Speed/Accuracy Tradeoff in ACT-R Models of the Concentration Game

Titus Barik, Arpan Chakraborty, Brent Harrison (tbarik | achakra | beharri5@ncsu.edu)
David L. Roberts, Robert St. Amant (robertsd | stamant@csc.ncsu.edu)
Computer Science Department
North Carolina State University
Raleigh, NC 53706, USA

Abstract

This paper describes the development of subsymbolic ACT-R models for the Concentration game. Performance data is taken from an experiment in which participants played the game under two conditions: minimizing the number of mismatches/ turns during a game, and minimizing the time to complete a game. Conflict resolution and parameter tuning are used to implement an accuracy model and a speed model that capture the differences for the two conditions. Visual attention drives exploration of the game board in the models. Modeling results are generally consistent with human performance, though some systematic differences can be seen. Modeling decisions, model limitations, and open issues are discussed.

Keywords: Cognitive modeling; ACT-R; Concentration game; memory game; speed-accuracy tradeoff

Introduction

The game of Concentration has been used for decades as an exercise in reasoning about probabilities (Kirkpatrick, 1954) and for testing theories of memory (Eskritt, Lee, & Donald, 2001). Recently it has also become a target for cognitive modeling.

The Concentration Game, also known as the Memory Game, is a classic solitaire card game in which cards are laid face down on a board or table. On each turn, the player turns over a first card and then a second card so that both are face up. If the two cards match (i.e., if they show the same symbol), the cards are removed from the board. If the cards are a mismatch, then they are turned face down again and the next turn proceeds. The object of the game is to turn over pairs of matching cards until all of the cards have been removed from the board. For every card on the board, there exists exactly one matching card.

In this paper we describe an experiment using an online version of Concentration with 16 cards arranged on a grid. 179 participants played the game under two conditions, accuracy and speed. We describe ACT-R models for each condition. We show that some aspects of a speed/accuracy tradeoff for this task can be reproduced in a simple way in conflict resolution for productions that retrieve information about cards from memory (Gerjets, Scheiter, & Tack, 2000). Modeling results are qualitatively similar to participant data.

To our knowledge, no ACT-R models of the Concentration game exist in the literature. Our results are preliminary, but we believe they provide useful insight into a rich and complex task for cognitive modeling research.

Related Work

Lavenex et al. (2011) describe game play as a procedure that retrieves card information from a fixed-size memory. On a given turn, if two cards in memory are a match, both are turned face up. The matching cards are then forgotten. Otherwise, the first card is chosen at random from those on the board but not in memory. If a match is found in memory, that becomes the second card; if not, a second card is turned face up at random. New cards are added to memory, with matches being preferred if memory capacity is reached.

Anderson et al. (2012) analyze a more complex variant of Concentration in which cards show math or verbal puzzles instead of symbols; these must be solved to determine whether cards are a match. An Imperfect Memory Model (IMM) records information about the board: "The model forgets the location of a matching card with probability p_f. However, even if the card is forgotten the model remembers that there was a matching card. If the model forgets the location of a card it tries one guess among the visited cards" (Anderson et al., 2012, p. 637). Participant choices are modeled by a 625-state Hidden Markov Model that uses the IMM to estimate transition probabilities. Combining the IMM with click timing and fMRI data, Anderson et al.'s model predicts the most probable path a participant takes through the game at better than 80% accuracy.

The ACT-R models developed in the body of this paper are simple, but they extend Lavenex et al. by incorporating cognitive constraints (e.g., they relax the assumption of perfect retrievals from a fixed-size memory, and they identify possible mechanisms for dealing with visual aspects of the game). While not approaching the sophistication of Anderson et al.'s work, our models explore a different aspect of behavior, speed/accuracy tradeoffs, which have been studied in only a few areas of the ACT-R architecture (e.g. Schneider & Anderson, 2012; Gerjets et al., 2000).

Method

We implemented a computer-based version of the Concentration game, as shown in Figure 1. The interface consists of a 4×4 grid of tiles, each tile 100 pixels square, to represent cards. When face down, a card is black. When face up, the card shows a white background with a single, centered black letter from the set A, B, C, E, H, I, P, and Q. Letters are presented in the Helvetica Neue LT Std 65 Medium typeface. We chose these letters and typeface to minimize letter confusion (Mueller & Weidemann, 2012).

In each trial, the participant turns over a card by clicking on it. When two cards are face up, their symbols are displayed for 1 second; at that point the cards are turned face down (in case of a mismatch) or cleared from the board (in case of a

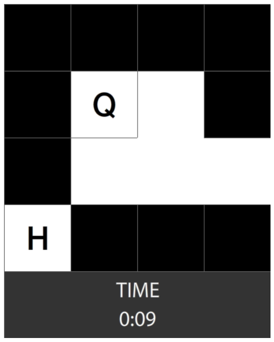

Figure 1: The Concentration interface, speed condition.

match). The participant may proceed without waiting for the system to turn cards back over, by clicking on any facedown card in the case of match, or by clicking on any card at all in the case of a mismatch. In either case, the clicked card then becomes the first card of the next turn. A trial is complete when all cards have been cleared from the board.[1]

To recruit participants for our study, we used snowball sampling, a technique in which initial participants help recruit additional acquaintances through various online social networks. After clicking on a recruitment message, participants completed a consent form followed by an optional survey asking for their age, gender, computer skills, and the type of pointing device that they would use for the study. The participant then played a small practice round with an in-game tutorial using a 2×2 board to become familiar with the game rules. Finally, the participants played the game.

A total of 179 out of 260 participants (69%) finished the experiment, though not all answered every survey question. We designed the experiment to take 10-15 minutes. The entire experiment lasted 9.6 minutes on average, not including times from 6 outlier players who took more than 1.5 hours to complete the experiment. 68 (38%) reported their gender as female, and 110 (61.5%) as male. Ages ranged from 16 to 67 with a mean of 30. With respect to computer skill, the breakdown among participants was 8 at the beginner level (4.5%), 41 intermediate (22.9%), 90 advanced (50.3%), and 39 expert (21.8%). For input devices, 116 used an external mouse (64.8%), 55 a touchpad (30.7%), 4 a pointing stick (2.2%), and 3 a trackball (1.7%). Because participants self-selected for the study, it is possible that the participant demographics are biased towards individuals who enjoy these types of games or have a better than average memory.

All participants were asked to play the game under two conditions: accuracy and speed. The conditions were distinguished by instructions and payoffs (Wickelgren, 1977). In the accuracy condition, participants were shown the number of mismatches on the screen during their game play, and they were scored on their ability to minimize mismatches. (The number of mismatches is a surrogate for the number of turns, because $turns = mismatches + matches$, where *matches* is always 8 to clear a 4×4 board; for consistency with the literature, we will discuss only turns in the remainder of this paper.) In the speed condition, participants were continuously shown the elapsed time in minutes and seconds on the screen, and they were scored on minimizing the time that it takes to clear the board. Participants played ten trials, each trial with a different board arrangement. Five of these trials were performed under one condition followed by five trials in the other condition. Within each condition, all participants played the same set of five boards in the same order. The order in which participants played each condition (accuracy rounds first or speed rounds first) was randomized. In brief, we have a repeated measures experiment, with participant being a random effect.

[1] Completing a turn requires just two clicks, unlike Anderson et al. (2012), in which a third click ends each turn; our game design does not allow us to distinguish between the end of one turn and the beginning of the next, which limits our analysis.

Table 1: Participant performance, mean (standard deviation).

Condition	Turns	Time
Accuracy	15.7 (4.0)	40.99 (16.12)
Speed	18.4 (4.9)	32.11 (11.96)

Results

To save space, we report only those results that are relevant to our modeling effort. Data cleaning reduced the 179 participants to 168, eliminating participants who had very long games (pauses of over 20 seconds). Our experiment results are as follows, with summary statistics in Table 1:

- Participants average fewer turns in the accuracy condition than in the speed condition (15.7 versus 18.4 turns, a difference of 2.7 turns; neglecting secondary variables, a within-subjects ANOVA gives $F = 82.26$, $p < 0.001$).
- Participants average less time in the speed condition than in the accuracy condition (32.11 versus 40.99 seconds, a difference of 8.88 seconds; neglecting secondary variables, a within-subjects ANOVA gives $F = 41.29$, $p < 0.001$).

The number of turns in the accuracy condition is comparable to values obtained in other experiments: 14.5 turns (Lavenex et al., 2011) and 16.9 turns (Anderson et al., 2012). 27 participants averaged below 13 turns on their five accuracy games, not far above Anderson et al.'s estimate of 12.24 turns for optimal performance. As expected, we also found significant variation across participants in both conditions, as shown in Figures 2 and 3; the distributions are truncated to better show their shapes. The range of turns was [10, 49] over all participants in the accuracy condition, [12, 47] in the speed condition. Completion times are similarly variable, with a range of [14.32, 129.44] in the accuracy condition and [12.60, 92.92] in the speed condition.

Game performance can be characterized in a number of ways beyond turns and time. Two interesting short-term strategies used by players are exploration—visiting

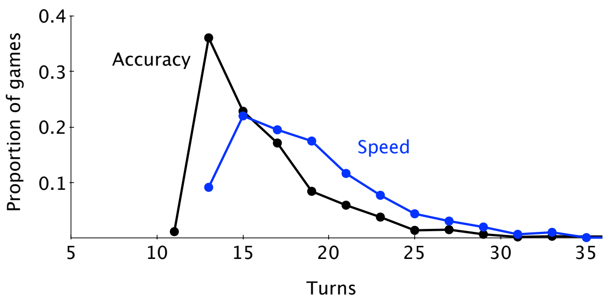

Figure 2: Participant turns per game.

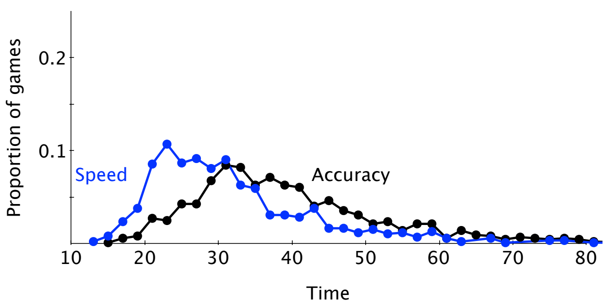

Figure 3: Participant time per game.

“new” cards, given memory limitations—and exploitation—successfully revisiting cards to find matches.

One relevant measure of exploration is the probability of not revisiting a card, especially on the first click of a turn; this is 66.4% in the accuracy condition and 54.2% in the speed condition. Thus, exploration is performed more often in the former case.

A useful measure of exploitation is the probability of choosing a match to a first card, given that the matching card has already been visited. Participants are successful 74.0% of the time in the accuracy condition (higher than Anderson et al.’s value of 62%, which we attribute to the greater cognitive demands of their game), and 62.4% in the speed condition. To summarize, players are better at both exploring the board as well as exploiting their limited memory in the accuracy condition, resulting in fewer mistakes, but lower completion times in the speed condition suggest that they trade off this accuracy for speed with some success.

Modeling

Our goal in modeling with ACT-R is to use the architecture to explain how our experimental results could arise. We worked within the basic architecture, without extensions, and relied on ACT-R parameters to fit our models to participant data. We built a baseline model for the accuracy condition, without subsymbolic processing, and then modified the model to handle the speed condition. Finally, we enabled subsymbolic computations, yielding a total of four models.

We started by translating Lavenex et al.’s description into

```
TAKE-TURN(memory)
    pair = RECALL-PAIR(memory) //
    if pair not recalled //
        first.location = RANDOM-NIM(memory)
        first.symbol = TURN-OVER(first.location)
        second.location = RECALL(memory, first.symbol)
        if second recalled
            TURN-OVER(second.location)
        else second.location = RANDOM-LOCATION()
            second.symbol = TURN-OVER(second.location)
    else TURN-OVER(pair.first.location) //
        TURN-OVER(pair.second.location) //
    if MATCH(first.symbol, second.symbol)
        REMOVE(memory, first)
        REMOVE(memory, second)
    else STORE(memory, first)
        STORE(memory, second)
```

Figure 4: Algorithm based on Lavenex et al. (2011).

an explicit algorithm, shown in Figure 4. With a memory size of 4, the algorithm generates a reasonable prediction for the number of turns under our accuracy condition. Simulation of 1,000 games, sampled from the ten board arrangements used in our experiment, produces a mean of 15.0 turns per game.

At a more detailed level, however, the algorithm does not quite reflect participant behavior. Specifically, consider when a pair of matching cards is recalled. This happens when the second card visited on a turn is added to memory; the pair will be taken off the board on the next turn. Because memory is limited, though, a pair is not always identified even if both cards have been seen before. In simulation, the second card on a turn matches an item that has been visited before in about 20% of turns. The algorithm retrieves such a pair with 80% probability; in the remaining cases, the matching card is no longer in memory. In contrast, in the same situation, participants choose the matching cards on the following turn only about 45% of the time. More commonly, 52% of the time, participants chose neither of the matching cards on the following turn. This suggests to us that identifying pairs in memory for the following turn is not critical in participants’ strategies. We thus simplified the algorithm in Figure 4 by removing the lines marked //.

The next issue we addressed was representation of information in memory. The algorithm records the location and symbol of a card together. Our model follows the algorithm in recording card locations and symbols in a single *card* chunk, with a *location* and a *symbol* slot, and new chunks are created with these slots filled when a card on the board is first visited.

A question for model design is whether card locations should be treated as positions in space or rather simply as unique identifiers, as the algorithm does. We found some spatial patterns in the participants’ choice of cards and their success in identifying matches. When a player chooses a card, sometimes its match has been visited before, but that matching card is not always chosen. Our intuition was that the locations of the mismatches would be closer to the matching card rather than uniformly distributed over the board. Figure 5 shows a representative distribution over all participant trials for accuracy and speed conditions, with the matching card in the top left corner.

0.803 (477) *0.618 (436)*	0.017 (10) *0.012 (7)*	0.017 (10) *0.016 (8)*	0.013 (8) *0.005 (3)*
0.015 (9) *0.009 (6)*	0.024 (14) *0.030 (21)*	0.007 (4) *0.023 (15)*	0.019 (11) *0.016 (11)*
0.017 (10) *0.021 (15)*	0.007 (4) *0.019 (14)*	0.010 (6) *0.007 (4)*	0.008 (5) *0.024 (14)*
0.013 (8) *0.005 (3)*	0.008 (5) *0.018 (9)*	0.015 (9) *0.009 (5)*	0.007 (4) *0.014 (8)*

Figure 5: Participant probability of correctly choosing a previously visited card, top left (with raw counts over all trials). Speed trials are italicized. The grid represents the 4x4 game board.

The mean probability of choosing a specific non-matching location in the accuracy condition when the matching card has been visited before is 1.7%, with a maximum of 5.7%. For the speed condition, the mean probability is 2.5%, with a maximum of 7.7%. There may be spatial patterns in the cases where the matching card is not chosen, dependent on its location, but we have not yet identified a way to predict them. Given these small probabilities, the models treat the locations of cards only as identifiers.

An important related issue is the interaction between vision and memory when exploring the board—the algorithm chooses the first card on a turn from those not in memory. To narrow the space of possible mechanisms, we examined the relationship between the duration between clicks and the number of cards seen; we find a near-zero correlation. We also find no relationship between click duration and the number of cards remaining on the board. This suggests to us that if memory is involved in the choice of new cards, it is not a simple serial elimination of cards that have been seen.

The algorithm can explicitly remove items from memory via the REMOVE operator (matching cards taken off the board should no longer be considered). The algorithm also has perfect memory, with failures due to capacity constraints. Finally, it can also determine whether a given card is *not* in memory (for the purpose of visiting a new card). ACT-R doesn't directly support removing items from memory. To choose cards not in memory in approximately constant time, the models rely on the visual attention capabilities of ACT-R. The models choose an unattended location for the first card on a turn and do the same when choosing a second card at random. Reliance on the visual system rather than memory retrievals to initiate a choice ensures that only cards still on the board will be considered.

The models implement the algorithm as productions in a straightforward way. The models begin by choosing an unattended card, clicking it, and reading the symbol. The chunk in the visual buffer is copied into the imaginal buffer for storage. The models then check for a second card in memory that matches the first. If a matching card is found (subject to memory limitations), that card is clicked. Otherwise, another unattended card is chosen, clicked, read, and stored in memory.

All of the above applies to the models for the accuracy condition. The models for the speed condition differ in a simple way: after the first card on a turn has been visited, the model may skip the memory retrieval of a possible match (via a different production) and instead choose an unattended card. The intuition is that participants speed up their actions by sometimes bypassing the time needed for memory retrievals, at the cost of missing matches. We implemented this by enabling randomness, so that the retrieval is skipped with a probability of 0.5. The speed model also contains adjustments to default motor module parameters (burst time and feature preparation time), to reduce the time to select targets.[2] Our analysis of motor issues has been minimal so far, in part due to the noisiness of movement in the participant data.

Results for these models on the Concentration game are shown in Table 2, in the ACT-R$_{baseline}$ rows. Model entries are based on 1,000 runs of the accuracy and speed models on board arrangements sampled from the ten used in the experiment, in a game simulation that uses ACT-R's interface-building facilities.[3] The baseline models fit human performance surprisingly well. The most obvious (and expected) discrepancy is the accuracy model's perfect memory of visited cards, shown in the Match column of the table. Other metrics show that the models revisit already-seen cards more often than participants, which increases both the number of turns and completion time, but this is offset by generally higher accuracy in memory retrievals than the participants.

Baseline model distributions of the number of turns and time per game are shown in Figures 6 and 7, overlaid with participant data. The means of the model distributions are close to the participant distributions, but there are clear differences, with the participant distributions having earlier peaks and longer tails. The models fail to capture some aspects of participant behavior, perhaps because participants apply a wider range of strategies than represented in the models.

To better reflect memory limitations, we enabled subsymbolic computations in the models and tuned ACT-R declarative parameters to approximate human performance. This contrasts with the algorithm's direct use of capacity constraints to enable memory failures and with Anderson et al.'s

[2]This corresponds to `(sgp :er t :motor-burst-time 0.01 :motor-feature-prep-time 0.01)`.

[3]The ACT-R game simulation is a simplification of the game used by participants in that all card symbols are visible at the start of the game; this allows us to avoid programming a one-second timeout to turn cards face down. Nevertheless, the models do not read any card symbol until after the card has been clicked and the state of the manual buffer is free.

Table 2: Mean participant and model performance across conditions, for turns and time per game; mean probability per game of visiting a previously seen card as the first card (Revisit$_1$), the second card when it is a match (Revisit$_{2+}$), and the second card when it is not a match (Revisit$_{2-}$); the mean probability of choosing a matching second card when it has been seen before.

Condition		Turns	Time	Revisit$_1$	Revisit$_{2+}$	Revisit$_{2-}$	Match
Accuracy	Participant	15.7	40.99	0.336	0.883	0.226	0.740
	ACT-R$_{baseline}$	15.1	33.50	0.338	0.842	0.276	1.000
	ACT-R$_{sc}$	15.9	36.60	0.383	0.848	0.315	0.805
Speed	Participant	18.4	32.11	0.458	0.900	0.358	0.624
	ACT-R$_{baseline}$	18.6	24.58	0.487	0.902	0.414	0.619
	ACT-R$_{sc}$	19.5	26.26	0.516	0.913	0.412	0.597

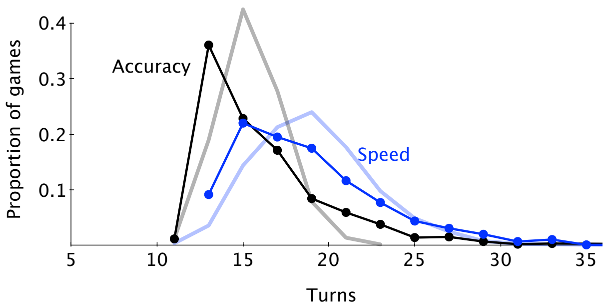

Figure 6: Baseline model turns per game (accuracy in gray, speed in light blue).

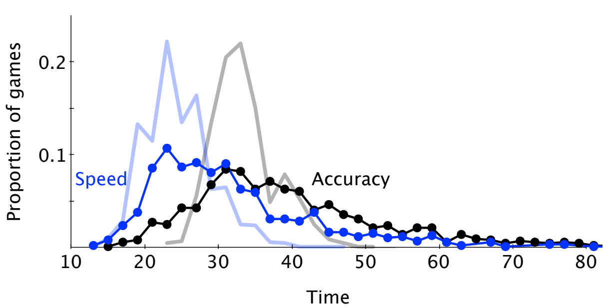

Figure 7: Baseline model time per game (accuracy in gray, speed in light blue).

p_f. We further modified the speed model, by associating utilities with the productions for retrieval and for choosing an unattended second card, and adding a small amount of noise.[4] Results are shown in the ACT-R$_{sc}$ rows of Table 2. The changes to the models alter the numbers but not the overall patterns in the relationship between the accuracy and speed models, or between the models and the participant data. In particular, for the new accuracy model, the probability of choosing a second card that matches the first card, given that the second card has been seen before, is a more reasonable 80.5%, though still above the participants' value of 74.0%. The distributions of turns and time for the new models are similar to those in Figures 6 and 7, shifted slightly to the right and with a wider spread. While the absolute estimates of performance are not perfect, all are within the bounds of participant performance. Further tuning of model parameters is possible (but changes should be motivated by specific hypotheses, which we are still developing). We also find that the metrics in Table 2 are very sensitive to small changes in the relationship between the baseline constant and the retrieval threshold.

To summarize, the accuracy and speed models give a reasonable match to participants on basic measures of performance, though we also find systematic differences.

Discussion

A number of caveats apply to our work. Decisions different from those we made are justifiable, and some would improve our models.

One modeling challenge was that of "choos[ing] at random a card from amongst those that are not in…memory" (Lavenex et al., 2011, p. 137), in approximately constant time. Our models rely on visual attention for this task. Given that the models revisit cards too often, and that the distributions of the number of turns per game consistently peak later than those of participants, we believe that the models are relying too much on the visual system for exploration, neglecting a memory component. A different approach could involve a memory representation that captures board information at a higher level than that of individual cards. Other possibilities are suggested by Johnson, Wang, and Zhang (2003), who have extended ACT-R to automatically encode relationships between visual objects in declarative memory, and by Lyon, Gunzelmann, and Gluck (2004), who model visuospatial working memory using similarity values in declarative memory. These may all be more cognitively plausible than our approach.

Our models ignore the spatial layout of cards, even though this clearly influences the participants' strategies. For exam-

[4] Parameter values were established by optimization (hill-climbing) over small neighborhoods, using turns for the objective function. Specifically, the accuracy model is parameterized with `(sgp :esc t :bll 0.5 :blc 3 :rt 2.4)`. The speed model adds `(sgp :ul t :egs 1 :er t)` and `(spp retrieve-second :u 10) (spp random-second :u 12) (spp find-match :reward 8)`.

ple, the cards most commonly visited early in the game are in the top row; these account for 49% of clicks before the first match is found. The first three card choices in a game typically walk the top row from left to right: the leftmost card first (67% of games), then the next card to the right (40%), and then again to the right (35%). This may allow for participants to exploit a form of distributed cognition: if a participant consistently follows a specific spatial pattern, then by remembering, "B was the first card I saw," the location of that card can be offloaded to the practiced procedure. Figure 1 suggests another possibility, that visual patterns in the aggregation of cards on the board may aid memory; we might suspect that the card directly under the top right card is more memorable than the others. A spatial component, as proposed by Gunzelmann and Lyon (2011), might support such retrievals.

A representation issue is raised by our models' storage of the location and symbol of a card in a single chunk. Following Altmann (2000), separating these items would allow for a more realistic representation of serial memory effects; also, as Anderson et al. (2012) note, Concentration players may remember having seen a card symbol before without remembering its location (Eskritt et al., 2001).

Another issue is our procedural approach to handling the difference between accuracy and speed trials, through conflict resolution between productions, which is a coarse approximation. A better approach could involve a memory model in which a longer period of time allocated for a retrieval would be more likely to return a result. The Retrieval by ACcumulating Evidence (RACE) model of memory by Van Maanen and Van Rijn (2007) suggests one possible avenue.

If our procedural approach were to generalize to other tasks, this would still leave open the question of how participants make the transition between the speed and accuracy conditions. The speed model incorporates a superset of the productions in the accuracy model, and it would be straightforward to include these as low- or zero-utility productions in the accuracy model. Learning new utility values poses challenging problems, however: participants were not given feedback per action or per turn about their performance, but rather a running total of either mismatches or time elapsed. They nevertheless adapted their performance successfully based on these cumulative measures. This is a non-trivial accomplishment; if we imagine that the fastest choices could be made by devoting no cognitive resources to memory at all, we would see purely random games in the speed condition that take more than 60 turns to finish.

Finally, as suggested earlier, it's important to recognize that our work has mainly been on model fitting. Our models constitute hypotheses about how Concentration is played and possible explanations for performance differences between the accuracy and speed conditions. Our models of accuracy and speed represent the behavior of a hypothetical "average" participant; we have not yet extended our work to the modeling of specific participants, and our analysis of patterns in their behavior is ongoing. Testing and validation of the assumptions behind our models remain for future work.

Acknowledgments

This work was funded by a grant from the Army Research Office [W911NF-08-1-0105]. Thanks to Erik Altmann, Dan Bothel, Duncan Brumby, Glenn Gunzelmann, Christian Lebiere, and Frank Ritter for advice, suggestions, and pointers to past modeling research.

References

Altmann, E. M. (2000). Memory in chains: Modeling primacy and recency effects in memory for order. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 31–36).

Anderson, J. R., Fincham, J. M., Schneider, D. W., & Yang, J. (2012). Using brain imaging to track problem solving in a complex state space. *NeuroImage*, *60*(1), 633–643.

Eskritt, M., Lee, K., & Donald, M. (2001). The influence of symbolic literacy on memory: Testing Plato's hypothesis. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, *55*(1), 39.

Gerjets, P., Scheiter, K., & Tack, W. H. (2000). Resource-adaptive selection of strategies in learning from worked-out examples. In *Proceedings of the 22nd Annual Conference from the Cognitive Science Society* (pp. 166–171).

Gunzelmann, G., & Lyon, D. R. (2011). Representations and processes of human spatial competence. *Topics in Cognitive Science*, *3*(4), 741–759.

Johnson, T. R., Wang, H., & Zhang, J. (2003). An ACT-R model of human object-location memory. In *Proceedings of the 25th Annual Cognitive Science Society* (p. 1361).

Kirkpatrick, P. (1954). Probability theory of a simple card game. *The Mathematics Teacher*, *47*(4), 245–248.

Lavenex, P. B., Lecci, S., Prêtre, V., Brandner, C., Mazza, C., Pasquier, J., & Lavenex, P. (2011). As the world turns: Short-term human spatial memory in egocentric and allocentric coordinates. *Behavioural Brain Research*, *219*(1), 132–141.

Lyon, D. R., Gunzelmann, G., & Gluck, K. A. (2004). Emulating a visuospatial memory field using ACT-R. In *Proceedings of the 6th International Conference on Cognitive Modeling* (pp. 368–369).

Mueller, S. T., & Weidemann, C. T. (2012). Alphabetic letter identification: effects of perceivability, similarity, and bias. *Acta Psychologica*, *139*(1), 19–37.

Schneider, D. W., & Anderson, J. R. (2012). Modeling fan effects on the time course of associative recognition. *Cognitive Psychology*, *64*(3), 127–160.

Van Maanen, L., & Van Rijn, H. (2007). An accumulator model of semantic interference. *Cognitive Systems Research*, *8*(3), 174–181.

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, *41*(1), 67–85.