Proceedings of the 12th International Conference on Cognitive Modeling, Ottawa, Canada.

Towards Adding Bottom-Up Homeostatic Affect to ACT-R

Christopher L. Dancy (cld5070@psu.edu)
Ryan Kaulakis (rmk216@ist.psu.edu)
The College of Information Sciences and Technology, The Pennsylvania State University
University Park, PA 16802

Abstract

Extending a cognitive architecture with a representation of physiology allows one to simulate the effects of homeostasis on cognitive processes. The underlying physiological substrate in the ACT-RΦ architecture allows one to model potential interactions between homeostatic affect and cognition. In this paper, we describe an extension of the ACT-R utility mechanism that uses information from the physiological representation provided by ACT-RΦ. We also discuss a model that uses this new system to simulate the effects of homeostatic-based thirst on rule utility. The model completes a variation of the ultimatum game that used a primary reward (water) as opposed to a monetary reward. In the experimental version of this game, researchers induced thirst in some participants by administering hypertonic saline via an intravenous line. Our model was able to represent both the physiological (i.e., hypertonic saline) and behavioral (i.e., forced selection of a primary reward) portions of the experiment. The model also can potentially provide predictions of behavior during the task that could not be observed due to the nature of the experimental condition.

Introduction

Extending a cognitive architecture with a representation of physiology allows one to model and simulate the effects of homeostasis on cognitive processes. Homeostasis, in this sense, describes the process of human physiological systems changing to remain in a stable state. This physiological change can result in changes of affect that in turn can alter decision-making and memory processes.

ACT-RΦ (Dancy, Ritter, & Berry, Accepted) is an extension of the ACT-R (Anderson, 2007) cognitive architecture that adds an underlying physiological substrate. In this paper, we describe a further extension of ACT-RΦ that adds a representation of homeostatic affect (subjective experience/feeling). Though there have been previous descriptions and implementations of cognitive and agent-based architectures that include a representation for homeostasis (e.g., Bach, 2009; Silverman et al., 2012), none have used a representation of bodily processes that attempts to model actual human physiology. This is an important step in human behavior representation as it allows a modeler to more closely map simulation changes and results with existing empirical data.

To keep in line with the theoretical nature of cognitive architectures, it is important that one has a theoretical account of how a physiological system may affect cognition. We have adapted theories from existing *SEEKING* (Panksepp, 1998) and incentive salience, or *wanting*, literature (Berridge, Robinson, & Aldridge, 2009); we use these theories to describe how some homeostatic processes affect cognitive processes. Taking this theoretical perspective is particularly useful in continuing to explore the question: “How can the human mind occur in the physical universe?” (i.e., Anderson, 2007).

SEEKING, Wanting, and Incentive Salience

The SEEKING system is a primitive affective system that is the driving force behind appetitive-approach actions (J. Wright & Panksepp, 2012). The neural substrates of this system are subcortical and evolutionarily older than cortical structures. Thus, the SEEKING system is seen as one of the most primitive affective systems.

Wanting and incentive salience (Berridge et al., 2009) are very similar to SEEKING, but are typically described in terms of a secondary learning process (e.g., conditional learning; Zhang, Berridge, Tindell, Smith, & Aldridge, 2009). Thus, incentive salience describes the process of learning that occurs when the wanting affective process brought on by unconditioned stimuli, is coupled to conditioned stimuli. This higher level of inquiry is important as it provides a connection between the underlying primitive affective circuitry (and associated behavior) of wanting (comparable to what is denoted by J. Wright & Panksepp, 2012, as the SEEKING system) and learning and decision-making behavior often studied in neuroscience and psychology.

Though the two theoretical takes are very similar, they inherently focus on different levels; this makes combing of SEEKING and incentive salience into a cohesive framework an attractive prospect. Taking both theoretical accounts as separate entities and combining them thus allows one to build a cohesive computational framework by combining fairly atomic representations. We have chosen to include these theoretical accounts in the ACT-RΦ architecture (Figure 1) as part of the Affect System.

ACT-RΦ: ACT-R with Physiology and Affect

The ACT-RΦ architecture (Dancy et al., Accepted) combines the existing ACT-R cognitive architecture with the simulation system HumMod (Hester et al., 2011); HumMod contains a top-down model of human physiology (e.g., representations of organs and hormones). ACT-RΦ (Figure 1) extends ACT-R with new modules to communicate with HumMod. These modules also use the physiological values obtained to modulate the architectures cognitive components. In turn, modelers can also simulate potential ways cognition can affect physiology.

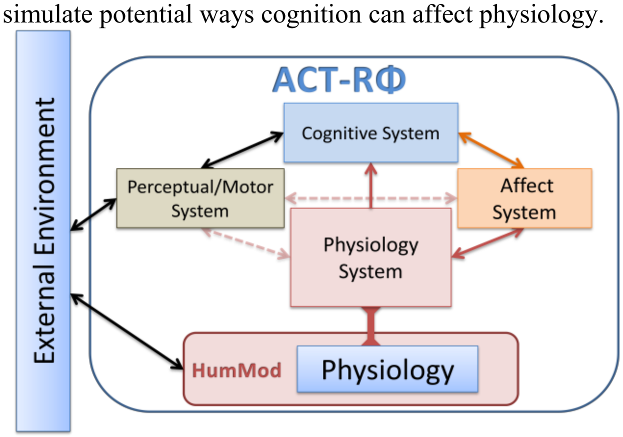

Figure 1: A high level diagram of the ACT-RΦ architecture. Dotted lines represent connections that are not yet completed in the software system.

The *physio* module (contained in the *physiology system* in Figure 1) facilitates communication between the ACT-R and HumMod components of the architecture so that they can affect each other in (simulated) real-time. This module can be used to retrieve and change variable values in HumMod. The module also contains any direct connections between a physiological variables and cognitive parameters (e.g., a function for epinephrine and memory noise; Dancy et al., Accepted).

The *SEEKING* module is meant to represent the functionality of the SEEKING system described in Panksepp (1998) and J. Wright and Panksepp (2012). This system (part of the *Affect System* in Figure 1) is modulated by physiological variables used to represent homeostatic-affect (that is, hunger, thirst, and skin temperature, Table 1) as well as a goal related value. These representations compete in a winner-take-all fashion and only one effect modulates behavior at any given time. The SEEKING system is updated every *SEEKING-delay* seconds and this parameter can be changed using the ACT-R *sgp* function.

Table 1: Homeostatic functionality present in ACT-RΦ and important related variables in HumMod

Function	HumMod
Hunger	GI-Lumen, Leptin, Glucose
Thirst	Osmo-Rec, BodyH20.vol
Skin Temperature	Skin-Temperature

The winning function affects existing production utility via an ACT-R utility-offset function. The utility offset adds the value found using equation 1; this is meant to match the equation described in Zhang et al. (2009). In equation 1, a variable *k* represents some physiological factor that dynamically modulates utility.

$$U_A = U + \log(k) \quad (1)$$

$$k = sVal * e^{reward_{max}} \; s.t. \; sVal \in [0,1] \quad (2)$$

$$sVal_{Thirst} = \frac{reward_{max} * (osmo - osmo_{base})}{osmo_{base} - osmo_{min}} + noise \quad (3)$$

The *sVal* variable in equation 2 is some normalized value to represent affect; $reward_{max}$ is a parameter that can be set by the modeler via the normal ACT-R *sgp* function that is used to change architectural parameters. In the model below, sVal is a normalized representation of thirst found using equation 3.

A separate *Affective-Associations* module is used to couple affective values with ACT-R productions. This not only affects the normal cached utility values, but also dynamically modulates the utility according to current physiological need. This functionality is contained within a separate module to reflect the existing neuroscience and psychology literature that indicates this affective associative functionality is likely separate from existing affect generating and memory systems (e.g., McGaugh, 2004; Seymour & Dolan, 2008). A user of the system can explicitly attach affect to a production rule using the spp-a function, an extended version of the ACT-R production-utility function; this extended version of the spp function allows a user to specify an affect-based function (e.g., thirst) that should be coupled with a production.

Using these modules together with a process model allows one to simulate different ways bottom-up homeostatic-affect may modulate human cognition and behavior. Though other architectures exist with representations of homeostatic-focused motivation/affect (e.g., Bach, 2009; Silverman et al., 2012), the representation of physiology used in ACT-RΦ is more easily mapped to physiological data found in experimental studies. We demonstrate this characteristic of the architecture with a model that displays a potential way thirst can affect cognition. We model an existing task and data that observes the effects of artificially induced thirst on a modified version of the ultimatum game (N. D. Wright, Hodgson, Fleming, Symmonds, Guitart-Masip, & Dolan, 2012).

Modeling the Effect of Thirst on Cognition

We developed a model based on an experiment conducted by N. D. Wright et al. (2012) in which primary rewards, i.e. water, were used as the bargaining object in an altered version of the ultimatum game. The ultimatum game is a task where one player (the proposer) proposes a specific division of some endowed capital to a second player (typically referred to as a responder). The responder can then accept or reject the offer; neither player receives any portion of the endowment if the responder rejects the proposed offer. In their version of the ultimatum game, participants were induced to be thirsty through the injection of a hypertonic saline solution via an IV, or injected with isotonic saline (control group). Participants then met in a room with two other participants, interacted, then were taken to their own room for testing. Each participant was then physically presented with two glasses of water, one that represented the proposer's offer and one that represented the proposer's keep. They were told that the proposer was another participant (that they had previously met), but the proposition was actually constant in the experiment; the proposition was for the participant to receive a glass of water that was 12.5% full, while the proposer received a glass of water that was 87.5% full. The subjects were then given 15s to respond by circling "accept" or "reject" on a form. Due to the inherent safety issues of the experimental manipulation (artificially induced thirst) participants were only asked to complete one round.

Their results are interesting, because the objective measure of thirst (osmolarity) did not directly correlate with the acceptance rate of the offered water. Rather, the subjective measure of thirst (1-10 scale) correlated strongly with the acceptance rate. This indicates (in their wording) a "dissociative" gap between objective and subjective perception, which is similar to an affective valuation in our terminology. For this reason, we chose their scenario for modeling affective moderation of cognitive decision-making.

The Model

This model implements the experimental condition of the Ultimatum Game with primary rewards (Wright et al. 2012). The model is "injected" with hypertonic saline, is presented with an offer, makes a consideration of the fairness of the offer, and then accepts or rejects it, all as in the normal game scenario. A few specific details of the model differ from the original experiment; for example, the model does not directly observe a glass of water, but is presented with a GUI which displays the numeric percentage being offered. This percentage is the ratio of the offered volume of water over the total water times 100 (12.5% for the purposes of this task).

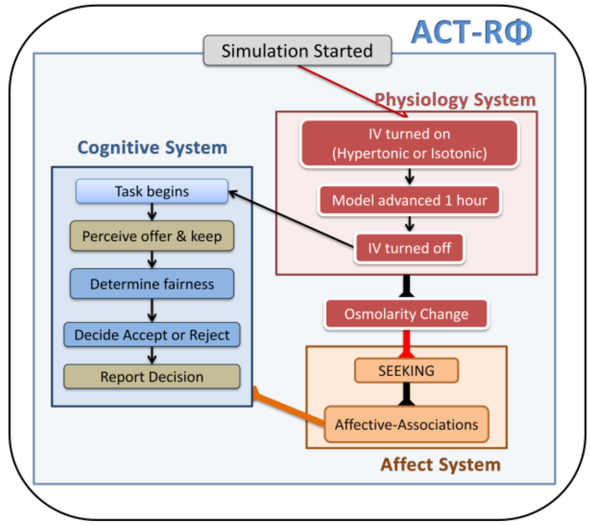

Figure 2: A diagram of the ACT-RΦ model. The thicker lines represent a continuous communication.

Once the offer has been processed through visual inputs, the fairness of the offer is evaluated by a pair of competing productions which determine the fairness of the offer. These productions compare the offered value to a computed fair value (at least an equal distribution of the water), and if the offered value is less than fair, one production rejects the offer, otherwise the other matches and accepts. However, these productions use the similarity and utility functionality in ACT-R, such that the application of the productions can be influenced by subjective valuation moderators. When moderated in this way, the effective level of fairness can be raised or lowered. The integration of the ACT-RΦ modulators discussed earlier allow an underlying homeostatic process to affect procedural memory selection. The accept and reject productions use the partial matching and utility mechanisms in ACT-R to determine rule selection with the accept rule receiving an "affective boost" when the model is thirsty (see figure 4). The model also did not have extraneous social knowledge about the other "player" in the Ultimatum Game, but the fairness appraisal productions were designed to replicate basic unfairness rejection behavior nonetheless, rather than being economically rational and accepting any offer.

Model Results

Running the model 1000 times yielded results similar to those exhibited in the original task (Table 2); running the model also yielded standard error mean (SEM) values that fit the criteria suggested by Ritter, Schoelles, Quigley, and Klein (2011). The model produced an average subjective thirst of 7.29 with a standard deviation of 1.71; these results were very close to the results found in the original experiment (7.3 and 1.60). The specific osmolarity of the *accept* and *reject* groups was not reported by N. D. Wright et al. (2012).

Table 2: Results from the ACT-RΦ model

Decision	Osmolarity (sd)	Subj. Thirst (sd)
$Accept_{experiment}$	Not Reported	8.90(1.7)
$Accept_{model}$	306.37(0.2)	7.94(1.2)
$Reject_{experiment}$	Not Reported	5.60(1.6)
$Reject_{model}$	305.86(3.74)	4.82(1.2)
$Both_{experiment}$	310(5.0)	7.30(1.6)
$Both_{model}$	306.27(1.7)	7.29(1.7)

Though mean subjective thirst results for the combined (accept and reject) results were similar between the experiment and model, the categorized subjective thirst values were lower on average in the model as compared to the experimental data (Table 2). An explanation may be the significantly higher number of data points from our model as compared to the original experiment. This illustrates a potential advantage of using ACT-RΦ to simulate an experiment that requires a high setup and run cost.

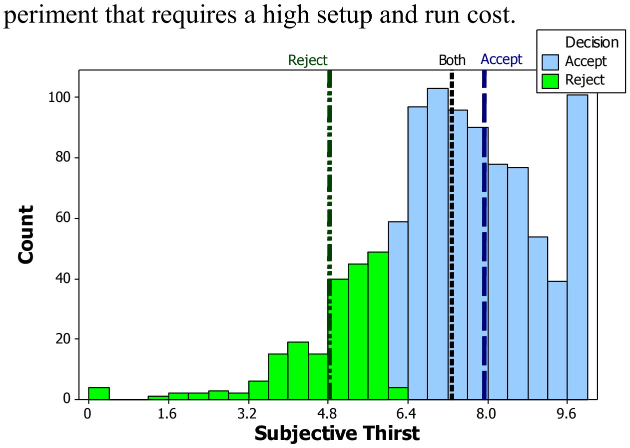

Figure 3: A histogram of the subjective thirst results produced by the model. The lines represent the mean for subjective thirst for models that accepted the offer, rejected the offer, and the two combined.

As expected (from equation 1) a plot of subjective thirst vs. utility of the acceptance production rule (figure 4) produces a logarithmic curve. Accordingly, as the model gets *thirstier*, the effect of homeostatic affect on cognition increases and accelerates.

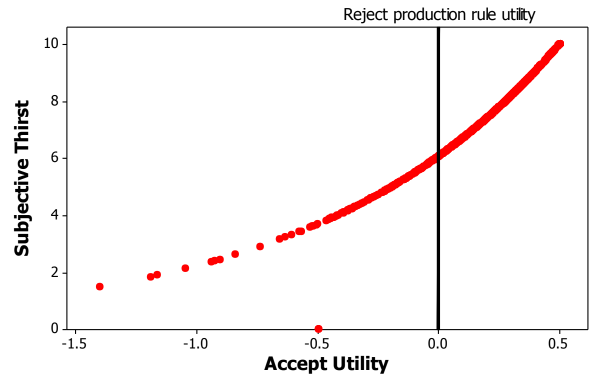

Figure 4: A plot of subjective thirst vs. acceptance production rule utility.

The model had a higher acceptance rate (79%) than the acceptance rate reported from the original experiment (50%). A likely reason for this disparity is that the *fairness* representation in the model does not have enough weight on the outcome. Developing a more complex fairness motivation representation will likely decrease the acceptance rate of the model. However, it is not clear whether the model should need to decrease to a 50% rejection rate as this number may also be representative of the small sample size and could change with an increase in participants.

Discussion and Conclusions

With this model, we were able to display one way that one can simulate bottom-up homeostatic affect. As opposed to explicitly changing goal behavior, the model featured an underlying change in utility based on a thirst that controlled an underlying motivational SEEKING or wanting system. Using a more realistic underlying physiological substrate to drive homeostatic-affective change gave the opportunity to explicitly model the thirst induction (intravenous hypertonic saline injection) portion of the experimental task protocol.

Though we have offered a partial explanation for some of the effects shown in the experiment conducted by N. D Wright et al. (2012), there are several areas that both the model and architecture can be expanded to provide a deeper explanation of this task and others. We offer a brief discussion on potential areas of exploration related to simulating homeostatic-affect and implications for such work.

Expanding the Ultimatum Model

Though we provide a new representation of implicit physiological value with our model, it could be expanded to better encompass the cognitive and social aspects of a person completing the ultimatum game. The model could have a more complex explicit valuation system to make the determination of whether to accept or reject an offer; we chose to exclude this in the current work as the original experiment only had participants make one choice (accept or reject) as a responder. Similarly, the model could provide a more expanded declarative representation of the task that would not only add to the cognitive portions of the model, but also to add homeostatic-based (affective) declarative memories. Previous models of similar tasks (e.g., Juvina, Lebiere, Martin, & Gonzalez, 2011; Lebiere, Wallach, & West, 2000) offer insights into the particulars of model expansion.

As detailed previously, participants were only asked to make one choice and only as the game responder. This is mainly due to the difficulty in artificially inducing thirst by physiological manipulation safely; it can be unsafe for the experimenter to induce thirst long enough to play an entire game. However, with an expanded model, our work will provide predictions for such an experimental protocol nonetheless as we are able to follow the exact physiological manipulations present in the experimental procedure. A particular advantage of providing a more realistic physiological substrate is that one can then explore potential experimental effects that would be dangerous in a normal human experiment. Such an expansion would also give one an opportunity to explore the implications of using the more *neuroscience-friendly* reward equation (Fu & Anderson, 2006) in conjunction with the incentive salience offset ($\mathbf{log}\ (\boldsymbol{k})$).

Explicit Homeostatic-Affect

With the model presented, we provide a representation for bottom-up implicit homeostatic affect. While the model and system implement connections important for affective memory, there is a declarative affective memory system that is not currently represented in this work. Adding an affective association to the declarative module should be explored in future work.

Past work with modeling emotional declarative memory will likely provide a good starting point for adding explicit homeostatic-affect. Cochran, Lee, and Chown (2006), for example, discuss representing the impact of emotion and arousal on declarative memory. Though that work would need to be expanded to encompass homeostatic-based affect, it nonetheless provides insights into the possible mechanisms for affective change.

Homeostatic-Affect and Performance

With these expansions one can begin to provide predictions of the performance change during a task that may accompany homeostatic need. Representing the effects of homeostasis on performance during normal daily task allows one to examine important points of time when the effects of homeostatic-need can severely impeded the completion of a task. One could for instance use results from existing studies (e.g., Lewis, Snyder, Pietrzak, Darby, Feldman, & Maruff, 2011; Tuk, Trampe, & Warlop, 2011) to simulate the effects of need to void on driving performance; a related model was made to represent the effects of sleep-loss on driving (Gunzelmann, Moore Jr, Salvucci, & Gluck, 2011).

A particular advantage of using this architecture to look at these effects is the ability to enact several different types of homeostatic need (and other physiological systems) at one time to see the inter-physiological effects and see how these connections potentially affect cognition. The best method to determine a *winning* homeostatic need as these physiological processes interact remains an open problem. Nonetheless, examining the interaction of these processes on a physiological and cognitive level remains important as these factors continuously affect human task performance over the course of a day.

Final Remarks

This work explores a way to integrate homeostatic-affect into an architecture with a unified theory of cognition and an underlying physiological substrate. Though, there are other manners in which homeostasis can affect cognition (e.g., explicit affective memory), this work still provides an important stepping stone towards including an integrated account of homeostatic-affect on cognition. The use of a computational architecture with more realistic representations of integrated physiology also gives one the opportunity to explore how interaction between physiological processes can modulate cognitive processes.

Acknowledgements

We thank members of the Applied Cognitive Science Laboratory at the Pennsylvania State University and Dr. Robert Hester for useful discussion related to this work. This work was partially supported by the Bunton-Waller Fellowship program.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: OUP.

Bach, J. (2009). *Principles of synthetic intelligence: PSI: an architecture of motivated cognition*. New York, NY: OUP.

Berridge, K. C., Robinson, T. E., & Aldridge, J. W. (2009). Dissecting components of reward: 'liking', 'wanting', and learning. *Current Opinion in Pharmacology, 9*(1), 65-73.

Cochran, R. E., Lee, F. J., & Chown, E. (2006). Modeling emotion: Arousal's impact on memory. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, 1133-1138:Vancouver, BC, Canada.

Dancy, C. L., Ritter, F. E., & Berry, K. (2012). Towards adding a physiological substrate to ACT-R. In

Proceedings of the 21st Conference on Behavior Representation in Modeling and Simulation, 78-85. The BRiMS Society:Amelia Island, FL.

Dancy, C. L., Ritter, F. E., & Berry, K. (Accepted). Using a cognitive architecture with a physiological substrate to represent effects of psychological stress on cognition. *Computational & Mathematical Organization Theory*.

Gunzelmann, G., Moore Jr, R. L., Salvucci, D. D., & Gluck, K. A. (2011). Sleep loss and driver performance: Quantitative predictions with zero free parameters. *Cognitive Systems Research, 12*(2), 154-163.

Hester, R. L., Brown, A. J., Husband, L., Iliescu, R., Pruett, D., Summers, R., et al. (2011). HumMod: A modeling environment for the simulation of integrative human physiology. *Frontiers in Physiology, 2*(12).

Juvina, I., Lebiere, C., Martin, J., & Gonzalez, C. (2011). Intergroup prisoner's dilemma with intragroup power dynamics. *Games, 2*(1), 21-51.

Lebiere, C., Wallach, D., & West, R. L. (2000). A memory-based account of the prisoner's dilemma and other 2x2 games. In *Proceedings of the Third International Conference on Cognitive Modeling*, 185-193. Universal Press:Veenendaal, Netherlands.

Lewis, M. S., Snyder, P. J., Pietrzak, R. H., Darby, D., Feldman, R. A., & Maruff, P. (2011). The effect of acute increase in urge to void on cognitive function in healthy adults. *Neurourology and Urodynamics, 30*(1), 183-187.

McGaugh, J. L. (2004). The Amygdala Modulates the Consolidation of Memories of Emotionally Arousing Experiences. *Annual Review of Neuroscience, 27*(1), 1-28.

Panksepp, J. (1998). *Affective neuroscience: The foundations of human and animal emotions*. New York, NY: OUP.

Ritter, F. E., Schoelles, M. J., Quigley, K. S., & Klein, L. C. (2011). Determining the number of simulation runs: Treating simulations as theories by not sampling their behavior. In S. Narayanan & L. Rothrock (Eds.), *Human-in-the-loop Simulations: Methods and Practice* (pp. 97-116). London: Springer-Verlag.

Seymour, B., & Dolan, R. (2008). Emotion, Decision Making, and the Amygdala. *Neuron, 58*(5), 662-671.

Silverman, B. G., Pietrocola, D., Nye, B., Weyer, N., Osin, O., Johnson, D., et al. (2012). Rich socio-cognitive agents for immersive training environments: case of NonKin Village. *Autonomous Agents and Multi-Agent Systems, 24*(2), 312-343.

Tuk, M. A., Trampe, D., & Warlop, L. (2011). Inhibitory Spillover. *Psychological Science, 22*(5), 627-633.

Wright, J., & Panksepp, J. (2012). An Evolutionary Framework to Understand Foraging, Wanting, and Desire: The Neuropsychology of the SEEKING System. *Neuropsychoanalysis, 14*(1), 5-39.

Wright, N. D., Hodgson, K., Fleming, S. M., Symmonds, M., Guitart-Masip, M., & Dolan, R. J. (2012). Human responses to unfairness with primary rewards and their biological limits. [10.1038/srep00593]. *Scientific Reports, 2*.

Zhang, J., Berridge, K. C., Tindell, A. J., Smith, K. S., & Aldridge, J. W. (2009). A Neural Computational Model of Incentive Salience. *PLoS Comput Biol, 5*(7), e1000437.