A Cognitive Model of a Temporal Binding Task

Laura Saad (laura.saad@rutgers.edu)
Department of Psychology, Rutgers University. New Brunswick, NJ

Alexander R. Hough (alexander.hough.1@us.af.mil)
Leslie M. Blaha (leslie.blaha@us.af.mil)
Air Force Research Laboratory, Wright-Patterson AFB

Christian Lebiere (cl@cmu.edu)
Department of Psychology, Carnegie Mellon University. Pittsburgh, PA

Abstract

Temporal binding (TB) is the subjective compression between a voluntary action and its associated outcome. It is regarded as an implicit measure of the sense of agency; however, an underlying mechanism has yet to be agreed upon. Previous research suggests memory as an alternative explanation for TB in two publicly available datasets. We test this idea by implementing a model within the ACT-R cognitive architecture and leveraging its existing memory and time perception mechanisms to simulate participants from these datasets. Our simulations provide evidence to suggest that memory and time perception mechanisms can explain the pattern of results. Implications for temporal binding and the sense of agency are discussed.

Keywords: Temporal Binding; ACT-R; Cognitive Models; Sense of Agency; Time Perception

Introduction

Temporal binding (TB) is assumed to be an implicit measure of the sense of agency. TB is defined as the perceived subjective compression of time between a voluntary action and its associated outcome (Haggard, Clark, & Kalogeras, 2002). In the seminal study by Haggard et al. (2002), participants were asked to press a button at a time of their choosing; a few hundred milliseconds later, there was an audible tone. Participants were then asked to estimate the timing of their button press and the tone. The key finding was that when the action is voluntary—as opposed to involuntary—participants subjectively estimated that their button press occurred later than it objectively did. Furthermore, participants also subjectively estimated that the tone occurred earlier than it objectively did. This compression, or underestimation, of the subjective time interval between the action and its outcome is what is known as *temporal binding*. Importantly, the opposite effect, or a repulsion of the subjective time interval, occurred for involuntary actions (e.g., finger twitch produced by transcranial magnetic stimulation of the motor cortex) and their outcomes. This difference in the pattern of results between these two conditions led to the conclusion that TB is an implicit marker for the sense of agency.

There is some theoretical debate in the literature over whether it is the presence of a voluntary action, and therefore intentionality (Haggard, 2005), or the perceived causality between events (Hoerl et al., 2020) that is necessary to elicit TB. One reconciliatory explanation for TB comes from cue integration theory (Ernst & Banks, 2002). This theory suggests the motor system optimally combines cues from different sources to reduce the overall variability of estimates. Cues are weighted by their reliability such that information from more reliable cues is more heavily weighted in the integration process. There has been one successful formal implementation of a Bayesian cue integration model in the context of TB (Legaspi & Toyoizumi, 2019); it remains unclear if this model can be applied to all timing estimation methods.

Though most TB tasks involve free recall, the role of memory has been largely underexplored in this literature. Recently, a memory process was proposed as a potential explanation for TB (Saad, Musolino, & Hemmer, 2022). In this paper, a regression pattern was revealed by re-plotting participant estimates from two publicly available datasets (Weller, Schwarz, Kunde, & Pfister, 2020) as the difference between the subjective responses and objective values (i.e., bias). Regression here refers to the bias in estimations such that participants, when making estimations, select a value closer to the mean of intervals observed in the task. This regression pattern replicated across conditions regardless of the agency manipulation. Saad et al. (2022) then successfully simulated participant estimates using a Bayesian rational memory model. This provided the first evidence that a memory mechanism could account for estimations at the aggregate level in a TB task.

Relatedly, the role of time perception in eliciting the TB effect has also been understudied in this literature. During encoding, participants perceive the timing of or the intervals between events. One mechanism that has been proposed to explain this is a pacemaker-accumulator process, where a pacemaker produces pulses at some rate, and these pulses are counted in an accumulator. The perceived length of the interval between two events is a function of how many pulses are in the accumulator; more pulses correspond to a longer perceived duration. This mechanism makes a prediction that a shortening of a perceived time interval (i.e., the compression characteristic of the TB effect) is a result of a slower pulse rate leading to fewer pulses in the accumulator.

Fereday, Buehner, and Rushton (2019) empirically investigated whether internal clock slowing is a viable explanatory mechanism for TB. In two experiments, the authors incorporated a temporal discrimination task where participants compared durations of causal (button press and a flash) and non-causal trials (two flashes) to a reference duration (black square presented on screen) and were asked to report which interval length was longer. The authors calculated point of subjective equality (PSE) values across conditions. The PSE

value represents the duration of the comparison interval (i.e., causal or non-causal) that is perceived as the same as the reference interval 50% of the time. The prediction is such that lower values of PSE correspond to more compression or underestimation (i.e., binding). The authors reported evidence to support this prediction in both experiments.

Although the pacemaker-accumulator process makes explicit predictions which can be tested empirically, a formal model of this process has never been implemented. Importantly, though much of the TB literature is focused on developing a theory of agency, there has not yet been an investigation into how memory and time perception processes may work together to influence or explain the temporal binding effect. We aim to do just this. We hypothesize the regression pattern in the human data results from a memory mechanism where participants estimate time intervals according to a pacemaker-accumulator process and then use estimates from previous trials during recall. We develop a cognitive model to test this hypothesis; the model specifies mechanisms for memory and time perception and is capable of simulating human performance in one condition in a TB task. We then explore how estimating different parameters related to memory mechanisms in the model affect simulation results at the aggregate and individual level.

We focus first on simulating one trial-type, action trials, because they represent the most frequently used trial-types in this literature. Additionally, our initial aim was to establish a cognitive model as a viable means for simulating human behavior in these tasks. These results lay the groundwork for future simulations using the same cognitive model to simulate the passive, comparison trial-type and therefore the entire temporal binding effect. We discuss this and other ideas for future work in the final section of the paper.

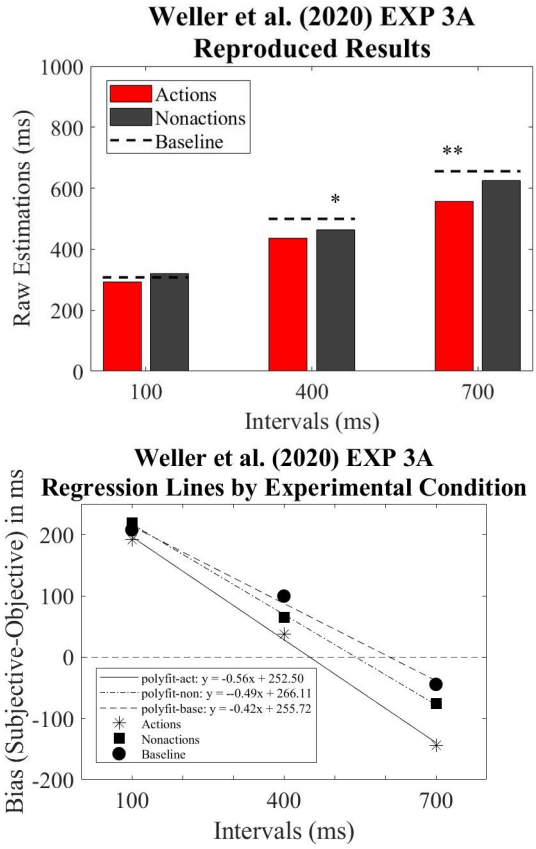

Figure 1: Reproduction of graphs from Weller et al. (2020) depicting mean raw estimations across three trial-types (top panel). Re-plotting these estimates as bias (bottom panel) reveals a consistent regression pattern across intervals. Data plotted here are from 27 participants. * $p < 0.05$, ** $p < 0.01$

Method

Data Set

We model the publicly available data from experiment 3A in Weller et al. (2020). Code for all analysis, figures, and supplementary material included in this paper are also publicly available (https://osf.io/6bkjp/). A detailed description of the experimental method and procedure can be found in the original Weller et al. (2020) paper.

We briefly describe the procedure for experiment 3A. The experiment included three trial-types: action, non-action, and baseline. At the beginning of each trial, participants were asked to choose between an action and a non-action which would each produce distinct outcomes. These trials were called operant trials. During non-action trials, participants chose not to act and a default outcome would occur; in the action trials, participants acted by pressing a button at a timing of their choosing to change the default outcome. In the baseline trials, there was no initial decision necessary, and participants passively watched two events unfold: a progress bar filled which ended in a "click" sound; then a ball launched in a pre-specified direction.

At the end of each trial, regardless of which type, participants were asked to recall and report their estimate of the interval between two events (i.e., either the keypress and ball launch in action trials, or the clicking sound and ball launch in non-action and baseline trials) in milliseconds using a slider on-screen. Three different time intervals were used between events: 100ms, 400ms, and 700ms. The presentation of these intervals was randomized across the different blocks of trials.

At the beginning of the experiment, participants completed a series of 20 practice trials, 10 baseline and 10 operant. Practice trials included time intervals between 100ms and 1000ms in steps of 100ms. Participants received feedback about the accuracy of their estimation at the end of each trial. Data was not collected or analyzed for practice trials. During the main experiment, no feedback was given.

Weller et al. (2020) reported statistical results comparing TB values across trial-type and delay (N=27), and reported two significant results for experiment 3A: more binding (or more compression) for actions compared to baseline at the 700ms interval and for non-actions compared to baseline at the 400ms interval. No other comparisons were significant. From these results, Weller et al. (2020) concluded that "temporal binding ha[d] [also] emerged for non-actions" (p. 8).

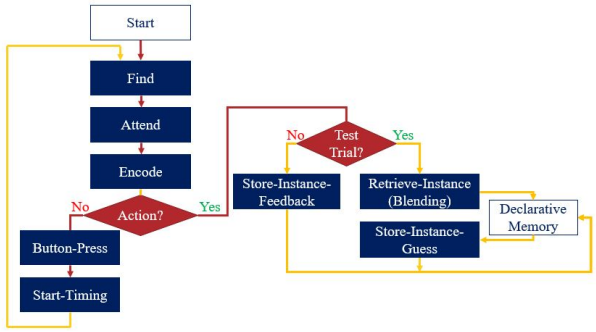

Figure 2: Visualization of TBM model processes.

Figure 1 reproduces the original visualization of Weller et al.'s results (top) and the regression effect (bottom) when participants' raw estimates are re-plotted as bias, or the difference between the average estimates and the actual length of each interval. When the estimates are re-plotted as bias, a clear regression pattern is revealed on all three trial-types.

The Temporal Binding Memory (TBM) Model

We developed a cognitive model called the Temporal Binding Memory (i.e., TBM) model based on the action trials from the Weller et al. (2020) TB task. The model was implemented in ACT-R, which is a hybrid cognitive architecture used to understand and simulate human cognition. ACT-R contains a set of modules which perform distinct cognitive functions and communicate via requests relayed through limited-capacity buffers. The TBM model used the imaginal, goal, vision, motor, procedural, declarative memory, and temporal modules.

The imaginal module holds the problem representation, and the goal module represents the model's current task focus. The vision module represents a visual attention system containing both a "what" and a "where" subsystem. The motor module represents two hands on a virtual keyboard. The procedural module, or production system, is a pattern matching system which constantly searches for productions matching the current state of the buffers using conditional statements (i.e., if-then rules). Only one production can be executed at a time. When a production is executed, or "fired", the state of the system changes, progressing the model through a task. Knowledge is represented in the form of chunks in declarative memory, which each have an activation value corresponding to the recency and frequency of the chunk with some noise. Chunks are retrieved via the retrieval buffer in the declarative memory module which searches through declarative memory to find a chunk with the highest activation value to satisfy the current request.

The temporal module, created to represent subjective time estimation between two events (Taatgen, Van Rijn, & Anderson, 2007), models time perception as a pacemaker-accumulator process. The pacemaker generates pulses, and the accumulator counts them. Tick lengths are noisy and increase in duration as time progresses, which means the temporal module is more accurate for shorter compared to longer time intervals. Tick lengths are based on the following equations for the n^{th} tick.

$$t_0 = \text{start} + \varepsilon_1 \tag{1}$$

$$t_n = a * t_{n-1} + \varepsilon_2 \tag{2}$$

The length of the first tick, t_0, is controlled by the *start* parameter (default = .011) with some noise. The a parameter (default = 1.1) affects the length of subsequent ticks. Noise is added to tick lengths using the act-r-noise command, and the s values for each are according to the following equations:

$$\varepsilon_1, s = b * 5 * \text{start} \tag{3}$$

$$\varepsilon_2, s = b * a * t_{n-1} \tag{4}$$

The b parameter is set to 0.015. The model recalls the intervals on each trial by accessing the current pulse value in the temporal buffer and reporting the tick count.

ACT-R models are cognitive models that specify mechanisms at the algorithmic level (Marr, 1982) and therefore require simulation of both the task and cognitive processes. As this is the first cognitive model of a TB task (to our knowledge), we aimed to replicate the major components of the experimental design (i.e., interval lengths, outcome modality, presence of feedback, and practice trials). However, we simplified stimuli and simulated action trials first, as these are the type of trial in which TB has been reported most frequently.

Figure 2 displays the steps the model completed to estimate time in our modified task. Each rectangle represents a separate production. Letters were used as cues for the beginning and end of the interval that was timed by the model. During both practice and test trials, the presentation of the first stimulus, "A", initiated the proceeding course of events. First, the model looped through a standard find-attend-encode loop by which the visual system located and then encoded the visual information on the virtual screen. After encoding, the model pressed the "A" key on the virtual keyboard, initiating the timing process by making a temporal buffer request to start timing in ticks. This triggered the presentation of another visual stimulus, "Z". The same find-attend-encode procedure was completed in response to this second visual stimulus before proceeding to the store-instance-feedback production. The tick count was stopped once the letter "Z" was perceived.

During each run, the model completed two different types of trials: 20 practice and 150 test. The key factor differentiating practice and test trials is that during practice trials, the model received feedback in the form of a real-time millisecond interval value which was then paired with the tick count from the temporal buffer and stored in declarative memory. During test, no feedback was given, and the tick count value on each trial was paired with a guess. This guess was informed by the chunks in memory that were formed during practice. We will now describe each trial type in detail.

Practice trials heavily influenced model performance as they generated chunks that were later retrieved by the model

and were the basis for the model's responses during test trials. During practice trials, feedback was provided after estimations to mimic the participants' experience in the original study (Weller et al., 2020). This feedback was used to build chunks which contained two slots: tick count (from the temporal buffer) and real-time delay which stored the actual delay given as feedback. Each chunk was encoded into declarative memory, and the entire process was repeated. The length of the interval between the action and outcome was manipulated depending on the trial type. During practice trials, we randomly presented twice each the intervals from 100ms to 1000ms in steps of 100ms, for a total of 20 practice trials.

During test trials, the process was similar except only three intervals were used (100, 400, and 700ms), and there was no feedback provided. The blending mechanism (Lebiere, 1999) was used to retrieve chunks from declarative memory. The blending mechanism computes a weighted average over chunks in memory learned during practice such that chunks with a higher likelihood of retrieval, determined by activation, carry more weight. The ACT-R activation equation,

$$A_i = B_i + S_i + P_i + \varepsilon_i \tag{5}$$

includes a: 1) base level term, B_i, for recency and frequency of use, 2) spreading term, S_i, for context effects, 3) partial matching term, P_i, for degree of match with retrieval cues, and 4) noise term, ε_i, for noise in memory. The blending mechanism uses the equation,

$$V = \underset{V}{\operatorname{argmin}} \sum_i P_i (1 - sim(V, V_i))^2 \tag{6}$$

to produce a value that minimizes the sum of all squared dissimilarities, $((1 - sim(V,V_i))^2$, of each chunk, i, between the consensus value V and the chunk value V_i, and weights it by its probability of retrieval,

$$P_i = \frac{e^{A_i/t}}{\sum_j e^{A_j/t}}. \tag{7}$$

The probability of retrieval is a function of the activation for a chunk, $e^{A_i/t}$, normalized by the activation of all retrieved chunks, $\sum_j e^{A_j/t}$.

During recall, the current tick count at the time the second stimulus appeared was used as a retrieval cue to find a match in declarative memory in the retrieve-instance-feedback production. We defined a linear similarity function for ticks in the temporal buffer which impacted how chunks were weighted during this retrieval process and how the blending average was computed. As ticks in the temporal buffer operate according to a log-scale (to reflect the scalar property of time estimations), we thought it appropriate to define ticks as having a relationship such that tick values were most similar to themselves with linearly decreasing similarity. At the end of each trial, a new chunk was created pairing the blended real-time value (i.e., the guess) and the tick count. This process was repeated over 150 trials. The accumulation of chunks during the experimental trials resulted in a gradual regression pattern in model estimates.

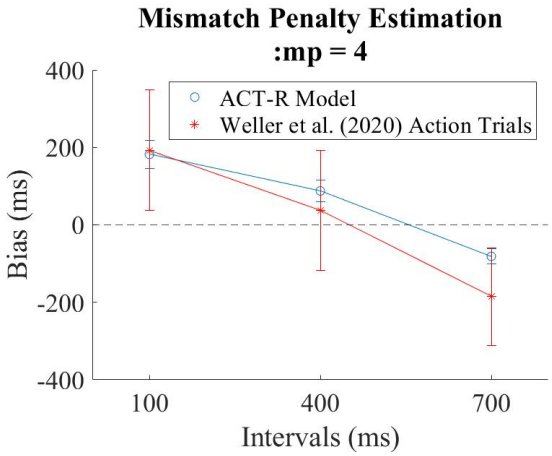

Figure 3: Best fitting model from analysis testing mismatch penalty (:mp) parameter values.

Before comparing model and human performance at the aggregate level, we completed a series of simulations to estimate the mismatch penalty (:mp) parameter as there is no default. The :mp parameter specifies the penalty, (P_i), in the activation equation and calibrates the degree of regression to the mean in the model. Lower values of :mp correspond to a wider range of chunks taken into account during retrieval. We tested values of :mp from 1 to 5 in steps of 0.5, keeping all other parameter at their default values, and we simulated 10 model runs per parameter value for a total of 90 model runs.

Results

The results from the :mp estimation are shown in Figure 3, where the best fitting model and human performance at the aggregate (across trials and individuals) are plotted together.

To determine the model with the best quantitative fit, we computed the root mean squared error (RMSE) for the difference between each model's simulated estimates compared to the human estimates. We also computed a Pearson correlation between the human and model estimates across the three interval lengths ($r = 0.99$, $p = 0.04$). The model with :mp = 4 produced the lowest $RMSE = 67.02$.

Though the adjustment of the :mp parameter improved the fit at the aggregate level, there was still a substantial difference in variance between the human data and model fit. This difference can be seen in Figure 4 which plots individual participant and model run data at the trial level using the parameter settings from the first analysis. When comparing across timing intervals (colored dots), it is clear that the variability in the model is substantially less compared to the human data. These results indicate that the model was not capturing individual level behavior.

To investigate this, we explored one parameter that represents a plausible way to account for the individual variability: blending temperature (:tmp). The :tmp parameter controls

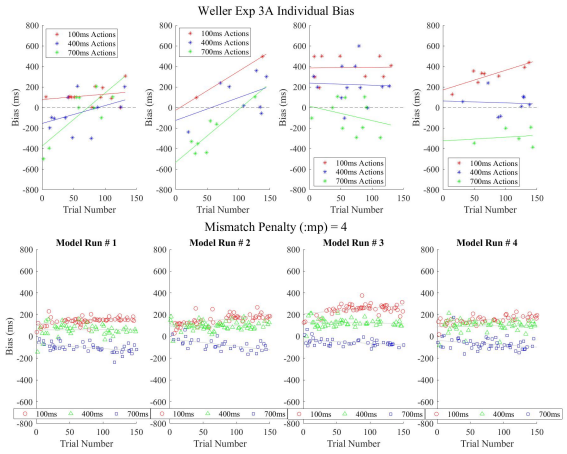

Figure 4: Individual variability and linear fits across human (top) and model simulation (bottom) data. Data points represent an estimate on a given trial across the interval lengths.

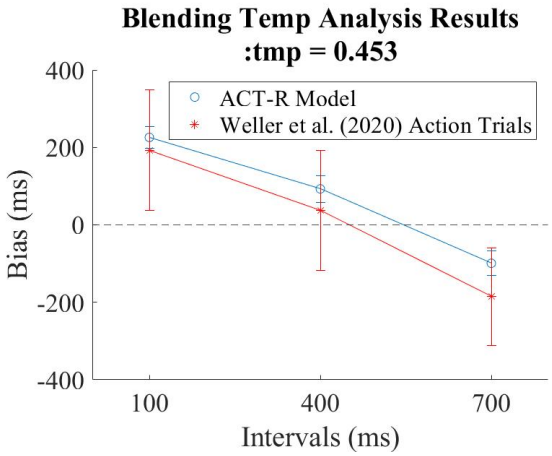

Figure 5: Best fitting model from the blending temperature (:tmp) parameter analysis.

the preference to blend chunks in memory. Higher values of :tmp correspond to more blending over chunks and, therefore, more regression to the mean of chunks. Lower values of :tmp correspond to a retrieval process closer to the best match, or "winner-take-all". We tested 30 values of :tmp sampled from a normal distribution ($\mu = 0.5$, $\sigma = 0.1$). For this analysis, we kept all other parameters at default, except for :mp $= 4$.

Figure 5 depicts the results from this analysis. We computed the RMSE to evaluate quantitative fit of the model simulation to the human data. Here, the model with the lowest average difference ($RMSE = 62.29$) used a :tmp value of 0.453. It produced only a marginal improvement from the first analysis ($RMSE = 67.02$) which did not set the :tmp parameter and only adjusted :mp to 4. We computed a correlation between the human data and model fit across the three interval lengths ($r = 0.99$, $p < 0.001$) which represented a good quantitative fit at the aggregate level. However, this value of :tmp did not improve fit to the individuals. As this was the aim of this analysis, these results indicate adjusting this parameter in future simulations may not be necessary.[1]

Discussion

Here we have developed and implemented the first cognitive model of a TB task. Using core components from the ACT-R architecture and default settings for all but one parameter (mismatch penalty), our model was able to simulate human time interval estimates in action trials from a TB dataset (Weller et al., 2020).

After defining a similarity function to specify how chunks were weighted during retrieval, our first aim was to estimate the appropriate value for the :mp parameter. Using the minimum RMSE value as the primary metric for quantitative fit, the simulations suggested only adjusting the :mp parameter to 4. This is relatively high in the range of values tested. Conceptually, this means that to simulate participant estimations in this task, the model needed a higher penalty against chunks in memory which limited the number of chunks that are averaged in memory during a single retrieval to those that were a close match to the current retrieval request.

In our second analysis, we evaluated sources of individual variability using another parameter affecting memory mechanisms, the :tmp parameter. The best fitting value provided a marginal improvement of model fit at the aggregate across three time intervals but did not improve variability comparable to human performance. Future work could allow variation in the :tmp parameter to represent individual differences in humans regarding whether they use more (i.e., blending) or fewer (i.e., single best match) previous instances in memory to inform current estimates of time intervals.

In the TBM model, parameters affecting memory mechanisms within ACT-R were the primary influence on model performance at the aggregate level. This provided some additional evidence to suggest that memory mechanisms are capable of capturing the patterns in human data from a TB task. Surprisingly, our investigation of the timing mechanism (:time-noise) did not appear to affect model performance at the individual level as much as expected (see Footnote 1 for more details). Increasing values of the :time-noise parameter increased individual level variability but worsened the fit at the aggregate level, indicating a trade-off. It may be the case that conducting an analysis similar to the one we suggest for blending temperature (i.e., altering :time-noise parameter value to fit individuals) may represent a feasible way forward

[1] We conducted an additional analysis to investigate how varying the amount of noise added between ticks, via the :time-noise parameter in the temporal buffer (b, in Equations 3 and 4), influenced variance in timing estimates. We assessed values of 0.005 to 0.1 in steps of 0.015 keeping all other parameters at default except :mp = 4. The best fitting value was :time-noise = 0.03, which is slightly higher than the default value 0.015. The minimum $RMSE = 51.11$ improved model performance at aggregate, but not substantially enough at the individual level to warrant adjusting the default setting. This represents an interesting area of investigation for future work. See https://osf.io/6bkjp/ for complete details.

in investigating sources of noise at the individual level.

It is also possible other aspects of the temporal buffer may affect model performance in this task. As mentioned in the introduction, according to Fereday et al. (2019) one might expect a slower clock rate when comparing action trials to passive ones. In the ACT-R architecture, the :time-mult parameter, which controls a multiplier constant applied to each tick, could be adjusted to formally investigate this hypothesis. Specifically, one might test a range of values higher and lower than the default value (1.1) to investigate whether this can capture any observed differences between conditions.

It is important to note that so far the majority of the analyses we describe here have been conducted at the aggregate level. However, there is evidence to suggest that the TB effect is not consistently present at the individual level, and currently this variability is not captured by this model. There are some reasons why this might be the case. For instance, one can interpret the aggregate model fit to represent one participant completing the experiment 10 times without variation. The observed variability in the human data then would potentially reflect an aggregation over different models using different parameter values. Future work should aim to investigate sources of variability across individuals and possible explanatory mechanisms.

Future work might also investigate the sources of individual variability in estimations. Currently the information in declarative memory is created in the same way across individual model runs (i.e., assuming no prior experience before beginning the task). This is due to the fact that we did not have access to practice trial data. Seeding each model run with actual practice trial data may better simulate the individual variability in the human data. In lieu of practice data, it may be useful to simulate a large number of practice trials at sub-second interval lengths to build a more realistic declarative memory store. Additionally, it may be useful to incorporate individual differences in the initial guesses the model makes during the task (e.g., Cranford et al. (2021)). Participants' initial guesses may be based on environmental priors (i.e., we expect our button press to lead to an outcome after a very short interval, typically less than 100ms in length). These expectations may vary across participants, which may lead to some of the variability present at the individual level.

As mentioned in the introduction, due to the preliminary nature of this work, our analysis did not include a simulation of the comparison trial-type that is used to determine the presence of the TB effect. In experiment 3A Weller et al. (2020), a baseline trial-type wherein participants passively observed two events, was the comparison of interest for the action trials. Baseline trials can be simulated by removing the initial voluntary action in our current task, so that the model passively observes and times the interval length between the presentation of the two visual stimuli (e.g., the letters "A" and "Z"). The estimations in this trial-type can then be compared to the action trial-type that we have developed here to determine whether the TBM model can account for the entire TB effect (i.e., more compression in the action trial-type compared to the baseline trial-type). We suggest first simulating both trial-types using the same core components and parameter values to evaluate whether additional specifications are necessary to produce the TB effect.

In conclusion, we have successfully developed and implemented the first cognitive model of a temporal binding task. An ACT-R model, using declarative memory and time perception mechanisms, provided a good qualitative fit to human data. These results, while still preliminary, add to the growing evidence that memory mechanisms can account for results from temporal binding studies. Future work should evaluate whether specifying an agency mechanism is necessary to account for the temporal binding effect.

Acknowledgments

Part of this work was completed while L.S. was an ORISE intern at AFRL, supported by the 711^{th} Human Performance Wing. C.L. was supported by the Center of Excellence for Trusted Human-Autonomy Teaming (AFRL/AFOSR award FA9550-18-1-0251). The views expressed are those of the authors and do not reflect the official guidance or position of the United States Government, the Department of Defense, or the United States Air Force. Distribution A. Approved for public release; distribution unlimited (AFRL-2023-1412).

References

Cranford, E. A., Gonzalez, C., Aggarwal, P., Tambe, M., Cooney, S., & Lebiere, C. (2021). Towards a cognitive theory of cyber deception. *Cognitive Science*, *45*(7), e13013.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.

Fereday, R., Buehner, M. J., & Rushton, S. K. (2019). The role of time perception in temporal binding: Impaired temporal resolution in causal sequences. *Cognition*, *193*, 104005.

Haggard, P. (2005). Conscious intention and motor cognition. *Trends in Cognitive Sciences*, *9*(6), 290–295.

Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, *5*(4), 382–385.

Hoerl, C., Lorimer, S., McCormack, T., Lagnado, D. A., Blakey, E., Tecwyn, E. C., & Buehner, M. J. (2020). Temporal binding, causation, and agency: Developing a new theoretical framework. *Cognitive Science*, *44*(5), e12843.

Lebiere, C. (1999). Blending: An ACT-R mechanism for aggregate retrievals. In *Proceedings of the Sixth Annual ACT-R Workshop.* George Mason University, Fairfax, VA, USA.

Legaspi, R., & Toyoizumi, T. (2019). A Bayesian psychophysics model of sense of agency. *Nature Communications*, *10*(1), 4250.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. MIT press.

Saad, L., Musolino, J., & Hemmer, P. (2022). Bayesian rational memory model simulates temporal binding effect. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 44).

Taatgen, N. A., Van Rijn, H., & Anderson, J. (2007). An integrated theory of prospective time interval estimation: the role of cognition, attention, and learning. *Psychological Review*, *114*(3), 577.

Weller, L., Schwarz, K. A., Kunde, W., & Pfister, R. (2020). Something from nothing: Agency for deliberate nonactions. *Cognition*, *196*, 104136.