Journal of Experimental Psychology:
Learning, Memory, and Cognition
1995, Vol. 21, No. 4, 848–862

0278-7393/95/$3.00

Role of Specific Instances in Controlling a Dynamic System

Zoltán Dienes and Richard Fahey
University of Sussex

This article examines the claim that the learning of a dynamic control task is mediated by a lookup table consisting of previously successful trials on the task. Consistent with the predictions of a lookup table, in 2 experiments participants tended to give the same response to situations in which they had previously been successful rather than unsuccessful. Further, in both experiments, participants knowledge did not generalize to new dissimilar situations, unless the dynamic control task was governed by a highly salient rule. A version of G. Logan's (1988) instance theory, which assumes that participants store each successful response as a separate instance linking the situation to the response, was able to quantitatively match a range of measures of participants' performance with one free parameter, except in the case in which the control task was governed by a salient rule. In a complementary way, an alternative rule-based model could only match participants' performance when the control task was governed by a highly salient rule.

One of the major frameworks for understanding human learning has been in terms of the storage and deployment of specific exemplars or instances (e.g., Broadbent, Fitzgerald, & Broadbent, 1986; Brooks, 1978; Estes, 1986; Hintzman, 1986; Kruschke, 1992; Logan, 1988, 1990; Medin & Schaffer, 1978; Nosofsky, 1984; Perruchet, 1994; for a recent contrasting approach see Nosofsky, Palmeri, & McKinley, 1994). According to models of this type (called *memory array models* by Estes), people can learn to assign exemplars to categories (Brooks, 1978; Kruschke, 1992; Medin & Schaffer, 1978; Nosofsky, 1984), to predict successive elements in a sequence (Perruchet, 1994), to develop automatic or skilled action (Logan, 1988, 1990), or to control complex systems (Broadbent et al., 1986) by storing individual instances in memory. Responses to new test items can then be made on the basis of similarity with the stored instances. According to these models, abstraction is not an active process that occurs mainly at the time of learning. Rather, rule-governed behavior emerges from the way that test items are compared with stored instances.

Memory array models have compellingly been applied to concept formation paradigms, and particularly those involving complex or ill-defined concepts. Medin and Schaffer (1978) showed how participants' classification of test items was sensitive to the similarity of the test items to specific training items and not just to the average structure (the prototype) of all training items. Further, they showed how memory array models, like participants, can sometimes be sensitive to average structure: A memory array model could classify the prototype of a set of training stimuli better than any of the training stimuli, not because a prototype was explicitly abstracted during learning but because the prototype was similar to many of the training stimuli. Indeed, Nosofsky, Kruschke, and McKinley (1992) found that the one model that accounted for both sensitivity to individual training exemplars and to their collective probability structure was a memory array model and not a model that abstracted prototypes. Brooks (e.g., 1978) has persistently shown how participants can learn to classify instances generated by a finite-state grammar, not by abstracting rules of the grammar but by comparing test items with stored training examplars. Whether this is the only or even the typical strategy of participants learning artificial grammars is currently a point of debate (see, for example, Dienes, 1992; Reber, 1989), but the view that in these experiments participants store simply exemplars and their fragments remains an appealing position (cf., for example, Mathews, 1991; Perruchet, 1994).

Logan (1988, 1990) argued that the formation of automatic skills in general could be understood in terms of the storage of specific instances. For example, on a lexical-decision task, the repeated exposure of a specific letter string results in progressively faster responses. On the first exposure to the letter string, the participant brings to bear whatever general strategies are available for classifying the string as a word or nonword. This exposure and the response given are then stored as an instance, according to Logan. On the second exposure to the same letter string, the same general strategies can start to apply, but in addition the instance encoding the previous exposure is retrieved and can control response. Whichever process (general strategy or instance retrieval) is fastest controls the response; thus, reaction times speed up on average. Logan argued that much skilled action can be understood in terms of simple memory retrieval in the same way.

One area in which memory array models have been suggested but not yet thoroughly explored is the control of complex systems (e.g., Berry & Broadbent, 1984, 1987, 1988). In the dynamic control tasks used by Berry and Broadbent, the

Zoltán Dienes and Richard Fahey, Laboratory of Experimental Psychology, University of Sussex, Brighton, United Kingdom. Richard Fahey is now at the Department of Psychology, University of Texas.

This research was supported by the Medical Research Council. We thank Gerry T. M. Altmann, Dianne Berry, Axel Cleeremans, Gordon Logan, Christopher Longuet-Higgins, Pierre-Jean Marescaux, Robert Mathews, and Al Parkin.

Correspondence concerning this article should be addressed to Zoltán Dienes at the Laboratory of Experimental Psychology, University of Sussex, Brighton BN1 9QG, United Kingdom. Electronic mail may be sent via Internet to dienes@epunix.biols.susx.ac.uk.

participant controls the level of one variable (e.g., workforce in a simulated sugar production factory) to reach target values on another variable (e.g., the amount of sugar production). Participants acquire considerable knowledge about how to control such systems, as indicated by their progressive ability to reach and maintain target values. There is evidence that this knowledge is not in the form of rules about how the system works. First, participants are unable to predict what a given change of workforce will have on sugar production (e.g., Broadbent et al., 1986), suggesting that participants do not have a working mental model of how the system behaves (although such a model does emerge with extensive practice; Sanderson, 1989). Second, asking participants to search for the rules underlying system behavior sometimes deteriorates performance (Berry & Broadbent, 1988; cf. Stanley, Mathews, Buss, & Kotler-Cope, 1989), suggesting that participants do not generally learn in an analytic way.

Broadbent et al. (1986) suggested that in learning a dynamic control task a participant could construct a *lookup table* that would determine the appropriate action by matching the current situation to the most similar of the entries already in the table. Lookup tables can be regarded as being on one end of a continuum of abstractness of knowledge induced by a learning mechanism. Learning by the acquisition of a lookup table means that participants learn by tending to give the same responses to old situations that were followed by a successful outcome. There is no absolute definition of *old situation.* The usefulness of the notion of a lookup table depends on there being some simple way of construing the task such that participants perform consistently on old situations and at chance on *new situations.* The less preprocessing that goes into determining whether a situation is old, the less abstract the knowledge is. The "purest" lookup table might store each situation separately, with each perceptual feature of the situation equally weighted (cf. attention weights in memory array models; Estes, 1986; Nosofsky, 1984). Situations can be defined in progressively more abstract ways, until the lookup table blends imperceptibly into a more general rule-based model. On the other end of the continuum of abstractness of knowledge would be a system that induces rules like "respond half way between target sugar production and its current value" and "start at the extremes and work towards the middle." We do not claim that there is a sharp distinction between lookup tables and other mechanisms; we do, however, attempt to show that participants' learning can be understood in terms of a lookup table in which situations are simply defined in terms of the features provided by the experimenter and that the induction and application of more abstract rules play a very minor role in the learning, at least over the numbers of trials considered in this article.

A lookup table was formally instantiated in a model by Cleeremans (reported in Marescaux, Luc, & Karnas, 1989). The model built a lookup table relating situations to specific responses. If the current situation was matched in the table, then the associated response was made; otherwise, different responses were given with some baseline probabilities (all responses were equilikely in Cleeremans's case). If a response led to the target output, then the response was entered in the table.

Marescaux et al. (1989) attempted to test several predictions of the model. One prediction was that the participant should perform well on a specific-situation task in which the participant is presented with hypothetical situations that he or she has come across before and given a correct response to (e.g., "If you had just employed 400 workers and if the sugar production was then 8,000 tons, what should you do next to bring the sugar production to target?"); the participant should not perform so well in other situations in which there is as yet no entry in the lookup table. This prediction was confirmed. Another prediction was that there should be consistency of response to the same situation. To quantify consistency of response, Marescaux et al. defined *concordance* as the percentage of times that participants gave the same response in the specific-situation task as in learning. Thus, Marescaux et al. predicted that concordance should be high for situations participants had been correct on during training; indeed, this concordance was 57%.[1]

Marescaux et al.'s (1989) evidence for a lookup table is suggestive but not conclusive. The finding that participants perform better on previous situations in which they reached the target rather than other sorts of situations would be expected from almost any theory. If the participant has partially valid knowledge (e.g., a partially valid rule), then selecting just those situations for which the knowledge worked well before will tend to lead to good performance again. The finding that the participants had a concordance of 57% for situations previously given a correct response needs to be compared with a baseline concordance expected if participants' responses were not sensitive to the situation and also with the concordance for situations previously given an incorrect response. For example, if participants always used a single response there would be a concordance of 100% without any sensitivity to the situation. Further, there needs to be a sizable difference between concordances for situations previously given correct and incorrect responses for a lookup table to work. A rule-based system, on the other hand, may produce very little difference between these concordances. For example, the consistent application of a partially valid rule may lead to equal concordances for situations given correct and incorrect responses; learning a new partially valid rule may produce incorrect responses to situations in which the participant had previously been correct. Although a lookup table requires a difference between concordances for situations given correct and incorrect responses, a rule-based system does not, as we demonstrate later. The finding of a difference between concordances for situations previously given correct and incorrect responses obviously does not in itself rule out more general rule-based systems. On the other hand, finding a close match between concordances to situations previously given correct and incorrect responses would be difficult for a

[1] Another prediction tested by Marescaux et al. (1989) was that knowledge obtained with one target level of sugar production should not transfer to another target; this prediction was also supported (for previous consistent results, see Berry & Broadbent, 1987; McGeorge & Burton, 1989, Experiment 1; for contrary results, see McGeorge & Burton, 1989, Experiment 2; Sanderson, 1989; the reasons for this discrepancy are not clear).

lookup table approach because it is that difference that drives learning.

Marescaux, DeJean, and Karnas (1990) suggested an additional prediction of the lookup table model: If participants learn purely by a lookup table, then their performance on new situations should be at chance. In fact, participants performed above chance on randomly selected situations, leading Marescaux et al. to conclude that there was some rule learning. However, chance was calculated by assuming purely random responses. Participants did not randomly respond on the control tasks, even on the first few trials before any learning could have taken place (Dienes, 1990; Marescaux et al., 1989); correcting for participants' response biases can substantially change chance predictions. Calculation of a correct chance level for new situations is important because this final prediction of a lookup table model is particularly discriminating: Use of a lookup table means that the knowledge only applies to old situations or those sufficiently similar.

The two experiments reported in this article addressed these problems by measuring concordances for situations previously given correct responses and also for situations previously given incorrect responses and by comparing both of these to an appropriate baseline. Participants' performance on new situations was also compared with an expected chance level that took into account participants' initial strategies. In Experiment 1, we explored the predictions with the sugar production task (Berry & Broadbent, 1984); in Experiment 2, we explored the predictions with two versions of the person interaction task (Berry & Broadbent, 1988). In the Computational Model section of the article we investigate whether the exact effect sizes obtained for learning and concordances are in the range required by a specific lookup table model (Logan, 1988). Also in this section, by way of contrast, the performance of an alternative rule-based model is compared with participants' data.

The Sugar Production Task

Participants in Experiment 1 were trained on the sugar production task introduced by Berry and Broadbent (1984). Participants were asked to imagine that they were in charge of a sugar production factory. They were told they could change the amount of sugar produced by changing the size of the workforce. Their goal was to achieve and maintain a target sugar output of 9,000 tons. The starting workforce was 600 workers, and starting level of sugar output was 6,000 tons. On each trial, participants entered a number between 1 and 12 on the computer keyboard to represent the number of hundreds of workers they wished to use on that trial. The level of sugar production on Trial n was determined by the equation $P_n = 2 \cdot W - P_{n-1} + N$, where P_n is the number of thousands of tons of sugar output on Trial n, W is the number of hundreds of workers used by the participant, and N is noise (N could be -1, 0, or $+1$ with equal probability). Note that the optimal response varies only according to the previous level of sugar production: it is given by $W = P_{n-1}/2 + 4.5$.

If the equation resulted in a sugar output of less than 1,000 tons, then the output was simply set at 1,000 tons; similarly, if the equation resulted in an output of greater than 12,000 tons, then it was set at 12,000 tons. Participants were aware of these lower and upper limits. Because of the noise, N, in the equation, participants' responses could be counted as correct if they resulted in an output on target or one level off, although participants were not aware of this loose method of scoring.

Experiment 1

Method

Participants. The participants were 24 paid volunteers aged between 18 and 35 years from the Sussex University Experimental Psychology participant panel. Eighteen participants attempted to control the sugar production factory (experimental participants); the remaining 6 simply described how they thought they would control the factory (strategy participants).

Procedure. Experimental participants were tested on two sets of 40 trials of the sugar production task. On each trial, participants saw a graph indicating the level of performance on all previous trials of that set. A horizontal line indicated the target performance. In addition, written information about the level of workforce and the level of sugar production for the last trial was presented above the graph (see Figure 1). This was the same display used by Berry and Broadbent (1984) and Marescaux et al. (1989).

The computer kept a record of all situations that the participant came across. A *situation* could be defined as the current level of sugar production, or the workforce entered on the last trial, or by the combination of the two. We tabulated the situations into those for which the participant entered a workforce that resulted in the target sugar output or only one level off (i.e., situations followed by a correct response) and those that were followed by an output more than one level from target (i.e., situations followed by an incorrect response). A loose method of scoring was used because of the noise, N, in the equation determining sugar output described above.

Next, participants performed the *specific-situation* task used by Marescaux et al. (1989). In this task, participants were shown possible situations consisting of written information about the current level of sugar output and the workforce used on the last trial; in addition, the three previous levels of production were graphically displayed (this was the procedure used by Marescaux et al., 1989). Participants were told to enter the next level of workforce they thought would achieve or maintain the target level of output, on the basis of their previous experience of controlling the factory. The same target level was used as in the training phase. Participants were told that after each situation, the next situation to be shown would be unrelated to the workforce they had just entered; it would simply be another possible situation, and thus, they would get no feedback on how successful they were. The old situations were all 40 situations experienced in the last block of training trials. In addition, 40 combinations of workforce and sugar production were presented that had not been experienced in either block. These additional combinations were separately determined for each participant by randomly selecting from all possible combinations of workforce and sugar production that had not been experienced as combinations by that participant.

As well as the specific-situation task, participants also performed a recognition task. The recognition task was not relevant for the issues addressed in this article and is analyzed in Dienes and Fahey (1994). The specific-situation and recognition tasks were tested in counterbalanced order.

Results

The degrees of freedom for some of the subsequent analyses were less than (total number of participants, minus one) because the raw data file for 1 participant was missing.

Initial learning. Participants' mean numbers of trials correct on the first and second blocks of the sugar production task are shown in the first two rows of Table 1, in the participants

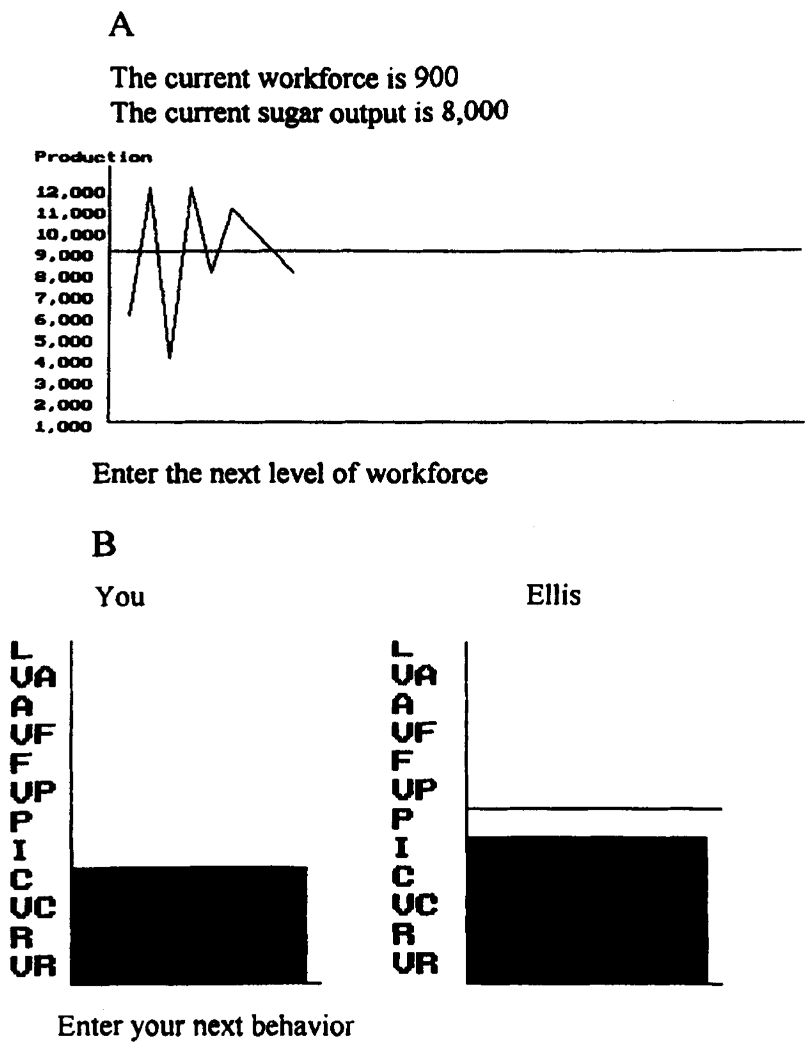

Figure 1. (A) Shows a typical display seen by participants in Experiment 1, and (B) shows a typical display seen by participants in Experiment 2. VR = very rude; R = rude; VC = very cool; C = cool; I = indifferent; P = polite; VP = very polite; F = friendly; VF = very friendly; A = affectionate; VA = very affectionate; L = loving.

column. An output one off target was counted as correct. The difference between the first and second blocks was significant, $t(17) = 4.78, p < .0005$.

Concordance. Concordance is the percentage of times that participants gave the same response in the specific-situation task as in learning. We calculated concordance for situations followed by a correct response by comparing the participants' response to a situation in the specific-situation task with the responses given to all occurrences of that situation in the training phase that were followed by the target or were one level of sugar production off. The proportion of times that the responses were the same was determined and averaged over all relevant situations in the specific-situation task. We calculated the concordance for situations followed by an incorrect response in a corresponding way.

It is not known a priori how situations were psychologically defined for participants. For example, participants may or may

Table 1
Sugar Production Task

	Participants		Logan		Rule based	
Condition	*M*	*SD*	*M*	*SD*	*M*	*SD*
Performance						
Block 1	7.9	3.7	8.1	3.8	8.0	4.7
Block 2	15.3	8.0	13.2	5.5	13.4	6.2
Concordance						
Baseline	.19	.08	.15	.04	.18	.04
Sugar output						
Correct	.33	.23	.32	.10	.31	.12
Incorrect	.21	.27	.13	.05	.13	.04
Workforce						
Correct	.38	.24	.37	.07	.31	.12
Incorrect	.27	.11	.26	.06	.24	.04
Combination						
Correct	.37	.26	.49	.13	.37	.15
Incorrect	.28	.16	.29	.08	.20	.20

not have been sensitive to the level of sugar production, and they may or may not have been sensitive to the level of workforce. Thus, concordances were defined separately for situations defined in terms of sugar production and for situations defined in terms of workforce. The means are displayed in Table 1 in the participants column.

We calculated a baseline concordance by determining the proportions of each of the 12 possible responses in the sugar production task overall and also in the specific-situation task overall. Multiplying together the proportions for the same response in the different tasks gives the probability of the participant using that response twice, assuming the responses are independent, that is, there is no relation with the situation. By adding all 12 such probabilities, an overall concordance baseline can be determined for each participant. The mean baseline concordance is shown in Table 1. Concordance for situations previously given a correct response was greater than baseline both when situations were defined in terms of sugar production, $t(16) = 2.70, p = .016$, and when situations were defined in terms of workforce, $t(16) = 3.41, p = .0056$. That is, participants' responses were associated with situations.

In terms of situations defined by the level of sugar production, the concordance was significantly greater for situations given a correct response (0.33) rather than an incorrect response (0.21), $t(16) = 2.12$, $p < .05$ (Wilcoxon $p = .06$, two-tailed; see Table 1). In terms of situations defined by the same sugar production but a different workforce, the concordance for situations given a correct response was 0.27 ($SD = 0.20$) and that for situations given an incorrect response was 0.16 ($SD = 0.12$). It can be seen that the difference between concordances for situations given correct or incorrect responses was virtually identical when the situation was defined by a different workforce rather than the same workforce (0.12 cf. 0.11, $t < 1$); that is, participants were sensitive to sugar production in the situation, regardless of the workforce.

We conducted a similar analysis for the participants' sensitivity to the workforce in the situation. When the situation was simply defined by the level of workforce, the concordance was significantly greater for situations given a correct response (.38) rather than an incorrect response (.27), $t(16) = 2.30, p < .05$ (Wilcoxon $p < .05$, two-tailed; see Table 1). The concordances for situations defined by the same workforce but a different sugar production were .39 ($SD = .30$) and .27 ($SD = .12$). Note that the difference between the concordances for situations given a correct response and for those given an incorrect response was virtually identical, whether the situation included a different sugar production or the same sugar production (.11 cf. .12, $t < 1$); that is, participants were sensitive to workforce in the situation, regardless of the sugar production.

Performance on new compared with old situations. This section compares participants' actual performance on new and old situations with those predicted by participants' explicit knowledge. The Appendix describes how we calculated explicit knowledge.

The percentage of loosely correct responses on the specific-situation task was determined for situations consisting solely of levels of sugar production the participant had not come across before and levels of workforce that the participant had not come across before.[2] A lookup table approach predicted participants would only perform as well as the explicit strategy allows on these new situations. Six participants had completely new situations of this sort. The mean performance is shown in Table 2 (participants column). Because these new situations were not necessarily a random sample from all possible situations, we calculated a predicted performance for just these situations, assuming that participants used the explicit strategy described in the Appendix. The actual performance was .00, and the predicted performance was .02.[3] That is, participants' performance was no better on new situations than that predicted by the explicit strategy.

The old row in Table 2 (participants column) shows participants' performance on situations containing a level of sugar production that the participants had come across before and given a correct response to (sugar production was used because it was the only aspect of the situation relevant to producing a correct response). All participants had situations of this sort, but the raw data file for 1 participant was missing. The performance on these old situations was significantly greater than performance predicted by the explicit strategy for these situations (.32 compared with .14), $t(16) = 5.11$, $p < .001$, and also significantly greater than performance on new situations, $t(5) = 8.80, p < .001$.

To summarize, participants performed better than the explicit strategy would predict only in old situations in which they had previously been correct and not in new situations.

Discussion

Experiment 1 addressed the question of whether participants learn to control the sugar production task by forming a lookup table. As we argued in the introduction, the usefulness of the notion of a lookup table depends on there being some simple way of construing the task such that participants perform consistently on correct old situations and at chance on new situations. Then we have a straightforward way of under-

[2] Situations that were one level different from old correct situations were excluded. This is because there may be generalization between similar situations. This possibility was analyzed by a further analysis of the increase in concordance for situations previously given a correct rather than an incorrect response. This increase was determined for situations in training that were one, two, or three levels (of sugar production or workforce) distant from a given situation in the specific-situation task. The increase declined significantly with distance, $F(2, 32) = 4.77, p < .05$. At a distance of one, the increase, .12, was marginally significant, $p < .10$, numerically smaller at a distance of two (.06), and negative at a distance of three (−0.05, nonsignificantly different from zero). The decline of the increase with distance did not interact with whether situations were defined by sugar production or workforce ($p > .10$).

[3] This result was replicated in the experiment reported by Dienes and Fahey (1994). They tested participants on a specific-situation task that presented participants only with levels of sugar production (i.e., no information was given about workforce). On new levels of sugar production, only 1 participant out of 14 scored above zero.

Table 2
Performance on New and Old Situations on the Specific-Situation Task

Condition	Participants		Logan		Rule based	
	M	*SD*	*M*	*SD*	*M*	*SD*
Sugar production task						
New	.00	.00	.13	.16	.37	.16
Old	.32	.15	.34	.21	.40	.14
Person U						
New	.23	.16	.24	.06	.28	.16
Old	.54	.26	.60	.32	.32	.22
Person S						
New	.41	.32	.25	.06	.48	.09
Old	.80	.20	.72	.26	.86	.19

standing participants' learning. The results of the experiment indicate that participants' learning of dynamic control tasks can indeed be understood in this way by defining situations in terms of the elements of the immediately preceding trial. First, participants were consistent in their responding to a situation for which they were previously correct, and to a greater extent than their consistency in responding to situations for which they were previously incorrect. Further, the data indicate that participants were not just sensitive to unique combinations of workforce and sugar production, but rather they were independently sensitive to both workforce and sugar production. Second, participants correctly responded at above baseline levels in the specific-situation task only for situations containing a level of sugar production that they had previously given a correct response to; participants responded at baseline levels to new situations. Participants did not appear to learn any rules that they could generalize to new situations.

Experiment 2

Our confidence in the generality of the findings of Experiment 1 could considerably be strengthened if it could be shown that participants still behave as if they were using a lookup table when controlling a system obeying a very simple rule. In Experiment 2, unlike Experiment 1, we used a task in which the correct response was the same no matter what the situation was. That is, the only rule participants needed to learn (for optimal performance) in Experiment 2 was to press repeatedly a certain key. In Experiment 2 we also investigated the conditions under which mechanisms other than lookup table learning may come into play.

Berry and Broadbent (1988) and Hayes and Broadbent (1988) introduced two versions of a *person interaction task* that they argued were learned in quite distinct ways. The person interaction task involves the participant telling a computer personality, Ellis, how friendly the participant is being to Ellis, and Ellis responds in turn with a certain level of friendliness. The participants' aim was to keep Ellis at a target level of friendliness. The two versions of the task (Person S and Person U) differed according to the rule linking Ellis's behavior to the participants'. Berry and Broadbent argued that the rule was salient in the case of Person S and nonsalient in the case of Person U. Consistently, they found that after initial experience with the task, instructions to search for the rule connecting Ellis's responses to the participants improved performance with Person S and deteriorated performance with Person U. Further, participants could predict what Ellis would do next reasonably accurately for Person S but not for Person U. Berry and Broadbent argued that Person U was learned by unselectively storing contingencies in a lookup table but that Person S was learned by formulating and testing more general rules. If this claim is true, then participants learning Person U (but not necessarily Person S) should show the characteristic difference between concordances for different types of situations, as predicted by a lookup table. Further, knowledge of Person S should generalize to new situations, but knowledge of Person U should not generalize to new situations dissimilar to old ones.

The Person Interaction Task

Friendliness varied along a 12-point scale ranging from (1), *very rude, rude, very cool, cool, indifferent, polite, very polite, friendly, very friendly, affectionate, very affectionate,* to *loving* (12).

Berry and Broadbent (1988; see also Green & Shanks, 1993; Hayes & Broadbent, 1988; Sanderson, 1990) used two equations for controlling Ellis's behavior, a "salient" equation and a "nonsalient" equation. The salient equation was $E_n = S_n - 2 + N$, where E_n is a number between 1 and 12 representing Ellis's behavior on the 12-point scale on Trial n, S_n is the participant's behavior on the 12-point scale on that trial, and N is noise (-1, 0, or $+1$ with equiprobability). The nonsalient equation was $E_n = S_{n-1} - 2 + N$. Because of the noise, a response of Ellis up to one off target was counted as correct. Note that for the nonsalient equation, Ellis's behavior depends not on how the participant just responded on that trial but on how the participant responded one trial back. The target value for both tasks was *polite,* so the optimal strategy for both tasks was simply to enter *friendly* no matter what the situation.

Method

Participants. The participants were 69 paid volunteers aged between 18 and 35 years from the Sussex University participant panel. Forty-eight were experimental participants who interacted with Ellis, and 21 were strategy participants who were only shown the instructions to the task.

Procedure. In the learning phase, experimental participants interacted with Ellis for one block of 30 trials for Person S and one block of 50 trials for Person U. Hayes and Broadbent (1988) found that if performance is taken as the number of correct responses in the last 10 trials, 50 trials of Person U is needed to produce the same performance as 30 trials of Person S; to be consistent with previous studies, we adopted this procedure. Participants in the visual learning condition saw two bar charts: one representing Ellis's behavior on the previous trial, and the other their behavior on the previous trial. A horizontal line in Ellis's chart indicated the target behavior (polite). The bars moved up and down according to Ellis's and the participant's respective behaviors (see Figure 1). Participants in the auditory learning condition did not see the computer screen. Zoltán Dienes simply said "Your behavior was X; Ellis's behavior was X." The scale of possible behaviors was placed in front of the participants to remind them. All participants entered their response by typing in the corresponding initials (e.g., VA for very affectionate).

The computer kept a record of all situations the participant came across. A *situation* could be defined as the current level of Ellis's behavior, the participant's behavior on the last trial, or a combination of the two. The situations were tabulated into those for which the participant entered a behavior that resulted in the target level of Ellis's behavior or were only one level off and those that were followed by a level of Ellis's behavior more than one level from target.

After the learning phase, participants performed the specific-situation task. In the specific-situation task, participants were shown possible situations consisting of Ellis's and the participant's behaviors on the preceding trial. Participants in the visual testing condition saw the information displayed as bar charts; participants in the auditory testing condition heard the experimenter read the information out loud. Participants were told to enter the behavior they thought would achieve or maintain the target level of Ellis's behavior, on the basis of their previous experience of interacting with Ellis. The target level was the same as that used in the training phase. Participants were told that after each situation, the next situation to be shown would be unrelated to the behavior they had just entered; it would simply be another possible situation, and thus, they would get no feedback on how successful they were being. Each participant was shown all combinations of Ellis's behavior and participant's behavior experienced as a combination in the training trials plus an equal number of new combinations. The new combinations were determined separately for each participant by randomly selecting from all possible combinations of Ellis's behavior and participant's behavior that had not been experienced as combinations by that participant.

As well as the specific-situation task, participants also performed a recognition task. The procedure and results for the recognition task are described in Dienes and Fahey (1994). The specific-situation and recognition tasks were tested in counterbalanced order.

Design. The main manipulation was person (Person S vs. Person U), but the modality with which information was presented was manipulated in both the training phase and the specific-situation task: Participants could read Ellis's friendliness or they could be told it. This manipulation is irrelevant to the point of this article and produced no main effects or interactions (further details are described in Dienes and Fahey, 1994).

In Experiment 2 we used a 2 × 2 × 2 (Person [Person S vs. Person U] × Learning Modality [visual vs. auditory] × Specific-Situations Task Modality [visual vs. auditory]) between-subjects design. Equal numbers of participants were allocated to each of these cells. In addition, there was a single group of strategy participants.

Results

Initial learning. Participants were scored for the number of trials correct in the first set of 10 trials and in the last set of 10 trials. Table 3 displays the means and standard deviations. A 2 × 2 × 2 × 2 (Block [first set of 10 trials vs. last set of 10 trials] × Person [Person S vs. Person U] × Learning Modality [visual vs. auditory] × Testing Modality [visual vs. auditory]) analysis of variance (ANOVA) on the number of trials correct indicated a significant effect of person, $F(1, 40) = 10.47, p < .005$, and of block, $F(1, 38) = 14.29, p = .0005$. That is, participants scored more trials correct for Person S (6.5) than for Person U (4.9), despite the fact that participants received more trials for Person U rather than Person S. Also, participants performed better on the last block (6.5) rather than on the first (5.0). The interaction of block with person was not significant, $p > .10$. The improvement from the first to the last block was 1.3 for Person U, $t(23) = 2.65, p < .05$, and 1.7, $t(23) = 2.92, p < .01$, for Person S.

Table 3
Person Interaction Task

Condition	Participants *M*	Participants *SD*	Logan *M*	Logan *SD*	Rule based *M*	Rule based *SD*
Person U						
Performance						
First 10 trials	4.3	2.5	4.3	2.4	3.2	1.9
Last 10 trials	5.5	2.1	5.0	2.7	3.0	2.3
Concordance						
Correct	.33	.17	.39	.18	.21	.14
Incorrect	.24	.14	.30	.16	.21	.10
Person S						
Performance						
First 10 trials	5.7	2.2	5.1	2.5	5.1	2.1
Last 10 trials	7.4	2.0	6.9	2.3	7.4	1.7
Concordance						
Correct	.40	.17	.41	.14	.42	.14
Incorrect	.24	.17	.31	.18	.26	.15

The scores on the last block for Persons S and U were 7.4 and 5.5, respectively. Hayes and Broadbent (1988) reported corresponding scores that were more closely matched (viz., 6.9 and 7.1, respectively). However, Sanderson (1990) reported corresponding scores of 7.6 and 5.7, respectively, closer to the values obtained in this article.

Concordance. Concordance is the percentage of responses to old situations in the specific-situation task that was the same as the response given originally to that situation. There were insufficient data to look at the features of the situations separately, so for this analysis, we consider a situation to be a combination of Ellis's behavior and the participant's behavior. The lookup table approach predicted that there would be a higher concordance for situations that were previously given a correct rather than an incorrect response in the learning phase. Table 3 displays the means and standard deviations. A 2 × 2 × 2 × 2 (Situation Type [loosely correct vs. incorrect] × person [Person S vs. Person U] × Learning Modality [visual vs. auditory] × Testing Modality [visual vs. auditory]) mixed-model ANOVA indicated a significant effect of situation type, $F(1, 40) = 13.26, p < .001$. That is, the concordance was greater for situations previously given a correct response (.37) rather than an incorrect response (.25). There were no other significant effects.

Performance on new and old situations in the specific-situation task. This section compares actual performance on old and new situations and compares actual performance with performance predicted by an explicit strategy. The Appendix describes how we calculated predicted performance. There were too few new participant or computer behaviors to have sufficient power to test individually whether participants were sensitive to each aspect of the situation. Thus, in the rest of this section, situations were defined in terms of combinations of participant's and computer's behaviors.

Table 4 shows the proportion of correct responses to

Table 4
Specific-Situation Task: Performance on New Situations

	n											
	1		2		3		4		5		6	
Condition	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Person												
U	.56	.38	.37	.27	.37	.31	.35	.35	.24	.34	.21	.28
S	.60	.20	.54	.26	.47	.24	.32	.27	.42	.38	.68	.42

Note. n is the difference (in levels of computer behavior, participants' previous behavior, or both) between the new situation and the nearest old correct situation. The maximum n can be is 22.

situations that had not occurred in training. We classified these new situations according to how many levels different (n) they were from the most similar old correct situation (different in participant's behavior, computer's behavior, or both). We calculated regression slopes for each participant to determine the rate of drop off of proportion correct with distance, $n = 0$ to 6. For Person U, the mean slope was $-.06$, which is significantly different from 0, $t(21) = 5.09$, $p < .0001$. For Person S, the mean slope was $-.06$, which is also significantly different from 0, $t(23) = 4.25$, $p < .0005$. This effect of similarity to old situations shows that participants had acquired knowledge sensitive to specific situations for both Persons U and S, despite the fact that the correctness of a response to either Person was not dependent on situation.

Table 2 shows performance on new situations more than four different from old correct ones (in the new row and participants column) and performance on situations that had only ever been loosely correct in training (in the old row). Crucially, for Person U, performance on these new situations was not significantly greater than predicted performance for these situations that were based on explicit knowledge (.23 compared with .30). Conversely, performance on old correct situations was significantly greater than that predicted for these situations by explicit knowledge (.54 vs. .32), $t(21) = 3.61$, $p < .01$. That is, consistent with the predictions of a lookup table, for Person U, the participant's acquired knowledge applied only to situations (combinations of participant's and computer's behavior) the participant had actually experienced or those that were sufficiently similar.

For Person S, unlike Person U, performance on new situations was significantly different from predicted performance that was based on explicit knowledge (.41 compared with .18), $t(23) = 2.80$, $p < .05$. Also, performance on old correct situations was different from that predicted by explicit knowledge (0.80 vs. 0.35), $t(23) = 7.90$, $p < .001$. Note also that for Person S performance on new situations more than four different from old correct ones (.41) was no different from performance on situations exactly four different from old correct ones (.42). That is, participants learning Person S may have acquired some knowledge that was relatively situation insensitive.

Discussion

The aims of Experiment 2 were, first, to determine whether the results of Experiment 1 generalized to a control task that could optimally be controlled by giving only one response and, second, to determine whether different versions of the control task that differ in the transparency of the underlying rule are learned in different ways. Participants learning both Persons S and U acquired situation-sensitive knowledge: Participants performed best on old previously correct situations rather than on new situations, and performance deteriorated on new situations the more dissimilar they were to old correct ones. Participants learning Person U but not Person S performed at baseline levels on new situations dissimilar to old correct ones by more than four levels. Interpretation of these results is confounded by the fact that initial learning was greater for Person S rather than for Person U, but one possibility is that participants learned Person U entirely by a lookup table and thus did not acquire any knowledge that applied to new situations sufficiently dissimilar to old correct ones. However, participants may have learned Person S perhaps partly by a lookup table but also by acquiring knowledge that could apply to many new situations. The idea that there could be two types of learning in general—one exemplar based and one rule based—is supported by the results of Nososfsy, Clark, and Shin (1989). They fitted an exemplar-based model and a rule-based model to the performance of participants classifying perceptual stimuli. The performance of participants asked to use rules was fitted better by the rule-based rather than the exemplar-based model; the performance of control participants, not so instructed, was fitted better by the exemplar-based rather than the rule-based model (see also Shanks & St. John, 1994).

We now consider whether a computational model of a lookup table can reproduce the pattern of data found for participants performing the sugar production and person interaction tasks, and whether it can do so in a more satisfactory way than an alternative rule-based model.

Computational Models

Logan's (1988) Instance Theory

Logan (1988, 1990) presented a theory in which automatization and repetition priming were construed as the acquisition of a domain-specific database formed of separate representations, or instances, of each exposure to the task. According to the theory, encoding into memory and retrieval from memory are obligatory consequences of attention. The theory assumes that each encounter with a stimulus is encoded and retrieved separately. When the participant performs a task, each stored episode relevant to the current situation races against the others and against any general problem-solving strategies (explicit knowledge) applied to the task; the first one past the finishing post controls performance. The race can be modeled by assuming that each episode has the same distribution of finishing times.

Logan (1988, 1990) showed that the model could give a good fit to the means and variances of reaction times to repeated presentations on a lexical-decision task. Logan (1992) investigated reaction times to perform an alphabet arithmetic task, in which participants were asked to verify equations of the form A + 2 = C (i.e., Is C two letters higher than A?). Participants initially counted through the alphabet to perform the task, but with practice they came to remember which equations were true and thus relied on memory retrieval. This change from a general purpose strategy to memory retrieval is the process that Logan defined as the development of automaticity. Logan (1992) also investigated a dot-counting task in which participants were shown a number of dots on a grid and were to report the numerosity. Participants improved on this task so long as the same instances were repeated. Further, the means and variances of the reaction time distributions of both the arithmetic and the counting task closely fitted the predictions of the theory.

The sugar production task. To model the sugar production task, we assumed that whenever an action led to a sugar production that was loosely correct,[4,5] the situation together with the response was stored. One instance was stored for the link between the current sugar production and the action, and a separate instance was stored for the link between current workforce and the action, consistent with the previous results that participants treat sugar production and workforce independently. Explicit knowledge was represented by a constant number, N, of instances that could be activated by any situation and that specified the explicit strategy described in the results for Experiment 1. Thus, on any given trial, all of the explicit instances and any specific instances that matched the current situation competed in a race. Because all of the situations had the same distribution of finishing times, the race was simply a random selection of one instance amongst the available instances, regardless of their distribution. That is, because responses rather than reaction times were being modeled, the distribution of finishing times was not important. Further, consistent with the theory, each encountered episode was stored separately, regardless of whether that episode had been encountered before.

When first controlling the sugar production factory, the model would only apply the explicit strategy; however, as specific episodes were stored, it became more and more likely that responding would be controlled by specific previous experiences. There is one parameter to vary: N, the number of episodes representing explicit knowledge. For a high value of N ($N > 15$), the explicit strategy was applied on virtually all of the 80 trials, and the model's learning was small compared with participants (see Figure 2). For a low N ($N < 5$), the concordance for situations previously given a correct response was considerably higher than participants because the explicit strategy did not provide much competition (also illustrated in

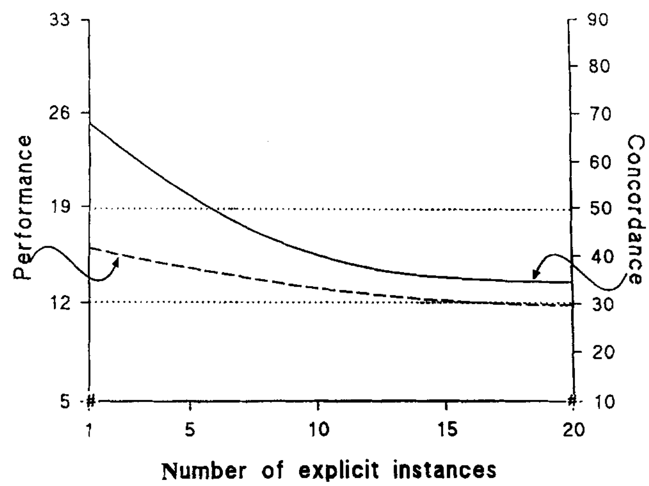

Figure 2. The effect of number of explicit instances on the Logan model simulating the sugar production task. For illustration, the figure shows performance on the second block of trials and concordance for situations that were defined by workforce and were previously given a correct response to. The dotted lines show the participants' 90% confidence intervals for both performance and concordance. When the number of explicit instances is small, concordance is too high; when the number of explicit instances is high, performance is too low.

Figure 2). For each value of N, the model was tested 50 times, each time starting with a different random seed. N was adjusted to give a good fit to the data. With $N = 10$, performance measures are shown in Tables 1 and 2. In Table 1, all nine means, as well as the degree of learning and the differences between corresponding concordances for situations previously given correct or incorrect responses, were within the 90% confidence intervals of participants' data. Table 2 shows the performance in the specific-situation task on old situations (defined by levels of sugar production to which the model only ever gave the correct response in training) and new situations (defined by levels of sugar production and workforce, each of which had not been experienced in training; see Footnote 2). The pattern was qualitatively similar to that of participants. The model's performance on old situations was within the 90% confidence intervals of participants' data. Participants had no variability on new situations, so we could

[4] When situation was defined by sugar production, the concordance for situations that were loosely correct without ever having been strictly correct (.27) was significantly greater than the concordance for situations that had been incorrect (.18 for these participants), $t(11) = 2.53, p < .05$. That is, being loosely correct appeared to be as good as being strictly correct in terms of reinforcing the response.

[5] All instances in which situations exactly matched the current situation, as well as all instances in which situations were exactly one off from the current situation, competed in the race. This is consistent with the results in footnote 2 that participants treat situations one off from the current situation in the same way as if they matched exactly, but those two off had only a minimal impact. This assumption was important to the success of the model. If instances entered the race only if they exactly matched the current situation, then for the same level of learning as participants, concordances were significantly greater than those of participants ($ps < .05$).

not calculate confidence intervals for participants. However, note that participants' performance on new situations closely matched that predicted by the explicit strategy for those situations (see the *Results* section for Experiment 1), and, from its assumptions, the model necessarily performed on new situations exactly as the explicit strategy predicted.

Predictions of the Logan model can more directly be tested on individual participants' data by considering the relation between concordance and repetition of a situation. Consider the concordance to situations that were correct in training. According to the Logan model, if there are N instances of the explicit strategy, and the model has experienced O occurrences of each feature of the current situation associated with the correct response, and D instances coding a different response, then the probability P of selecting an instance representing the correct response is given by

$$\mathrm{P} = O/(N + O + D). \quad (1)$$

If an explicit instance is selected, the concordance is that expected on the basis of the explicit strategy; if a situation instance is selected, then the same response as the old situation is given. In general, if CE is the concordance expected by using the explicit strategy, then the expected concordance is

$$\langle \text{concordance} \rangle = \mathrm{P} + \mathrm{CE} \cdot N/(N + O + D). \quad (2)$$

The effect of O on concordance was measured in the training phase of participants' data. For each trial of the training phase, the number of previous occurrences, O, of each of its features that were followed by the target (or a level of sugar production one level off) was determined.[6] For $O \leq 3$, a restriction was that D should equal zero: This was to ensure that the predicted concordance would increase maximally with O. Participants did not have data for $O > 3$ given this restriction, so for $O = 5$ the restriction was that $D \leq 1$, and for $O = 9$ the restriction was $D \leq 3$. (Note that the only reason D was varied was to allow there to be data to be modeled; D values were chosen before model fits were determined.) The concordance was then measured for the trial. For each participant, concordance was then averaged separately for each level of O. Figure 3 shows how concordance varied with O.

We calculated estimates for the concordance expected by using the explicit strategy for each situation and response by performing 50 simulations of the explicit strategy and by determining the probability that it would produce that response in the situation. Thus, for each participant, the expected concordance for each situation could be calculated and averaged over situations by using Equation 1. N was set at 10 for each participant. These model predictions are also shown in Figure 3. There were no significant differences between participants' data and the models' prediction for any O (t tests with arcsine transformation and with a Wilcoxon signed-ranks test). The correlation between participants' means and the model's means was .91, $p < .02$.

Finally, the importance to the modeling of assuming that only correct situations were stored was tested by determining

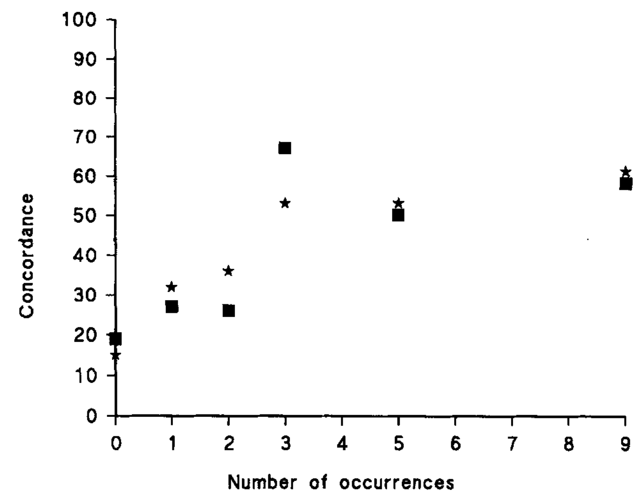

Figure 3. For the sugar production task, what is shown in the figure is how concordance varies with number of previous occurrences of the features of the situation that were followed by the target (or were one level of sugar production off). Participants' data are plotted as squares, and the predictions of the Logan model are plotted as stars.

the performance of the model when it stored all instances. If a situation was incorrect, retrieval of that instance would lead to application of the explicit strategy, ensuring only that the response stored in the instance was not given. This model could not perform the task as well as participants: The irrelevant workforce situations provided too much noise by proscribing responses that were in fact appropriate.

The person interaction tasks. Logan's theory was applied to Person U in the same way as it had been applied to the sugar production task, except in this case situations were defined in terms of combinations (of participants' behavior and Ellis's behavior), and an old situation previously given a correct response could be activated by any current situation up to four levels different from it, consistent with the previous data. The explicit strategy was represented by N instances, as before. We conducted sets of 100 simulations for $N = 1$ to 15. For $N = 7$, the performance of the model is shown in Tables 2 and 3. In Table 3, all four simulation means were within the 95% confidence interval of participants' data, as were the difference between performance on the first and last 10 trials and the difference in concordances. In Table 2, performances on new and old situations, as well as the difference between them, were within the 90% confidence interval of participants' data.

Figure 4 shows how concordance varied with O for both the participants and the model (D was always zero for these data). There were no data for $O > 3$. The predictions of the Logan model were all nonsignificantly different from the participants' data (by using t test with arcsine transformation or with a Wilcoxon signed-ranks test). The correlation between participants' means and model means was .94.

---

[6] To allow for generalization (see footnote 2), we included previous occurrences of a feature that were one level off (and were followed by the target or were only one level off) in determining O.

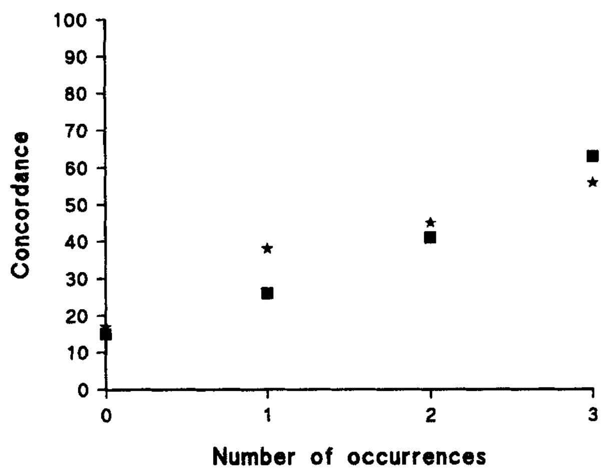

Figure 4. For Person U, what is shown in the figure is how concordance varies with number of previous occurrences of the situation that were followed by the target (or were one level of sugar production off). Participants' data are plotted as squares, and the predictions of the Logan model are plotted as stars.

With the same parameter value, the simulation scores for Person S are also given in Tables 2 and 3. *N* was varied between 1 and 15, but at no value did the simulation means fit the participants. For higher *N*, the explicit strategy dominated the model's performance, and there was little learning. For all *N*, the model's performance on new situations lay below the 95% confidence intervals of participants' data. From its assumptions on new situations, the model necessarily performed exactly as the explicit strategy predicted. On new situations, participants performed significantly higher than the explicit strategy predicted.

A Rule-Based Model

The performance of our implementation of Logan's model was compared with a rule-based model that learned in a similar way. The rule-based model started with a number of rules. The rules competed amongst themselves in a race to control performance. The first rule past the finishing post was the one that determined the response. If the application of a rule led to the target, or one level of sugar production off, another token of that rule was added to the set, increasing the chances of that rule winning in the next race. All tokens of all rules competed in every situation. This model was constructed to be as similar to our implementation of Logan's model as possible, but with general rules controlling behavior rather than specific situation–response links.

The sugar production task. The strategy participants in Experiment 1 did not describe a sufficient variety of rules to allow learning in the rule-based model. Therefore, we were inspired by Sanderson (1990), who identified four different types of rules that participants might use to control a dynamic system. First, the task could be modeled as a closed-loop system in which a participant always responds to the last computer output in a manner proportional to system error (error being target − output) about a pivot, that is, workforce_{t+1} = pivot + gain* (target − output_t) + noise. If the pivot is the target, and the gain is −0.5, then this strategy leads to optimal performance in the sugar production task: workforce_{t+1} = 0.5* (target + output_t). We called this the *optimal rule.* If (target + output_t) is odd, then workforce_{t+1} is randomly rounded up or down.

The second rule suggested by Sanderson (1990) is to use the workforce on the last trial as the pivot and use a gain of one: workforce_{t+1} = workforce_t + (target − output_t). This is similar to the rule described by the strategy participants in Experiment 1, namely, to increase work force according to the error. We called this the *workforce rule,* after its pivot.

The third rule suggested by Sanderson (1990) is to ignore the error, so that control is open loop. For example, workforce = 9, whatever the situation. We called this the *Open Loop 9 rule.* Workforce = 3 is the Open Loop 3 rule (if you had to give just one response, 9 would lead to relatively high scores and 3 to relatively low scores).

The final rule is to take workforce as the pivot but to set the gain to zero: $\text{workforce}_{t=1}$ = workforce_t + noise. We set the noise to be randomly −1, 0, or +1. Note that changing workforce by a small amount corresponded to the other rule mentioned by the strategy participants. We called this the *small-change rule.*

The relative number of tokens of the different rules needs to be adjusted to get the right level of learning. Thus, there are five free parameters, one for each of the optimal, workforce, Open Loop 9, Open Loop 3, and small-change rules. For each parameter value tried, we determined performance by averaging over 50 simulated subjects. Performance for the two learning blocks and the different concordance measures was fitted to participants' values. The results for one token of optimal and four of each of the others are shown in Tables 1 and 2. In Table 1, the model's results closely match those of participants. In contrast, as shown in Table 2, the model's performance on new situations was quite unlike that of participants: The model's performance on new situations was not substantially smaller than its performance on old situations. The Logan model, on the other hand, performed on new situations only as well as the explicit strategy predicted.

The results show that the rule-based model can produce the required pattern of concordances. The application of abstract rules leads to consistent responding in the same situations, and the relative weeding out of invalid rules can lead to greater consistency in old correct rather than incorrect situations. On the other hand, the tendency to give the same response in old correct rather than incorrect situations is the mechanism of lookup table learning, and so this tendency must be of a certain magnitude for a given lookup table to learn. As we now see, however, in the rule-based system, the size of the tendency can be masked to an arbitrary degree. The rule only needs to produce more correct than incorrect situations; it does not need to produce greater consistency to correct than to incorrect situations for learning to occur.

The rule-based model can give qualitatively different patterns of concordance, depending on parameter values. When optimal was given 1 token, workforce 15 tokens, and the rest 0 tokens, the numbers of trials loosely correct in Blocks 1 and 2 were 8.9 and 13.0, respectively, closely matching that of participants. For situations defined by workforce, concordance was reliably greater for situations previously given a correct response (.31) rather than an incorrect response (.19), consistent with the participants' data. However, for situations defined by a combination of workforce and sugar production, concordance was reliably greater for situations previously given an incorrect response (.54) rather than a correct response (.48).

When optimal was given one token, Open Loop 3 twelve tokens, and the rest zero tokens, the numbers of trials loosely correct in Blocks 1 and 2 were 6.8 and 17.1, respectively, closely matching that of participants. For situations defined by a combination of workforce and sugar production, concordance was reliably greater for situations previously given a correct response (.58) rather than an incorrect response (.31). However, for situations defined by workforce, concordance was reliably greater for situations previously given an incorrect (.34) rather than a correct response (.14).

In summary, the rule-based model can allow almost any pattern of concordances; lookup tables need a certain positive difference between concordance correct and incorrect. Participants' data fell in the range required by lookup table models. Note also that for the rule-based model, unlike the Logan model, there was no simple way of predicting the change in concordance for each participant with O (see Figure 2).

Person interaction task. Because strategy participants in Experiment 2 gave a wide range of rules, these could directly be used in the rule-based model. If all 21 rules were given one token, the chances of a good rule being selected were too low to allow adequate learning for either Person U or Person S. Thus, we assumed that for any given participant, there were only a few rules competing. Randomly selecting different subsets of rules for different participants did not produce good learning, and so more specific combinations of rules were explored.

It was relatively easy to produce good learning for Person S simply by combining a good rule with a bad rule. When the worst rule and the best rule were combined (one token each), the results for 100 simulated subjects are given in Tables 2 and 3. All seven means, the amount of learning, and the difference between correct and incorrect concordances were within the 90% confidence interval of participants' data. Note that the rule model, unlike the Logan model, can match participants' performance on both new and old situations for Person S.

For Person U, each of the good rules (i.e., a rule that, if applied by itself, would lead to greater than 5 trials correct out of 10 on average) was randomly combined with one of the bad rules (a rule that is not good). No combinations of rules were found that produced good learning for Person U. Tables 2 and 3 display the results for Person U by using the same rule combination as had been used for Person S. The rule-based model had problems learning Person U because, on the one hand, the consequence of a rule did not occur on the next trial but the trial after; on the other hand, which rule token was applied on 1 trial was independent of the last trial. Thus, there was no way of ensuring that tokens of successful rather than unsuccessful rules would be increased. The Logan model was equivalent in that which token was applied on one trial was also independent of the last trial. However, which response was given on one trial was not independent of the last trial: The assumed rules of participants allowed only small changes in responses. Because the Logan model stored specific responses (linked to situations) rather than rules, it could learn Person U.

In summary, the rule-based model performed in a complementary way to the lookup table model. The lookup table model could fit the data for Person U but not for Person S; the rule-based model could fit the data for Person S but not for Person U.

Discussion

For both the sugar production task and Person U, a one-parameter implementation of a lookup table could closely match participants' results. We believe that this finding is the most compelling support for the lookup table approach. Experimental data can always be interpreted in a number of ways, but preference must surely be given to the simplest way. An alternative rule-based model could not match the data so closely as the lookup table model, despite having more parameters.

Participants learning Person S outperformed the lookup table model on new situations. This result is interesting because previous authors have claimed that learning Person U can occur by an automatic process of linking situations to actions (Berry & Broadbent, 1988), but learning of Person S involves noticing and testing more general rules. Consistently, the results of the modeling show that participants' learning of Person U can be understood simply in terms of the storage of specific exemplars but there may be a need to postulate a more powerful induction process for Person S. A simple rule-based model could fit the level of learning for Person S but not for Person U. (A more sophisticated framework for looking at rule learning is provided by Holland, Holyoak, Nisbett, & Thagard, 1989; see Druhan & Mathews, 1989, for an application to a different learning paradigm.)

The modeling was guided by participants' data in how situations were represented. There was apparent inconsistency across sugar production and person interaction experiments in terms of whether the features of the preceding trial should be represented independently (sugar production) or configurally as combinations (person interaction). Medin and his colleagues (e.g., Medin, Altom, Edelson, & Freko, 1982; Medin & Schaffer, 1978) have shown that participants are sensitive to the configural properties of instances in classification tasks, so it would not be wise to propose a general learning mechanism that loses configural information. When learning the dynamic control tasks, participants may store configural information in instances (as suggested by the person interaction data) but only partial overlap of the instance with the current situation may be needed to activate the instance at least partially (as suggested by the sugar production data). Activation by partial overlap is perfectly consistent with theories such as Medin's; in fact, it is the use of partial overlap that allows exemplar theories to generalize in classification tasks. When the Logan model simulated controlling the person interaction task with individual features rather than with combinations stored, the means in Table 3 could be fitted (for Person U), so storing combinations or features was not crucial to the success of the model. Another possibility is that different cover tasks (sugar production or person interaction) induce participants to encode situations differently (featurally or configurally); this could be tested by redoing the experiments with the cover tasks switched.

General Discussion

This article has demonstrated how participants can learn to control dynamic systems simply by storing the appropriate responses for different situations. Experiments 1 and 2 showed that, for two different dynamic control tasks, participants tended to repeat the same response for situations in which they have previously been correct rather than incorrect. Further, participants performed at chance on new situations dissimilar to old ones. We can understand participants' learning on these tasks simply in terms of the storage of specific exemplars; there is no need to postulate a more powerful induction process.

Participants' performance could be matched by a lookup table that was based on Logan's (1988) instance theory. According to our application of Logan's theory, if a response was loosely correct, the situation would be stored with the response as an instance. All instances were encoded, including repetitions. Instances competed with each other and with the explicit strategy by racing to control performance. With one free parameter, representing the speed of the explicit strategy, participants' data from two experiments were separately fitted to within their 95% confidence intervals. The model could not fit participants' performance on a task with a highly transparent rule; in this case, and only this case, a rule-based model could match participants' performance.

So far in this article we have argued that a very simple lookup table model can account for learning nonsalient rules over a modest number of trials. Clearly, a more general account of human learning will add complexities to this picture. For example, the models reported in this article successfully accounted for participants' behavior by assuming previous knowledge could be represented by the equivalent of less than a dozen instances, although there must have been many thousands of instances that might have been construed as relevant (e.g., real person interactions in the person interaction task). Somehow people seemed able to limit their search to a defined set of instances. A similar phenomenon arises in the study of people's semantic memory for facts (Anderson, 1983): Usually, the more facts participants are taught about a fictional person, the slower they are to retrieve any one of them (the fan effect). However, participants can group by topic so that the fan effect only emerges within each topic. That is, participants appear to be able to restrict search to those instances that are most relevant. Future research could usefully explore how these restrictions are determined.

There is also evidence that participants who are not learning by hypothesis testing can nonetheless learn to selectively attend to dimensions that are relevant (Kruschke, 1993) and generalize around training instances in ways not well predicted by exemplar models (Dienes, 1992). It is likely that the modest number of trials and the simple nature of the situations in the dynamic control tasks (values along a dimension) allowed only simple types of generalization of knowledge (i.e., to values of the dimension numerically close to old stored values). Dynamic control tasks using situations with different types of structure may lead to more interesting generalizations around old situations.

This article has supported Broadbent et al.'s (1986) claim that learning the dynamic control tasks can simply occur by forming a lookup table. More generally, our article is consistent with accounts of human learning that stress the deployment of stored exemplars (e.g., Medin & Schaffer, 1978) and the use of fragmentary knowledge (Perruchet, 1994) in solving complex tasks.

References

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

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

Berry, D. C., & Broadbent, D. E. (1987). The combination of explicit and implicit learning processes. *Quarterly Journal of Experimental Psychology, 39,* 585–609.

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

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. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale, NJ: Erlbaum.

Dienes, Z. (1990). *Implicit concept formation.* Unpublished doctoral dissertation, University of Oxford.

Dienes, Z. (1992). Connectionist and memory array models of artificial grammar learning. *Cognitive Science, 16,* 41–79.

Dienes, Z., & Fahey, R. (1994). *The role of implicit memory in controlling a dynamic system.* Unpublished manuscript.

Druhan, B., & Mathews, R. (1989). THIYOS: A classifier system model of implicit knowledge of artificial grammars. *Proceedings of the 11th annual conference of the Cognitive Science Society* (pp. 66–73). Hillsdale, NJ: Erlbaum.

Estes, W. K. (1986). Memory storage and retrieval processes in category learning. *Journal of Experimental Psychology: General, 115,* 155–174.

Green, R., & Shanks, D. (1993). On the existence of independent explicit and implicit learning systems: An examination of some evidence. *Memory & Cognition, 21,* 304–317.

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.

Holland, J., Holyoak, K., Nisbett, R., & Thagard, P. (1989). *Induction: Processes of inference, learning, and discovery.* Cambridge, MA: MIT Press.

Kruschke, J. K. (1992). An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 22–44.

Kruschke, J. K. (1993). Human category learning: Implications for backpropagation models. *Connection Science, 5,* 3–36.

Logan, G. (1988). Towards an instance theory of automatization. *Psychological Review, 95,* 492–527.

Logan, G. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology, 22,* 1–35.

Logan, G. (1992). Shapes of reaction time distributions and shapes of learning curves: A test of the instance theory of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 883–914.

Marescaux, P. -J., Luc, F., & Karnas, G. (1989). Modes d'apprentissage selectif et nonselectif et connaissances acquies au control d'un processes: Evaluation d'un modele simule [Selective and nonselective modes of learning and acquiring knowledge in order to control systems: Evaluation of a simulation model]. *Cahiers de Psychologie Cognitive, 9,* 239–264.

Marescaux, P. -J., DeJean, K., & Karnas, G. (1990). Acquisition of specific or general knowledge at the control of a dynamic simulated system: An evaluation through a static situations questionnaire and a transfer control task (Tech. Rep. No. 2PR2GK). Brussels Free University, Belgium.

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

McGeorge, P., & Burton, A. (1989). The effects of concurrent verbalization on performance in a dynamic systems task. *British Journal of Psychology, 80,* 455–465.

Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8,* 37–50.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85,* 207–238.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 104–114.

Nosofsky, R. M., Clark, S. E., & Shin, H. J. (1989). Rules and exemplars in categorization, identification, and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 282–304.

Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 211–233.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101,* 53–79.

Perruchet, P. (1994). Learning from complex rule-governed environments: On the proper function of conscious and unconscious processes. In C. Umilra and M. Moscovitch (Eds), *Attention and Performance XV: Conscious and nonconscious information processing* (pp. 811–835). Cambridge, MA: MIT Press.

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

Sanderson, P. M. (1989). Verbalizable knowledge and skilled task performance: Association, dissociation, and mental models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 729–747.

Sanderson, P. M. (1990). *Implicit and explicit control of a dynamic task: Empirical and conceptual issues* (Tech. Rep. No. EPRL-90-02). Engineering Psychology Research Laboratory, University of Illinois.

Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioural and 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, 41,* 553–577.

(Appendix follows on next page)

Appendix

Simulating Explicit Knowledge

The strategy participants in Experiments 1 and 2 were given the instructions to the task and shown the starting situation. They were asked to describe as fully as possible what strategies they would use in performing the task. The purpose of collecting these data was to provide a means of simulating how participants would perform on the tasks in the absence of any learning. This simulation can then (a) be used to assess if the experimental participants did better on different types of situations than that predicted in the absence of any learning, and (b) be used in computational modeling of lookup tables and rule models both as a means of generating responses when the lookup table has no entries and as a means for generating rules for the rule model.

Sugar Production Task

All 6 strategy participants reported as their main strategy an algorithm that is sensitive to both the sugar production and workforce in the situation: If the sugar production was below target, they would increase the work force; if the sugar production was above target, they would decrease the work force. Four of the 6 participants made it clear that the changes in workforce would be in steps on the order of about 100. Four of the 6 participants reported that they would start with a workforce of 900; one of the other participants said 800.

These strategies could be modeled by the following assumptions: First, if the sugar production is below (above) target, then respond with a workforce that is different from the previous one by an amount of 0, +100, or +200 (0, −100, or −200). Second, for the very first trial, start with a workforce of 700, 800 (consistent with the last assumption), or 900. Finally, we added the assumption that if the sugar production is on target, then respond with a workforce that is different from the previous one by an amount of −100, 0, or +100 (with equal probability). We scored the first 10 trials of the original participants' performance for consistency with these assumptions: 86% (SD = 17%) of responses were in fact consistent, suggesting that the original participants did initially largely use the strategy that most new participants said they would use. Further, consistent with the stated strategy, the average magnitude of change in successive responses for participants was 0.97 (SD = 0.45); that is, participants did not like to change their responses by much more than one level of workforce.

Person Interaction Task

The 21 strategy participants gave a number of different types of responses. In general, participants said that they would respond in some region around polite or move above or below polite depending on Ellis's behavior. Each of the 21 strategies was implemented as a separate algorithm (if a given strategy was named more than once, it was separately implemented as many times as it was named). Each of the strategies was applied to all of the situations (combinations of Ellis's behavior and participants' previous behavior) actually experienced by each participant on the specific-situation task. The mean proportion correct (M = 0.29, SD = 0.03) was the same for Person U and Person S.

Similarly to the sugar production task, the average magnitude of change in successive responses was scored for the first five trials of each participant; the mean was 1.53 (SD = 0.96). That is, as for the sugar production task, participants did not like to change their response by much more than one or two steps. Also, participants tended to stay with the same response if Ellis was loosely on target. However, participants did not appear to change their response when Ellis was not on target in the same way as the sugar production task; there seemed to be little relation between the direction of change and whether Ellis was above or below target. Rather, the tendency seemed to be to simply move toward the polite region. When Ellis was loosely on target, the percentage of responses that were exactly the same was 52%; when Ellis was not on target, participants gave the same response only 23% of the time. When Ellis was not on target and participants gave a different response, they moved toward polite 74% of the time. This behavior was simulated in the following way. When Ellis was loosely on target, the simulation repeated the same response 50% of the time and gave a response one different (randomly up or down) the remaining 50% of the time. For the trials that Ellis was not on target, the simulation gave the same response 25% of the time. When Ellis was not on target and the simulation gave a different response, 75% of the time the simulation gave a response in a region around polite (polite plus-or-minus two steps inclusive; all responses within that region were equally likely); the rest of the time the simulation moved away from polite randomly by one or two steps. We applied this algorithm to all of the situations actually experienced by participants on the specific-situation task. The mean proportion correct was .33 (SD = .08) for Person S and .29 (SD = .06) for Person U. These means are not significantly different from the proportion correct obtained with the 21 individual strategies. We used this algorithm, which matched what participants actually did in the first few trials for simulating explicit knowledge in the specific-situation task and for simulating explicit knowledge in the Logan model that was previously described in this article. The use of an explicit algorithm that involves only small changes in successive responses is important for allowing the model to learn Person U. Tending to give the same response again if Ellis was on target allowed the model to reinforce the right response. Tending to stay around the polite region gave the model a high probability of trying the right response.

Received April 9, 1993
Revision received August 10, 1994
Accepted September 16, 1994 ■