

Artificial Motivation for Cognitive Software Agents

Ryan J. McCall

MCCALL.RYAN@GMAIL.COM

*Computer Science Department and Institute for
Intelligent Systems, University of Memphis,
Memphis, TN 38152 USA*

Stan Franklin

FRANKLIN@MEMPHIS.EDU

*Institute for Intelligent Systems
University of Memphis
365 Innovation Drive
Memphis, TN 38152, USA*

Usef Faghihi

JFAGHIHI@YAHOO.COM

*University of Indianapolis
1400 E Hanna Ave
Indianapolis, IN 46227 USA*

Javier Snaider

JAVIERSNA@GMAIL.COM

*Computer Science Department and Institute for
Intelligent Systems, University of Memphis,
Memphis, TN 38152 USA*

Sean Kugele

SKUGELE@MEMPHIS.EDU

*Computer Science Department and Institute for
Intelligent Systems, University of Memphis,
Memphis, TN 38152 USA*

Editors: Joscha Bach, Eva Hudlicka, Stacy Marsella

Abstract

Natural selection has imbued biological agents with motivations moving them to act for survival and reproduction, as well as to learn so as to support both. Artificial agents also require motivations to act in a goal-directed manner and to learn appropriately into various memories. Here we present a biologically inspired motivation system, based on feelings (including emotions) integrated within the LIDA cognitive architecture at a fundamental level. This motivational system, operating within LIDA's cognitive cycle, provides a repertoire of motivational capacities operating over a range of time scales of increasing complexity. These include alarms, appraisal mechanisms, appetite and aversion, and deliberation and planning.

Keywords: Artificial Motivation, Emotion, LIDA

1. Introduction

Cognitive architectures model minds, which are thought of as control structures for autonomous agents (Franklin, 1995, p. 412). An autonomous agent is a system situated within, and a part of, an environment that senses that environment and acts on it, over time, in pursuit of its own agenda, and so as to affect what it may sense in the future (Franklin & Graesser, 1997). Given this definition, how can the agenda of an autonomous agent be implemented in a cognitive architecture? How does such agenda arise from basic cognitive processes? In other words, what motivates an agent to choose one action over another, or to choose any action at all? What is its motivation to act? How is this motivation to act accomplished? These are the foundational questions we address in this work on artificial motivations.

Motivation is a broad term guiding everything from simple reflex actions to decade-long career commitments. We discuss an assortment of motivational processes covering a range of temporal scales, and having varying complexity. We focus primarily on motivation based on feelings. Several researchers have addressed this topic to date (Bach, 2009; Canamero, 2003; Franklin & Ramamurthy, 2006; Gmytrasiewicz & Lisetti, 2002; Lang & Davis, 2006; Sloman & Croucher, 1981). Here we present a motivational system, based on feelings, and integrated within the broad systems-level LIDA cognitive architecture (Franklin et al., 2016) at a fundamental level; this is an approach to motivation in the entirety of cognition. It is biologically inspired but does not aim to model the neuroscience side of human or animal motivation *per se*. However, we hold that only by building upwards from solid conceptual primitives can we provide precise systems-level explanations of such complicated phenomena. Our goal is to be able to do so.

James and Lange pioneered early emotions research (as summarized in Cannon, 1927). More recently, Picard's seminal work on "affective computing" (1997; 2003) motivates much current research concerned with artificial emotions in computer systems. Fellous (2004) suggests focusing on the functions of emotions rather than what they are. He argues that one of the main functions of emotions is to achieve "the multi-level communication of simplified but high impact information," and that it may be crucial to understand emotions as "dynamical patterns of neuromodulations, rather than patterns of neural activity." Given that LIDA's perceptual nodes, including those for "emotions" (see Section 3.6) are hypothesized to correspond to particular basins of attraction in the dynamics of neuron populations (Franklin, Strain, Snaider, McCall, & Faghihi, 2012), feeling nodes in LIDA (see Section 3.2) can be viewed as being close to Fellous' vision. Cañamero (1997) produced an autonomous agent with basic drives and an action selection mechanism modulated by its "emotions." Almost all cognitive architectures have motivation mechanism(s). For instance, SOAR (Laird, J, Newell, & Rosenbloom, 1987) has explicit goals, and an impasse mechanism that creates sub-goals for complex tasks. More recently emotions have been integrated into SOAR (Marinier, Laird, & Lewis, 2009), and used to drive reinforcement learning (Marinier & Laird, 2008), but not as motivators. The role of emotions in cognition has also been studied in ACT-R, including its affect on rationality (Fum & Stocco, 2004), decision-making (Belavkin, 2001b), and the "inverted U" effect (Belavkin, 2001a).

Some cognitive architectures have sophisticated motivational mechanisms. For example, CLARION uses a bipartite motivational representation where "explicit goals", e.g. finding food, are generated based on past and current "internal drive states", e.g. being hungry (Sun, 2009). Another architecture with a rich motivational mechanism is Psi (Dorner & Hille, 1995) and its progeny Micro-Psi (Bach, 2009). Psi distinguishes between demands, urges, motives, cognitive modulators, affect, and directed emotions. The Psi model provides inspiration for some of the concepts in the expanded motivation model we are here adding to LIDA.

To be clear, the LIDA cognitive architecture consists of the conceptual LIDA model and one or more computational LIDA frameworks, which individually facilitate the development and computational implementation of various LIDA-based agents. At any point in time, the current LIDA model can be found in the most recent summary paper (for example, Franklin, et al., 2016), together with more specific details in earlier publications, and with conceptual additions from newer papers such as this one. A LIDA framework is a computational system that expedites the development of various LIDA-based agents, by providing simplified beginning computational versions of most of the LIDA modules. Currently, there is one operational LIDA framework (Snaider, McCall, & Franklin, 2011). LIDA-based agents come in two varieties, those that perform real world tasks (Franklin, Kelemen, & McCauley, 1998), and those that simulate psychological experiments. All LIDA-based agents must include basic versions of critically important LIDA modules. However, not every LIDA module in the conceptual model need be included in every LIDA-based agent. But, every part of the conceptual LIDA model is computationally implementable in theory.

The organization of the rest of the paper is as follows. We present a brief overview of the LIDA model (Section 2). Next, we discuss a series of concepts critical to describing the role of motivation within LIDA in preparation for the expanded motivational system to be added here. (Section 3). Section 4 describes the fundamental expanded motivational mechanisms in the LIDA model. Section 5 is concerned with how feelings, including emotions, modulate learning in LIDA. More complex motivational mechanisms are described in Section 6.

2. The LIDA Architecture and its Cognitive Cycle

The LIDA architecture is grounded in the LIDA cognitive cycle. Every autonomous agent (Franklin & Graesser, 1997), be it human, animal, or artificial, must frequently sample (sense) its environment and select an appropriate response (action). Sophisticated agents, such as humans, process (make sense of) the input from such sampling in order to facilitate their action selection (Hoffman, Singh, & Prakash, 2015). The agent's "life" can be viewed as consisting of a continual sequence of these cognitive cycles. Each cycle consists of three phases, an understanding phase, an attending phase, and an action selection/learning phase. It is commonly referred to as the action-perception cycle (Dijkstra, Schöner, & Gielen, 1994; Freeman, 2002; Fuster, 2004; Neisser, 1976). A cognitive cycle can be thought of as a cognitive "moment." Higher-level cognitive processes are composed of many of these cognitive cycles, each a cognitive "atom".

Just as atoms have inner structure, the LIDA architecture includes a rich inner structure for its cognitive cycles (Baars & Franklin, 2003; Franklin, Baars, Ramamurthy, & Ventura, 2005; Franklin, et al., 2016). During each cognitive cycle a LIDA agent first makes sense of its current situation as best as it can, by updating its representation of both external (coming through the external senses) and internally generated features of its outer and inner world. This is the *understanding* phase of the cycle. By a competitive process (described below), it then decides what portion of the represented situation is most in need of attention, that is, is most salient. This portion is broadcast to the rest of the system, making it the current contents of consciousness. This is the *attending* phase. The contents of a broadcast facilitate the recruitment of internal resources, that is, potential actions, from which the action selection mechanism chooses. It also enables learning in all the available memories (see Figure 1). This is the *action selection/learning* phase.

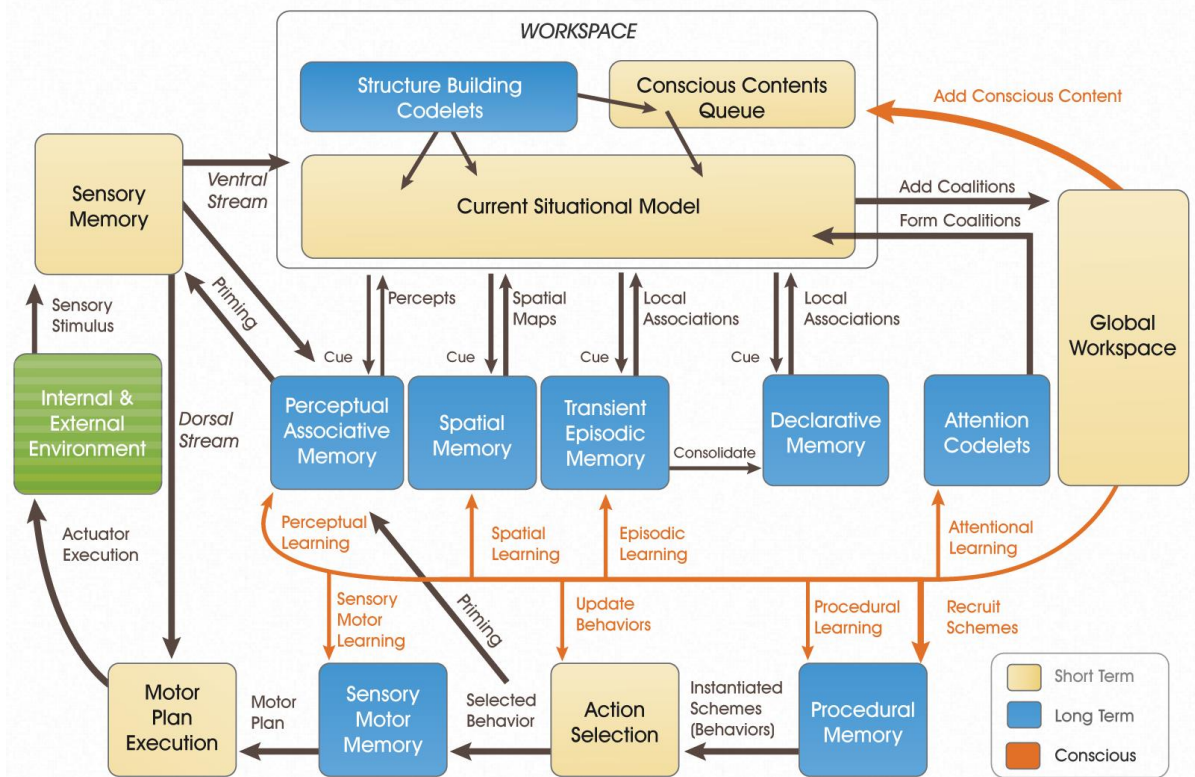


Figure 1. The LIDA Cognitive Cycle Diagram.

Figure 1 shows this highly parallelized processing in more detail. It starts in the upper left corner and proceeds roughly clockwise. The cycle begins with sensory stimuli from sources in the agent's external and internal environment being interpreted in Sensory Memory (e.g., the iconic buffer). Low-level feature detectors in Sensory Memory begin the process of making sense of the incoming stimuli. Next, Perceptual Associative Memory (PAM) processes these low-level features, recognizing higher-level features, such as objects, categories, relations, actions, feelings, events, situations, etc. When recognized preconsciously¹, these entities comprise the current percept, which enters the Workspace asynchronously. Here a preconscious model of the agent's current situation, termed the Current Situational Model, is continually assembled (updated). This percept and items from the Current Situational Model serve to cue transient episodes (Conway, 2001; Franklin, et al., 2005) and declarative memories (autobiographical and semantic). Responses to the cue consist of local associations, which may include remembered events and relevant facts that were associated with the various elements of the cue. Contents of Spatial Memory serve to flesh out events from the episodic memories. In addition to the current percept and the Current Situational Model, the Workspace contains the Conscious Contents Queue, a queue of the recent contents of consciousness (Franklin, et al., 2016).

¹ In the LIDA architecture preconscious representation is currently not conscious, but has the potential to become conscious (Franklin & Baars, 2010).

An updated version of the Current Situational Model (CSM) is assembled from the percepts, the associations, and the undecayed parts of the previous model. This assembling process will typically require structure-building codelets.² These are small, special-purpose processors, each of which has some particular type of structure it knows how to build. These codelets continually monitor the Workspace for opportunities to fulfill their particularly specified task. They may draw upon recognition memory, the Conscious Contents Queue, and even Sensory Memory to enable the recognition of relations and situations. The newly assembled model constitutes the agent's best understanding of its current situation within its world. This completes the understanding phase for this cognitive cycle.

For an agent operating within a complex, dynamically changing environment, the contents of the Current Situational Model (CSM) may be much too large for the agent to cope with all at once. It needs to selectively attend to a salient portion of the CSM. Portions of the CSM compete for attention. These portions take the form of coalitions of structures from the CSM. Such coalitions are formed by special-purpose attention codelets, whose function is to bring certain structures of concern to the particular attention codelet into the Global Workspace (hence the name Global Workspace Theory (Baars, 1988)). The most salient (important, urgent, insistent, novel, threatening, promising, arousing, unexpected, bright, loud, moving, etc.) coalition wins the competition. In effect, the agent has decided on what to attend to at the moment. The contents of the winning coalition are then broadcast globally, bringing its contents to consciousness and, thereby, completing the attending phase of the cycle. Thus, the architecture's functional consciousness mechanism acts as an attention filter. The purpose of all this attentional processing is twofold, first to help the agent choose what to do next, and second to select what to learn into each of its several memories. Learning and action selection take place concurrently in LIDA. We will describe action selection first, and learning later in Section 5.

Though the contents of this conscious broadcast are available globally, a primary recipient is Procedural Memory, which stores templates of possible actions including their contexts and possible results. It also stores an activation value for each such template that attempts to measure the likelihood that an action taken within its context will produce the expected result. Templates whose contexts sufficiently intersect the contents of the conscious broadcast instantiate specified copies of themselves. These instantiated copies are bound to the specifics of the current situation. Instantiated templates remaining from previous cycles may also continue to be available. These instantiations are next processed by the Action Selection mechanism, which chooses a single action from one of these instantiations. The selected action then goes to Sensory-Motor Memory where it is executed by an appropriate algorithm (motor plan) (Dong & Franklin, 2014; Franklin, et al., 2016). The action taken affects the environment, and thus completes that cognitive cycle.

In the LIDA architecture, all cognitive processing takes place via a continual iteration of such cognitive cycles. While it is tempting to imagine each cognitive cycle as a distinct, temporally-disjoint "atom" in a sequence of such atoms, this conceptualization is not strictly true. Due to the asynchronous and parallel nature of most of LIDA's processes, phases from consecutive cognitive cycles can overlap; for example, the action phase from one cycle and the understanding phase from the next cycle can occur simultaneously. The serial nature of consciousness, which is necessary to maintain a stable, coherent interpretation of the world, emerges from the synchronous operation of the Global Workspace. Action selection is also synchronous, enabling the well-ordered execution of sequences of actions. Apart from the Global Workspace and Action

² In the LIDA architecture, the term codelet refers generally to any small, special-purpose processor or running piece of computer code. Codelets correspond to processors in Global Workspace Theory (Baars, 1988)

Selection modules, all other LIDA processes operate asynchronously (that is, without the benefit of a system clock to “synchronize” their operations). This generally asynchronous mode of processing is one of LIDA’s conceptual commitments (Franklin, Strain, McCall, & Baars, 2013).

2.1 Event Representation in LIDA

A recent development in the LIDA model involves the representation of events (Zacks, Speer, Swallow, Braver, & Reynolds, 2007; Zacks & Tversky, 2001). Previously, events were represented by an action node, a perceptual symbol, with thematic-role links connected to role-filling nodes (McCall, Franklin, & Friedlander, 2010). More recently we suggest maintaining the use of action nodes with thematic roles; however, an additional more abstract event node should be added encompassing the action node and its roles, and representing the entirety of the event (Franklin, et al., 2016). For example, suppose the event “Charles takes the pen” is recognized in PAM, and becomes part of the current percept. Then it should appear in the Current Situational Model of the Workspace as shown in Figure 2

With this representation it is clear how events can be recreated by passing activation “top-down” from the event node. Furthermore, since the single event node effectively encapsulates the entire event, the event can become a part of a higher-level structure in a straightforward manner. For example, a link between two event nodes can express a temporal relationship between the two events. For the purposes of this paper we will assume that learned or newly recognized events, while involving several nodes, are effectively encapsulated, and can be referenced by their event node.




Figure 2. The representation of the recognized event “Charles takes a pen.” Double-headed arrows indicate bidirectional information passing. The labels are for illustrative purposes and play no role in LIDA’s dynamics.

2.2 Scheme Representation in LIDA

LIDA’s Procedural Memory serves as a critical bridge from the motivational primitives introduced in this paper to action selection. Procedural Memory provides long-term storage of

ideas of what to do in various circumstances. Each such datum is called a scheme, inspired by Drescher’s schema mechanism (1991). A scheme in LIDA, referred to as a template in Section 2 above, consists of a “context” and a “result” along with an “action,” which connects to “an appropriate network in sensory-motor memory” (D’Mello, Ramamurthy, Negatu, & Franklin, 2006). A scheme also maintains a base-level activation, and, when instantiated as a behavior, a current activation as well. The current activation measures “the relevance of the scheme to the current situation (environmental conditions, goals, etc.),” while its base-level activation is an estimate of “the likelihood of the result of the scheme occurring as a result of taking the action within its context” (D’Mello, et al., 2006).

3. Motivational Concepts and the Role they Play in LIDA

This paper is devoted to a detailed description of how motivation is to be constructed within the LIDA conceptual model. In previous conceptual work, we gave only abstract accounts of using feelings as the basic motivators (for example, Franklin, et al., 2016 Section 4.8) or arguments that their use could be consistent with appraisal theory (Franklin, Madl, D’Mello, & Snaider, 2014). Likewise, most of our agent implementations lacked explicit motivations (for example, Madl et al., 2011; Madl & Franklin, 2012; Faghihi, McCall, & Franklin, 2012), and those that did include motivational constructs, such as feelings or emotions, implemented them in a simplistic way (for example, Ait Khayi & Franklin, 2018). Bear in mind that the LIDA model, as a systems-level model, is intended to be able, in principle, to model the entirety of human cognition. Motivational mechanisms are not added with specific tasks in mind.

In this section we provide several definitions related to motivation by other researchers, and then explain how each concept is implemented within the LIDA conceptual model. We hope to bridge LIDA to existing concepts related with motivation in order to facilitate comparison with existing work. We do not intend to offer LIDA as a comprehensive model of motivation and emotion, but rather as a useful tool for thinking about motivation and emotion, and as a possible architecture for implementing them in software agents.

3.1 Homeostatic Receptor

Biological agents use *homeostatic control mechanisms* to regulate their internal environment, keeping internal variables within stable ranges (Cannon, 1929). More recently they have been used to implement drives in robots (Breazeal, 1998). Each homeostatic control mechanism has at least three interdependent components for the variable being regulated. The *receptor* is the sensing component that monitors and responds to changes in the internal environment. When the receptor senses a stimulus, it sends information to a “*control center*,” the component attempting to keep the variable within a stable range. The control center then sends signals to *effectors*, which can be muscles, organs, or other structures, that receive signals from the control center (Marieb & Hoehn, 2007). Therefore, homeostatic receptors are sensors that detect the value of critical attributes in an agent’s bodily state. Examples of these in humans include sensors for low levels of oxygen or glucose in the blood. The detection of low oxygen and/or high carbon dioxide in the body sends a signal to the brain (control center). From this point, the brain must interpret the signal. Such interpretation gives rise to a *feeling* of, for example, the need to breathe more rapidly. A sensor may be implemented in many ways, for instance, with distributed inputs from all over the body or by taking inputs from a localized bodily region.

In the LIDA conceptual model, homeostatic sensors are implemented as sensors, and are considered to be a part of Sensory Memory (see Fig. 1). A particular feature detector begins the control process detecting an incoming internal stimulus, and activating an associated feeling node in Perceptual Associative Memory. The control center concept is effectively described throughout the rest of this paper including the aforementioned mechanisms of LIDA’s cognitive cycle.

3.2 Feelings

Damasio viewed feelings as somatic markers (Damasio, 1999), the perception of a certain state of the body along with the perception of a certain mode of thinking, and thoughts with certain themes (Damasio, 2003). Franklin and Ramamurthy (2006) also adopted this view writing:

Feelings in humans include hunger, thirst, various sorts of pain, hot or cold, the urge to urinate, tiredness, depression, etc... One feels feelings in the body. Implemented biologically as somatic markers, feelings typically attach to response options and, so, bias the agent’s choice of action...Feelings, including emotions, are nature’s means of implementing motivations for actions in humans and other animals. They have evolved so as to adapt us to regularities in our environments. These general preferences derived evolutionarily from regularities can be viewed as values. Thus feelings become implementations of values in biological agents, providing a common currency for quick and flexible action selection.

We adopt this notion of feelings for this work. Feelings are implemented as PAM nodes in the LIDA conceptual model. They always have exactly one valence sign, either positive or negative. We distinguish two types of feeling nodes:

- 1) **Drive feeling nodes**, each representing a drive of the agent
- 2) **Interpretive feeling nodes**, each representing a perceptual interpretation of a stimulus

The former is a representation of a current internal state of the agent’s body, independent of the current external environment, while the latter is the agent’s evaluation of a stimulus regardless of whether it is external or internal to the agent. For example, a “hungry” feeling node represents a drive, and has negative valence sign, while a “satiated” feeling node also represents a drive but with positive valence sign. A “sweet” feeling node represents a perceptual interpretation of some stimulus with positive valence sign, while a “bitter” feeling node is similar, but with negative valence sign. Intuitively, one may be hungry or satiated in a variety of external situations, and one may interpret a variety of different stimuli as sweet or bitter.

3.3 Affective Valence

Affective valence is the hedonic niceness or nastiness (“liking” versus “disliking”) of reward or aversion that is essential to emotions (Berridge & Kringelbach, 2008). For LIDA, we propose that each feeling node has a *valence sign*, having either a positive or negative value. With a positive value, the sign stands as the agent’s basic representation of niceness or liking, while a negative value represents basic nastiness or disliking. For example, the feeling node “thirst” would have a negative valence, while the feeling node “relief-from-thirst” would have a positive valence. A

feeling node's valence sign plus its current activation specifies its *affective valence*, with the activation representing the degree of "liking" or "disliking."

Like all PAM nodes, feeling nodes must be activated from a sensory stimulus to be instantiated in the Workspace. Said in another way, feature detectors in the feeling node's receptive field must detect features of the sensory stimulus. Either by direct recognition enabled by earlier perceptual learning, or by close temporal association, feeling nodes may become part of preconscious event representations in the Workspace. Such events will be referred to as *feeling events*, a distinction likely true of most all events. The affective valence of a feeling event is the sum of the affective valence of its feeling nodes. The feeling node can be aptly thought of as a feature of the event. This suggests that recalled or imagined feeling events may involve a non-zero degree of affective valence. By having different feeling nodes carrying positive or negative affective valence, it is also possible to both like and dislike the same event structure. This can happen when two feeling nodes with positive and negative affective valence respectively become part of the same event structure.

3.4 Liking vs. Wanting and Disliking vs. Dreading

Kringelbach and Berridge (2009) distinguish "wanting" from "liking." By analogy, "dreading" should be distinguished from "disliking." "Wanting" is an objective hedonic reaction, measured behaviorally or neurally. Core "liking" reactions result from activity in identifiable brain systems that paint hedonic value on a sensation such as sweetness (Berridge & Kringelbach, 2008). There is evidence of separate neural representations for wanting and liking (Smith, Berridge, & Aldridge, 2011). Dread appears to be the opposite of wanting or desire, as neurotransmitters in the nucleus accumbens have been shown to modulate whether the area generates desire (wanting) or dread (Richard & Berridge, 2011).

Incentive Salience. Incentive salience is a motivational "wanting" attribute of stimuli, transforming the representation of a stimulus into an object of attraction. This "wanting" is unlike "liking" in that liking is produced by a pleasure immediately gained from consumption or other contact with stimuli, while the "wanting" of incentive salience is the motivational magnet quality of a stimulus that makes it currently a desirable and attractive goal, transforming it from a mere sensory experience into something that commands attention, induces approach, and increases the likelihood of its being sought out (Berridge & Robinson, 1998). Roughly, stimuli have incentive salience, i.e., are "wanted," only if they have been previously "liked." However, it has also been demonstrated in biological agents that physiological signals, e.g., hunger, are *integrated* with a stimulus' incentive salience (Tindell, Smith, Berridge, & Aldridge, 2009). Thus, overall incentive salience is adaptable, as its learned component is modulated by current physiological signals, a view originally articulated by Bindra (1978) and Toates (1986), and later referenced in the incentive salience model (Berridge & Robinson, 1998).

Inspired by all of this, we include incentive salience as a distinct scalar attribute of event nodes including base-level³ and current components. This attribute is similar to, but separate from, node activation. Activation represents a node's salience, which may involve relevance, importance, urgency, insistence, novelty, unexpectedness, loudness, brightness, motion, etc.

³ PAM is recognition memory, and a PAM node having some base-level incentive salience is not sufficient for recognition. However, base-level incentive salience is a fundamental attribute of nodes, and, as such, is considered a part of PAM.

While activation represents an event’s situational salience, incentive salience represents the value the agent currently places on the event actually occurring. Our version of incentive salience has its range from a positive value, (e.g. 1), where it represents the degree that an event is “wanted,” to a negative value, (e.g. -1), representing the degree of aversion to the event occurring. Just as a node’s total activation is a function of its base-level activation (learned usefulness) and current activation (Franklin, et al., 2016), a node’s total incentive salience combines its base-level incentive salience and current incentive salience. Base-level incentive salience is modified in perceptual learning (see Section 5.2); events gain incentive salience by being associated with feelings with positive affective valence. Likewise, they lose incentive salience when associated with feelings having negative affective valence. Current incentive salience is determined, in part, by activated feeling nodes that are connected to events by incentive salience links (see Section 5.5). Additionally, current incentive salience may come from expected future events (see Section 6.4). Both the total activation and the total incentive salience of an event contribute to its salience, making it more likely to come to consciousness.

3.5 Option

An option for action (option) (Keller & Ho, 1988; Klein, Wolf, Militello, & Zsombok, 1995; Raab, de Oliveira, & Heinen, 2009; Ward, Suss, Eccles, Williams, & Harris, 2011) is an action representation an agent can select for execution. It is a candidate for action or a “choice alternative” (Kalis, Kaiser, & Mojzisch, 2013).

So far, we have discussed feeling events, events with one or more associated feeling nodes. Feeling events have an overall affective valence, which is based on the aggregate affective valence(s) of the feeling node(s) associated with (linked to) the event. We have also proposed that events could have incentive salience. In fact, we hypothesize that most every event will involve some feelings and some non-zero amount of incentive salience. Events in PAM can have base-level incentive salience; however, they do not have current incentive salience before coming to the Workspace.

Given this, we define a *virtual event* as an event that is instantiated without the presence of its corresponding sensory signature from the agent’s sensors. For example, recalled episodic memories, plans, imaginings, options, and hallucinations are all virtual events. Next, we define an *option* to be an instantiated virtual event in the Workspace with some non-zero total incentive salience. Thus, options are unique in that their salience is based on an event’s total incentive salience, and, possibly, the event’s activation (e.g., from stimuli) as well. Options can be thought of as possible “choices” that have not yet been made. When an option helps instantiate a scheme that is likely to bring it about, that is, a behavior with the option for its result, and if such a behavior is selected, then the option becomes an *immediate goal*. Given that a behavior stream is a stream of behaviors that can be thought of as a partial plan of action, we define a *pending goal* as an option that is a part of an instantiated behavior stream, one of whose behaviors has been selected, but the underlying event cannot be immediately achieved.

3.6 Emotion

The usage of the term “emotion” has been imprecise and, for many researchers, the concept is not easily definable (Alvarado, Adams, & Burbeck, 2002; Thompson & Madigan, 2007). Various definitions of emotion have been given, and very important functions have been attributed to emotion; however, there is currently no consensus for any one definition. Some well-known work

on emotions specifies six basic emotions: surprise, fear, disgust, anger, happiness, and sadness (Ekman, Sorenson, & Friesen, 1969). They are particular to individuals (Picard, 2003), and influence cognition directly (Phelps, 2006; Squire & Kandel, 2000). More recently some researchers deny that there is anything “basic” about these six (Camras, 2011). Westen (1999) describes emotions as a mix of several biochemical, sociocultural, and neurological factors. Purves et al. (2008) divide emotions into three processes: 1) a behavioral action (e.g., escape), 2) a conscious experience of an event or situation (e.g., regret or shame), 3) a physiological expression (e.g., paleness, blushing, palpitations, uneasiness). It is not specified how these three processes are related. Diener (1999) suggested that a theory of emotions should account for several components: 1) subjective experience (how it feels to be in an emotional state), 2) physiological processes (neural, neurochemical and physiological mechanisms that facilitate emotion), 3) expressive behavior (facial and bodily expression, movement patterns, the modulation of voice and breath etc.), and 4) cognitive evaluations.

The LIDA model does not address subjective experience; rather its consciousness is only postulated to be functional⁴ (Franklin, 2003). While “cognitive evaluations” are addressed in the next section, the physiological and expressive aspects of emotions are beyond the scope of this current work.

It is our view that emotions exist only within the context of an organism, or an agent. It is a feeling about *something* or *someone* that is an emotion. For example, the emotion of *feeling* sad because one has suffered a loss, or the emotion of *feeling* angry at someone because they interfere with something you want to do. We adopt such a definition of emotion for the LIDA model (Franklin & Ramamurthy, 2006):

Emotions, such as fear, anger, joy, sadness, shame, embarrassment, resentment, regret, guilt, etc., are taken to be feelings with cognitive content (Johnston, 1999). One cannot simply feel shame, but shame at having done something. The something done constitutes the cognitive content.

A feeling node in PAM would represent shame. A currently active event node, the cognitive content, with the shame node, within the Workspace would constitute an instance of the emotion shame. Once again, such an emotion may, or may not, come to consciousness. Shame, of course, may come in multiple varieties, each represented by its own node. Note that every emotion is a feeling, but there are feelings, e.g. thirst or a pain, that need not be emotions. Thus emotions constitute a subset of feelings.

3.7 Appraisal

Appraisal theory is described as providing a cognitive interpretation of what we sense or perceive (Lazarus, 1991). Moreover, this theory explains our evaluation of specific external (environmental) or internal (within ourselves) stimuli that give rise to emotions. Roseman and Smith (2001) explained how human motives and goals play an important role for the evaluation

⁴ Global Workspace Theory, which LIDA implements and fleshes out, is a high-level theory of the role of consciousness in cognition. It assumes that phenomenal consciousness occurs in conjunction with each conscious broadcast. Wishing to apply our LIDA model to both biological and artificial agents, the model remains agnostic on the question of phenomenal consciousness.

of a specific situation. Lazarus (1991) identified two appraisal methods: 1) primary appraisal, for the organism to determine the importance and/or meaning of an event, and 2) secondary appraisal, for the organism to determine whether or not it can manage or handle the event and its consequences.

Lazarus also discusses two types of appraisal models: the structural appraisal model and the process model. The structural model has three components: 1) a relational aspect, involving the person-environment interactions, 2) a motivational aspect, involving the assessment of a current situation's relevance in relation to one's current goals, and 3) a cognitive aspect involving one's appraisal of the situation. Accordingly, one's evaluation of the situation leads to specific emotions.

The process model suggests that person-environment interactions might go through an appraisal-coping-reappraisal cycle (Lazarus, 1991). For example, imagine a father becomes angry (appraisal) after hearing of his daughter's disobedience. He then has a heated argument with her (coping). Later on, he may feel guilty about his behavior during the argument (reappraisal), which, in turn, may lead to an apology (coping). This, finally, leads to a feeling of satisfaction at having apologized (reappraisal). However, the process model of appraisal fails to account for automatic emotional response, a fast and direct pathway for action in critical situations. We will discuss how this occurs in the LIDA model in detail in Section 4.1 below on *alarms*.

In LIDA model terms, appraisal is an initially preconscious process that occurs automatically with perception, but may also involve episodic and semantic recall, imagination, and consciousness. For this work, we consider the appraisal process as one that evaluates a situation (event) along dimensions such as novelty, goal-alignment, agency, coping potential, and availability of a plan, thereby determining the activation level of relevant feeling nodes. In this sense the appraisal process culminates in the attachment of one or more activated feeling nodes to the cognitive content, typically an event, producing a corresponding emotion. This is typically accomplished in LIDA's Workspace by the preconscious association of an event and the emotion(s) activated during the event's appraisal. Such association is performed by structure-building codelets concerned with particular appraisal dimensions. The attachment of the feeling node(s) modulates the current activation of the event. If this newly appraised structure, including the feeling node(s) is brought to consciousness, this emotion, this feeling event, is learned into PAM (long-term recognition memory). Once learned, emotions, like other feelings can be recognized in PAM without additional appraisal. Often times the appraisal of emotion is a multi-cyclic affair requiring at least several cognitive cycles. However, we hypothesize that the initial appraisal of emotion can also occur during a single LIDA cognitive cycle.

3.8 Reinforcement Learning

Reinforcement learning (RL) is a form of machine learning in which software agents learn to optimize their behavior through trial-and-error exploration of their environments. Agents choose actions based on their beliefs about their environment's current state, and then observe the subsequent states and rewards following the execution of those actions. Rewards are formulated as scalar values, where the rewards with the greatest numerical value are the most desirable. In general, this reward signal serves as a proxy for the fulfillment of an agent's goals, and any sequence of actions that maximizes the expected total reward accumulated by an agent can be interpreted as an optimal solution for achieving those goals. Note that with respect to the application of reinforcement learning in LIDA agents (and autonomous agents in general), all

“rewards” are derived from that agent’s perceptions rather than an external reward signal; more specifically, they are determined based on the activation of built-in motivational structures. If an agent’s perceptions include appraisals with positive affective valence, that is, if an event is liked, then the presence of that affective valence can be viewed as a positive “reward” with a value proportional to the magnitude of the affective valence.

In many cases, an agent’s actions may not directly result in an immediate reward, but, instead, contribute to the acquisition of future rewards. These delayed rewards result in a credit-assignment problem where partial credit should be attributed to actions that contributed to eventual rewards. Without the assignment of partial credit to these unsung supporting actions, an agent’s behavior will tend to be myopically focused on immediate gratification. Many reinforcement learning algorithms solve this credit assignment problem by introducing explicit value (or utility) functions; for example, functions that approximate the expected cumulative reward associated with a given state or action taken from that state. These approximations are conditioned on the assumption that the agent behaves as it did in the past (or, more precisely, that the agent follows the same behavioral policy). Value functions are used as the basis for many popular *model-free* RL algorithms; for example, Q-learning (Watkins, 1989) and Deep Q-Networks (Mnih et al., 2015).

In the context of cognitive science, Lucantonio, Stalnaker, Shaham, Niv, and Schoenbaum (2012) described model-free learning as “a set of reinforcement learning methods that use prediction errors to estimate and store scalar cue or action values from experience” (p. 360). These stored values represent the predicted total reward an agent expects to receive if an action or cue were pursued. The name “model-free” is based on the fact that actions are evaluated and selected using these stored values, without the need for an internal model of the environment. Lucantonio et al. characterized such a decision strategy as “simple but inflexible” because “the values are simply scalar numbers, separated from the identity of the expected future outcomes themselves or the specific events that will ensue en route to obtaining the outcomes” (p. 360). Furthermore, these values do not immediately change if the outcomes change (e.g., a cue, typically followed by a reward, suddenly begins to lead to a punishment). Instead, the values must be gradually relearned through repeated experience and prediction error-based updates. Due to this inflexibility, many have suggested that model-free decision making underlies habitual behavior (Fishbach, Roy, Bastianen, Miller, & Houk, 2005; Lucantonio, et al., 2012).

In Section 5.3 of this paper, we introduce a new model-free motivational learning mechanism based on temporal-difference learning. The cached values updated by this learning mechanism are the base-level incentive saliences of event nodes, which are modified during conscious learning. These learned base-level incentive salience values support LIDA’s *consciously mediated action selection* (among other things), which is a form of model-free decision making in which actions are selected for execution based primarily on the total activation and incentive salience of their corresponding instantiated schemes (see Section 2.2). The term “consciously mediated action selection” refers to the fact that conscious access is required for this mode of action selection; however, actions are chosen impulsively and without the benefit of deliberative thought. For example, based solely on consciously mediated action selection, an agent having an “open-the-refrigerator-door” event node with a strong positive base-level incentive salience may habitually open the refrigerator door when hungry, even when the refrigerator is empty, because this event has historically led to the consumption of food—an event with positive incentive salience.

An alternate approach to reinforcement learning, referred to as *model-based* RL, learns internal models that predict the effect of an agent’s actions on its environment, and uses those models in service of decision making. In the context of cognitive science, Lucantonio, Stalnaker, Shaham, Niv, and Schoenbaum (2012) describe model-based learning as “a set of reinforcement learning methods in which an internal model of the environment is learned and used to evaluate available actions or cues on the basis of their potential outcomes” (p. 360). Unlike model-free learning, which is based on cached values, model-based learning methods evaluate and select actions by simulating their consequences, and choosing the most desirable predicted outcomes. This approach is computationally expensive, but flexible, allowing rapid adaptations to an agent’s behavior given an accurate and up-to-date model of the agent’s environment. It was suggested by Lucantonio et al. that goal-directed behavior relies on model-based learning.

In LIDA’s Workspace, the Current Situational Model (CSM) contains relevant portions of the agent’s current model of its environment, and its beliefs about its current environmental state. Over multiple cognitive cycles, LIDA agents may mentally simulate the results of possible actions, or entire action plans, and the total incentive salience associated with those outcomes may be used to bias the outcome of LIDA’s *volitional (deliberative) action selection* (Franklin, 2000). Section 6 describes volitional action selection in more detail, and, in particular, describes how the motivational constructs introduced in this paper (such as incentive salience) can be used in support of deliberation and model-based control.

4. Single-Cycle Motivational Processes in LIDA

In Section 3, we introduced several motivational enhancements to LIDA’s conceptual model. In this section, we describe how they can influence LIDA’s single-cycle processes. Most mechanisms make a trade off between speed and the amount of memory brought to bear on the situation. This makes each best suited for particular situations; however, implemented together in an integrated architecture they provide a flexible set of faculties well-suited for a range of situations.

To help frame this section, we first present two distinctions used in LIDA. The first is a distinction between levels of conscious access. LIDA identifies three levels of conscious access of mental processes or representations: 1) *Never conscious* means unable to ever become conscious, 2) *Preconscious* means potentially, but not currently, conscious, and 3) *Conscious* means it is part of the content of a conscious broadcast (Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Franklin & Baars, 2010). LIDA also distinguishes between four types of action selection: volitional, consciously mediated, automatized, and alarm. In *volitional (deliberative) action selection (decision making)* much of the selection process itself is conscious, that is, a part of successive contents of consciousness (Franklin, 2000). *Consciously mediated action selection* may occur in each of LIDA’s cognitive cycles, when a never-conscious selection process consults the current contents of consciousness during its choosing of the next action (Franklin & Baars, 2010; Franklin, et al., 2012). *Automatized action selection* occurs when one action calls the next via a never-conscious process, and without use of the current contents of consciousness (Franklin, et al., 2012; Negatu, 2006). Finally, we discuss action selection during an alarm in the next section.

4.1 Alarms

In dangerous situations or during certain athletic occasions, agents must react as quickly as possible. The notion of “alarms” was previously articulated in the context of agent architectures by Sloman (1998) and enjoys neuroscientific support (LeDoux, 2006; Liddell et al., 2005). In terms of the LIDA model, an *alarm situation (alarm)* is one where an agent instantiates a scheme, selects its action, and executes an associated motor plan in rapid succession in response to a dangerous or athletic event prior to, or simultaneous with, consciousness of the event. This would correspond to an action selection and execution before (or while) the event is brought to consciousness by attention codelets. For example, an alarm occurs with the unconscious turning of a car’s steering wheel in response to suddenly being cut off by another vehicle, or a professional baseball batter swinging at a pitch. During an alarm, some part(s) of the Current Situational Model in the Workspace directly instantiates a scheme from Procedural Memory. The ensuing behavior is sufficiently activated for the immediate selection of its action and, in close succession, the execution of an associated motor plan. All of this occurs before, or simultaneously with, consciousness. Note that alarms require the existence of some learned or built-in scheme. They also require that a scheme be strongly “activated.” Here, we suggest the activation *and the incentive salience* both play a role in the rapid recruitment of scheme(s) during alarms. Note that every alarm directly instantiates a scheme before it can come to consciousness. However, later or simultaneously, an agent can become conscious of the situation causing the alarm, and of the action it took in direct response to the alarm. After an action is taken during an alarm situation, the result will likely come to consciousness and become involved in the procedural learning for the scheme selected and executed during the alarm. (LeDoux, 2006; Liddell, et al., 2005; Sloman, 1998)

4.2 Motivational Effects of Feeling Nodes

The presence of activated feeling nodes, be they drive feelings or interpretive feelings, in a LIDA agent’s Workspace has several motivational effects. Feelings and emotions are part of the cue to Transient Episodic and Declarative Memory (Franklin & Ramamurthy, 2006) as well as to LIDA’s nascent Spatial Memory (Madl, Franklin, Chen, & Trappl, 2013). The resultant local associations can contain records of the agent’s past feelings and emotions in associated situations. Present and past feelings and emotions influence the competition for consciousness in each cognitive cycle, by affecting the activation of structures in the Current Situational Model, as well as the activation of coalitions containing such structures, strengthening a coalition’s chances of winning the competition for consciousness. This implies conscious content can include affective portions, and play a role in conscious learning (see *Perceptual Associative Memory Learning* below). If modeling human performance, then affect should modulate learning following an inverted U curve according to the Yerkes-Dodson law (Yerkes & Dodson, 1908): increasing affect causes increased learning up to a point, after which more affect results in less learning.

4.3 Incentive Salience of Nodes Contributes to Motivation

We have proposed that event nodes in PAM have a base-level incentive salience while their current instantiations in the Current Situational Model additionally have a current incentive salience. The total incentive salience of events contributes to the salience of events, and as a result, can play several roles within a single cognitive cycle. A coalition involving events with

high incentive salience is more likely to win the competition for consciousness, and become the object of conscious learning. If broadcasted, events having a large incentive salience better recruit schemes from Procedural Memory, and can instantiate more salient behaviors, increasing their likelihood of selection. This influence of incentive salience on action selection is akin to the influence of “goals” on behaviors in Maes’ behavior network model of action selection (Maes, 1989).

4.4 Action Selection and Consciousness

The LIDA model distinguishes action selection from motor plan selection. The former is the selection of a behavior (an instantiated scheme) from Procedural Memory, while the latter is the selection of an instantiated motor plan template from Sensory-Motor Memory in the service of action execution. Procedural Memory is built using structures originating in Perceptual Associative Memory, i.e., it is a grounded representation. This implies that it may be primed preconsciously, allowing for preconscious priming of actions. Does this suggest that action selection occurs without consciousness? The answer is that there can be several possibilities: The unconscious selection of alarms can occur with a sudden, strong instantiation of a learned scheme that leads to a motor plan execution before consciousness of the event can occur (see Section 4.1). However, in the vast majority of situations, action selection follows the conscious broadcast. Such consciously mediated action selection (see Section 4) can be viewed as a “fast” decision-making using cached base-level activation and base-level (as well as current) incentive salience values to select options. Such a mechanism has also been referred to as model-free control (Daw, Niv, & Dayan, 2005) as well as “System 1” (Faghihi, Estey, McCall, & Franklin, 2015; Kahneman, 2003, 2011). Finally, the action selection process can itself be partially conscious, as is the case in volitional decision-making (see “Volitional Decision-making in LIDA” below). Volitional decision-making has also been referred to as model-based control (Daw, et al., 2005) and Kahneman’s “System 2” (2011).

5. Conscious Learning in LIDA

LIDA’s conscious learning involves several memories. In Perceptual Associative Memory (recognition memory), new concepts, and associations between them, are learned while existing ones are reinforced. In Spatial Memory, new cognitive maps are stored and existing ones updated. In Transient Episodic Memory broadcasted events containing what, where, and when are encoded. *Schemes*, procedural memories of what to do when, are learned into Procedural Memory. It is also possible for attention codelets, as well as structure building codelets, to be learned or reinforced. Since affective valence plays a role in determining saliency, its presence among the content of a broadcast affects conscious learning, in addition to influencing scheme recruitment and action selection. Following Tindell et al. (2009) affective valence is hypothesized to strongly influence learning based on feelings. We present the details of this below and then discuss how such learning may lead to appetitive or aversive behavior.

5.1 Perceptual Associative Memory Learning

Whenever an event is part of a conscious broadcast the nodes and links of the event are learned into PAM (perceptual learning). If the node or link is not currently present in PAM, then it is added to PAM, and given an initial base-level activation according to some function of the node's activation. If the node or link is already present in PAM, then its base-level activation is updated according to the same function.

5.2 Base-Level Incentive Salience Learning

Previous work on feelings as motivators (Franklin & Ramamurthy, 2006), stipulated that feeling nodes, each with an activation, play a role in LIDA's conscious learning, i.e., add to a containing coalition's activation which, in turn, modulates the learned base-level activation of coalition nodes. Here we additionally propose that feeling nodes have a valence sign, positive or negative. We now specify how affective valence *additionally* affects perceptual learning. We propose that for feeling events having affective valence, an additional kind of perceptual learning occurs in PAM. The base-level incentive salience of broadcast event nodes is positively updated as a function of the overall affective valence of the event, that is, the sum of the affective valences of all feeling nodes (drive feelings or interpretive feelings) currently associated with the event. If the event has an overall positive affective valence, then base-level incentive salience is increased, while if it has an overall negative affective valence its base-level incentive salience is decreased. The magnitude of the overall affective valence determines the strength of this base-level incentive salience learning. In summary, perceptual learning updates the base-level incentive salience of event nodes. After such learning, events can be directly recognized as typically desirable or not. Below, in the section on incentive salience links, we discuss a learning sub-type affecting how event nodes in the Current Situational Model obtain *current* incentive salience.

5.3 Temporal Difference Learning

Temporal difference (TD) learning (Sutton & Barto, 1998) is an approach to learning how to predict a quantity, which depends on future values of a given signal. TD is a reinforcement learning technique, which considers not just the current reward but the expected future rewards as well. The name TD derives from its use of the changes, or differences, in predictions over successive time steps to drive the learning process. In general, the process mandates that the prediction of a quantity (e.g., the base-level incentive salience as a prediction of the affective valence) at any given time step be updated to bring it closer to the prediction of the same quantity at the next time step. TD algorithms are often used in reinforcement learning to predict a measure of the total amount of reward expected over the future, which, in the context of this paper, corresponds to base-level incentive salience. A number of neuroscientific findings support the hypothesis that the brain performs TD learning (Hollerman & Schultz, 1998; Montague, Dayan, & Sejnowski, 1996; O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003). The mechanisms underlying the changes in valuation appear to occur during learning itself, and "bias decisions without effortful retrieval at the time of decision" (without deliberation, which is effortful in LIDA) (Wimmer & Shohamy, 2012).

Let us denote the t^{th} conscious event (or state) in LIDA by s_t , and the event's accompanying overall affective valence by r_t . Also, let the *discount rate*, γ , be a real number from the closed interval, $[0,1]$. This rate defines the degree to which the rewards (overall affective valence) of

future events are discounted or de-emphasized, with a value of 1.0 corresponding to no discounting and 0.0 corresponding to complete discounting. Given these definitions, we define the *complete return*⁵, R_t , for conscious event s_t as:

$$R_t = r_t + g r_{t+1} + g^2 r_{t+2} + \dots = r_t + g(r_{t+1} + g r_{t+2} + \dots) = r_t + g R_{t+1} \quad [1]$$

It is then assumed that the long-term value of a state (here, conscious event), $V(s)$, is directly equivalent to the expected return under an unspecified action policy⁶, π :

$$V(s) = E_{\rho}(R_t | s_t = s) \quad [2]$$

In the context of this work, we denote the base-level incentive salience of the event node in memory (PAM) underlying the i^{th} conscious event s_i at the time t , by $V(s_i, t)$. Thus $V(s_i, t) = V(s)$, where s is s_i at the time t . Given a learning rate, α , the temporal difference update for the base-level incentive salience of event s_i at the time $t+1$ is given by:

$$V(s_i, t+1) = V(s_i, t) + \alpha[r_t + gV(s_{t+1}, t) - V(s_i, t)] \quad [3]$$

This equation says that the update for the base-level incentive salience of event s_i is a function of the difference between the event's previous base-level incentive salience, $V_t(s_i)$, and a new estimate of the complete return of s_i ; namely, $r_t + gV(s_{t+1}, t)$. Notice that the r_t term corresponds to the affective valence-based learning. Thus, temporal difference learning additionally suggests using the discounted base-level incentive salience of the next encountered event when learning base-level incentive salience of an event.

5.4 Event Sequence Learning

Sequences are a large, critical part of Perceptual Associative Memory, connecting patterns having temporal regularity. Consider the multitude of melodies, word phrases, and action sequences that one may recognize immediately; these are all event sequences in PAM. As mentioned in the earlier section on schemes, event sequences are integral to schemes. Such temporal connections may be initially "hypothesized" by structure-building codelets as temporary temporal links in LIDA's preconscious Workspace. If two events and their temporal link win the competition for consciousness then such a link between them would be learned in PAM. The role of event sequences in volitional decision making will be explored in the Section 6.

⁵ Since LIDA considers the overall affective valence (reward) as coming with a given conscious event (state), the complete return is defined starting at t , the time of the event, instead of $t+1$ as is done in the reinforcement learning literature.

⁶ Please see http://www.scholarpedia.org/article/Reinforcement_learning

5.5 Incentive Saliency Link Learning

Drive feeling events and their affective valence play another role in perceptual learning. This learning is critical in determining which event nodes (whether instantiated or not) receive current incentive saliency from active drive feelings. In the proposed model, current incentive saliency is transmitted from drive feelings by incentive saliency links, that is, PAM links that transmit only incentive saliency, not activation. Let us first consider an example of learning an incentive saliency link. Suppose an agent becomes conscious of an active “thirst” drive feeling node having a negative affective valence. Despite this awareness, the agent takes a long walk, and returns home with the “thirst” feeling node, now strongly activated, carrying a large negative affective valence. Due to a preconscious temporal association by a structure building codelet in the Workspace, the agent becomes conscious of the fact that the walk led to the increase in the thirst node’s activation. Consequently, the agent learns an incentive saliency link with a negative weight from the “thirsty” node to the “walking” event. (Note that this incentive saliency link learning would occur in addition to the learning of other links, such as a temporal link running from the walking event node to the strong feeling of thirst event.) Having learned this incentive saliency link, whenever the “thirsty” node becomes activated, “walking” (and any other similarly implicated events) receive a negative amount of current incentive saliency. In contrast, suppose that the agent had instead drunk water in response to being conscious of the feeling “thirst,” and consequently reduced the drive feeling’s activation. If it becomes conscious of this temporal relationship, an incentive saliency link with a positive weight would be learned, which would promote drinking when thirsty.

How can a LIDA agent learn from experiences involving a change in a drive feeling? In general, these situations involve a temporal link in the Workspace from one event to another event involving *change in activation*, across the two events, of a *drive feeling node*. (Note that such a preconscious association could occur between instantiated representations from various places in LIDA’s memory.) In such situations, a structure building codelet can introduce an incentive saliency link that has the *feeling node as its source* and the event implicated in the feeling node’s change as its sink. Supposing that the entire structure, consisting of both nodes connected by the temporal link, wins the competition for consciousness, then the proposed incentive saliency link would be learned into PAM. Links learned in this manner are assigned a weight proportional to the *change* in the drive feeling’s activation, a , multiplied by the feeling’s valence sign, v_s . Equation 4 expresses how this link weight, w , is proposed to be calculated:

$$w = v_s \cdot f(Da) \quad [4]$$

After such links are learned they transmit some amount of *current incentive saliency* to their sink events proportional to 1) the activation of the source (the drive feeling node) and 2) the link’s weight. This, by itself, may lead to the instantiation of the sink. Intuitively, the idea is to send positive incentive saliency to events leading to positive changes to homeostasis, be they a decrease in a negative affective valence or increases in positive affective valence. Similarly, negative incentive saliency is sent to events leading to negative changes to homeostasis under this approach. Figure 3 illustrates the example.

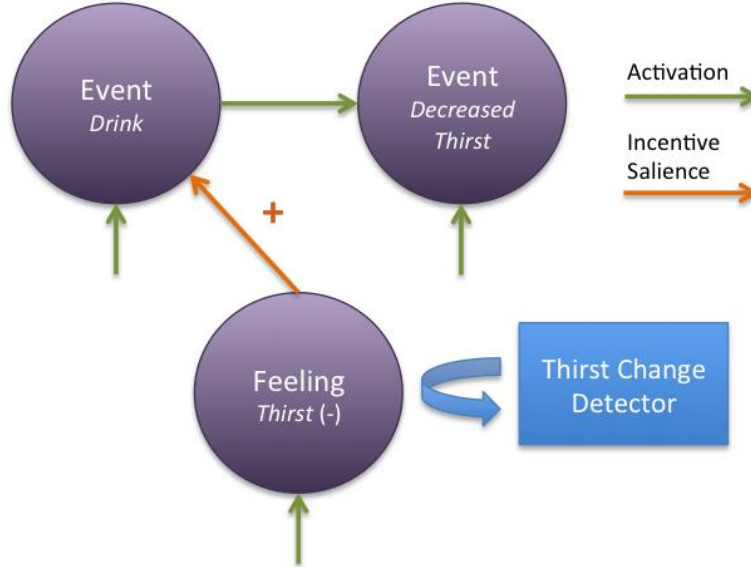


Figure 3. A previously learned incentive salience link in PAM. The orange-colored link depicts an incentive salience link having a positive weight in (long-term) PAM. The link's source is the *thirst* drive feeling having a negative valence sign. The link's sink is the event implicated in the *decreased thirst* event, here, the *drink* node. The thirst-change detector, a structure building codelet, can recognize decreased thirst by monitoring changes in activation of the *thirst* feeling node. It creates the Decreased Thirst event and the temporal (causal) association link from the Drink event. The learning process updates the weight of the incentive salience link according to Equation 4. In this example, a feeling, having a negative valence sign, decreased in activation, necessitating a positive link of drinking to reduce thirst. After an incentive salience link with positive weight is learned, the link transmits positive current incentive salience, based on the activation of the “thirst” node, to the “drinking water” event.

5.6 Action Selection and Consciousness

The LIDA model distinguishes action selection from motor plan selection. The former is the selection of a behavior (an instantiated scheme), a procedural memory, while the latter is the selection of an instantiated motor plan template, a sensory-motor memory (Dong & Franklin, 2015). As mentioned in the previous section, Procedural Memory is built using structures originating in perceptual memory, i.e., it is a grounded representation. This implies that it may be primed preconsciously, allowing for preconscious priming of actions. Does this suggest that action selection occurs without consciousness? The answer is that there can be several possibilities: The unconscious selection of alarms can occur with a sudden, strong instantiation of a learned scheme that leads to a motor plan execution before consciousness of the event can occur. However, in the vast majority of situations, action selection follows the conscious broadcast (consciously mediated action selection). Consciously mediated action selection can be viewed as a “fast” decision making using cached base-level activation and base-level incentive salience values (and current values too) to select options (see (Franklin, et al., 2016) for a detailed description). Such a mechanism has also been referred to as model-free control (Daw, et al., 2005) as well as “System 1” (Kahneman, 2003, 2011). Finally, the action selection process can itself

be partially conscious, as is the case in volitional decision making. Volitional decision making has also been referred to as model-based control and Kahneman's "System 2."

6. Multi-Cyclic Motivational Processes in LIDA

So far, we have discussed the role of motivational processes within a single cognitive cycle. However, many important aspects of cognition are implemented by, and occur over, multiple cognitive cycles. Recall the previous distinction (Section 3.5) between 1) an option, an instantiated virtual event having nonzero total incentive salience, 2) an immediate goal, an option that is the expected result of a selected behavior, and 3) a pending goal, a goal that is part of an instantiated behavior stream, one of whose behaviors have been selected, that cannot be immediately achieved. Now, suppose that an option comes to mind, that is, is a part of the content of a conscious broadcast. As a direct result, several schemes in Procedural Memory will typically be instantiated as behaviors with the option bound to their results. In some situations, an agent may immediately select an action to achieve this option in which case the option becomes the agent's immediate goal. In other situations, the agent enters into a volitional decision-making process. More generally, we can speak of *deliberation* (Sloman, 1999), a kind of cognitive process that builds upon reactive processes. For each option, a mental simulation will likely occur, resulting in new virtual events in LIDA's Current Situation Model. The total incentive salience of those virtual events will bias the results of deliberation. The LIDA conceptual model implements deliberation via ideomotor theory, which Franklin (2000) described in detail.

6.1 Volitional Decision Making in LIDA

What conditions lead to volitional decision making? We hypothesize that it may be initiated by a behavior with considerable activation, but not enough that it can be currently selected. There are a few main reasons why a behavior would fit this billing: 1) competition with other behaviors, 2) having an unsatisfied context or 3) a result with low incentive salience. In Maes' behavior network (1989), a competition situation arises between two behaviors when one undoes a precondition of the other. Alternatively, competition may arise when there is insufficient difference in behavior activation to prefer one over the other. For instance, evidence from the basal ganglia suggests that action selection performance improves when competing alternatives inhibit each other requiring an action to have a "statistically significant" greater activation than its competitors to be selected (Bogacz, Usher, Zhang, & McClelland, 2007). Alternatively, a behavior may have strong activation due to its underlying result and/or action event having positive incentive salience; however, its context might not be satisfied. Similarly, the context may be satisfied, but the result may have a high negative incentive salience. In each of these cases, a deliberative process may be initiated. What initiates such a process? In LIDA a competing behavior with an *internal* action enters into Action Selection with every instantiated behavior that was produced by an option or a drive. If this special behavior is selected, its internal action starts an ideomotor theory volitional process (see next section).

6.2 Ideomotor Theory in LIDA

William James first introduced the ideomotor theory of volition (James, 1890; Shin, Proctor, & Capaldi, 2010). James postulated proposers, objectors, and supporters as actors in the drama of

acting volitionally. He might have suggested the following scenario in the context of dealing with a feeling of thirst: The idea of drinking orange juice “pops into mind,” that is, propelled to consciousness by a proposer motivated by the drive feeling, “thirst,” and a “liking” for the sweetness of orange juice, it becomes the contents of consciousness. “No, it’s too sweet,” asserts an objector. “How about a beer?” says a different proposer. “Too early in the day,” says another objector. “Orange juice is more nutritious,” says a supporter. With no further objections, drinking orange juice is volitionally selected.

Baars incorporated ideomotor theory directly into his Global Workspace Theory (1988, Chapter 7). The LIDA model fleshes out volitional decision-making via ideomotor theory within Global Workspace Theory (Franklin, 2000) as follows. An idea “popping into mind” in the LIDA model is accomplished by the idea being part of the conscious broadcast of a cognitive cycle. Thus, the conscious broadcast implements the characters in James’ scenario, with some broadcasts acting as proposers, others as objectors, and others as supporters, the content of each “popping into mind” if it wins the competition, and is broadcast.

But, how does the thought “Let’s drink orange juice” come to consciousness? It must arise from an option, a virtual “drink-orange-juice” event node with positive total incentive salience in the Workspace. According to this motivational extension to LIDA, an option may arise from an event having nonzero base-level incentive salience, e.g., seeing a glass of one’s favorite orange juice. Alternatively, or additionally, an active drive feeling, e.g., thirst, may provide sufficient current incentive to the drink-orange-juice event node in PAM to instantiate it in the Workspace as an option. Like every higher-order cognitive process in the LIDA model, volitional decision-making occurs over multiple cycles. In this example, “drinking orange juice” may come to consciousness for the reasons just described. If this proposal fails to immediately lead to an external action selection, another proposal, or an objection, could come to consciousness. LIDA fleshes out ideomotor theory by the use of a “timekeeper” (Franklin, 2000) which as the result of an internal action, starts a timer as the result of a proposal, stops it when confronted by an objection, and restarts it after the occurrence of some support. Volitional decision-making ends when time runs out on some proposal (option), that is no further objection arises, and an action is thus selected.

The LIDA model hypothesizes that neural correlates of volitional decision-making include the ventromedial prefrontal cortex, the striatum, the basal ganglia and the ventral anterior cingulate cortex (MacDonald, 2008; Pasquereau et al., 2007; Rowe, Hughes, Eckstein, & Owen, 2008).

6.3 Proposers & Objectors from Mental Simulations

How might LIDA produce James’ proposers and objectors? During volitional decision-making, instantiated events in the Workspace, including options and drives, repeatedly cue LIDA’s perceptual and episodic memories. This cueing retrieves relevant local associations, and, drawing upon knowledge stored by the temporal links of event sequences, includes potential future events. Potential future events are grounded virtual representations of potential consequences of pursuing the options. Over multiple cognitive cycles, such virtual actions can produce an instantiated, “chain” of events. Such chains can be included in coalitions that may compete for consciousness. Critically, each event in the coalition contributes its total activation and total incentive salience. Recall that total incentive salience includes a base level component (updated by TD learning) and a current component (modulated by current drive feelings). Thus, the incentive salience of future

events, proposers and objectors in ideomotor theory, would contribute to a coalition's overall strength or salience. The net effect is to bias those behaviors or behavior streams believed to lead to a "desirable" future set of events. For instance, considering going outside in the rain unprepared may lead you to imagine getting wet, which leads you to grab an umbrella, or imagining being robbed on a walk in the city may compel you to take pepper spray along. Thus, appetitive or aversive behavior may also result from an anticipation, imagination, or recollection of an event with incentive salience.

6.4 Deliberation and Model-Based Control

It is our view that deliberation in general should involve the kind of mental simulations described in the previous section. For each alternative a mental simulation must occur with total incentive salience biasing the results of deliberation. Huys et al. (2012) have found that humans adopt a Pavlovian strategy for pruning decision trees. During mental evaluation of a sequence of choices, they curtailed any further evaluation of a sequence as soon as they encountered a large loss. In terms of the LIDA model high negative incentive salience may cut off one "branch" in a tree of possibilities. For example, while planning a route to the airport you realize the plan will involve taking the interstate, which you view as dangerous. This leads you to abandon that specific plan, and come up with another route taking safer roads. Conversely, sufficient positive incentive salience may prompt a decision, leading to the end of the deliberation altogether.

Recall that model-based control refers to a set of reinforcement learning methods in which an internal model of the environment is learned and used to evaluate available actions or cues on the basis of their potential outcomes (see Section 3.10). At this point it is clear that volitional decision making modulated by the total incentive salience of anticipated or potential events is an implementation of model-based control in the LIDA conceptual model's Current Situational Model. Deliberative cognitive processes are by nature slower than reactive ones. However, as we have suggested, deliberation in LIDA can produce model-based control, whereas reactive or consciously mediated action selection implements less normative model-free control.

7. Comparison between LIDA and Psi

As we mentioned in the introduction, the Psi model (Bach, 2009) provided inspiration for some of the motivational concepts introduced in this paper in the context of LIDA. Table 3 provides an overview of each Psi motivational construct and its corresponding instantiation in LIDA.

Psi distinguishes between demands, urges, motives, cognitive modulators, affect, and directed emotions. *Demands* are built-in drives that give rise to goals while *urges* are the signals that make demands apparent. In the LIDA conceptual model, *homeostatic sensors* and the sensors' associated *drive feeling nodes* implement drives. Psi theory suggests: "an abrupt increase of an urge corresponds to... a 'displeasure or distress signal'... while a decrease of an urge — its satisfaction — yields ... 'pleasure signals'" (Bach, 2012). In Psi, a *motive* is defined as the combination of an urge and an event associated with the satisfaction of the urge. LIDA refers to motives as events with nonzero incentive salience. Psi theory asserts that goals depend on the nature of a given situation and the existing demands.

Higher-level emotions such as jealousy and anger are said to arise in Psi from the combination of motives and the agent's perception of a situation. Fellous' work (2004) has been interpreted as providing a biological underpinning for the use of emotions in Psi (Lee-Johnson &

Carnegie, 2009). Psi maintains that emotions should be distinguished from motivational phenomena: while hunger is cognitively represented (as an urge) and implies a valence and an effect on the allocation of mental and physical resources of an organism, it is not an emotion, but a *motivator*.

Concept	Psi Term	LIDA Term
Built-in homeostatic mechanism	Demand (demand sensor)	Homeostatic sensor
Cognitive representation of a drive	Urge	Drive feeling node
Valence (Psychology)	Valence (kind of modulator)	Affective valence
Arousal	Arousal (kind of modulator)	Sum activation of feeling nodes
Desired (or aversive) event	Motive	Event with nonzero incentive salience
Emotion	Motives combined with perceptual representation	Feelings with cognitive content

Table 3. Approximate overlap of motivation concepts in the Psi and LIDA models.

Finally, Bach et al (Bach, 2012) adopts the general notion that motivation determines *what* has to be done, while emotion determines *how* it has to be done. In our view such an assertion is more problematic than elucidating. One difference between the extended LIDA model introduced above and the Psi model is the way drive feeling nodes (urges in Psi) affect perception, action selection and action execution. In Psi urges interact with perception via modulators (Bach, 2003, Fig. 1; 2012, Fig. 6). In LIDA feelings can immediately modulate perception and action selection, given that conscious learning has previously occurred. Some feeling nodes are likely part of most coalitions, and affect learning if they win the competition for consciousness.

Emotional valence is considered a representation of desirability (Bach, 2012), while in this model the two are distinguished, and play different functional roles (see sections “Affective Valence” and “Incentive Salience” above). Furthermore, it appears that in Psi, valence is a “reinforcement signal that emanates from changes in the perceived demands” (Bach, 2012). However, in our model, valence can also arise from intrinsically positive or negative stimuli, e.g., a sweet feeling node with positive valence. Lastly, the Psi model does not integrate cognitive neuroscience evidence for the “short route of emotion” (Cannon, 1927; Faghihi, Nkambou, Poirier, & Fournier-Viger, 2009; LeDoux, 2006), which have been referred to here as *alarms*.

8. Future Work

The motivational concepts and mechanisms outlined in this paper were implemented as part of the LIDA Framework (Snider, et al., 2011), a Java-based software framework that simplifies the creation of LIDA-based software agents; however, more computational experiments are needed to verify that these concepts and mechanisms accurately model human and animal behavior. McCall

(2014) made progress in this direction by replicating a reinforcer devaluation experiment (Gallagher, McMahan, & Schoenbaum, 1999). In particular, McCall simulated an experiment testing the effects of orbitofrontal lesions on the representation of incentive value in associative learning. The orbitofrontal cortex (OFC) is thought to contain representations of the motivational significance of cues (conditioned stimuli), and the incentive value of expected outcomes. The significance of the reinforcer devaluation task is that normal performance depends on the ability of a conditioned stimulus (CS) to gain access to the motivational properties of an upcoming unconditioned stimulus (US). Moreover, it illustrates the difference between the fast acting, slow-adapting model-free control (consciously mediated action selection in LIDA) and the slow-acting, fast-adapting model-based control (volitional decision-making in LIDA). In order to replicate this experiment, it was necessary to implement two LIDA agents—a “lesioned OFC” LIDA agent and an “intact OFC” LIDA agent—and the experimental environment in which the agents were situated. Both LIDA agents were implemented using the LIDA Framework, and the experimental environment was implemented as a separate framework module that was accessible to the agent through its sensors and effectors. McCall’s results were consistent with the original experiment, lending credibility to the general applicability of these motivational mechanisms.

Many opportunities exist for enhancing the learning mechanisms introduced in Section 5. For example, our temporal difference-based learning of base-level incentive salience assumed that the discount factor, γ , was fixed. However, a dynamic discount factor that is modulated by changes in the agent’s physiological state (for example, stress) may be more appropriate for cognitive modeling. Another avenue of future research could take advantage of *eligibility traces*, which have been shown to expedite temporal difference learning by assigning credit back further in time than the immediately previous event. Eligibility traces can be viewed as providing a short-term memory of multiple previous events so that several of these previous events may all be updated as each new observation arrives. Eligibility traces are usually implemented by an exponentially decaying memory trace, with decay parameter λ (Barto, 2007).

9. Conclusion

In this paper, we explored how the agenda of an autonomous agent could be specified in LIDA—a systems-level cognitive architecture. To this end, we presented extensions to LIDA’s conceptual model covering a range of motivation-related concepts, including feelings, affective valence, incentive salience, emotion, appraisal, reinforcement learning, and model-free and model-based learning. We discussed the role of motivations in a single LIDA cognitive cycle, as well as over multiple cycles. These new capabilities have been implemented in the LIDA Framework, and we refer interested readers to McCall (2014) for details on the application of these motivational mechanisms to the computational replication of a reinforcer devaluation experiment (Gallagher, McMahan, & Schoenbaum, 1999). Finally, we described the relationship between some of these motivational concepts and similar counterparts in the Psi model, which was used as a source of inspiration for the present work.

References

- Alvarado, N., Adams, S. S., & Burbeck, S. (2002). The role of emotion in an architecture of mind. *IBM Research*.
- Baars, B. (1988). *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.
- Baars, B., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Science*, 7, 166–172.
- Bach, J. (2003). *The micropsi agent architecture*. Paper presented at the Proceedings of ICCM-5, international conference on cognitive modeling, Bamberg, Germany.
- Bach, J. (2009). *Principles of Synthetic Intelligence: Psi: An Architecture of Motivated Cognition*. Oxford: Oxford University Press.
- Bach, J. (2012). Modeling Motivation and the Emergence of Affect in a Cognitive Agent *Theoretical Foundations of Artificial General Intelligence* (pp. 241-262): Springer.
- Barto, A. G. (2007). Temporal difference learning. *Scholarpedia*, 2(11), 1604.
- Belavkin, R. V. (2001a). Modelling the inverted-U effect with ACT-R. In Erik M. Altmann, Wayne D. Gray, A. Cleeremans & Christian D. Schunn (Eds.), *Proceedings of the 2001 Fourth International Conference on Cognitive Modeling* (pp. 296). Hillsdale, NJ Lawrence Erlbaum Associates.
- Belavkin, R. V. (2001b). *The role of emotion in problem solving*. Paper presented at the Proceedings of the AISB'01 Symposium on emotion, cognition and affective computing, Heslington, York, England.
- Berridge, K. C., & Kringelbach, M. L. (2008). Affective neuroscience of pleasure: reward in humans and animals. *Psychopharmacology*, 199(3), 457-480. doi: 10.1007/s00213-008-1099-6
- Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Research Reviews*, 28(3), 309-369.
- Bindra, D. (1978). How adaptive behavior is produced: a perceptual-motivational alternative to response reinforcements. *Behavioral and Brain Sciences*, 1(01), 41-52.
- Bogacz, R., Usher, M., Zhang, J., & McClelland, J. L. (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philos Trans R Soc Lond B Biol Sci*.
- Breazeal, C. (1998). *A Motivational System for Regulating Human-Robot Interaction*. Paper presented at the AAAI98, Madison, WI.
- Camras, L. A. (2011). Differentiation, dynamical integration and functional emotional development. *Emotion Review*, 3(2), 138-146.
- Cañamero, D. (1997). *Modeling motivations and emotions as a basis for intelligent behavior*. Paper presented at the Proceedings of the first international conference on Autonomous agents.
- Canamero, Lola D. (2003). Designing Emotions for Activity Selection in Autonomous Agents. In R. Trappl, P. Petta & S. Payr (Eds.), *Emotions in Humans and Artifacts* (pp. 115-148). Cambridge, MA: MIT Press.

- Cannon, W. B. (1927). The James-Lange theory of emotions: A critical examination and an alternative theory. *The American Journal of Psychology*, *39*(1/4), 106-124.
- Cannon, W. B. (1929). Organization For Physiological Homeostasis. *Physiol Rev.* , *9*, 399-431.
- Conway, M. (2001). Sensory–perceptual episodic memory and its context: autobiographical memory. *Philos. Trans. R. Soc. Lond B.*, *356*, 1375–1384.
- D'Mello, S., Ramamurthy, U., Negatu, A., & Franklin, S. (2006). A Procedural Learning Mechanism for Novel Skill Acquisition. In T. Kovacs & James A. R. Marshall (Eds.), *Proceeding of Adaptation in Artificial and Biological Systems, AISB'06* (Vol. 1, pp. 184–185). Bristol, England: Society for the Study of Artificial Intelligence and the Simulation of Behaviour.
- Damasio, A. (2003). *Looking for Spinoza: Joy, Sorrow and the Feeling Brain*. New York: Harcourt.
- Damasio, A. (1999). *The Feeling of What Happens*. New York: Harcourt Brace.
- Daw, N., Niv, Y., & Dayan, P. (2005). Actions, policies, values, and the basal ganglia. In E. Bezdard (Ed.), *Recent Breakthroughs in Basal Ganglia Research*.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. [Research Support, Non-U.S. Gov't]. *Nature Neuroscience*, *8*(12), 1704-1711. doi: 10.1038/nn1560
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, *10*, 204–211.
- Diener, E. (1999). Introduction to the special section on the structure of emotion. *Journal of personality and Social Psychology*, *76*(5), 803.
- Dijkstra, T. M. H., Schöner, G., & Gielen, C. C. A. M. (1994). Temporal stability of the action-perception cycle for postural control in a moving visual environment. *Experimental Brain Research*, *97*(3), 477-486.
- Dong, D., & Franklin, S. (2014). *Sensory Motor System: Modeling the process of action execution*. Paper presented at the Proceedings of the 36th Annual Conference of the Cognitive Science Society.
- Dong, D., & Franklin, S. (2015). A New Action Execution Module for the Learning Intelligent Distribution Agent (LIDA): The Sensory Motor System. *Cognitive Computation*. doi: 10.1007/s12559-015-9322-3.
- Dorner, D., & Hille, K. (1995). *Artificial souls: motivated emotional robots*. Paper presented at the IEEE International Conference on Systems, Man and Cybernetics, Vancouver, BC , Canada.
- Drescher, Gary L. (1991). *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. Cambridge, MA: MIT Press.
- Ekman, P., Sorenson, E. R., & Friesen, W. V. (1969). Pan-cultural elements in facial displays of emotion. *Science*, *164*(3875), 86-88.
- Faghihi, U., McCall, R., & Franklin, S. (2012). A Computational Model of Attentional Learning in a Cognitive Agent. *Biologically Inspired Cognitive Architectures*, *2*, 25-36.

- Faghihi, U., Estey, C., McCall, R., & Franklin, S. (2015). A Cognitive Model Fleshes Out Kahneman's Fast and Slow Systems. *Biologically Inspired Cognitive Architectures*, *11*, 38-52.
- Faghihi, U., Nkambou, R., Poirier, P., & Fournier-Viger, P. (2009). *Emotional Learning and a Combined Centralist-Peripheralist Based Architecture for a More Efficient Cognitive Agent*. Paper presented at the 7th IEEE International Conference on Industrial Technology (ICIT 2009).
- Fellous, J.-M. (2004). From human emotions to robot emotions. *Architectures for Modeling Emotion: Cross-Disciplinary Foundations*, American Association for Artificial Intelligence, 39-46.
- Fishbach, A., Roy, S. A., Bastianen, C., Miller, L. E., & Houk, J. C. (2005). Kinematic properties of on-line error corrections in the monkey. *Experimental Brain Research*, *164*(4), 442–457.
- Franklin, S. (1995). *Artificial Minds*. Cambridge, Ma: MIT Press.
- Franklin, S. (2000). Deliberation and Voluntary Action in 'Conscious' Software Agents. *Neural Network World*, *10*, 505–521
- Franklin, S. (2003). IDA: A Conscious Artifact? *Journal of Consciousness Studies*, *10*, 47–66.
- Franklin, S., & Baars, B. (2010). Two Varieties of Unconscious Processes. In E. Perry, D. Collerton, H. Ashton & F. LeBeau (Eds.), *New Horizons in the Neuroscience of Consciousness* (pp. 91–102). Amsterdam: John Benjamin.
- Franklin, S., Baars, B. J., Ramamurthy, U., & Ventura, M. (2005). The Role of Consciousness in Memory. *Brains, Minds and Media*, *1*, 1–38.
- Franklin, S., & Graesser, A. C. (1997). Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents *Intelligent Agents III* (pp. 21–35). Berlin: Springer Verlag.
- Franklin, S., Kelemen, A., & McCauley, L. (1998). IDA: A Cognitive Agent Architecture *IEEE Conf on Systems, Man and Cybernetics* (pp. 2646–2651). Menlo Park, CA: IEEE Press.
- Franklin, S., Madl, T., D'Mello, S., & Snaider, J. (2014). LIDA: A Systems-level Architecture for Cognition, Emotion, and Learning. *IEEE Transactions on Autonomous Mental Development*, *PP*(99), 1 doi: 10.1109/TAMD.2013.2277589
- Franklin, S., Madl, T., Strain, S., Faghihi, U., Dong, D., Kugele, S., . . . Chen, S. (2016). A LIDA cognitive model tutorial. *Biologically Inspired Cognitive Architectures*, 105-130. doi: 10.1016/j.bica.2016.04.003
- Franklin, S., & Ramamurthy, U. (2006). Motivations, Values and Emotions: Three sides of the same coin *Proceedings of the Sixth International Workshop on Epigenetic Robotics* (Vol. 128, pp. 41–48). Paris, France: Lund University Cognitive Studies.
- Franklin, S., Strain, S., Snaider, J., McCall, R., & Faghihi, U. (2012). Global Workspace Theory, its LIDA model and the underlying neuroscience. *Biologically Inspired Cognitive Architectures*, *1*, 32-43. doi: 10.1016/j.bica.2012.04.001
- Franklin, S., Strain, S., McCall, R., & Baars, B. (2013). Conceptual Commitments of the LIDA Model of Cognition. *Journal of Artificial General Intelligence*, *4*(2), 1-22. doi:10.2478/jagi-2013-0002

- Freeman, W. J. (2002). The limbic action-perception cycle controlling goal-directed animal behavior. *Neural Networks*, 3, 2249-2254.
- Fum, D., & Stocco, A. (2004). *Memory, Emotion, and Rationality: An ACT-R interpretation for Gambling Task results*. Paper presented at the ICCM.
- Fuster, J. M. (2004). Upper processing stages of the perception–action cycle. *Trends in Cognitive Sciences*, 8(4), 143-145.
- Gallagher, M., McMahan, R. W., & Schoenbaum, G. (1999). Orbitofrontal cortex and representation of incentive value in associative learning. *The Journal of neuroscience*, 19(15), 6610-6614.
- Gmytrasiewicz, P. J., & Lisetti, C. L. (2002). Emotions and personality in agent design and modeling *Game theory and decision theory in agent-based systems* (pp. 81-95): Springer.
- Hoffman, D. D., Singh, M., & Prakash, C. (2015). The interface theory of perception. *Psychonomic bulletin & review*, 22(6), 1480-1506.
- Hollerman, J., & Schultz, W. (1998). Dopamine Neurons Report an Error in the Temporal Prediction of Reward during Learning. *Nature Neuroscience*, 1, 304-309.
- Huys, Q. J., Eshel, N., O’Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Comput Biol*, 8(3), e1002410.
- James, W. (1884). II.—What is an emotion? *Mind*(34), 188-205.
- James, W. (1890). *The Principles of Psychology*. Cambridge, MA: Harvard University Press.
- Johnston, Victor S. (1999). *Why We Feel: The Science of Human Emotions*. Reading MA: Perseus Books.
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *The American economic review*, 93(5), 1449-1475.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kalis, A., Kaiser, S., & Mojzisch, A. (2013). Why we should talk about option generation in decision-making research. *Front. Psychol*, 4(555), 10.3389.
- Keller, L. R., & Ho, J. L. (1988). Decision problem structuring: Generating options. *Systems, Man and Cybernetics, IEEE Transactions on*, 18(5), 715-728.
- Klein, G., Wolf, S., Milotello, L., & Zsombok, C. (1995). Characteristics of skilled option generation in chess. *Organizational Behavior and Human Decision Processes*, 62(1), 63-69.
- Kringelbach, M. L., & Berridge, K. C. (2009). Towards a functional neuroanatomy of pleasure and happiness. *Trends in cognitive sciences*, 13(11), 479-487.
- Laird, John E., J, E., Newell, A., & Rosenbloom, Paul S. P. S. (1987). SOAR: An Architecture for General Intelligence. *Artificial Intelligence*, 33, 1–64.
- Lang, P. J., & Davis, M. (2006). Emotion, motivation, and the brain: Reflex foundations in animal and human research. In G. E. M. J. J. K. S. Anders & D. Wildgruber (Eds.), *Progress in Brain Research* (Vol. Volume 156, pp. 3-29): Elsevier.
- Lazarus, R. (1991). *Emotion and adaptation*. New York: Oxford University Press.
- LeDoux, J. E. (2006). Emotional Memory: In Search of Systems and Synapses. *Annals of the New York Academy of Sciences*, 702(1), 149-157.

- Lee-Johnson, C. P., & Carnegie, D. A. (2009). Robotic Emotions: Navigation with Feeling. In J. Vallverdú & D. Casacuberta (Eds.), *Handbook of Research on Synthetic Emotions and Sociable Robotics* (pp. 88-117): IGI Global.
- Liddell, B. J., Brown, K. J., Kemp, A. H., Barton, M. J., Das, P., Peduto, A., . . . Williams, L. M. (2005). A direct brainstem–amygdala–cortical 'alarm' system for subliminal signals of fear. *NeuroImage*, *24*(1), 235-243.
- Lucantonio, F., Stalnaker, T. A., Shaham, Y., Niv, Y., & Schoenbaum, G. (2012). The impact of orbitofrontal dysfunction on cocaine addiction. *Nature Neuroscience*, *15*(3), 358-366.
- MacDonald, K. (2008). Effortful Control, Explicit Processing and the Regulation of Human Evolved Predispositions. *Psychological Review*, *115*(4), 012–1031.
- Madl, T., Baars, B. J., & Franklin, S. (2011). The Timing of the Cognitive Cycle. *PLoS ONE*, *6*(4), e14803.
- Madl, T., & Franklin, S. (2012). A LIDA-based Model of the Attentional Blink. *Proceedings of the 11th International Conference on Cognitive Modelling*, 283-288.
- Madl, T., Franklin, S., Chen, K., & Trappl, R. (2013). Spatial Working Memory in the LIDA Cognitive Architecture. In R. West & T. Stewart (Eds.), *Proceedings of the 12th International Conference on Cognitive Modelling* (pp. 384-390). Ottawa, Canada: Carleton University.
- Maes, P. (1989). How to do the right thing. *Connection Science*, *1*, 291–323.
- Marieb, E. N., & Hoehn, K. (2007). *Human Anatomy & Physiology* (Seventh ed.). San Francisco, CA: Pearson Benjamin Cummings.
- Marinier, R., & Laird, J. E. (2008). Emotion-driven reinforcement learning. *Cognitive science*, 115-120.
- Marinier, R. P., Laird, J. E., & Lewis, R. L. (2009). A computational unification of cognitive behavior and emotion. *Cognitive Systems Research*, *10*(1), 48-69.
- McCall, R., Franklin, S., & Friedlander, D. (2010). *Grounded Event-Based and Modal Representations for Objects, Relations, Beliefs, Etc.* Paper presented at the FLAIRS-23, Daytona Beach, FL.
- McCall, R. J. (2014). *Fundamental motivation and perception for a systems-level cognitive architecture*. The University of Memphis.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Petersen, S. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of neuroscience*, *16*(5), 1936-1947.
- Negatu, A. (2006). *Cognitively Inspired Decision Making for Software Agents: Integrated Mechanisms for Action Selection, Expectation, Automatization and Non-Routine Problem Solving: Ph.D.* Dissertation, The University of Memphis, Memphis TN USA.
- Neisser, U. (1976). *Cognition and Reality: Principles and Implications of Cognitive Psychology* San Francisco: W. H. Freeman.

- O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, *38*(2), 329-337.
- Pasquereau, B., Nadjar, A., Arkadir, D., Bezdard, E., Goillandeau, M., Bioulac, B., . . . Boraud, T. (2007). Shaping of motor responses by incentive values through the basal ganglia. *Journal of Neuroscience*, *27*, 1176-1183.
- Phelps, E. A. (2006). Emotion and Cognition: Insights from Studies of the Human Amygdala. *Annual Review of Psychology*, *57*(1), 27-53. doi: doi:10.1146/annurev.psych.56.091103.070234
- Picard, R. (1997). *Affective Computing*. Cambridge MA: The MIT Press.
- Picard, R. W. (2003). Affective computing: challenges. *International Journal of Human-Computer Studies*, *59*(1,Äi2), 55-64. doi: 10.1016/s1071-5819(03)00052-1
- Purves, D., Brannon, E. M., Cabeza, R., Huettel, S. A., LaBar, K. S., Platt, M. L., & Woldorff, M. G. (2008). *Principles of cognitive neuroscience* (Vol. 83): Sinauer Associates Sunderland, MA.
- Raab, M., de Oliveira, R. F., & Heinen, T. (2009). How do people perceive and generate options? *Progress in brain research*, *174*, 49-59.
- Richard, J. M., & Berridge, K. C. (2011). Nucleus accumbens dopamine/glutamate interaction switches modes to generate desire versus dread: D1 alone for appetitive eating but D1 and D2 together for fear. *The Journal of neuroscience*, *31*(36), 12866-12879.
- Roseman, I. J., & Smith, C. A. (2001). Appraisal theory: Overview, assumptions, varieties, controversies *Appraisal processes in emotion: Theory, methods, research* (pp. 3-19). New York: Oxford University Press.
- Rowe, J., Hughes, L., Eckstein, D., & Owen, A. M. (2008). Rule-Selection and Action-Selection have a Shared Neuroanatomical Basis in the Human Prefrontal and Parietal Cortex. *Cerebral Cortex*, *18*, 2275--2285. doi: [10.1093/cercor/bhm249](https://doi.org/10.1093/cercor/bhm249)
- Schoenbaum, G., Takahashi, Y., Liu, T. L., & McDannald, M. A. (2011). Does the orbitofrontal cortex signal value? *Annals of the New York Academy of Sciences*, *1239*(1), 87-99.
- Shin, Y. K., Proctor, R. W., & Capaldi, E. J. (2010). A review of contemporary ideomotor theory. *Psychological Bulletin*, *136*(6), 943-974. doi: 10.1037/a0020541
- Sloman, A. (1998). Damasio, Descartes, Alarms and Meta-management *Proceedings Symposium on Cognitive Agents: Modeling Human Cognition*. San Diego: IEEE.
- Sloman, A. (1999). What Sort of Architecture is Required for a Human-like Agent? In M. Wooldridge & A. S. Rao (Eds.), *Foundations of Rational Agency* (pp. 35-52). Dordrecht, Netherlands: Kluwer Academic Publishers.
- Sloman, A., & Croucher, M. (1981). Why robots will have emotions.
- Smith, K. S., Berridge, K. C., & Aldridge, J. W. (2011). Disentangling pleasure from incentive salience and learning signals in brain reward circuitry. *Proceedings of the National Academy of Sciences*, *108*(27), E255-E264.
- Snaider, J., McCall, R., & Franklin, S. (2011). *The LIDA Framework as a General Tool for AGI*. Paper presented at the Artificial General Intelligence (AGI-11), Mountain View, CA.

- Squire, L. R., & Kandel, E. R. (2000). *Memory: From mind to molecules*: Macmillan.
- Sun, R. (2009). Motivational representations within a computational cognitive architecture. *Cognitive Computation*, *1*(1), 91-103.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Thompson, R. F., & Madigan, S. A. (2007). *Memory*. Princeton, NJ: Princeton University Press.
- Tindell, A. J., Smith, K. S., Berridge, K. C., & Aldridge, J. W. (2009). Dynamic computation of incentive salience: 'wanting' what was never 'liked'. *The Journal of Neuroscience*, *29*(39), 12220-12228.
- Toates, F. M. (1986). *Motivational systems*: CUP Archive.
- Ward, P., Suss, J., Eccles, D. W., Williams, A. M., & Harris, K. R. (2011). Skill-based differences in option generation in a complex task: A verbal protocol analysis. *Cognitive processing*, *12*(3), 289-300.
- Watkins, C. J. C. H. (1989). *Learning from Delayed Rewards*. Ph.D. thesis, Cambridge University.
- Westen, D. (1999). The Scientific Status of Unconscious Processes: Is Freud Really Dead? *Journal of the American Psychoanalytic Association*, *47*(4), 1061-1106. doi: 10.1177/000306519904700404
- Wimmer, G. E., & Shohamy, D. (2012). Preference by Association: How Memory Mechanisms in the Hippocampus Bias Decisions. *Science*, *338*(6104), 270-273. doi: 10.1126/science.1223252
- Yerkes, R. M., & Dodson, J. D. (1908). The Relationship of Strength of Stimulus to Rapidity of Habit Formation. *Journal of Comparative Neurology and Psychology*, *18*, 459-482.
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event Perception: A Mind-Brain Perspective. *Psychological Bulletin*, *133*(2), 273-293.
- Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological bulletin*, *127*(1), 3.