A Model of Motivation and Effort Allocation in the ACT-R Cognitive Architecture

Yuxue C. Yang (chery@uw.edu)
Department of Psychology, University of Washington, Seattle, WA 98105, USA

Andrea Stocco (stocco@uw.edu)
Department of Psychology, University of Washington, Seattle, WA 98105, USA

Abstract

Motivation is the driving force that influences people's behaviors and interacts with many cognitive functions. Computationally, motivation is represented as a cost-benefit analysis that weighs efforts and rewards in order to choose the optimal actions. Shenhav and colleagues (2013) proposed an elegant theory, the Expected Value of Control, which describes the relationship between cognitive efforts, costs, and rewards. In this paper, we propose a more fine-grained and detailed motivation framework that incorporates the principles of EVC into the ACT-R cognitive architecture. Specifically, motivation is represented as a specific slot in *Goal* buffer with a corresponding scalar value, *M*, that is translated into the reward value R_t that is delivered when the goal is reached. This implementation is tested in two models. The first model is a high-level model that reproduces the EVC predictions with abstract actions. The second model is an augmented version of an existing ACT-R model of the Simon task, in which the motivation mechanism is shown to permit optimal effort allocation and reproduce known phenomena. Finally, the broader implications of our mechanism are discussed.

Keywords: Motivation, Cognitive Control, Effort, Computational Modeling, Cognitive Architecture

Introduction

Observable behavior in cognitive tasks is affected by the degree to which a participant puts effort into the task. The driving force behind this effort allocation is usually called motivation and represents a significant obstacle in properly inferring individual characteristics from observations. For example, a participant performing poorly in an N-back task might be poorly motivated to perform the task, rather than having limited working memory capacity. Despite its importance, motivation is rarely modeled or accounted for in cognitive models. In this paper, we outline a theory of motivation implemented in the ACT-R cognitive architecture and demonstrate its application.

To understand motivation from a cognitive modeling perspective, it is necessary to clarify the definition and relationships between several important constructs. *Motivation* is not directly observable. It is usually described as a driving force or invigorating impact on behavior or cognition that initiates a goal-oriented behavior. That is to say, we can only infer one's motivation from his behavior and cognition. *Effort* refers to how many cognitive resources one would allocate to a particular activity in order to achieve the goal. According to Inzlicht, Shenhav, and Olivola (2018), *Motivation* specifies both direction and intensity of goal-oriented behavior, while effort only indicates the intensity of any possible action, without reference to any goal. *Demand* is different from *Effort* in that it is the descriptive property of the task or environment, while *Effort* indicates the magnitude of the force that an individual might apply toward the environment. Other cognitive states such as mental fatigue, curiosity, and high arousal may interact with motivation in certain ways to have crucial impacts on learning, memory, and other cognitive control functions therefore, cognitive modeling gives us a unique opportunity to parse apart the specific effect of motivation alone.

Expected Value of Control Theory

Although several attempts have been made to capture motivation within a computational framework (e.g., Niv, 2007), the current dominant theory is the Expectancy Value Theory. It was first proposed by Voom in the 1960s and recently expanded into a formal theory known as the Expected Value of Control (EVC) model by Shehav et al. (2017). The EVC model assumes that individuals would evaluate cost-benefit tradeoffs in order to maximize gains and minimize costs in deciding how much cognitive effort one would allocate to the chosen action, as shown in Fig 1(A, B).

According to the EVC model, the expected value of control is determined by the expected reward and efficacy of the task. The expected reward indicates the expected outcomes of achieving the goal (e.g., monetary incentives) and efficacy refers to how likely the goal will be achieved by allocating a certain amount of control and expending a certain amount of effort (time). Computationally, the EVC model specifies that cognitive effort is allocated based on two dimensions: 1) identify the object (what to attend); 2) intensity (how much effort to allocate, compared to default level). A key assumption of this model is that intrinsic cost would be associated with higher control intensity. At the neural level, the translation between the expected value (i.e., the difference between expected rewards and costs) and corresponding effort allocation is mediated by the dorsal anterior cingulate cortex (dACC), a region that is known to play a critical role in linking adjustments in performance (Botvinik et al., 1999) with task feedback (Holroyd et al., 2004), error learning (Yeung et al., 2004) and with expected rewards (Adcock et al. 2006).

Thus defined, the EVC is an elegant, comprehensive, but highly abstract framework: it does not provide a direct mechanism by which costs and rewards are computed and associated to specific cognitive steps, nor it does make specific predictions about how motivation would precisely shift how an individual performs a task. To do so, we need a more fine-grained and detailed theory of human cognition.

One such prominent theory is the ACT-R cognitive architecture (Anderson 2007).

ACT-R

ACT-R is the most prominent and successful cognitive architecture in psychology and neuroscience (Kotseruba and Tsostsos, 2020). Surprisingly, despite the high relevance of motivation to other cognitive functions and the apparent potential and an ACT-R model of motivation, the interaction between motivation and cognitive control has been largely overlooked in ACT-R literature. Several modeling attempts have been made in order to incorporate motivation-related constructs into ACT-R, such as intrinsic motivation (Nagashima et al., 2020), emotion (Smith et al., 2021), mental fatigue (Herlambang et al., 2021; Halverson et al., 2021), and depression (van Vugt and van der Velde, 2018).

In ACT-R, knowledge is represented in two fundamental formats: chunks and production rules. A *Chunk* is a vector-like structure that stores semantic or episodic memories. A *Production rule* (or simply *production*) is a basic action unit that represents procedural knowledge as an "IF-THEN" conditional statement. Productions and chunks interact through a set of modules which represent different cognitive processes. For example, a *Visual* module encodes visual information as chunks, and a *Motor* module transforms chunks into motor outputs. Most critical to this paper are the *Goal* module (holding current goal information), the *Declarative* module (storing all declarative memories and managing their availability for retrieval), and the *Procedural* module (maintaining, updating, and selecting productions).

Each chunk is associated with a scalar value, called activation, which represents the odds of a chunk being needed in the future (Anderson, 1998). Similarly, each production has a utility value which represents the expected future rewards associated with the execution of that production. In ACT-R, only one chunk can be retrieved and only one production can be fired at any time; thus, computing chunks are selected on the basis of their activations, and production rules are selected on the basis of their utility. Utilities are learned through experience. At any time point t, the utility U of production p is calculated based on Reinforcement Learning using Eq 1, where α denotes the learning rate, R_t denotes the reward the production received at time t; s denotes the noise parameter.

$$U_t = U_{t-1} + \alpha(R_t - U_{t-1}) + s \qquad (1)$$

In ACT-R, rewards and costs are represented in time units. For instance, if the model fires a production P1 at t_1 and it receives a reward $R_{delivered}$ at t_2. The utility learning discounted the reward by the time it passed: The received amount of reward is: $R_{received} = R_{delivered} - (t_2 - t_1)$.

It should be noted that, in ACT-R, the above-mentioned *Goal* module is putatively associated with dACC (Anderson, 2007), but has no relationship to rewards and is, in fact, used only as a way to add additional information to select between competing productions. This violates established findings in neuroscience and is incompatible with the EVC. It is also a major departure from early versions of the ACT-R architecture (e.g., Anderson & Lebiere, 1998), in which goals were associated with specific values, and values were explicitly used to rank productions on the basis of a cost-benefit analysis. This older framework was, in principle, much more compatible with the EVC theory. One of our objectives is to propose a framework that conserves the current RL-based utility mechanisms but connects it with explicit goal values, re-introducing some of the most desirable features of the previous implementations.

Present study

The goal of this paper is to outline a general framework of goal-oriented motivation in ACT-R that is consistent with the EVC theory and can be implemented and deployed in any ACT-R model. This framework assumes that the goal module assigns value to chunks representing goals, with this value representing the subjective reward associated with accomplishing the goal. This is implemented by adding to the current *Goal* chunk a special motivation slot that contains a numeric value M. Once the goal is achieved, M is interpreted as the amount of reward delivered in the end. At that moment, M is automatically translated into the R_t value that is used in Eq 1., and propagates back to previous productions. Because in ACT-R, rewards are represented in time units, the value M can be interpreted in two ways: as the subjective value associated with reaching a goal, and as the maximum amount of time the model is willing to spend on a particular goal. The first interpretation is consistent with the current interpretation of the reward value R_t, while the second is consistent with the original interpretation of the goal value G in previous versions of ACT-R. By incorporating this mechanism, the *Goal* buffer is not only a passive recorder of task status, but an active power behind adaptive behaviors. Crucially, our model also attempts to account for where the intrinsic reward R_t comes from, and how motivation value M alters one's behavior by affecting the calculation of expected reward and effort.

We compare our motivation model to another well-known model of effort allocation and motivation proposed by Shenhav et al. (2013). We argue that ACT-R's procedural system provides an equivalent way of calculating the expected value of control as proposed in the EVC model. To prove that, we develop a simple effort allocation model in ACT-R, showing that ACT-R is capable of selecting the optimal strategy by weighing costs and rewards when making a decision, in line with EVC Theory. Further, we extend this simple model to a more complicated and realistic computational model of a cognitive interference task (the Simon task), augmenting it with the new motivation component. The result demonstrates that the proposed framework is compatible with the EVC model, and it helps us understand why cognitive systems vary widely in making decisions for engaging in effortful activities. Moreover, we propose a modeling approach for future

ACT-R modelers that incorporates costs, rewards, and motivational components into cognitive function. All of the model and simulation codes and data are freely available at https://github.com/UWCCDL/ACTR-Motivation

Computational Models

Motivation and Effort Allocation in an Abstract Model

To demonstrate the relationship between EVC theory and the proposed ACT-R motivational framework, we first present a simple, abstract ACT-R model and simulate the expected value of control predicted by the EVC model. To translate the continuous effort allocated in the EVC model, the abstract model assumes that different amounts of effort correspond to ten possible productions, indicated as P1, P2 … P10. The pre-conditions of these 10 productions are the same to guarantee that they are competing with each other. When the model starts running, only one of these 10 productions is selected based on the highest utility. Following this, an END production delivers a certain amount of reward at the end.

The 10 productions represent ways to perform the task that is associated with different rewards and costs in terms of mental effort. The reward is represented in terms of the probability of achieving the goal. The cost of production is represented by the time it takes to execute, which is controlled by a production-specific *:AT* parameter (for "Action time") in ACT-R. This parameter represents the effort associated with each production and, in the EVC framework, the associated cost of cognitive control. By default, it takes 0.05 seconds to fire a production, in this simple model, we assign different *:AT* to 10 productions in ascending order (0.01-0.1). Larger *:AT* suggests that the model needs to allocate more resources (time) in order to achieve the goal, while smaller *:AT* suggests that it could quickly finish the task, without spending more time on it. Specifically, production P1 is assigned to the smallest *:AT*, and P10 is assigned to the largest *:AT.*

To model expected payoffs, we set various amounts of rewards for 10 productions, in ascending order (0 - 10). P1 is assigned to the lowest reward, while P10 is assigned to the highest reward. Following Musslick, S., Shenhav, A., Botvinick, M. M., & Cohen (2015)'s suggestion, we varied the costs of the different productions according to an exponential function and varied each production's probability of receiving a reward as a sigmoid function. Thus, assigned cost increases from P1 to P10 exponentially, and the delivered rewards increase from P1 to P10 following the sigmoid function. Simply put, P1 spends the least cognitive resources but also has the lowest payoff, while P10 spends the most cognitive resources but has the highest payoff.

Two experimental conditions were simulated, corresponding to the two theoretical cases discussed by Shehav et al (2006). The first is the effect of increased task difficulty. This was simulated by decreasing each production's probability of obtaining the reward. The second was increasing the payoff. This was done by increasing the absolute value of *M* and, therefore, of the reward associated with accomplishing the task. In our framework, this is equivalent to simulating higher intrinsic motivation. We simulated 100 times per parameter set and 100 seconds (in ACT-R time) for each trial. During each trial, we recorded the counts of each production being selected to estimate the probability of selecting production. For each selected production, the received reward was also recorded to estimate the payoff.

It was expected that the probability of production being selected would show the same pattern predicted by the EVC model. Specifically, we hypothesized that, under different combinations of rewards and costs, the model would assign the greatest utility (and, therefore, the greatest probability of being selected) to the production that maximizes the difference between rewards and costs. both low-cost low-payoff productions (P1, P2), and high-cost high-payoff productions(P9, P10) are less likely to be selected than optimal cost-reward balanced productions (P6, P7, P8).

Results

Fig 1(C) and (D) demonstrate the relationship between cost, reward, and expected value of control in the abstract ACT-R model. As expected, our simple mental effort allocation model generated identical patterns of cognitive control allocation as the EVC model does. It selected the optimal production by weighing costs and rewards through utility learning in Reinforcement Learning. At a medium level of difficulty and a medium level of payoff (Fig 1C), ACT-R selects the P7 production most frequently because the utility of P7 is the highest after subtracting costs from payoffs. As the task difficulty increases, ACT-R moves to selecting production P9 most frequently. In terms of the EVC model, ACT-R now allocates more resources (a more costly production) to obtain rewards. If, on the other hand, the task difficulty decreases, ACT-R switches to a less effortful production rule (P6), which guarantees similar rewards but with less costs (shorter times).

We observe similar patterns when the Payoff is manipulated (Fig 1D). In lower payoff, ACT-R chooses P4 most frequently. As payoff increases, ACT-R tends to allocate more resources to gain more rewards by selecting higher-cost higher-reward production P5.

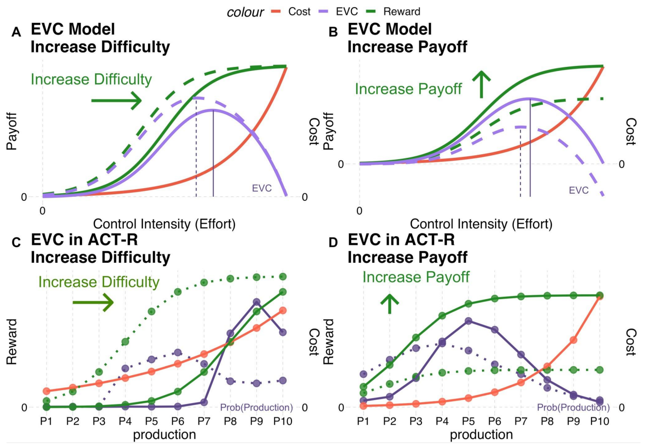

Fig 1. The recreated plot of the EVC Model (A, B); the probability of production in the simple ACT-R model (C, D). In A) and B), x-axis represents control intensity, red curves describe the costs increase exponentially as control intensity, and green curves describe the payoffs curve. The purple curve describes the expected value of control. In C) and D), x-axis represents 10 productions assigned with increasing costs and increasing rewards; Green dots denote the reward each production received; Red dots denote the cost (:AT) of each production; Purple dots denote the probability of each production being selected given the reward and cost parameter. Line types denote two conditions: increase the difficulty (A, C) and increase payoff (B, D). In (A) and (C), dashed curves describe the lowest task difficulty, while solid curves describe the highest task difficulty. In (B) and (D), dashed curves describe the lowest payoff (received rewards), and the solid curves describe the highest payoff (received reward).

Motivation and Effort Allocation in a Realistic Task

With the simple ACT-R model of effort allocation, it is safe to say that ACT-R provides a mechanistic implementation of the EVC framework. This case, however, was highly stylized: the ten productions do not represent specific cognitive operations and their costs do not realistically reflect cognitive times; this level of detail is, by contrast, the very strength of ACT-R. To examine whether the motivation framework outlined above could be translated into a realistic ACT-R model of a cognitive task of effort we applied it to Stocco et al.'s (2017) model of the Simon task. The model was chosen because it is freely available (at github.com/UWCCDL/PSS_Simon) and is the same task used by Boksem et al (2006) to study motivation. The Simon task requires participants to respond to visual stimuli by pressing a leftward button to one shape (e.g., a circle) and a right button to another (e.g., a square). Congruent trials are where the stimulus is displayed on the same side as the rule dictates, while incongruent trials are on the opposite side. This paradigm was widely used in neuropsychological studies to assess the ability to inhibit cognitive interference that occurs when the processing of a particular visual property hinders the simultaneous processing of a second stimulus property.

Fig 2. provides a complete overview of the model. It is composed of 4 main steps: 1) Encoding visual stimulus 2) Retrieving a Simon rule 3) Responding and 4) Monitoring performance. The model starts by encoding a cue stimulus, and then it selects which dimension of the Simon stimulus to attend to, color or shape. Followed by stimulus processing, it retrieves the corresponding rule. The attended dimension provides spreading activation that facilitates the retrieval of the associated rule (a feature common to other response interference models in ACT-R: Lovett 2001; van Rijn 2009). The equation below describes the activation of chunks calculated with a base-level learning function (B_i), which reflects the recency of previous retrievals, as well as a spreading activation component that reflects the degree to which a chunk matches the contextual components, i.e., the values of every slot j in every buffer k (Eq 2).

$$A_i = B_i + \sum_k \sum_j W_{kj} S_{ji} + \varepsilon \qquad (2)$$

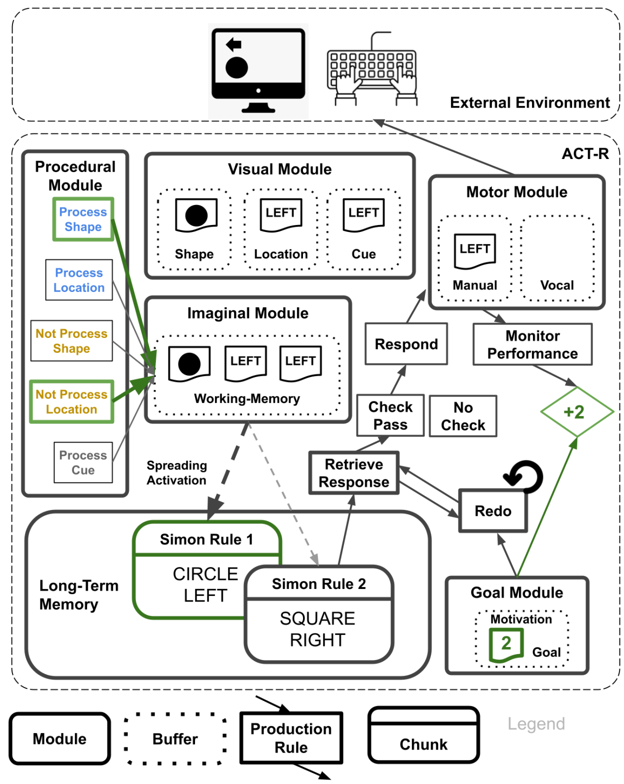

Fig 2. The flowchart of the motivation model in ACT-R

Before a response is made, a "check" production executes a final verification step and, if it finds that the response is incorrect, attempts to re-allocate attention and retrieve another answer. The model contains additional assumptions (about the nature of competing processes in attention) that are not relevant to the goals of this paper, and will not be discussed. What is relevant, instead, is the nature of the "check" process. In the original paper, it was constrained to occur only once. In our extended simulation, we removed this limitation and allowed the model to check as many times as possible. Because the number of checks performed corresponds to how much effort is used to control attention, it provides a natural way to model the cognitive control effort in the task.

In addition, we incorporated a motivation value *M* in the *Goal* chunk and added a self-monitor production that assesses whether the response was correct and, if so, triggers a reward of magnitude *M*. Like the reward and cost parameter in ACT-R, M is also in time units, representing how many seconds the model is willing to continue working on the task. Note that the model will continue checking only if it finds the current response incorrect. Thus, if M is set high, the model would have more opportunities to correct its response. On the contrary, if M is small, it would have less opportunities to refocus attention.

To increase the task difficulty, Boksem et al., (2006)'s paradigm added cues stimulus, where 80% of the cues were valid. They identified an interesting post-error slowing effect in which participants tended to respond more carefully and slowly after they thought they made a mistake. This process is believed to reflect adjustments in the allocation of mental effort, which is key to the EVC and our motivation framework. Critically, we verified that post-error slowing is *not* produced by Stocco's original model; thus, any success in reproducing this effect must be due to our additional changes.

It was predicted that this model would be able to change strategies based on the probability of gaining rewards and costs. If it never checks, the likelihood of gaining rewards will become low because of many errors. If the model checks a lot, the expected payoff will also decline because of the increasing costs. Therefore, the model should weigh costs and rewards to decide the attempts of checking optimally.

We varied Motivation parameter M, the task difficulty parameter VC (which represents the percentage of cues that are valid cues) as well as the initial cognitive control costs through the *action time* (*AT*) parameter in ACT-R, which determines the time (and, thus the effort) needed to execute each production The parameter space is shown in the table below.

Parameter	Value	Meaning
M	0.5 - 10	Motivation
VC	0 - 1	Task difficulty
AT	0.01 - 0.1	Cost of control at T_0

Results

To test the validity of our model, we first compared simulated data to Boksem et al.'s (2006)'s empirical data. To focus on Simon effects, we fixed the difficulty parameter VC = 0.5, cost parameter *AT* = 0.05, and limited motivation parameter to a medium-range (0.1 - 2). Fig. 3 confirmed that our model still reproduced the main Simon effects. Incongruent trials were associated with lower accuracy, and longer response time than congruent trials, same for invalid trials.

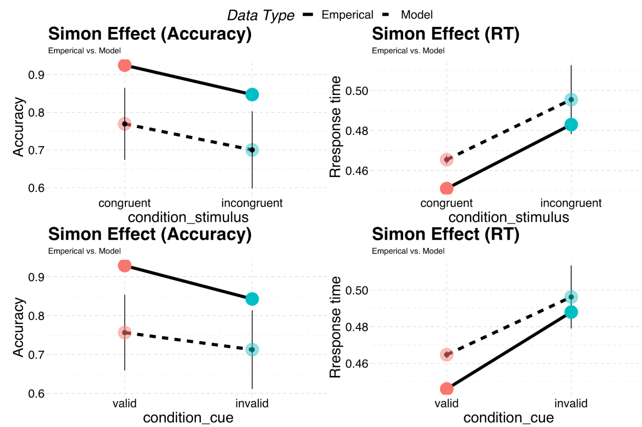

Fig 3. Model simulation results of Simon effect vs. empirical findings of Simon effects (Boksem et al., 2006). Solid lines denote effects of empirical data and dashed lines denote effects simulated by our model.

In addition, our model could reproduce the post-error slow effect observed in empirical data (Fig 4). Note that this effect could not be reproduced by the original model (Stocco et al., 2017) under any combination of parameters; thus, it is a unique feature of the added motivation mechanism. Specifically, post-error slowing is a consequence of the model adjusting control after a mistake is made.

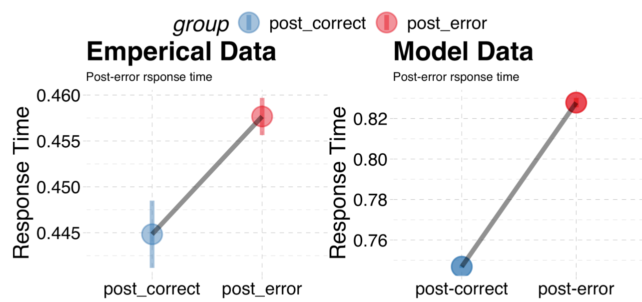

Fig 4. The post-error slowing effects in both empirical data and ACT-R model simulation data (across all parameter combinations). The standard error for both empirical data and model data is shown in the plot.

In the Simon model, the degree of cognitive control is determined by how often the CHECK production is employed before a response is made. Additional firings of the CHECK productions result in repeated allocations of attention and, thus, more time spent before making a response. As hypothesized, we found that a model with a lower *M* ($M < 2.5$) would check only once or never check, while a model with a high *M* ($M > 7.5$) tends to check more. For example, when $M < 2.5$, the model performs an average of 0.54 checks, when $M < 7.5$, the model performs 0.81 checks, and when $M >= 7.5$, it performs on average 1.01 checks.

As predicted by the EVC theory, the relationship between motivation, number of checks, and utility of the CHECK production are complex and nonlinear. To examine this relationship, we fixed the parameter ValidCue% to 0.5. Fig.

5 represents the resulting relationship between the costs, rewards, and allocated control. In the figure, the x-axis represents the intensity of control as the number of firings for the Check production, and the y-axis represents rewards, costs, and utilities in time units. The cost curve (red line) is represented by the total response time the model takes as a function of the count of checking. Moreover, in our self-monitoring process, once the model verifies that the response was correct, a reward equal to M is delivered. The utility of the CHECK production (purple line) represents the expected value of control in the EVC model. In line with the EVC model, our ACT-R model predicts that the model will be encouraged to invest more effort if expecting a higher payoff, but as costs increase, the expected reward decreases and the model decides to stop investing more effort by reserving effort. Note that, although the model *could* achieve greater performance through greater control, it naturally sets to an estimated value of one check per trial because, at this level, the payoff is maximal: additional checks have many diminishing returns. Incidentally, this is precisely the number of checks that were determined to yield optimal results in Stocco et al (2017) and Lovett (2005).

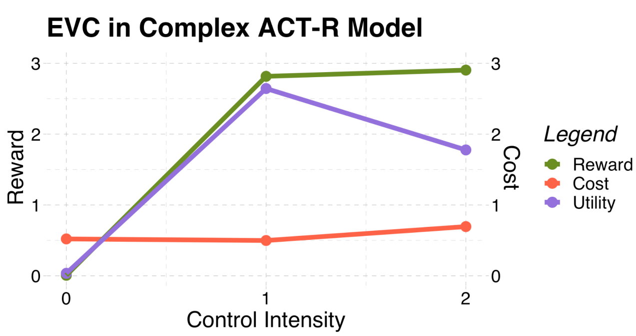

Fig 5. Expected value of control in the Simon Task model. Control intensity is expressed as the number of firings of the CHECK production. Note that, even when higher rewards would be possible at a higher level of control, the model naturally shifts to the amount of control that maximizes the difference between rewards and costs.

Discussion

In this paper, we have proposed a mechanistic interpretation of motivation within the ACT-R cognitive architecture. Specifically, we propose that motivation can be modeled by assigning a value M to the current model's goal and translating this value as the reward R_t that is triggered when the goal is accomplished. With this mechanism in place, ACT-R's utility learning mechanism then provides a way to adjust the specific combination of productions that are used to perform a task. Because in ACT-R, rewards and time spent on a task are expressed on the same scale (and rewards are adjusted by the time elapsed), the motivation parameter M can be equivalently expressed as the subjective reward associated with accomplishing the goal and the maximum amount of time that the model is willing to spend on the task. We first demonstrated, using a simple abstract model, that this mechanism is equivalent to the EVC theory. We then showed how this mechanism can be easily implemented in an existing model of a common laboratory task (the Simon task) and used to account for experimental effects that would otherwise go unmodeled, such as post-error slowing, the effect of difficulty, and even fatigue. All of these effects can be understood as ways in which the model flexibly copes with changes in task demands.

A number of limitations must be acknowledged. First, the level control intensity is quantified by the counts of checking attempts, as a discrete variable. Future work will be needed to address these issues and expanding our model to represent the control intensity with a continuous variable could be the next step of research. Moreover, individual variability in motivation could be examined in future modeling work, specifically how motivation affects the response time rather than accuracy for individuals putting different priorities in speed vs. accuracy tradeoffs (Boksem et al., 2006).

These limitations notwithstanding, we believe that our results are noteworthy for three reasons. In addition to providing a way to implement motivation into ACT-R, our framework provides a more complete view of the role of the Goal module in ACT-R. Currently, the model's capabilities make it distinguishable from the Imaging module. By connecting it to the amount of reward that is generated, this framework provides an interpretation that is more in line with neuroscientific data. It also provides a connection to the original interpretation of the goal in previous versions of ACT-R, as well as the original production selection mechanisms. Finally, it provides a way to better fit models at the individual levels, decoupling the effects of individual capacity and motivation.

References

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron, 50*(3), 507–517.

Anderson, J. R., & Lebiere, C. J. (1998). The atomic components of thought. *Psychology Press.*

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: Oxford University Press. ISBN 0-19-532425-0.

Boksem, M. A., Meijman, T. F., & Lorist, M. M. (2006). Mental fatigue, motivation and action monitoring. *Biological psychology, 72*(2), 123-132.

Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S., & Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature, 402*(6758), 179-181.

Halverson, T., Myers, C., Gearhart, J., Linakis, M. W., & Gunzelmann, G. (2021). Physio-cognitive modeling: explaining the effects of caffeine on fatigue. Paper presented at Virtual MathPsych/ICCM 2021. Via mathpsych.org/presentation/619.

Herlambang, M. B., Cnossen, F., & Taatgen, N. A. (2021). The effects of intrinsic motivation on mental fatigue. *PloS one, 16*(1), e0243754.

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R. B., Coles, M. G., & Cohen, J. D. (2004). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nature neuroscience*, *7*(5), 497-498.

Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The effort paradox: Effort is both costly and valued. *Trends in Cognitive Sciences, 22*(4), 337–349.

Kotseruba, I., & Tsotsos, J. K. (2020). 40 years of cognitive architectures: core cognitive abilities and practical applications. *Artificial Intelligence Review, 53*(1), 17-94.

Kouneiher, F., Charron, S., & Koechlin, E. (2009). Motivation and cognitive control in the human prefrontal cortex. *Nature Neuroscience, 12*(7), 939-945.Musslick, S., Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2015). A computational model of control allocation based on the expected value of control. *In The 2nd multidisciplinary conference on reinforcement learning and decision making.*

Nagashima, K., Morita, J., & Takeuchi, Y. (2020). Modeling intrinsic motivation in ACT-R: Focusing on the relation between pattern matching and intellectual curiosity. *Proceedings of the 18th International Conference on Cognitive Modelling*

Niv, Y., Daw, N., & Dayan, P. (2005). How fast to work: Response vigor, motivation and tonic dopamine. *Advances in neural information processing systems, 18*.

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron, 79*(2), 217–240.

Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a rational and mechanistic account of mental effort. *Annual Review of Neuroscience, 40*(1), 99–124.

Smith, B. M., Thomasson, M., Yang, Y. C., Sibert, C., & Stocco, A. (2021). When fear shrinks the brain: A computational model of the effects of posttraumatic stress on hippocampal volume. *Topics in Cognitive Science, 13*(3), 499-514.

Stocco, A., Murray, N. L., Yamasaki, B. L., Renno, T. J., Nguyen, J., & Prat, C. S. (2017). Individual differences in the Simon effect are underpinned by differences in the competitive dynamics in the basal ganglia: An experimental verification and a computational model. *Cognition, 164*, 31-45.

Stocco, A. (2018). A Biologically Plausible action selection system for cognitive architectures: Implications of basal ganglia anatomy for learning and decision-making models. *Cognitive Science*

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. https://doi.org/10.1111/tops.12318

Vroom, V. H. (1964). Work and motivation. *New York: Wiley.*

Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychological review, 111*(4), 931.