Time Interval Estimation: Internal Clock or Attentional Mechanism?

Niels Taatgen[1,2] (taatgen@cmu.edu), John Anderson[1] (ja@cmu.edu),
Daniel Dickison[1] (danieldickison@cmu.edu) and Hedderik van Rijn[2] (hedderik@ai.rug.nl)
[1]Carnegie Mellon University, Department of Psychology, 5000 Forbes Av.
Pittsburgh, PA 15213 USA
[2]University of Groningen, Department of Artificial Intelligence, Grote Kruisstraat 2/1
NL-9712 TS Groningen, Netherlands

Abstract

The human ability to accurately estimate time intervals in the order of 0 to 20 seconds can be explained by two seemingly incompatible theories: the internal clock and the attentional counter theory. Based on a dual timing task experiment we conclude that a symbiosis of both theories is necessary to explain all the phenomena found in our experiment and in the experiments we review. This conclusion is supported by the computational models we present of the experiments.

Interval Estimation

The human ability to routinely estimate short time intervals plays an important role in everyday life. Time estimates play a role in situations where we take an action and expect some response, for example when we click on a link in a web-browser, in real-world decisions, for example judging whether we should brake for a yellow traffic light or not, and in multi-tasking situations where we have to strategically switch between tasks, for example using a mobile phone in a car.

There are at least two theories that address interval estimation, the *attentional counter theory* and the *internal clock theory*. The attentional counter theory (Hick et al, 1977; Thomas & Weaver, 1975) assumes that there is a cognitive timer that counts subjective time events. Increasing the counter is a process that requires attention. As a consequence, if there are other processes competing for attention, the counter is increased less often, "stretching" time. For example, in an experiment by Zakay (1993), participants had to estimate a 12 second interval. In one condition, they had to estimate the interval while doing a secondary task, and had to reproduce it while doing nothing else. In the other condition, they had to estimate the interval while doing nothing else, but had to reproduce it while doing a secondary task. The secondary tasks were, in increasing level of complexity:

- No secondary task (ET, empty time)
- Reading color words (printed in black) (W)
- The Stroop task: Naming the color names of color words printed in incongruent ink (CW)
- Color-word associations: like the Stroop task, but now participants had to name a word associated with the ink color (CWA)

Figure 1 shows the results of this experiment. In the relatively easy ET and W conditions there is no effect of the secondary task, but in the more demanding CW and CWA tasks, time estimates are affected. According to the attentional counter theory, when the participants have to do a demanding secondary task during presentation of an interval, they can devote less attention to keeping track of time, resulting in a lower cognitive count and a shorter

reproduced interval. On the other hand, if they have to do a demanding task during the reproduction of the interval, their cognitive counting is slower resulting in a longer interval.
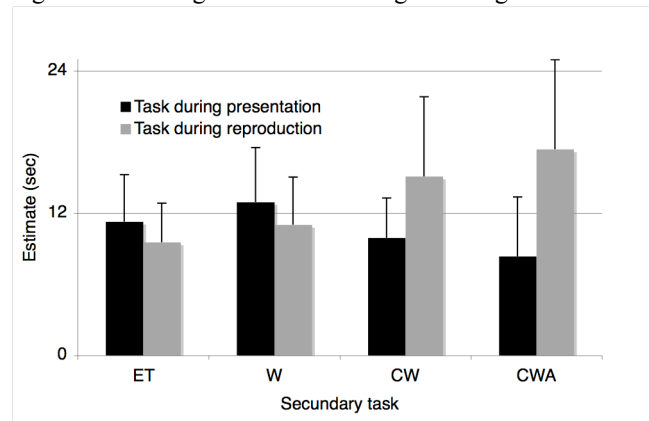

Figure 1: Results of Zakay's (1993) experiment.

The internal clock theory (i.e., Matell & Meck, 2000) states that the brain has devoted several areas that implement a time estimation system. The general idea is that certain stimuli can synchronize neurons in a certain area of the cortex, effectively acting as a starting sign. As each of the neurons produces its own particular pattern of activation over time, each interval is associated with a unique pattern of activation. These patterns can then be read out by other brain areas in the basal ganglia. Contrary to the attentional counter theory, the timing system does not require any attention, and errors in time estimation are mainly due to noise. In a typical interval timing experiment (Rakitin et al, 1998) participants were presented with intervals of 8, 12 and 21 seconds, which they then had to reproduce. Figure 2 shows the results of this experiment in terms of the distributions of the responses. Although the variance increases for larger intervals, the peaks of each of the distributions are exactly at the duration of the interval that the participants had to estimate.

An important difference between the Zakay (1993) and Rakitin et al. (1998) experiments is that in the Zakay experiment each participant produced exactly one time interval, while in the Rakitin et al. experiment they produced 80 intervals with feedback on the true duration every few trials. Although in the Rakitin task there is a secondary task to prevent participants from counting, this task is always the same and doesn't produce the distortions in time perception in the Zakay experiment. When just looking at these two experiments, the attentional counter theory is consistent with both, but the internal clock theory only with the Rakitin experiment. On the other hand, many

practical examples of time perception seem to be highly automated (for example many timing aspects of driving a car), giving some credibility to an internal clock mechanism. The Zakay experiment mainly proves there is some effect of attention, but not necessarily that the role of attention is to keep an explicit cognitive count. The internal clock theory by itself does not really deal with attention, or other aspects of cognition. With this in mind we (Taatgen, van Rijn & Anderson, 2004) designed an internal clock module for the ACT-R architecture. This module can not only model timing experiments like the Rakitin et al. (1998) experiment, but can also shed some light on how timing interacts with other aspects of cognition, including attention.
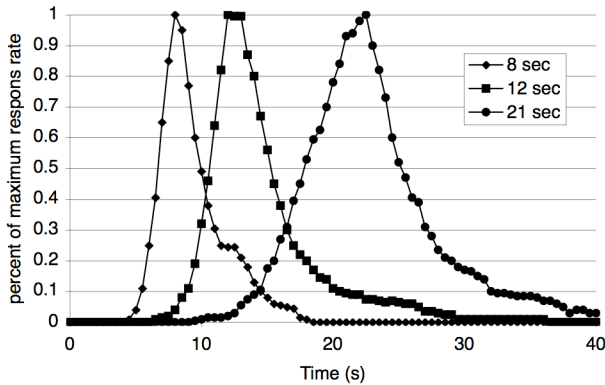

Figure 2: Distribution of responses adapted from Rakitin et al. (1998) for intervals of 8, 12 and 21 seconds.

The goal of our present study is to reconcile the experimental results from the attentional and internal clock theories. Our model will use an internal clock that will keep track of time on its own account. To keep track of time, however, this clock has to be explicitly attended, at least in the initial stages of learning.

The Temporal Module

The general idea, based on Matell and Meck (2000), is that an internal timer can be started explicitly to time the interval between two events. A reset event sets an integer counter to zero, after which it is increased as time progresses. The temporal module acts like a metronome, but one that starts ticking slower and slower as time progresses. The interval estimate is based on the number of ticks the metronome has produced. More precisely, the duration of the first tick is set to some start value:

$$t_0 = \text{starttick}$$

Each tick is separated from the previous tick by an interval that is a times the interval between the previous two ticks. Each interval has some noise drawn from a logistic distribution added to it. The distribution of this noise is determined by the current tick-length.

$$t_{n+1} = at_n + \text{noise}(mean = 0, sd = b \cdot at_n)$$

Suppose we want to reproduce a time interval, as represented in the first horizontal bar in Figure 3 that is

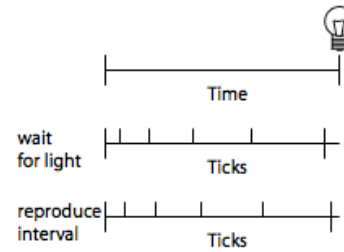defined as the time between the start of a trial and the moment a light comes on.

Figure 3: Illustration of the temporal module

The timer is initiated at the start of the trial. When the light comes on, the value of the timer (5 in the example) is read and stored. When the interval has to be reproduced, the value of the timer perceived earlier is used. We have estimated values for the parameters in these equations to obtain an optimal fit to the Rakitin et al. (1998) experiment of interval estimation (reported in Taatgen, van Rijn & Anderson, 2004). We found 11 ms for start tick, 1.1 for a, and 0.015 for b. These values also provided excellent fits to the other experiments discussed in that paper.
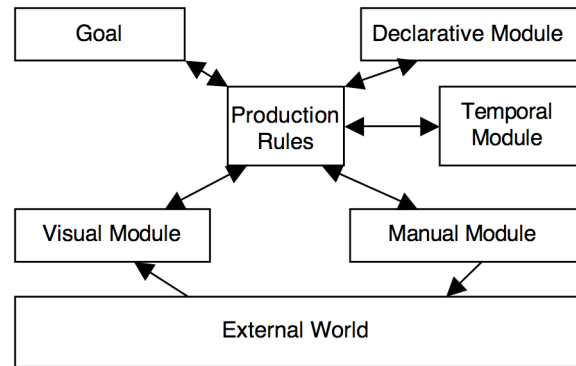

Figure 4: Outline of ACT-R

Figure 4 illustrates the role of the temporal module within the ACT-R architecture (Anderson et al., 2004). ACT-R is a cognitive architecture based on production rules. What is of importance in relationship to the issue of attention is that production rules have to start and read the timer in the temporal module. This means that errors in time estimates are not only due to noise in the temporal module itself, but also due to the production system not initiating or reading the module at the right moment. This latter type of "noise" is especially important in the initial stages of skill acquisition, when the length of the interval has not been established yet. In order to investigate the various contributions to time estimation, we designed the following experiment.

Experiment

Method
Participants 17 males and 15 females from the Carnegie Mellon student population volunteered to participate in the experiment
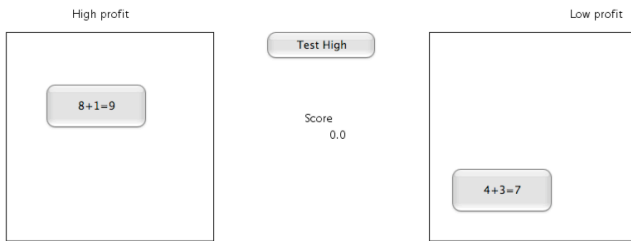
Experimental task

Figure 5: the dual timing task

Figure 5 outlines the task. The display consists of two boxes, a high profit box on the left, and a low profit box on the right. In each of these boxes stimuli can appear, to which the participant has to respond. Stimuli are buttons with, depending on the condition, an addition with one-digit numbers or a letter. Additions can either be correct or wrong by one, and letters are either "A" or "B". Participants have to respond to correct additions and "A"'s by clicking on them when they are in the right box, or by pressing space on the keyboard with the left hand when they are in the left box. In the right box stimuli appear for 1200 ms, with between 0 and 2000 ms in between. Stimuli in the left box do not appear by themselves: they have to be requested by the participant by pushing the "Test High" button. Stimuli in the left box are available in certain time periods, basically six seconds on and seven seconds off. So at the start of the experiment there are six seconds in which there are stimuli available, then seven seconds without available stimuli, then six seconds again with stimuli, etc. The end of a six second period is always marked in the left box with a brief appearance of the word "End". Whenever the participant presses the "Test High" button during a period with stimuli, stimuli actually appear in the left box for the remainder of that period. When the participant presses the "Test High" button during the seven second period without stimuli, nothing happens. Optimal behavior is therefore to press the "Test High" button right at the beginning of the six second period, which is exactly seven seconds after the word "End" appeared in the left box. Stimuli in the left box appear for 1200 ms with 300 ms in between. A trial lasts 120 seconds. Participants score 30 points for each correct response in the right box, 100 points for each correct response in the left box, they pay 10 points for each click on the "Test High" button, receive 30 penalty points for an incorrect response in the right box, and 100 penalty points for an incorrect response in the left box. Participants were instructed on all the aspects of the task, except the durations of the intervals, which they had to determine by trial-and-error.

Design The experiment has four conditions: 15 trials of 120 seconds with the letter task (LL), 15 trials with addition task (AA), 10 trials with the letter task followed by 5 trials with the addition task (LA), and 10 trials with the addition task followed by 5 trials with the letter task (AL).

Results

The two solid lines in Figure 6 plot the distributions of the moments at which participants first click the test button. These moments are defined as the deviation from the optimal time, that is, the time at which new high profit

stimuli become available. The data are averaged over the two conditions that start with the letter task and the two conditions that start with the addition task, and are plotted separately for trial 1-5 (block 1) and trial 6-10 (block 2). The dotted line plots the distribution that would be expected if this were a pure interval estimation experiment like the Rakitin et al. (1998) experiment plotted in Figure 2. These plots show that participants do worse than that ideal, but also show that the task they are doing, letter or addition, has only a small impact.
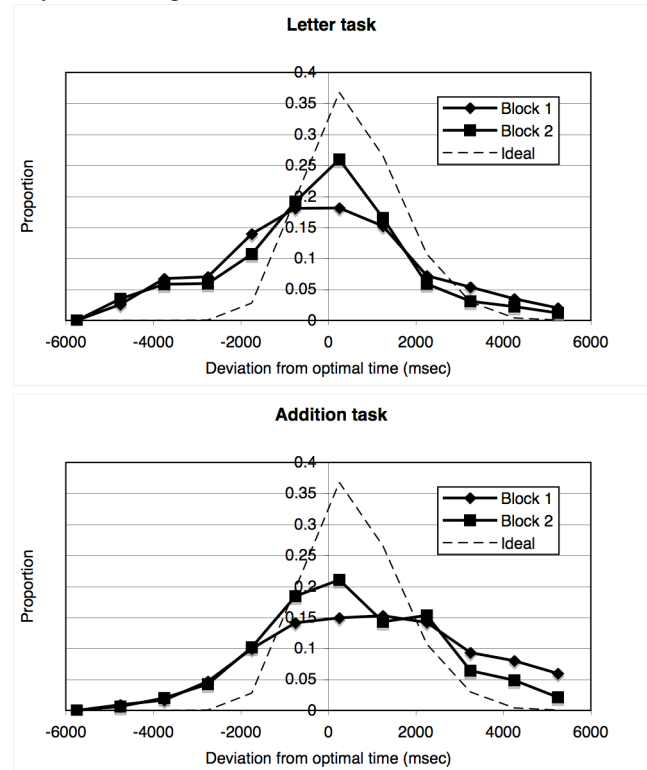

Figure 6: Distribution of first-click times for both tasks compared to an expected distribution of a pure interval estimation experiment.

To get a more precise idea of the impact of learning and condition, we fit a linear mixed-effect model (Laird & Ware, 1982) with the absolute deviation of the time of first click from the optimal moment to click as dependent variable, and condition and trial number as independent variables. Averages for these absolute deviation values, aggregated in three blocks, are plotted in Figure 7. The analysis revealed a significant learning effect (p=0.001), no effect of initial task (Letter or Addition), no effect of changing from the Addition to the Letter task (condition AL), but significant slowing effect in the Letter to Addition condition (condition LA, p=0.013). Similarly, we looked at the average times relative to the click moment. Although the average time offers no indication of performance, it can reveal shifts in timing due to change in task. Figure 8 shows the basic outcomes. Again we fit a linear mixed-effect model to these data, revealing no learning effect, an effect of initial task with Addition producing later responses than Letter (p=.0011), no effect of changing from Addition to Letter, but a significant later response of changing from Letter to Addition (p<0.0001).
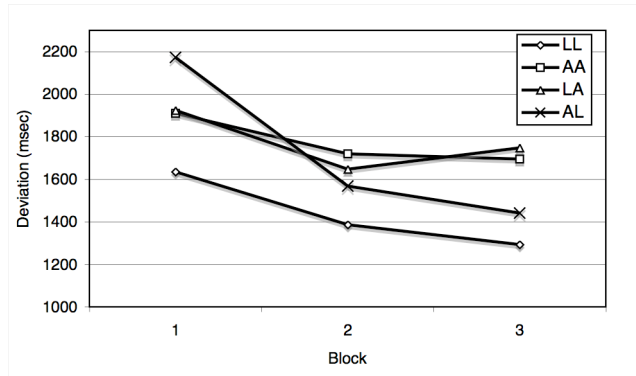
Figure 7: Average absolute deviation from optimal click moment for the four conditions. Note that the change of task for the LA and AL conditions is between block 2 and block 3
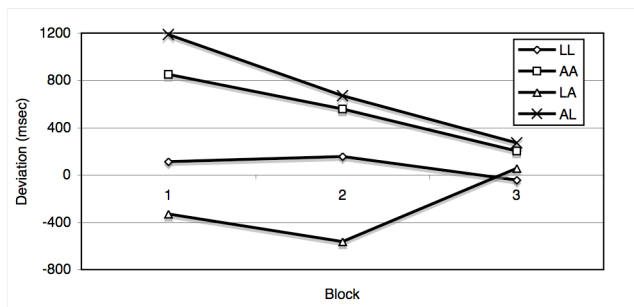

Figure 8: Average deviation from optimal click moment for the four conditions
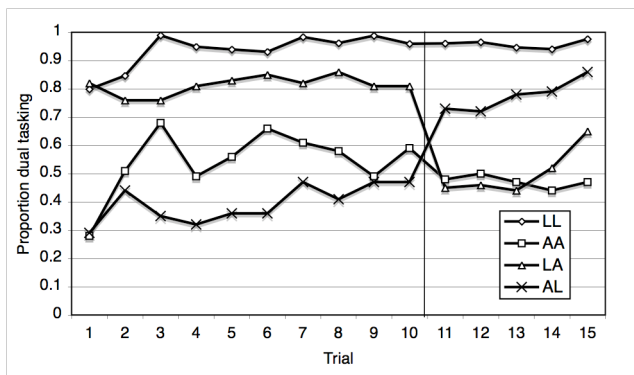

Figure 9: Proportion dual-tasking for the four conditions

To summarize, the effect of task difficulty seems to be rather small. If the attentional mechanism theory would be right, we would expect to see a lower accuracy for the Addition task than the Letter task, because less attention can be devoted to keeping track of time, and we would also expect significant shifts in timing after a task change (in the LA and AL conditions). Such a shift can only be found in the LA condition, and it is relatively small compared to the deviations found in the Zakay (1993) experiment. A possible explanation is that both tasks are just too easy: Zakay only found an effect in the more difficult secondary tasks. It is therefore useful to look at the amount of dual tasking that participants manage to do at intervals that there are stimuli in both the left and the right box. A measure of

dual-tasking can be obtained by looking at the periods that there are stimuli in both the left and the right box. As stimuli in the left box produce higher scores, we assume that people will only react to stimuli in the right box if they have spare capacity to do so. We therefore took as a measure of dual tasking the proportion of target stimuli in the right box to which the participant responded while there were also stimuli in the left box. Figure 9 shows the results: participants turn out to be able to achieve a high level of dual-tasking in the Letter task, but only around 50% in the Addition task. This shows that the Addition task does indeed require more attention than the Letter task, making a simplicity-based explanation unlikely.

Our eventual result is consistent with neither the attentional mechanism theory (which would predict much larger impacts of the task-difficulty manipulations) nor the internal clock theory (which would predict no influence at all). However, it is consistent ACT-R's temporal module, which predicts small influences of difficulty, because the temporal module cannot be attended as often as necessary when the task is more difficult. In order to show this in more detail, we have constructed a computational model of the task.

The Model

The model builds on earlier models of time estimation (Taatgen, van Rijn & Anderson, 2004), dual tasking (Anderson, Taatgen & Byrne, submitted) and skill acquisition (Taatgen & Lee, 2003). We will explain the model at a fairly global level.

Time Estimation

Because the duration of the interval is unknown, the model has to determine it by trying out various intervals. When the model sees "End" in the left box, signaling the start of the interval, it starts the internal clock, which starts generating time ticks as illustrated in Figure 3. Whenever the model has some slack time to think about time (the details of which we will discuss in the next section), it attempts to retrieve a previous experience of pressing the test button at approximately the present time. If a successful experience is retrieved, the model will initiate a click on the test button. If a failed experience is retrieved, the model will do nothing. Finally, if no experience at all is retrieved for the present time, the model randomly decides to press the button or not. After the button has been pressed, the model judges whether the button-press was successful. If stimuli appear in the left box, the present time is stored as successful, but if nothing happens the present time is stored as a failure. Note that late test-button presses are judged as successful, even if they are for example 4 seconds late, but that early presses are judged as failures, even if they are early by only 100 ms. As the model accumulates more experiences it will become more accurate at estimating the right interval, but only within the boundaries of the accuracy of the temporal module itself (i.e., what is depicted by the dotted pure interval estimation distribution in Figure 6).

Multi Tasking

The model has to divide its attention between three tasks: attending and responding to the left box, attending and

responding to the right box, and attend the time. Only two of these tasks are relevant at the same time: either the left box and the right box have to be attended, or the right box and the time. The model is mainly event-driven, and will respond to changes on the screen. When a new stimulus appears on the screen, the model will attend to it and initiate a response. The only exception is when the model is busy with a stimulus on the left box, because then it will ignore stimuli in the right box until it is done with the stimulus in the left box. A stimulus in the left box can on the other hand interrupt processing in the right box. Attending to the time interval is initiated whenever the model has no stimulus to process. However, since retrieving a past experience takes time, especially when these experiences are relatively new and still have a low activation, attending the time can be interrupted if a new stimulus appears on the screen.

Skill acquisition

ACT-R's rule learning mechanism, production compilation, will gradually learn rules that correspond to cognitive actions that are repeated often. For the present model, the main rules of importance are rules that are associated with retrieving previous experiences from memory. Initially, the process of judging whether or not to click at a certain time requires three steps:

1. A production rule makes a retrieval request to declarative memory for a past experience about the present time (50 ms)
2. Declarative memory tries to find a matching experience (can take up to 680 ms, depending on the activation of the experience)
3. A rule that acts upon the retrieved experience fires (50 ms)

At any moment this process can be interrupted by a new visual stimulus. However, once a certain experience has been retrieved often enough, the three steps are collapsed into a new rule specialized for that experience, for example:

IF the current time is 41 time ticks
THEN press the test-button

This new rule only takes 50 ms to execute, and has therefore a much higher probability to fire at the appropriate time.

Model results

We ran the model 200 times for each of the four conditions, and then averaged the results. The model produces time distributions that are quite similar to the distributions found in the experiment (Figure 10). Although the differences between the two distributions are subtle, it shares two characteristics with the experimental data: the peak of the distribution is slightly higher for the letter task, and there is a slight shift to the right in the distribution of the Addition task. These two aspects become also clear if we look at the average response times and average deviations from the optimal time. Figure 11 shows the average deviations from the optimal response time, and Figure 12 the average response times relative to the optimal response time. Although the model results are far less noisy than the data, we can see the same effects that we saw in the data: a learning effect in the deviations but not in the response times, a small effect of task difficulty, a small effect of

making the task harder, and an even smaller (and in the data insignificant) effect of making the task easier.

The model's changes in behavior due to task difficulty are mainly due to the fact that the model has less (or more) time to occasionally check the time. Although the internal timer might have a good estimate of the time, this will not help if there is no production rule that reads it at the right moment. Later in the experiment, when the model starts learning production rules, this problem becomes smaller because it starts to learning specific timing rules that do not require retrievals from declarative memory, and are therefore less susceptible to interruption.
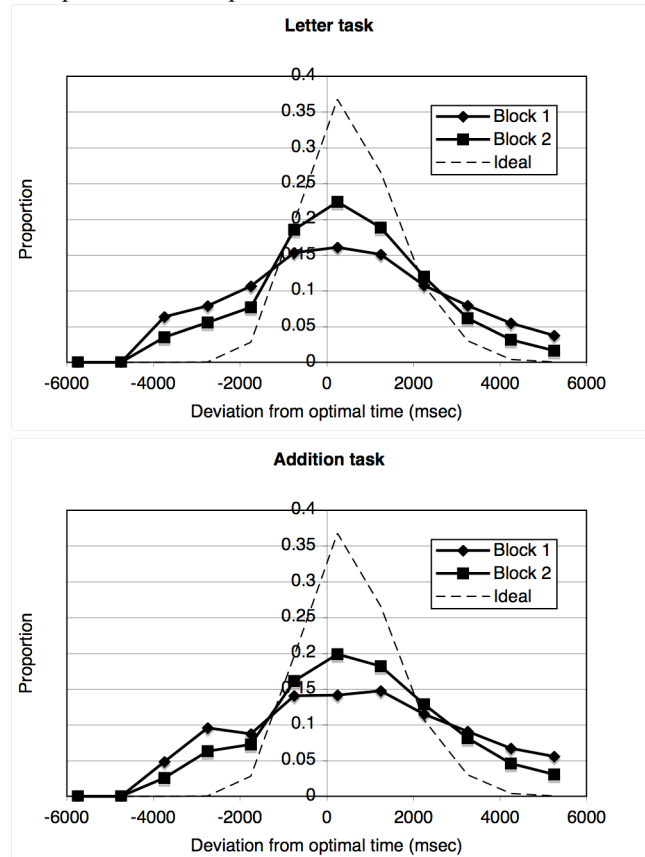

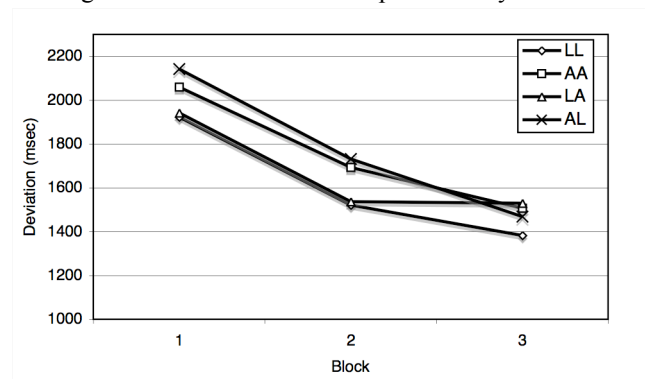Figure 10: Time distributions produced by the model

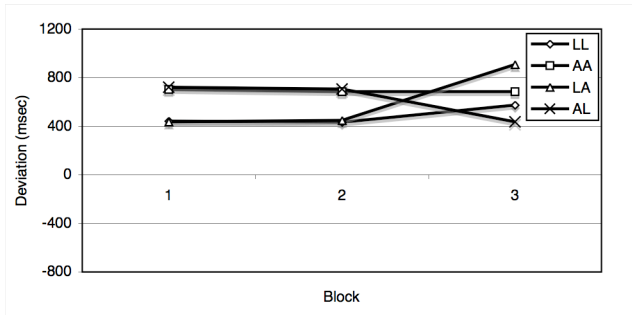Figure 11: Average absolute deviation from optimal click moment for the model

Figure 12: Average deviation from optimal click moment for the model

Figure 13, finally, shows the dual-tasking performance by the model, which is quite consistent with the dual-task performance by the participants.
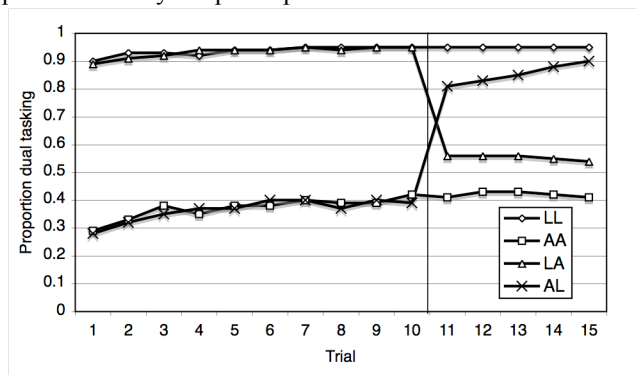

Figure 13: Proportion of dual tasking by the model

Discussion

The model shows the same small effects of task difficulty as the participants in the experiments. In the case of the model, the generally slightly longer times for the Addition task can be explained by the fact that the Addition task itself needs more time to finish, leaving smaller intervals in between stimuli in which the time-retrieval process tries to retrieve old experiences, and thus increasing the probability that this process is interrupted.

Another interesting aspect of the model is that the model learns production rules that handle timing, reducing the need for attention in situations where time intervals are well practiced. This aspect is of importance we want to use it to model situation in which a sense of timing is automated (e.g., the traffic light situation).

A final issue is how we can explain Zakay's (1993) estimation effects with a clock model, as these effects are often quoted as supportive of an attentional counter theory. To do this, we have to make one extra assumption: given the fact that the participants estimate this interval just once, they are prone to making all sorts of "startup" mistakes. One possible mistake is that the temporal module used for one of the secondary tasks, effectively resetting it to zero, and the probability for this becomes larger as the secondary task becomes more demanding. Based on this assumption, the model results in Figure 14 can be produced (other assumptions could probably produce similar predictions in combination with the temporal module).
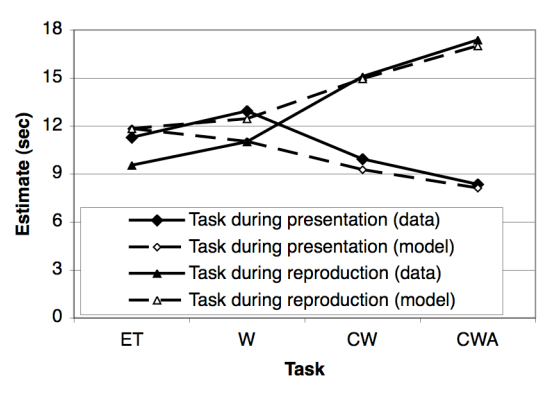

Figure 14: Zakay task, comparison between model and data

In summary, a model consisting of an internal clock combined with a general cognitive architecture attending this clock provides explanations for both the in general accurate human capacity for timing intervals and effects of attention due to secondary tasks.

Acknowledgments

References

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

Anderson, J.R., Taatgen, N.A. & Byrne, M.D. (submitted). Learning to Achieve Perfect Time Sharing: Architectural Implications of Hazeltine, Teague, & Ivry (2002)

Hick, E.R., Miller, G.W., Gaes, G. & Bierman, K. (1977). Concurrent processing demands and the experience of time in passing. *American Journal of Psychology, 90,* 413-446.

Laird, N.M. & Ware, J.H. (1982). Random-Effects Models for Longitudinal Data, *Biometrics, 38,* 963-974.

Matell, M.S. & Meck, W.H. (2000). Neuropsychological mechanisms of interval timing behavior. *BioEssays, 22,* 94-103.

Rakitin, B.C., Gibbon, J., Penney, T.B., Malapani, C., Hinton, S.C. and Meck, W.H. (1998). Scalar Expectancy Theory and Peak-Interval Timing in Humans. *Journal of Experimental Psychology: Animal Behavior Processes, 24,* 15-33.

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

Taatgen, N.A., van Rijn, H. & Anderson, J.R. (2004). Time perception: beyond simple interval estimation. *Proceedings of the sixth international conference on cognitive modeling* (pp. 296-301). Mahwah, NJ: Erlbaum.

Thomas, E.A.C. & Weaver, W.B. (1975). Cognitive processing and time perception. *Perception & Psychophysics, 17,* 363-367.

Zakay, D. (1993). Time estimation methods–do they influence prospective duration estimates? *Perception, 22,* 91-101.