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

Maarten van der Velde (m.a.van.der.velde@rug.nl)¹, Florian Sense (f.sense@rug.nl)¹, Jelmer P. Borst (j.p.borst@rug.nl)², Leendert van Maanen (l.vanmaanen@uu.nl)³, and Hedderik van Rijn (d.h.van.rijn@rug.nl)¹

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

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

³Dept. of Experimental Psychology, Utrecht University, the Netherlands

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 (LBA), 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. Existing parameter estimation methods for the LBA can be used, instead of the computationally expensive parameter sweeps that are traditionally done. 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. In addition, we find that the mapping between ACT-R and LBA lends a more concrete interpretation to ACT-R's latency factor parameter, namely as a measure of response caution. This work contributes to a growing movement towards integrating formal modelling approaches in cognitive science.

Keywords: Memory; dynamic performance; individual differences; cognitive modelling; 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, Hout, & 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.

Traditionally, research in cognitive modelling has taken many different forms, often with relatively little overlap in methodology and terminology, making it challenging to integrate and compare approaches (Gluck, Bello & Busemeyer, 2008). Two of the dominant paradigms are the cognitive architecture approach and the mathematical modelling approach. On one hand, there is a rich history of model development in the context of cognitive architectures like ACT-R (Anderson, 2007), which provide a general-purpose framework suitable for constructing detailed, integrative models of complex behaviour, ranging from multi-tasking and skill acquisition to reasoning and mind wandering (Gluck & Pew, 2005; Salvucci & Taatgen, 2008; van Vugt, van der Velde & ESM-MERGE Investigators, 2018). On the other hand, mathematical models of

Keywords:

Memory; dynamic performance; individual differences; cognitive modelling; ACT-R; linear ballistic accumulator.

Corresponding author:

Maarten van der Velde (m.a.van.der.velde@rug.nl)
Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands
<https://orcid.org/0000-0003-4849-2676>

cognition—such as the family of evidence accumulation models, which includes the linear ballistic accumulator (Brown & Heathcote, 2008) and the drift diffusion model (Ratcliff, 1978)—have proven invaluable in modelling performance on a wide array of cognitive tasks in a statistically rigorous and scalable way (Heathcote et al., 2019). While work in both paradigms has mostly happened in parallel, there have been attempts at integrating the two (e.g., van Maanen, van Rijn & Taatgen, 2012; Nicenboim & Vasishth, 2018; Fisher et al., 2020). As Fisher et al. (2020) point out, such integration of approaches can lead to new insights, and can help promote a common language across modelling traditions. Here, we explore how a convergence of the ACT-R cognitive architecture and the linear ballistic accumulator helps in modelling dynamic performance in the domain of memory retrieval.

Memory retrieval as evidence accumulation

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 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 et al., 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. As we will discuss more extensively later, 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 an

inverse measure of response caution—the larger the distance, the more evidence needs to be 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's declarative memory and the LBA. We then demonstrate how the LBA can be used to recover ACT-R parameters in a simulation study. Finally, we fit 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 the response process, such as perceptual and motor functions. There are two sources of variability between trials: the starting point k , which is typically drawn from a uniform distribution $[0, A]$, and the drift rate v for each response option i , which is often drawn from a normal distribution $N(v_i, s_i)$. The LBA assumes a constant rate of evidence accumulation over a trial, so the time required for an accumulator i to reach its boundary on a trial j is the distance $d - k_j$ divided by the drift rate, plus non-decision time:

$$RT_{ij} = \frac{d - k_j}{v_{ij}} + t_0 \quad (1)$$

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

$$\mathbb{E}(RT_i) = \frac{d - A/2}{v_i} + t_0 \quad (2)$$

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. Just as in 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 at the time of the request—and therefore the lowest retrieval time—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. This term fulfils the same role as the t_0 parameter in the LBA.

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.

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 parameter¹:

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

By default, ACT-R represents F as a constant. However, to match the structure of the LBA, we assume here that F is variable between trials. Similar to the starting point k in the LBA, F is drawn from a uniform distribution $[a_F, b_F]$ in every trial. In addition, following Lebière, Anderson and Reder (1994), we assume normally distributed trial-to-trial variability, or noise, in the activation of chunks: $A_i \sim N(\mu_i, \sigma_i)$. The time required to retrieve chunk i on trial j and respond accordingly therefore becomes:

$$RT_{ij} = F_j * e^{-A_{ij}} + t_{er} \quad (4)$$

Across trials, the average latency factor is $\bar{F} = \frac{a_F + b_F}{2}$ and the average activation is μ_i , so the expected finishing time is:

$$\mathbb{E}(RT_i) = \bar{F} * e^{-\mu_i} + t_{er} \quad (5)$$

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

$$\mathbb{E}(RT_i) = \frac{\bar{F}}{e^{\mu_i}} + t_{er} \quad (6)$$

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

$$\bar{F} = d - A/2 \quad (7)$$

$$\mu_i = \ln(v_i) \quad (8)$$

$$t_{er} = t_0 \quad (9)$$

With this mapping, ACT-R's mean latency factor ($\bar{F}$) corresponds to the average distance between starting point and boundary in the accumulator model, often conceptualised as the response caution. This means that, given a constant activation, a higher value of $\bar{F}$ means that more evidence is required to complete a retrieval. Fluctuations in this parameter between trials can be thought of as representing moment-to-moment changes in response caution (Boehm, van Maanen, Forstmann & van Rijn, 2014). By default, ACT-R assumes a fixed F parameter, which means that all variability in retrieval time over the duration of a task is ascribed to activation and activation noise. This assumption can lead to incorrect conclusions in scenarios where we cannot assume consistent response caution between trials, for example due to post-error slowing (Dutilh et al., 2012) or shifts in the speed-accuracy trade-off (van Maanen et al., 2011). By assuming a variable, rather than fixed, F parameter, we can distinguish such fluctuations in response caution from general activation-related variability.

The mapping relates the mean activation μ_i of a chunk to the logarithm of the mean drift rate v_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 v_i is equivalent to e^{μ_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)

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

and f (red), competing for retrieval over multiple trials. In each trial, both accumulators race to cover the vertical distance to the boundary. The winner gets retrieved in the time it takes to reach the boundary. The vertical distance is shared among the accumulators but varies between trials, following a uniform distribution. Furthermore, there is normally distributed trial-to-trial variability in the activation of the chunks, and therefore in the rate at which each chunk accumulates evidence. Given that activation has a logarithmic relation to drift rate in the proposed mapping, drift rates now follow a lognormal distribution (see Terry et al., 2015, for more detail on non-normal drift rate distributions). 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 (7)–(9). 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, respectively². In each trial, the response time was the

time (in seconds) required for the winning accumulator to reach the decision boundary. 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 as the original ACT-R parameters, we fixed the standard deviation of the activation of the correct response (σ_c) to 1, both in ACT-R and in the LBA³.

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 *ridists* 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 (7)–(9).

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.

Figures 2B and 2C show how recovery accuracy changed as a function of data set size, both in terms of the absolute error of recovered parameter values relative to the original values (Figure 2B) and in terms of the correlation between original and recovered values (Figure 2C). 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.

Given a sufficient number of observations, σ_f and t_{er} could be recovered with very high accuracy. The mean latency factor $\bar{F}$ and the activation means μ_c and μ_f remained more difficult to recover, suggesting a degree of trade-off between parameters: variability in one parameter could to some extent also be

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	$\sigma_f \sim N_{+}(1.5, .5)$
$\bar{F}$	Mean latency factor	$\bar{F} \sim N_{+}(1, .5)$
R_F	Range of latency factor ($b_F - a_F$)	$R_F \sim N_{+}(.1, .05)$
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.

²The model presented here does not explicitly account for possible retrieval failures; these are currently lumped together with incorrect retrievals in a single accumulator. Provided that the data being modelled allow for a distinction between these different types of error, an additional accumulator could be added, representing ACT-R's retrieval threshold τ (Lebière et al., 1994).

³See Brown and Heathcote (2008) for alternative solutions to the scaling problem in accumulator models. Note that the choice of scaling constraint affects the interpretation of the model parameters (van Maanen & Miletic, 2021).

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. The grey bar marks the data set shown in sub-figure A. **C**: Correlation (Pearson's r) between original and recovered parameter values for data sets with different numbers of trials per participant. Coloured points show the correlation per parameter, and black points show the mean correlation across parameters. The grey bar marks the data set shown in sub-figure A. **D**: Correlations (Pearson's r) between pairs of recovered parameters in the data set with 50,000 trials per participant.

captured by adjustments to another parameter (see also, e.g., Miletić, Turner, Forstmann & van Maanen, 2017). Figure 2D confirms that there were indeed correlations between some of the recovered parameters in the largest data set we tested, which contained 50,000 trials per participant: the mean activation estimates of the two chunks were positively correlated, and the mean activation of the incorrect chunk correlated

negatively with the activation noise of that chunk.

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 mean-

ingful constructs, we fitted the LBA to empirical data from a multi-session retrieval practice task.

Data

We used 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 fictitious safe houses on a map (van der Velde et al., 2020; van der Velde, Sense, Borst, den Hartigh et al., 2021). The data were collected as part of a larger project; see Huijzer et al. (2022) for a more detailed description of the participants and procedure. The task consisted of a sequence of retrieval practice trials. On first presentation, a safe house 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 (Sense, Behrens, Meijer & van Rijn, 2016; 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 material. 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 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. As such, the data set provides an excellent test case for our method.

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 session⁴ 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 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.

⁴Session 2 was used as the reference level (intercept). Coefficients and corresponding test statistics for Sessions 1 and 3 were estimated relative to this intercept.

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.

The inferred ACT-R parameters for all participants 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. The estimates of the mean latency factor ($\bar{F}$) are similar to values identified by others (e.g., see Table 4 in Anderson, Bothell, Lebiere & Matessa, 1998). We found a non-retrieval time (t_{er}) of about 1 second, which we consider a plausible estimate of the time needed to encode a complex visual stimulus and execute a response.

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 sessions⁵. 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. Others have reported a similar reduction in this parameter as a result of practice (e.g., Dutilh, Vandekerckhove, Tuerlinckx & Wagenmakers, 2009).

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 is the result of a higher mean activation of the correct answer and a greater difference in activation between the correct

⁵Aside from the reported effects, there was some evidence for a decrease in the $\bar{F}$ 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 in starting values.

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.

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 (variable) measure of response caution. The mapping requires that this latency factor is variable between retrieval attempts instead of fixed, as it would be by default in ACT-R. As Fisher et al. (2020) suggest, adapting the assumptions made by ACT-R can be appropriate if doing so facilitates the integration of modelling approaches. Furthermore, we argue that this change makes for a better model of retrieval, as it enables ACT-R to separate between-trial variability in memory activation from variability in response caution. This makes it possible to capture trial-to-trial shifts in participants' speed-accuracy trade-off, which is a key advantage of evidence accumulation models that has

been difficult to replicate in ACT-R (Brown & Heathcote, 2008). Previously, Schneider and Anderson (2012) modelled the speed-accuracy trade-off in ACT-R by assuming an additional deadline-induced guessing process, but such a mixture of processes appears to lack empirical support (van Maanen, 2016). Modelling speed-accuracy trade-off through a variable F parameter may therefore be a good alternative.

An important limitation of our 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., in sentence processing; Nicenboim & Vasishth, 2018), or when retrieval attempts are timed such that they occur whenever a particular activation is reached (e.g., with adaptive scheduling, as used in our empirical example).

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.

The mapping between ACT-R's declarative memory and the linear ballistic accumulator represents a specific instance in which two different modelling paradigms converge. Such convergences provide an opportunity for modellers from different traditions to express their ideas in a common language, and to benefit from the possibilities offered by a different paradigm to their own. For instance, in contrast to ACT-R, mathematical models like the LBA are well-suited

for Bayesian and hierarchical methods of parameter estimation (Heathcote et al., 2019). We currently used a relatively simple procedure for fitting the LBA, but 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 data sets with fewer observations per participant, where our current approach still struggles.

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

Acknowledgements

We thank Maurits Baatenburg de Jong, Ruud den Hartigh and Rik Huijzer for their help in collecting the Commando Corps data, and we thank Bruno Nicenboim for his helpful comments on an earlier version of this manuscript.

References

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford, UK: Oxford University Press.
- Anderson, J. R., Bothell, D., Lebiere, C. & Matessa, M. (1998). An Integrated Theory of List Memory. *Journal of Memory and Language*, 38, 341–380. doi: 10.1006/jmla.1997.2553
- Boehm, U., van Maanen, L., Forstmann, B. & van Rijn, H. (2014). Trial-by-trial fluctuations in CNV amplitude reflect anticipatory adjustment of response caution. *NeuroImage*, 96, 95–105. doi: 10.1016/j.neuroimage.2014.03.063
- 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
- Dutilh, G., Vandekerckhove, J., Forstmann, B. U., Keuleers, E., Brysbaert, M. & Wagenmakers, E.-J. (2012). Testing theories of post-error slowing. *Attention, Perception, & Psychophysics*, 74(2), 454–465. doi: 10.3758/s13414-011-0243-2
- Dutilh, G., Vandekerckhove, J., Tuerlinckx, F. & Wagenmakers, E.-J. (2009). A diffusion model decomposition of the practice effect. *Psychonomic Bulletin & Review*, 16(6), 1026–1036. doi: 10.3758/16.6.1026
- 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
- Gluck, K. A., Bello, P. & Busemeyer, J. (2008). Introduction to the Special Issue. *Cognitive Science*, 32(8), 1245–1247. doi: 10.1080/03640210802473582
- Gluck, K. A. & Pew, R. W. (Eds.). (2005). *Modeling human behavior with integrated cognitive architectures: Comparison, evaluation, and validation*. Mahwah, N.J: Lawrence Erlbaum.
- Heathcote, A., Lin, Y.-S., Reynolds, A., Strickland, L., Gretton, M. & Matzke, D. (2019). Dynamic models of choice. *Behavior Research Methods*, 51(2), 961–985. doi: 10.3758/s13428-018-1067-y
- Huijzer, R., Jeronimus, B. F., Reehoorn, A., Blaauw, F., Baatenburg de Jong, M., de Jonge, P. & den Hartigh, R. (2022). Personality traits of special forces operators: Comparing commandos, candidates and controls. *Sport, Exercise and Performance Psychology*.
- Lebiere, C., Anderson, J. R. & Reder, L. M. (1994). Error Modeling in the ACT-R Production System. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (pp. 555–559). Routledge. doi: 10.4324/9781315789354-96
- Miletić, S., Turner, B. M., Forstmann, B. U. & van Maanen, L. (2017). Parameter recovery for the Leaky Competing Accumulator model. *Journal of Mathematical Psychology*, 76, 25–50. doi: 10.1016/j.jmp.2016.12.001
- 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
- Salvucci, D. D. & Taatgen, N. A. (2008). Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review*, 115(1), 101–130. doi: 10.1037/0033-295X.115.1.101
- Schneider, D. W. & Anderson, J. R. (2012). Modeling fan effects on the time course of associative recognition.

- Cognitive Psychology*, 64(3), 127–160. doi: 10.1016/j.cogpsych.2011.11.001
- Sense, F., Behrens, F., Meijer, R. R. & van Rijn, H. (2016). An Individual's Rate of Forgetting Is Stable Over Time but Differs Across Materials. *Topics in Cognitive Science*, 8(1), 305–321. doi: 10.1111/tops.12183
- 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
- Terry, A., Marley, A., Barnwal, A., Wagenmakers, E.-J., Heathcote, A. & Brown, S. D. (2015). Generalising the drift rate distribution for linear ballistic accumulators. *Journal of Mathematical Psychology*, 68–69, 49–58. doi: 10.1016/j.jmp.2015.09.002
- 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., den Hartigh, R., Baatenburg de Jong, M. & van Rijn, H. (2020). Cognition at Special Forces Boot Camp: Does High-Intensity Physical Exercise Affect Memorisation? In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 42, p. 1857). Toronto, Canada.
- van der Velde, M., Sense, F., Borst, J., den Hartigh, R., Baatenburg de Jong, M. & van Rijn, H. (2021). Memory Performance in Special Forces: Speedier Responses Explain Improved Retrieval Performance after Physical Exertion. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 43, p. 3471). Vienna, Austria.
- 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*, 4, 231–249. doi: 10.1007/s42113-021-00101-6
- van Maanen, L. (2016). Is There Evidence for a Mixture of Processes in Speed-Accuracy Trade-Off Behavior? *Topics in Cognitive Science*, 8(1), 279–290. doi: 10.1111/tops.12182
- van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E.-J., Ho, T., Serences, J. & Forstmann, B. U. (2011). Neural Correlates of Trial-to-Trial Fluctuations in Response Cautiousness. *Journal of Neuroscience*, 31(48), 17488–17495. doi: 10.1523/JNEUROSCI.2924-11.2011
- van Maanen, L. & Miletic, S. (2021). The interpretation of behavior-model correlations in unidentified cognitive models. *Psychonomic Bulletin & Review*, 28(2), 374–383. doi: 10.3758/s13423-020-01783-y
- 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 Vugt, M. K., van der Velde, M. & ESM-MERGE Investigators. (2018). How Does Rumination Impact Cognition? A First Mechanistic Model. *Topics in Cognitive Science*, 10(1), 175–191. doi: 10.1111/tops.12318