# The Mind According to LIDA – A Brief account

By Stan Franklin and the Cognitive Computing Research Group

## The LIDA Model and its Cognitive Cycle

The LIDA model is a systems-level, conceptual and computational model covering a large portion of human cognition[1]. Based primarily on Global Workspace Theory (Baars, 1997; Baars, Franklin, & Ramsøy, 2013; Baars, 1988), the model also implements and fleshes out a number of other psychological and neuropsychological theories including situated (embodied) cognition (de Vega, Glenberg, & Graesser, 2008; Glenberg & Robertson, 2000; Varela, Thompson, & Rosch, 1991), perceptual symbol systems (Barsalou, 1999; 2008), working memory (Baddeley & Hitch, 2007), memory by affordances (Glenberg, 1997), long-term working memory (Ericsson & Kintsch, 1995), and transient episodic memory (Conway, 2002a). Unlike most cognitive models, LIDA possesses an explicit attention mechanism, emotions that motivate action as well as modulate learning, and multiple distinct modes of learning including perceptual, transient episodic, declarative, attentional, procedural, and spatial memories. The LIDA computational architecture is derived from the LIDA cognitive model.

The LIDA model and its ensuing architecture are 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). More sophisticated agents, such as humans, make sense of the input from such sampling in order to facilitate their decision making. Neuroscientists and psychologists call this three part process the action-perception cycle (Dijkstra, Schöner, & Gielen, 1994; Freeman, 2002; Fuster, 2002; Neisser, 1976; Swenson & Turvey, 1991). The agent's "life" can be viewed as consisting of a continual sequence of these cognitive cycles. Each cycle constitutes a unit of sensing, attending and acting. A cognitive cycle can be thought of as a moment of cognition, a cognitive "moment." Higher-level cognitive processes are composed of many of these cognitive cycles, each a cognitive "atom."

We will now very briefly describe what the LIDA model hypothesizes as the rich inner structure of the LIDA cognitive cycle. More detailed descriptions are available elsewhere (Baars & Franklin, 2003; Franklin, Baars, Ramamurthy, & Ventura, 2005b; Franklin, Strain, Snaider, McCall, & Faghihi, 2012). During each cognitive cycle the LIDA agent first makes sense of its current situation as best as it can by updating its representation of its current situation, both external and internal. By a competitive process, as specified by Global Workspace Theory, it then decides what portion of the represented situation is most in need of attention. Broadcasting this

---

[1] "Cognition" is used here in a particularly broad sense, so as to include perception, feelings and emotions, and actions.

portion, the current contents of consciousness[2], enables the agent to chose an appropriate action and execute it, completing the cycle.

Thus, the LIDA cognitive cycle can be subdivided into three phases, the understanding phase, the attention (consciousness) phase, and the action selection and learning phase. Figure 1 should help the reader follow the description. It starts in the upper left corner and proceeds roughly clockwise. Beginning the understanding phase, incoming stimuli activate low-level feature detectors in Sensory Memory. The output is sent to Perceptual Associative Memory (recognition memory) where higher-level feature detectors feed in to more abstract entities such as objects, categories, actions, feelings, events, etc. The resulting percept moves to the Workspace where it cues both Transient Episodic Memory and Declarative Memory producing local associations. These local associations are combined with the percept to generate a Current Situational Model; the agent's understanding of what is going on right now.

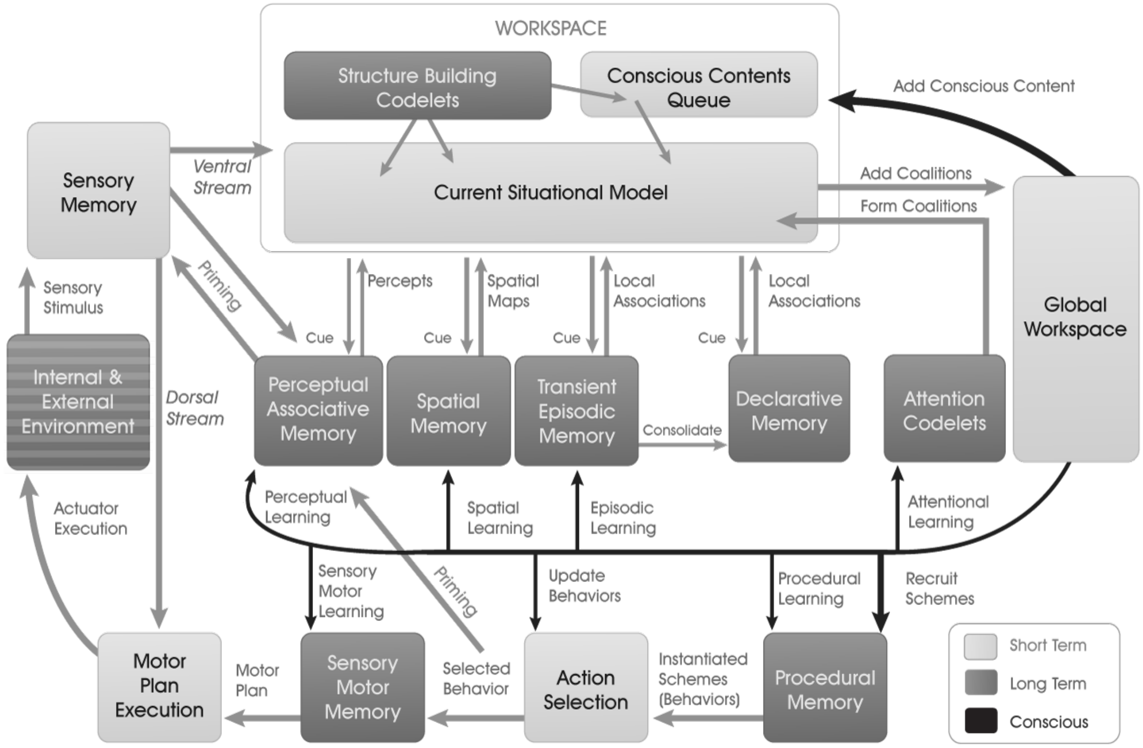

Figure 1. Cognitive Cycle Diagram.

Attention Codelets[3] begin the attention phase by forming coalitions of selected portions of the Current Situational Model and moving them to the Global Workspace. A competition in the Global Workspace then selects the most salient coalition, e.g., the one with the most relevant, important, urgent, novel, unexpected, loud, bright, moving, etc. content, whose contents become the content of consciousness. These conscious contents are then broadcast globally, initiating the action selection phase.

---

[2] Here "consciousness" refers to functional consciousness (Franklin 2003). We take no position on the need for, or possibility of, phenomenal consciousness.

[3] A codelet is a small piece of code that performs a specific task in an independent way. It could be interpreted as a small part of a bigger process, similar to an ant in an ant colony.

The action selection phase of LIDA's cognitive cycle is also a learning phase in which several processes operate in parallel (see Figure 1). New entities and associations, and the reinforcement of old ones, occur as the conscious broadcast reaches Perceptual Associative Memory. Cognitive maps are created or updated in Spatial Memory. Events from the conscious broadcast are encoded as new memories in Transient Episodic Memory. Possible action schemes, together with their contexts and expected results, are learned into Procedural Memory from the conscious broadcast. Older schemes are reinforced. In parallel with all this learning, and using the conscious contents, possible action schemes are recruited from Procedural Memory. A copy of each such is instantiated with its variables bound, and sent to Action Selection, where it competes to be the behavior selected for this cognitive cycle. The selected behavior triggers Sensory-Motor Memory to produce a suitable motor plan for the execution of the behavior. Its execution completes the cognitive cycle.

The Workspace requires further explanation. Its internal structure includes the Current Situational Model and the Conscious Contents Queue. The Current Situational Model is where the structures representing the actual current internal and external events are stored. Structure building codelets are responsible for the creation of these structures using elements from the various sub-modules of the Workspace. The conscious contents queue holds the contents of the last several broadcasts and permits LIDA, using codelets, to understand and operate upon time related concepts (Snaider, McCall, & Franklin, 2009).

## LIDA's Memory Modules and their Processes

The LIDA cognitive cycle emerges as an apparent sequence in the temporal organization of the interaction between LIDA's processes. However, the processes themselves are asynchronous in the sense of computer science in that they unfold continually and independently of one another. Before detailing these processes, it is crucial to review the nature of the representations on which these processes operate.

These representations are managed by memory modules (see Figure 2). At the most general conceptual level, memory refers to our capacity for encoding information, storing and retaining it, and perhaps subsequently retrieving it as needed. From a computational perspective, memory can be thought of as a database of information together with processes for encoding and storing new items, and for retrieving existing items.

Humans and other animals regularly employ several different memory systems during their cognitive processing. There is relatively little agreement among scientists as to how to partition memory into various systems, and how to name these systems. We wish to sidestep this debate, and include definitions of terms, not to impose them on others, but to assist the reader in understanding how they are used here. The diagram in Figure 2 contains our current nomenclature. Most leafs of that hierarchy correspond to memory modules that are part of the LIDA Model. The diagram is organized from left to right beginning with very short-term memory and increasing with retention time. We will briefly describe the several memory systems following the diagram from left to right.

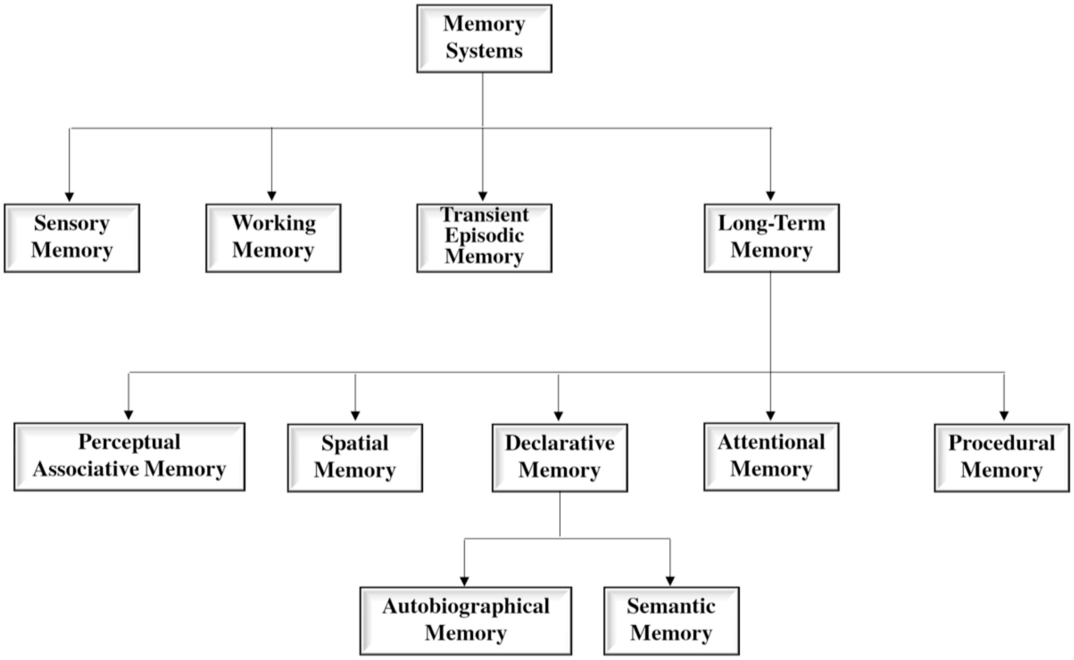

Figure 2. Memory Systems

Sensory Memory allows an animal to retain impressions of sensory data detected by the sensory receptors after the original stimulus has ceased. Sensory Memory has a large capacity for as yet only slightly processed data but can only hold accurate images of visual scenes or other sensory modalities momentarily, that is, for a few tens of milliseconds.

Working Memory (Baddeley & Hitch, 1974) refers to the structures and processes used for producing information from sensory and perceptual data, and temporarily storing and manipulating that information. Working Memory decays in a few tens of seconds.

Episodic Memory records, or points to, the what, where, and when of unique personal experiences (Baddeley, Conway, & Aggleton, 2001). More specifically, such individual memories include events, times, places, actions, associated feelings and emotions, etc. Episodic Memory comes in two versions, Transient Episodic Memory (Conway, 2002b) and long-term episodic memory (also called Declarative Memory). Transient Episodic Memory is so named because it typically decays within a few hours or a day. Where I parked my car, or what I had for breakfast might be recalled from Transient Episodic Memory, while where I parked my car last Tuesday, or what I had for breakfast that day could not.

Long-term memories (like Declarative Memory) may last for decades or a lifetime. On the other hand, if not reinforced or not initially provided with a strong emotional charge, they may decay away relatively quickly[4]. The diagram of Figure 2 classifies long-term memories into Perceptual Associative Memory, Spatial Memory, Declarative Memory (including Autobiographical Memory and Semantic Memory), Attentional Memory, and Procedural Memory.

Perceptual Associative Memory is often referred to as recognition memory (Neath & Surprenant, 2003). We may think of it as providing the ability to recognize elements in the surrounding environment such as objects or faces or places or actions or events. Humans and other animals need to recognize individuals, such as family members or conspecifics, and to be able to navigate from one place to another by recognizing landmarks. Prey animals need to be able to tell when something new is present in their environment, or to be able to learn about new sources of food. Perceptual Associative Memory seems to be ubiquitous in the animal kingdom.

Declarative Memory is a common term for long-term episodic memory. It comes in two forms: Autobiographical Memory, the memory for individual events, and Semantic Memory, the memory for facts divorced from the place and time they were acquired (Sims & Gray, 2004).

Procedural Memory is usually defined as the long-term memory of skills and procedures, or of "how to" knowledge. Here we use the term somewhat differently to refer to memory of what actions (skills) are useful under what circumstances to produce what results. In the LIDA Model, action selection is distinguished from action execution. Procedural Memory, as we use the term, is concerned with action selection, as we will see below.

Attentional learning is learning what to attend to (Kruschke, 2010), that is, learning what kind of events are relevant in a particular situation and should be attended to. Attentional learning can bias sensory competition for visual attention (Vidnyánszky & Sohn, 2003).

Containing a hierarchy of cognitive maps, Spatial Memory is responsible for the acquisition, organization, utilization and revision of knowledge of spatial environments (Hirtle & Jonides, 1985; Jacobs & Schenk, 2003; Manns & Eichenbaum, 2009). Each cognitive map is a rough, allocentric topographical map representing objects of a physical environment, including for each object its identity, location, shape, size and orientation, as well as the direction and distance between any two objects. Spatial Memory may also employ egocentric maps.

## Inside LIDA’s Cognitive Cycle

As was said previously, LIDA’s cognitive cycle emerges as an apparent sequence in the temporal organization of the interaction between LIDA’s processes. However, the processes themselves are asynchronous in that they unfold continually and independently of one another, and they are local in that, except for the broadcast from the Global Workspace, they affect only one module, or, perhaps, two. There’s only the one global process, the broadcast.

---

[4] This assertion is still controversial. Some memory researchers maintain that long-term memory traces remain for a lifetime, but mechanisms for recalling them are lost. For a recent review, see (Cansino, 2009).

The cycle begins with sampling the environment, either external or internal to the agent. Low-level feature detectors in Sensory Memory begin the process of making sense of the incoming stimuli. These low-level features are passed on from Sensory Memory to Perceptual Associative Memory, or directly to the Workspace, where higher-level features, such as objects, categories, relations, actions, events, feelings, etc. are recognized. These entities, which have been recognized preconsciously, make up the percept that is passed to the Workspace, where a small model of the agent's current situation is assembled at a pre-conscious level[5]. This percept and other Workspace contents serve to cue the two forms of episodic memory, transient and declarative. Responses to the cue are recalled into the Workspace in the form of local associations; that is, remembered events from these two memory systems that were associated with the various elements of the cue. In addition to the current percept, the Workspace contains recent percepts and other items assembled from previous cycles that haven't yet decayed away.

Many processes in LIDA are implemented as codelets.[6] Codelets are small code sequences with very circumscribed objectives. They operate independently and asynchronously; they can also organize into cooperative or competitive structures such as coalitions or behavior streams (see below).

A new model of the agent's current situation, the Current Situational Model, is assembled (updated) from the current percepts from Perceptual Associative Memory, the local associations retrieved from Transient Episodic Memory and Declarative Memory, a current cognitive map from Spatial Memory, and the undecayed parts of the previous model present in the Workspace. This assembling process will typically require structure-building codelets, each of which has some particular type of structure it is designed to build. To fulfill their task these codelets may draw downhill upon Perceptual Associative Memory, and even Sensory Memory, to enable the recognition of relations and situations. The newly updated Current Situational Model constitutes the agent's understanding of its current situation (writ small[7]) within its world. It has made sense of the incoming stimuli.

For an agent "living" in a complex, dynamically changing environment, this Current Situational Model may well be much too much for the agent to consider all at once in deciding what to do next. It needs to select a portion of the model to be attended to. Which are the most relevant, important, urgent or insistent structures within the model? Which are novel, unexpected, bright, loud or moving? Attention codelets claim portions of the Current Situational Model matching their concerns so as to form coalitions, which are moved into the Global Workspace to compete for consciousness. One of these coalitions wins the competition. In effect, the agent has decided on what to attend. The purpose of all this processing is to help the agent decide what to do next. To this end, a representation of the contents of the winning

---

[5] LIDA's Workspace plays something of the role of working memory, but all the operations on Baddeley's working memory are conscious (1992).

[6] The codelet concept is essentially the same as Baars' (psychological) processors (1988), Minsky's agents (1985), Jackson's demons (1987), or Ornstein's small minds (1986). The term was borrowed from Hofstadter and Mitchell (1995).

[7] An agent's current situation during a single cognitive cycle (~300 ms) will be quite small. This writer's at a recent moment likely consisted of little more than a view of a piece of a computer monitor, the sound of furnace fan, and the feel of a chair, not all conscious.

coalition is broadcast globally from the Global Workspace to numerous recipients in the cognitive system (See Figure 1).[8] First, this global broadcast transfers conscious information back to the Conscious Contents Queue in the Workspace where it remains available to structure building codelets for a few tens of cognitive cycles. Second, it provides content for perceptual, spatial, procedural, attentional and episodic learning (see Figure 1). Third, it sends the conscious broadcast to Procedural Memory, to which we turn next.

Though the contents of this conscious broadcast are available globally, a primary recipient is Procedural Memory, which stores templates, called schemes, of possible actions. A scheme has four components: the action to be performed, a context in which it might be triggered, an expected result of the action if performed in that context, and the likelihood that the action would produce the expected result in that context. An action can be of three types: it can be a primitive action, a set of actions happening in parallel, or a stream of actions (i.e., a set of actions whose organization is more complex). Schemes whose contexts or results intersect with the contents of the conscious broadcast instantiate copies of themselves to the Action Selection mechanism. Copies are instantiated in the sense that they become specific to the current situation (i.e., instant) being considered. Instantiated schemes not yet decayed away from previous cycles may also continue to be available for selection. The Action Selection mechanism then chooses a single action from one of these instantiations. The chosen action is sent to Sensory-Motor Memory, where it picks up an appropriate skill (procedure, algorithm, motor plan) by which it is then executed. The executed action affects the environment, and the cycle is complete.

The LIDA Model hypothesizes that all human cognitive processing is via a continuing iteration of such cognitive cycles. These cycles emerge as an interleaved sequence, with each cognitive cycle taking roughly 300-600 milliseconds (Koivisto & Revonsuo, 2010; Madl, Baars, & Franklin, 2011). The cycles cascade; that is, overlapping cycles have processes that run in parallel. This cascading must, however, respect the way consciousness processes information serially in order to maintain the stable, coherent image of the world with which consciousness endows us (Franklin, 2005b; Merker, 2005). It must also respect the seriality of action selection. This cascading allows a rate of cycling in humans of five to ten cycles per second. A cognitive "moment" is thus quite short! For instance, for cascading cycles of 300 milliseconds each offset by 100ms, cycles would occur at ~10 Hz. There is considerable empirical evidence from neuroscience suggestive of and consistent with such cognitive cycling in humans (Doesburg, Green, McDonald, & Ward, 2009b; Massimini et al., 2005; Sigman & Dehaene, 2006; Uchida, Kepecs, & Mainen, 2006; Willis & Todorov, 2006).

## Feelings & Emotions in the LIDA Model

Feelings in humans include hunger, thirst, various sorts of pain, hot or cold, the urge to urinate, tiredness, depression, etc. Damasio (1999) views feelings as somatic markers (i.e., locations in the body). An individual perceives feelings in the body.

---

[8] In brains, the Global Workspace's broadcast results from the ignition of a neuronal assembly into a thalamocortical core, in the sense of Edelman & Tononi (2000). (Baars, et al., 2013)

Implemented biologically as somatic markers, feelings typically attach to response options and, so, bias the agent's choice of action.

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 cognitive content. Similarly, one must be angry at someone for something, that someone and something being the cognitive content. 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.

Feelings and emotions give us the ability to make an almost immediate assessment of situations. They allow us to determine whether a given state of the world is beneficial or detrimental. For humans, emotions are the result of millions of years of evolution; a blind trial and error process that has given us default responses to common experiences. Unlike a reflexive action alone, however, feelings and emotions temper our responses to the situation at hand. Simple though such a response may be, it allows us to adapt to a new situation in a quick and non-computationally intensive way. Our lives as humans are filled, moment-to-moment with the complex interplay of emotional stimuli both from the external world and from our internal selves (Damasio, 1994).

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 (Franklin & Ramamurthy, 2006).

Every autonomous agent must be equipped with primitive motivators, drives that motivate its selection of actions. In humans, in animals, and in the LIDA model, these drives are implemented by feelings (Franklin & Ramamurthy, 2006). Such feelings implicitly give rise to values that serve to motivate action selection. This section is devoted to an explication of how feelings are represented in the LIDA model, the role they play in attention, and how they act as motivators, implicitly implementing values. (Feelings also act as modulators to learning: see below.) Reference to Figure 1 will prove helpful to the reader.

Every feeling has a valence, positive or negative, associated with pleasure or pain. Also, each feeling must have its own identity; we distinguish between the pains of a pinprick, a burn or an insult. From a computational perspective it makes sense to represent the valence of a single feeling as either positive or negative, that is, as greater or less than zero. In biological creatures, feelings typically have only positive valence or negative valence (Heilman, 1997). For example, the feeling of distress at having to over-extend holding one's breath at the end of a deep dive is a different feeling from the relief that ensues with the taking of that first breath. Such distress is implemented with varying degrees of negative valence, and the relief with varying positive valence. Each has its own identity. However, multiple feelings can be present simultaneously,

Feelings are represented in the LIDA Model as nodes in its Perceptual Associative Memory. Each node constitutes its own identity, for example, distress at not enough oxygen is represented by one node, relief at taking a breath by another. Each feeling node has its own valence, always positive or always negative. The current activation of the node measures the arousal. Those feeling nodes with sufficient activations, along with their incoming links and object nodes, become part of the current percept and are passed to the Workspace.

Like other Workspace structures, feeling nodes help to cue the two episodic memories. The resulting local associations may also contain feeling nodes associated with memories of past events. These feeling nodes play a major role in the assigning of activation to coalitions of information to which they belong, helping them to compete for attention within the Global Workspace. Any feeling nodes that belong to the winning coalition become part of the conscious broadcast, the contents of consciousness. Thus the LIDA agent becomes conscious of those feelings.

Any feeling node in the conscious broadcast that also occurs in the context of a scheme in procedural memory adds to the current activation of that scheme, increasing the likelihood of it instantiating a copy of itself into the Action Selection mechanism It is here that feelings play their first role as implementation of motivation by adding to the likelihood of a particular action being selected. That feeling in the context of the scheme implicitly increases the value of the result of taking that scheme's action.

In the Action Selection mechanism, the activation of a particular scheme, and thus its ability to compete for selection and execution, depends upon several factors. These factors include how well the context specified by the scheme agrees with the current and very recently past contents of consciousness, that is, with the current situation. As mentioned earlier, the activation of this newly arriving behavior also depends on the presence of feeling nodes in its context, and their activation as part of the conscious broadcasts. Thus feelings contribute motivation for taking action to the activation of newly arriving behavior schemes.

On the basis of the resulting activation values, a single instantiated scheme (behavior) is chosen by the Action Selection mechanism. The action ensuing from this selected behavior represents the agent's current intention in the sense of Freeman (1999) (p. 96), that is, what the agent intents to do next. The expected result of that scheme can be said to be the agent's current goal. Note that the selection of this action was affected by its relevance to the current situation (the environment), the nature and degree of associated feelings (the drives), and its relation to other actions, some of these being prerequisite for the action.

The selected instantiated scheme, including its feelings, is then passed to Sensory-Motor Memory for execution. There the feelings modulate the execution of the action (Zhu & Thagard, 2002). A feeling in the conscious broadcast emitted by the Global Workspace also plays its role in modulating the various forms of learning.

## Learning in LIDA

Edelman (1987) usefully distinguishes two forms of learning. Selectionist learning requires selection from a redundant repertoire that is typically organized by some form of reinforcement learning. A repertoire of, say, possible actions, is redundant if slightly different actions can lead to roughly the same result. In reinforcement learning (Kaelbling, Littman, & Moore, 1996) a successfully executed action is reinforced, making it more likely to be chosen the next time the result in question is needed. In Edelman's system little used actions tend to decay away. In contrast, instructional learning allows the learning of new representations of, say, new actions, that is, actions not currently in the repertoire.

Global Workspace Theory postulates that learning requires only attention (Baars, 1988 pp. 213-218). In the LIDA Model this implies that learning must occur with each cognitive cycle, because whatever enters consciousness is being attended to. More specifically, learning occurs with the conscious broadcast from the Global Workspace during each cycle.

Learning is both a function of attention and of affect. Feelings in the conscious broadcast modulate learning. Up to a point, the higher the affect (arousal of a feeling) the greater the learning. Beyond that point, more affect begins to interfere with learning. Thus learning rate in LIDA varies with affect according to the venerable Yerkes-Dodson law (Diamond, 2005; Yerkes & Dodson, 1908), graphed as an inverted U-curve.

Learning in the LIDA Model follows the tried and true artificial intelligence principle of generate and test. New representations are learned in a profligate manner (the generation) during each cognitive cycle. Those that are not sufficiently reinforced during subsequent cycles (the test) decay away. Three modes of learning – perceptual, episodic and procedural – employing distinct mechanisms (Franklin, Baars, Ramamurthy, & Ventura, 2005a; Nadel, 1992), have been designed for the LIDA Model and are in various stages of implementation. The design of a fourth mode, attentional learning, is underway but not yet completed. Yet another mode, spatial learning, is currently being designed with mechanisms inspired by empirical neuroscience studies. Two other forms of learning, the learning of structure building codelets and the learning of action execution skills, are also just being considered.

These different modes of learning, each feeding into its own memory module, correspond to memory modules identified in humans. Each memory module likely requires its unique form of data representation, and unique processes for encoding, retrieving, etc. Most of these modes of learning into memory modules will be described in the following paragraphs.

Perceptual learning enables an agent to recognize features, objects, categories feelings, actions, relations, events, situations, etc. In the LIDA Model what is learned perceptually is stored in Perceptual Associative Memory (Franklin, 2005a, 2005c; McCall, Franklin, & Friedlander, 2010). Motivated by the Slipnet from the Copycat architecture (Hofstadter & Mitchell, 1995), the LIDA Perceptual Associative Memory is implemented as a collection of nodes and links with activation passing between the nodes. Nodes represent features, individuals, categories, actions, feelings, events, and more complex structures. Links, both excitatory and inhibitory, represent roles, relations, etc. Each node and link has both a current and a base-level activation. The base-level activation measures how useful the node or link has been in the past, while the current activation depends on its relevance in the current situation. The percept passed on to the Workspace during each cognitive cycle is composed of those nodes and links whose total activation is over its threshold. Perceptual learning in its selectionist form modifies base-level activation, and in its instructionalist form creates new nodes and links in Perceptual Associative Memory. One or the other or both may occur with the conscious broadcast during each cognitive cycle.

Episodic learning refers to the memorization of events – the what, the where and the when (Baddeley, et al., 2001; Tulving, 1983). In the LIDA Model such learned events are stored in Transient Episodic Memory (Conway, 2002b; Franklin, et al., 2005a), and in the longer-term Declarative Memory (6) (Franklin, et al., 2005a). Both are implemented using sparse distributed memory (Kanerva, 1988), which is both associative and content addressable, and has other desirable psychological properties. In particular, it knows when it doesn't know, and exhibits the tip of the tongue phenomenon. Episodic learning in the LIDA model (Franklin, et al., 2005a; Ramamurthy, D'Mello, & Franklin, 2005; Ramamurthy, D'Mello, & Franklin, 2004) is also a matter of generate and test, with such learning occurring at each conscious broadcast. Episodic learning is initially directed only to Transient Episodic Memory. At a later time and offline, the undecayed contents of Transient Episodic Memory are consolidated (Nadel & Moscovitch, 1997; Stickgold & Walker, 2005) into Declarative Memory, where they still may decay away or may last a lifetime depending on subsequent reinforcement or the lack thereof.

Procedural learning refers to the learning of the possible selection of new tasks to be accomplished and the improvement of the selection of old tasks. In the LIDA model such learning is accomplished in Procedural Memory (D'Mello, Ramamurthy, Negatu, & Franklin, 2006), which is implemented via a scheme net motivated by Drescher's schema mechanism (1991). Each scheme in Procedural Memory is a template for an action, consisting of a context, an action and a result, together with a base-level activation intended to measure how likely the result would be to occur were the action taken within its specific context. Once again, the LIDA Model's procedural learning is via a generate and test mechanism, using base-level activation as reinforcement, as well as through the creation of new schemes. These new schemes can support multiple actions, both parallel and sequential.

Attentional learning, that is, the learning of what to attend to (Estes, 1993; Vidnyánszky & Sohn, 2003) has been relatively little studied by neuroscientists or cognitive scientists but see (Kruschke, 2003; Yoshida & Smith, 2003). To our knowledge it has been totally ignored by AI researchers, no doubt because few of their systems contain mechanisms for both attention and learning. In the LIDA Model attentional learning would involve attention codelets (see Figure 4), small processes whose job it is to focus the agent's attention on some particular portion of its Current Situational Model. When designed, we envision the LIDA Model's attentional learning mechanism involving modulating the base-level activation of attention codelets, as well as the creation of new ones.

Unlike attentional learning, spatial learning has been much studied by neuroscientists who have described mechanisms for producing and updating cognitive maps using subiculum boundary vector cells (Lever, Burton, Jeewajee, O'Keefe, & Burgess, 2009), entorhinal border cells (Solstad, Boccara, Kropff, Moser, & Moser, 2008), entorhinal head direction cells (Rolls & Stringer, 2005), entorhinal grid cells (Hayman, Verriotis, Jovalekic, Fenton, & Jeffery, 2011), hippocampal place cells (Moser, Kropff, & Moser, 2008), (para)hippocampal spatial view cells (Stringer, Rolls, & Trappenberg, 2005). Implementing spatial learning in the LIDA Model will involve designing computational mechanisms inspired by the corresponding neuroscience. Early attempts in this direction have been made by others (e.g.,(O'Keefe, 1990)).

## Higher-level Cognitive Processes in the LIDA Model

Higher-level cognitive processing in humans might include imagination, deliberation, volitional decision-making, metacognition, reasoning, planning, scheduling, problem solving, understanding language, and language generation. In the LIDA model such higher-level processes require multiple cognitive cycles. Every higher-level cognitive process can be implemented by one or more behavior streams[9], that is, streams of instantiated schemes (behaviors) and links from Procedural Memory.

Cognitive processes have differing levels of control. Sloman (1999) distinguishes three levels that can be implemented by the architecture of an autonomous agent – the reactive, the deliberative, and the metacognitive[10]. The reactive, is the level we would typically expect of many insects, that is, a relatively direct connection between incoming sensory data and the outgoing actions of effectors. The key point is the relatively direct triggering of an action once the appropriate environmental situation occurs. Though direct, such a connection can be almost arbitrarily intricate, requiring quite complex algorithms to implement in an artificial agent. In the LIDA model we refer to such actions as being consciously mediated, that is using conscious information via a never conscious process (Franklin & Baars, 2010). Thus the action selected at the end of any single cognitive cycle is done so at the consciously mediated level, though other reactive selection of actions may require several cognitive cycles.

The reactive level is perhaps best defined by what it is not. "What a purely reactive system cannot do is explicitly construct representations of alternative possible actions, evaluate them and choose between them, all in advance of performing them." (Sloman, 1999). Reactive control alone is particularly suitable for agents occupying relatively simple niches in reasonably stable environments, that is, for agents requiring relatively little flexibility in their action selection. Such purely reactive agents typically require relatively few higher-level, multi-cyclic cognitive processes.

On the other hand, deliberative control typically employs such higher-level cognitive processes as reasoning, planning, scheduling and problem solving. Such deliberative processes in humans, and in some other animals[11], are typically performed in an internally constructed virtual reality. Such deliberative information processing and decision-making allows an agent to function more flexibly within a complicated niche in a complex, dynamic environment. An internal virtual reality for deliberation requires a short-term memory in which temporary structures can be constructed with which to "mentally" try out possible actions without actually executing them. In the LIDA Model the Workspace serves just such a function. It is essentially a preconscious working memory in the sense of Baddeley (1992). The action selected during several cognitive cycles may consist of building, or adding to,

---

[9] A behavior stream is a sequence schemes with its order only partially specified. Some actions in a stream may be taken in either order.

[10] Sloman speaks of meta-management rather than metacognition. We prefer the more common psychological term.

[11] Deliberation has been demonstrated in apes (Mulcahy & Call, 2006), birds (Werdenich & Huber, 2006), and even in arachnids (Tarsitano, 2006; Wilcox & Jackson, 2002).

some representational structures in the Workspace during the process of some sort of deliberation. Structure-building codelets, the sub-processes that create such structures, modify or compare them, etc., are typically implemented as internal reactive processes. Deliberation builds on reaction. In the LIDA Model, deliberation is implemented as a collection of behavior streams, each selected action (behavior) of which is an internal reactive process (Franklin, 2000).

As deliberation builds on reactions, metacognition[12] typically builds on deliberation. Often defined as "thinking about thinking," metacognition in humans and animals (Smith & Washburn, 2005) involves monitoring deliberative processes, allocating cognitive resources, and regulating cognitive strategies (Flavell, 1979). As are other higher-level cognitive processes, metacognition in LIDA can be implemented by a collection of appropriate behavior streams, each with its own metacognitive task. Metacognitive control adds yet another level of flexibility to an agent's decision making, allowing it to function effectively in an even more complex and dynamically changing environmental niche.

## The LIDA Computational Framework

The LIDA Framework is a generic and customizable computational implementation of much of the LIDA model, programmed in Java (Snaider, McCall, & Franklin, 2011). Its primary goal is to provide a generic implementation of the LIDA model, easily customizable for specific problem domains, so as to allow for the relatively rapid development of LIDA controlled software agents and/or robots.

The Framework is intended to be ready customizable at several levels depending upon the required functionality. At the most basic level, developers can use an XML file to customize their applications. Several small pieces in the Framework can also be customized by implementing particular versions of them. For example, new strategies for decaying activations or types of codelets can be implemented. Finally, more advanced users can also customize and change the internal implementation of whole modules. In each case, the Framework provides default implementations that greatly simplify the customization process.

The main components of the Framework interconnect elements that represent modules in the LIDA model. Each main component of the LIDA cognitive model has a corresponding module in the Framework. For example, the Sensory Memory, Workspace and Action Selection are all modules in the Framework. Each module has its own application programming interface (API) that defines the functionality for this particular module. Modules can have submodules, that is, module nested inside another module. For example, the Workspace has several submodules, such as the Current Situational Model submodule.

Most modules in the Framework are domain independent. For each of these modules, the Framework provides a default implementation. For example, the Transient Episodic Memory is implemented as sparse distributed memory (Kanerva, 1988) and the Action Selection Module as a behavior net (Maes, 1989). However,

---

[12] An early implementation of metacognition in the "Conscious" Mattie software agent (Zhang, Franklin, & Dasgupta, 1998) shows that metacognition can occur in the absence of deliberation, that is, in an otherwise reactive agent. Modeled after Minsky's B-brain (Minsky, 1985), this metacognitive system was implemented independently of the underlying architecture.

some modules must be domain specific. In particular, Sensory Memory and Sensory-Motor Memory have to be specified on the basis of the domain that the Framework is being applied to. Nevertheless, the Framework supplies base implementations from whence the developer can implement domain specific functionality.

Modules need to perform several tasks in order to achieve their specific functionalities. The Framework provides Tasks, encapsulations of small processes. A module can create several Tasks to help it perform its function. A Task can run one time or repeatedly. A Task that passes activation is an example of the former, while a structure-building codelet is an example of the latter. The Task Manager controls the execution of all Tasks in the Framework. Tasks can be executed on separate threads by the Task Manager, achieving parallel execution in a way that is approximately transparent to the user.

Summing up, the Framework allows the creation of new applications and simulations based on the LIDA model. Its design and implementation aims at simplifying this process by permitting the developer to concentrate on the specifics of the application, while hiding the complexities of the generic parts of the model. Use of the Framework also enforces good software practices that simplify the creation of complex architectures. It achieves a high level of abstraction permitting several ways and levels of customization with a low level of coupling among modules. Supplemental tools like a complete GUI and logging support are also provided. The result is a powerful and customizable tool with which to develop LIDA controlled software agents and robots.

## LIDA-Based Software Agents

We have developed several cognitive software agents that replicate experiment data from human subjects (Faghihi, McCall, & Franklin, 2012; Madl, et al., 2011; Madl & Franklin, 2012) in order to show how the computational LIDA architecture can model human cognition in basic psychological tasks. Our main goals with these agents were to substantiate some of the claims of the LIDA model, and to take a first step towards identifying a set of internal parameters. Ideally, these internal parameters will remain constant when disparate datasets from different experiments conducted on human subjects are reproduced with LIDA agents. Finding such a set of parameters would provide substantial evidence of the accuracy and usefulness of the LIDA conceptual cognitive model.

Basic values for the parameters governing mechanisms in LIDA were derived from neuroscience data (Madl, et al., 2011). For example, visual feature detectors in LIDA agents have to take about 30ms to run, derived from neuronal conduction delays in area V1 in the human visual cortex (Huang & Paradiso, 2008; Kirchner, Barbeau, Thorpe, Regis, & Liegeois-Chauvel, 2009). These basic parameters were first tested in a simple reaction time task (LIDA Reaction Time agent), and verified in an experiment designed to investigate perceptual simultaneity and continuity (LIDA Allport agent), and two experiments examining the properties of attention (the LIDA Attention and Attentional Blink agents). The latter three agents were also motivated by the goal of validating some of the claims of Global Workspace Theory (GWT) of consciousness underlying the LIDA model. GWT posits that consciousness is discrete, which is consistent with some recent neuroscientific evidence (Doesburg, Green, McDonald, & Ward, 2009a; Raffone & Srinivasan, 2009; Thompson & Varela, 2001).

## A Guide to the LIDA Literature

This concludes our brief account of the mind according to LIDA. A persistant reader should have gained enough insight to begin to be able to use the LIDA model as a cognitive prosthesis with which to think about thinking, with which to understand and explain various common mental processes. For the occasional reader who has been enticed to want a deeper understanding of the model, there is a literature consistings of more than a hundred published papers describing LIDA and its predecessor IDA, as well as numberous slide presentations. A guide to this considerable literature will be forthcoming.

## References

Baars, B. (1997). In the theatre of consciousness. Global Workspace Theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies, 4*, 292–309.

Baars, B., Franklin, S., & Ramsøy, T. (2013). Global workspace dynamics: Cortical "binding and propagation" enables conscious contents. *Frontiers in Consciousness Research, 4*, 200. doi: 10.3389/fpsyg.2013.00200

Baars, Bernard J. (1988). *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

Baars, Bernard J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Science, 7*, 166–172.

Baddeley, A. D. (1992). Consciousness and working memory. *Consciousness and Cognition, 1*(1), 3-6.

Baddeley, A. D., Conway, M., & Aggleton, J. (2001). *Episodic memory*. Oxford: Oxford University Press.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. A. Bower (Ed.), *The Psychology of Learning and Motivation*. New York: Academic Press.

Baddeley, A. D., & Hitch, G. J. (2007). Working memory: Past, present…and future? In N. Osaka, R. Logie & M. D'Esposito (Eds.), *Working Memory - Behavioural & Neural Correlates*. Oxford: Oxford University Press.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577–609.

Barsalou, L. W. (2008). Grounded Cognition. *Annual Review of Psychology, 59*, 617–645.

Cansino, S. (2009). Episodic memory decay along the adult lifespan: A review of behavioral and neurophysiological evidence. *International Journal of Psychophysiology, 71*(1), 64-69.

Conway, Martin A. (2002a). Sensory-perceptual episodic memory and its context: autobiographical memory. In A. Baddeley, M. Conway & J. Aggleton (Eds.), *Episodic Memory* (pp. 53–70). Oxford: Oxford University Press.

Conway, M. A. (2002b). Sensory–perceptual episodic memory and its context: Autobiographical memory. In A. Baddeley, M. Conway & J. Aggleton (Eds.), *Episodic Memory*. Oxford: Oxford University Press.

D'Mello, S. K., Ramamurthy, U., Negatu, A., & Franklin, S. (2006). *A procedural learning mechanism for novel skill acquisition.* Paper presented at the Workshop on Motor Development: Proceeding of Adaptation in Artificial and Biological Systems, AISB'06, Bristol, England.

Damasio, A. R. (1994). *Descartes' Error*. New York: Gossett/Putnam Press.

Damasio, A. R. (1999). *The Feeling of What Happens*. New York: Harcourt Brace.

de Vega, M., Glenberg, A., & Graesser, A. (Eds.). (2008). *Symbols and Embodiment: Debates on meaning and cognition*. Oxford: Oxford University Press.

Diamond, D. M. (2005). Cognitive, endocrine and mechanistic perspectives on non-linear relationships between arousal and brain function. *Nonlinearity in Biology, Toxicology, and Medicine, 3*(1), 1-7.

Dijkstra, T., Schöner, G., & Gielen, C. (1994). Temporal stability of the action-perception cycle for postural control in a moving visual environment. *Experimental Brain Research, 97*(3), 477-486.

Doesburg, S., Green, J., McDonald, J., & Ward, L. (2009a). Rhythms of Consciousness: Binocular Rivalry Reveals Large-Scale Oscillatory Network Dynamics Mediating Visual Perception. *PLoS ONE, 4*(7), e6142. doi: 10.1371/journal.pone.0006142

Doesburg, S. M., Green, J. J., McDonald, J. J., & Ward, L. M. (2009b). Rhythms of consciousness: binocular rivalry reveals large-scale oscillatory network dynamics mediating visual perception. *PLoS ONE, 4*(7), e6142.

Drescher, G. L. (1991). *Made-up minds: a constructivist approach to artificial intelligence*. Cambridge, MA: The MIT Press.

Edelman, G. (1987). *Neural Darwinism: The theory of neuronal group selection*. New York: Basic Books.

Edelman, G., & Tononi, G. (2000). *A Universe of Consciousness: How Matter Becomes Imagination*. New York: Basic Books.

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102*, 211–245.

Estes, W. K. (1993). *Classification and Cognition*: Oxford University Press.

Faghihi, U., McCall, R., & Franklin, S. (2012). A Computational Model of Attentional Learning in a Cognitive Agent. *Biologically Inspired Cognitive Architectures, 2*, 25-36.

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry. *American Psychologist, 34*(10), 906-911.

Franklin, S. (2000). Deliberation and Voluntary Action in 'Conscious' Software Agents. *Neural Network World, 10*, 505-521.

Franklin, S. (2005a). *Cognitive robots: Perceptual associative memory and learning*. Paper presented at the 14th Annual International Workshop on Robot and Human Interactive Communication (ROMAN 2005).

Franklin, S. (2005b). Evolutionary pressures and a stable world for animals and robots: A commentary on Merker. *Consciousness and Cognition, 14*(1), 115-118.

Franklin, S. (2005c). *Perceptual memory and learning: Recognizing, categorizing, and relating.* Paper presented at the American Association for Artificial Intelligence (AAAI), Stanford University.

Franklin, S., & Baars, B. (2010). Two Varieties of Unconscious Processes. In E. Perry, D. Collerton, H. Ashton & F. LeBeau (Eds.), *New Horizons in the Neuuroscience of Consciousness* (pp. 91–102). Amsterdam: John Benjamin.

Franklin, S., Baars, B. J., Ramamurthy, U., & Ventura, M. (2005a). The role of consciousness in memory. *Brains, Minds and Media, 1*(1), 1-38.

Franklin, S., Baars, B. J., Ramamurthy, U., & Ventura, M. (2005b). 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., & Ramamurthy, U. (2006). *Motivations, Values and Emotions: 3 sides of the same coin.* Paper presented at the Sixth International Workshop on Epigenetic Robotics, Paris, France.

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

Freeman, W. J. (1999). *How Brains Make Up Their Minds*. London: Weidenfeld & Nicolson General.

Freeman, W. J. (2002). *The limbic action-perception cycle controlling goal-directed animal behavior.* Paper presented at the International Joint Conference on Neural Networks, Honolulu, HI.

Fuster, J. Ã. M. (2002). Physiology of executive functions: The perception-action cycle. In D. T. Stuss & R. T. Knight (Eds.), *Principles of frontal lobe function* (pp. 96–108). New York: Oxford University Press.

Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences, 20*, 1–19.

Glenberg, A. M., & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language, 43*, 379–401

Hayman, R., Verriotis, M. A., Jovalekic, A., Fenton, A. A., & Jeffery, K. J. (2011). Anisotropic encoding of three-dimensional space by place cells and grid cells. *Nature Neuroscience, 14*(9), 1182-1188.

Heilman, K. M. (1997). The Neurobiology of Emotional Experience. *Journal of Neuropsychology and Clinical Neurosciences, 9*, 439-448.

Hirtle, S. C., & Jonides, J. (1985). Evidence of hierarchies in cognitive maps. *Memory & Cognition, 13*(3), 208-217.

Hofstadter, D. R., & Mitchell, M. (1995). The Copycat Project: A model of mental fluidity and analogy-making. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol. 2: logical connections* (Vol. 2). Norwood, N.J.: Ablex.

Huang, X., & Paradiso, M. A. (2008). V1 response timing and surface filling-in. *Journal of Neurophysiology, 100*(1), 539-547. doi: 10.1152/jn.00997.2007

Jackson, J. V. (1987). Idea for a Mind. *ACM SIGART Bulletin, 101*, 23-26.

Jacobs, L. F., & Schenk, F. (2003). Unpacking the cognitive map: The parallel map theory of hippocampal function. *Psychological review, 110*(2), 285-315.

Johnston, V. S. (1999). *Why We Feel: The Science of Human Emotions*. Reading, MA: Perseus Books.

Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research, 4*, 237-285.

Kanerva, P. (1988). *Sparse Distributed Memory*. Cambridge MA: The MIT Press.

Kirchner, H., Barbeau, E. J., Thorpe, S. J., Regis, J., & Liegeois-Chauvel, C. (2009). Ultra-rapid sensory responses in the human frontal eye field region. *The Journal of neuroscience : the official journal of the Society for Neuroscience, 29*(23), 7599-7606. doi: 10.1523/JNEUROSCI.1233-09.2009

Koivisto, M., & Revonsuo, A. (2010). Event-related brain potential correlates of visual awareness. *Neuroscience & Biobehavioral Reviews, 34*(6), 922-934.

Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science, 12*(5), 171-175. doi: 10.1111/1467-8721.01254

Kruschke, J. K. (2010). Models of attentional learning. In E. M. Pothos & A. Willis (Eds.), *Models of Learning and Categorization*. Cambridge: Cambridge University Press.

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *The Journal of Neuroscience, 29*(31), 9771-9777. doi: 10.1523/JNEUROSCI.1319-09.2009

Madl, T., Baars, B. J., & Franklin, S. (2011). The Timing of the Cognitive Cycle. *PLoS ONE, 6*(4), e14803. doi: 10.1371/journal.pone.0014803

Madl, T., & Franklin, S. (2012). *A LIDA-based Model of the Attentional Blink.* Paper presented at the International Conference on Cognitive Modelling.

Maes, P. (1989). How to do the right thing. *Connection Science, 1*, 291–323.

Manns, J. R., & Eichenbaum, H. (2009). A cognitive map for object memory in the hippocampus. *Learning and Memory, 16*(10), 616-624. doi: 16/10/616 [pii]10.1101/lm.1484509

Massimini, M., Ferrarelli, F., Huber, R., Esser, S. K., Singh, H., & Tononi, G. (2005). Breakdown of cortical effective connectivity during sleep. *Science, 309*(5744), 2228-2232.

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.

Merker, B. (2005). The liabilities of mobility: A selection pressure for the transition to consciousness in animal evolution. *Consciousness and Cognition, 14*(1), 89-114.

Minsky, M. (1985). *The Society of Mind*. New York: Simon & Schuster.

Moser, E. I., Kropff, E., & Moser, M. B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience, 31*, 69-89.

Mulcahy, N. J., & Call, J. (2006). Apes save tools for future use. *Science, 312*(5776), 1038-1040.

Nadel, L. (1992). Multiple memory systems: What and why. *Journal of Cognitive Neuroscience, 4*(3), 179-188.

Nadel, L., & Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology, 7*(2), 217-227.

Neath, I., & Surprenant, A. M. (2003). *Human memory: An introduction to research, data, and theory*. Pacific Grove, CA: Brooks/Cole.

Neisser, U. (1976). *Cognition and Reality: Principles and Implications of Cognitive Psychology* San Francisco: W. H. Freeman.

O'Keefe, J. (1990). A computational theory of the hippocampal cognitive map. In J. Z. J. Storm-Mathisen & O. P. Ottersen (Eds.), *Progress in Brain Research* (Vol. 83, pp. 301-312): Elsevier.

Ornstein, R. (1986). *Multiminds: A new way to look at human behavior*. Boston: Houghton Mifflin and Co.

Raffone, A., & Srinivasan, N. (2009). An adaptive workspace hypothesis about the neural correlates of consciousness: insights from neuroscience and meditation studies. *Progress in Brain Research, 176*, 161-180.

Ramamurthy, U., D'Mello, S., & Franklin, S. (2005). *Role of Consciousness in Episodic Memory Processes*. Paper presented at the Ninth Conference of the Association for the Scientific Study of Consciousness - ASSC9, Pasadena, CA.

Ramamurthy, U., D'Mello, S. K., & Franklin, S. (2004). *Modified sparse distributed memory as transient episodic memory for cognitive software agents*. Paper presented at the The International Conference on Systems, Man and Cybernetics.

Rolls, E. T., & Stringer, S. M. (2005). Spatial view cells in the hippocampus, and their idiothetic update based on place and head direction. *Neural Networks, 18*(9), 1229-1241.

Sigman, M., & Dehaene, S. (2006). Dynamics of the Central Bottleneck: Dual-task and Task Uncertainty. *PLoS Biology, 4*(7), e220.

Sims, C. R., & Gray, W. D. (2004). *Episodic versus semantic memory: An exploration of models of memory decay in the serial attention paradigm*. Paper presented at the 6th International Conference on Cognitive Modeling -- ICCM2004, Pittsburgh, PA. http://128.113.2.9/~grayw/pubs/papers/2004/SG04_ICCM.pdf

Sloman, A. (Ed.). (1999). *What Sort of Architecture is Required for a Human-like Agent?* Dordrecht, Netherlands: Kluwer Academic Publishers.

Smith, J. D., & Washburn, D. A. (2005). Uncertainty monitoring and metacognition by animals. *Current Directions in Psychological Science, 14*(1), 19-24.

Snaider, J., McCall, R., & Franklin, S. (2009). *Time Production and Representation in a Conceptual and Computational Cognitive Model*. Paper presented at the AAAI Fall Symposium on Biologically Inspired Cognitive Architecture, Washington, DC.

Snaider, J., McCall, R., & Franklin, S. (2011). *The LIDA Framework as a General Tool for AGI.* Paper presented at the The Fourth Conference on Artificial General Intelligence (Springer Lecture Notes in Artificial Intelligence), Mountain View, California, USA.

Solstad, T., Boccara, C. N., Kropff, E., Moser, M. B., & Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science, 322*(5909), 1865-1868. doi: 10.1126/science.1166466

Stickgold, R., & Walker, M. P. (2005). Memory consolidation and reconsolidation: what is the role of sleep? *Trends in Neurosciences, 28*(8), 408-415.

Stringer, S. M., Rolls, E. T., & Trappenberg, T. P. (2005). Self-organizing continuous attractor network models of hippocampal spatial view cells. *Neurobiology of learning and memory, 83*(1), 79-92. doi: 10.1016/j.nlm.2004.08.003

Swenson, R., & Turvey, M. T. (1991). Thermodynamic Reasons for Perception--Action Cycles. *Ecological Psychology, 3*(4), 317-348.

Tarsitano, M. (2006). Route selection by a jumping spider (*Portia labiata*) during the locomotory phase of a detour. *Animal Behavior, 72*(6), 1437-1442.

Thompson, E., & Varela, F. J. (2001). Radical embodiment: neural dynamics and consciousness. *Trends in Cognitive Sciences, 5*(10), 418-425.

Tulving, E. (1983). *Elements of episodic memory*. Oxford: Clarendon Press.

Uchida, N., Kepecs, A., & Mainen, Z. F. (2006). Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nature Reviews Neuroscience, 7*(6), 485-491.

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind*. Cambridge, MA: MIT Press.

Vidnyánszky, Z., & Sohn, W. (2003). Attentional learning: learning to bias sensory competition. *Journal of Vision, 3*(9), 174.

Werdenich, D., & Huber, L. (2006). A case of quick problem solving in birds: string pulling in keas, *Nestor notabilis*. *Animal Behaviour, 71*(4), 855-863.

Wilcox, S., & Jackson, R. (Eds.). (2002). *Jumping spider tricksters: deceit, predation, and cognition*. Cambridge, MA: MIT Press.

Willis, J., & Todorov, A. (2006). First impressions:  Making Up Your Mind After a 100-Ms Exposure to a Face. *Psychological Science, 17*(7), 592-599.

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).

Yoshida, H., & Smith, L. B. (2003). Known and novel noun extensions: Attention at two levels of abstraction. *Child Development, 74*(2), 564-577.

Zhang, Z., Franklin, S., & Dasgupta, D. (1998). *Metacognition in software agents using classifier systems*. Paper presented at the Fifteenth National Conference on Artificial Intelligence.

Zhu, J., & Thagard, P. (2002). Emotion and action. *Philosophical Psychology, 15*, 19–36.