Capturing Dynamic Performance in a Cognitive Model: Estimating ACT-R Memory Parameters with the Linear Ballistic Accumulator

Maarten van der Velde[1] **(m.a.van.der.velde@rug.nl), Florian Sense**[1] **(f.sense@rug.nl), Jelmer Borst**[2] **(j.p.borst@rug.nl), & Hedderik van Rijn**[1] **(d.h.van.rijn@rug.nl)**

[1]Dept. of Experimental Psychology & Behavioural and Cognitive Neuroscience, University of Groningen, the Netherlands

[2]Bernoulli Institute, Dept. of Artificial Intelligence, University of Groningen, the Netherlands

Abstract

The parameters governing our behaviour are in constant flux. Accurately capturing these dynamics in cognitive models poses a challenge to modellers. Here, we demonstrate a mapping of ACT-R's declarative memory onto the linear ballistic accumulator, a mathematical model describing a competition between evidence accumulation processes. We show that this mapping provides a method for inferring individual ACT-R parameters without requiring the modeller to build and fit an entire ACT-R model. We conduct a parameter recovery study to confirm that the LBA can recover ACT-R parameters from simulated data. Then, as a proof of concept, we use the LBA to estimate ACT-R parameters from an empirical data set. The resulting parameter estimates provide a cognitively meaningful explanation for observed differences in behaviour over time and between individuals.

Keywords: Memory; dynamic performance; individual differences; ACT-R; linear ballistic accumulator.

Introduction

Cognitive architectures such as ACT-R (Anderson, 2007) provide a framework for developing models of cognition. A challenge commonly faced by modellers is to accurately capture changes in cognitive performance over time, as well as individual differences between people, in the parameters of such models. Current approaches tend to rely on computationally expensive and statistically sub-optimal methods like parameter sweeps to identify the best-fitting parameter values. Mathematical modelling methods can serve as a more efficient and rigorous alternative (Fisher, Houpt & Gunzelmann, 2020). In this paper, we contribute to previous efforts to connect cognitive architectures and mathematical modelling by using the linear ballistic accumulator (Brown & Heathcote, 2008) to infer ACT-R parameters governing memory retrieval.

Retrieval of information from memory can be viewed as a process of evidence accumulation, in which internal and external cues contribute evidence to candidates in memory that are competing for retrieval (Ratcliff, 1978; Anderson, 2007). The first candidate to accumulate enough evidence to cross a boundary wins the race and is retrieved. The dynamics of this process are determined by the amount of evidence each candidate needs to accumulate to cross the boundary, and the rate at which this evidence accumulates.

While such evidence accumulation models have seen most use in the domain of decision making (e.g., Ratcliff, Smith, Brown & McKoon, 2016; Smith & Ratcliff, 2004; Usher & McClelland, 2001; Brown & Heathcote, 2008), there have been some applications in the domain of memory retrieval. Van Maanen et al. showed that a leaky competing accumulator model could explain performance in picture-word interference tasks (van Maanen & van Rijn, 2007; van Maanen, van Rijn & Taatgen, 2012). In this model, memory chunks accumulate activation by receiving positive and negative spreading activation from other chunks. More recently, Nicenboim and Vasishth (2018) and Fisher et al. (2020) implemented the ACT-R model of declarative memory in a lognormal race model (LNR; Rouder, Province, Morey, Gomez & Heathcote, 2015), in which the rate at which evidence for a chunk accumulates depends on its activation.

Here, we extend this formalisation of ACT-R memory retrieval as an LNR to a more flexible linear ballistic accumulator model (LBA; Brown & Heathcote, 2008). Unlike the LNR, the LBA is able to estimate the rate of accumulation separately from the distance accumulators need to travel to reach the decision boundary. This is useful, because both accumulation rate and distance to the boundary have natural counterparts in ACT-R: the accumulation rate corresponds to the activation of the chunk, while the distance can be linked to the latency factor (F) parameter. As such, the LBA provides a cognitively meaningful interpretation of ACT-R's F parameter as a measure of response caution—the larger the distance, the more evidence needs to collected before a response is made—and offers a method by which it can be estimated from response data.

In the following sections, we first describe the formal link between ACT-R and the LBA. We then demonstrate how the LBA can be used to recover ACT-R parameters in a simulation study. Finally, we fitted the LBA to an empirical data set, showing how it can offer insight in the mechanisms underlying changes in retrieval performance over time.

Casting ACT-R's Declarative Memory as a Linear Ballistic Accumulator

The linear ballistic accumulator model (Brown & Heathcote, 2008) proposes that response behaviour can be explained through a race between accumulators. Each accumulator has a certain amount of starting evidence k that increments linearly at a drift rate v until it reaches a decision boundary d. The first accumulator to reach the boundary determines the response choice and latency. A constant non-decision time t_0 is also added, representing the time required for other components of

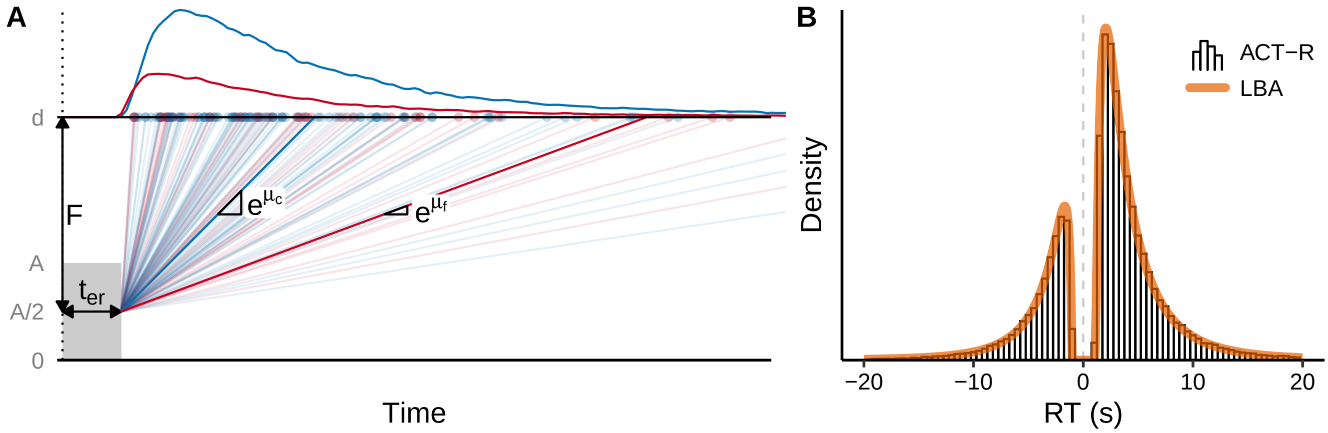

Figure 1: Casting ACT-R memory retrieval as a linear ballistic accumulator. **A**: ACT-R retrieval with two competing chunks visualised as an LBA, with marginal RT distributions shown at the top. See the main text for details. **B**: RT distributions of an ACT-R model (histogram) and the equivalent LBA model (orange curve). Error responses are shown as negative RTs.

the response process, such as perceptual and motor functions. There are two sources of variability between trials: the starting point $k \sim U(0,A)$, and drift rate $v_i \sim N(\mu_i, \sigma_i)$ for each response option i. The LBA assumes a constant rate of evidence accumulation over a trial, so the time required for an accumulator to reach its boundary on a trial j is the distance $d - k_j$ divided by the drift rate, plus non-decision time:

$$RT_j = \frac{d - k_j}{v_j} + t_0 \tag{1}$$

Across trials, the average starting point is $A/2$ and the average drift rate is μ_i, so the expected finishing time for an accumulator is:

$$\mathbb{E}(RT) = \frac{d - A/2}{\mu_i} + t_0 \tag{2}$$

We can map the LBA parameters onto ACT-R memory parameters. ACT-R models declarative memory as a set of symbolic chunks, each with a sub-symbolic activation that decays over time and is subject to noise (Anderson, 2007). The time required to retrieve a chunk depends on its activation: the more active the chunk, the faster its retrieval will be, but like the LBA, the time course of memory retrieval is deterministic once the starting values are known. If multiple chunks match a retrieval request, the chunk with the highest activation—and therefore the lowest retrieval time—at the time of the request wins. A full response also involves non-memory operations, such as stimulus encoding and response execution, which can be captured by adding a term t_{er} to the retrieval time.

ACT-R defines the full time required to retrieve a chunk i with an activation A and respond accordingly by the following equation, in which F is the latency factor, a positive scaling parameter[1]:

$$RT_i = F * e^{-A_i} + t_{er} \tag{3}$$

We can rewrite this equation in a similar form to (2):

$$RT_i = \frac{F}{e^{A_i}} + t_{er} \tag{4}$$

The mapping between ACT-R's parameters (left) and those of the LBA (right) then becomes straightforward:

$$F = d - A/2 \tag{5}$$

$$A_i = \ln(\mu_i) \tag{6}$$

$$t_{er} = t_0 \tag{7}$$

With this mapping, ACT-R's latency factor (F) is equivalent to the average distance between starting point and boundary in the accumulator model, often conceptualised as the response caution: given a constant activation, a higher value of F means that more evidence is required to complete a retrieval. The mapping relates the activation A_i of a chunk to its drift rate μ_i, meaning that a highly activated chunk can be seen as accumulating evidence more rapidly than one with a lower activation. Put differently, the drift rate μ_i is equivalent to e^{A_i}, the odds of needing the chunk. Finally, there is a direct equivalency between the non-retrieval time (t_{er} and t_0) in both models.

Figure 1A visualises the ACT-R retrieval process in the style of an accumulator model. It shows two chunks, c (blue) and f (red), competing for retrieval over multiple trials. In each trial, both accumulators race to cover the vertical distance F to the boundary. The winner gets retrieved in the time it takes to reach the boundary. There is normally distributed trial-to-trial variability, or noise, in the activation of the chunks, and therefore in the rate at which each chunk

[1] An additional parameter f may be used to scale the activation: $RT = F * e^{-f * A_i} + t_{er}$. This parameter is typically held constant at 1, and we make the same simplification here as it has no bearing on the outcomes.

accumulates evidence: $A_i \sim N(\mu_i, \sigma_i)$. As such, drift rates follow a lognormal distribution. The resulting RT distributions can be shown to be lognormal too:

$$RT_i \sim LN(\mu_i + \ln(F), \sigma_i) + t_{er} \tag{8}$$

Figure 1B demonstrates that ACT-R and the LBA generate identical response time distributions for a given set of parameters when using the mapping in equations (5)–(7). Interactive versions of these figures, in which the model parameters can be freely adjusted, are available at `https://cogmod.shinyapps.io/actr-lba/`.

Simulation: Recovering ACT-R Parameters

Given this mapping, it should be possible to identify ACT-R memory parameters from response data (RT and choice) using existing methods for fitting the LBA. Therefore, we performed a simulation study with two goals: to investigate whether the LBA can recover ACT-R memory parameters from a typical participant sample completing a reasonable number of trials, and to ensure that parameter recovery works regardless of specific parameter values. The code required to reproduce this simulation study is available at `https://osf.io/wpvj7/`.

Data

ACT-R was used to simulate 25 distinct model participants, each performing a sequence of retrieval trials. Retrieval was modelled as a competition between two chunks, c and f, representing a correct and incorrect response to a retrieval cue, respectively. For each model participant, ACT-R parameters were sampled randomly from plausible distributions, listed in Table 1. To ensure that parameters recovered by the LBA would all be on the same scale, we fixed the standard deviation of the activation of the correct response (σ_c) to 1, both in ACT-R and in the LBA[2].

We repeated the process with differently sized data sets, ranging from 25 to 50,000 trials per participant, to gauge the effect of data set size on recovery accuracy.

Model fitting

The LBA was fitted separately to each model participant's responses using the *nlminb* optimiser in R (version 3.6.3; R Core Team, 2020). We ran this optimiser 250 times with randomly generated starting values, and kept only the parameter values that yielded the highest summed log-likelihood across all runs. The *dLBA* density function from the *rtdists* package (version 0.11-2; Singmann, Brown, Gretton & Heathcote, 2020) served as the objective function. For each model participant, we derived individual ACT-R parameters from the best-fitting LBA using the mapping in equations (5)–(7).

Results

The results of the parameter recovery process are shown in Figure 2. As Figure 2A indicates, original parameter values could already be recovered with reasonable accuracy from a data set with 100 trials per participant. Some parameters (e.g., σ_f and t_{er}) appear easier to recover than others, but even the larger errors do not appear to show systematic over- or underestimation.

Figure 2B shows how recovery accuracy changed as a function of data set size. Recovery accuracy, measured as the absolute error of recovered parameter values relative to the original values, is shown separately per parameter (coloured points) as well as across parameters (black points). Unsurprisingly, recovery accuracy improved when there were more trials constraining the fit, though the current fitting method reached a plateau once there were at least 250 trials per participant.

Table 1: ACT-R parameters used in the simulation study.

	Description	Distribution
μ_c	Mean activation of correct answer	$\mu_c \sim N_-(-.5, .5)$
μ_f	Mean activation of incorrect answer	$\mu_f \sim N_-(-1.5, .5)$
σ_c	SD of activation of correct answer	$\sigma_c = 1$
σ_f	SD of activation of incorrect answer	$\sigma_f \sim N_+(1.5, .5)$
F	Latency factor	$F \sim N_+(1, .5)$
t_{er}	Non-retrieval time	$t_{er} \sim N_+(.75, .5)$

Note: N_+ and N_- are truncated normal distributions, limited to positive and negative values, respectively.

Example Application: Modelling Changing Retrieval Performance in Empirical Data

To demonstrate how the method may be used to explain dynamic memory performance in terms of cognitively meaningful constructs, we fitted the LBA to empirical data from a multi-session retrieval practice task.

Data

We use data from a retrieval practice task completed by recruits of the Commando Corps, Royal Netherlands Army (*Korps Commandotroepen*), in which participants learned the names of ficticious safehouses on a map. On first presentation, a safehouse was shown with its name, while subsequent repetitions required participants to select the correct name themselves from a set of four answer options. Participants completed three 8-minute sessions over the course of a week. They studied a different map in every session, and maps were counterbalanced between participants. The task was presented within an adaptive learning system that schedules each item to be repeated whenever its activation is expected to hit a threshold (van Rijn, van Maanen & van Woudenberg, 2009; van der Velde, Sense, Borst & van Rijn, 2021). As such, we could expect the activation of the chunks being retrieved to be fairly stable across trials, despite the novelty of the materials. Response accuracy and response time were recorded in every trial.

Session 1 was scheduled on the first day of the week, while the second and third sessions took place several days later and were scheduled immediately before and after a high-intensity loaded speed march of about 40 minutes. We

[2]See Brown and Heathcote (2008) for alternative solutions to the scaling problem in accumulator models.

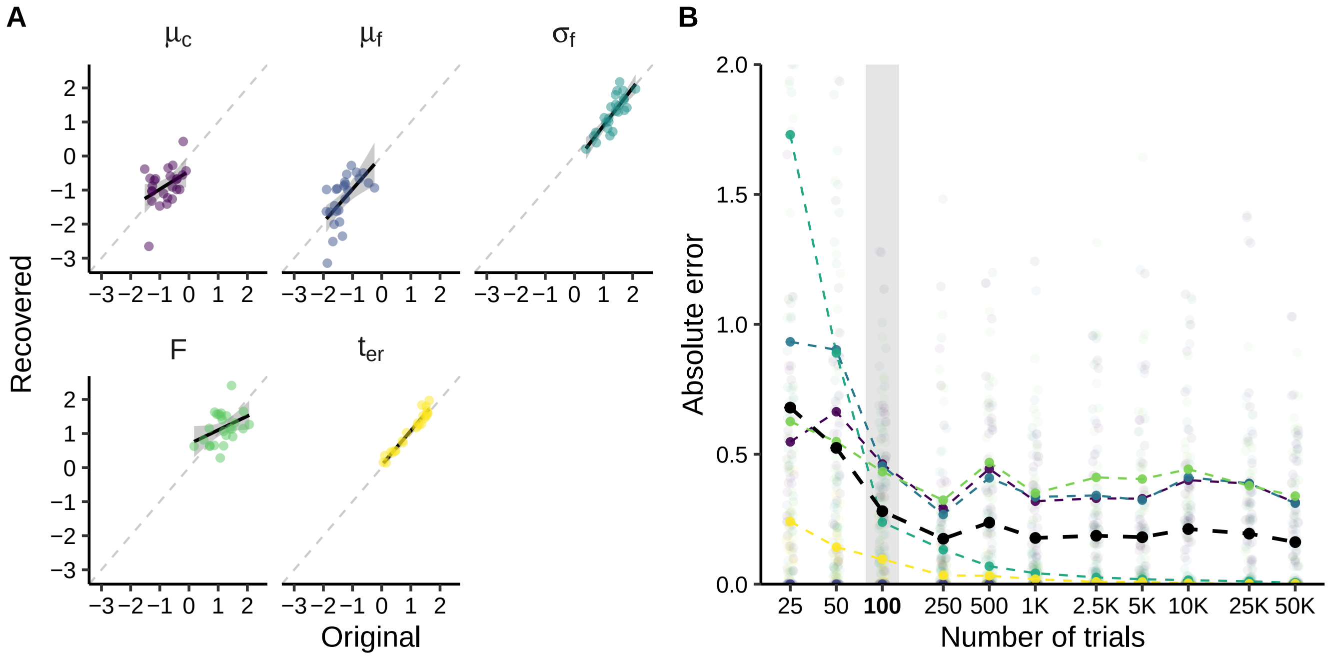

Figure 2: Recovery of ACT-R parameters using the LBA. **A**: Original versus recovered parameter values for a data set with 100 trials per participant. Parameter descriptions are given in Table 1. **B**: Recovery accuracy (absolute error) for data sets with different numbers of trials per participant. Light coloured points show individual errors, dark coloured points show the mean error per parameter, and black points show the mean error across parameters.

expected performance to change for two reasons: increased familiarity with the task might lead to better performance after the first session, and the physical exertion of the speed march might affect performance in the third session.

For the analysis, we removed the first trial for each item (in which the answer was shown on screen), trials in which participants did not respond within 30 s, and trials in which the recorded response time was lower than 300 ms. Since the simulation study showed that recovery was worse in small data sets, participants had to have completed at least 50 practice trials per session to be included. In addition, we required that participants made at least 5 error responses per session, to give the model a chance of fitting the error response distribution. These criteria struck a balance between ensuring a sufficient number of observations per participant and including as many participants as possible. They yielded a data set with 12,568 usable observations (out of 29,441) from 50 (out of 127) participants.

Model fitting

We fitted the LBA separately to each of the three retrieval practice sessions for each participant. The fitting procedure was the same as in the simulation study. The analysis code is available at `https://osf.io/wpvj7/`.

Results

Figure 3 shows participants' performance on the task over the three sessions. Despite the task difficulty being the same in all three sessions, performance improved in two ways. Firstly, response accuracy increased and then plateaued: a logistic mixed-effects model with a main effect of session and random intercepts for participants showed that accuracy increased from the first to the second session ($z = -4.680$, $p < .001$), but found no evidence for a change from the second to the third session ($z = -0.253$, $p = .8$). Secondly, responses became faster: a generalised mixed-effects model with a Gamma link function and with a main effect of session and random intercepts for participants found a decrease in response times on correct trials from session 1 to session 2 ($t = 2.250$, $p = .0244$), and from session 2 to session 3 ($t = -7.182$, $p < .001$).

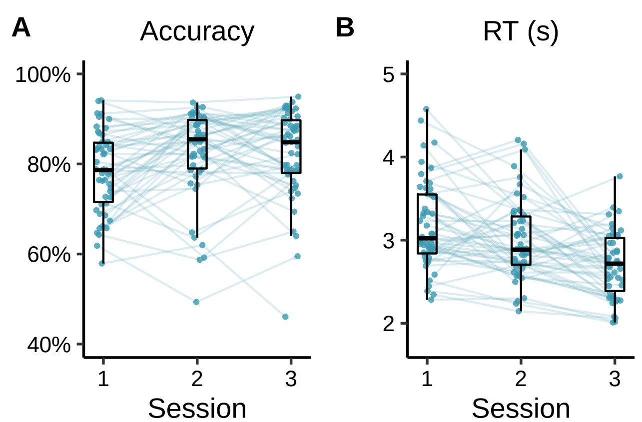

Figure 3: Performance on the retrieval practice task by participant. **A**: percentage correct responses per session. **B**: median response time on correct responses per session.

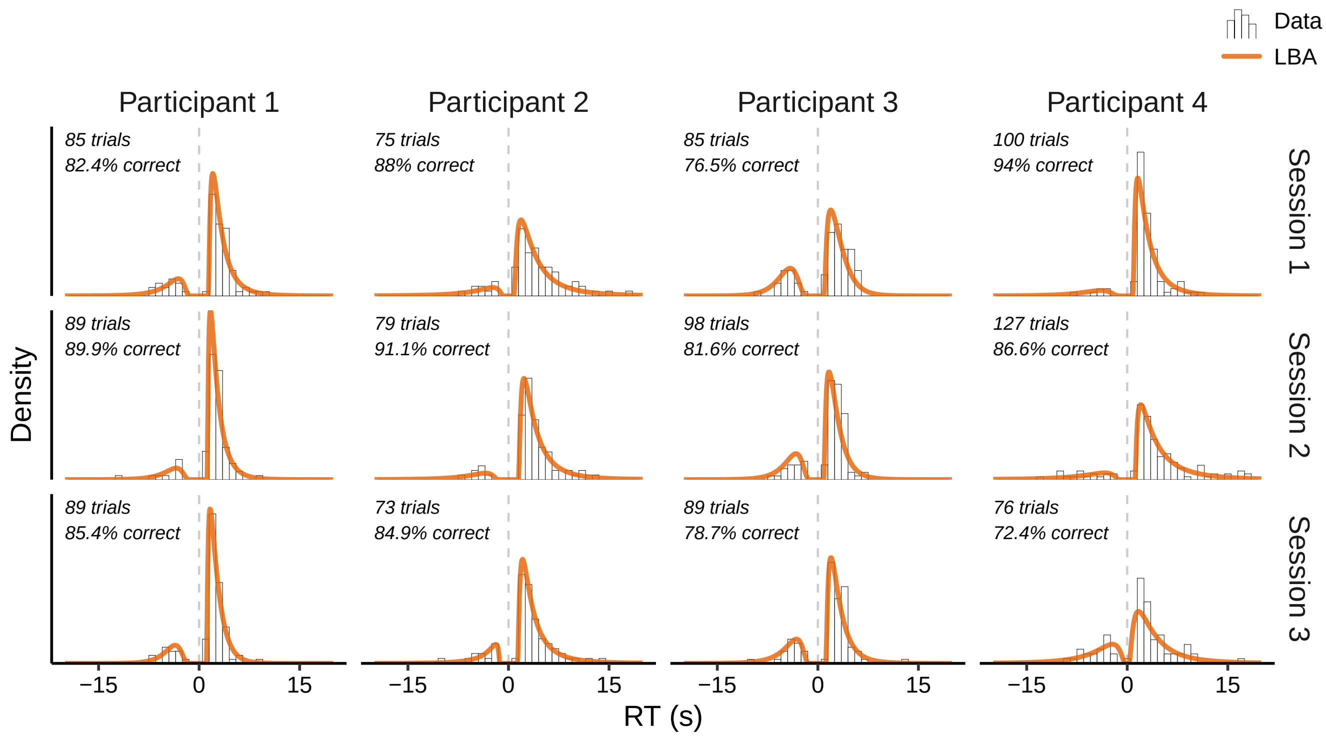

Figure 4: Best fits of the LBA to the response data of four participants over three retrieval practice sessions. Error responses are shown as negative RTs. The number of available trials and the response accuracy are shown in the top left corner of each plot.

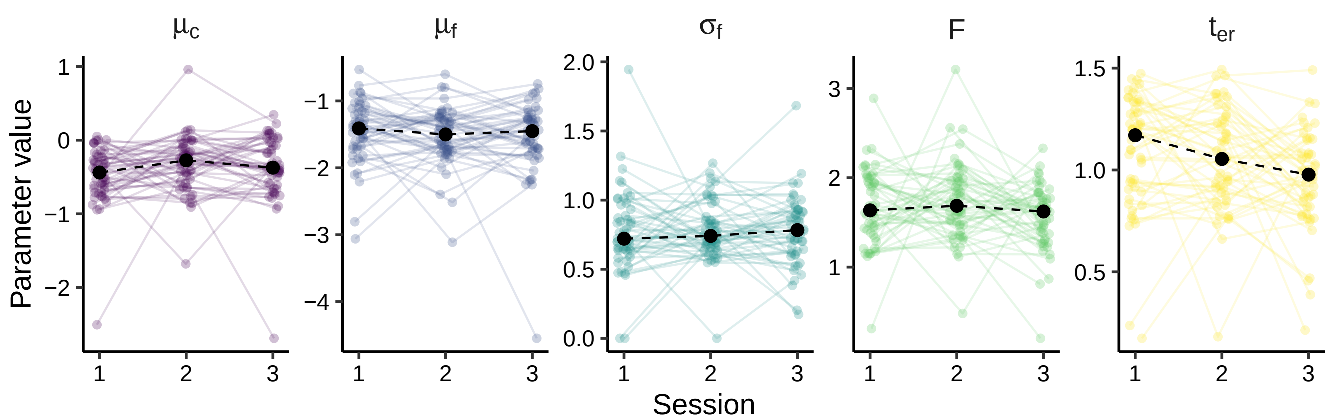

Figure 5: ACT-R memory parameters inferred from the data. Coloured points show individual estimates; large black points indicate the median value across participants. Parameter descriptions are given in Table 1. Note: Y-axes differ between plots.

Figure 4 shows the best fit of the LBA to the response time distributions of four randomly selected participants. These examples suggest that the model captured the shape of the data quite well, although the low number of trials and high response accuracy did make it challenging to fit the error responses.

The inferred ACT-R parameters are shown in Figure 5. There is substantial variation in the parameter values for individual participants, but they are nonetheless clustered quite neatly around the sample averages. As one would expect, the activation of the correct answer (μ_c) tended to exceed the activation of the incorrect answer (μ_f), reflecting participants' better-than-chance performance. To explore possible changes in parameter values over time, we fitted separate linear mixed-effects models to each parameter, testing whether there was a session effect on the parameter value, with random intercepts for participants. These models suggested that the parameters generally remained fairly constant between sessions[3].

[3] Aside from the reported effects, there was some evidence for a decrease in the F parameter from session 2 to session 3, though the corresponding model failed to converge. More generally, these results should be interpreted with a degree of caution, as repeated LBA fits yield slightly different parameter estimates due to random variation.

However, the activation of the correct answer (μ_c) did appear to increase from session 1 to session 2 ($t(98) = -2.050$, $p = .043$). Furthermore, since the outcome of the retrieval depends on which of the two candidate chunks has the highest activation, rather than on the individual activation of either chunk, we also fitted a linear mixed-effects model with the difference in activation $\mu_c - \mu_f$ as the dependent variable. This model similarly suggested that the activation difference was higher in session 2 than in session 1 ($t(98) = -3.133$, $p = .00228$), indicating that, on average, participants' chances of retrieving the correct answer improved. Finally, the non-retrieval time t_{er} showed a significant decrease from session 2 to session 3 in particular ($t(98) = -2.351,\ p = .0207$), reflecting a speed-up in perceptual and/or motor functions.

In conclusion, exploratory analysis of the inferred ACT-R parameter estimates suggests that the observed increase in accuracy and response speed from session 1 to session 2 could be the result of a higher mean activation of the correct answer and a greater difference in activation between the correct and incorrect answer, while the drop in response times from session 2 to session 3 may be attributable to a decrease in non-retrieval time t_{er}.

Discussion

We have demonstrated a mapping of the parameters of the linear ballistic accumulator onto parameters governing declarative memory retrieval in ACT-R. By fitting the LBA to retrieval data and mapping the inferred LBA parameters onto ACT-R memory parameters, we can arrive at a mechanistic explanation for observed performance changes, without needing to build and fit an ACT-R model directly. The resulting ACT-R parameters—activation of chunks, duration of non-retrieval processes, and latency factor—have cognitively meaningful interpretations within the wider context of the architecture, enhancing the interpretation that could be given by the LBA alone. The mapping extends upon an earlier mapping between the lognormal race model and ACT-R (Nicenboim & Vasishth, 2018), by adding the ability to fit the latency factor. From a theoretical standpoint, ACT-R benefits from this connection to the LBA too: the latency factor is given a more concrete meaning, namely as a measure of response caution.

The method described here allows one to disentangle several factors contributing to memory retrieval performance. In many settings, inside and outside the laboratory, the parameters governing our behaviour are inevitably in flux: we learn and forget, we become tired or impatient, our goals and desires change, we let our minds wander. There is clear explanatory power in being able to capture such changes within a mathematical model. Linking the terms of that mathematical model to constructs defined in a cognitive architecture can further aid the interpretation of observed behaviour.

An important limitation of this method is that it assumes that the distribution of drift rates—and therefore the activation of memory chunks—remains constant across a block of trials. This assumption is most likely to be met when information is so ingrained that there is no appreciable decay in its activation (e.g., sentence processing; Nicenboim & Vasishth, 2018), or when retrieval attempts are timed such that they occur whenever a particular activation is reached (e.g., adaptive scheduling, as used in our empirical example).

We used a relatively simple procedure for fitting the LBA. Extending this approach to a hierarchical Bayesian LBA may be beneficial (e.g., Nicenboim, 2018). It would enable modelling multiple participants and sessions simultaneously, improving the ability to estimate and compare participant-level and group-level effects, while also capturing the uncertainty in those estimates (Fisher et al., 2020). This could be particularly valuable in smaller data sets, where our current approach still struggles.

In summary, we have demonstrated how a mapping between the linear ballistic accumulator and the ACT-R cognitive architecture can aid in capturing dynamic performance in a cognitive model, thereby contributing to growing efforts to integrate formal modelling approaches.

Acknowledgements

We thank Leendert van Maanen for his helpful guidance on accumulator modelling, and Maurits Baatenburg de Jong and Ruud den Hartigh for their help in collecting the empirical data.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford, UK: Oxford University Press.

Brown, S. D. & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178. doi: 10.1016/j.cogpsych.2007.12.002

Fisher, C. R., Houpt, J. W. & Gunzelmann, G. (2020). Developing memory-based models of ACT-R within a statistical framework. *Journal of Mathematical Psychology*, *98*, 102416. doi: 10.1016/j.jmp.2020.102416

Nicenboim, B. (2018). The implementation of a model of choice: The (truncated) linear ballistic accumulator. In *StanCon.* Helsinki, Finland. doi: 10.5281/zenodo.1465990

Nicenboim, B. & Vasishth, S. (2018). Models of retrieval in sentence comprehension: A computational evaluation using Bayesian hierarchical modeling. *Journal of Memory and Language*, 34. doi: 10.1016/j.jml.2017.08.004

R Core Team. (2020). *R: A language and environment for statistical computing.* R Foundation for Statistical Computing. Vienna, Austria.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. doi: 10.1037/0033-295X.85.2.59

Ratcliff, R., Smith, P. L., Brown, S. D. & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, *20*(4), 260–281. doi: 10.1016/j.tics.2016.01.007

Rouder, J. N., Province, J. M., Morey, R. D., Gomez, P. & Heathcote, A. (2015). The lognormal race: A cognitive-process model of choice and latency with desirable psychometric properties. *Psychometrika*, *80*(2), 491–513. doi: 10.1007/s11336-013-9396-3

Singmann, H., Brown, S., Gretton, M. & Heathcote, A. (2020). *Rtdists: Response time distributions.*

Smith, P. L. & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*(3), 161–168. doi: 10.1016/j.tins.2004.01.006

Usher, M. & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592. doi: 10.1037/0033-295X.108.3.550

van der Velde, M., Sense, F., Borst, J. & van Rijn, H. (2021). Alleviating the cold start problem in adaptive learning using data-driven difficulty estimates. *Computational Brain & Behavior*. doi: 10.1007/s42113-021-00101-6

van Maanen, L. & van Rijn, H. (2007). An accumulator model of semantic interference. *Cognitive Systems Research*, *8*(3), 174–181. doi: 10.1016/j.cogsys.2007.05.002

van Maanen, L., van Rijn, H. & Taatgen, N. (2012). RACE/A: An architectural account of the interactions between learning, task control, and retrieval dynamics. *Cognitive Science*, *36*(1), 62–101. doi: 10.1111/j.1551-6709.2011.01213.x

van Rijn, H., van Maanen, L. & van Woudenberg, M. (2009). Passing the test: Improving learning gains by balancing spacing and testing effects. In *Proceedings of the 9th International Conference on Cognitive Modeling* (pp. 110–115). Manchester, UK.