UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Outcome Evaluation and Procedural Knowledge in Implicit Learning

Permalink
https://escholarship.org/uc/item/7758m01h

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 25(25)

ISSN
1069-7977

Authors
Fum, Danilo
Stocco, Andrea

Publication Date
2003

Peer reviewed

Outcome Evaluation and Procedural Knowledge in Implicit Learning

Danilo Fum (fum@units.it)
Department of Psychology, University of Trieste
Via S.Anastasio 12, I-34134, Trieste, Italy

Andrea Stocco (stocco@units.it)
Department of Psychology, University of Trieste
Via S.Anastasio 12, I-34134, Trieste, Italy

Abstract

Although implicit learning has been considered in recent years as a declarative memory phenomenon, we show that a procedural model can better elucidate some intriguing and unexpected data deriving from experiments carried out with Sugar Factory (Berry & Broadbent, 1984), one of the most popular paradigms in this area, and can account for other phenomena reported in the literature that are at odds with current explanations. The core of the model resides in its adaptive mechanism of action selection that is related to outcome evaluation. We derived two critical predictions from the model concerning: (a) the role of situational factors, and (b) the effect of a change in the criterion adopted by participants to evaluate the outcome of their actions. We tested both predictions with an experiment whose results corroborated the model. We conclude the paper by emphasizing the role played by procedural knowledge and evaluation mechanisms in explaining some implicit learning phenomena.

Introduction

One of the current hot topics in cognitive science is represented by implicit learning, that is the ability to acquire new knowledge even when apparently unaware of it, or unable to express it. Research on implicit learning, has focussed mostly on three phenomena. First, people that in a learning phase have experienced a set of grammatical strings generated from artificial grammar automata are surprisingly good in the test phase at discriminating new grammatical strings from ungrammatical ones (Reber, 1967). Second, participants seem able to abstract the structure of temporal event series, as shown, for instance, by an above-chance ability to predict the successive events (Kushner, Cleeremans & Reber, 1991). Third, participants improve their capacity to control a dynamic system, albeit being unable to report significant information on how the system worked (Berry & Broadbent, 1984).

In all these domains implicit learning was first explained by supposing the existence of an unconscious and autonomous learning system capable of extracting knowledge in general, abstract, and transferable rule-like form. Starting from the 90s, however, it has been suggested that many results could be explained by supposing that participants acquire concrete, fragmentary, declarative knowledge of the task represented through instances. Instance-based knowledge has indeed been shown to underlie or significantly affect performance on artificial grammar learning (Perruchet & Pacteau, 1990), sequence prediction (Perruchet, 1994) and dynamic system control tasks (Dienes & Fahey, 1995; Lebiere, Wallach & Taatgen, 1998, Taatgen & Wallach, 2002).

In this paper we will focus on a widely adopted dynamic system control task known as Sugar Factory (Berry & Broadbent, 1984). In this task, henceforth SF, participants are asked to imagine themselves in chief of a sugar-producing factory. Their goal is to reach and maintain the production of sugar (*P*) at a specified target level by modifying the amount of workers (*W*) allocated to the production. Unbeknown to participants, *P* and *W* are related by the equation:

$$P_t = 2W_t - P_{t-1} + e$$

where P_t is the quantity of sugar produced at trial *t*, just after the allocation of W_t workers, P_{t-1} represent the amount of sugar produced in the previous trial, and *e* is a random variable that can assume with equal probability one of the values in the set $\{-1, 0, +1\}$. The value for both *W* e *P* ranges from 1 to 12; resulting values of *P* less than 1 are simply set to 1, and values exceeding 12 are set to 12. For a more realistic interpretation, values of *W* are multiplied by 100 (hundreds of workers), and values of *P* by 1,000 (tons of sugar).

In controlling SF for two subsequent blocks of trials, participants generally exhibit an improvement in their performance from the first to the second phase demonstrating that some learning occurred, even if they were actually unaware of the relation between *W* and *P*.

Performance in SF has been generally measured by the total amount of hits, i.e., trials in which participants reached a production that was at most one unit above or below the target level. This loose criterion was introduced to take into account the role of the stochastic noise *e*, but is important to notice that participants were

not informed about the adopted criterion, and they were kept to believe that the only correct responses were those capable of obtaining the exact target production level.

Two main computational models have been put forward to explain the results from the Sugar Factory task: one developed by Dienes & Fahey (1995), and the other by Wallach and coworkers (Lebiere, Wallach & Taatgen, 1998; Taatgen & Wallach, 2002). The former is loosely based on Logan's theory of automatization (Logan, 1988), the latter is grounded on the ACT-R cognitive architecture (Anderson & Lebiere, 1998) and takes advantage of many of its features. Both models rely on the retrieval from memory of previously stored instances of the interaction with SF, although they differ in the way this information is represented, and in the retrieval mechanism adopted.

In the paper we illustrate a new procedural model that can better elucidate some intriguing and unexpected data obtained in previous experiments with SF (Fum & Stocco, 2003), and that can account for other phenomena reported in the literature that are at odds with current explanations. We present a new experiment testing two critical predictions of the model concerning (a) the role of situational factors, and (b) the effect of a change in the criterion adopted by participants to evaluate the outcome of their actions in which we obtained results that corroborated the model. We conclude the paper by emphasizing the role played by procedural knowledge and evaluation mechanisms in explaining implicit learning phenomena.

A Procedural Model for Sugar Factory

The model we developed for the SF task is grounded on the assumption that participants interacting with SF can exploit a set of very simple strategies in choosing the workforce value for the SF task. By combing the SF literature, and by looking at the interaction traces of the participants, we could identify some of these strategies:

Choose-Random: Choose randomly a value between 1 and 12.

Repeat-Choice: Repeat the value of W chosen in the previous interaction episode.

Stay-on-Hit: Whenever the previous W choice resulted in a success, repeat it. This strategy can be considered as a more selective variant of the previous one.

Pivot-Around-Target: Choose for W the value of the target, plus/minus one.

Jump-Up/Down: If the resulting P is lower than the target, increase the value of W; if it is higher, diminishes it. There exist several possible variants of this strategy. The one employed in the model was the *Jump-on-Middle,* i.e., choose as the new W a quantity that lies midway between the previous value and the upper/lower limit of the distribution of possible value (i.e., 1 when decreasing, and 12 when increasing).

Our procedural model is also built on top of ACT-R architecture (Anderson & Lebiere, 1998) but, differently from Wallach's model (Lebiere, Wallach & Taatgen, 1998; Taatgen & Wallach, 2002), it relies almost exclusively on the subsymbolic procedural learning mechanism provided by that architecture.

ACT-R is based on the assumption that the knowledge utilized by the cognitive system can be distinguished between declarative and procedural. The former encodes knowledge related to remembered facts, perceived events or sensorial input, and is represented through *chunks*, i.e., frame-like structures constituted by labeled slots with associated filler values. Procedural knowledge represents processes and skills used by the cognitive systems to pursue its goals. It is represented through production rules, or *productions*, whose conditions specify the patterns of declarative elements that are to be active for the production to apply. The action part of a production specifies some modifications to be brought to declarative chunks, or some actions to be performed in the environment. Differently from chunks, productions are supposed not to be conscious or reportable, and thus are implicit in nature. All the strategies encoded in our model have been implemented as ACT-R productions that are let compete for execution.

The process of selecting which production to execute among those that have their conditions satisfied is based in ACT-R on the concept of *expected utility*. The utility U_i of a production i is a noisy, continuously varying quantity defined by the formula $U_i = P_iG - C_i$, where P_i represents the probability that production i will reach the goal it is meant to achieve, G is the value of the goal itself, and C_i is the estimated cost of applying i. The higher the expected utility of a production, the higher the probability that the production will be chosen for execution, even if a randomly distributed noise prevents the mechanism from being completely deterministic.

Subsymbolic procedural learning is accomplished in ACT-R through the adjustment and refinement of the production parameters. The estimated value P_i of a production i, in particular, is given by the ratio between the number of successes obtained in the past by applying the production (i.e., the number of times the production has succeeded in reaching its goal) and the total number of times (i.e., the sum of the successes and failures) the production has been executed. After each production firing, the probability P_i is updated to reflect the statistical structure of environment, according to a Bayesian updating framework. Because probability P is initially not defined (and after the very first production firing its estimate will be 1 or 0 according to whether it was a success of a failure), it is common practice to start the learning process by providing each production with values that represent prior estimates of the production's successes and failures.

This learning schema proved to be adaptive, and it enabled to replicate some interesting phenomena like the capability of humans to choose an option proportionally to its probability of being correct (probability matching), or to increasingly use those problem solving operators

that were learned to be the most successful ones (Anderson & Lebiere, 1998).

Whenever the subsymbolic ACT-R procedural learning mechanism is activated, successful applications of a production increase its expected utility, and therefore augment the likelihood that the production will be chosen for execution in a next occasion.

Successful Applications of the Model

Our model allows to explain the common learning effect obtained in SF as well as to better elucidate some phenomena reported in the literature that are at odds with current explanations, and eventually to account for the puzzling results we found Fum & Stocco (2003).

It is important to note that almost all of the experiments carried out with this paradigm adopted 9,000 tons of sugar as target value for the production. In our first experiment we contrasted the performance obtained in the standard condition with that deriving from a target of 3,000 tons, and we found that the latter condition made the task of controlling the SF system significantly easier (see first two rows of Table 1). The finding that the target value affects performance is difficult to explain for the models that rely on the retrieval of stored interaction instances. A priori all the target levels have an equal likelihood to be obtained, and there is no sensible reason, in fact, why remembering the W values leading to a target P of 3,000 tons should be easier than remembering the values associated with 9,000 tons.

Even more difficult to account are the results obtained in Fum & Stocco's (2003) second experiment, concerning the effects of a change in the value of the target production between the first and the second phase. In the experiment, one group switched from a starting target value of 3,000 tons to a value of 9,000 tons in the second phase. Another group experienced the same target levels but in the reversed order. A general effect of the phase was found, providing evidence for a performance improvement together with an interesting Target x Phase interaction (see second two rows of Table 1). According to the instance based models, a change in the target value should destroy the possibility of taking advantage in the second phase from previous experience because almost all the stored instances refer to episodes that are useless in the new condition. As a consequence, no positive transfer between the two phases should be obtained. These predictions were falsified by our experiment.

Table 1: Results from Fum & Stocco (2003) compared with model's predictions

Condition	First Phase		Second Phase	
	Particip.	Model	Particip.	Model
3000-3000	10.57	9.35	13.18	12.30
9000-9000	7.26	7.02	9.26	8.83
3000-9000	9.39	9.32	8.70	9.04
9000-3000	6.86	6.96	12.84	12.12

An important preliminary insight deriving from our model is that, because of the interaction between the procedural strategies and the system dynamics (e.g., the stochastic noise, or the way it deals with the out-of-range values), some production target levels are inherently easier to obtain than others. Figure 1 illustrates this idea by reporting the results of a simulation carried out with 2,500 runs of the model for each target level between 2,000 and 11,000 tons. It can be shown that each production rule, implementing a separate solution strategy, has a different probability of success (i.e. of leading to the target) in the different conditions. For instance, the *Repeat-Choice* rule has an averaged (between phases) success probability of 15.18% in the 3,000 tons condition vs. 11.32% in the 9,000 one. Analogously, the *Jump-on-Middle* (up and down) rule has success probabilities of 2.44% vs. 1.22% in the 3,000 and 9,000 tons conditions, respectively. This is a first important fact that should be taken into account.

If we run the model with the ACT-R procedural learning mechanism off and let the productions compete for execution led only by chance and stochastic noise, the result is affected only by situational factors, i.e. by the target. If we change the target level between the two phases, the model produces the results illustrated on the left panel of Figure 2.

Another important insight deriving from the model is that not only the success probability of a production is different in different target conditions, but also that the success probabilities of *different* productions in the *same* target condition are different. In other words, some productions in the SF domain are inherently better that others, independently from the condition in which they are applied. By looking at the execution traces, for instance, it is possible to realize that *Stay-on-Hit* is always better than *Jump-on-Middle,* or that *Repeat-Choice* is always better than *Choose-Random*.

At the very beginning, all productions are considered as equal by the model that does not use any a priori estimates of success for them. When the ACT-R subsymbolic production learning mechanism is activated, however, it increasingly tends to prefer the most successful productions, and the more frequent firing of the these productions results in the performance improvement that is obtained from the first to the second experiment phase.
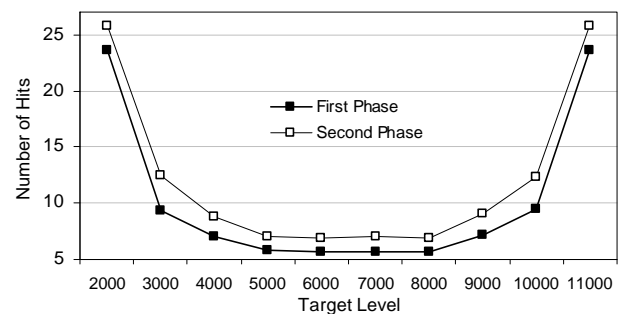

Figure 1. Model's performance for different target levels.

Just to give an example, the percentage of firings of the *Choose-Random* production decreases in the 3,000 tons condition, from the 16.63% in the first phase to the 11.54% in the second phase while the more successful *Repeat-Choice* has an increase from the 29.07% in the first phase to the 42.37% in the second.

The results of the first experiment are thus easily explained: the difference between the 3,000-3000 and the 9,000-9,000 conditions is caused by the fact that the same production has a different success rate in the two conditions; the learning effect is cause by the fact that the ACT-R subsymbolic learning mechanism stimulates the firing of the most successful productions while discouraging the least successful ones. The question of what is learned by participants in the SF task has therefore a simple answer: they are increasingly led to use those productions that are more likely to obtain the goal.

The results obtained in the second experiment derive from the combined influence of situational (Target) and learning (Phase) factors. Figure 2 illustrates the idea. Without learning, only situational factors are involved. In this case the model predicts a completely symmetrical situation for the 3,000-9,000 and the 9,000-3,000 conditions. When the learning mechanism is activated, it interacts with the situational factors originating the results obtained in the second experiment.

By using only the pure subsymbolic learning mechanism of ACT-R, without any need for parameter fitting or a priori estimates, the model is capable of explaining our experimental results.

Further Explanations

Besides explaining previous results from the SF task and the new findings from our experiments, our model can in fact clarify some data from Marescaux (1997) that have been considered as supporting the instance-based viewpoint. Marescaux reports evidence from an unpublished experiment, where participants first interacted with the system and then were asked to complete a written questionnaire. The items in the questionnaire were constituted by hypothetical interaction episodes, and participants had to figure out the proper value of *W* in order to reach (a) the same production target level adopted in the interaction phase (i.e., the standard 9,000 tons), or (b) a new target level (i.e., 6,000 tons). Marescaux found that performance was lower for the (b) items, and took this finding as indicative of lack of generalization, and therefore corroborating instance-based learning. Leaving aside the questionable choice of adopting the written modality to test the participants' knowledge, Marescaux did not take into account the differential effect of the target levels that we found in our experiment. As we will see below, our model predicts a difference in performance exactly in the same direction.

Another widely known, but hitherto never modeled result, comes from a study of Berry (1991). The author showed that material interaction with the system is necessary for learning to occur. Berry (1991, Experiment 1) divided the participants into two groups: the Observation group simply observed the experimenter controlling SF for the first block of trials, and directly interacted with the simulator in the subsequent phase. A Control group interacted with the system in both periods.

Participants in the Observation condition were impaired in the second phase in which they had to interact with the system: their performance was just at the same level of that provided by Control participants in the first phase, as if they had never encountered the system before. Moreover, performance during the interaction phase remained unchanged even when participants were previously given *two* observation phases (Berry, 1991, Experiment 3).

It is somehow possible to account for such an effect even in an instance-based framework. However, the explanation given by our model is straightforward: during observation, participants had no change to experience the relative success rate of each production. Lack of experience corresponds to the inhibition of subsymbolic learning in the ACT-R architecture.
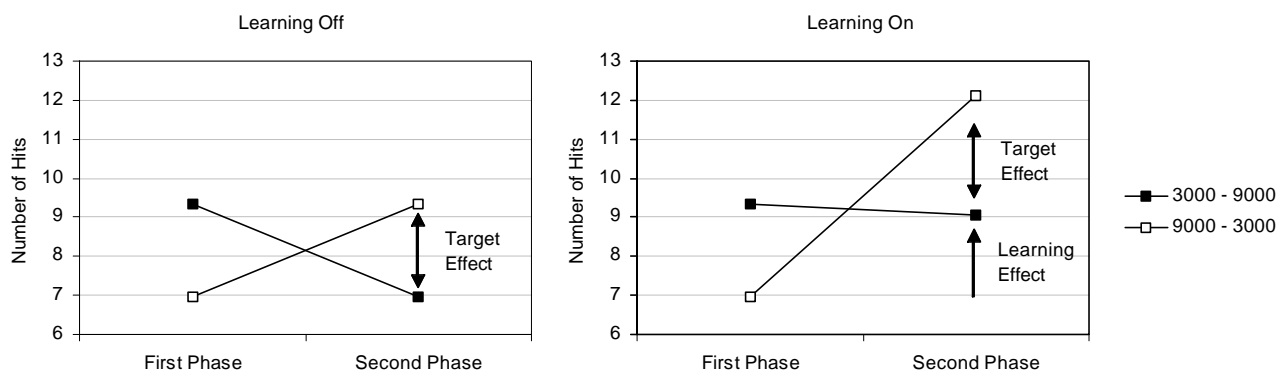

Figure 2: Model's predictions for the second experiment of Fum & Stocco (2003). Predicted performance (right) is the sum of the effect of target level (left) and the acquisition of transferable knowledge.

A New Experiment

Beyond providing an explanation for previous results, our model, being a zero-parameter model, allows to make experimentally testable predictions. Here we take into account here two of them. The first concerns the possible effect of a change in the criterion adopted by participants to evaluate the outcome of their actions. The second concerns, again, the role of situational factors.

As it has been previously mentioned, due to existence of random noise that makes the task of controlling the sugar production quite difficult, it is common practice in SF to consider a participant successful in reaching the target whenever the production falls within an interval of ±1 unit from it. It is interesting to wonder what would happen if we make the participants aware of this scoring criterion by saying that, when the target is set for instance to 6,000 tons, their goal is to keep the production level between 5,000 and 7,000 tons. According to our model, making the scoring criterion explicit would cause, through a rich-gets-richer effect, a more frequent application of the most successful productions leading to an overall performance improvement. It is important to emphasize that this result should be obtained not by changing the scoring criterion adopted by the experimenter but by changing the criterion used by participants to evaluate their own performance. The experiment here described was designed just to test such critical prediction.

It is worth noticing that instance-based approaches cannot make clear predictions about the expected results: performance could be positively affected by the opportunity of storing a larger amount of successful instances, but, on the other hand, such instances have a greater probability of leading to an unsuccessful result.[1] Moreover, a larger amount of instances is likely to negatively affect the retrieval by augmenting memory workload.

In the experiment we decided to adopt a target production level of 6,000 tons in order to test whether this situational factor would affect the performance in the way predicted by our model. We already found that the 3,000 tons condition was significantly easier than the 9,000 one, (Fum & Stocco, 2003, Experiment 1), but this result would be consistent with other possible functions, e.g. by assuming that performance would be inversely related to the target level.

Our hypotheses were therefore that: (a) adopting a loose scoring criterion should improve the participants' performance; and (b) the performance with the new production level would comply to the predictions of our model, i.e. should be lower than that obtained with both the previous target levels.

Method

Participants Participants were 93 students (29 males and 64 females) aged 19 to 30 (median 19) all recruited from a General Psychology course.

Design The main independent variable was the scoring criterion. Participants in the Narrow group were instructed to maintain the sugar production at exactly 6,000 tons, while those in the Broad group were told to keep it between 5,000 and 7,000 tons. The main dependent variable was, as usual, the number of hits. The experiment adopted a 2 (Criterion: Narrow vs. Broad) x 2 (Phase: First vs. Second) mixed design, with Criterion as between subjects and session Phase as within subjects factors.

Procedure All participants were tested individually in single sessions comprising two blocks of 40 trials each. Participants interacted with SF through a computer program, written in Python and running on a Dell PC provided with a 15' LCD flat screen. The GUI of the program presented, for each trial, the current amounts of workers (W) and production (P). Participants could change the former by pressing one of keys F1…F12 on the computer keyboard, each key corrisponding to one of the possible values for W. After a 1s delay the program showed the resulting value for P, and was ready again for a new input value. Each time a participant reached the correct production level (6,000 for the Narrow, a value between 5,000 and 7,000 for the Broad condition), the system played a beep through the speakers. Initial values for the first trial for both W and P were randomized.

Results and Discussion

Table 2 reports both the results from the experiment and our model predictions (obtained through 2,500 runs).

Table 2: Experiment results and model's predictions

Condition	First Phase		Second Phase	
	Particip.	Model	Particip.	Model
Narrow	6.70	5.85	7.68	6.93
Broad	7.52	7.55	9.96	8.50

An ANOVA computed on participants data revealed as significant the main effects of Criterion ($F(1,91)=5.97$, $MSE=111.36$, $p=.016$) and of the Phase ($F(1,91)=12.26$, $MSE=134.22$, $p=.007$), but not their interaction. As predicted by the model, the performance in the Narrow condition was significantly lower than in the Broad condition, while both groups improved their ability of controlling the system from the first to the second phase.

In order to test our second hypothesis concerning the existence of the U-shaped curve relating target level and performance (Figure 1), we carried an ANOVA comparing the performance of the Narrow group with that

---

[1] In fact, it is possible to demonstrate that, due to the noise parameter e, the probability for a retrieved instance to lead to success falls from 0.78 in the narrow condition to 0.60 in the broad one.

obtained by the 9,000-9,000 group in first experiment of Fum & Stocco (2003), that was run under comparable conditions. We found, again, a significant effect of phase ($F(1,88)=8.28$, $MSE=96.8$, $p=.005$), but the effect of target level (6,000 vs. 9,000) resulted marginally significant ($F(1,88)=3.59$, $MSE=50.88$, $p=.06$). However, we decided to accept this result as a corroboration for our hypothesis, given the fact that the predicted magnitude of the effect was small, and that a one-tailed t test carried out on the results of the second phase yielded a significant difference ($t(88) = 1.94$, $p=.028$).

Conclusions

We developed a model for the SF dynamic system control task that explains far more results than any other instance-based model and that relies on procedural knowledge unconsciously selected on the basis of its expected utility. We provided further evidence for the model by testing two critical predictions: the existence of a U-shaped performance curve, and the effect of broadening the range of perceived successful results. Our predictions turned out to be correct.

It is worth to emphasize that we do not assume that participants are unable to store any memory of their interactions with the system nor that instance-based models are useless to explain human cognition. We simply state that, contrary to what is posited by these models, instances do not play a significant functional role in learning to control a simple, but not trivial, dynamic system like SF, at least for the time scale covered by the experiments. As a matter of fact, both Dienes & Fahey (1998), and Schoppek (2002) reported no evidence for instance-based behavior even when participants were able to show some kind of declarative knowledge. Finally, instance-based explanations are somewhat at odds with results by Squire & Frambach (1990), who found amnesic patients performing at the same level of normal control participants when interacting with SF.

We believe that research on implicit learning has somehow neglected the role of outcome evaluation mechanisms that are essentially procedural in nature and adaptive in function. By putting these mechanism again into focus we could hope to bridge again research on implicit learning with evidence from animal cognition and neuropsychological studies.

References

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

Berry, D. C. (1991), The role of action in implicit learning. *The Quarterly Journal of Experimental Psychology, 43A(4)*, 881-906.

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

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

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

Fum, D., & Stocco, A. (2003). Instance vs. rule-based learning in controlling a dynamic system. In F. Detje, D. Dörner & H. Schaub (Eds.), *Proceedings of the 5th International Conference on Cognitive Modelling.* Bamberg, Germany: Universitäts-Verlag, 105-110.

Kushner, M., Cleeremans A., & Reber, A. S. (1991) Implicit detection of event interdependecies and a PDP model of the process. In K. J. Hammond & D. Gentner (Eds.), *Proceedings of the 13th Annual Conference of the Cognitive Science Society*, 215-220.

Lebiere, C., Wallach, D., & Taatgen N. A. (1998). Implicit and explicit learning in ACT-R. In F. Ritter & R. Young (Eds.), *Proceedings of the Second European Conference on Cognitive Modelling.* Nottingham, UK: Nottingham University Press, 183-193.

Logan, G. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492-528.

Marescaux, P.-J. (1997). Can dynamic control task knowledge be communicated? *Psychologica Belgica, 37*, 51-68.

Perruchet, P. (1994). Learning from complex rule-governed environments: On the proper functions of nonconscious and conscious processes. In C. Umiltà e M. Moscovitch (Eds.), *Attention and Performance XV. Conscious and nonconscious information processing.* Cambridge, MA: MIT Press.

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General, 119*, 264-275.

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

Schoppek, W. (2002). Stochastic independence between recognition and completion of spatial patterns as a function of causal interpretation. In W. D. Gray & C. Schunn (Eds.), *Proceedings of the 24th Annual Conference of the Cognitive Science Society*, Mahwah, NJ: Erlbaum, pp. 804-809.

Squire, L. R., & Frambach, M. (1990). Cognitive skill learning in amnesia. *Psychobiology, 18*, 109-117.

Taatgen, N.A., & Wallach, D. (2002). Whether skill acquisition is rule or instance based is determined by the structure of the task. *Cognitive Science Quarterly, 2,* 163-204.