A Standard Model of Mind:
AAAI Technical Report FS-17-05

Proposal to Add Emotion to the Standard Model

Robert L. West
Institute of Cognitive Science, Carleton University, Ottawa, Canada
robert.west@cunet.carleton.ca

Jeremy T. Young
Institute of Cognitive Science, Carleton University, Ottawa, Canada
jeremyyounggrad@gmail.com

Abstract

This paper outlines a proposal to add emotion to the Standard Model. Part 1 describes how emotional factors impact decision-making and Part 2 outlines a way to augment the standard model to include these factors.

Introduction

The standard model is a model of deliberate action, or choice, occurring in the 50 to 100 millisecond range of human cognition. To incorporate emotion into the standard model we need an understanding of how emotion participates in this type of choice. However, there is no agreed upon standard for emotion. Instead, there is a lot of diverse research and differing frameworks for interpreting the results. This has led to a diversity of computational models of emotion (see Lin et al, 2012, for a review of cognitive models incorporating emotion). The position taken in this paper is to look for major points of intersection between emotion and deliberate actions, and to interpret them in terms of the standard model. That is, to develop a theory of emotion *as it relates to the standard model*.

Part 1 of this paper outlines the major points of intersection and how they can be interpreted in terms of the standard model. Part 2 describes how these areas of intersection can be formally incorporated into the standard model. This research is based on modelling emotion in ACT-R so the discussion will be mainly informed by ACT-R, but it is applicable to the standard model. Neurologically, the focus will be on the amygdala as the nexus for emotional participation in symbol driven choice. However, no strong neurological claims are made, except for the fact that the amygdala is a reliable neural correlate for this type of processing.

Part 1: Points of Interaction

Neural and Symbolic

Anderson's concept of symbolic and sub symbolic structures is particularly important for understanding the boundary between the neural and cognitive levels (see Anderson and Lebiere 1998). Symbolic structures in ACT-R, such as chunks and productions, are composed of symbolic information combined with associated magnitude information related to sub symbolic processing. For example, productions are associated with a utility value and chunks are associated with an activation value. The division between the symbolic and sub symbolic is strict, symbols cannot access their own sub symbolic values. This strict divide makes sense if sub symbolic values represent the neural process underlying the symbolic representation.

Emotional processes at the neural level can be described through their effect on sub symbolic algorithms and values. One way of architecturally conceptualizing this is in terms of an overlay, which adjusts sub symbolic processing to account for factors such as fatigue or stress (see Ritter et al. 2006) for review and discussion). This is a good approach for understanding how state wide emotional factors influence the performance of tasks, but it is limited in terms of accounting for the role of emotion in symbol based decision making.

Involving emotions in deliberative decision making minimally requires (1) that the symbol system can access a representation of the agent's emotional state, and (2) that the emotional system can understand the symbolic representations of the current state in working memory.

To access state wide effects modelled with overlays (e.g., stress, fatigue) the symbol system can infer the emotional state by observing its own reactions and behaviors (as in the Two Factor theory of emotion, Schachter and Singer 1962). For example, this method has been used in ACT-R to model how people estimate their confidence in a judgement by equating longer decision times with lack of confidence (Drewitz and Thuring 2009). This mechanism would allow a limited type of emotional decision making (e.g., I am stressed so I choose to quit).

However, it is also important to account for the fact that emotions occur in response to symbolic information, for example, numerical rewards or a letter informing you of a tax audit, etc. Corresponding to this, there is evidence that the amygdala responds to symbolic information (e.g. Jenison et al. 2011). There is also evidence that the amygdala can communicate alarm judgements to the symbolic system (e.g. Fox et al. 2015), that it is sensitive to reward amounts (e.g. Gospic et al. 2014), and that it is capable of learning (e.g. Damasio 1994). This is different from what can be modelled with an overlay. To incorporate this type of emotional activity, a separate module is needed and the relative contributions of symbolic and sub symbolic processing need to be considered.

Rewards and Punishments

Problem space search is an important component of the standard model. In ACT-R, Time Delayed (TD) learning is used to learn pathways through problem spaces by adjusting the utility of productions in procedural memory (see also Kelly and West 2013, for a holographic mechanism). However, regardless of the mechanism, the amount of the rewards must be set by the modeler. Having different award amounts reflects the obvious fact that some rewards are higher than others. Also, humans can value the same rewards differently, even when they are numerically identical (West and Ward 1998). Research shows that the mechanism for this valuation of rewards involves the amygdala (e.g. Gospic et al. 2014). However, to model this we need to know what the evaluation is based on. This is covered below in the section in values.

Somatic Markers

Damasio, in his book, *Descartes Error* (1994), elucidated the relationship between emotion and rational thought by showing that emotional learning plays a role in rational decision making. He did this by demonstrating that people with brain damage blocking communication between the amygdala and frontal cortex could not learn the Iowa Gambling Task (IGT), although they could learn other types of rational tasks (Damasio 1994). The IGT involves selecting cards from four decks to reveal rewards and punishments. Two of the decks are "bad" decks that will lead to loosing in the long term, and two of the decks are "good" decks that will lead to winning in the long term. Initially, people tend to choose the bad decks but over time they slowly switch to the good decks, but people with damage to the connections between the amygdala and the frontal cortex seem unable to learn this. Damasio's (1994) interpretation of this result was that rational thinking in the frontal cortex initially leads people to perform poorly in the IGT, but this is overcome by emotional learning occurring in the amygdala.

Damasio explains the IGT results using his theory of somatic markers. Somatic markers are emotional tags attached to the representation of an object or person, indicating overall desirability. In theory, somatic markers are blocked by damage to the connections between the amygdala and the frontal cortex. Damasio has insisted that emotional learning and somatic markers do not involve the use of symbols. However, this is widely disputed as it has been shown that IGT players are aware of the knowledge that they need to win (Maia and McClellan 2004).

ACT-R has already been successfully used to model somatic markers and paradigms related to the IGT. In fact, there are at several different ways to get ACT-R to mimic this type of learning. Lebiere (1999) created blending, which can operate in ACT-R or independently (Gonzalaz et al. 2003), and Rutledge-Taylor et al. (2004, 2015) proposed varying the number of memory saves to reflect the amount of reward. Similar to the proposal in this paper, Stocco et al (2005) and Juvina et al (2017) have proposed adding a module to ACT-R to model these effects. All of these operate by somehow encoding reward amounts into an instance based memory model.

Alarm

In addition to being involved with somatic markers, there is wide agreement that the amygdala plays an important role in our ability to detect threats (Andrew et al. 2015). This is important because threat detection, arguably, represents a computational shortcoming in the standard model. Both ACT-R and the standard model have a serial bottleneck in the action cycle. In ACT-R, the bottleneck is produced because only one production can fire at a time. Computationally, the bottleneck is important because it functions as a control system, but it also limits flexibility in the face of unexpected interruptions.

This is best illustrated with an example. Consider a technician executing a series of well-practiced actions to fix an elevator. The series of actions would be executed by a chain of productions, each one changing the buffer conditions so that the right production will fire next. If the fire alarm rings, this series of productions will continue to fire unless there is a higher utility production that fires if the fire alarm is ringing. This approach assumes that our ability to be interrupted by perceived threat is based on a set of high utility productions in procedural memory. However, there are several reasons why this approach is problematic.

Humans are sophisticated in their response to alarms. Consider the elevator technician's response to the alarm. If they are in the middle of an important operation they might finish it first if the consequence of not doing so is serious (e.g., the elevator falls). Or, they might know that the fire alarm is not working properly and choose to ignore it. Alternatively, their knowledge that there is a large amount of dynamite stored in the building could cause them to drop everything and run.

Humans are able to intelligently modify their response based on the alarm type and the context. Humans are also able to resume their task later without difficulty. High utility productions can fire at any time, which means that the buffer states can be altered at any point. To produce human like abilities under these conditions would require a system for buffering the response, understanding the context, choosing the right response, and remembering where the interruption occurred so the task can be resumed. Putting all of this through the procedural memory bottleneck would cause relatively large time delays compared to human responses, which can be quite fast in experts.

The standard model is not committed to the ACT-R production utility mechanism so this problem could potentially be dealt with by using a more complex production selection mechanism. However, the neural evidence indicates that the human brain solves this problem by using the amygdala as a parallel threat-screening mechanism. As proposed below, this provides a simple, flexible solution to this problem for the standard model.

Emotional Knowledge

Emotional Intelligence (EQ) is a collection of different emotional skills, many of which can be explicitly learned. Likewise, Cognitive Behavioral Therapy (CBT) has demonstrated that emotional disorders such as depression, anxiety, and anger management, can be effectively treated by teaching people to replace poor emotional skills with more effective emotional skills. What this tells us is that knowledge, in terms of cognitive routines, plays a key role in how emotions are processed. This type of knowledge can be modelled directly in the standard model using procedural and declarative memory, but there needs to be a source of emotional information.

Values

The standard model is driven by goals. Each time the system moves to a new state in the task the working memory representation is updated by the procedural memory system. Thus a goal hierarchy is implicit in the procedural memory representations that alter working memory. More globally, the procedural and declarative representations in the model can be considered to embody a rational analysis of how to achieve the goal. However, as the pragmatists have pointed out, a rational analysis is only possible in the context of knowing the goal (James 1907). Without a predetermined goal we run into the naturalistic fallacy (Moore 1903), which states that we cannot rationally go from what is, to what we ought to do. The consequence of this is that you cannot use rationality to choose the top goal for a model.

Top goals are ultimately based on values. Rationality can be used to create better medicine or better weapons, or it can be completely ignored. It depends on one's values. In the standard model values are arguably connected to reward amounts, the more a goal is valued, the higher the reward when it's achieved. To understand the relationship between values, goals, and rationality consider Asimov's three laws of robotics:

> A robot may not injure a human being or, through inaction, allow a human being to come to harm. A robot must obey orders given it by human beings except where such orders would conflict with the First Law. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

Given a task, the robot uses rational analysis to create a goal structure that will not violate these three laws. The laws function as the values of the robot. Although the values can be treated as high level goals, they were not created by rational analysis. Instead the values were placed there by humans. Likewise, in humans, a fixed set of values is often provided by society. However, unlike Asimov's robots, humans can change their values.

The top diagram in Figure 1 illustrates how the standard model currently works. Goals produce results and the results feedback to update the goals. The reward amounts as well as the choice of goals are hand coded by the programmers. Once the goals and associated rewards are set, learning algorithms can be used to work out, or learn, efficient pathways through the problem space to the goal. The bottom diagram shows how the standard model would work if it had values; the agent uses its values to set its own goals and rewards, and updates its values based on feedback.

Values don't appear out of nowhere, you need a starting point. Initial values are provided to us by society and possibly by evolution. The process of selecting goals (i.e., thinking about what one ought to do) has a strong knowledge component. The process could involve quite extensive deliberation or it could be heuristically based. This distinction can be mapped to Kahneman's (2011) fast and slow thinking theory. In either case, the standard model could process it (Thomson et al. 2015). Triggering goal selection can be accomplished by a high utility production of the type - if a top goal has been achieved, set a goal to choose the next goal based on the context. Values represent the values of goals, and can change due to feedback from results.

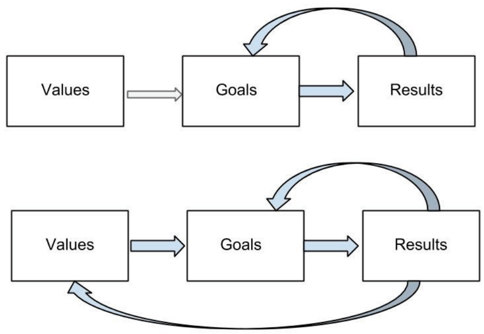

Figure 1. Values are needed to choose goals.

Neural imaging indicates that values are processed through the amygdala (Gospic et al. 2014). There is evidence (Kanai 2011) that different value sets (conservatism versus liberal) are correlated with relative size differences in the amygdala and the anterior cingulate cortex, which is associated with control and executive processes. Conservatives have a relatively larger right amygdala and liberals have a relatively larger anterior cingulate cortex. In terms of the model described in this paper, this would indicate that processing conservative values depends more on information stored or processed through the amygdala, while processing liberal values relies more on top down executive control. The standard model should be able to model this difference.

Part 2: An Emotional Architecture

Our proposal for adding emotion to the standard model is based on defining how emotion participates in the creation of deliberate actions at the cognitive level. It is not a neural model nor does it involve speculation on the fundamental structures or representations underlying emotion. However, it is broadly consistent with research on these topics.

The standard model derives its rational intelligence from the procedural and declarative memory systems. The procedural system is based on choosing what to do next given the current context. The declarative system provides further information relevant to the context, when requested by the procedural system. The goals in designing the emotional system were (1) to incorporate it into the standard model in the simplest possible way, (2) deliver the functionality described in Part 1, and (3) not to break existing, non emotional models. This system has been partially constructed and tested in ACT-R but will be described here in more general terms, suitable to the standard model.

The Emotional Module

The proposal for an emotional module is based on research on the amygdala. However, the proposed emotional module is not meant as a model of the amygdala. Rather, the claim is that the amygdala functions as the nexus of a control structure for mediating the relationship between emotions and rational, symbol-based decision-making. Other brain areas are certainly involved in the sub symbolic processes related to the emotional module (e.g., see Juvina et al, 2017).

Computationally, the emotion module is similar to the procedural module, in that it reacts to the contents of the working memory buffers. However, rather than producing actions, the emotional module produces emotional evaluations, both symbolic and sub symbolic. The buffer structure is shown in figure 2. The emotional module reacts to the same buffers as the procedural module but it is limited in its responses to a single buffer, called the emotional buffer. This limitation prevents the emotion module from changing the buffer conditions unexpectedly, in a way that could confuse the procedural module, which relies on stable buffer conditions for chaining actions together. At the same time, though, the contents of the emotional buffer are available to the procedural module and can alter its behavior. An example is provided below.

In terms of its interactions with other modules, the emotion module produces a parallel emotional evaluation of the current buffer contents that is used by both the procedural and declarative modules. This is illustrated in Figure 3. The input to the declarative module is for managing somatic markers. One expedient way to do this is to create a chunk representing all of the buffer conditions and store it in declarative memory along with the emotional evaluation. As noted above there are a number of ways to implement somatic markers, anyone of them could work.

In terms of the procedural module the evaluation from the emotional module is used as the amount of reward or punishment for supervised learning. Feed back to the production system already exists in a number of architectures that fit the standard model, however, the reward amounts must be specified. Using the emotional module to set the reward amounts extends the standard model to cover models of how we derive our internal rewards.

Emotional Knowledge

For this system to work there needs to be a fair amount of emotional knowledge in procedural memory. Although this is not what ACT-R is typically programmed with, research on EQ skills and on CBT clearly shows that people have different levels of expertise in dealing with their emotions and the emotions of others. To research this we have been looking at how this type of expert knowledge interacts with more traditional forms of expert knowledge.

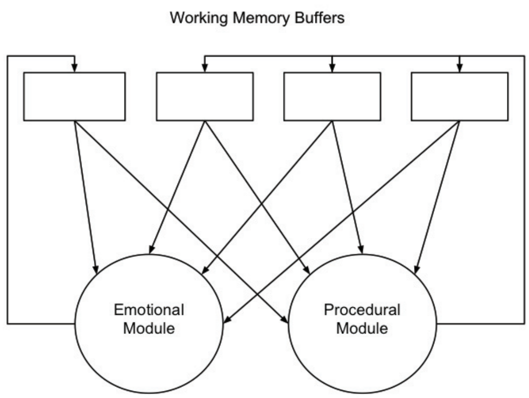

Figure 2. Emotion module buffer configuration.

If each different area of expertise (e.g., nurse, parent, hockey coach, martial artist, camper, video game player) is organized differently then each one requires a customized set of emotional productions to interact with them. While this is theoretically possible, it seems unlikely and inefficient. Instead, we assume that all expert knowledge is organized according to SGOMS (West et al. 2007, 2009, 2015), which is a version of GOMS (Card, Moran, and Newell 1983). Thus each activity is organized according to operators, methods, unit tasks, and, in the case of SGOMS, planning units. By assuming this as a generic structure we can make use of generic emotional knowledge. Imagine that someone is afraid of spiders. This fear will not change during different activities, but the response could. For example, while camping one might immediately attempt to deal with a spider but during an emergency hospital procedure one might ignore it. This can be modelled as follows - in both cases the amygdala recognizes the spider and places a medium threat chunk in the emotion buffer. In the camping case this is enough to trigger an interruption to the roast marshmallow planning unit, but in the medical case it is not enough to interrupt the evaluate emergency room patient planning unit.

As part of the SGOMS model, each operator, method, unit task, and planning unit contains expert knowledge of how important it is and thus what level of threat is needed to interrupt it. For example, a nurse might complete the unit task for measuring blood pressure but then briefly interrupt the planning unit for processing the patient to stomp on the spider. However, if a lion charged through the door the nurse would likely drop everything and try to escape. The important thing to note is that the expert system for nursing does not need to know anything about spiders or lions, that knowledge is processed in parallel in the amygdala and communicated as a threat level.

We have found that this approach allows ACT-R to model the sort of quick but intelligent interruption abilities displayed by human experts. We have used it to model interruptions in first person shooter video game play and expert mediation in training scenarios with actors (MacDougal et al. 2014). This approach separates task expertise and threat evaluation into parallel threads, which is arguably more efficient and scalable.

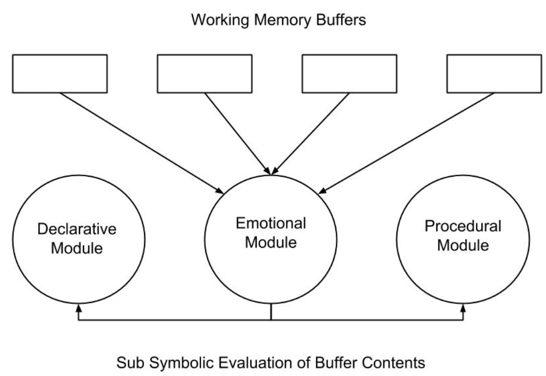

Figure 3. The same sub symbolic values from the emotion module are used in the procedural and declarative modules.

Implementation

To implement the emotional module, we created it as a second production system in Python ACT-R. However, this is not a parallel production system. The emotional module cannot send instructions to other modules and it can only alter the contents of the emotional buffer. It is more of a shadow production system that offers emotional commentary on the ongoing activity. The serial production bottleneck still controls activity. Implementing it in this way allows the emotional module to react to the buffer conditions, and it can model emotional learning through production utility. We are currently using this system to model how individuals change, or fail to change their values when exposed to different information.

Conclusion

The emotional module provides a simple, direct method for modelling a variety of emotional effects that is consistent with research on the amygdala and the architectural assumptions of the standard model. Moreover, having a module devoted to emotional processing creates a clear, unambiguous way of exploring different ways to model the effects of emotion as they apply to the standard model.

References

Anderson, J. R. & Lebiere, C. 1998. *The Atomic Components of Thought.* Mahwah, NJ: Erlbaum.

Fox, A. S., Oler, J, A., Tromp, D. P. M., Fudge, J. L., & Kalin, N. H. 2015. Extending the amygdala in theories of threat processing. *Trends Neuroscience*. 2015 May; 38(5): 319–329.

Card, Stuart; Thomas P. Moran; Allen Newell 1983. *The Psychology of Human Computer Interaction*. Lawrence Erlbaum Associates.

Damasio, A. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain,* Putnam Publishing.

Drewitz, U. & Thüring, M. 2009. Modeling the Confidence of Predictions: A Time Based Approach. In Proceedings of the 9th International Conference of Cognitive Modeling (paper 222), Manchester, United Kingdom.

Gonzalez, C., Lerch, J. F., & Lebiere, C. 2003. Instance-based Learning in Dynamic Decision Making. *Cognitive Science*, 27(4), 591-635.

Gospic, Sundberg, Maeder, Fransson, Petrovic, Isacsson, Karlström, Ingvar. 2014. Altruism Costs-the Cheap Signal from Amygdala. Soc Cogn Affect Neurosci. 2014 Sep;9(9):1325-32.

James, William. 1907. *Pragmatism and Other Writings*. New York, NY: Penguin Books. (republished in 2000.)

Jenison, R. L., Rangel, A Oya, H., Kawasaki, H., and Howard, M. A. 2011. Value Encoding in Single Neurons in the Human Amygdala during Decision-making. J *Neuroscience.* 2011 Jan 5; 31(1): 331–338.

Juvina, I., Larue, O., & Hough, A. 2017. Modeling valuation and core affct in in a cognitive architecture: the impact of arousal and valence on memory and decision-making. Cognitive Systems Research.

Kahneman, D. 2011. *Thinking Fast and Slow.* Random House of Canada.

Kelly, M, & West, R. L. 2013. Decision Making in a Dynamically Structured Holographic Memory Model: Learning from Delayed Feedback. Proceedings of the International Conference on Cognitive Modelling.

Lebiere, C. 1999. Blending: An ACT-R Mechanism for Aggregate Retrievals. Sixth Annual ACT-R Workshop at George Mason.

Lin, J., Spraragen, M., & Zyda, M. 2012. Computational Models of Emotion and Cognition. Advances in Cognitive Systems 2 (2012) 59-76.

MacDougall, W. K., West, R. L., & Hancock, E. 2014. Modeling Multi-Agent Chaos: Killing Aliens and Managing Difficult People. Proceedings of Cognitive Science.

Maia, T., & McClelland, J. 2004. A Reexamination of the Evidence for the Somatic Marker Hypothesis: What Participants Really Know in the Iowa Gambling Task. *National Academy of Science of the USA*, 101, 16075-80.

Moore, George Edward 1903. *Principia Ethica*. Cambridge: Cambridge University Press.

Ritter, F. E., Reifers, A. L., Klein, A. C., & Schoelles, M. J. 2006. Lessons from Defining Theories of Stress. In W. Gray (Ed.) *Integrated Models of Cognitive Systems*. New York, NY: Oxford University Press.

Rutledge-Taylor M., & West, R. L., 2004. Cognitive Modeling Versus Game Theory: Why Cognition Matters. The proceedings of the 7th International Conference on Cognitive Modeling. Pittsburgh.

Rutledge-Taylor, M. F., Kelly M. A., West, R. L., & Pyke, A. A. 2015. Dynamically Structured Holographic Memory. *Biologically Inspired Cognitive Architectures*.

Kanai, R. Feilden, T., Firth, C., & Rees, G. 2011. Political Orientations Are Correlated with Brain Structure in Young Adults. Curr Biol. 2011 Apr 26; 21(8): 677–680.

Schachter, S.; Singer, J. 1962. Cognitive, Social, and Physiological Determinants of Emotional State. *Psychological Review*. 69: 379–399.

Stocco, A., Fum, D., & Zalla, T. 2005. Revising the Role of Somatic Markers in the Gambling Task: A Computational Account for Neuropsychological Impairments. In B. Bara, L. Barsalou & M. Bucciarelli (Eds.), Proceedings of the 27th Annual Conference of the Cognitive Science Society. Mahwah, NJ: Lawrence Erlbaum Associates.

Thomson, R., Lebiere, C., Anderson, J. R., & Staszewski, J. 2015. A General Instance-Based Learning Framework for Studying Intuitive Decision Making in a Cognitive Architecture. *Journal of Applied Research in Memory and Cognition*; *Special Issue on Modeling and Aiding Intuitions in Organizational Decision Making*.

West, R. L., & Ward, L. M. 1998. The Value of Money: Constrained Scaling and Individual Differences. Fechner Day: Proceedings of the Fourteenth Annual Meeting of the International Society for Psychophysics. Quebec, Canada: International Society for Psychophysics.

West, R. L., & MacDougal, K. 2015. The Macro Architecture Hypothesis: Modifying Newell's System Levels to Include Macro Cognition. *Biologically Inspired Cognitive Architectures*.

West, R. L., & Pronovost, S. 2009. Modeling SGOMS in ACT-R: Linking Macro- and Microcognition. *Journal of Cognitive Engineering and Decision Making*, 3 (2), 194-207.

West, R. L., & Nagy, G. 2007. Using GOMS for Modeling Routine Tasks Within Complex Sociotechnical Systems:Connecting Macrocognitive Models to Microcognition. *Journal of Cognitive Engineering and Decision Making*.