

Choice and Learning under Uncertainty: A Case Study in Baseball Batting

Christian Lebriere (cl@cmu.edu)
Human Computer Interaction Institute
Carnegie Mellon University

Dario Salvucci (salvucci@cs.drexel.edu)
Department of Computer Science
Drexel University

Abstract

This paper describes the modeling of human performance in a real-world, embodied, stochastic task: baseball batting. Experimental results were gathered in a virtual reality setup and a Markov model of performance, especially errors, was developed. The focus of this paper is on a model of the task developed in the ACT-R cognitive architecture, most specifically of the critical subtask of generating an expectation for the next pitch. The model required no parameter tuning and provides an *a priori* account of the results based on the architectural constraints of declarative memory. The Markov and ACT-R models are briefly compared. The broader relevance of the task is discussed and possible applications are suggested.

Introduction

Uncertainty is an ever-present factor that affects countless choices people make on an everyday basis — investing in the stock market is only one obvious example. A much broader class of activities, however, involves trying to find patterns in uncertainty (even and perhaps especially when none exist) that allow for improving one's decisions in the face of incomplete and uncertain information.

One domain in which pattern recognition under uncertainty plays a major role is that of two-player games. For instance, the classic game of paper rocks scissors (PRS) can be viewed as the attempt to find patterns in an opponent's sequence of moves. We recently developed models of PRS play (Lebriere & West, 1999; West & Lebriere, 2001) based on a model of sequence learning (Lebriere & Wallach, 1998).

The focus of this paper is baseball batting, fundamentally a two-player game between batter and pitcher (plus catcher). Baseball batting brings to bear many of the same interesting phenomena of other two-player games such as paper rocks scissors. In this work, we utilize the basic sequence-learning approach to model baseball batter expectations and behavior and demonstrate how this approach captures batters' cognitive processes as elucidated by the temporal error of their swings.

We begin with a brief introduction to baseball for the benefit of non-American (and non-Japanese) readers. We then describe two experimental studies, the first examining the effects of pitch speed on batter performance, the second examining the effects of both pitch speed and pitch count on batter performance. We will also describe the original

Robert Gray (robgray@asu.edu)
Applied Psychology Department
Arizona State University

Robert West (robert_west@carleton.ca)
Departments of Cognitive Science and Psychology
Carleton University

Markov model developed to account for that data. We then provide an overview of the ACT-R cognitive architecture, emphasizing the points most critical to our models of baseball batting, and describe the models themselves. Both models emphasize parameter-free modeling in which the simplest possible models are developed and run to generate *a priori* predictions, and then compared to empirical data with no subsequent data fitting. We conclude by comparing the Markov and ACT-R models and discuss some possible implications and applications.

Empirical Results and Markov Model

Baseball

Baseball is a complex game but the subset at the heart of this study is both fundamental and relatively simple. Pitcher and batter face each other 90 feet apart, the former standing on a dirt mound and the latter near home plate. The pitcher throws the ball and the batter attempts to hit it with a bat to put it in play and reach base safely. If the ball crosses the strike zone, i.e. the area above home plate at a height between the batter's knees and shoulders, or if the batter swings at the ball without putting it in play, the play is ruled a strike. Otherwise, i.e. the ball does not cross the strike zone and the batter doesn't swing, the play is ruled a ball. If the count for a given batter reaches four balls, the batter is given a walk and is allowed to reach first base safely. If the count first reaches three strikes, the at-bat is ruled a strike out and the batter is retired.

At some level, this can seem and indeed sometimes is a matter of pure physical skill: the pitcher throws the ball as fast as he can, thereby reducing the batter's time to react, and the batter swings as hard as he can, attempting to hit the ball out of the park for a home run. However, at the highest levels the game is considerably more complex and mentally demanding. A pitcher can often throw the ball in several different ways, resulting in pitches of varying speed and trajectories. Because of the limited time to react, the batter often needs to formulate an expectation of the type and location of the pitch before it is thrown, resulting in poorer performance when that expectation is not correct. At that level, baseball becomes a mental game of outguessing the opponent.

To examine the impact of cognitive processing on baseball batting performance, Gray (2001) designed a

virtual reality experiment in which six college-level baseball players batted against simulated pitches. A position tracker installed on the bat recorded its motion and allowed it to be compared against the simulated flight of the ball, yielding a measure of error between the time the ball and the bat crossed the plate. Two separate experiments were run.

Experiment 1 – Effect of pitch sequence

In experiment 1, a sequence of pitches was constructed from a random distribution of fast and slow pitches. To quantify the effect of pitch sequence and pitch count on batting performance we measured the absolute temporal error for each swing, i.e. the difference between the times at which the ball and the bat cross the same point above home plate. The dark bars in Figure 2 and 3 plot absolute temporal errors (for batter 1 and for all batters, respectively) for fast pitches that were preceded by different pitch sequences. It is clear that the prior sequence of pitches had a strong influence on the temporal error in the swing. For example, when a fast pitch was preceded by three consecutive slow pitches (S, S, S) the absolute temporal error was 59 ms higher than when a fast pitch was preceded by three consecutive fast pitches (F, F, F). This is a substantial difference given that the temporal margin for error in hitting has been estimated at ± 9 ms (Watts & Bahill, 1990).

The effect of pitch sequence was modeled in Gray (2001) using an extension of the finite-state Markov model developed by Falmagne et al. (1975). In this model, the hitter can be in one of two states: expecting a fast pitch ($EX_n=F$) or expecting a slow pitch ($EX_n=S$) and the hitter's expectation can change from pitch to pitch based on two simple strategies. If the speed of the current pitch matches the hitter's current expectation we assume that the hitter maintains the same expectation for the next pitch. But whenever the hitter's expectation is incorrect, there is some probability that the expectation will be changed from fast to slow (a_s) or from slow to fast (a_f) for the next pitch. Moreover, the temporal error for a fast pitch when the hitter is expecting a fast pitch (T_f) should be less than when expecting a slow pitch (T_s). These transition rules can be used to predict the temporal error (TE) as function of the sequence of prior pitch speeds. The four model parameters (a_f , a_s , T_f and T_s) were estimated simultaneously using the STEPIT procedure (Chandler, 1969). This procedure uses a least squares criterion to minimize the difference between observed and predicted temporal errors for all pitch sequences. The model provided good fits, with R^2 values ranging between 0.51 and 0.96 for the 6 batters.

Experiment 2 – Effect of pitch count

To examine the effect of pitch count (i.e. the current number of balls and strikes in a given at-bat) on hitting performance we varied the horizontal location of the simulated pitches so that some of the pitches were strikes and some of the pitches were balls. Balls and strikes were randomly distributed with the same probability, as were slow and fast pitches, except for “ahead” counts of 0-2 and 1-2 (where slow pitches

occurred with a probability of 0.65) and “behind” counts of 2-0, 3-0 and 3-1 (where fast pitches occurred with a probability of 0.65).¹ Visual feedback was given for the pitch call (ball or strike), total pitch count, walks and strikeouts. The dark bars in Figure 4 show the mean temporal errors for all subjects for each pitch count. Clearly pitch count had a large effect on temporal error and its effect was similar to a well-known aspect of real baseball: batting performance was better when the hitter is ahead in the count (1-0, 2-0 or 3-0) then when the hitter is behind in the count (0-1, 0-2 and 1-2). To model the effect of pitch count on hitting performance, two additional transition rules (and associated parameters) were added to the basic model described above. The model provided good fits, with R^2 values ranging between 0.59 and 0.83 for the 6 batters. For additional details, see Gray (2001).

ACT-R Cognitive Architecture

ACT-R (Anderson & Lebiere, 1998) is a hybrid architecture of cognition, which combines a production system to capture the sequential, symbolic structure of cognition, together with a subsymbolic, statistical layer to capture the adaptive nature of cognition. ACT-R 5.0 (Figure 1) is a modular, neurally plausible, architecture that is decomposed as a set of localized modules (e.g. long-term memory, visual, motor, etc) interacting through buffers connected to a central production pattern-matching module.²

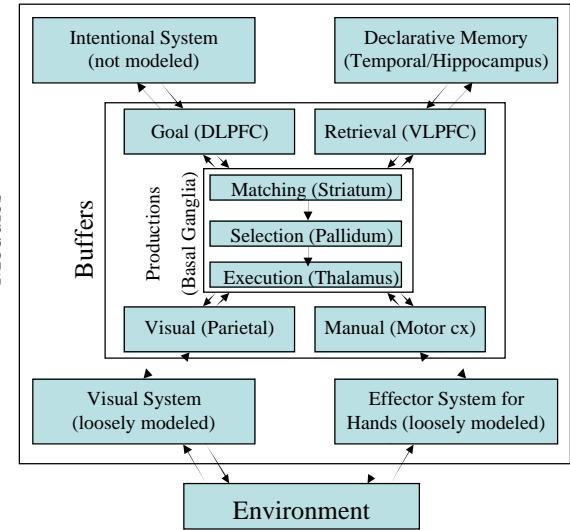


Figure 1: ACT-R 5.0 Architecture

¹ “Ahead” and “behind” here refers to the situation of the pitcher. The first number in each pair is the number of balls and the second is the number of strikes. With more strikes, a pitcher is ahead in the count and thus more likely to throw slow pitches such as curveballs which are harder to hit but also tend to be harder to control and are more likely to result in balls. With more balls, a pitcher is behind in the count and thus more likely to throw faster pitches such as fastballs which are easier to control but also to hit.

² Our current hypotheses as to the neural location of the various modules and buffers are indicated in parenthesis.

The module of foremost importance in this paper is declarative memory. Memory is organized as a set of chunks that are retrieved according to their activations. Activation, as defined in the following equation, is composed of a base-level term, a context-sensitive term and a stochastic component respectively:

$$A_i = B_i + \sum_j W_j \cdot M_{ji} + N(0, s) \quad \text{Activation Equation}$$

The base-level term is learned as a function of experience and captures the Power Law of Forgetting (Rubin & Wenzel, 1996) as well as the Power Law of Learning (Newell & Rosenbloom, 1981):

$$B_i = \ln \sum_j t_j^{-d} \quad \text{Base Level Learning Equation}$$

t_j is the time elapsed since each rehearsal of that chunk and d is the decay rate. The context-sensitive term that reflects the current task and provides a similarity-based partial matching similar to connectionist distributed representations. The stochastic component, normally distributed with mean 0 and magnitude s provides transitional variations in performance. For more details, see (Anderson & Lebiere, 1998).

Study 1: Effect of Pitch Speed

Model

Simplicity was an overriding goal in the development of the model. A simple model, directly based on assumptions and mechanisms of the architecture, limits free parameters and degrees of freedom in representation and processes. This is consistent with the intuitive nature of the task and the lack of sophisticated strategies reported by the subjects. In declarative memory, we simply represented each possible event, i.e. a fast or slow pitch, as a chunk, yielding a grand total of two chunks. Each new trial led to an activation boost for the chunk corresponding to the pitch, as determined by the base-level learning equation. This essentially corresponds to the batter being aware of each pitch, specifically whether it was slow or fast. No other declarative structure was represented.

There is a similar simplicity in the procedural module. It is generally acknowledged that, because of the very short time available between the release of the pitch and the time to swing the bat, batters anticipate the next pitch. Consequently, there was only one production rule responsible for decision making. That production simply retrieved the most active pitch in memory (i.e. fast or slow), and that was the one that it anticipated. That cognitive operation can be performed in a fraction of a second³, and its result is determined solely by the activation of the

competing chunk, which is computed according to the Activation Equation described above. Since the retrieval is essentially an unconstrained free association, the context part of the activation equation is not relevant. The base-level part of the activation is determined by the base-level learning equation, which itself is a function of past events. The memory decay rate d is fixed at 0.5, a value used in almost all ACT-R models. The noise part of the equation is determined by its magnitude s , which is fixed at a value of 0.25, also used in many ACT-R models.⁴

The model described above is essentially a decision-making model predicting which pitch will be anticipated next. Unfortunately, the motor module currently implemented in ACT-R is more concerned with typing at a keyboard and moving a mouse than swinging a baseball bat. Thus, we had to specify the mean temporal error (MTE) resulting from a given anticipatory guess and the actual pitch speed:

$$MTE = \frac{dist}{v_g} - \frac{dist}{v_a} \quad \text{Mean Temporal Error Equation}$$

$dist$ is the (simulated) distance from the mound to home plate, v_g is the anticipated (guess) speed and v_a is the actual pitch speed. In other words, the error is equal to the difference between the estimated time for the ball to get to the plate and the actual time. This essentially means that the batter will swing to hit the ball at the time that it expects it to cross the plate.

Results

Figure 2 presents the results for the average of 100 model runs compared to the results for subject 1, the best of the 6 subjects. Performance is measured in terms of mean temporal error for fast pitches (fastballs) as a function of all previous sequences of pitches of length 1, 2 and 3.

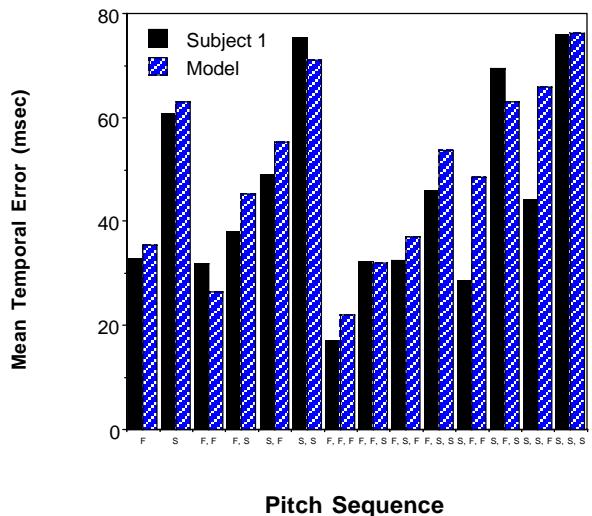


Figure 2: Experiment 1 Data and Model Results

³ ACT-R predicts the time to retrieve a chunk as a (negative exponential) function of its activation, but that is not relevant here since batters had sufficient time between at-bats that the latency of a single memory retrieval is well within that limit.

⁴ A repository of ACT-R models can be found at <http://act-r.com>.

The model closely reproduces the features of the data. The first two data points on the left establish that the mean temporal error (for a fast pitch) is much lower following a fast pitch than a slow pitch. This results directly from the activation calculus: the previous fast pitch gives a boost of activation to the corresponding chunk, which makes it more likely to be retrieved. This tends to result in a correct guess for the next pitch when a fast pitch is followed by another fast pitch. Conversely, if the fast pitch is followed by a slow pitch, the guess is more likely to be another slow pitch, and thus wrong. Those probabilities of guessing by the model directly (and linearly) map into mean temporal error to be comparable to the subject data.

The same pattern holds for longer sequences of prior pitches. The next four data points are for sequences of length two and the next eight for sequences of length three. A number of additional patterns can be observed. First, the range of performance widens with the length of the sequence. In other words, the MTE for a prior sequence of “F, F” is lower than for “F” (and even lower for “F, F, F”) and the MTE for a prior sequence of “S, S” is higher than for “S” (and even higher for “S, S, S”). For subjects, the obvious interpretation is that the more pitches in a row of a particular kind, the more one expects that pitch the next time. For the model, that is a direct consequence of the frequency effect captured by the base-level learning equation: the more occurrences of an event, the higher the activation boost, and thus the larger the probability of being selected for retrieval. Second, performance is sensitive to the order of recent pitches. For example, a prior sequence of “S,F” leads to a higher MTE (for a fast pitch) than a sequence of “F,S” because for the latter the fast pitch is the more recent of the two whereas for the former the slow pitch is the more recent. The model captures that effect because the base-level learning equation decays activation boosts as a function of time. Thus the more recent the rehearsal, the larger the activation boost and thus the higher the probability of being retrieved. However, the frequency effect is stronger than the recency effect. Thus the MTE is lower for a prior sequence of “S, F, F” than for a sequence of “F, S, S” because having two out of three fast pitches has more of an effect than having the most recent of the three (listed first) be the fast one.

The predictions of the model are both qualitative and quantitative. Qualitatively, the patterns described above will hold for the model no matter which motor transfer equation is used to map the guessing probabilities generated by the retrieval process from the declarative module because it will preserve the monotonic relations between guessing probabilities for the various conditions. As is the case, those patterns hold for each individual subject. It should be emphasized that those patterns are direct, no-parameter predictions of the ACT-R architecture, specifically the activation equations, together with the default values for the relevant architectural parameters used in many other models. However, the results of Figure 2 are also direct quantitative predictions of the model. We had to add the

additional assumption of the Motor Transfer Equation, but that equation is a mere expression of the experiment, specifically the distance between pitcher and batter and speed of the baseball. As such, it simply assumes that the batter is incapable of incorporating his perception of the pitch in his batting performance, but that his pitch anticipation is translated in a perfectly timed swing. Using perception of the pitch to adjust one’s swing could lead to arbitrarily good performance while an imperfect translation of pitch anticipation into swing could lead to arbitrarily poor performance. The Motor Transfer Equation we adopted is merely the neutral case. As such, the quantitative correspondence to subject 1 is remarkable.⁵ The fit for prior sequences of length 1 is particularly good. The increased variation for longer sequences is a direct result of the fact that the number of observations per data point decreases with the length of the sequence. To limit that variation, Figure 3 presents the same data averaged across all subjects. To account for the presence of subjects with poorer skills, we introduced a multiplicative factor of 2 in the MTE equation. In other words, we assumed that the average subject is about twice as poor as subject 1. As can be seen, all the trends present in the individual data are reproduced.

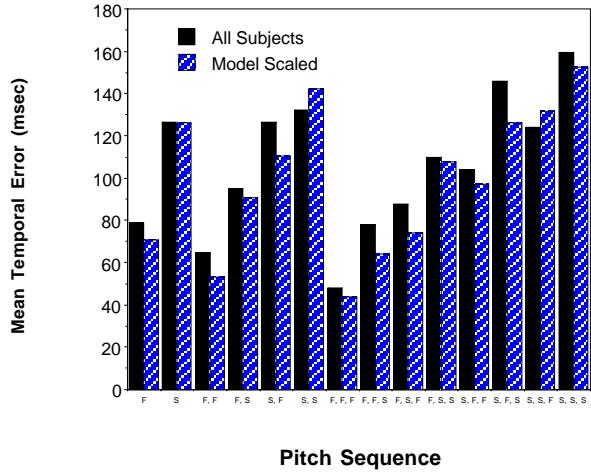


Figure 3: Experiment 1 All Subjects Data and Model

Study 2: Effect of Pitch Speed and Count

Model

The basic model is unchanged, except to account for the addition complexity of the pitch count and location. The pitch count essentially provides the strategic context of each pitch, and as such each pitch (fast or slow) is stored in memory in a chunk that also records the current number of balls and strikes. Thus, rather than only two chunks (fast and slow pitch), there are now 24 possible chunks: 2 speeds

⁵ As it turns out, subject 1 is the best of the six subjects. Thus it is reasonable to assume that the subjects have less than perfect swings and little ability to adjust their swings on the fly. That is consistent with their level of experience.

$\times 4$ ball counts $\times 3$ strike counts. For example, one such chunk is “fastball on a 3 balls and 2 strikes count”. The base level of each chunk will be learned as in experiment 1 as a result of each batting experience. Similarly, to take into account the current context, the production that retrieves the most active chunk to anticipate the pitch speed constrains its matching to the current ball and strike count. However, this is mitigated by the partial matching component of the activation equation. For example, if the current count is 3 balls and no strikes, one might retrieve a recent (and thus highly active) outcome for a 2 balls and no strikes count because the counts are sufficiently similar. To fully specify the partial matching process, we defined exponentially decreasing similarities M_{ji} between the components of the count, i.e. the number of balls and strikes. For example, the similarity of 3 balls to 2 balls is 0.5, to 1 ball it is 0.25 and to no balls it is 0.125. These similarities were used in other models (e.g. Lebiere, 1998) and do not constitute additional degrees of freedom. Finally, the location of each pitch is stored in separate chunks in the same way as their speed and is retrieved by a production rule in the same manner.

The model’s numerical parameters were left unchanged from experiment 1. The Motor Transfer Equation was also left unchanged, with the multiplicative parameter left at the value of 2 used to match the average subjects in experiment 1. The architectural parameter W_j that scales the context-sensitive term as a function of degree of mismatch, was left at its default architectural value of 1.5.

Results

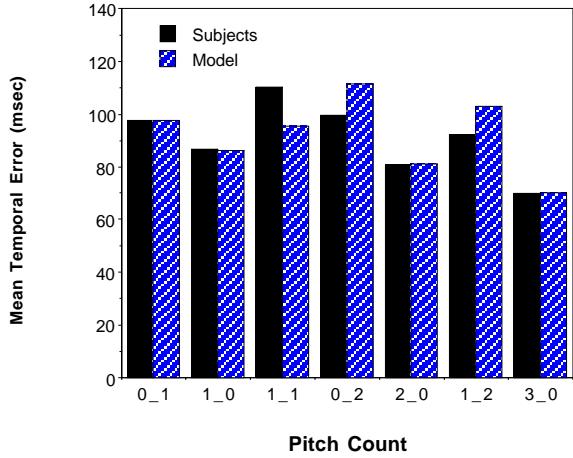


Figure 4: Experiment 2 All Subjects Data and Model

Figure 4 presents the match between the average of the 6 subjects and the model. The correspondence is excellent, especially considering that no parameters were adjusted to improve the fit. Both models and subjects display low MTEs for favorable counts such as 2-0 and 3-0 in which a fastball is more likely. That result arose in the model through the base-level learning equation that raised the activation of fastballs on those counts. Similarly, both models and subjects display high MTEs for unfavorable

counts such as 0-2 and 1-2, in which fastballs were less likely.⁶ But both models and subjects generalize those expectations to counts for which the distribution of fastballs and softballs was actually even. For example, the MTE is higher for the somewhat unfavorable count of 0-1 than for the somewhat favorable count of 1-0. In the model, this results from the similarities between balls and strikes numbers, which generalize retrieval to similar values.

ACT-R vs. Markov Modeling

Gray (2001) presented a Markov model of these results. A comparison between a statistical model such as a Markov model and a cognitive architecture, such as ACT-R, raises a number of points. The Markov model is composed of two states: expecting a fast pitch or a slow pitch. This echoes the two corresponding chunks in the ACT-R model. However, in the Markov model the state representation is purely binary whereas in ACT-R it includes the continuous activation values for each chunk and thus can represent fine-grained state distinctions. The second part of the Markov model is the transition probabilities between states as a function of the external events (slow or fast pitches). Those probabilities are estimated from the data to maximize the likelihood of producing the particular sequence. This is a major difference between the two models: the Markov model does not make a priori predictions while the ACT-R model predicts the probabilities of guessing based on the architectural constraints on memory. In the absence of accurate models of motor movements, both models map guessing probabilities into mean temporal errors using what is essentially a linear weighting formula. When generalizing to experiment 2, they adopt slightly different strategies. The Markov model introduces additional transition rules and associated probabilities to reflect the influence of the count. The ACT-R cognitive model, on the other hand, represents the count explicitly and then relies on partial matching to generalize between similar counts.

Discussion

Why study games? From the game theory perspective, much of human behavior, as well as the process of evolution, can be understood as games. The skills involved in batting represent a fundamental game-playing skill - anticipating what your opponent will do next. This skill was, no doubt, critical for the survival of our species. In baseball, the player tries to anticipate the pitch to increase his chance of hitting the ball. The same skill would have been used by our ancestors to anticipate the movements of their opponents or prey, and to increase their chance of striking them (possibly with something like a bat, which may partially explain the appeal of baseball). Because of the fundamental importance of game-playing skills in the process of evolution, we argue that game-playing is a critical test of any cognitive architecture, and that our results provide strong support for the claim that the ACT-R

⁶ Recall that the MTE is computed for fastballs only.

architecture represents a rational adaptation to the environment, through the process of evolution.

Since von Neumann & Morgenstern (1944) created the field of game theory there have been many different models of how humans play games, so it is worthwhile to ask what is special about our model. First, it is based on two well known experimental results: (1) the fact that people are very poor at randomizing (see Tune, 1964, and Wagenaar, 1972 for reviews), and (2) the fact that people compulsively search out and use sequential dependencies in tasks involving sequential guessing (e.g., Anderson, 1960; Estes, 1972). The vast majority of game-playing models rely on randomness and ignore sequential dependencies, putting them at odds with these findings. Second, our model is essentially a zero-parameter model, so the fit is not based on parameter tweaking. Third, this model has been used to account for human behavior in a variety of different game situations, without any modifications (e.g. Lebiere, Wallach & West, 2000; Lebiere & West, 1999). Finally, because the model is based on a cognitive architecture, we can relate the behavior of the model to other models based on the same mechanism. For example, the model is directly based on Lebiere and Wallach's (1998) model of implicit learning, indicating that game playing and implicit learning are based on the same underlying mechanism.

Conclusions

These results demonstrate that cognitive modeling is applicable in the real world in uncontrolled, naturalistic settings. A common criticism of cognitive modeling is that it is merely a parameter-fitting exercise (Roberts and Pashler, 2000). The models presented here, especially the one for experiment 1, are essentially zero-parameter predictions of the fundamental equations of the declarative memory module, themselves grounded in the rational analysis of cognition (Anderson, 1990). The ACT-R model predicts the hitter's expectations using its theory of declarative memory, in particular the activation calculus, rather than having to estimate those from the data. Being able to reliably predict the game of expectations between pitcher and batter raises the prospects of practical applications, especially in the domain of training. However, while we have been successful at modeling aggregate error probabilities, additional research is needed to look into the detailed, short-term sequential patterns of pitch selection and expectation in order to generate maximally accurate predictions, not just in the average but for each pitch.

References

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

Anderson, J. R (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.

Anderson, N. H. (1960). Effect of first-order probability in a two choice learning situation, *Journal of Experimental Psychology*, 59, 73-93.

Chandler, J. P. (1969). STEPIT: Finds local minima of a smooth function of several parameters. *Behavioral Science*, 14, 81-82.

Estes, W. K. (1972). Research and theory on the learning of probabilities. *Journal of the American Statistical Association*, 67, 81-102.

Falmagne, J. C., Cohen, S. P., & Dwivedi, A. (1975). Two-choice reactions as an ordered memory scanning process. Rabbit, P., & Dornic, S. (Eds.). *Attention and Performance V*. NY: Academic Press.

Gray, R. (2001). Markov at the bat: A model of cognitive processing in baseball batters. *Psychological Science*, 13(6), 542-547.

Lebiere, C. (1998). The dynamics of cognition: An ACT-R model of cognitive arithmetic. Ph.D. Dissertation. *CMU Computer Science Dept Technical Report CMU-CS-98-186*. Pittsburgh, PA.

Lebiere, C., & Wallach, D. (1998). Implicit does not imply procedural: A declarative theory of sequence learning. *Forty First Conference of the German Psychological Society*, Dresden, Germany.

Lebiere, C., Wallach, D., & West, R. L. (2000). A memory-based account of the prisoner's dilemma and other 2x2 games. In *Proceedings of the 3rd International Conference on Cognitive Modeling*.

Lebiere, C., & West, R. L. (1999). Using ACT-R to model the dynamic properties of simple games. In *Proceedings of the Meeting of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Newell, A. & Rosenbloom, P.S. (1981). Mechanisms of skill acquisition and the power law of practice. In J.R. Anderson (Ed.). *Cognitive skills and their acquisition* (pp. 1-56). Hillsdale, LEA.

Roberts, S. & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review* 107(2), 358-367.

Rubin, D.C. & Wenzel, A.E. (1990). One hundred years of forgetting: A quantitative description of retention. *Psychological Review*, 103, 734-760.

Tune, G. S. (1964). A brief survey of variables that influence random generation. *Perception and Motor Skills*, 18, 705-710.

Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behaviour*. Princeton, N. J.: Princeton University Press.

Wagenaar, W. A. (1972). Generation of random sequences by human subjects: A critical survey of the literature. *Psychological Bulletin*, 77, 65-72.

Watts, R. G. & Bahill, A. T. (1990). *Keep Your Eye on the Ball: the Science and Folklore of Baseball*. New York: W.H. Freeman and Company.