Kennedy, W.G. and Patterson, R.E. (2012) Modeling Intuitive Decision Making in ACT-R. In Proceedings of the 11th International Conference on Cognitive Modeling (ICCM). Berlin, Germany. 12-15 April 2012.

Modeling Intuitive Decision Making in ACT-R

William G. Kennedy (wkennedy@gmu.edu)
Krasnow Institute for Advanced Study, 4400 University Drive
Fairfax, VA 22030 USA

Robert E. Patterson (Robert.Patterson@wpafb.af.mil)
Air Force Research Laboratory,
Wright-Patterson AFB, OH 45433 USA

Abstract

One mode of human decision-making is considered intuitive, i.e., unconscious situational pattern recognition. Implicit statistical learning, which involves the sampling of invariances from the environment and is known to involve procedural (i.e., non-declarative) memory, has been shown to be a foundation of this mode of decision making. We present an ACT-R model of implicit learning whose implementation entailed a declarative memory-based learner of the classification of example strings of an artificial grammar. The model performed very well when compared to humans. The fact that the simulation of implicit learning could not be implemented in a straightforward way via a non-declarative memory approach, but rather required a declarative memory-based implementation, suggests that the conceptualization of procedural memory in the ACT-R framework may need to be expanded to include abstract representations of statistical regularities. Our approach to the development and testing of models in ACT-R can be used to predict the development of intuitive decision-making in humans.

Keywords: implicit learning; cognitive models; unconscious learning; ACT-R theory.

Introduction

The vast majority of cognitive models discussed at this conference are models of rational or analytical cognition. The architectures used are primarily ACT-R (R for rational) or Soar, and the best papers compare a computational implementation of a theory, i.e. a model, to human behavior observed in careful laboratory experiments. The authors of these papers then claim that the model under consideration is a plausible theory of the cognitive process behind the observed behavior. This approach is advancing the understanding of cognitive processes. However, modeling consciously rational behavior addresses only part of human cognition and it ignores the ubiquitous influence that implicit processing and intuitive decision making has on human behavior.

In the dual-process framework of reasoning and decision making (e.g., Evans, 2008; Patterson, Pierce, Bell, Andrews & Winterbottom, 2009; Sloman, 1996), one mode of decision making is called intuitive. Intuitive decision making refers to implicit situational pattern recognition that is not thought to involve symbolic rules (Klein, 1998). The other mode of decision making is called analytical, which is generally accepted to entail symbolic rules. Intuitive decision making, which falls under the rubric of 'System 1' processing in this literature, is typically described as unconscious, fast, and effortless decision making. Analytical decision making, which falls under the rubric of 'System 2' processing, is described as conscious, rational, slow, and effortful. Evans (2008) provides a review of the evidence supporting the Dual-Process theory. Analytical decision making is relatively simple to study because it is easy to create tasks for testing and recording behavior during rational performance. Intuitive decision making, on the other hand, is difficult to study because it is hard to artificially create environmental patterns with sufficient fidelity to study situational pattern recognition.

Recently, Patterson and colleagues (Boydstun, Patterson, Pierce, Park & Tripp, 2011; Covas-Smith, Patterson, Pierce, Cooke & Homa, 2011; Patterson et al., 2009) have investigated the development of intuitive decision making in a simulated real-world environment. These authors had human participants experience simulated flight over a synthetic terrain with a sequences of objects (e.g., house; vehicle) positioned on the terrain along the flight path. Each object sequence was derived from paths taken through a finite-state algorithm, which defined a grammar for constructing the content of the scene. The use of a finite-state grammar for creating object sequences was analogous to the way in which finite-state grammars have been used for studying the implicit learning of artificial letter strings (e.g., Reber, 1967). Patterson and colleagues tested the conjecture that implicit learning (Cleeremans, Destrebecqz & Boyer, 1998; Perrachet & Pacton, 2006) could be one way in which intuitive decision making is developed.

Patterson and colleagues found that naive participants could implicitly learn the object sequences quite easily. Moreover, the implicit learning of the sequences provided a foundation for intuitive decision making about the underlying structure of the sequences: following training with the artificial object sequences, the participants were successful in recognizing novel sequences taken from the same grammar during test. That is, the human participants implicitly learned to recognize situational patterns.

The ACT-R architecture (Anderson, 2007; Anderson, et al., 2004) has been used before to model implicit learning. In particular, Wallach and Lebiere (2003) reviewed the theoretical approaches to implicit learning and observed "a major shortcoming of these models is their failure to also account for explicit learning and for the difference between implicit and explicit learning" (pg 217). They then presented

Kennedy, W.G. and Patterson, R.E. (2012) Modeling Intuitive Decision Making in ACT-R. In Proceedings of the 11th International Conference on Cognitive Modeling (ICCM). Berlin, Germany. 12-15 April 2012.

ACT-R models of two well-known implicit and explicit learning tasks and specifically linked explicit learning with the learning of declarative chunks and implicit learning with ACT-R's sub-symbolic learning of the activations of those chunks. We will use the same approach here, using the sub-symbolic representation associated with declarative memory as the basis of our model of intuitive decision making.

This paper presents an approach to studying intuitive decision making (i.e., System 1 cognition) that exposes the cognitive process to computational modeling and experimental testing of theories implemented as models. An ACT-R model was developed and compared to human subject data on an intuitive decision-making task used by Patterson, Pierce, Boydstun, Park, Shannan, Tripp and Bell (submitted).

Patterson et al. Study

Patterson, Pierce, Boydstun, Park, Shannon, Tripp, and Bell (submitted) investigated whether implicit learning can be a process by which intuitive decision making is acquired. One form of implicit learning entails the learning of spatial or temporal patterns without full awareness of what is learned (Cleeremans, Destrebecqz & Boyer, 1998; Perrachet & Pacton, 2006). Implicit learning is likely to be a key process by which individuals learn situational patterns on which intuitive decisions are based (e.g., Patterson et al., 2009).

Patterson et al. extended the classic paradigm by Reber (1967) used for studying implicit learning, which entailed the learning of a synthetic grammar produced by a finite state algorithm that generated artificial letter strings. Patterson et al. instead investigated the implicit learning of passively viewed, structured object sequences presented in a simulated real-world immersive environment used for simulating locomotion (Figure 1). In doing so, they used a finite state algorithm that created an artificial grammar for generating the object sequences and thus the content of the environment (Figure 2). For comparison, Patterson et al. also investigated the implicit learning of *memorized* static letter strings presented on a flat display, as has been done in the past (Reber, 1967) (See Figure 3).

The finite-state diagram of the grammar shown in Figure 3 has also been used in many other studies (Cleeremans, Destrebecqz, & Boyer, 1998; Matthews, et al., 1989; Perrachet & Pacton, 2006). It produces 44 valid structured strings of length 8 or shorter. Participants are trained by being presented a series of example strings from the grammar and are then tested by being asked if a test string is legal or not.

During training, Patterson et al. had human participants (1) passively view structured sequences of objects presented on a dynamic terrain seen in perspective view (the 'immersive display' condition), or (2) memorize structured strings of letters presented on a static flat display. Following training, participants were tested for implicit learning by making intuitive pattern-recognition judgments of novel structured object sequences or letter strings versus random sequences or strings.

By training participants on the structured object sequences or letter strings, and then testing recognition of structured versus random sequences or strings, the participants performed an *'anomaly recognition'* test. The random sequences or strings effectively served as an anomaly to be recognized because the participants were never trained on random sequences or strings.

Figure 1. Photograph showing the simulated real-world environment. The scene underwent expansive optic flow motion, which simulated passive movement by the participant in the forward direction toward the horizon.

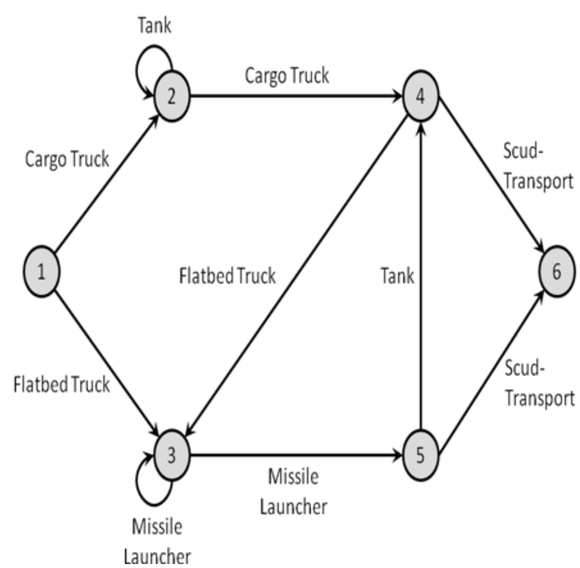

Figure 2. Depiction of finite state algorithm that defined the grammar employed for generating the structured sequences of objects used in Patterson, Pierce, Boydstun, Park, Shannon, Tripp, and Bell (submitted).

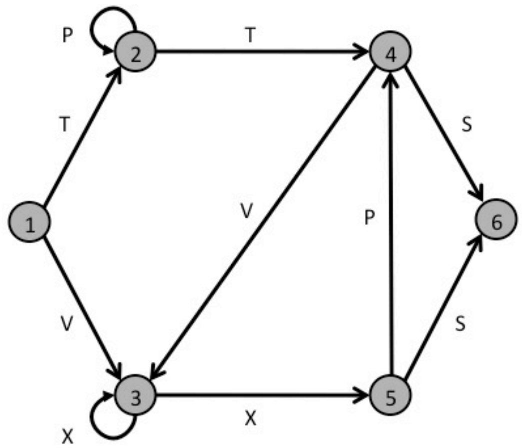

Figure 3. Finite state algorithm that defined a grammar of letter strings with the same structure as the sequences of objects of Figure 2. (From Reber, 1967.)

Results. Figure 4 depicts results obtained for the simulated real-world environment and for the static flat display, as reported by Patterson et al. Passive viewing of object sequences (third bar from the left) resulted in an average accuracy of intuitive decision making that was equivalent to the average recognition performance that was obtained when letter strings presented on a flat display were memorized (middle bar). An a-priori *t*-test showed that the difference between these two conditions was not significant, $t(14) = 0.4$, $p = 0.7$. The two training conditions were significantly higher than the no-training control condition (left bar), which was at chance-level performance.

During debriefing, the human participants had trouble explicitly verbalizing all of what they had learned during training and that a number of their decisions made during testing were from a feeling "in the gut". Thus, the training methods produced a significant level of implicit learning that was a foundation for the pattern-recognition-based (intuitive) decision making.

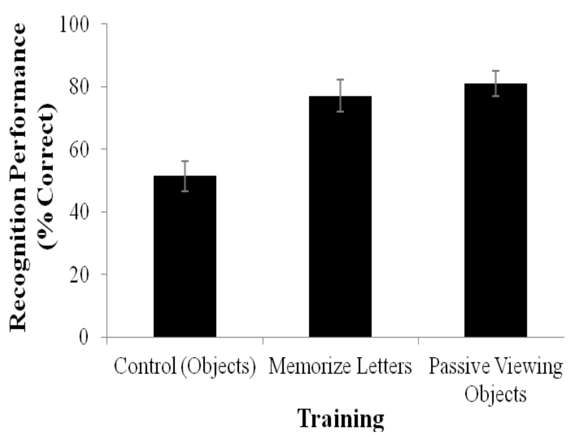

Figure 4. Accuracy of intuitive decision making, as measured by recognition performance during the test phase. During training, the object sequences were passively viewed, and the letter strings were memorized. The control group, which involved object sequences, entailed no training. Each data point is the mean of eight human participants; error bars depict ± 1 standard error of the mean (SEM). The data was from one portion of Patterson et al. (submitted).

ACT-R Model

Our model is based on ACT-R (Anderson, 2007; Anderson, et al., 2004). ACT-R is a rule-based architecture representing cognitive processes symbolically and sub-symbolically. Its declarative memory holds chunks of declarative facts with an activation level based on the recency and frequency of use. IF-THEN rules are held in a long-term procedural memory. ACT-R models can learn by adjusting the activation of accumulated declarative chunks, by adjusting the relative measure of rules, or by combining sequential rules into new rules.

Statistical learning is sometimes modeled as the tuning of the relative measures of rules and that approach could have been used here. However, to study different strategies believed to be used in the implicit learning of abstract grammars, the model developed here uses the activation of declarative chunks of memory, with each chunk representing a bigram of letters. (The task modeled was the letter string version of the implicit learning task. An analogous model would apply to the object sequence version of the task.)

Our ACT-R model uses both the declarative and procedural modules to passively learn and then respond to this task. The rules are fixed during the run of the model. Declarative memory chunks are added based on experience during training and are recalled to make the valid/invalid evaluation during testing.

During training, rules direct the system to read the string letter by letter, left to right. The system then forms declarative chunks and saves them as an internal representation of the grammar based on observed training strings. Each declarative fact is a representation of observed bigrams indicating which letter was seen before another, i.e., a first letter and a predicted second letter.

During testing, a representation of the intuitive decision-making process determined whether all the bigrams in a test stimulus have been seen before. To respond, the system reads the string left to right and attempts to recall bigrams predicting the next letter. A successful retrieval increases the activation of that declarative chunk and the ACT-R architecture returns the one declarative chunk for an attempted recall operation. If the retrieved bigram does not match the second letter, a second retrieval is attempted using both the first and second letter. If successfully recalled, the evaluation continues. If not successfully recalled, the test string is evaluated as invalid.

This approach implements a form of predictive and evaluative behavior. Other approaches that could have been studied include recalling the first few letters of a string (primacy), recalling the last few letters of a string (recency), recalling both the first few and the last few letters, or simply deciding by noting whether the number of pairs of letters recallable was above a threshold. We could also have tested trigram representations, or other representations. The model discussed here based its decisions on whether predictive bigrams were recallable.

Replicating Human Subject Experiments

To compare the ACT-R model to data from Patterson et al (submitted), we replicated the passive training and testing protocols for strings of letters generated by the artificial grammar shown in Figure 3. The human participants and the model were trained and tested the same way.

Training used 18 unique strings drawn from the 44 valid strings of length 8 or less. Training was organized as six blocks of three unique valid strings, which were presented 16 times with each string presented for 5 seconds and a blank screen shown for 0.6 seconds between strings.

The testing process also replicated that used with human participants. The system presented 88 strings, 22 valid strings that were not used for training, and 22 foils, each presented twice in a random order. The foils used the same letters, but in a random order and were of length 6, 7, or 8.

For this model, only default ACT-R parameters were used, except the retrieval threshold (:rt) for declarative memory and the activation noise variable (:ans).

Model's Performance

The model reports whether a test string is valid or not for each of the 88 trials. Response accuracy was our only performance measure because the response time for the human participants was not collected in the Patterson et al. study. Figure 5 shows the performance results of the model together with the human participants' test performance as reported by Patterson et al. (submitted).

The plotted results for the human participants are the means and standard error for 8 individuals as described in Patterson et al. (submitted). The plotted model results are the means and standard error for 30 runs of the model varying the retrieval threshold (:rt) parameter from -1.5 to +3.0. The noise parameter, :ans, was 0.1. The :rt parameter sets the threshold for successful retrievals from memory base on the activation level of chunks of memory. Lower values of the parameter allow the model retrieve more instances and higher values restrict retrievals to the most activated memories.

A two-tailed, equal variance t-test found that the difference in mean accuracy between the humans and the model was not significant for the 30 runs with :rt =2.0 ($t(29) = 15.74$, $p < 0.001$). Note, however, that the mean accuracy for the humans and model were very similar and within a few percentage points of one another for several values of :rt. Therefore, we are encouraged by the closeness of the model behavior to that of humans on this implicit learning task.

Because we are using a computational model, it is relatively easy to collect additional information on the performance of the model and compare it to human data. Table 1 provides the human and model performance in more detail than the summary information shown in Figure 5.

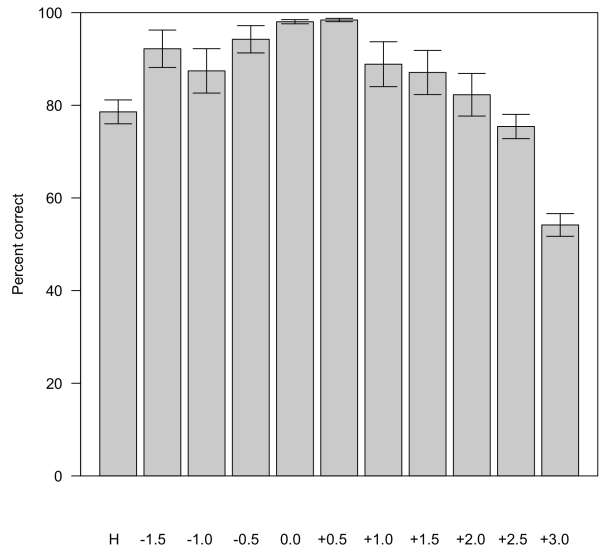

Figure 5. Model and Human Performance. The human performance (H column) and the model's performance for retrieval threshold parameter (:rt) from -1.5 to +3.0 on strings of letters learned by passive viewing. Error bars in this figure depict ±1 standard error of the mean (SEM)

Table 1: Human and Model Performance in Detail.

Trial/Response Type	Human	Model :rt=2.0
Hits	33/44	34/44
Correct Rejections	36/44	39/44
Misses	11/44	10/44
False Alarms	8/44	5/44

In the table, "Hits" refers to the number of grammatical strings that were detected as grammatical, "Correct Rejections" refers to the number of ungrammatical strings that were detected as ungrammatical, "Misses" refers to the number of grammatical strings that were incorrectly seen as ungrammatical, and "False Alarms" refers to the number of ungrammatical strings that were incorrectly seen as grammatical. In a signal detection analysis, one can compare the hit rate and false alarm rate to get an estimate of the level of criterion that is being used for detection: a high hit rate coupled with a high false alarm rate would suggest that the detection system is overly responsive and that its actual sensitivity is not particularly high. However, a high hit rate coupled with a low false alarm rate would suggest that the system is responding selectively to a signal and that its sensitivity is high.

While the data shown in Table 1 is insufficient for a formal signal detection analysis, it is clear that both the human data and the model data reveal a very similar pattern of high hit rates coupled with low false alarm rates. This would suggest that both systems, human and model, possess a similar high level of sensitivity for implicit learning.

Discussion

The results show that the ACT-R model can replicate the human performance when the retrieval threshold parameter is tuned to account for the training protocol. Interestingly, the model appears better at recognizing foils than hits like humans. Therefore, the results imply we have a reasonable model of the intuitive decision making process.

Intuitive decision making and its development through implicit learning depend upon a form of non-declarative memory called procedural memory. According to Squire (2004, 2009), human memory can be subdivided into two basic systems. One system is a declarative memory system, which entails conscious recollection about facts and events. The other system is a non-declarative memory system, one form of which is procedural memory, which involves memory relating to the ability to extract common elements and patterns from separate events (Knowlton, Ramus & Squire, 1992; Knowlton & Squire, 1993; 1996), as well as memory supporting the development of skill-like abilities.

Procedural memory is involved in much more than motor skill. Rather, procedural memory is also involved in the recognition of invariant properties within patterns of information that unfold over time (Patterson, et al., 2009). Because procedural knowledge is highly implicit and does not require full conscious processing to be evoked and used, it is especially useful in situations where the traditional analytical (conscious) processing of information, which is slow and limited by working memory capacity, would burden a person already stressed within a dynamic, time-pressed task environment.

ACT-R implements a formalized representation of declarative memory and non-declarative procedural memory systems and both systems have sub-symbolic components. In ACT-R, the declarative memory contains facts and events, but the retrieval of facts is not always a conscious process in that it is based on the sub-symbolic activation, which is based on the history of use of the memory.

ACT-R's procedural memory is activated by recognizing stimuli, i.e., matching the "IF" parts and then initiates actions in one or more of the architecture's modules, such as changing the current description of the goal, initiating the recall of a declarative chunk of memory, initiating a motor action, or moving the focus of the eyes. This is a different concept of "procedural" memory that discussed above.

We used ACT-R's declarative memory for facts and its sub-symbolic activation associated with those memories along with simple productions to represent the intuitive decision making process. We can produce both the overall performance as well as the different performance on hits and correct rejections implying we are modeling the cognitive processes involved. Our modeling formalization and data available raises research questions concerning the intuitive decision making process and the appropriate architectural approach. Further research will be needed to determine if another strategy for the learning and use of the learned knowledge would also perform well compared with human data.

Modeling the non-declarative knowledge that was investigated in the present study is a challenge because this kind of procedural knowledge is more abstract than ACT-R's simple symbolic chunks, their activations, or productions, yet it is not declarative. This means that the conceptualization of non-declarative procedural memory in ACT-R may need to be expanded to include abstract representations of statistical regularities and invariances sampled from the environment.

Our ability to match available human performance data does not mean that we have proven the ACT-R model is necessarily an explanation of the underlying human cognitive processes. Humans can make the translation of their learning in one environment to another, as demonstrated by Patterson, et al. Their participants could learn pattern independent of the specific items in the sequence. However, the ACT-R model is not able to do that because the declarative chunks learned are specific to the letters in the stimuli and the knowledge are not generalizable. Further work in this area may justify extending ACT-R to represent implicit patterns more abstractly.

Conclusions

This work demonstrates some of the reasons for building computational cognitive models. First, we are able to replicate human performance on this implicit learning and intuitive decision making task. This was accomplished by implementing a model of a cognitive learning and evaluation process that, while consistent with the ACT-R theory of cognition (Anderson, 2007), was inconsistent with the intuitive nature of procedural memory in humans. The strategy implemented was to build a memory of bigrams of sequential letters and then evaluating a test string by checking that each bigram had been seen before. However, other strategies may be similarly successful.

Second, cognitive modeling supports formally exploring alternative explanations for observed behavior. The model could be modified to test whether learning trigrams in the training strings could yield similar results. It could also be modified to test if recognizing only the first few and/or the last few letters, i.e., primacy or recency, can match the human participants' performance. A third alternate strategy is simply a voting strategy where recognized bigrams are counted and if above a threshold, the model would report a match. These strategies have not yet been tested, but with a cognitive modeling environment, they can be.

Third, this work also demonstrates that at least some System 1 as well as System 2 forms of cognition can be replicated within the current ACT-R architecture, but not necessarily all. This demonstration included implicit learning and intuitive decision making. From the work of Patterson et al. (submitted), there is data on the performance of human participants who memorize training strings rather than passively viewing them. This may be a nice example of an effortful System 2 learning strategy rather than the far less effortful System 1 passive learning strategy. The characteristics of each system need more study. To address other examples of System 1 cognition, ACT-R may need to be extended to include introspective factors representing emotional aspects of cognition such as current arousal, general mood, and temperament.

Finally, cognitive modeling advances our understanding of cognitive processes by providing a framework to represent and explore the explanation of behaviors, such as intuitive decision making, that seem to be driven by cognition that is “beyond rational”.

Acknowledgments

This work was supported in part by AFOSR/AFRL grant FA9550-10-1-0385 and the George Mason University Center of Excellence in Neuroergonomics, Technology, and Cognition (CENTEC).

References

Anderson, J.R. (1996). Implicit memory and metacognition: Why is the glass half-full? In L. M. Reder (Ed.) Implicit memory and metacognition, 123-136. Mahwah, NJ: Erlbaum.

Anderson, J. R. (2007). How Can the Human Mind Occur in the Physical Universe? Oxford University Press.

Anderson, J. R. & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General, 128*, 186-197.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglas, S., Lebiere, C., and Qin, Y. (2004). An integrated theory of mind. *Psychological Review, 111*, 1036-1060.

Boydstun, A.S., Patterson, R.E., Pierce, B.J., Park, L.M., & Tripp, L.M. (2011). Articulatory suppression affects situational pattern recognition in immersive environments. *Proceedings of the 10th Annual Conference on Naturalistic Decision Making* (pp. 157-159) Orlando, FL.

Covas-Smith, C.M., Patterson, R., Pierce, B.J., Cooke, N., & Homa, D. (2011). Diversity of experience and intuitive decision making in an immersive environment. *Proceedings of the 10th Annual Conference on Naturalistic Decision Making* (pp. 53-156). Orlando, FL.

Cleeremans, A., Destrebecqz, A., and Boyer, M. (1998) Implicit learning: News from the front. *Trends in Cognitive Sciences, 2*, 406-416.

Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment and social cognition. *Annual Review of Psychology, 59*, 255-278.

Klein, G. (1998). Sources of power: How people make decisions. Cambridge, MA: MIT Press.

Knowlton, B.J., Ramus, S.J. & Squire, L.R. (1992). Intact artificial grammar learning in amnesia: Dissociation of classification learning and explicit memory for specific instances. *Psychological Science, 3*, 172-179.

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

Knowlton, B.J. & Squire, L.R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science, 262*, 1747-1749.

Lebiere, C. & Wallach, D. (2001). Sequence Learning in the ACT-R Cognitive Architecture: Empirical Analysis of a Hybrid Model. In R. Sun & C. L. Gilles (Eds.). Sequence Learning: Paradigms, Algorithms, and Applications (pp. 188-212). Berlin: Spinger Lecture Notes in Computer Science.

Lebiere, C. & Wallach, D. (2002). Explicit or implicit learning? An integrative theory of sequence learning. In *Proceedings of the 43rd conference of the German Psychological Association*.

Matthews, R.C., Buss, R.R., Stanley, W.B., Blanchard-Fields, F., Cho, J.R., and Druhan, B. (1989) Role of implicit and explicit processes in learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 1083-1100.

Patterson, R., Pierce, B.J., Bell, H., Andrews, D. & Winterbottom, M. (2009). Training robust decision making in immersive environments. *Journal of Cognitive Engineering and Decision Making, 3*, 331–361.

Patterson, R.E., Pierce, B.J., Boydstun, A.S., Park, L., Shannan, J., Tripp, L., and Bell, H. Implicit statistical learning creates intuitive decision making in a simulated real-world environment. *Journal of Cognitive Engineering and Decision Making*, submitted.

Perrachet, P. and Pacton, S. (2006) Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in Cognitive Sciences, 10*, 233-238.

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

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*, 3–22.

Squire, L.R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory, 82*, 171–177.

Squire, L.R. (2009). Memory and brain systems: 1969 – 2009. *The Journal of Neuroscience, 29*, 12711–12716.

Wallach, D. and Lebiere, C. (2003) Implicit and explicit learning in a unified architecture of cognition. In L. Jimenez (Ed.) Advances in Consciousness Research, 215-250. Amsterdam: John Benjamins Publishing Company.