Memory & Cognition
2008, 36 (1), 157-169
doi: 10.3758/MC.36.1.157

Facilitative interactions of model- and experience-based processes: Implications for type and flexibility of representation

Sean M. Lane, Robert C. Mathews, Bill Sallas, and Robert Prattini
Louisiana State University, Baton Rouge, Louisiana

and

Ron Sun
Rensselaer Polytechnic Institute, Troy, New York

People are often taught using a combination of instruction and practice. In prior research, we have distinguished between model-based knowledge (i.e., acquired from explicit instruction) and experience-based knowledge (i.e., acquired from practice), and have argued that the issue of how these types of knowledge (and associated learning processes) interact has been largely neglected. Two experiments explore this issue using a dynamic control task. Results demonstrate the utility of providing model-based knowledge before practice with the task, but more importantly, suggest how this information improves learning. Results also show that learning in this manner can lead to "costs" such as slowed retrieval, and that this knowledge may not always transfer to new task situations as well as experientially acquired knowledge. Our findings also question the assumption that participants always acquire a highly specific "lookup" table representation while learning this task. We provide an alternate view and discuss the implications for theories of learning.

The goal of education and training, whether in the classroom, on the athletic field, or in the conference room, is to raise people from a basic state of knowledge to a more advanced one. Yet, progressing from the condition of a novice to that of an expert on a complex task is typically a long, tedious process (e.g., Ericsson, Krampe, & Tesch-Römer, 1993). Given that typical forms of training include some mixture of instruction and experience, one possible path for optimizing learning involves discovering how best to integrate such activities.

A wide variety of evidence has been used to argue that complex mental skills are learned and deployed through the use of two different and complementary types of processes and resulting representations (e.g., Anderson, 1982; Berry & Dienes, 1993; Mathews et al., 1989; Reber, 1993; Sun, 2002; for arguments against see, e.g., Shanks & St. John, 1994). Although a variety of theoretical terms have been used to describe these two types of processes (e.g., Anderson, 1982; Knowlton & Squire, 1996; Reber, 1967), we refer to these two categories as *experience-based* and *model-based* processing (see Mathews et al., 1989; Sallas, Mathews, Lane, & Sun, in press).[1] In this conception, experience-based knowledge is acquired relatively directly from the environment, and features of family resemblance are abstracted over multiple encounters with category members (e.g., Brooks, 1987; Estes, 1986; Hintzman, 1986; Medin & Schaffer, 1978). This type of processing is fairly error-tolerant (e.g., Sun & Mathews, 2005), and the resulting knowledge is often difficult to articulate (e.g., Lewicki, Czyzewska, & Hoffman, 1987; Reber, 1989). More broadly, important features of stimuli may be learned without the intention of doing so. In contrast, model-based processing involves using a mental model (e.g., Johnson-Laird, 1982) or other type of explicit task representation to guide performance; the resulting knowledge is communicated to others relatively easily. This knowledge can vary from quite abstract representations of a task (e.g., an equation) to more concrete representations of how to accomplish the goal of a task (e.g., a set of instructions or a recipe). When one performs a task using model-based knowledge, the kind of thinking required can place a substantial demand on people's limited capacity memory and attention (e.g., Hayes & Broadbent, 1988), and often leads to slow, but accurate, performance (e.g., Domangue, Mathews, Sun, Roussel, & Guidry, 2004).

Much of the research on this topic has attempted to identify tasks that isolate a specific type of processing, often under the assumption that during learning individuals rely exclusively on one process rather than on another (e.g., Lewicki et al., 1987). However, a growing number of researchers have begun to acknowledge that

S. M. Lane, slane@lsu.edu

both types of processes operate in nearly all tasks (e.g., Mathews, 1997; Reber, 1989; Seger, 1994; Willingham, Nissen, & Bullemer, 1989). Because most tasks involve both processes, one important but insufficiently studied question is how these processes interact (Mathews, 1997; Mathews et al., 1989; Sun & Mathews, 2005). The purpose of the following experiments is to explore this interaction and its implications, particularly for the type and flexibility of knowledge representations which result from these processes. We examine this question in the context of a dynamic control task—a task that can be learned and performed quite well before participants have acquired explicit knowledge of the underlying system (e.g., Berry & Broadbent, 1984, 1988; Stanley, Mathews, Buss, & Kotler-Cope, 1989). In other words, participants often learn the task in a relatively experiential manner. In the following studies, we manipulate whether participants receive model-based knowledge about the task prior to task experience, and we assess such knowledge's influence on performance and the nature of the resulting knowledge representation acquired from training. In the remainder of the introduction, we first briefly describe the dynamic control task and then consider relevant research on the interaction of model- and experience-based knowledge. Subsequently, we discuss the widely held notion that learning in the dynamic control task consists of building up instances of successful interactions with the system (a *lookup table*; Broadbent, Fitzgerald, & Broadbent, 1986; Dienes & Fahey, 1995, 1998; Marescaux, Luc, & Karnas, 1989), and note its implications for the potential interaction of model- and experience-based processes. Finally, we describe our specific hypotheses and the general procedure of our studies.

Dynamic Control Task

In the following experiments, we explore learning processes using the dynamic system control task developed by Berry and Broadbent (1984). In our version of the task, participants control the temperature of a nuclear reactor by varying the number of fuel pellets fed into it. The target temperature was 6,000º. Production is affected by the number of pellets according to the formula $T = 20p - T_{tr-1} + N$. In this equation, T = current reactor temperature in thousands; p = number of pellets in hundreds, T_{tr-1} = reactor's temperature on previous trial; and N = random noise (+1,000, −1,000, or 0 with equal probability). Participants receive immediate and accurate feedback (with small deviations for noise) on each trial, in the form of the resulting production level. During practice, participants complete blocks of 10 trials, and each block begins with a new randomly chosen temperature level. Participants receive extensive practice with the task (multiple sessions).

Although the task is fairly straightforward, it has a number of features which characterize learning in many real-world situations. In particular, we note that (1) it is a dynamic task, in that the existing state of the system keeps changing; (2) noise affects the precision of the feedback received by participants; (3) performance is enhanced by extensive practice (e.g., Stanley et al., 1989); and (4) mental models develop late in learning and are sometimes inaccurate (e.g., Stanley et al., 1989). To the extent that participants are able to articulate their understanding of the task, they never discover or state the underlying formula. Rather, their reports tend to consist of inputs to enter for certain output states.

A brief example makes clear that these features are not uncommon. Imagine that you are trying to improve the performance of employees under your supervision. The effectiveness of a given intervention with an employee will often depend on factors such as the employee's stage of development (a dynamic context; i.e., an intervention will work better at some stages than at others). Further, a given intervention may have slightly different effects, depending on the day on which it is applied (i.e., there is noise in the feedback). Becoming a good manager can sometimes take years, and an individual can become a good manager without necessarily being able to articulate the rules that are the basis of his or her performance. In short, the dynamic control task has a number of features that make it useful for studying the process of learning.

The Interaction of Model- and Experience-Based Processes

In considering the nature of the interaction of model- and experience-based processes, we first argue that—although both processes are typically involved in most tasks—it is possible to manipulate conditions in such a way as to emphasize one type over another (Sun & Mathews, 2005). In the following experiments, we examine one specific type of interaction—namely, the effect of providing partial or full model-based information before training on participants' ability to learn to control a dynamic system (Berry & Broadbent, 1984). We note that participants in prior studies using this task learned to control the system quite well without being given explicit instructions about the underlying nature of the system (e.g., Berry & Broadbent, 1984, 1988; Dienes & Fahey, 1995, 1998) and achieved this level of performance before they were able to articulate a (relatively) accurate description of the system's behavior (Stanley et al., 1989). Thus, we argue that participants can learn to control a dynamic system using primarily experience-based processing. Most previous attempts to study the impact of model-based processing on learning this task have focused on the effect of instructing participants to try to discover the rules governing the system (e.g., Berry & Broadbent, 1988; Sun & Mathews, 2005), and the result has often been performance decrements rather than facilitation.

A different approach was taken by Roussel (1999). Specifically, he attempted to facilitate learning of the dynamic control task by providing a hint (valid inputs for three output states) to some participants. The motivation for this manipulation came from the notion of a "lookup" table (Broadbent et al., 1986; Dienes & Fahey, 1995, 1998; Marescaux et al., 1989). According to this concept, participants learn to control the dynamic control task by constructing a table of successful output–input responses (e.g., "When the temperature is 1,000º, input 400 fuel pellets") while they are learning the task. As learning progresses, participants are able to rely on their memories for previous experiences to make responses to current system states. Further, Dienes and Fahey (1995, 1998) provided arguments and evidence that these representations are acquired implicitly rather than explicitly. Because of this, Roussel (1999) predicted that providing information relevant to a lookup table representation should enhance performance over that of participants who simply practice the task without this information. This model-based knowledge could enhance performance either by "filling in" entries in the table or by enhancing attention to appropriate aspects of the task. Roussel's results confirmed this prediction, since participants who received the "hint" substantially outperformed participants who learned the task experientially; in other words, the provision of model-based knowledge of this nature allowed participants to acquire more valid knowledge about the system.

Although Roussel's (1999) findings demonstrate that providing model-based knowledge can improve performance, they nevertheless leave open a number of questions about the locus of this effect. First, Roussel's design did not definitively rule out the possibility that improved performance in the hint condition was simply a function of participants doing better when they encountered the "hint" states; in other words, the question is whether the hint helped change the manner in which participants learned the task, or simply allowed them to use the knowledge provided to act appropriately in certain specific situations. Second, Roussel's hint provides two different types of information to participants, either one of which could facilitate learning. Specifically, the hint tells participants that they should look for a relationship between prior output states and the input to achieve the goal state, and it informs participants of specific valid inputs for 3 output states (out a total of 12). As will be seen, both of these issues are addressed in Experiment 1. A third broad issue concerns the nature of the knowledge representation that results from receiving model-based knowledge and extensive practice with the task. Although learning in such a manner has a performance benefit, whether or not knowledge acquired under such conditions differs in any way from knowledge acquired only from experience is a question worth asking. For instance, it has been argued that model-based (explicit) knowledge can be transferred to new situations much more readily than experience-based (implicit) knowledge can (e.g., review by Dienes & Berry, 1997; but see Willingham, 1997). One implication of this claim is that explicitly provided lookup table information may transfer more effectively (e.g., because participants can adjust table values) than can similar information acquired only through experience (e.g., because it is not easily accessible to awareness for modification). Finally, there is research suggesting that provided model-based knowledge may lead to improved accuracy in relatively implicit tasks but that such knowledge often has a cost in terms of speed of retrieval (e.g., Domangue et al., 2004; but see Sallas et al., in press). Both Experiments 1 and 2 will address these issues by providing data on the nature of the acquired knowledge representation.

Current Experiments

In the following experiments, we had participants learn the dynamic control task over a series of sessions (i.e., they received extensive practice). We varied the type of information participants received prior to and during training. In each experiment, an experiential condition served as a baseline. These participants were given task instructions and the goal of maintaining the system at a given level (6,000º), but were given no additional information about the task. Participants of primary interest were provided information about the correct response to make in a given system state (i.e., the proper input to achieve the task goal) for either a subset (Experiment 1; hint+quiz condition) or all system states (Experiment 2; table condition) before they began training. Acquired knowledge was assessed using computerized performance tests, in which participants attempted to attain a goal state, and, in Experiment 1, a final paper-based "table" test in which participants were questioned about their knowledge of valid inputs for each output state.

Experiment 1 was designed to test between alternative explanations of the effectiveness of providing a task hint to participants and to establish whether the hint has a narrow or a broad impact on learning the dynamic control task. There were three conditions. Experiential (control) participants were simply told to achieve and maintain the system at the goal state during training. A second group (quiz-only) was given the same instructions and, during every two blocks of training, was also quizzed about their knowledge of the appropriate inputs to make for each output state. The goal of the quiz was to emphasize to participants that they should consider the relation between prior output states and inputs when trying to learn how to attain the goal state (note that the quiz also alerts participants that there is a single correct answer for each output state). A third group (hint+quiz) received, in addition to task instructions and quizzes, partial lookup table information in the form of three valid output–input pairs (e.g., "If temperature is 1,000º, then input 400 fuel pellets") before training. This design allowed us to test several hypotheses. First, if providing specific examples to participants leads them simply to apply those "rules," any advantage of the hint+quiz condition in overall performance over the remaining groups should disappear when only "nonhint" states are considered. However, if the specific examples lead participants to more effectively learn the task as a whole, any advantage should remain. Second, if prior findings of the effectiveness of the hint (Roussel, 1999) are simply a function of focusing participants on the importance of the relation between prior output states and inputs, performance in both quiz-only and hint+quiz groups should exceed that of experiential participants and should not differ from each other. If the provided examples are playing a critical role, hint+quiz participants' performance should exceed that of the quiz-only condition. Finally, we examined whether there is any "cost" to provided model-based knowledge by evaluating response times (RTs) during the performance test. On the basis of prior research (e.g., Domangue et al., 2004), we predicted hint+quiz participants would be slower than participants who learned the task experientially would be.

To anticipate: The results of Experiment 1 suggest the importance of provided model-based knowledge. In Experiment 2, we utilized a group of "table" participants who received full lookup table information (all correct output–input pairs) before training. We examined the nature of knowledge acquired from practice by assessing the flexibility of such knowledge using a transfer test. We discuss these issues in more detail in the introduction to that experiment.

EXPERIMENT 1

Method

Participants and Design. Eighty-five undergraduate students enrolled in introductory psychology courses at Louisiana State University voluntarily participated in return for extra credit. There were three training conditions: experiential (control); quiz-only; and hint+quiz. Although participants were randomly assigned to conditions, some participants did not attend all three sessions, and their data was excluded. Altogether, there were 29 participants in the experiential condition, 31 participants in the quiz-only condition, and 25 in the hint+quiz condition.

Task. The reactor control task used in the present work is a computer-based task in which participants imagine they are in charge of a nuclear reactor. The participants attempt to achieve and maintain a specified level of an output variable—reactor temperature—by controlling the number of fuel pellets consumed by the reactor. Participants were given the goal of maintaining production at 6,000º. Task trials were grouped into blocks of 10 trials and each block began with a randomly selected temperature level. On each task trial, the computer presented the current temperature. Participants saw a display that depicted temperature and the number of fuel pellets input. Reactor temperature was allowed to vary from 1,000º to 12,000º. Participants could select a number of fuel pellets ranging from 100 to 1,200 in multiples of 100. The participants responded by choosing and entering the number of pellets to be fed into the reactor; the computer then updated and displayed the new temperature level. As noted in the introduction, the formula determining the output includes a random noise element (+ or −1,000º). Thus, 1/3 of the time a given input would result in a corresponding output (e.g., 6,000º), 1/3 of the time the output would be 1,000º higher, and 1/3 of the time it would be 1,000º lower. The task was constructed in such a way that if the input resulted in an output of less than 1,000º or more than 12,000º, it was set to 1,000º and 12,000º, respectively. Participants were made aware of the lower and upper limits of production levels. At the end of each block (i.e., every 10 trials), the display was cleared and a new graph displayed for the next block of trials. The main dependent measure was the mean unsigned deviation from target production, in degrees. Because the target production level was always 6,000º, the dependent measure could vary from a minimum of zero to a maximum of 6,000º. RT was also recorded.

Each session consisted of a study phase and a test phase. The duration of the study phase was 20 min, and participants worked at their own pace. At the end of every second block (20 trials) of the study phase, participants in the quiz-only and hint+quiz conditions were given a quiz shown in a pop-up window that completely occluded the visual display of the primary task. For each of the 12 positions on the temperature scale (1,000º–12,000º in increments of 1,000º), participants were asked, "If the reactor's temperature is (insert temperature)º, how many fuel pellets should you enter to move the temperature to 6,000º?" The presentation order of an item (output state) within a quiz was chosen randomly without replacement. Participants selected a response from the same levels of fuel pellets used in the main task (100–1,200 in intervals of 100), and no feedback was given. The quiz was not present during the test phase.

Participants in the hint+quiz condition also received a printed hint before the study phase. They were told that the hint would help them learn to control the reactor's temperature. The hint, which was taken away at the beginning of the test phase for each session, was as follows:

If temperature is 1,000º, then input 400 fuel pellets.
If temperature is 4,000º, then input 500 fuel pellets.
If temperature is 7,000º, then input 600 fuel pellets.

Procedure. Participants were tested in groups ranging from three to eight. Each group was randomly assigned to one of the three conditions. Regardless of condition, all participants completed three sessions, one per day with one day between each session (i.e., Monday, Wednesday, and Friday of the same week). Each session consisted of a 20-min training phase and a test phase.

Participants were instructed to take on the role of manager of a nuclear reactor, and were told that this job entailed learning how to achieve and maintain a target temperature level (6,000º) by interacting with the simulation. They were further informed that the only variable they could control was the number of fuel pellets entered into the reactor. Thus, their task was to learn the relationship between number of fuel pellets and reactor temperature. After receiving instructions, participants were given 20 min to practice the task. As described above, participants in the quiz-only and hint+quiz conditions were quizzed after every other block. Participants in the hint+quiz condition had the hint available to them during the study phase.

Each day, after practicing for 20 min, all participants completed the performance test, which consisted of 50 blocks of 10 trials of the reactor control task. The participants were informed that their goal was to maintain a temperature level of 6,000º and they were excused from the session after they had all finished the test phase during the first two sessions. After completing the test phase of the third session, participants completed a table test (a paper version of the same questions asked during practice in the quiz groups), which required them to provide the correct number of fuel pellets to use for each output state (from 1,000º–12,000º in increments of 1,000º) to achieve a temperature of 6,000º.

Results and Discussion

All analyses were conducted with an α level of .05. The primary dependent variable of interest was performance as indicated by the average absolute deviation from target goal; the lower the score, therefore, the better the performance. Means and standard errors (*SE*s), in the form of error bars, for this measure are depicted in Figure 1. We report means and *SE*s for RT and table test measures in the text below.

The effect of type of training on task performance. We examined test performance across the three sessions using a 3 × 3 mixed factorial ANOVA with training type (experiential, quiz-only, and hint+quiz) as a between-subjects variable and session (1, 2, 3) as a within-subjects variable. Performance improved across sessions [M = 1,525, 1,120, and 946; $F(2,162) = 43.1$, MS_e = 7,502,055, $\eta_p^2 = .347$], and subsequent comparisons revealed that performance on each session was better than on the previous session. Performance also differed significantly as a function of training type [$F(2,81) = 219.0$, MS_e = 369,929,625, $\eta_p^2 = .730$]. Subsequent comparisons revealed that participants in the hint+quiz condition performed more accurately than those inthe quiz-only and experiential conditions, which did not significantly differ

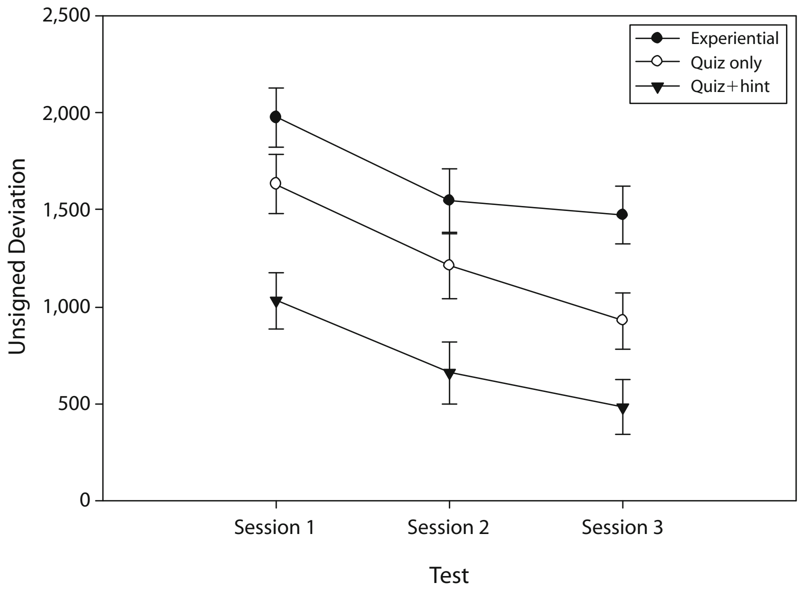

Figure 1. Test performance in Experiment 1 by training condition. The dependent measure is the absolute deviation from target, and lower scores indicate greater accuracy. Error bars are in standard error units.

from each other ($M = 725, 1{,}257, 1{,}665$). The interaction between training group and session was not statistically significant ($F < 1$). Thus, facilitation from the provided hint appears to be primarily a function of receiving valid output–input knowledge, rather than of simply increasing awareness of output–input relationships. However, such model-based knowledge could simply help performance in ways that are narrowly limited to specific provided examples. Thus, we also examined performance restricted to "nonhint" output states. A parallel ANOVA to the one described above was conducted and revealed that, even with these trials removed, the hint+quiz condition still showed significantly better performance than did the quiz-only and experiential conditions [$M = 742, 1{,}426, 1{,}589$; $F(2{,}81) = 8.6$, $MS_e = 18{,}163{,}413$, $\eta_p^2 = .172$]. Thus, our results suggest that the facilitation observed in the hint+quiz condition is relatively broad, because participants' performance improved even for system states not specified in the hint. It appears that the provided information allowed these participants to infer the remaining states and thus to reduce the problem space to be searched (Newell & Simon, 1972).

We next examined the speed of participants' responses at test using a 3×3 ANOVA as described above, with median RT as our dependent measure. Our results revealed only a main effect of training condition [$F(2{,}81) = 12.9$, $MS_e = 3{,}281{,}285$, $\eta_p^2 = .236$], with the hint+quiz and quiz-only participants significantly slower than participants in the experiential conditions ($M = 1{,}004, 864, 625$ msec; $SE = 42, 62$, and 60, respectively).[2] Thus, consistent with prior research (e.g., Domangue et al., 2004), the superior accuracy of the hint+quiz condition appeared to have a cost.

The effect of training on the table test. Finally, we measured participants' task knowledge on the paper-based table test, which was completed by all participants at the end of the third session. As with the performance test, the dependent measure was the average absolute deviation from target goal. A one-way ANOVA with training type (hint+quiz, quiz-only, and experiential) as a between-subjects variable revealed significant differences between groups [$M = 1{,}145, 1{,}917$, and $2{,}663$, $SE = 169, 259$, and 288, respectively; $F(2{,}81) = 10.6$, $MS_e = 16{,}026{,}743$, $\eta_p^2 = .986$]. Results of subsequent comparisons showed that the hint+quiz condition performed significantly better than the quiz-only condition, which performed significantly better than the experiential condition. However, when the "hint" states were excluded from the analysis, the hint+quiz and quiz-only groups did not differ significantly, although both groups were significantly more accurate than the experiential condition [$M = 1{,}344$, $1{,}859$, and $2{,}987$, $SE = 221, 272$, and 311, respectively; $F(2{,}82) = 9.951$, $MS_e = 19{,}093{,}185$, $\eta_p^2 = .201$]. Thus, the participants' receiving the quizzes regularly during training did improve the accuracy of their acquired model-based knowledge of the task compared with participants who simply practiced the task. However, receiving the example states did not confer any significant additional advantage beyond those of the specific examples (although hint+quiz participants were nominally better, even on nonhint states).

To return to our original hypotheses, our results suggest that providing model-based knowledge about valid inputs for a subset of output states improved performance by changing how participants learned from training rather than by simply teaching them to respond appropriately when they encountered the specific states described in the hint. Hint+quiz participants' performance was superior to those in other conditions on the performance test, even when only "nonhint" states were considered. The effects of the "hint," therefore, were broad rather than narrow, in that the provided knowledge allowed them to infer the appropriate inputs for the nonhint states. Second, facilitation resulting from the task hint appears to come primarily from the provided model-based information (and its resulting effects on learning), rather than from highlighting the relation between prior output states and inputs, or highlighting that there is a single correct answer for each output state. This was demonstrated by the clear superiority on the performance test of hint+quiz over the quiz-only participants. Third, even with extensive practice on the task, there was a "cost" to utilizing model-based knowledge to improve performance. Although hint+quiz participants improved their accuracy, they were slower to respond.

EXPERIMENT 2

The results of Experiment 1 reveal that provided model-based knowledge in the form of a subset of valid output–input states improves task performance in a nontrivial manner, since participants showed superior performance even for states not directly specified in the hint (although there was a cost in speed of response). In Experiment 2, we sought to extend the findings of Experiment 1. In particular, we were interested in potential differences in the flexibility of knowledge acquired experientially rather than from provided (model-based) task information. To do so, we more strongly manipulated reliance on these two types of knowledge by having a group of participants learn a full, accurate lookup table prior to training on the task. Specifically, these table condition participants learned to criterion the appropriate inputs for achieving the goal state from each of the 12 output states (1,000º–12,000º). This could be conceived as a "super-hint," in that it provides an optimal (and highly specific) guide to performing the task. Thus, these participants are likely to rely fairly heavily on provided model-based knowledge, particularly compared with participants who simply practice the task (experiential condition). More broadly, this manipulation allows us to ask about the nature of knowledge acquired from practice with the task.

In Experiment 2, participants in the table and experiential conditions completed a series of tests. During the first and second sessions, participants completed a test that used the same goal as during training (6,000º; standard test). Obviously, one would predict that participants in the table condition would do much better on these tests. More importantly, participants completed two other tests during the second session. The first test, using the same goal, assessed participants' ability to control the system under time pressure (speeded test). If the "cost" of utilizing provided model-based knowledge observed in participants from Experiment 1 is a general one (see Domangue et al., 2004), one would predict it would be more difficult to access model-based knowledge to perform the task. Thus, the prediction would be that table participants would be more impaired on the speeded test (relative to the unspeeded test) than would experiential participants. However, an alternative hypothesis is that the "cost" in speed incurred by hint+quiz participants in Experiment 1 resulted from not having the opportunity to practice deploying their model-based knowledge (because they needed to acquire knowledge of "nonhint" states during practice). In this case, one might predict that the table participants would be better able to convert their model-based knowledge into procedural (i.e., experience-based) knowledge and thus should show little impairment on the speeded test (e.g., Anderson, 1982).

As a means of assessing knowledge flexibility, the last and most important test required participants to maintain a different goal (8,000º; the transfer test) than the one they had during training. We note that competing theoretical assumptions lead to different predictions about performance on the transfer test. As noted in the introduction, the issue of flexibility of representation has been much discussed with respect to implicit and explicit learning. Specifically, it has been assumed that explicit (model-based) knowledge transfers more easily than implicit (experience-based) knowledge (e.g., Dienes & Berry, 1997; but see Willingham, 1997). One implication of this claim is that explicitly provided lookup table information (table condition) would be predicted to transfer more effectively than would similar information acquired through experience (experiential condition). This seems a reasonable possibility, since knowledge of the lookup table for the old goal could be transferred to the new goal simply by adding a constant to each entry.

A similar prediction would be made by theories of the dynamic control task that assume that experiential participants learn a goal-specific implicit lookup table. These theories (e.g., Cleereman's, described in Marescaux et al., 1989; Dienes & Fahey, 1995) assume that participants acquire specific lookup table entries through interaction with the task. When participants successfully achieve the goal from a given output state, they encode this information into memory and later rely on the entries to make responses when they face previously encountered system states. When participants face situations not previously encountered (as would be the case on a transfer test), responses will be made with some baseline probability (in Cleereman's theory, these responses are equiprobable). Thus, these theories suggest that learning in the dynamic control task acquired with one set of parameters should not transfer well when those parameters are changed. More specifically, theories which posit specific implicit lookup table entries (e.g., Dienes & Fahey, 1995; Marescaux et al., 1989) would predict poor transfer for the experiential condition, and would be agnostic on the outcome of the table condition.

Finally, this prediction about lack of transfer only comes from theories which claim that specific output–input pairs are represented in the lookup table. Dienes and Fahey (1995, 1998) note that, although their results are consistent with specific lookup representations, such specificity may not characterize the knowledge of participants with extensive task experience. Given that our participants receive considerably more training than do participants in typical dynamic control task studies, this is a clear possibility: Experiential participants in Experiment 1, for example, completed an average of 3,410 study trials across three sessions, versus 40–80 trials for participants in typical studies (Dienes & Fahey, 1995). In principle, representations in a lookup table can also define task situations in a more abstract way such that they approximate general rules ("when output is high, do X"). Theories of learning that allow for more abstract representations (e.g., Mathews et al., 1989; Reber, 1993) would predict transfer when task parameters are changed, to the extent that the general rules are still applicable. If acquired lookup table representations are more general than provided lookup-table (model-based) representations, this would predict better transfer in the experiential condition and poorer transfer in the table condition (relative to their performance on the same goal test).

Having three control groups allowed us to more precisely assess the effects of our experimental conditions. One control group had no training, but simply took the standard test (no training–standard); another had no training and took the transfer test (no training–new target); and a third first memorized the table, then took the transfer test (they were told the table applied to a goal of 6,000º and that the test goal would be 8,000º; table control–new target). The two no-training conditions provide a strong baseline for assessing the impact of task practice on test performance. The table control condition allows for an assessment of how much model-based knowledge can transfer without experience in the task.

Method

Participants. Participants were undergraduate introductory psychology students at Louisiana State University who voluntarily participated for extra credit. A total of 127 participants were randomly assigned to conditions. Altogether, there were 21 participants each in the experiential and table conditions. We also had 30 participants in the no training–new target test condition; 30 participants in the no training–standard test condition; and 25 participants in the table control–new target test condition.

Materials and Procedure. In this experiment, we used the same nuclear reactor cover task. All participants in the experimental conditions completed two 1-h sessions (spaced 48 h apart).

First session. In the first session, all participants were told about the simulated nuclear reactor task and were informed that their goal was to discover how to achieve and maintain a target level of temperature (6000º) by interacting with the simulation. Participants were told that they were in control of only a single input variable (i.e., the number of fuel pellets).

Participants in the experiential condition began performing the task immediately after reading its description and did not receive any further instruction. Participants in the table condition were provided with a table listing all the potential beginning states (previous output) and the correct corresponding inputs, and were instructed to study the table for 2 min for a recall test. After the 2-min study period, table participants were given two recall tests. Each test consisted of all the possible output states of the nuclear reactor and a space to indicate the correct input responses. The two versions of the recall test contained the same information (all possible outputs), but in two different orders. Participants were required to get all 12 input responses correct in both orders before beginning the practice phase. If they failed to meet criterion, they had to restudy the table and were retested until they did.

Following instructions, participants had 20 min of practice time using a self-paced version of the task. Participants next completed a self-paced test, consisting of 50 blocks of 10 trials each (i.e., 500 trials). As in the practice phase, they were given the goal of maintaining the system at 6,000º.

Second session. At the beginning of the second session, participants received a copy of their group's instructions (table participants were given the table to review before practicing), and completed 5 min of self-paced practice. Following the practice period, participants completed another self-paced test (6,000º goal) consisting of 50 blocks. We refer to these tests in Sessions 1 and 2 as *standard tests*, in that both were self-paced and used the same target level as used in practice.[3] We also tested a control condition (no training–standard) which received no training but completed the 50-block standard test with a goal of 6,000º. This baseline condition allowed for a comparison with experimental conditions to determine performance improvement from training.

After completing the standard test, participants completed the remaining tests in a single, standard order. First, they completed a fast-paced version of the standard test that we term the *speeded test*. Participants in this version had to enter their input for each trial within 1.5 sec (1,500 msec). If they failed to make a response within the time limit, the computer entered a random value, which typically made the next trial more difficult. Thus, the test reinforced quick responses. As in the standard test, participants completed 50 blocks. Finally, participants received the transfer test. On this *new target goal* test, participants were told to maintain temperature at 8,000º instead of at 6,000º. All other aspects matched the standard test (i.e., 50 blocks, self-paced).

Besides the experimental conditions described above, we had two control conditions for the new target goal test. In the no training–new target condition, participants studied task instructions and began the test. In the second control condition (table control), participants exactly replicated the Session 1 procedure for the table condition (i.e., they memorized the table and were required to recall it twice); following this phase, they were immediately given the new goal transfer test without any practice. They were informed that the temperature goal of the new test would be 8,000º, and that they had just learned the table which provided correct answers for a goal of 6,000º. The performance of this control group is useful when compared against the table condition, because it allows us to separate transfer effects resulting from practice from effects resulting from good knowledge of the table. Note that knowledge of the lookup table could be transferred to the new goal simply by adding a constant to each entry.

Results and Discussion

As in previous experiments, the α level was set to .05 and the primary dependent measure was the absolute deviation from target. Means for the conditions in each of the tests are portrayed graphically in Figure 2.

Standard Tests. We first examined the effect of training condition on test performance using a 2×2 mixed factorial ANOVA with training condition as a between-subjects factor and session (standard tests 1 and 2) as a within-subjects factor. The analysis revealed there was a significant main effect of type of training [$F(1,40) = 45.7$, $MS_e = 785{,}157$, $\eta_p^2 = .53$]. As is clear from the figure, learning a complete lookup table before practice led to dramatically better test performance than did simply practicing the task. Although the main effect of session did not reach statistical significance [$F(1,40) = 3.34$, $MS_e = 83{,}273$, $p < .08$, $\eta_p^2 = .08$], the session $\times$ training interaction was significant [$F(1,40) = 8.25$, $MS_e = 83{,}273$, $\eta_p^2 =$

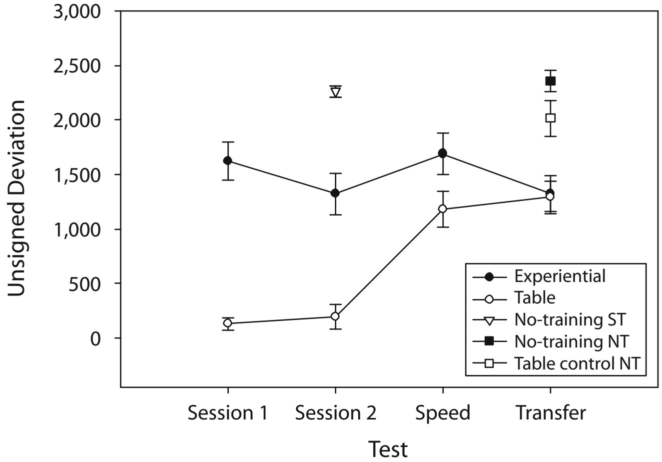

Figure 2. Test performance in Experiment 2 by training condition and type of test. The dependent measure is the absolute deviation from target, and lower scores indicate greater accuracy. Error bars are in standard error units. Session 1 refers to the standard test (same goal as during training) administered during the first session, Session 2 refers to the standard test administered during the second session, Speed refers to the speeded test given during the second session which utilized the same target goal as the standard test, and New Target refers to the transfer test during the second session which had a different target goal. Apart from the experiential and table conditions, we had control participants who did not practice the task before taking the standard test (no-training ST), participants who did not practice the task before taking the new goal test (no-training NT) and participants who memorized the table of correct inputs for the standard test before taking the new goal test (table control NT).

.17]. This reflected the fact that the performance in the experiential condition improved from Session 1 to Session 2 [$F(1,20) = 8.1$, $MS_e = 114{,}095$, $\eta_p^2 = .29$], but the table condition did not ($F < 1$).

Using a Dunnett's test, we next compared performance of the experimental conditions on the second session test with the no-training control (standard) condition. Results revealed that both the table and experiential conditions were significantly more accurate than the no-training control ($p < .01$). In short, both table and experiential groups learned from practice on the task, but the additional explicit task information provided to table participants allowed them to reach very high levels of performance.

Speeded test. We next examined performance in a test environment which was designed to restrict the amount of time participants had available to recall table entries and thus potentially affect reliance on model-based knowledge. Using a one-way ANOVA, we assessed performance on the speeded test. The effect of training condition was just shy of statistical significance (see Figure 1 [$F(1,40) = 4.0$, $p = .051$, $MS_e = 662{,}575$, $\eta_p^2 = .09$]. Although table participants had nominally better performance than experiential participants on this test, a comparison of performance across the standard and speeded tests also suggests that this group was disproportionately impaired by the speeded conditions. We sought further evidence of this effect by performing a 2 × 2 mixed factorial ANOVA with training type (table and experiential) as a between-subjects factor and type of test (standard test and speeded test from Session 2) as a within-subjects factor. Of greatest relevance is the obtained group × test type interaction [$F(1,40) = 14.4$, $MS_e = 140{,}593$, $\eta_p^2 = .27$]. Simple effects analyses revealed that, although both groups were significantly less accurate on the speeded than the standard test, this drop in performance was greater for table participants [$F(1,20) = 62.24$, $MS_e = 164{,}524$, $\eta_p^2 = .76$] than on for experiential participants [$F(1,20) = 12.1$, $MS_e = 116{,}661$, $\eta_p^2 = .38$]. This suggests that even with the opportunity to repeatedly practice deploying model-based knowledge, table participants were much less accurate when they had restricted time to use their knowledge. Thus, the results are consistent with Experiment 1 and other research (e.g., Domangue et al., 2004) in suggesting that model-based knowledge is often highly accurate but may require more time to apply. However, we note that the speeded test also had (less pronounced) negative effects on performance in the experiential condition. We believe this suggests either that these participants may acquire some model-based knowledge during training or that, even though experientially acquired knowledge is generally retrieved more quickly (e.g., Domangue et al., 2004), our test was fast enough to interfere with retrieval on some attempts.

New target goal test–transfer. Performance on the transfer test with a new target level (8,000º) was analyzed with a one-way ANOVA. As is clear from Figure 2, the effect of type of training was not significant ($F < 1$). This finding suggests that a similar amount of generalizable knowledge was learned by participants in the two conditions. To address the issue of transfer, we first compared the groups with the no training–new target condition, using a Dunnett's test. Results revealed that performance in each of the conditions was significantly better than in the control group ($p < .01$). Thus, we see transfer to a new target level in both conditions.

Although both conditions benefitted from practice, what is performance on the new goal relative to the same goal? We analyzed performance using a 2 × 2 mixed factorial ANOVA with type of test (Session 2 standard test vs. new goal test) as a within-subjects variable, and type of training (table vs. experiential) as a between-subjects variable. The finding of greatest interest was a significant interaction [$F(1,40) = 39.31$, $MS_e = 160{,}114.88$, $\eta_p^2 = .50$]. As is obvious from Figure 2, experiential participants showed nearly the same performance on the transfer test as on the standard test ($M = 1{,}324$ and $1{,}325$, respectively, $F < 1$), whereas table participants' performance was significantly worse on the transfer test [$M = 1{,}293$ and 198; $F(1,20) = 63.67$, $MS_e = 197{,}526.44$, $\eta_p^2 = .76$] than on the standard test.

These results have important implications for clarifying what type of knowledge transfers in the table condition. As we have seen, the advantage of the table condition over the experiential condition on the standard test is eliminated in the transfer test. However, on the transfer test, performance in the table condition does not get worse than in the experiential condition, but rather returns to the same level. This suggests two things. First, table condition participants appear to also acquire experience-based knowledge through training. When test conditions make it difficult to use the provided model-based knowledge, they appear to rely more on experience-based knowledge. Second, when considering transfer test performance, it appears that participants' experientially acquired knowledge transferred, whereas their explicit (model-based) lookup table knowledge did not. This possibility can be directly addressed by comparing the performance of participants who learned the table and practiced the task with the performance of participants who learned the table but did not practice the task before taking the transfer test (table control–new goal). We made this comparison in the context of comparing performance in this control condition with that in the experiential and table conditions. A Dunnett's test revealed that performance in both conditions was significantly more accurate than in the table control condition ($p < .01$).[4] Thus, table participants do show transfer to a new goal, but the knowledge used for this transfer appears to be largely acquired by experience rather than from the explicit lookup table knowledge provided beforehand.

To return to our original questions: The results of Experiment 2 revealed that, although provided model-based knowledge greatly facilitated performance on a test that mirrored training conditions, there were "costs." Although a speeded test did not eliminate the advantage of table over experiential participants, their performance was disproportionately affected. More strikingly, the advantage of table participants was completely eliminated in a transfer test which required them to achieve and maintain a new goal, even though they could have performed successfully simply by modifying their table entries by adding a constant. Importantly, experiential participants showed evidence of transfer to the new goal, as their performance was nearly identical in the standard and new goal tests. The fact that table and experiential participants performed similarly on the transfer test suggests that both groups acquired experience-based knowledge from practicing the task, and that in the case of the table condition, this knowledge transferred more easily than did provided model-based knowledge. Finally, the pattern of findings on the transfer test appears most consistent with the notion that typical (i.e., experiential) participants acquire general rather than specific lookup table representations with extensive practice.

GENERAL DISCUSSION

Two experiments confirm the facilitative effects of provided model-based knowledge on learning to control a dynamic system (Roussel, 1999). More importantly, our results clarify the impact of providing such knowledge and the nature of the resulting knowledge representation. In Experiment 1, providing accurate task knowledge (three correct output–input pairs) to participants before practice increased those participants' accuracy at test, compared with those who simply practiced the task or those who practiced and were repeatedly quizzed about the appropriate inputs for each output state. Results revealed that the impact of the provided model-based knowledge was broad rather than narrow, since performance in this condition was superior even when hint states were excluded. However, hint+quiz participants' speed of responding at test was significantly slower than that of participants in the experiential condition. Experiment 2 showed the dramatic impact of providing a complete and accurate lookup table to participants before practice. Although this model-based knowledge led to very high levels of accuracy when tested using an unspeeded test, performance was reduced in a speeded test. In addition, Experiment 2 provided evidence of transfer of experience-based knowledge, as experiential participants performed similarly on standard goal and new goal tests. The results also suggest that experience-based knowledge transferred better to a new goal than did provided lookup table (model-based) knowledge. We next discuss the implications of these findings for a number of important issues.

Effects of Providing Model-Based Knowledge

In both experiments, participants who received model-based knowledge before training were demonstrably better at controlling the system than were participants who relied more heavily on experience-based knowledge (i.e., those participants who simply practiced the task), when the task parameters mirrored those encountered in training. The results of Experiment 1 showed that such improvement can be nontrivial, because performance was improved even for states for which answers were not explicitly provided. Thus, receiving task information appears to change how participants learn from training rather than simply allow them to deploy the knowledge for those specific states. These results also suggest that the utility of the "hint" resided primarily in the provided information, since participants repeatedly quizzed about the table entries did not do significantly better than participants who only practiced the task. Most likely, the provided information allowed hint+quiz participants to infer the remaining states. Put another way, the hint allowed participants to reduce the problem space to be searched (Newell & Simon, 1972).

Despite the advantages of model-based processing for learning, there do appear to be costs. In Experiment 1, hint+quiz participants' test performance was more accurate than those in other conditions, but they were also the slowest at making responses (see also Domangue et al., 2004). One likely possibility is that participants were deliberately attempting to retrieve correct table entries (or exemplar states from training) from memory, whereas those who simply practiced did not. There is also evidence for slowing in Experiment 2. Although table participants learned the table to criterion before training (and thus had the opportunity to practice retrieving these entries), their performance was substantially impaired when the test required speeded responses (however, their accuracy remained superior to experiential participants).

A second intriguing cost concerned the ability of participants to transfer their knowledge to a new goal state. Experiment 2 provided clear evidence of transfer, since participants who learned the task experientially performed similarly under old and new task goals. Although table participants' performance on the new task goal was significantly poorer than on the old task goal, their performance was similar to that of participants who only practiced the task. In other words, experience-based knowledge acquired during practice appeared to show *greater* flexibility than did provided model-based knowledge of a full lookup table. Thus, our results also call into question the notion that experience-based representations are always less flexible than model-based ones are (e.g., Berry & Dienes, 1997). One might argue that this particular result is not surprising, since the model-based knowledge we provided participants was highly specific to a particular goal state; but this is exactly the point, since we do not claim that experience-based knowledge will always be more flexible than model-based knowledge, or that model-based knowledge will not transfer (see, e.g., Kieras & Bovair, 1984). Rather, we argue that both types of knowledge can show flexibility or inflexibility depending on the relationship between learning and testing conditions (see also Kolers & Roediger, 1984; Willingham, 1997). In most studies using "implicit learning" paradigms (e.g., artificial grammar, dynamic control tasks), participants have relatively few training trials and the variability of exemplars within a category is low. Because of this, such training can lead to hyperspecific knowledge that shows low levels of transfer (Mathews, 1997). Our training procedures (see also Stanley et al., 1989) included substantially more practice, and our version of the dynamic control task began each block of trials with a random current output value, forcing participants to experience all of the potential problem space. Given our results, we believe the flexibility of experientially acquired knowledge is dependent on the range of variability inherent in training, and more broadly, on the match between training and performance conditions (Mathews, 1997; Willingham, 1997).

The transfer test results also have interesting implications for the deployment of model- and experience-based processes during learning. The table condition presents a situation in which participants could rely exclusively on model-based knowledge to perform the dynamic control task. After all, these participants were explicitly provided with all the information they needed to do the task correctly. Even in this extreme case, they nevertheless appeared to have acquired experience-based knowledge during training. Specifically, when their ability to use model-based knowledge to perform the task was disrupted by having to maintain a new goal at test, performance in the table condition returned to the level of experiential participants. Further, on the new goal test, participants' performance in the table condition was significantly higher than was that of control participants who received the table without additional training. Thus, even in situations where we might expect to see heavy reliance on model-based knowledge, participants appear to use both model- and experience-based processes during training.

Theoretical Implications

Our results have clear implications for the hypothesis that participants normally learn the dynamic control task by acquiring a lookup table (e.g., Broadbent et al., 1986). Although our results are consistent with the notion of an implicit lookup table representation, they are not consistent with the idea that entries in this table are always highly specific representations of successful output–input pairs (e.g., Dienes & Fahey, 1995; Marescaux et al., 1989). Rather, in line with Dienes and Fahey's (1995, 1998) speculation, the representations acquired during extensive practice with the dynamic control task do not appear to be as specific as do those observed with less practice.

We believe two findings are particularly difficult to reconcile with theories that argue that participants rely on highly specific representations of previously successful outcomes (i.e., achieving the goal state) to make task responses (e.g., Dienes & Fahey, 1995; Marescaux et al., 1989). First, we find transfer of learning to a new system goal. If experiential participants were simply relying on specific table entries to guide performance, one would expect that performance would be poor when the system goal was changed because the table entries change. In contrast to this prediction, performance in Experiment 2 was quite similar in the experiential condition for both the standard and new goal tests, showing substantial transfer. Second, participants who had extensive practice with the task nevertheless showed poor knowledge of the appropriate inputs to make in a given situation. Performance of experiential participants on the table test in Experiment 1 (M = 2,663 deviation units) was substantially poorer than their task performance on a test taken minutes beforehand (M = 1,497 deviation units). This finding stands in contrast to prior research (e.g., Dienes & Fahey, 1995) that found good memory for previously successful outcomes, compared with situations in which participants had not experienced success during a short training period. Further, if one assumes that participants rely on specific lookup table entries, one should see a very strong correlation between "table" test performance and task performance. We computed this correlation for experiential participants in Experiment 1 and obtained [$r(28) = .44, p < .05$]. The r^2 of .20 indicates that output–input knowledge can only account for a small proportion of the variance associated with task performance. Here we must emphasize that the nature of our procedure provides extensive experience with system states relative to most previous work (e.g., 80 trials, Experiment 1; Dienes & Fahey, 1995). In our procedure, participants not only practice for many trials (e.g., in the experiential condition of Experiment 1, $M = 3,410$ trials), but in each block (i.e., every 10 trials) they receive a new random starting output (1,000–12,000º). In the majority of prior research (e.g., Berry & Broadbent, 1984), participants begin from the same starting point on each block (for an exception, see McGeorge & Burton, 1989). Thus, an average experiential participant in our Experiment 1 would have seen each potential output state on their first trial a minimum of 28 times during practice (341 blocks/12 potential states) and likely many more times during the remaining 9 trials they encountered in each block. Further, these participants had experienced success in achieving the goal from most or all the output states before taking the table test.[5] Under these conditions, these theories would predict very good memory for specific output–input situations. The fact that we did not obtain this result would seem to require at least a consideration of alternatives.

One alternative view suggests that participants do not need to store instances in memory (e.g., output–input situations) to learn the dynamic control task, but instead to use successful outcomes to "tune" a small set of strategies over time (Fum & Stocco, 2003a, 2003b). This theory of the dynamic control task utilizes the procedural subsystem of ACT–R[6] (Anderson & Lebiere, 1998). Specifically, participants start with a small group of general strategies (e.g., *pivot around target*; *stay-on-hit*) that are represented as productions. As training proceeds, the more often a strategy leads to success (i.e., achieving the goal state), the greater the increase in a parameter corresponding to its expected utility. "Good" strategies are used with greater frequency and "poor" strategies are used less. Strategies are chosen according to updated expected utilities in a stochastic rule. Strategies are thus tuned to the task goal (e.g., 6,000º), but are not tuned differently for different output states. Although such an approach can account for the observed transfer performance of experiential participants in Experiment 2 (because the expected utility of the strategies is similar for the old and new goals), we believe it has more difficulty accounting for several other findings. First, it appears that providing information about correct inputs for output states has a big impact on learning the task. The performance of participants who received a task hint in Experiment 1 was better than that of experiential participants, even for states not covered by the hint. Thus, this result appears inconsistent with the argument that instances do not play a significant role in learning the dynamic control task (Fum & Stocco, 2003b, p. 431). Second, table participants in Experiment 2 would seem to have little need to use the strategies hypothesized by Fum and Stocco, since they could recall the optimal response for a given situation; yet, as evidenced by their transfer test performance, they too acquired experiential knowledge at levels similar to those of experiential participants. In other words, when table participants could not rely on the model-based knowledge, there should have been little tuning of the strategies, and performance would be expected to be extremely poor. Although our study was not designed to test this theory (and this therefore remains a question for future research), we believe our results appear more consistent with the notion that participants do store instance-based information about their interaction with the system. We discuss this alternative below.

If participants do not acquire specific output–input information while learning the dynamic control task, but they do encode information about previous interactions with the system, what type of knowledge do they acquire? Although they endorsed a relatively specific lookup table representation, Dienes and Fahey (1995) noted that the lookup table conception can be extended from highly specific instances at one end of a continuum to more abstract representations at the other. Further, Dienes and Fahey (1995) suggested that representations may become more general with increased learning:

> . . . implicit learning itself may become less well-approximated by a lookup table as learning proceeds, and subjects may progressively extrapolate and interpolate to greater degrees around the situations trained on. (Dienes & Fahey, 1998, p. 609).

This suggests that, early in learning, participants rely upon specific instances where they made the appropriate response, but as the amount of experience with the task increases, this knowledge is represented in more abstract ways such as "when output is high, choose about 800 workers." Such a shift is also consistent with the "forgetting algorithm" proposed by Mathews (1991), which asserts that forgetting of distinguishing features of individual instances is adaptive as the commonality between instances (or "gist") is retained. This view can account for the previous findings of good performance on "specific situation" tests at short intervals (e.g., Dienes & Fahey, 1998). Further, it suggests that interference increases and there is greater forgetting as participants see more situations. This process allows commonality between similar situations to be detected and used by participants. After substantial training, these representations become more abstract and allow for greater flexibility of application (i.e., transfer). Thus, we suggest that participants who simply practice the task develop general, but contextually relevant, representations of which actions should be performed in given output–input situations.

Application to Real-World Learning

Much research effort has been expended to identify "pure" tasks which isolate implicit or explicit learning processes (Mathews, 1997). Although there has been much theoretical progress in the field, we feel this emphasis has led researchers to overlook the important issue of how learning processes interact. This is a particularly important oversight, since nearly every real-world task of any complexity (e.g., radiologist, fighter pilot, truck driver) relies upon both experience- and model-based knowledge for its performance. Our results join those of prior work (e.g., Sallas et al., in press; Stanley et al., 1989) in suggesting that the two types of processing can sometimes be fruitfully combined to facilitate learning (but not always; Domangue et al., 2004; Sun & Mathews, 2005).

To the extent that the characteristics of our dynamic control task reflect real-world learning, our findings also have potentially broader implications. For example, our hint+quiz condition is similar to the situation in which students are taught a subset of important facts about a task, and then they focus on important variables during subsequent practice (e.g., guided learning; see Mayer, 2004). Such guidance appears to be more beneficial than simply practicing the task without guidance would be (although we note that even our experiential practice is well structured). However, our results suggest that improvement in accuracy may sometimes be accompanied by slower responding. In fields where speed of execution is critical (e.g., air traffic control, emergency room medicine), this may be an important issue. Recent work from our laboratory in the artificial grammar paradigm suggests that slowing may often be a problem when model-based knowledge is given before training (Domangue et al., 2004), but that slowing can be overcome in situations in which model-based knowledge is provided during practice precisely when it is needed to perform the training task (Sallas et al., in press). Our results also suggest the importance of structuring practice as a means of enhancing the flexibility of knowledge. One critical aspect of our practice procedure is that it involves extensive exposure to all the potential states of the system (i.e., search space). Practice conditions that lack variability and are not a good match to "test" conditions are likely to lead to knowledge that transfers poorly (Mathews, 1997; Willingham, 1997). Finally, some good news for educators and coaches is that learners who are provided with, and use, specific task information during learning can show great improvement and still acquire useful experiential knowledge that can generalize to new situations. These findings, and the new questions they suggest, support the claim that pursuing a greater understanding of the interaction between these processes has the potential to illuminate the mechanisms involved in the development of expertise.

AUTHOR NOTE

This research was partially supported by the Army Research Institute Grant DASW01-00-K-0012 to R.S. and R.C.M. We thank Janet McDonald, Jason Hicks, Emily Elliott, and members of the LSU Memory Research group for their comments on an earlier version of this article. Correspondence concerning this article should be addressed to S. M. Lane or R. C. Mathews, Department of Psychology, Louisiana State University, 236 Audubon Hall, Baton Rouge, LA 70803 (e-mail: slane@lsu.edu or psmath@lsu.edu).

REFERENCES

Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review*, **89**, 369-406.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Berry, D. C., & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology*, **36A**, 209-231.

Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and implicit–explicit distinction. *British Journal of Psychology*, **79**, 251-272.

Berry, D. C., & Dienes, Z. (1993). *Implicit learning: Theoretical and empirical issues*. Hove, U.K.: Erlbaum.

Broadbent, D. E., Fitzgerald, P., & Broadbent, M. H. P. (1986). Implicit and explicit knowledge in the control of complex systems. *British Journal of Psychology*, **77**, 33-50.

Brooks, L. R. (1987). Decentralized control of categorization: The role of prior processing episodes. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 141-174). New York: Cambridge University Press.

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

Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 848-862.

Dienes, Z., & Fahey, R. (1998). The role of implicit memory in controlling a dynamic system. *Quarterly Journal of Experimental Psychology*, **51A**, 593-614.

Domangue, T. J., Mathews, R. C., Sun, R., Roussel, L. G., & Guidry, C. E. (2004). Effects of model-based and memory-based processing on speed and accuracy of grammar string generation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **30**, 1002-1011.

Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, **100**, 363-406.

Estes, W. K. (1986). Array models for category learning. *Cognitive Psychology*, **18**, 500-549.

Frensch, P. A., Haider, H., Runger, D., Neugebauer, U., Voigt, S., & Werg, J. (2003). The route from implicit learning to verbal expression of what has been learned: Verbal report of incidentally experienced environmental regularity. In L. Jimenez (Ed.), *Attention and implicit learning* (pp. 335-366). Amsterdam: John Benjamins.

Fum, D., & Stocco, A. (2003a). Instance vs. rule-based learning in controlling a dynamic system. In F. Detje, D. Dörner, & H. Schaub (Eds.), *Proceedings of the Fifth International Conference on Cognitive Modeling* (pp. 105-110). Bamberg, Germany: Universitäts-Verlag.

Fum, D., & Stocco, A. (2003b). Outcome evaluation and procedural knowledge in implicit learning. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 426-431). Mahwah, NJ: Erlbaum.

Hayes, N. A., & Broadbent, D. E. (1988). Two modes of learning for interactive tasks. *Cognition*, **28**, 249-276.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, **93**, 411-428.

Johnson-Laird, P. N. (1982). Ninth Bartlett memorial lecture: Thinking as a skill. *Quarterly Journal of Experimental Psychology*, **34A**, 1-29.

Kieras, D. E., & Bovair, S. (1984). The role of a mental model in learning to operate a device. *Cognitive Science*, **8**, 255-273.

Knowlton, B. J., & Squire, L. R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 169-181.

Kolers, P. A., & Roediger, H. L., III (1984). Procedures of mind. *Journal of Verbal Learning & Verbal Behavior*, **23**, 425-449.

Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 523-530.

Marescaux, P. J., Luc, F., & Karnas, G. (1989). Modes d'apprentissage sélectif et nonsélectif et connaissances acquises au contrôle d'un processus: Évaluation d'un modèle simulé. *Cahiers de Psychologie Cognitive*, **9**, 239-264.

Mathews, R. C. (1991). The forgetting algorithm: How fragmentary knowledge of exemplars can abstract knowledge. *Journal of Experimental Psychology: General*, **120**, 117-119.

Mathews, R. C. (1997). Is research painting a biased picture of implicit learning? The dangers of methodological purity in scientific debate. *Psychonomic Bulletin & Review*, **4**, 38-42.

MATHEWS, R. C., BUSS, R. R., STANLEY, W. B., BLANCHARD-FIELDS, F., CHO, J. R., & DRUHAN, B. (1989). Role of implicit and explicit processes in learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1083-1100.

MAYER, R. E. (2004). Should there be a three-strikes rule against pure discovery learning? The case for guided methods of instruction. *American Psychologist*, **59**, 14-19.

MCGEORGE, P., & BURTON, A. M. (1989). The effects of concurrent verbalization on performance in a dynamic systems task. *British Journal of Psychology*, **80**, 455-465.

MEDIN, D. L., & SCHAFFER, M. M. (1978). Context theory of classification learning. *Psychological Review*, **85**, 207-238.

NEWELL, A., & SIMON, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall.

PERSAUD, N., MCLEOD, P., & COWEY, A. (2007). Post-decision wagering objectively measures awareness. *Nature Neuroscience*, **10**, 257-261.

REBER, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning & Verbal Behavior*, **6**, 855-863.

REBER, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, **118**, 219-235.

REBER, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. New York: Oxford University Press.

ROUSSEL, L. G. (1999). *Facilitating knowledge integration and flexibility: The effects of reflection and exposure to alternative models*. Unpublished doctoral dissertation. Louisiana State University, Baton Rouge.

SALLAS, B., MATHEWS, R. C., LANE, S. M., & SUN, R. (in press). Developing rich and quickly accessed knowledge of an artificial grammar. *Memory & Cognition*.

SEGER, C. (1994). Implicit learning. *Psychological Bulletin*, **115**, 163-196.

SHANKS, D. R., & ST. JOHN, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioural & Brain Sciences*, **17**, 367-448.

STANLEY, W. B., MATHEWS, R. C., BUSS, R. R., & KOTLER-COPE, S. (1989). Insight without awareness: On the interaction of verbalization, instruction, and practice in a simulated process control task. *Quarterly Journal of Experimental Psychology*, **41A**, 553-577.

SUN, R. (2002). *Duality of the mind: A bottom up approach toward cognition*. Mahwah, NJ: Erlbaum.

SUN, R., & MATHEWS, R. (2005). *Exploring the interaction of implicit and explicit processes to facilitate individual skill learning*. Technical Report TR-1162, Army Research Institute for the Social and Behavioral Sciences, Arlington, VA.

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

WILLINGHAM, D. B., NISSEN, M. J., & BULLEMER, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1047-1060.

NOTES

1. We have adopted these terms (as opposed to other terms such as *implicit* or *explicit*; see Reber, 1993) for two major reasons. First, although experiential knowledge is largely nonconscious, people can and often do consciously access this information by observing the cues they are attending to while performing a task (see Frensch et al., 2003, for a related but somewhat different approach). In other words, participants can become explicitly aware of knowledge that they nevertheless learned in an experiential manner (e.g., Persaud, McLeod, & Cowey, 2007). Second, although we argue that these two categories are psychologically meaningful, we also assume that both types of processes operate in nearly all tasks (e.g., Mathews, 1997; see also Reber, 1989; Seger, 1994; Willingham, Nissen, & Bulmer, 1989; but see Lewicki, Czyzewska, & Hoffman, 1987), and thus the potential interaction of these processes should be considered.

2. Quiz-only participants were significantly slower to respond than experiential participants. We assume that the quiz manipulation encourages model-based processing, such that participants may attempt to discover the appropriate inputs for the output states (i.e., engage in hypothesis testing). Thus, they too may attempt to retrieve the "correct" inputs for an output state that they learned during training.

3. Experiment 2 used a version of the program that did not record response time (but was the same in all other respects). Thus, our only speed-related measure in this experiment is performance on the speeded test.

4. Although not of primary interest, the data also allow us to ask whether participants who learn the lookup table before taking the new goal test perform better on the test than participants who simply took the test. When compared directly, performance in the table control condition was nominally, but not significantly, more accurate on the new goal test than in the no training condition [$F(1,53) = 3.4$, $MS_e = 460{,}891$, $p < .07$, $\eta_p^2 = .06$].

5. We examined the number of times that experiential condition participants in Experiment 1 successfully achieved the goal from each of the 12 output states during Test 3 (which occurred just prior to the table test). According to Dienes & Fahey's (1995) theory, if participants had been successful at least once for a given output state, they should be able to store this information and later produce it on the table test. Our analysis revealed that, on average, participants achieved the goal at least once for 10.56 of the 12 output states. Note that this calculation does not include either the practice trials completed earlier in that session, or the practice and tests from the previous two sessions. Thus, we feel safe concluding that, according to Dienes and Fahey (1995), participants should have done much better on the table test if they had been storing specific output–input representations whenever they successfully reached the goal state.

6. Note that the argument that follows only applies to Fum and Stocco's (2003b) strategy-based view, not ACT–R. ACT–R could accommodate our results, for instance, by encoding general lookup table representations (e.g., "when output is high, choose about 800 workers") into production rules.

(Manuscript received April 23, 2007;
revision accepted for publication July 24, 2007.)