

Acquisition context and the use of causal rules

RICHARD A. CARLSON

The Pennsylvania State University, University Park, Pennsylvania

and

WALTER SCHNEIDER

University of Pittsburgh, Pittsburgh, Pennsylvania

Judgmental asymmetries in using causal knowledge (e.g., for prediction or diagnosis) have been attributed to the inherent directionality of causal knowledge. The present study examines the effect of acquisition context—representations used for initial instruction, and the type of judgment required during acquisition—on judgments using causal rules. In contrast to traditional concept formation research, this paradigm examined the development of procedures for using rules, rather than rule induction. College-student subjects learned to use causal rules describing digital logic gates, receiving instruction with either verbal rules or truth tables, and practicing either predicting or verifying logic-gate outputs. After 200 trials of practice with each rule, subjects were transferred to the untrained judgment task. Transfer was strongly asymmetrical. Subjects trained to make prediction judgments were slowed substantially by transfer to the verification task, while subjects trained to make verification judgments had little difficulty with transfer to the prediction task. Truth-table representations resulted in superior performance, especially for verification judgments. Contrary to prediction, verification judgments always required more time. The results demonstrate that acquisition context may be partly responsible for judgmental asymmetries, and imply that examining conditions of acquisition is important for understanding how causal knowledge is used.

Much of our knowledge can be analyzed as hierarchies of causal relations, and in many domains these relations can be represented by deterministic rules. Rules, of course, may be represented in a variety of notations (e.g., Johnson-Laird, 1983), and causal knowledge may be used in several ways (e.g., Einhorn & Hogarth, 1986), including reasoning from cause to effect (*prediction*), from effect to cause (*diagnosis*), and from a cause-effect combination to confirmation or disconfirmation of a causal rule (*verification*). In the present study, we examined how individuals learn to use causal rules. Specifically, we considered whether the initial representation provided by instruction and the type of judgment required during initial practice affect the procedures that subjects develop for using causal rules. The results indicate that these factors have a substantial impact on the acquisition and use of causal knowledge.

Most research on causal thinking has used paradigms in which subjects reason with unpracticed verbal rules, as in causal-rule versions of the Wason card-selection task (e.g., Cheng & Holyoak, 1985), new and essentially arbitrary causal relations (e.g., Schustak & Sternberg, 1981), or assumed prior or normative knowledge (e.g., Tversky

& Kahneman, 1980). These researchers and many others (see Einhorn & Hogarth, 1986, for a review) have concluded that there is something special about the representation of causal knowledge that affects reasoning processes that operate on it. Tversky and Kahneman (1980, p. 51), for example, attributed inferential asymmetries in thinking about causal relations to the operation of a directional "causal schema" that "evolves from causes to consequences," whereas Einhorn and Hogarth (1986) suggested that diagnostic thinking is more difficult because considering an effect first provides a temporal cue that conflicts with the cause-to-effect ordering built into knowledge of a causal relation. These suggestions are difficult to evaluate because rule learning and the acquisition of cognitive procedures have not been considered in the context of questions about causal knowledge. It is plausible, however, that typical task demands during acquisition of causal knowledge encourage individuals to develop procedures that are optimized for prediction. If so, it should be possible to encourage the development of procedures optimized for other types of thinking with the same causal content.

There has been a great deal of research on the learning of concepts defined by simple logical rules (e.g., Bourne, 1970; Bruner, Goodnow, & Austin, 1956; Levine, 1975). Experimental tasks used in this research fall into two categories: tasks in which the subjects are required to identify relevant stimulus attributes given a rule (for example, disjunction), and tasks in which the subjects are required to identify the correct rule for classification, given the relevant attributes (Bourne, 1970). In either case, the task requires induction on the basis of a series of positive

This research was supported in part by the Army Research Institute, under Contract MDA903-86-C-0149 to Walter Schneider. We would like to thank Rondall Khoo and Robin Yaure for their assistance in analyzing the data, and several anonymous reviewers for helpful comments on earlier drafts. Correspondence should be addressed to Richard Carlson, Department of Psychology, The Pennsylvania State University, 613 Moore Building, University Park, PA 16802.

and negative instances of a concept. Although they are not explicitly described as "causal" (either theoretically or to the subjects), these paradigms can be construed as examples of a general paradigm that requires subjects to learn causal rules describing feedback contingencies (Dulany, 1974). That is, stimulus attributes "cause" the classification of a stimulus according to a logical rule.

Previous research on conceptual rule learning supports the hypothesis that the representation with which subjects are instructed affects rule acquisition. For example, Dodd, Kinsman, Klipp, and Bourne (1971) demonstrated that training or instruction concerning truth tables resulted in positive transfer to rule learning. Bourne (1970) suggested that experience with rule-learning problems led subjects to develop a general truth-table schema for such problems.

In the present study, we were concerned with the procedures that people develop for using causal knowledge. Concept-learning research has generally employed discovery paradigms in which the subject must induce rules or relevant attributes from a series of instances (Bourne, 1970), and the acquisition procedure is stopped once the subject has demonstrated accurate knowledge of the concept. Practice in these paradigms has been operationalized as solution of a series of induction problems in which rules are applied to new attributes on each problem, rather than repeated use of the same attribute-rule combinations. The theoretical focus of previous research has therefore been on the development and use of problem-solving strategies, such as the truth-table strategy (Bourne, 1970) and hypothesis testing (Levine, 1975). Our theoretical focus led us to use an alternative paradigm, in which subjects received instructions about both relevant stimulus attributes and rules and practiced using those rules for a substantial number of judgments. In addition, we explicitly described the rules as causal and presented them in an experimental context that emphasized their causal nature.

In the present study, the subjects learned causal rules that describe the input-output functions of digital electronic components known as *logic gates*. A logic gate is a "building block" component that performs a simple logical function on binary symbols; for example, an AND gate may be described as producing a binary 1 whenever it receives a conjunction of 1s as inputs. These rules have several advantages for investigating the learning and use of causal rules in general. First, they embody a set of primitive logical functions that presumably are capable of representing causal relations in many content domains. Second, the content of these rules is unfamiliar to most college-student subjects, allowing us to study the acquisition of procedures for using these rules. Third, the hierarchical nature of digital electronic circuits provides a good experimental model for causal knowledge domains. Finally, the rules are deterministic, and their binary outputs make collection of latency data straightforward. Figure 1 shows symbols for the five logic gates that subjects learned in the present study, along with truth tables and verbal rules representing the input-output function of each gate.

The subjects received instruction about these rules with one of two representations (verbal rules or truth tables) and practiced making either prediction or verification judgments. We expected both manipulations to affect the procedures that the subjects developed to make judgments using these rules. First, the verbal statements of causal relations have an explicit cause-to-effect form. Here, verbal rules had the form "if input, then output," where "input" stands for a value or combination of values, and "output" also stands for a value. Thus verbal rules should be more appropriate for prediction than for verification judgments. Truth tables, on the other hand, do not have an inherent causal direction and should be equally appropriate for both types of judgment. Second, practice with prediction judgments should encourage the subjects to develop procedures for considering information in causal, input-to-output sequence. Verification judgments, on the other hand, can be made by considering the entire pattern of input, symbol, and output, and should encourage the development of procedures without an inherent cause-to-effect directionality. Thus there should be no lasting advantage of prediction over verification judgments. Also, transfer between judgment types should be asymmetric: subjects trained to make prediction judgments should find verification difficult, whereas subjects trained to make verification judgments should have little difficulty with prediction.

The causal nature of the present task was established by describing the rules as governing the operation of electronic components. In other research (for examples, see Einhorn & Hogarth's, 1986, review), verbal description has been established as a strong manipulation of subjects' interpretation of data as causal. We had two reasons for contrasting verification, rather than diagnosis (inferring inputs on the basis of a given output), with prediction. First, verification judgments allowed unambiguous binary responses comparable to the prediction judgments. Second, many "diagnosis" tasks that have been studied really involve judgments in which input state (cause), output state (effect), and (sometimes implied) causal relation are presented simultaneously (as in the examples presented by Tversky & Kahneman, 1980).

Besides the hypothesized directionality, other properties of the alternative representations may also influence the course of learning. Verbal rules are more concise and thus more easily held in working memory, and may drop out of verbal rehearsal earlier in learning, whereas a truth table initially puts a greater burden on working memory. Pilot work and subjects' comments made it clear that the verbal rules were initially more memorable. Also, the greater difficulty of deriving verification judgments from verbal rules should initially trade off with the memory advantage. Therefore, we expected that subjects who received verbal rule instruction and made verification judgments would initially need to refresh their memory for the rules more often than would other groups. After declarative knowledge of appropriate procedures was developed, however, the need to refresh memory should

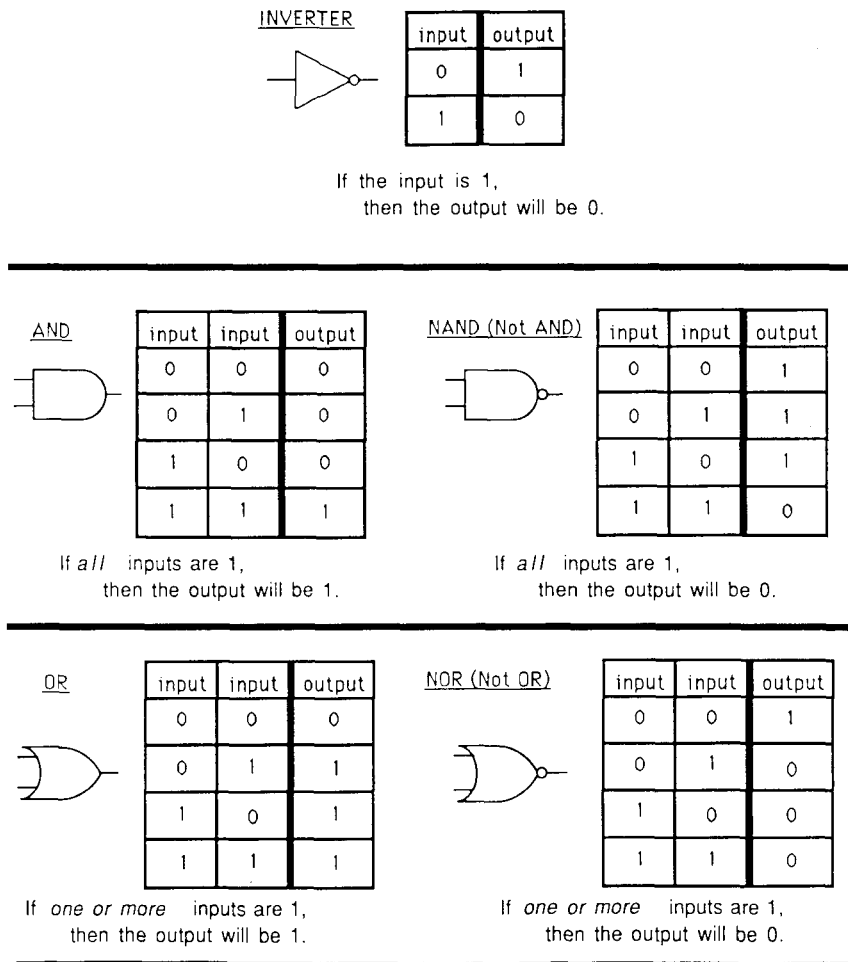


Figure 1. Logic-gate symbols, truth tables, and verbal rules.

decline faster for the subjects given verbal rules than for those given truth tables.

METHOD

Subjects

Forty-eight University of Illinois students volunteered to serve as subjects in return for payment. All subjects were right-handed. All reported that they had normal or corrected-to-normal vision, and that they had had no previous experience with digital electronics or computer programming.

Design

The experimental design was a 2x2 factorial, with judgment type during learning (prediction or verification) and type of representation (verbal rule or truth table) as between-subjects factors. Practice block and gate type were varied within subjects.

Experimental Tasks

Each subject completed several tasks, in two sessions. The tasks are described in this section, and the sequence of events is detailed in the following section.

Logic-gate judgments. An IBM PC/XT microcomputer with an IBM monochrome display was used to present practice trials for the logic-gate judgments. Each trial began with the display of a ready message. When the subject pressed the space bar on the computer

keyboard, a logic-gate symbol appeared centered on the screen. For prediction judgments, the subjects responded by pressing keys indicating that the output should be 0 or 1. For verification judgments, the subjects pressed keys indicating that the displayed output was correct or incorrect. Response keys were located in the center row of the computer keyboard, and carried labels marked "0" and "1" for the prediction task, and "Y" (correct) and "N" (incorrect) for the verification task. Figure 2 shows sample displays for prediction (Panels A and C) and verification (Panels B and D) judgments. The display was approximately 5° wide and 3° high at the viewing distance of approximately 55 cm.

Both two- and three-input gates appeared during practice. Pilot research had demonstrated that subjects had little difficulty in generalizing to the three-input cases, regardless of the representation used for instruction.

The "help" displays, also shown in Figure 2, presented the representations that were used for initial instruction. These representations were not displayed at the beginning of a trial. The subjects could, however, request help at any time during the trial (i.e., before or after a response was given).

When the subject responded, accuracy feedback was presented below the gate symbol. During the second experimental session response time feedback in milliseconds was also displayed. This display remained visible until the subject pressed the space bar, a which time it was replaced by the ready message for the next trial. On every 20th trial, summary feedback was provided, indicating the proportion of correct responses on the last 20 trials.

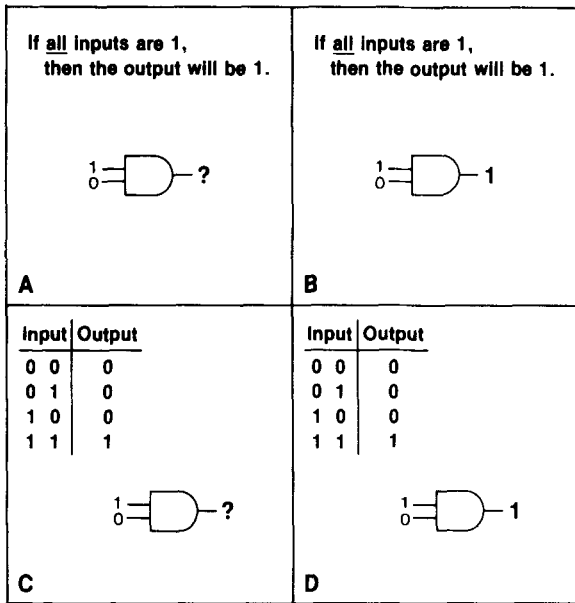


Figure 2. Sample displays for logic-gate judgments: (A) prediction/verbal rule; (B) verification/verbal rule; (C) prediction/truth table; (D) verification/truth table.

Problem solving. The subjects attempted to solve prediction and diagnosis problems in networks modeled after actual digital circuits. Prior to practicing with logic gates, the subjects received simplified versions of these problems, in which abstract components replaced logic gates. After the logic-gate practice, similar problems with networks of logic gates were presented. These problems were originally intended to assess the effects of different logic-gate training conditions. However, few differences were detected and variability was very great; therefore, data from these problems are not presented here.

Reproduction of representations. The subjects were asked to reproduce the representations they had been given for gate functions. Each subject was given a booklet with the symbol for a gate drawn at the top of each page. The experimenter asked each subject to draw the truth table or write the rule given at the beginning of the experiment. In addition, the experimenter told the subject that many people developed their own rules, and that any personal rules should be described in as much detail as possible.

Procedure

The subjects were assigned to experimental conditions in a counterbalanced order as they arrived for the experiment. Each subject participated in two individual sessions of approximately 2 h each, separated by about 48 h. Table 1 describes the overall order of events.

Three experimenters assisted with the data collection. The experimenters were counterbalanced over experimental conditions.

Session 1. On arrival at the lab, the subjects heard a brief description of the experiment and were asked about their previous experience with electronics or computer programming. They then performed a brief problem-solving task involving networks of abstract components.

After completing the problem-solving task, the subjects were introduced to the functions of the five logic gates, in the constant order (1) inverter, (2) AND, (3) OR, (4) NAND, (5) NOR. To introduce each gate, the experimenter first started the computer program to display the gate symbol and verbal rule or truth table. The experimenter named the gate and described the function, either by

reading the verbal rule aloud or by reading each row of the truth table (e.g., "If both inputs are 0, the output is 0"). Although the truth tables were thus given verbal descriptions, it was emphasized that the table as a whole represented the function. The experimenter then described the practice procedure and use of the help key (see above). The subjects practiced each gate type for 48 trials before they were introduced to the next gate.

After this procedure had been repeated for each gate type, the subjects practiced for four blocks of 100 trials. In each block, the gates appeared in a random order, with the gates mixed in proportion to the possible input combinations for each gate. Because both two- and three-input examples of AND, OR, NAND, and NOR gates were used, this resulted in blocks consisting of 24 examples of each of these gate types and 4 examples of inverters. The subjects were asked to try to answer correctly on each trial. For these practice blocks, accuracy but not response time feedback was provided.

Session 2. The second experimental session began with four more blocks of practice with the logic-gate judgments. The procedure was identical to that of Session 1, except that the subjects were instructed to be fast as well as accurate, and response time feedback (in milliseconds) was provided on each trial.

Following these practice blocks, each subject attempted to solve eight problems involving networks of logic gates. After the subjects had completed these problems, the experimenter described the untrained gate judgment task (i.e., verification for subjects trained on prediction judgments, and vice versa). The subjects performed 100 trials of the untrained judgment task, with instructions to be fast as well as accurate. Gate type varied randomly within these transfer blocks. Help was not allowed on these trials.

Finally, the subjects attempted to reproduce the representations they had been given when the logic gates were introduced. They were then debriefed and paid for their participation.

RESULTS

Logic-Gate Learning

Introduction to logic gates. The subjects had little difficulty in understanding and using the rules of operation for individual gates. On the initial practice trials, with each rule appearing in a separate block of 48 trials, the subjects were very accurate (>96% correct) and very rarely used help (<2% of trials). The mean response time for correct responses on trials on which help was not used was 1,271 msec.

Mixed logic-gate practice, Session 1. Mixing the types of gates made the task much more difficult. The subjects'

Table 1
Sequence of Events (Approximate Time)

Session 1	
Introduction	(5 min)
Simplified problem solving	(15 min)
Introduction to logic gates (48 trials each)	(30 min)
Inverter, AND, OR, NAND, NOR	
Practice with logic gates randomized	(60 min)
4 blocks of 100 trials, accuracy feedback only	
Session 2	
Practice with logic gates randomized	(45 min)
4 blocks of 100 trials, accuracy and latency feedback	
Problem solving	(35 min)
Transfer to untrained judgment	(15 min)
Reproduction of rules	(5 min)
Debriefing	(5 min)

comments, as well as their objective performance, demonstrated this difference in difficulty. The average proportion of trials on which help was requested increased from less than .02 in the introductory blocks to .44 in the first block of mixed gates, proportion correct for trials on which help was not used fell from .96 to .90, and the mean response time for trials on which help was not used increased from 1,271 to 3,251 msec. This increase in difficulty indicates that discriminating among the gate symbols and accessing the appropriate rules was a major source of difficulty early in learning. This effect is similar to the contextual interference effect observed in verbal learning (e.g., Battig, 1979) and motor learning (e.g., Shea & Morgan, 1979), and is discussed in more detail below.

The proportion of trials on which subjects requested help provides an index of the difficulty of retrieving and using rules, given different representation and judgment types. Figure 3 shows the proportion of trials on which help was used by practice block and experimental condition. As expected, the use of help declined significantly with practice [$F(3,132) = 56.74, p < .0001, MS_e = .0190$]. This decline was more pronounced for groups using verbal rules, as reflected in a significant representation \times practice block interaction [$F(3,132) = 3.90, p = .011, MS_e = .0190$]. On early practice blocks, help was used more often by the verification-rule group than by any other group. This effect is reflected in a significant practice \times representation \times judgment interaction [$F(3,132) = 4.09, p < .01, MS_e = .0190$]. This result supports the hypothesis that it is more difficult to derive a verification judgment from the directional representation provided by a verbal rule. The use of help did not vary with gate type ($F < 1$), and no other effects were significant (all $ps > .05$).

Applying the rules may be a significant source of difficulty, even when the representation is perceptually available. The overall proportion correct was .96 when help was used. Fifty-eight percent of the subjects,

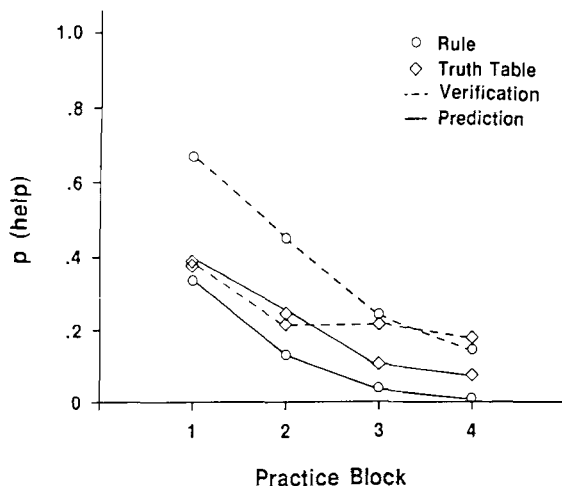


Figure 3. Use of help as a function of practice and condition.

however, made at least one error after using help, and some subjects' performance was as poor as .75 correct when using help. Accuracy when help was used did not differ as a function of experimental conditions ($F < 1$). When help was not used, accuracy was also nearly constant over experimental conditions ($F < 1$), with an overall mean of .90 correct.

There were no differences among experimental conditions in the latency to request help ($p > .25$). The overall mean time from the beginning of a trial until the help key was pressed was just over 5 sec (5,057 msec). This relatively long latency suggests that subjects requested help only after they had failed to retrieve a representation in which they had adequate confidence. Time to derive an answer once the help key was pressed was also relatively long ($M = 3,092$ msec), and did not differ among experimental conditions ($p > .25$).

Logic-gate practice, Session 2. Although the subjects were allowed to use help after response time feedback was introduced, the proportion of trials on which help was requested continued to decline, from about .10 in the last block of Session 1 to about .06 in Session 2. Trials on which help was used were excluded from the following analyses of judgment accuracy and latency.

Accuracy on the first block of speeded practice declined to .86 correct, from about .91 in the last block of Session 1. Accuracy on Blocks 2, 3, and 4 was .89, showing no further improvement. The improvement from Block 1 to subsequent blocks resulted in a reliable main effect of practice [$F(3,132) = 5.10, p = .002, MS_e = 9,758$]. Accuracy also varied with gate type [$F(4,176) = 5.99, p < .001, MS_e = 20,190$]. The proportion of correct responses was .91, .90, .89, .85, and .86 for inverter, AND, OR, NAND, and NOR gates, respectively. No other effects on accuracy were significant (all $ps > .20$).

Response times for correct judgments were analyzed with times greater than three standard deviations from each subject's mean ($< 1\%$ of the data) excluded. Response times declined over blocks of practice, with means of 1,714, 1,441, 1,334, and 1,289 msec for Blocks 1 through 4, respectively [$F(3,132) = 60.97, p < .0001, MS_e = 143,423$]. No other variables interacted with practice (all $ps > .10$).

Verification judgments were considerably slower ($M = 1,678$ msec) than prediction judgments [$M = 1,210$ msec, $F(1,44) = 26.73, p < .0001, MS_e = 1,967,620$]. Subjects who had used truth tables were faster ($M = 1,340$ msec) than subjects who had used verbal rules [$M = 1,549$ msec, $F(1,44) = 5.35, p = .026, MS_e = 1,967,620$]. Figure 4 shows that the advantage of truth-table over rule representations was primarily for verification judgments, although the judgment \times representation interaction is only marginally significant [$F(1,44) = 3.79, p = .058, MS_e = 1,967,620$].

Response times differed for the different types of gates [$F(4,176) = 55.00, p < .0001, MS_e = 252,057$]. As shown in Figure 5, inverters were judged most rapidly, AND and OR gates somewhat more slowly, and NAND

and NOR gates most slowly. Figure 5 also shows that these differences were greater for verification judgments. This is reflected in a significant gate type \times judgment type interaction [$F(4,176) = 5.83, p = .0002, MS_e = 252,057$]. No other effects on response time were significant (all $ps > .10$).

Transfer to untrained judgment task. According to the experimental hypotheses, transfer to the untrained judgment task should be asymmetric. That is, subjects trained to make verification judgments should have little difficulty with prediction judgments, whereas those trained to make prediction judgments should have trouble reversing their representations to make verification judgments. This prediction is supported by the data, as shown in Figure 6. This figure displays mean latency for correct judgments in the last block of training (trained task), and in the transfer (untrained) task, as a function of experimental condition. Overall accuracy for the transfer trials was .898, and did not differ among experimental groups [$F(1,44) < 1, MS_e = .00044$]. Although subjects trained on the verification task were actually faster (by approximately 200 msec) on the untrained prediction task, subjects trained on the prediction task were slower (by 800 msec) on the untrained verification task. This interaction of training judgment (prediction vs. verification) and task (trained vs. transfer) is significant [$F(1,44) = 104.94, p < .001, MS_e = 111,374$]. The type of repre-

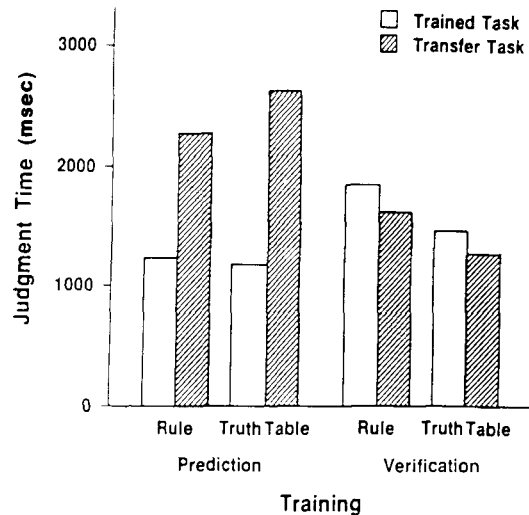


Figure 6. Response time as a function of experimental condition, comparison of final block of trained task and untrained transfer task.

sentation (verbal rule or truth table) did not interact with these variables ($p > .17$).

Reproduction of representations. All groups of subjects very accurately recalled the representations they had been given at the beginning of the experiment. Recall of verbal rules was scored as correct if the subject produced a recognizable paraphrase of the rule for each gate. Truth-table recall was scored as correct if all input patterns were listed together with the correct outputs. Both the subjects in the prediction condition and those in the verification condition correctly reproduced rules 95% of the time. The subjects in the prediction condition correctly reproduced truth tables 90% of the time; almost all of their errors involved omitting one or more input combinations. The subjects in the verification condition correctly reproduced truth tables 100% of the time.

About 45% of the subjects in the truth-table group reported that they had used verbal rules. Ninety-eight percent of these rules were correct. The subjects in the prediction-table group were almost twice as likely as those in the verification-table group to report verbal rules (58% vs. 31%). This finding provides further support for the view that verbal rules are especially appropriate for making forward, prediction judgments.

DISCUSSION

The present study demonstrated that the context in which causal rules are acquired—the representations used for initial instruction and the type of judgment required in initial practice—has a substantial impact on judgments using those rules, even after fairly extensive practice. The most striking result is that transfer between judgment types was strongly asymmetrical: although subjects trained to make prediction judgments performed much more poorly on transfer to verification judgments, subjects trained to make verification judgments had little difficulty in making predictions. This finding supports our hypothesis that

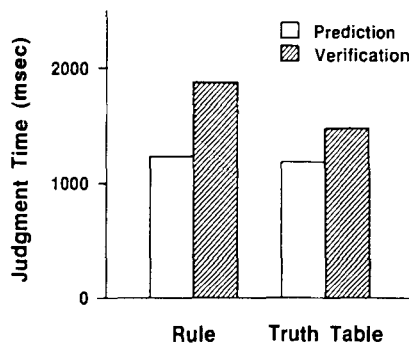


Figure 4. Response time as a function of experimental condition, Session 2.

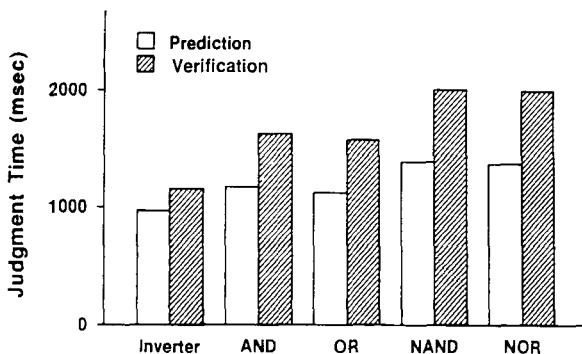


Figure 5. Response time as a function of judgment type and gate type, Session 2.

judgmental asymmetries in using causal knowledge are due at least in part to the context in which that knowledge is acquired.

The type of representation used to instruct the subjects also affected judgments. Overall, the subjects instructed with truth tables made faster judgments, consistent with previous results demonstrating the benefits of truth-table training for rule induction (Dodd et al., 1971). As pointed out by an anonymous reviewer, there are several alternative interpretations of this finding (e.g., greater completeness of the truth-table representation, as well as its lack of explicit directionality). More importantly, several aspects of the data demonstrate an interaction of representation and judgment types. The advantage of truth-table representations was apparent primarily for verification judgments (see Figure 4). Furthermore, the subjects who made verification judgments on the basis of verbal rules initially made much greater use of the help facility to refresh their memories for the rules, suggesting greater difficulty in deriving verification judgments from verbal rules. Finally, the subjects who were instructed using truth tables were much more likely to report using their own verbal rules if they practiced making prediction judgments. Taken together, these results provide strong support for the hypothesis that alternative representations affect both the initial acquisition of and procedures for using causal knowledge. Specifically, the explicitly cause-to-effect form of verbal rules is inappropriate for verification judgments.

Contrary to our initial hypothesis, verification judgments always required more time than predictions, even after substantial practice. This might be attributed to the simple addition of a step (verification) to the procedures used to make prediction judgments. Two aspects of the data suggest, however, that the actual situation is more complex than this. First, judgment type and negation showed an overadditive interaction, such that the additional time required for verification was greater for negated gates. This contrasts with the results found in a superficially similar situation, the comparison of production and verification latencies in simple arithmetic problems. In the arithmetic task, problem difficulty shows an underadditive interaction with judgment type, which has been attributed to interference between the displayed answer and retrieval of the correct answer (Campbell, 1987). Second, the additional time required for verification was greater for the subjects instructed with verbal rules. Although we are not prepared to provide a detailed explanation of these effects, this finding has several implications. First, practice with verification judgments was not sufficient to eliminate the difference between verification and prediction, placing a limit on our suggestion that asymmetries in causal judgment are due to lasting effects of acquisition context. Second, restructuring theories of learning (e.g., Rosenbloom & Newell, 1987), which appear to share our initial prediction that practice would

result in the subjects' developing single-step pattern-recognition strategies for verification, seem to be challenged by this result. Finally, any sequential processing model of these judgments must allow for indirect effects on latency, in which the output of each stage can vary in quality as well as latency, thus slowing succeeding inferences and generating overadditive interactions (McClelland, 1979, p. 317).

Negated rules resulted in longer latencies than affirmative rules, paralleling results from studies of linguistic negation (e.g., Clark & Chase, 1972; Wason & Jones, 1963). Processing theories of linguistic negation have generally explained this effect by postulating additional processing steps (e.g., Carpenter & Just, 1975), and Trabasso, Rollins, and Shaughnessy (1971) applied a similar approach to modeling negation in rule processing. The effect of negation observed here supports Trabasso et al.'s general conclusion that a common mechanism is involved in negation effects in a variety of situations (see also Carlson, 1989). However, the interaction with judgment type described above rules out a simple additive model.

Implications for the Use of Causal Knowledge

The present results demonstrate that knowledge of a causal rule acquired in a prediction context is relatively difficult to apply in a verification context. With more complex or less practiced causal relations, acquisition context might also have lasting effects on the likelihood or confidence judgments collected by Tversky and Kahneman (1980) and others. Consider an example used by Tversky and Kahneman to illustrate the asymmetry of inferences about causal relations. In this example, subjects were asked to judge which event is more likely: (1) that a girl has blue eyes if her mother has blue eyes, or (2) that the mother has blue eyes if her daughter has blue eyes. Although these events are normatively of equal likelihood, far more subjects chose the "predictive" statement (1) than the "diagnostic" statement (2). Like the verification problems in the present study, such examples present an input state, an output state, and a (strongly implied) causal relation—here, genetic inheritance. It seems likely that knowledge of genetic inheritance is acquired in a prediction context, perhaps with the aid of directional verbal rules. Indeed, the fact that we cannot in general know the context in which everyday causal knowledge is acquired suggests the need for caution in postulating a general, directional causal schema that constrains procedures for reasoning about causal knowledge.

It is important to distinguish tasks in which input state, output state, and causal relation are presented simultaneously—as in the logic-gate judgment task and many others—from open-ended diagnosis tasks in which individuals must construct or search complex causal models. The latter situation occurs, for example, when a technician diagnoses the cause of a fault symptom in an electronic circuit. In troubleshooting tasks not reported

here, subjects attempted to solve problems based on networks of logic gates. In these tasks, we could not detect any reliable effects of acquisition context on problem-solving performance. There was enormous variability among subjects in this problem-solving task, emphasizing the need to distinguish between performance on tasks that directly evoke knowledge of particular causal relations (and are therefore appropriate for addressing issues concerning representation and use of particular causal knowledge) and performance on tasks that require much more complex problem-solving processes (and are therefore subject to many sources of variability in addition to the use of causal knowledge).

Acquisition of Procedures for Using Rules

Performance declined dramatically when subjects shifted from practicing single rules in a blocked procedure to practicing with the rules mixed in a randomized procedure. The average proportion of trials on which help was requested increased from less than .02 in the introductory blocks to .44 in the first block of mixed gates, the proportion correct for trials on which help was not used fell from .96 to .90, and the mean response time for trials on which help was not used increased from 1,271 to 3,251 msec. This phenomenon is similar to the contextual interference effect found in verbal learning (e.g., Battig, 1979) and motor learning (e.g., Shea & Morgan, 1979), in which randomized presentation results in poorer acquisition performance but superior retention. Hiew (1977) demonstrated a similar effect in the rule-induction task (Bourne, 1970). Although we did not contrast blocked and random acquisition in terms of retention or transfer, the present results suggest that previous findings on contextual interference might be extended to the domain of procedural skill.

The difficulty of dealing with multiple rules does suggest that learning simple cognitive skills severely taxes working memory. The ease of blocked practice with a single rule may be a result of subjects' maintaining the rule in short-term memory. Alternatively, the rule may be maintained by the context in the form of an episodic memory or rapidly available reminding (e.g., Ross, 1984). When a new rule is used in the same context, however, the previous rule becomes very difficult to recover. This is consistent with a model of working memory proposed by Schneider and Detweiler (1988). They suggest that new procedures can be maintained in working memory by maintaining a set of codes within one context. If a second code is stored in the same context, however, retroactive interference greatly reduces the availability of the first procedure. As the subject practices the task, direct associations between the stimuli (e.g., the logic symbol and the inputs) and the response develop, making it unnecessary to recall the rule in declarative form. This learning, however, is slow in comparison to the context-learning mechanism.

CONCLUSION

The present study demonstrated that the type of judgment required and the type of representation used for instruction during acquisition have substantial effects on the ability to use causal rules, even after extensive practice. The results suggest that the asymmetries observed in causal judgment may result in part from lasting effects of acquisition context, although some asymmetry may be inherent in the requirements of alternative judgment tasks (such as prediction and verification). This research extends previous work on rule learning (e.g., Bourne, 1970) by examining the development of skill in using rules rather than rule induction, and by examining the use of rules that represent causal relations. Understanding how particular pieces or types of knowledge can be used requires consideration of how that knowledge is acquired.

REFERENCES

- BATTIG, W. F. (1979). The flexibility of human memory. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing in human memory* (pp. 23-44). Hillsdale, NJ: Erlbaum.
- BOURNE, L. E., JR. (1970). Knowing and using concepts. *Psychological Review*, *77*, 546-556.
- BRUNER, J. S., GOODNOW, J. J., & AUSTIN, G. A. (1956). *A study of thinking*. New York: Wiley.
- CAMPBELL, J. I. D. (1987). Production, verification, and priming of multiplication facts. *Memory & Cognition*, *15*, 349-364.
- CARLSON, R. A. (1989). Processing nonlinguistic negation. *American Journal of Psychology*, *102*, 211-224.
- CARPENTER, P. A., & JUST, M. A. (1975). Sentence comprehension: A psycholinguistic processing model of verification. *Psychological Review*, *82*, 45-73.
- CHENG, P. W., & HOLYOAK, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, *17*, 391-416.
- CLARK, H. H., & CHASE, W. G. (1972). On the process of comparing sentences against pictures. *Cognitive Psychology*, *3*, 472-517.
- DODD, D. H., KINSMAN, R. A., KLIPP, R. D., & BOURNE, L. E., JR. (1971). Effects of logic pretraining on conceptual rule learning. *Journal of Experimental Psychology*, *88*, 119-122.
- DULANY, D. E. (1974). On the support of cognitive theory in opposition to behavior theory: A methodological problem. In W. B. Weimer & D. S. Palermo (Eds.), *Cognition and the symbolic processes* (pp. 43-56). Hillsdale, NJ: Erlbaum.
- EINHORN, H. J., & HOGARTH, R. M. (1986). Judging probable cause. *Psychological Bulletin*, *99*, 3-19.
- HIEW, C. C. (1977). Sequence effects in rule learning and conceptual generalization. *American Journal of Psychology*, *90*, 207-218.
- JOHNSON-LAIRD, P. N. (1983). *Mental models*. Cambridge, MA: Harvard University Press.
- LEVINE, M. (1975). *A cognitive theory of learning*. Hillsdale, NJ: Erlbaum.
- MCCLELLAND, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*, 287-330.
- ROSENBLUM, P., & NEWELL, A. (1987). Learning by chunking: A production system model of practice. In D. Klahr, P. Langley, & R. Neches (Eds.), *Production system models of learning and development* (pp. 221-286). Cambridge, MA: MIT Press.
- ROSS, B. H. (1984). Reminders and their effects in learning a cognitive skill. *Cognitive Psychology*, *16*, 371-416.
- SCHNEIDER, W., & DETWEILER, M. (1988). A connectionist/control architecture for working memory. In G. H. Bower (Ed.), *The psychol-*

- ogy of learning and motivation* (Vol. 21, pp. 53-119). New York: Academic Press.
- SCHUSTAK, M. W., & STERNBERG, R. J. (1981). Evaluation of evidence in causal inference. *Journal of Experimental Psychology: General*, **110**, 101-120.
- SHEA, J. B., & MORGAN, R. L. (1979). Contextual interference effects on acquisition, retention, and transfer of a motor skill. *Journal of Experimental Psychology: Human Learning & Memory*, **5**, 179-187.
- TRABASSO, R., ROLLINS, H., & SHAUGHNESSY, E. (1971). Storage and verification stages in processing concepts. *Cognitive Psychology*, **2**, 239-289.
- TVERSKY, A., & KAHNEMAN, D. (1980). Causal schemas in judgment under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (Vol. 1, pp. 49-72). Hillsdale, NJ: Erlbaum.
- WASON, P. C., & JONES, S. (1963). Negatives: Denotation and connotation. *British Journal of Psychology*, **54**, 299-307.

(Manuscript received October 19, 1987;
revision accepted for publication September 12, 1988.)