

A Standard Model of the Mind Needs a Body

Christopher L. Dancy

Department of Computer Science
Bucknell University
christopher.dancy@bucknell.edu

Frank E. Ritter

College of Information Sciences and Technology
The Pennsylvania State University, University Park
frank.ritter@psu.edu

Abstract

A standard model of the mind with realistic representation of human-like behavior needs a body to represent the interactions of the mind with the world. Though several cognitive processes continue to be studied in the pursuit of a better understanding of the mediators of human-like behavior, the study of interactions between non-cognitive and cognitive human processes remains sparse. We present two aspects that have not been often seen as part of a standard model of the mind, but that appear to be essential to representing at least the human mind. These are a body and a way to interact with the world. The body supports cognition through energy and adaptation, and the brain helps run the body, committing significant resources to this control. The body also provides a way to interact with the world, through vision primarily, as well as motor output. Without these two aspects, the standard model will just be a brain in a vat.

Introduction

A standard model of the mind with realistic representation of human-like behavior needs a body. Though cognitive processes continue to be studied in the pursuit of a better understanding of the mediators of human-like behavior, the study of interactions between non-cognitive and cognitive processes remains sparse. Particularly processes that may be described on the physiological-level, but nonetheless have behavioral effects, are important to the study of the human-mind and human-like minds. Thus, it follows that having a representation for a body is important to a standard model of the mind.

It is also important for the mind to be able to interact with the world, otherwise it is really just a brain in a rather impoverished vat. Thus, the model of the senses, particularly vision and some motor output, need to be specified in

enough detail to be used routinely. A colleague of ours noted that they “budgeted 25% of each project to support interaction”, which suggests that interaction is not a solved, and is an important part of a standard model to be specified in more detail than it has so far.

The representation of physiology of a body is important for both interactions with cognitive processes, and also processes associated with emotional behavior. Indeed, one may think of a continuum of behavior between physiological, affective, and cognitive processes that all interact to mediate intelligent behavior over several time-scales. The body is also useful as a mind design heuristic; physiological processes serve as a guide towards intelligence behavior. Functionally, we may not need to find a cognitive answer to a question that evolution chose to answer by interacting a mind with a body. Lastly, representing the body also provides an opportunity to understand the process as opposed to just the state of factors that affect sub symbolic representations in the mind, the latter of which we currently see in models of cognition and the mind. Having a model for these processes and how they interact with typically modeled mind processes is useful for studying the functional complexity and behavior of underlying systems.

In the sections following we discuss some moderators related to human physiology that are known to have vast effects on human-behavior and thus likely are important for human-like intelligence. These moderators also are such that physiological processes (i.e., a body) are needed at some level of representation to tractably model how these effects combine.

We then provide an overview of an example model of the mind with a representation of physiology. After also discussing the importance of interacting with the world, we

conclude with a discussion and suggestions for a standard model that includes interaction with the environment.

Moderators and modulators

The human mind is implemented within a body. It follows that one may leverage this knowledge to build a standard model of the mind. This especially applies when one considers that several behavioral bands (Newell, 1990) can be mapped to physiological processes that run somewhat in parallel and interact with the behavior seen in these bands. These bands include biological, cognitive, rational, and social.

Including a model of these bands will help support modeling stress and affect and emotion in a more principle way. Previous work on modeling an architecture with an overlay of changes (Ritter, Reifers, et al., 2007) showed that this approach could be productively used to model how cognition changed for a single moderator. The aforementioned approach breaks down when multiple moderators have to be modeled and combined.

Stress

Human minds are affected by various stressors across the day causing changes in several physiological systems. These changes cause several pervasive effects on cognition, affecting our ability to retrieve memories and select actions given various contexts (Sara & Bouret, 2012; Schwabe & Wolf, 2013). The cognitive effects also appear to interact with time course of processing, operating over several time-scales (Joëls & Baram, 2009). In addition, many of the systems directly affected by stressors also show cycles or rhythms that modulate their levels across various timescales (Walker, Terry, & Lightman, 2010).

It follows that these changes to the processing mechanisms (or their physiological implementation) can be a form of adaptive response, affecting behavior in a way that allows a mind to accomplish its goals given a limited (and possibly threatened) set of resources. These insights of adaptation are not only important for a mind in a physiological body, but any mind with limited resources. The representation of a body (particularly a model of physiology) gives a straightforward, tractable way to explore insights for these adaptations over the various timescales that human minds can be examined at. This also holds true for understanding behavior that can be particularly maladaptive (e.g., a cycle of perseverations).

Affect and emotion

Though several theories of emotion and affect exist, models and frameworks that describe systems at the primal levels (e.g., LeDoux, 2012; Panksepp, 1998) can be particularly useful for describing various aspects of emotional behavior

from an architectural point of view. These models can be connected both to cognitive systems and physiological systems, representing a wide-range behavior across various timescales. Figure 1 gives a high-level picture of how these affect representations may integrate with other systems.

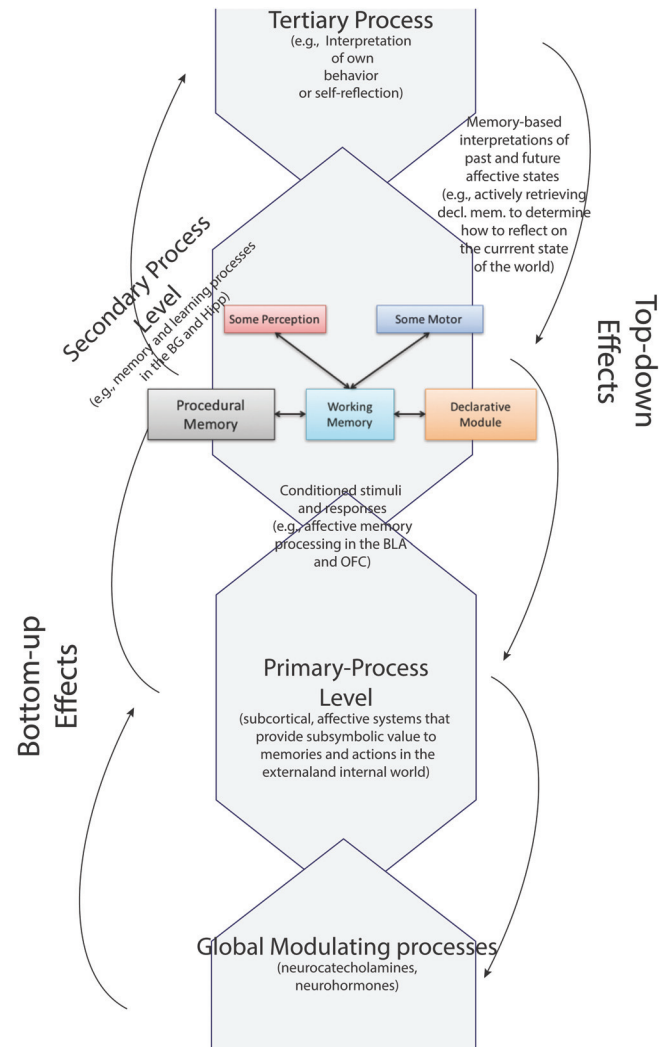


Figure 1. Levels of behavioral processes, modified from Panksepp et al. (2011).

In Figure 1, global modulating processes (e.g., those found in peripheral physiology, like cortisol) have bottom-up effects on a more primal-process level affect, which then feeds up to a more cognitive level. These processes also cause top-down effects on “lower” processes. Cognitive/secondary level systems can be modeled as interacting with these affective systems (e.g., see LeDoux, 2012; Phelps, 2006). This allows a fairly straightforward integration into existing research into models of the mind. Previous work has already explored connections between affect and cognitive systems from a computational perspective (e.g.,

Zhang et al., 2009), albeit mostly without an architecture perspective. Approaches like Zhang et al. (2009) that integrate within existing learning paradigms that are in many cognitive architectures (reinforcement-learning in the aforementioned case) can be particularly useful. This functional integration also creates a hierarchy of systems, something which can be useful for intelligent systems (Simon, 1996). The hierarchical interactions between these systems may provide a more general heuristic for the design of intelligent systems, human and human-like.

Physiology as a design heuristic

Though one may consider the more straightforward reasons for having a representation of physiology in a standard model, it is also useful to consider physiology for the use as a design heuristic of such a model. Similar to an argument made previously by Langley (2012) we may use physiology to limit the standard model. Indeed, some matches between (central) physiology and cognitive architectures have already been carried out that can be used to inform mechanism and model design (Anderson et al., 2008).

Nonetheless, many aspects of physiology remain limited in their representation of models of the mind. Epinephrine can surge while facing unexpected maladaptive situations, and variables associated with the HPA-axis (e.g., cortisol and corticotropin-releasing hormone) show circadian and ultradian rhythms (Walker et al., 2010). These forms of cycles across time modulate the way we think and behave and are often adaptive to the environment (e.g., the effect of sunlight on circadian rhythms, Leproult et al., 2001). They serve as another form of adaptation, and can be seen as another form of memory of the environment. Though many of these physiological processes are evolutionarily old, they serve as a useful design heuristic for a standard model as the mind grew out of these adaptations over time.

We have explored both of these ideas of extending the cognitive architectures, by both including a model of a body and how it supports and interacts with the mind. We also have extended interaction directly to an environment, attempting to create a cognitive architecture level of reusable capabilities. To explore this approach, we have created an example model of the mind with a body.

An example model with physiology

We have connected a model of physiology, HumMod (Hester et al., 2011), to a model of cognition, ACT-R (Anderson, 2007), resulting in the ACT-R/ Φ architecture (Dancy, 2013). This architecture uses a theoretical framework to tie the cognitive and physiology together from work in emotion and affective neuroscience (Panksepp, 1998; Panksepp & Biven, 2012), which acts as a functional layer between the

two systems and guides connections between physiological and cognitive processes. ACT-R/ Φ is more so a functional extension of ACT-R and thus focuses on the ability to represent moderators and modulators in a coherent manner rather than using physiology as a design heuristic. Figure 2 gives a high-level picture of ACT-R/ Φ .

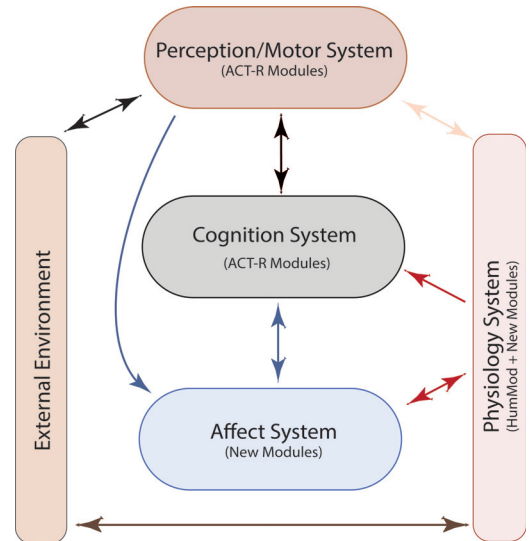


Figure 2. A high-level schematic of the ACT-R/ Φ architecture.

ACT-R/ Φ extends the ACT-R architecture with an Affect and Physiology system. The Physiology system is composed mainly of the HumMod physiological model and simulation system, and a module. The Physio module keeps the cognitive and physiological systems synced in time, and contains functionality to affect subsymbolic quantities in the cognitive system directly.

The Affect system implements two systems from primal affect theory (which broadly represent an appetitive motivational system and defense system) and a central memory system that associates affective state with memory elements. When combined with the physiological system, the affect system can represent interesting results like changing utilities and rewards due to homeostatic physiological change (e.g., see Dancy & Kaulakis, 2013). The architecture also includes other representations to match human-behavior (e.g., Dancy et al., 2015, discusses another set of connections showing how the two systems interact with each other), but these are out of the current scope.

Though the ACT-R/ Φ architecture does not implement all of the many likely affective systems that modulate cognitive (and physiological) systems, it provides a good start and is the first of such architecture to attempt to bring together these aspects. These representations were built in separate disciplines, but nonetheless provide lessons and can be used to come closer to a cohesive standard model of the mind. The system does not yet have the particular advantage of the

forementioned design heuristic, but it is functionally useful and could be used for this purpose in the future.

Interaction abilities

Models within an architecture need a way to interact with the world. This includes senses and motor output. Agents within the standard model deserve to be able to interact with the world (Ritter et al., 2000). Current architectures (e.g., Soar, ACT-R, EPIC) have capabilities in the architecture for interaction noted as modules, but often these modules for interaction are fairly thinly implemented, somewhat abstract, and rather rarely used in their entirety. For example, in the case of ACT-R, while it has a capability for interaction, modelers often revert to using rules directly to interact with a more local, limited, abstracted simulation. In this case, set inputs are placed into memory with the associated aspects of active vision (Findlay & Gilchrist, 2003) including the time and knowledge needed to look around and understand an object, and where actions are merely reported or their impacts are realized rather than performed and checked (we have been guilty of this ourselves on numerous occasions, because the interactions would increase the programming effort required from 25% to perhaps 300% more work).

We would like to suggest that architectures should include a more complete model of interaction and that computer screens are a minimum first step. St. Amant first put forward that computer screens have some particular value because they are regular, ubiquitous, and available to most models (St. Amant & Riedl, 2001).

Thus, we suggest that the standard model include simple interaction capabilities like those provided by SegMan (Ritter, Kukreja, & St. Amant, 2007; St. Amant, Horton, & Ritter, 2007; St. Amant & Riedl, 2001), to read letters and words, dialogue items and colors, and icons, and to move the mouse and click and to type. The features of SegMan worth carrying forward are that it is relatively architecturally neutral (it has been used by Soar, ACT-R, and agent models), that it provides interaction with any task that can be put onto a screen, and that it can be extended (and will need to be).

Conclusion

Models of human minds will find including a model of the human body increasingly necessary. It will particularly be necessary to model how behavioral moderators interact. Simple changes to the high-level architecture can model single moderators, but do not provide a principled way to combine moderators.

A body is also required to model interaction with the world. We are not arguing for robots per se, but arguing for

modeling interaction on a detailed level at least with computer screens.

General AI systems that have any realization as hardware, such as a robot or an agent on a power-limited computer, will also have to be tied to their physical realization. That is, they will not only have to use the vision and motor modules released as hardware, but they will be required to monitor the body lest they, to include a recent robot example in the news, fall into a lake and drown, or more likely, fail to modify their behavior in time to get recharged. The complex processes interacting within this system will, similar to the human body, provide another representation of memory and adaptation that interacts with the normally represented architectural systems.

Including physiology and direct interaction mechanisms will both help constrain and validate the standard model. These additions may prove key to understanding how the human mind can occur in the physical universe and, indeed, how an intelligent mind can occur within any universe.

References

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: OUP.
- Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. *Trends in Cognitive Sciences*, 12(4), 136-143.
- Dancy, C. L. (2013). ACT-RΦ: A cognitive architecture with physiology and affect. *Biologically Inspired Cognitive Architectures*, 6(1), 40-45.
- Dancy, C. L., & Kaulakis, R. (2013). Towards adding bottom-up homeostatic affect to ACT-R. *In proceedings of the 12th International Conference on Cognitive Modeling*, Ottawa, Canada, 316-321.
- Dancy, C. L., Ritter, F. E., Berry, K. A., & Klein, L. C. (2015). Using a cognitive architecture with a physiological substrate to represent effects of a psychological stressor on cognition. *Computational and Mathematical Organization Theory*, 21(1), 90-114.
- Findlay, J. M., & Gilchrist, I. D. (2003). *Active vision: The psychology of looking and seeing*. Oxford, UK: OUP.
- Hester, R. L., Brown, A. J., Husband, L., Iliescu, R., Pruett, D., Summers, R., & Coleman, T. G. (2011). HumMod: A modeling environment for the simulation of integrative human physiology. *Frontiers in physiology*, 2(12).
- Joëls, M., & Baram, T. Z. (2009). The neuro-symphony of stress. *Nature Review in Neuroscience*, 10(6), 459-466.
- Langley, P. (2012). Intelligent behavior in humans and machines. *Advances in Cognitive Systems*(2), 3-12.
- LeDoux, J. E. (2012). Rethinking the emotional brain. *Neuron*, 73(4), 653-676.
- Leproult, R., Colecchia, E. F., L'Hermite-Balériaux, M., & Van Cauter, E. (2001). Transition from dim to bright light in the morning induces an immediate elevation of cortisol levels. *The Journal of Clinical Endocrinology & Metabolism*, 86(1), 151-157.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, Massachusetts: Harvard University Press.

- Panksepp, J. (1998). *Affective neuroscience: The foundations of human and animal emotions*. New York, NY: OUP.
- Panksepp, J., & Biven, L. (2012). *The Archeology of Mind: Neuroevolutionary Origins of Human Emotions*. New York, NY: W.W. Norton & Company.
- Phelps, E. A. (2006). Emotion and cognition: Insights from studies of the human amygdala. *Annual Review of Psychology*, 57, 27-53.
- Ritter, F. E., Baxter, G. D., Jones, G., & Young, R. M. (2000). Supporting cognitive models as users. *ACM Transactions on Computer-Human Interaction*, 7(2), 141-173.
- Ritter, F. E., Kukreja, U., & St. Amant, R. (2007). Including a model of visual processing with a cognitive architecture to model a simple teleoperation task. *Journal of Cognitive Engineering and Decision Making*, 1(2), 121-147.
- Ritter, F. E., Reifers, A. L., Klein, L. C., & Schoelles, M. J. (2007). Lessons from defining theories of stress for cognitive architectures. In W. D. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 254-262). New York, NY: OUP.
- Sara, S. J., & Bouret, S. (2012). Orienting and reorienting: The locus coeruleus mediates cognition through arousal. *Neuron*, 76(1), 130-141.
- Schwabe, L., & Wolf, O. T. (2013). Stress and multiple memory systems: From 'thinking' to 'doing'. *Trends in Cognitive Sciences*, 17(2), 60-68.
- Simon, H. A. (1996). *The science of the artificial* (3rd ed.). Cambridge, MA: MIT press.
- St. Amant, R., Horton, T. E., & Ritter, F. E. (2007). Model-based evaluation of expert cell phone menu interaction. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 14(1), 1-24.
- St. Amant, R., & Riedl, M. O. (2001). A perception/action substrate for cognitive modeling in HCI. *International Journal of Human-Computer Studies*, 55(1), 15-39.
- Walker, J. J., Terry, J. R., & Lightman, S. L. (2010). Origin of ultradian pulsatility in the hypothalamic-pituitary-adrenal axis. *Proceedings of the Royal Society B: Biological Sciences*, 277(1688), 1627-1633.
- 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.