 0.79; $F(1,117) = 8.61$, $MS_e = 0.03$, $p < .01$], accounted for by a significant linear trend of those blocks [$F(1,39) = 12.73$, $MS_e = 0.06$, $p < .01$]. There was no main effect of awareness on BD performance [$F(1,39) = 1.02$, $MS_e = 0.31$, $p = .46$] and no significant interaction between BD block trend and awareness ($F < 1$).

Given a significant tendency for response tenacity, we hypothesized that those with higher awareness might change their minds less frequently when choosing B or D (i.e., 2 participants with an equal proportion of correct BD responses could have all their incorrect responses amassed early on, whereas those with lower certainty of their guesses might change their responses throughout the experiment). We thus did a post hoc analysis of *reversals*, which simply counts the number of times each participant changed responses from B to D or vice versa. The correlation of BD performance and reversals was
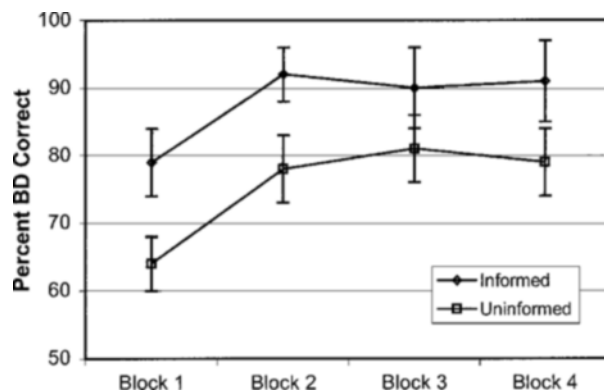

Figure 3. In Experiment 2, BD performance across blocks is better for informed than for uninformed participants. For both groups, performance nears an asymptote as early as Block 2.

not different from zero [$r = 0.015$, $R^2 = .0003$; $t_r(39) = 0.10$, $p = .79$].

Taking all these results together, it is somewhat puzzling that explicitly informing participants to look for a transitive relationship prior to training appears to facilitate performance on the BD task, whereas serendipitous awareness of a transitive relationship is unrelated to BD performance. Although volition in memory and awareness have been found to be dissociable (Richardson-Klavehn & Gardiner, 1995, 1998), it is still puzzling that performance is enhanced only when participants know at the outset that a hierarchy exists. It may be that instruction facilitates performance because it consciously constrains searches, whereas (at least for early learning) serendipitous awareness proceeds in parallel with task performance and thus offers no such useful constraint. Whether awareness corresponds to performance after significant time has passed certainly bears further investigation.

EXPERIMENT 2B

One factor limiting any strong claims about the lack of influence of awareness on performance is that we have thus far been limited to null effects in this regard. However, it is possible to design a manipulation in such a way that a relationship between awareness and performance corresponds to the null hypothesis and a lack of relationship between performance and awareness corresponds to the alternative hypothesis. By interrupting training as soon as BD performance reaches its asymptote and assessing awareness of the hierarchy at that point, it may be possible to reject the notion that awareness precedes performance.

If awareness of the hierarchy does facilitate BD performance, a high level of awareness should necessarily precede the asymptotic level of performance. We chose a criterion of at least seven out of eight correct BD responses on a given test for ending training, because that score constitutes an individual's above-chance performance and is roughly equivalent to group mean performance observed in Experiment 2A. Thus, if awareness at the time of training interruption (Experiment 2B) is comparable with postexperimental awareness (Experiment 2A), this would suggest that awareness is high at least as early as when BD performance reaches asymptotic levels. However, if awareness is significantly lower following interruption than following completed training, this would constitute positive evidence against any notion that hierarchical awareness is a prerequisite for BD performance.

Method

The 41 participants were from the same population and recruited in the same manner as those in the previous experiment. The participants were all assigned to the uninformed condition. The materials, design, and procedure were identical to those used for the uninformed participants in Experiment 2A, with one exception. The

participants were stopped immediately following the first testing block, in which BD performance was greater than or equal to seven of eight correct responses, and then were given the postexperimental awareness assessment. If BD performance did not reach seven of eight correct, the experiment continued to completion, as in Experiment 2A.

Results and Discussion

Of 41 participants, 30 reached criterion after the first training block, 6 after the second block, 3 after the third block, and only 2 did not reach the criterion, so that the experiment ended after Block 4. Consistent with the intent of the experiment, final BD performance (uninformed only) between Experiment 2A (.79) and Experiment 2B (.78) did not differ [$t(80) = 0.12$, $SEM = 0.056$]. The awareness ratings resulted in an average score of 2.82, with a standard deviation of 1.47, a median of 2.82, and a mode of 5.00. Interrater reliability was high and significantly greater than zero (Pearson's multiple $R = 0.943$; $F(3,37) = 85.20$, $MS_e = 0.327$, $p < .01$].

BD performance did not differ between Experiments 2A and 2B; however, as can be seen from Figure 4, interrupting training did significantly diminish postexperimental awareness [$t(80) = 3.12$, $SEM = 0.29$, $p < .01$]. Although the majority of the participants who completed training (Experiment 2A) were relatively aware, those who were not overtrained (Experiment 2B) were largely unaware. Analysis of effect size ($\omega^2 = 0.06$) shows a moderate effect (see Keppel, 1991), indicating substantial differences in awareness for complete versus incomplete training despite comparable scores in the BD probe trials.

Thus, differential awareness between Experiments 2A and 2B suggests at least two interpretations. First, when training continues after BD performance has neared an asymptote, awareness continues to increase. Second, the result constitutes additional convergent evidence for the dissociation between performance and task awareness.
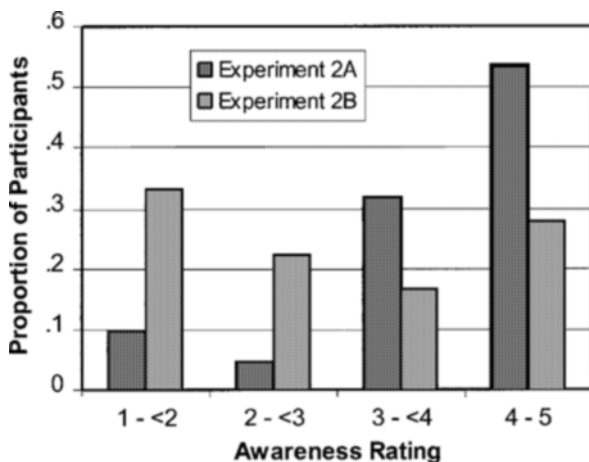

Figure 4. Proportion of participants by level of awareness. In Experiment 2A ($N = 41$), the participants completed all four training blocks, whereas in Experiment 2B ($N = 36$), the participants were stopped when performance reached a criterion of at least 7/8 correct on BD pairs.

GENERAL DISCUSSION

The central finding of this paper is that performance on the inferential judgment (BD) does not depend on knowledge or awareness of a superordinate hierarchy, nor does performance correlate with serendipitous awareness of that hierarchy. However, Experiment 1 shows that being informed of the existence of a hierarchy at the outset does improve performance on the inferential pair. For those participants who were not informed prior to training, postexperimental assessment reveals that many participants discovered the hierarchy on their own but that a substantial proportion remained unaware of any hierarchical structure. Performance on the inferential pair was significantly above chance, regardless of the level of awareness. Furthermore, the correlation between postexperimental awareness and BD performance was near zero. Experiment 2A replicates the main findings of Experiment 1 and further shows that the course of acquisition does not vary with postexperimental awareness. In addition, performance on the BD pairs reaches near asymptotic levels as early as the second block of training. Experiments 2A and 2B together show that awareness does not reach high levels prior to BD's reaching asymptotic levels. Experiment 2B also shows that increasing the amount of training beyond that which is required for high-BD performance may increase the level of awareness. It should be noted that our conclusions are based largely on correlational data. Since both postexperimental awareness and BD performance are necessarily dependent measures, our within-subjects measures are limited to a correlational analysis. However, because we are proposing that there is no relationship between awareness and performance, we are not limited by the inability to infer causal direction normally inherent in correlational designs. In addition, although some of our conclusions are based on null findings (e.g., zero correlations), we do demonstrate positive results by dissociating awareness and performance between Experiments 2A and 2B.

Overall, the results suggest that there is not a necessary relationship between awareness of a transitive relationship and transitive performance. This result is interesting for several reasons. First, it is fairly widely held that implicit tasks lack the flexibility of declarative tasks (e.g., Dienes & Berry, 1997). Because the TI task is a relational task and clearly requires a great deal of representational flexibility and organization, this finding constitutes further evidence that implicit tasks may rely on flexible representations (e.g., Reinitz & Demb, 1994; Srinivas et al., 1997; for a review, see Willingham, 1997). Second, although declarative memory and relational learning are arguably highly similar in function and structure (e.g., Cohen & Eichenbaum, 1993), they are not interchangeable constructs.

This distinction between relational learning and conscious memory converges with neuropsychological findings. Some studies have raised the possibility that hippocampal-system-dependent tasks necessarily in-

volve awareness (e.g., Clark & Squire, 1998). However, recent studies show that certain relational learning tasks do require the hippocampal system but do not require awareness (Chun & Phelps, 1999; Schacter et al., 1995). In addition, it is clear that such tasks as rote memorization may be declarative in that they are fully conscious but that they lack relational properties, such as flexibility and generalizability (e.g., Anderson, 1974). The confluence of all these considerations provides converging evidence for our finding that relational learning and declarative memory are dissociable.

A second way of interpreting the present findings, as well as related findings (Chun & Phelps, 1999; Rajaram & Coslett, 2000a, 2000b; Schacter et al., 1995) is that relational learning may well entail awareness but that proper organization of the contingencies and awareness follow a different time course. Clearly, the notion that awareness precedes performance must be abandoned, but it may well be that the formation of appropriate associations manifests itself behaviorally prior to full awareness. A third possible explanation for the present findings is that TI does not actually constitute a higher order relational task. For instance, some have proposed (e.g., von Fersen, Wynne, Delius, & Staddon, 1991) that somewhat simple associations (i.e., *value transfer* from end items) could account for TI performance in animals, and by extension, it may be reasonable to suggest that TI in humans could be accounted for by simple, nonrelational mechanisms. However, the tenability of value transfer theory remains an open question.

Another intriguing result of the present experiments is that the informing of participants at the outset improves performance over both serendipitous awareness and remaining unaware. One explanation may be that prior awareness of a hierarchical structure provides constraints that restrict errors and guide learning in appropriate ways from the outset. Serendipitous awareness, on the other hand, must occur after appropriate relationships have been discovered and their validity tested. Although this explanation is consistent with our observations and has some intuitive appeal, it has not yet been subject to empirical scrutiny. The finding that performance is not improved by serendipitous awareness but is improved by task instructions raises many questions about the role of awareness in learning. Clearly, more empirical work on this issue is needed before a satisfactory account can be proposed.

REFERENCES

ANDERSON, J. R. (1974). Verbatim and propositional representations of sentences in immediate and long-term memory. *Journal of Verbal Learning & Verbal Behavior*, **13**, 149-162.

ANDERSON, J. R. (1981). Effects of prior knowledge on memory for new information. *Memory & Cognition*, **9**, 237-246.

BOWER, G. H., CLARK, M. C., LESGOLD, A. M., & WINZENZ, D. (1969). Hierarchical retrieval schemes in recall of categorized word lists. *Journal of Verbal Learning & Verbal Behavior*, **8**, 323-343.

BRADSHAW, G. L., & ANDERSON, J. R. (1982). Elaborative encoding as an explanation of levels of processing. *Journal of Verbal Learning & Verbal Behavior*, **21**, 165-174.

BRYANT, P. E., & TRABASSO, T. (1971). Transitive inferences and memory in young children. *Nature*, **232**, 255-257.

CHUN, M. M., & PHELPS, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nature Neuroscience*, **2**, 844-847.

CLARK, R. E., & SQUIRE, L. R. (1998). Classical conditioning and brain systems: The role of awareness. *Science*, **280**, 77-81.

CLARK, R. E., & SQUIRE, L. R. (1999). Human eyeblink classical conditioning: Effects of manipulating awareness of the stimulus contingencies. *Psychological Science*, **10**, 14-18.

COHEN, N. J., & EICHENBAUM, H. (1993). *Memory, amnesia, and the hippocampal system*. Cambridge, MA: MIT Press.

CRAIK, F. I. M., & TULVING, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, **104**, 268-294.

DE BOYSSON-BARDIES, B., & O'REGAN, K. (1973). What children do in spite of adults' hypotheses. *Nature*, **246**, 531-534.

DIENES, Z., & BERRY, D. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin & Review*, **4**, 3-23.

DUSEK, J., & EICHENBAUM, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences*, **94**, 7109-7114.

EICHENBAUM, H. B. (1999a). Conscious awareness, memory and the hippocampus. *Nature Neuroscience*, **2**, 775-776.

EICHENBAUM, H. B. (1999b). The hippocampus and mechanisms of declarative memory. *Behavioural Brain Research*, **103**, 123-133.

EICHENBAUM, H. B., STUART, C., & MORRIS, R. G. (1990). Hippocampal representations in place learning. *Journal of Neuroscience*, **10**, 3531-3542.

HEBB, D. O. (1949). *The organization of behavior*. New York: Wiley.

HOLYOAK, K. J., & SPELLMAN, B. A. (1993). Thinking. *Annual Review of Psychology*, **44**, 265-315.

JACOBY, L. L., & DALLAS, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General*, **110**, 306-340.

KALLIO, K. D. (1982). Developmental change on a five term transitive inference. *Journal of Experimental Child Psychology*, **33**, 142-164.

KEPPEL, G. (1991). *Design and analysis: A researcher's handbook* (3rd ed.). Englewood Cliffs, NJ: Prentice-Hall.

KROLL, N. E. A., KNIGHT, R. T., METCALFE, J., WOLF, E. S., & TULVING, E. (1996). Cohesion failure as a source of memory illusions. *Journal of Memory & Language*, **35**, 176-196.

MCNAMARA, T. P., HARDY, J. K., & HIRTLE, S. C. (1989). Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 211-227.

MENZEL, E. W., JR. (1969). Responsiveness to food and signs of food in chimpanzee discrimination learning. *Journal of Comparative & Physiological Psychology*, **68**, 484-489.

MILNER, B. (1972). Disorders of learning and memory after temporal lobe lesions in man. *Clinical Neurosurgery*, **19**, 421-446.

MOSCOVITCH, M. (2000). Theories of memory and consciousness. In E. Tulving & F. I. M. Craik (Eds.), *The Oxford handbook of memory* (pp. 609-625). New York: Oxford University Press.

NGUYEN, D. B., & REVLIN, R. (1993). Transitive inferences from narrative relations. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 1197-1210.

O'KEEFE, J., & NADEL, L. (1978). *The hippocampus as a cognitive map*. Oxford: Oxford University Press, Clarendon Press.

PEARCE, J. M., ROBERTS, A. D. L., & GOOD, M. (1998). Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. *Nature*, **396**, 75-77.

PIAGET, J. (1928). *Judgment and reasoning in the child*. London: Kegan, Paul, Trench, & Trubner.

RAJARAM, S., & COSTLETT, H. B. (2000a). Acquisition and transfer of new verbal information in amnesia: Retrieval and neuroanatomical constraints. *Neuropsychology*, **14**, 427-455.

RAJARAM, S., & COSLETT, H. B. (2000b). New conceptual associative

learning in amnesia: A case study. *Journal of Memory & Language*, **43**, 291-315.

REINITZ, M. T., & DEMB, J. B. (1994). Implicit and explicit memory for compound words. *Memory & Cognition*, **22**, 687-694.

RICHARDSON-KLAVEHN, A., & BJORK, R. A. (1988). Measures of memory. *Annual Review of Psychology*, **39**, 475-543.

RICHARDSON-KLAVEHN, A., & GARDINER, J. M. (1995). Retrieval volition and memorial awareness in stem completion: An empirical analysis. *Psychological Research*, **57**, 166-178.

RICHARDSON-KLAVEHN, A., & GARDINER, J. M. (1998). Depth-of-processing effects on priming in stem completion: Tests of the voluntary-contamination, conceptual-processing, and lexical-processing hypotheses. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 593-609.

RICKARD, T. C., & GRAFMAN, J. (1998). Losing their configural mind: Amnesic patients fail on transverse patterning. *Journal of Cognitive Neuroscience*, **10**, 509-524.

ROEDIGER, H. L., III (1990). Implicit memory: Retention without remembering. *American Psychologist*, **45**, 1043-1056.

ROEDIGER, H. L., III, & SRINIVAS, K. (1993). Specificity of operations in perceptual priming. In P. Graf & M. Masson (Eds.), *Implicit memory: New directions in cognition, development and neuropsychology* (pp. 17-45). Hillsdale, NJ: Erlbaum.

SCHACTER, D. L. (1998). Memory and awareness. *Science*, **280**, 59-60.

SCHACTER, D. L., CHURCH, B., & BOLTON, E. (1995). Implicit memory in amnesic patients: Impairment of voice-specific priming. *Psychological Science*, **6**, 20-25.

SIEMANN, M., & DELIUS, J. D. (1993). Processing of hierarchic stimulus structures has advantages in humans and animals. *Biological Cybernetics*, **71**, 531-536.

SIEMANN, M., & DELIUS, J. D. (1996). Influences of task concreteness upon transitive responding in humans. *Psychological Research*, **59**, 81-93.

SIEMANN, M., DELIUS, J. D., & WRIGHT, A. A. (1996). Transitive responding in pigeons: Influences of stimulus frequency and reinforcement history. *Behavioural Processes*, **37**, 185-195.

SQUIRE, L. R., & KNOWLTON, B. J. (1999). The medial temporal lobe, the hippocampus, and the memory systems of the brain. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., pp. 765-779). Cambridge, MA: MIT Press.

SQUIRE, L. R., ZOLA-MORGAN, S., & CHEN, K. S. (1988). Human amnesia and animal models of amnesia: Performance of amnesic patients on tests designed for the monkey. *Behavioral Neuroscience*, **10**, 210-221.

SRINIVAS, K., GREENE, A. J., & EASTON, R. D. (1997). Visual and tactile memory for 2-D patterns: Effects of changes in size and left–right orientation. *Psychonomic Bulletin & Review*, **4**, 535-540.

SUTHERLAND, R. J., & RUDY, J. W. (1989). Configural association theory: The role of the hippocampal formation in learning, memory, and amnesia. *Psychobiology*, **17**, 129-144.

VON FERSEN, L., WYNNE, C. D., DELIUS, J., & STADDON, J. E. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **17**, 334-341.

WEAVER, J. E., STEIRN, J. N., & ZENTALL, T. R. (1997). Transitive inference in pigeons: Control for differential value transfer. *Psychonomic Bulletin & Review*, **4**, 113-117.

WILLINGHAM, D. B. (1997). Implicit and explicit memory do not differ in flexibility: Comment on Dienes and Berry (1997). *Psychonomic Bulletin & Review*, **4**, 587-591.

WYNNE, C. D. L. (1995). Reinforcement accounts for transitive inference performance. *Animal Learning & Behavior*, **23**, 207-217.

NOTE

1. The set of experiments presented here do not include memory-impaired participant groups. However, the neuropsychology literature contains parallel and convergent findings for our claims; thus, some background is reviewed in this paper.

APPENDIX A
Preexperimental Briefing

All participants. "In this experiment, two figures will appear simultaneously on the computer screen. You are to select the "correct" figure. At first, this will be by trial and error; however, with practice, you will find that the correct figure is easily learned."

Informed participants only. "We are specifically interested in transitive inference. Simply put, transitive inference is a form of logical reasoning, which allows us to know things we have not specifically seen. For example, suppose you see Al and Bob side by side, and Al is taller. Later, you see Bob and Carl, and Carl is taller. You can infer correctly that Al is taller than Carl without ever seeing them together.

"In this task, we have randomly ordered five shapes. Much like the height task with Al, Bob, and Carl, you will be presented with all four adjacent pairs of shapes and will learn which is correct from each pair. Later, two of the shapes which are not adjacent (like Al and Carl) will be paired for the first time, and you should be able to infer which is correct even though you were never taught which one was correct. In other words, you must discover the ordering to solve the task."

APPENDIX B

Postexperimental Questionnaire (for uninformed participants only).

(page breaks, with instructions not to return to a prior page, followed Questions 3, 7, and 8.)

1. What did you think we were trying to find out with this experiment?

2. What did you think was the point of the trials where you were not told if you were right or wrong (the no feedback condition)?

3. Regarding question 2, were all the pairs of shapes in the no feedback condition the same as pairs you had already learned when feedback was given?___Yes ___Not Sure___No

If no, do you think there was a correct answer? ___Yes___ Not Sure___No

If you believe there was a correct answer, explain why:

4. Were you familiar with these symbols before today's testing? ___Yes___ No If yes, from where?

5. You were given the following pair of shapes several times, but never told an answer.

ま　め
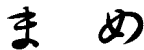
Circle the shape you believe is correct (guess if necessary):

6. Regarding question 5, what reason (if any) did you have for your choice (check one):

___ There is a logically correct choice because (explain):

___ One just seemed right but I can't explain why.

___ I guessed: There may be a correct shape, but I don't know what it is.

___ I made a random choice because there is no correct choice.

7. What strategy (if any) did you use to learn the shapes (check one):

___ I gave them names.___ I memorized each shape.___ I memorized part of the each shape.

___ I just watched and eventually got it. ___I used their similarity to familiar shapes. ___ No strategy.___ Other strategy (please describe in the space below)

8. Based on your understanding of the relationships among these shapes, arrange the shapes appropriately. Use the numbers provided to stand for the shapes.

め か
ゑ せ ま

9. If applicable, please indicate how much knowledge you have of formal logic, syllogism, or transitive inference:

Awareness Ratings

The questionnaire was designed to give increasingly leading questions in order to assess how explicitly the hierarchy among the characters was understood. The following criteria were agreed upon as having been important, following all raters' evaluations. Questions 1–3, 6: some assertion that there was a hierarchy, serial ordering, and so on. Confidence was evaluated according to the response for Question 7. Correctness of responses for Questions 5 and 8 was explicitly disregarded because of the potential to artificially inflate the awareness-accuracy correlations.

Ratings on a scale of 1–5 were agreed to represent the following: 5 = *knowledge of hierarchy, knowledge of logical BD choice*; 4 = *some knowledge of hierarchy, possible knowledge of logical BD choice*; 3 = *possible knowledge of hierarchy, possible or vague knowledge of logical BD choice*; 2 = *highly questionable knowledge of hierarchy, confusion or incorrect assertion about logical BD choice*; 1 = *no knowledge of hierarchy, no indication of any logical processes governing BD choice.*