# From 1000ms to 650ms:
# Why Interleaving, Soft Constraints, and Milliseconds Matter


**Bella Z. Veksler (zafrib@rpi.edu)**
**Wayne D. Gray (grayw@rpi.edu)**
**Michael J. Schoelles (schoem@rpi.edu)**
Rensselaer Polytechnic Institute
110 8th St., Troy, NY 12180



### Abstract

Evaluating and modeling human performance on even simple tasks requires a great deal of attention to millisecond-level cognitive and perceptual-motor operations. Modeling human performance in a task often requires that special care be taken to understand how these millisecond level operations are interleaved and how they evolve during the execution of the task. In modeling a simple decision-making task, we found that human subjects improved their routine speed as they became more familiar with the task. Modeling was conducted using the ACT-R architecture (Anderson & Lebiere, 1998). Refinements of the model indicated that interleaving of millisecond-level perceptual-motor and cognitive operators was crucial in accounting not only for the strategy shift as per soft constraints, but also in the marked speedup in performance over the course of several trials.


## Introduction

Milliseconds matter in understanding human performance (Gray & Boehm-Davis, 2000). The soft constraints hypothesis (Gray, Sims, Fu, & Schoelles, 2006) implies that in the course of routine interactive behavior, the cognitive controller tends to select interactive routines that shave milliseconds off of task performance. Unfortunately, this local optimization may not result in optimal global performance. Hence, even in tasks that are thought of as involving higher-level cognition, such as decision-making, global performance may be suboptimal due to nearsighted, local optimization of interactive routines.

From the perspective of the soft constraints hypothesis, computational models of decision-making must encompass a full accounting of the costs of information exploration and exploitation (Fu, 2007). Hence, an initial task for the modeler is to account for the perceptual-motor costs of skilled performance. As we show in this paper, this initial task brings to the foreground the interleaving of cognitive, perceptual, and motor operations that is characteristic of skilled performance.

We describe an exploratory effort to model the interleaving of cognitive, perceptual, and motor operations required for information exploration/exploitation in a table-based, decision-making task (Lohse & Johnson, 1996). The constraints of the model/framework are examined at the level of milliseconds, contrasted against human data, and the differences are analyzed with respect to the ACT-R framework (Anderson & Lebiere, 1998).

## The Task

The experimental environment used in this research was designed to study and model how information access influences the way in which a decision is made – specifically what information is considered and how it is integrated given the environmental constraints and accessibility of information. We used a simple table task (see Figure 1) in which each of six alternatives (arranged in rows) had a value on each of six attributes (arrayed in columns). The value of the alternative was derived by summing the attribute scores so that the higher the value, the better the alternative. This environment allowed us to manipulate the way information was accessed in order to determine the cognitive and perceptual-motor tradeoffs involved.

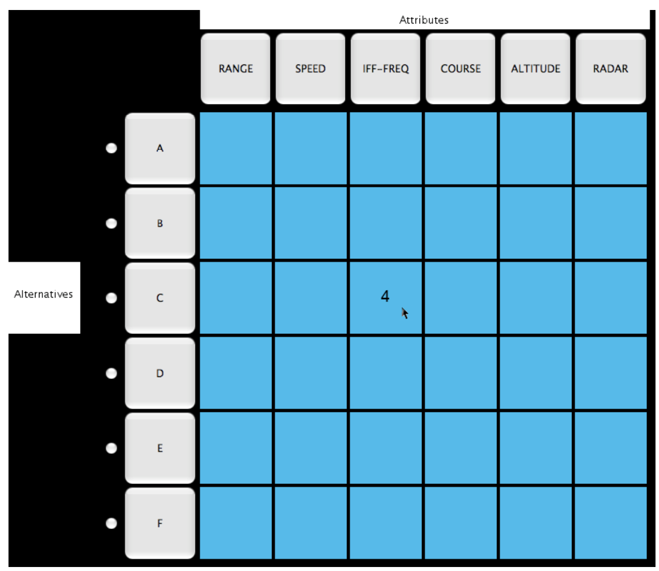


**Figure 1:** The *Table Task* environment for decision-making. The figure shows a subject clicking on the IFF-FREQ attribute (column) for alternative C (row).

We predicted that performance would vary based on exploration/exploitation costs that variations in the task environment imposed on the decision maker. In particular, we expected different costs to result in differences in the time to make a decision as well as the amount of information considered during the trial (i.e., information exploration). We also predicted that when participants were transferred to conditions with different environmental constraints, that the transfer of old strategies or the adoption of new ones would be influenced by exploration/exploitation costs of the old strategies applied to the new task environment. (Gray, 2000; Gray, Veksler, & Fu, 2004)

Although these general predictions are validated in the Results section, this analysis is beyond our current modeling effort. Exploratory modeling of this simple task revealed the necessity to focus our scope of analyses on the basic motor components prior to taking the next step into modeling the higher-level experimental effects.

## Human Data

### Method

The table task environment consisted of 6 alternatives arranged as rows and 6 attributes arranged as columns in a grid. Alternatives were military targets with attributes that contributed to their overall threat. There were values in the corresponding grid cells and it was the task of the participant to select the alternative whose corresponding attribute values summed to the highest value (see Figure 1).

There were a total of four conditions that varied how the values in the grid could be accessed. For purposes of this paper and the models presented, we only cover the "by cell" condition (CE). In this condition, participants accessed information one cell at a time by clicking on the grid cell corresponding to the value of an attribute for a particular alternative.

Each trial consisted of the participant checking the values in the grid and selecting the alternative with the highest overall value. Feedback on the number of correct answers was provided at the conclusion of the experiment. Participants completed 30 trials in this manner and for the CE condition included 18 participants.

### Results & Discussion

Our interest lays in modeling the millisecond level interactive routines of each trial in addition to the changes in information exploration/exploitation that occurred in performance within and across trials. The trial duration analysis below is intended as a benchmark for the subsequent model's performance.

**Total Trial Duration**

Total trial duration averaged 23.77s, StErr = 478.32ms. However, trial durations across the 30 trials follow a power-law of learning (Figure 2). It is thus important to note that trial duration decreased from the first (M = 36.92s; StErr = 3.66s) to the last trial (M = 21.72s; StErr = 1.8s).

**Number of Cell Clicks**

Participants were presented with a 6x6 grid of cells for a total of 36 cells that would need to be checked to have perfect information during a trial. Is there any indication that participants saved time by not checking every cell? Although there was some variability in the number of cells clicked across the trials, participants clicked an average of 35.81 cells. Therefore, participants roughly clicked on each cell once. The subsequent model therefore also clicks on each cell once during a trial.

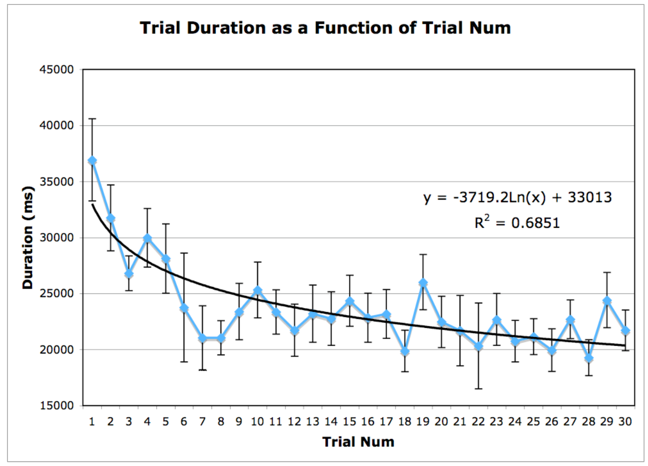


**Figure 2:** Power law of learning in trial duration

**Inter-Cell Click Interval**

In addition to trial duration and number of cells clicked, we assessed how participants were spending their time during task performance. In particular, we analyzed how long they spent between cell clicks. We will call this the "inter-cell click interval". Given that not all participants clicked the same number of cells during each trial, the following analysis only shows data from the first 36 cell clicks. As will be discussed later, the inter-cell click interval provides insights into how strategies evolve over time and how we can modify our models to match human performance. It also provides insights into the cognitive and perceptual-motor shortcuts that people take and that a cognitive model needs to account for. Essentially, these are the millisecond-level operations that are crucial in many repetitive or well-practiced tasks.

Figure 3 illustrates how inter-cell click intervals changed over the course of the whole task. Initially, inter-cell click intervals averaged ~950ms whereas by the end of the 30 trials, they had decreased to ~550ms. This trend is analogous to the trend of the overall trial duration we observed in Figure 2 and is one of the ways we can determine how the strategy that the participants employed evolved.

Furthermore, within a trial (see Figure 4), we witness variability in inter-cell click intervals with respect to cell click number. The more cells are clicked within a trial, the shorter the inter-cell click interval becomes. Notice also that there is a seesaw pattern such that every $7^{th}$ inter-cell click interval is longer than the surrounding ones. This is accounted for by the fact that each $7^{th}$ click was a row switch. One explanation for this is that at the end of a row, participants updated their current highest value and therefore took longer transitioning to the next alternative.

Within a row, the inter-cell click interval also showed a slight increase presumably explained by the increase in cognitive load as participants added more values to their running total of the alternative's value. Figure 4 shows average data across all 30 trials.

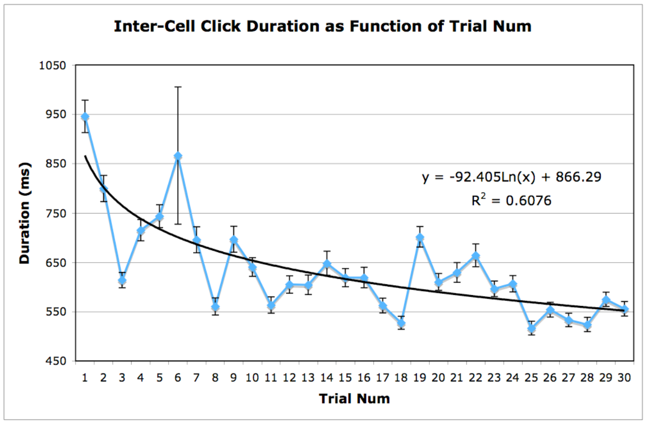


**Figure 3:** Power law of learning in inter-cell click interval

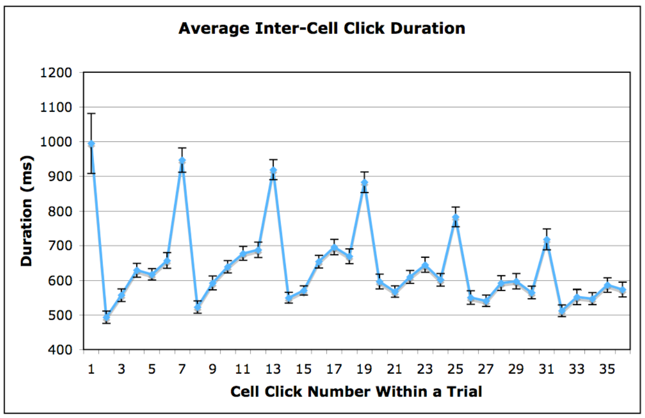


**Figure 4:** Average human inter-cell click interval within a trial. Peaks represent transitions between alternatives (rows)

### Transitioning Between Rows

Another important consideration for task performance is encompassed by the soft constraints hypothesis (Gray et al., 2006). Soft constraints guide the selection of interactive routines at the millisecond level to minimize performance costs as measured in terms of time. The utilization of soft constraints is reflected in apparent strategy shifts as performers become familiar with the task. Although initially certain biases may have caused the performer to use one set of interactive routines, the cost of exploration/exploitation ultimately shifts performance towards more efficient strategies.

This shift in strategy is most clearly seen in the transitions between alternatives (rows). Whereas initially participants were biased to "read" the values in the rows from left to right (Figure 5A), after several trials a more efficient strategy emerged. The new strategy had participants alternating the direction in which they clicked the cells based on their final position in a particular row (Figure 5B). Figure 5C shows that across trials participants increased their use of strategy B by 10%.

## The Model(s)

To model human performance on this task, we used the ACT-R cognitive architecture (Anderson et al., 2004). ACT-R is a modularized production system with a subsymbolic memory module. It has visual and motor modules to embed it in the task environment. It also has declarative memory and a procedural module. In addition, it has imaginal and goal buffers to store its working memory and goal chunks, respectively. Thus, it serves as a good framework to model human performance on this simple table task.

Several ACT-R models were developed in order to model the various components of human speed increases during this task. The essential structure of all of the models is the same: each model simply goes through each alternative, uncovers each cell value, sums the cells and updates its memory of the highest value after comparing it with the previous highest value. The differences between the models primarily lie in how they execute this list of perceptual-motor and cognitive operations.

At present, we have deliberately avoided implementing different decision-making strategies and have focused our modeling effort on getting the interleaving of cognitive, perceptual, and motor operators right. As discussed below, we do not know how to account for the obvious adaptations in interleaving that humans undergo. We view our lack in this regard as a comment on the state of the art in

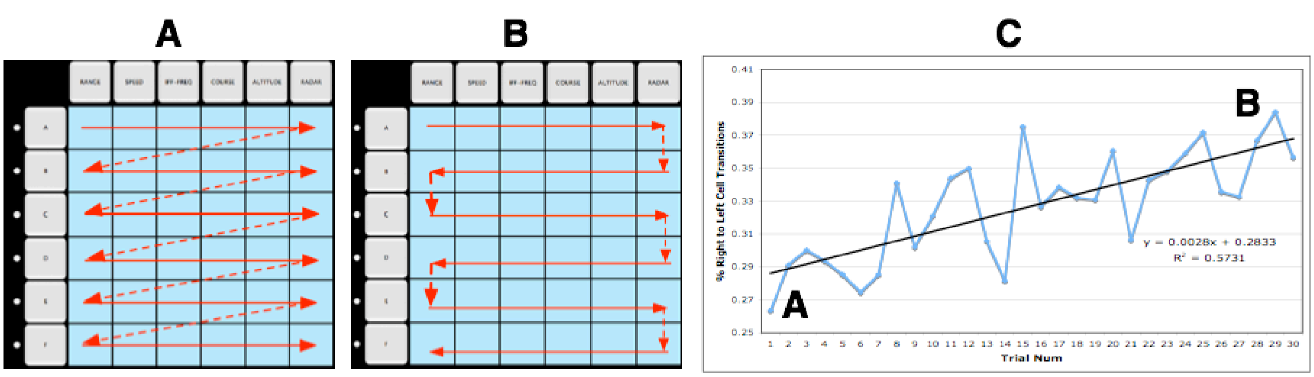


**Figure 5:** Different strategies in uncovering cell values. (A) Reading values left to right (B) Reading values alternating left-to-right and right-to-left (C) Percent of right-to-left cell click transitions as a function of trial number

interleaving which has not advanced much since the pioneering EPIC-Soar work of Chong in the late 90's (Chong, 1998a, 1998b; Chong & Laird, 1997). One method that has touched on interleaving of perceptual-motor and cognitive operators since then is Cognitive Constraint Modeling (Lewis, Howes, & Vera, 2004). Cognitive Constraint Modeling provides a description of behavior derived via constraint satisfaction. However, unlike Chong's work, this method is not at all concerned with how human interleaving strategies adapt through experience.

The absence of a mechanism that interleaves cognitive operators has led us to build models that do not change over trials but which bracket human performance (Gray & Boehm-Davis, 2000; Kieras & Meyer, 2000). Understanding the differences between these models offers some insight into how perceptual-motor and cognitive mechanisms might evolve across trials.

## Model 1: Non-Interleaved

This was an "out-of-the-box" model, composed of sequential productions that can roughly be divided into four categories. The first set of productions (Figure 6A) started each trial and switched between alternatives. The second set of productions (Figure 6B) was the workhorse of the model. This set of productions initiated the perceptual-motor operations of moving the mouse and visual attention to the various cells. It was also responsible for adding the values in the cells. It did this in a systematic left-to-right fashion for all alternatives. Thus, this model employed strategy A from Figure 5. The third set of productions (Figure 6C) compared a current alternative's value to the highest value so far and updated the model's memory of the highest alternative seen. The fourth set of productions (Figure 6D) only fired after each alternative's value had been computed and the model was ready to select its answer.

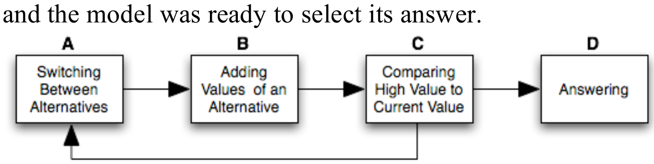


**Figure 6:** Workflow of the models

The model had declarative knowledge of addition facts and number relation facts. Thus, whenever it needed to compare whether a current alternative's value was greater than the highest value so far, it would search in its declarative memory for a relation fact involving those values.

The model also had a goal buffer that kept track of the alternative that it was currently scrutinizing, and an imaginal buffer that kept track of the highest value seen so far. Since it is beyond the scope of this paper to model human accuracy on this task, it sufficed for the model to hold the highest value and alternative in its imaginal buffer at all times, disallowing for any forgetting errors to occur.

This was the simplest model that encoded the task and for this reason, we did not expect its performance to match well with human data. It is termed non-interleaved because the perceptual-motor and cognitive operations were done largely sequentially and not interleaved with each other. We found that although this simple model failed to match duration times on the majority of trials, it did match duration times as compared to the first trial of human data (Figure 8, Model 1: Non-Interleaved).

## Model 2: Interleaving Cognitive with Perceptual-Motor Operations (I-CPM)

Examining the time plot of ACT-R's various modules over the course of a single inter-cell click interval (Figure 7), we noticed that gaps between production firings could be used to interleave perceptual-motor and cognitive operations. The interleaving was accomplished by firing productions that added the value of the last cell to the running total as the motor module was moving the mouse to the next cell. This interleaving saved ~100ms during each inter-cell click interval and decreased total trial duration by about 2.8 seconds from ~34.6s to ~31.8s, matching human duration times from Trial 2 (Figure 8, Model 2: I-CPM).

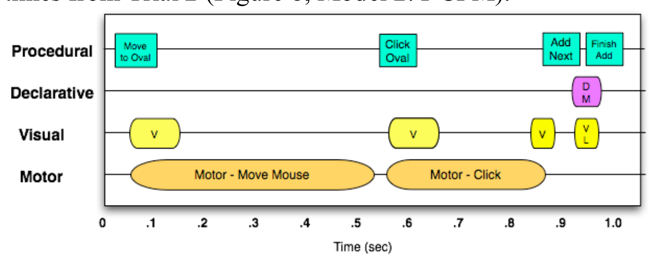


**Figure 7**: Time graph of between cell clicks in ACT-R for the non-interleaved model.

## Model 3: Interleaving Motor Preparation Time (I-CPM+MP)

Figure 7 also shows that the motor component (moving the mouse to the cell and then clicking it) comprised 80% of the time of the entire duration (786ms out of 931ms). What this means is that *just the motor component alone takes more time than the entire inter-cell click interval in the human data*.

During the course of a single trial, the repetitive sequence of moving the mouse to a cell and then clicking is done many times. In the human data, this practiced motor sequence became increasingly faster as attested by the decrease in the inter-cell click interval (see Figure 4). We decided to account for this increase in speed by taking advantage of ACT-R's motor module mechanisms.

When a motor command is issued to ACT-R's motor module, that command is executed in three phases: preparation, initiation, and execution. In cases where the model can tell ahead of time what movement will follow, it is beneficial to begin "preparing" the next movement before the current movement is finished executing. Thus, to account for the learning effects we observe in repeating the same two motor commands over and over, we allowed the model to begin preparing the next motor command prior to the finish of the current command. For example, while the move-mouse command was executing, the model already began preparing the mouse-click that would inevitably follow.

This motor preparation interleaving refinement of the model drastically decreased inter-cell click interval and, consequently, trial duration. The improvement decreased individual trial time by about 7.3 seconds, from ~31.8s to ~24.5s. This is a marked improvement over Models 1 and 2 and brought the model closer to the average human trial time of ~23.8s (Figure 8, Model 3: I-CPM+MP).

### Model 4: Alternating Transitions Between Rows (I-CPM+MP+R)

In a task in which interactive routines are on the order of hundreds of milliseconds, it is important to be able to determine where exactly it was that the model was incurring a large time cost. We therefore compared the inter-cell click interval analyses for human and model data. This comparison revealed that the major difference between human and model inter-cell click intervals was during the transitions between alternatives (where each alternative is a row in Figure 1).

As discussed earlier, participants' strategies changed over the course of the task. Initially, they clicked on cells in a left to right fashion whereas later they alternated the direction depending on their ending position in a given row. We thus incorporated this alternating behavior into the model thereby decreasing the distance the mouse had to move when a new alternative was encountered. Since move-mouse execution time in ACT-R is closely related to the distance that the mouse must move, as per Fitts' Law (Fitts, 1954; MacKenzie, 1992), this feature allowed the model to transition faster between alternatives (compare Figure 5 A and B).

This refinement in the model decreased total trial duration time by about 900ms from ~24.5s to ~23.6s (Figure 8, Model 4: I-CPM+MP+R). Although this was not a large difference, incorporating this component into the model makes it more cognitively plausible especially given that we see human participants exhibiting this shift in strategy.

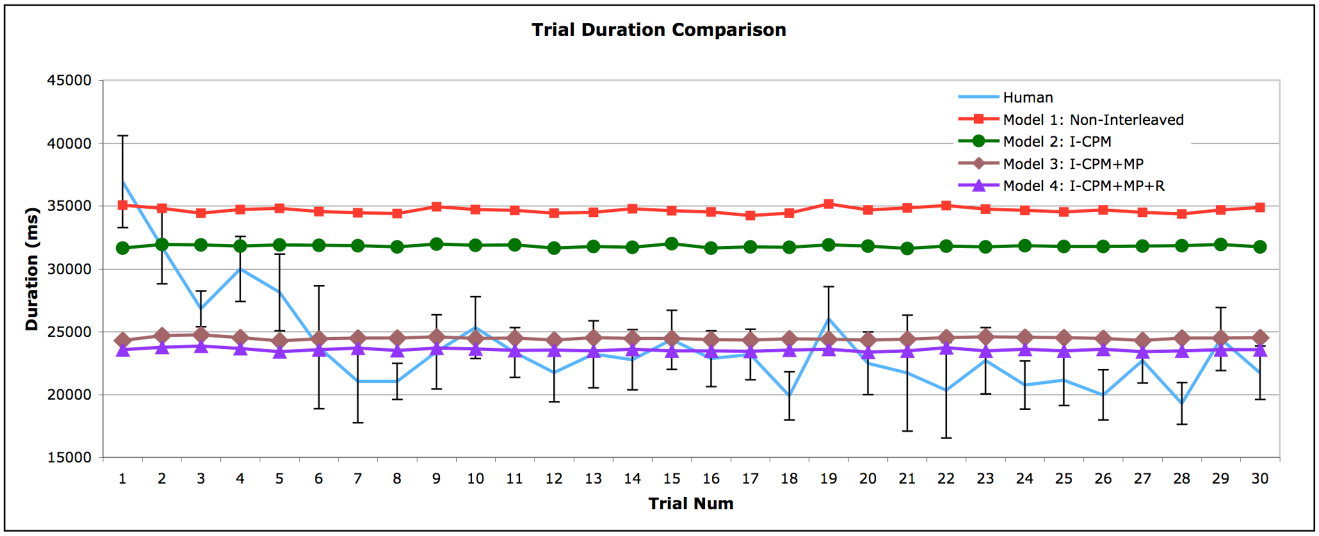


**Figure 8:** Comparison of Trial Duration Times Between the 4 Models and Human Data. Model 1: Non-Interleaved – Purely sequential model; Model 2: I-CPM – Interleave cognitive, perceptual, and motor operations; Model 3: I-CPM+MP – Interleaved motor preparation (MP) time added; Model 4: I-CPM+MP+R – Alternating transitions between rows

This final refinement of the model had the best fit to the asymptote performance in human data. Future work will include the model learning to choose between the two strategies.

## Conclusion & Future Work

Sometimes one can learn more from a modeling effort when the model does not fit the data than when it does. In fact, the lack of fit can tell us a lot about not only the limitation of the model itself and how to proceed to modify it but also about the limitation and error of the constraints with which the model was implemented. In this case, we wanted to investigate where and how the speedup in performance in humans occurs and in particular what it was about the "out-of-the-box" model that prevented it from matching human times.

### Where Does the Time Go?

According to human data, the inter-cell click interval varies as a function of trial number (Figure 3) and cell click number within a trial (Figure 4). We can see that the majority of inter-cell click intervals fall within the 600ms range. The first trials have longer durations as compared to the last trials, and the first few cell clicks in a trial take longer than subsequent cell clicks. However, the "out-of-the-box" model (Model 1) performs considerably slower in all cases, an average of around 950ms per inter-cell click interval.

One way we can speed up the model's performance is to interleave the cognitive and perceptual-motor components.

This results in at most a speedup of ~100ms per inter-cell click interval. However, if we look at human data particularly towards the end of the 30 trials (Figure 3), we see times of 520-600ms, which is considerably faster than the model's motor component alone, as per Figure 7, would allow.

Another way to speed up the model's performance is to interleave the motor preparation times with execution times. Since ACT-R does not do production compilation across perceptual and motor commands, there does not seem to be any other way of incurring this speedup in performance (Taatgen & Lee, 2003). The speedup afforded by this preparation interleaving results in a decrease of ~200ms per inter-cell click interval.

As per the soft-constraints hypothesis, a further refinement of the model altered how transitions between rows occurred. This resulted in an additional savings of ~30ms per inter-cell click interval.

Taken together, this modeling effort demonstrates the importance of millisecond-level considerations operating under even the simplest of tasks. The current model was intended to address the most perceptually motor intensive condition of the study. As such, it has led us to discover the crucial nature of interleaving and soft-constraints in attaining skilled performance.

The table task environment is a rich test bed for exploring how interactive routines in an information exploration/exploitation task evolve to produce skilled performance. Future modeling work of this task will explore how the different experimental conditions affect this evolution of interactive routines, and how these interactive routines influence performance in the decision-making task.

## Acknowledgements

Thanks to Hansjörg Neth and Christopher W. Myers for their invaluable contributions to this project. This work would not have been possible if not for the many useful discussions with Vladislav D. Veksler. The writing of this chapter was supported by a grant from the Air Force Office of Scientific Research AFOSR #FA9550-06-1-0074, Dr. Jerome Busemeyer, Program Officer.

## References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglas, S., Lebiere, C., & Quin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111*(4), 1036-1060.

Anderson, J. R., & Lebiere, C. (Eds.). (1998). *Atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Chong, R. S. (1998a). Modeling dual-task performance improvements with EPIC-Soar, *The Twentieth Annual Conference of the Cognitive Science Society* (pp. 1209). Hillsdale, NJ: Lawrence Erlbaum Associates.

Chong, R. S. (1998b). Modeling single-task performance improvement using EPIC-Soar, *Ninth Midwest Artificial Intelligence and Cognitive Science Conference* (pp. 10–16). Menlo Park, CA: AAAI Press.

Chong, R. S., & Laird, J. E. (1997). Identifying dual-task executive process knowledge using EPIC-Soar. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 107–112). Palo Alto, CA: Lawrence Erlbaum Associates.

Fitts, P. M. (1954). The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *Journal of Experimental Psychology, 47*(6), 381–391.

Fu, W.-T. (2007). A rational–ecological approach to the exploration/exploitation trade-offs: Bounded rationality and suboptimal performance. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 165-179). New York: Oxford University Press.

Gray, W. D. (2000). The nature and processing of errors in interactive behavior. *Cognitive Science, 24*(2), 205-248.

Gray, W. D., & Boehm-Davis, D. A. (2000). Milliseconds Matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experimental Psychology: Applied, 6*(4), 322–335.

Gray, W. D., Sims, C. R., Fu, W.-T., & Schoelles, M. J. (2006). The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior. *Psychological Review, 113*(3), 461-482.

Gray, W. D., Veksler, V. D., & Fu, W. T. (2004). *Probing the Paradox of the Active User: Asymmetrical Transfer May Produce Stable, Suboptimal Performance.* Paper presented at the Twenty-Sixth Annual Meeting of the Cognitive Science Society.

Kieras, D. E., & Meyer, D. E. (2000). The role of cognitive task analysis in the application of predictive models of human performance. In J. M. Schraagen, S. F. Chipman & V. L. Shalin (Eds.), *Cognitive task analysis* (pp. 237–260). Mahwah, NJ: Lawrence Erlbaum Associates.

Lewis, R.L., Howes, A., & Vera, A. (2004). A constraint-based approach to understanding the composition of skill. *International Conference on Cognitive Modeling,* Pittsburgh, 2004.

Lohse, G. L., & Johnson, E. J. (1996). A comparison of two process tracing methods for choice tasks. *Organizational Behavior and Human Decision Processes, 68(1),* 28-43.

MacKenzie, I. S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction, 7*(1), 91–139.

Taatgen, N.A. & Lee, F.J. (2003). Production Compilation: A simple mechanism to model Complex Skill Acquisition. *Human Factors*, 45(1), 61-76.