

University of Memphis

University of Memphis Digital Commons

CCRG Papers

Cognitive Computing Research Group

2012

A Computational Model of Attentional Learning in a Cognitive Agent

U. Faghihi

R. McCall

S. Franklin

Follow this and additional works at: https://digitalcommons.memphis.edu/ccrg_papers

Recommended Citation

Faghihi, U., McCall, R., & Franklin, S. (2012). A Computational Model of Attentional Learning in a Cognitive Agent. [10.1016/j.bica.2012.07.003](https://digitalcommons.memphis.edu/ccrg_papers/10.1016/j.bica.2012.07.003)

This Document is brought to you for free and open access by the Cognitive Computing Research Group at University of Memphis Digital Commons. It has been accepted for inclusion in CCRG Papers by an authorized administrator of University of Memphis Digital Commons. For more information, please contact khggerty@memphis.edu.

A Computational Model of Attentional Learning in a Cognitive Agent

Usef Faghihi, Ryan McCall, and Stan Franklin

Computer Science Department, University of Memphis, TN, USA

ABSTRACT. Biologically inspired cognitive architectures should faithfully model the high-level modules and processes of cognitive neuroscience. Also, they are expected to contribute to the BICA “challenge of creating a real-life computational equivalent of the human mind.” One important component of the mind is attention and attentional learning. In this paper, we describe conceptual and computational model of attention and attentional learning for intelligent software agents in the context of the broad-based biologically inspired cognitive architecture, LIDA. In LIDA attention is defined as the process of bringing content to consciousness. Implementing Global Workspace Theory, the mechanism of consciousness consists of a continuing sequence of broadcasts of the most salient current contents to all of cognition. We argue that the term attention describes the selection of conscious contents and should be distinguished from mechanism of consciousness itself. Attentional learning, the learning of to what to attend, has been relatively little studied by memory researchers. Here we describe a mechanism for attentional learning using the LIDA architecture. A basic implementation of such an attentional learning mechanism in a LIDA-based agent is presented. The agent performs a psychological attention experiment and produces results comparable to human subjects. The agent’s contribution in determining internal parameters for the LIDA architecture is also described. Our model of attentional learning distinguishes different aspects of selectionist and instructional learning. Attentional learning has not received its deserved attention in cognitive architecture research. This work represents a first step toward implementing the full range of cognitive faculties associated with attention and attentional learning in the LIDA cognitive architecture.

Keywords— Attention, Attentional learning, cognitive agents, LIDA cognitive architecture, Global Workspace Theory, alarms.

INTRODUCTION

Since the early days of psychology, researchers have given various definitions of attention and have ascribed different functions to it. William James defined attention in the following way:

“Everyone knows what attention is. It is the taking possession of the mind in clear and vivid form of one out of what seem several simultaneous objects or trains of thought.” (James, 1890). Attention has been described as a threshold-based process that determines whether or not we become aware of a stimulus (Thorne et al., 2003). Some researchers have attributed additional roles to attention: as a filter on what is perceived¹, and as a stabilizing force on sensory mechanisms (S Franklin, 2005; Merker, 2005; Purves et al., 2008; Ratey, 2001; Thorne, et al., 2003). Posner (Michael I Posner et al., 2004) suggests three separate functions of attention with distinct underlying brain networks. These functions include 1) *alerting*: “maintaining an alert state”; 2) *orienting*: “focusing our senses on the information we want” (e.g., your focus on reading this document); and 3) *executive attention*: “the ability to manage attention towards goals and planning”. Baars (1988, p299) defines attention as *the process of bringing content to consciousness*, a definition that is adopted by LIDA model and this paper as well.

Attentional learning is learning to what to attend, or “learning to attend” (Estes, 1993; Gelman, 1969; Kruschke, 2010; Vidnyánszky et al., 2003; Yoshida et al., 2003). It is also the human ability to acquire knowledge about the predictive relationships between events, which allows us to use past experiences to predict the future (Iordanova et al., 2006). Attentional learning can also bias sensory competition for visual attention (Vidnyánszky, et al., 2003). Autonomy and experience play key roles in the development of attention in humans (Leach, 1997).

Alarms play an important role in the survival of any species by bypassing attention and signaling the presence of dangers or warning sounds (Bliss, 2003; Lachman et al., 1961; Xiao et al., 2004). When an alarm situation occurs, we do not merely attend to it — we act on it. Thus alarms bypass attention (Häkkinen, 2010; Mateo, 2010; S. M. Miller et al., 2007).

In this paper, we do not address all aspects of attention discussed in the psychology, philosophy, and neuroscience literature. Rather, we seek to address specific aspects of attention, attentional learning, and alarms in terms of the broad cognitive architecture, LIDA. Many cognitive scientists have explored the relationships between attention and

¹ In terms of our LIDA model, “what is perceived” actually refers to “what comes to consciousness”.

learning often producing narrowly focused models (Kruschke, 2010; Reynolds et al., 2009). Our aim in this paper is not to present a model of attention or attentional learning having state-of-the-art performance. Rather here we are taking a first step towards integrating attentional learning into the LIDA cognitive architecture. It also must be noted that our aim for attentional learning in this paper is not to suggest a mathematical model. This is a conceptual and computational model for attentional learning. Thus, we are interested in finding what attentional learning does and whether it is compatible with the data obtained from neurological experiments.

In section two, we introduce the Global Workspace Theory of consciousness and cognition and, in section three; we briefly describe the LIDA model. Section four describes attention in the context of the LIDA model. Section five describes attentional learning and section six the process of alarms. Finally we describe a LIDA-based agent implementation that performs an attention manipulation task. The agent successfully replicates the human data for this task.

GLOBAL WORKSPACE THEORY AND THE LIDA ARCHITECTURE

LIDA implements Global Workspace Theory (GWT) (Bernard J Baars, 1988; B.J Baars, 1997), which has become perhaps the most widely accepted psychological and neurobiological theory of consciousness (B. J Baars, 2002; Dehaene et al., 2001; Kanwisher, 2001; Shanahan, 2010).

GWT postulates that human cognition is implemented by a multitude of relatively small, special purpose processors, almost always unconscious (Edelman, 1987; Jackson, 1987; Minsky, 1986; Ornstein, 1986). Processors are comparatively simple, and communication between them is relatively rare, occurring over a narrow bandwidth. A coalition of such processors is a collection that works together to perform a specific internal cognitive task. Coalitions normally perform routine actions, in pursuit of sensory, motor, or other problem-solving tasks. GWT suggests that the brain supports a global workspace capacity that allows for the distribution of conscious content in a single broadcast and integration of conscious content over multiple broadcasts (for neuroscience evidence, see (B. J Baars, 2002; Schneider et al., 2003)). A coalition of processors that gains access to the global workspace and wins the competition there can broadcast a message to all the unconscious processors, in order to recruit new components to join in interpreting a novel situation, or in solving the current problem, as well as to enable several forms of learning.

LIDA is a proof-of-concept model that fleshes out GWT. Many of the tasks in LIDA are accomplished by codelets (Hofstadter et al., 1994) that implement the processors in GWT. Codelets are small pieces of code (processes), each running independently. A class of codelets, called attention codelets, serves, along with LIDA's *Workspace* and the *Global Workspace*, to implement attention. Each attention codelet attempts to bring the contents of its coalition to the "consciousness" spotlight by winning the competition in the Global Workspace. The content of the winning coalition is then broadcast to all the processors in the system to recruit resources to respond to the current situation, and to learn.

It must be noted that LIDA is not a model of biological brains; it is a model of mind which means that the model should be able to replicate data from cognitive neuroscience's experiments but not necessarily MRI or EEG data.

THE LIDA ARCHITECTURE

The need for using broad systems-level cognitive architectures has been championed in the past by several researchers. As social psychologist Kurt Lewin so succinctly pointed out "There is nothing so practical as a good theory" (Lewin, 1951, p. 169) . Artificial intelligence pioneer Allen Newell strongly supported the need for systems-level theories and architectures, asserting that "You can't play 20 questions with nature and win" (1973). Echoing Newell in decrying the reliance on modeling individual laboratory tasks, memory researcher Douglas Hintzman (2011) wrote "Theories that parsimoniously explain data from single tasks will never generalize to memory as a whole..." Hintzman's arguments, which rest upon the need for systems-level cognitive architectures in memory research, carry over into the realm of intelligent agents, again calling for systems-level architectures. In their review article, Langley, Laird and Rogers (Langley et al., 2008) argue that "Instead of carrying out micro-studies that address only one issue at a time, we should attempt to unify many findings into a single theoretical framework, then proceed to test and refine that theory." They are calling for a broad-based, systems-level architecture, such as the LIDA architecture, which already makes headway toward many of the problems they cite as open (S Franklin et al., in review)(Franklin, et al., in review). In a table allowing ready comparison of properties of some twenty-six "biologically inspired cognitive architectures" (Samsonovich, 2010), LIDA compares quite well indeed in terms of

modeling a complete cognitive system, and also in terms of being truly biologically inspired. Thus the LIDA architecture is embodied.

(For further information on situated or embodied cognition, please see (Anderson et al., 2003; de Vega et al., 2008; Harnad, 1990; Ziemke et al., 2007).) The mechanisms used in implementing LIDA's several modules have been inspired by a number of different "new AI" techniques (Brooks, 1986; Drescher, 1991; Hofstadter, et al., 1994; Jackson, 1987; Maes, 1989). LIDA can be viewed as operating through its cognitive cycles (Figure 1), which occur five to ten times per second (S Franklin, 2005), and depend upon saliency determination in the agent. A cognitive cycle starts with sensing and usually ends with selection of an internal or external action. The cognitive cycle is conceived of as an active process that allows interactions between the different components of the architecture. Thus, cognitive cycles are always ongoing.

LIDA's Primary Modules and Processes.

1) Perceptual Associative Memory (PAM): This module corresponds to later areas of the different sensory cortices in humans (visual, auditory and somatosensory) as well as to integrative areas such as the entorhinal cortex. PAM is implemented as a modified slipnet (Hofstadter, et al., 1994) and allows the agent to distinguish, classify and identify external and internal information. There are connections (links) between PAM nodes. Activated segments of the slipnet are instantiated into the agent's Workspace as part of the current percept. PAM implements what is more commonly referred to as recognition memory.

2) Workspace: This module includes human-like preconscious buffers of working memory and additional submodules. It contains recent percepts that have not yet decayed away, local associations from the episodic memories, and recent contents of consciousness. All of these representations are combined with the help of structure-building codelets to generate a Current Situational Model; a buffer containing the agent's understanding of what is going on right now. Information written in the Workspace may reappear in different cognitive cycles until it decays away.

3) Episodic memories: These are the memories for events (what, where and when). When the consciousness mechanism broadcasts, its contents are encoded into Transient Episodic Memory (TEM), and may be later consolidated into LIDA's Declarative Memory (DM) (S Franklin, 2005).

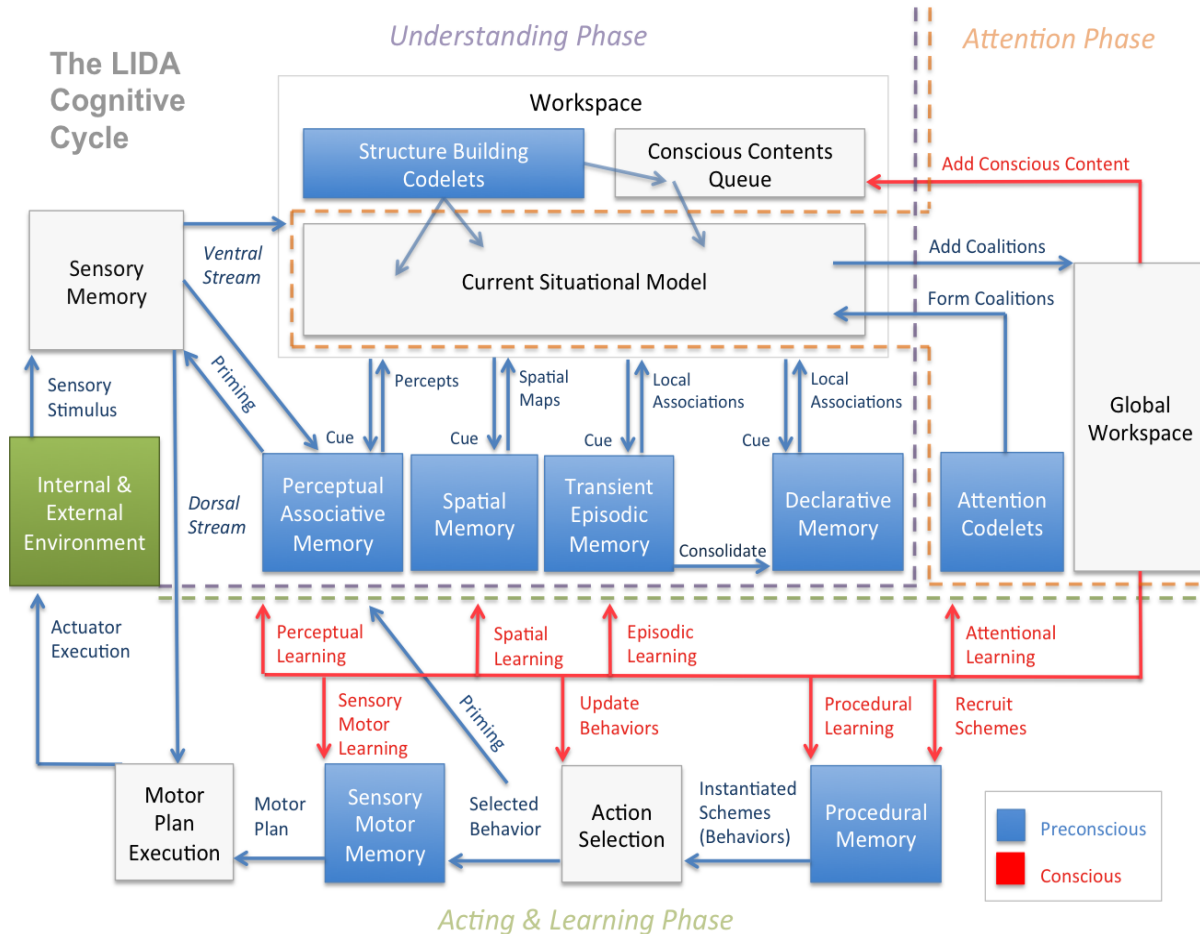


Figure 1 LIDA's Architecture

4) **Attentional Memory (ATM):** ATM consists of a collection of a particular kind of codelet called an attention codelet. All attention codelets are tasked with finding structures reflecting their own specific content of concern in the Current Situational Model (CSM) of the Workspace. For example, one codelet may look for a node representing fear. When an attention codelet finds its concern, it creates a coalition containing this structure and related content. The coalition is added to the Global Workspace to compete for consciousness. Each attention codelet has the following attributes: 1) *concern*: that content, whose presence in the CSM, can trigger the codelet to act; 2) a base-level activation, a measure of the codelet's usefulness in bringing information to consciousness; and 3) a current activation which measures the degree of intersection between its concern and the Current Situational Model. This measures the current saliency of its concern.

We propose that ATM includes at least four kinds of attention codelets. The *default attention codelet* observes the Current Situational Model in the Workspace, trying to bring its most energetic content to the Global Workspace. Thus it can be concerned with a broad spectrum of content, but its maximum activation is low. *Specific attention codelets* are codelets with specific concerns that have been learned. Each tries to bring particular Workspace content to the Global Workspace. *Expectation codelets*, created during action selection, attempt to bring the result (or non-result) of its recently executed action to consciousness. *Intention codelets* are attention codelets that bring to consciousness any content that can help the agent reach its current goal. When the agent makes a volitional decision, an intention codelet is generated.

5) **Procedural Memory:** LIDA's procedural memory begins the process of deciding what to do next. It is similar to Drescher's schema mechanism but with many fewer parameters (Drescher, 1988, 1991). The scheme net is a directed graph in which each of its nodes, called schemes, has a context, an action, results, and a base-level activation. As a result of the conscious broadcast, schemes from Procedural Memory are instantiated and put into the Action Selection mechanism. The Action Selection mechanism then chooses an action and Sensory-Motor Memory

executes the action (Figure 1). LIDA uses Maes' Behavior Network (Maes, 1989), with some modifications, as its Action Selection mechanism (Negatu et al., 2002).

Thus, in LIDA's architecture, while Procedural Memory and Action Selection mechanism are responsible for deciding what will be done next, Sensory-Motor Memory is responsible for deciding how a task will be performed. Thus, each item in Sensory-Motor Memory requires a distinct mechanism.

More detailed descriptions of various components of the LIDA architecture can be found in prior publications (S Franklin & F. G. Jr Patterson, 2006; S Franklin & U Ramamurthy, 2006).

ATTENTION IN THE LIDA MODEL

Now let us proceed to a careful discussion of the term *attention* (Michael I Posner, 2004; Michael I Posner, 2009; Michael I Posner, et al., 2004). The LIDA model adopts a psychological definition of attention, but it also includes the conscious broadcast in its definition. In the LIDA model, attention is the process of bringing Workspace content to consciousness (S Franklin, 2005). As described below, automatic, bottom-up processes unconsciously, and without effort bring items to the Workspace which may later, through attention, become conscious. Attention may also occur voluntarily and requires effort in a conscious, goal-directed way, over multiple cycles (S Franklin, 2005). A prevalent notion in psychology is that "attention" is a limited resource, a bottleneck. Our definition of attention includes consciousness, which is hypothesized to be the location of that bottleneck. More specifically a winner-take-all competition takes place in LIDA's Global Workspace each cognitive cycle between coalitions involving different perceptual content.

In the LIDA model, any Workspace content requires an attention codelet to add it to the competition for consciousness. In humans, sensory percepts appear to have a "royal road" to consciousness (Bernard J Baars, 1988). In LIDA this is implemented by a default attention codelet, which can potentially bring any sufficiently active sensory content to the Global Workspace.

In order to bridge our model to existing concepts of attention we will now identify several key aspects of attention currently accepted in the literature, and explain how they are implemented in the LIDA model. This is not meant to be an exhaustive list of such concepts.

Relation to Consciousness. Attention is also intimately related to consciousness as it determines what becomes the content of consciousness (Bernard J Baars, 1988; Michael I Posner, 2009). The contents of consciousness are learned (Bernard J Baars, 1988).

In the LIDA model attention codelets help determine the content of coalitions that competes for consciousness. They also influence the activation of the coalition itself.

Kinds of Attentional Processes. Attentional processes come in two forms: 1) bottom-up, and 2) top-down (Connor et al., 2004; Koch et al., 1985; Sarter et al., 2001). Bottom-up attention is an unconscious, involuntary process that produces conscious content. For example, if you are typing at the computer and a colorful chat notification flashes in one corner of the screen, your eyes automatically move to that flashing area of the screen, regardless of whether you volitionally choose to look away from your typing. This is an example of bottom-up attention, where you do not consciously (volitionally) control your attention to the stimulus.

On the other hand, top-down attention is voluntary. At any given moment, we can deliberately focus on an urgent or important situation around us. For example, the attention one uses to focus on a speaker at a conference is top-down.

In the LIDA model, bottom-up attention is a hard-wired reflex action. Top-down attention is achieved by a selected action, which has the agent focus on content in the Workspace (e.g., a person). Such an action creates an attention codelet that focuses on the person.

Influence of Emotion. Emotions can influence attention, and vice-versa (Damasio, 1999; Dolan, 2002; Stan Franklin et al., 2006; Haerich, 1994; Phelps, 2006; Rolls, 2000). The influence of emotion on attention can be illustrated by the following example. Imagine that a person at a subway station sees a brief glimpse of an attractive member of the opposite sex wearing a long, navy blue coat boarding the train. The two never meet, but for the next few minutes, or hours (or weeks?), the attention of the victim of cupid's arrow is drawn to those wearing similar

navy blue coats. In the opposite direction, the influence of attention on emotion can be illustrated by the following example. Suppose a person crossing a road sees an oncoming truck and feels some fear. The person becomes conscious of the feeling of fear. A second later they see a child step out in front of the truck. The empathetic fear for the endangered child may very well add to the fear they are already experiencing.

Using the LIDA model, the first example can be explained as follows. The attractive conspecific is perceived in PAM with a strong emotional affect, i.e., the node for the person is instantiated in the Workspace with strongly-activated emotion nodes attached, giving the entire representation a high activation. This representation, because of the high activation, easily wins the competition for consciousness; it is strongly learned in PAM. Thus the PAM nodes associated with this broadcast, e.g. person, navy blue coat, the subway location, etc., gain high base-level activation such that a small excitation from the current stimulus may put such nodes over the percept threshold. Being over threshold such nodes are likely to cue related episodic memories. Both can come to consciousness.

The second example can be explained as follows: by seeing the truck the person's PAM nodes related to the truck looming and the resulting fear gain activation, resulting in the fear becoming conscious. After one second these nodes are still highly activated so that when the child steps out the fear node's activation will increase beyond a normal amount and the person is conscious of the heightened fear

Selectivity and Saliency Determination. Thorne referred to our ability to select the most salient stimuli among different internal/external stimuli as *selectivity* (Michael I Posner, 2009; Michael I Posner et al., 1980; Thorne, et al., 2003). The brain regions associated with the selectivity process are referred to as the “spotlights.” They are associated with the superior parietal and frontal regions and the temporo-parietal region along with the superior colliculi and pulvinar area, modulated mostly by the neurotransmitter acetylcholine (Purves, et al., 2008). *Saliency determination* is defined as the ability to pick out the most salient stimulus among other stimuli (Thorne, et al., 2003). Importance is only one dimension of saliency, others including urgency and insistence (Sloman, 1987). Still others are movement, novelty, unexpectedness and goal relatedness.

LIDA considers the definition of selectivity to be exactly that of attention; namely, the process of bringing content to consciousness (S Franklin, 2005). For instance, in LIDA, if the information coming to the Workspace contains emotional valences, they will have more activation and, therefore, salience. Activation, in this case, is intended to measure saliency. Attention codelets could also be concerned with shiny objects or specific colors, etc. (Selectivity and saliency determination are discussed in detail in the next section.)

Orienting. Orienting refers to actions that focus our senses on information that is of current interest (Michael I Posner, et al., 2004). In humans, this most typically involves vision. Furthermore, orienting or focal maintenance refers to the maintaining of attentional focus for the accomplishment of an action. For instance, a mathematics or physics problem usually requires more time and mental resources to be solved than reading a page of a novel does. The parts of the brain that control our ability to maintain mental energy for a sufficient amount of time are associated with the anterior cingulate gyrus, anterior insula and parts of the basal ganglia (Dayan et al., 2000; Fan et al., 2005; Michael I Posner, 2009; Michael I Posner et al., 2007).

In the LIDA model, attention codelets whose contents are sensitive to movement, location, or spatial information help bring content to consciousness. If such content wins the competition for consciousness an orienting action may ensue.

Previewing and Planning. Previewing refers to the involvement of attention in deliberation and planning. Given a situation, before selecting an action, we estimate the outcome of several possible actions in order to decide what the best action is for the current situation. This corresponds to James' ideomotor theory (Bernard J Baars, 1988 ch.7; Franklin 2000; James, 1890).

Deliberation and planning in LIDA take place over multiple cognitive cycles in the Current Situational Model, particularly in the virtual window of the perceptual scene in the Workspace.

Self-Monitoring and Self-Regulation. Self-monitoring (or *executive attention*) refers to the regulation of a variety of processes, such as monitoring a task's procedures, detecting errors, resolving conflicts and modifying a task's procedures when necessary. It also involves planning and decision-making (Michael I Posner, et al., 2004). Self-monitoring also refers to our estimation of the time required to complete a given task, and the level of concentration and resources needed to accomplish such a task. For example, while driving a car, we are careful to

respect all road signs. Self-regulation processes are associated with the anterior cingulate gyrus, the anterior insula and part of the basal ganglia (Dayan, et al., 2000; Fan, et al., 2005; Michael I Posner, 2009; Michael I Posner, et al., 2007). Posner (Michael I Posner et al., 1998) has also suggested that the aforementioned brain regions are involved when a task includes conflicting elements and requires what he calls executive attention. However, the detailed functionality of the aforementioned parts of the brain during self-regulation process is yet to be found.

To describe executive attention we quote an example from Miller and Cohen (2001). “In the USA, we usually look left when first crossing the road. This is not true when we go to UK where we must first look right. Thus, in the case of being in UK we must deal with the following parameters, the stimulus (the road) with a context (UK) to cue a behavior (look right).”

In LIDA, the road stimulus is recognized by a feature detector in PAM which instantiates a “road” node. The context (UK) was also recognized by PAM and would be part of the content in the Current Situational Model in the Workspace. Supposing that both the road node and the UK node entered the Global Workspace and won the competition for consciousness, these nodes would be broadcast throughout the system including to Procedural Memory. Appropriate schemes in Procedural Memory could be instantiated, e.g. a scheme with context “road” and “UK”, and action of “look- right”.

In the LIDA model, the expectation codelets play an important role in resolving conflicts. Expectation codelets are created when a behavior is selected. An expectation codelet’s content of concern comes from result of the selected behavior that created it. Expectation codelets are a kind of attention codelet that try to bring the result (or non-result) of a previously executed action to consciousness. This helps the system detect errors and consequently search for solutions to fix the errors. Consider a US citizen in the UK wanting to cross a road. They select an action look left with the expectations of the car coming from the left. However, no car comes from left. Then the scheme that contains the action (look left before crossing the road) has its base-level activation decreased.

In the remainder of this paper, we present a computational model of attention and attentional learning in LIDA.

ATTENTIONAL LEARNING IN LIDA

Here we give a conceptual explanation of attentional learning in LIDA. For more details of attention in the LIDA model please see (Stan Franklin et al., 2005). Recall that attention codelets run within LIDA as part of Attentional Memory (ATM). All attention codelets have concerns, which are typically template-like as the concerns are usually expressed in relatively abstract terms. A particular attention codelet’s concern may include an interest in birds, for example. So the appearance of a *robin* node with attached *bird* node in the Current Situational Model may induce it to create a coalition from this content of concern, the bird node, as well as nodes directly connected to the concern, in this case, the robin node.

Two kinds of attentional learning may occur each time a conscious broadcast comes to ATM. In *selectionist learning*, the attention codelet that brought the winning coalition to consciousness has its base-level activation reinforced. In *instructionalist learning*, a brand-new attention codelet is created with a specific content of concern. The LIDA model employs a general principle in each of its several modes of learning, namely, the principle of profligacy. The current contents of each conscious broadcast lead to learning, both selectionist and instructional, in each mode. Thus, learning occurs with the least provocation, but learned entities decay away unless they are later reinforced.

We will first consider instructional attentional learning. During LIDA’s cognitive cycles, percepts from PAM and local associations from TEM and DM continually enter the Workspace. Such Workspace content can be acted upon by attention codelets, which detect salient events. The default attention codelet is a primitive, built-in attention codelet, which tries to bring the most energetic Workspace content to consciousness. When a coalition created by this attention codelet wins the competition for consciousness, ATM will determine that it was produced by the default attention codelet. ATM’s attentional learning mechanism then creates a new *specific attention codelet*. This new codelet’s concern is set to be the most highly activated part of the winning coalition. The new specific attention codelet will have an initial base-level activation based on the default attention codelet’s base-level activation and the coalition’s current activation. In this way, an attention codelet is created in ATM for each broadcast of conscious content for which there is not already a dedicated attention codelet. Consider, for example, a person who is unafraid of dogs. One day, he is bitten by a dog, and feels great pain. From that point on, anything in his environment resembling a dog will very likely come to his consciousness.

Let’s walk through the events that would occur in a LIDA agent in this example. We will start with the event of

the person being bitten by the dog. The person is simultaneously conscious of a great pain and of the dog. A conscious broadcast of these nodes results in the perceptual learning of a relation between the *dog* node and *fear* node in his PAM. The dog node has its base-level activation reinforced. At this point it is very likely a coalition with dog and fear is created by the default attention codelet. Thus, from this point forward, he readily perceives dogs and, by association, is fearful of the dog. Instructionalist attentional learning of the conscious content of “*fearing dogs*” can lead to having an attentional codelet that would bring dogs to consciousness.

Selectionist learning occurs for an existing attention codelet when a coalition it has added to the Global Workspace wins the competition for consciousness. Such reinforcement learning is implemented by increasing the base-level activation of attention codelet. Consider again the example of the same man who is now fearful of dogs. When he perceives a dog one of his attention codelets creates a coalition with the *dog* node in it. If this coalition wins the competition for consciousness then the attention codelet that created it has its base-level activation reinforced.

Expectation codelets have their base-level activation reinforced whenever the coalition they add to the Global Workspace wins the competition for consciousness. If no similar attention codelet exists already then this expectation codelet is learned as a new attention codelet. Otherwise the existing attention codelet will be reinforced. If an expectation codelet confirms the presence of its content of concern, and its coalition wins the competition for consciousness, then the codelet’s associated scheme in Procedural Memory is positively reinforced. Similarly if the codelet does not find its content of concern, and a coalition containing this information wins the competition for consciousness, then its associated scheme is negatively reinforced.

Selectivity and Saliency Determination in LIDA

Saliency in LIDA, measured by total (current plus base-level) activation, is the currency of the competition for consciousness, and thus directly influences selectivity. The process by which selectivity and saliency are determined in LIDA can be illustrated by the following example. Imagine a man gets into his car and turns the keys to start the engine. He puts the car in gear and presses down on the gas pedal, but the car does not move. His attention shifts from the task of starting the car to finding out why the car is not moving. Looking around he discovers that the car’s emergency hand break had not been released.

What does the LIDA model say is occurring in the man’s mind in this situation? For many cognitive tasks, humans are able to recognize situations immediately (within about 100ms) and react appropriately in a short time frame (about 300ms). However, more time is sometimes needed to respond to a situation, such as in the above example. Note that a LIDA agent may not always have an external action to take; however, internal actions can still be taken. Internal actions can start a new process, for example, initiating a timekeeper at the beginning of volitional decision making (Franklin 2000). They can also affect the content of the Workspace, for example, adding a new scene in an ongoing daydream to the Current Situational Model.

We can assume that the man in the example has selected a behavior (instantiated scheme) to press the gas pedal. Such a behavior has the result “car moves forward.” Thus an expectation attention codelet is spawned with “car moves forward” as its content of concern. The more time that passes without the man perceiving that the car is moving forward the higher the activation of this codelet becomes, and thus the coalitions it creates grow more salient. Eventually a coalition containing the content “car is not moving” wins the competition. This recurring process of creating and adding coalitions to the Global Workspace can easily lead to many conscious broadcasts of content about the car not moving. These repeated broadcasts continually instantiate schemes relevant to determining why the car does not move. These instantiated schemes are highly activated and likely to be selected in Action Selection. Once the car does begin to move the expectation attention codelet concerned with the car moving no longer finds its concern in the content of the CSM. Thus it ceases to create and add coalitions about the movement of the car.

ALARMS IN LIDA

In dangerous situations, agents must react as quickly as possible. An *alarm* situation is one where an agent selects, and executes action(s) to respond to a dangerous situation prior to becoming conscious of the situation. In an alarm situation, the percepts are not brought to consciousness by attention codelets; rather percepts directly activate schemes in Procedural Memory. Such schemes will be highly activated and will likely be selected for execution. Whether a percept initiates an alarm depends on its activation. If the activation of a node or link is above

an alarm threshold, it will initiate an alarm. Emotions also play an important role in alarms by providing additional activation to perceptual representations.

Alertness² is also involved in alarms. Learning alarms, that is, learning to bypass attention, is a form of attentional learning (Chun et al., 1999; Häkkinen, 2010; Liddell et al., 2005; Mateo, 2010; S. M. Miller, et al., 2007; Ogawa et al., 2002). Once we have learned that a situation is dangerous, it can influence our decision-making, reaction time and intensity, and attentional process (LeDoux, 2000; Rolls, 2000; Sloman, 1998; Squire et al., 2000). The events that take place in LIDA during an alarm situation are illustrated by the following example.

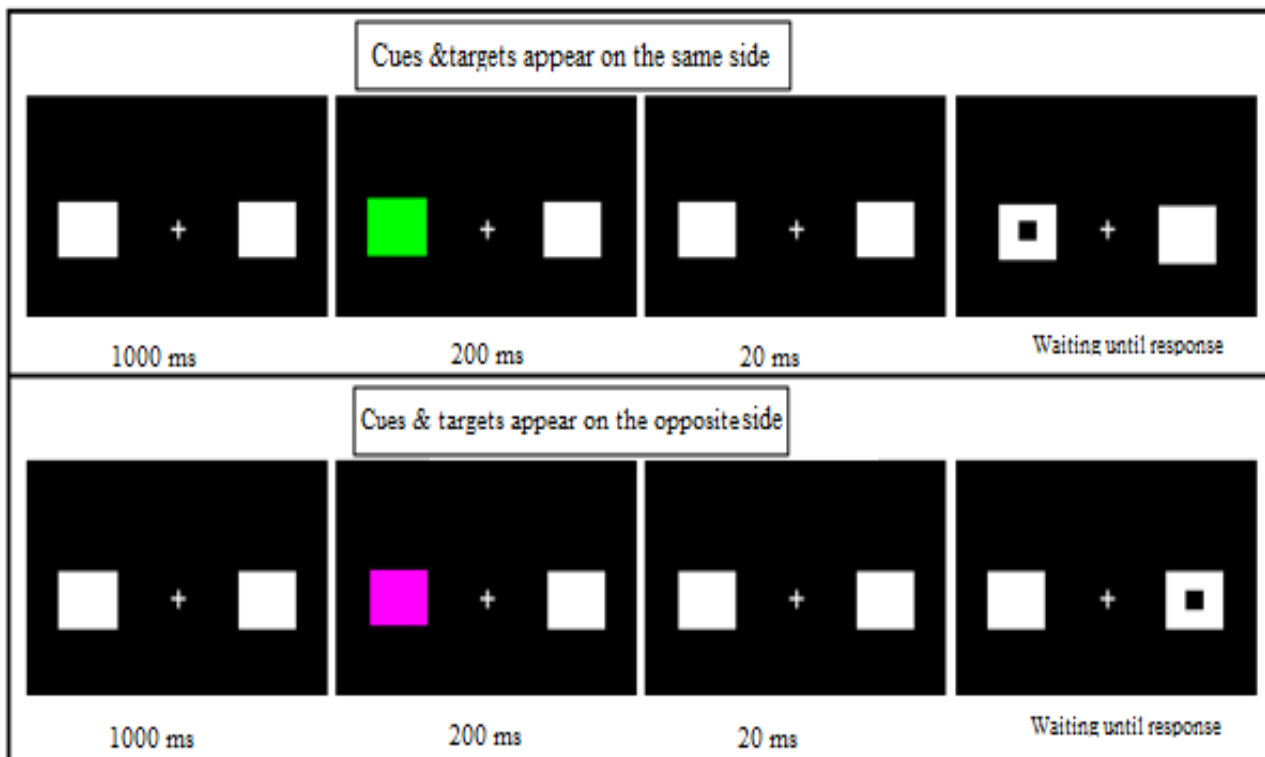


Figure 2 Van Bockstaele's Attention experiment

Suppose that you are driving a car and suddenly another car cuts in front of you. You respond by quickly pressing down hard on the brake pedal and turning the steering wheel before you “even realized what happened”. While perceiving an oncoming car is a complex perceptual task, we will assume that the agent (you) in this situation is capable of recognizing the danger in this situation. This preconscious recognition takes place in Perceptual Associative Memory. Since, a strong fear is associated with this percept; its activation goes over the alarm threshold. In this situation the nodes and links related to this alarm directly prime the schemes of Procedural Memory along the “alarm pathway,” bypassing both the Workspace and the Global Workspace (Figure 1). This representation has high activation so the schemes it excites also become highly activated. Action Selection will likely select such a scheme for execution. Before its action is executed, the instantiated scheme spawns a new expectation codelet. After the action is executed, this newly created expectation codelet focuses on changes in the environment as a result of the action being executed.

At the same time that activation is passing along the alarm pathway, the percept of the oncoming car enters the Workspace as normal. It may be brought to consciousness by attention codelets. Thus the agent may still choose a consciously mediated action after the alarm pathway has already produced the selection and execution of an action.

The agent can become conscious of the action it took in direct response to the alarm. Based on the alarm response action and the actual result of the action that the agent perceives, the agent can comprehend how effective this action was. If it becomes conscious of this appraisal it will undergo procedural learning for the scheme selected during the

² State of being watchful for possible danger or emergency (Merriam-Webster Online Dictionary).

alarm.

In terms of the example, by rehearsing this near-death event consciously, you become conscious of the fact that breaking hard and turning when a car pulls in front of you has saved you from an accident. Procedural learning will then occur for this scheme as a result of a conscious broadcast of this event. This learning will cause the agent respond in this way more reliably in response to similar alarm situations in the future.

ATTENTION AGENT EXPERIMENT

To test the attention mechanism we replicated part of a psychological experiment that manipulated subjects' attention. We adapted Van Bockstaele's experiment (Van Bockstaele et al., 2010) for use with a LIDA software agent (attention agent). To simulate this task using the attention agent, we replicated the same environment (Figure 2). Each of the eight large black squares in Figure 2 represents the state of the computer monitor at various stages in the experiment. As in the human experiment, the agent was required to respond to a target that appeared after a cue on a computer monitor. The cue and the target appeared on either the right or left side of the screen against a black background. The target appeared on either the same side as the cue (congruent) or on the opposite side (incongruent). In each trial a fixation target was presented for 1000ms surrounded by two white rectangles (150 by 150 pixels). Then, one of the white rectangles was randomly replaced by the cue, a colored rectangle (either green or pink, at random) for 200ms. Next two white rectangles were presented again for 20ms removing the colored rectangle. Then, the target, a black rectangle (about 30 by 30 pixels) was randomly presented inside one of the white rectangles until the subject responded. The attention agent could respond "left" if the target appeared on the left or "right" if the target appeared on the right. In the human experiment, subjects' reaction times for congruent trials (cue and target appearing on the same side) were 360ms on average whereas the average reaction time was 380ms for incongruent trials (cue and target appearing on different sides). The experimenters concluded that the 20ms difference in reaction time was due to the fact that the cues attract attention and thus targets appearing at the same location will elicit a faster reaction time than targets at the opposite location (Van Bockstaele, et al., 2010).

Now we describe the details of the attention agent implementation. The agent has rudimentary perceptual faculties in the form of feature detectors that were hand-crafted to detect the presence of cues and targets, and to excite the appropriate built-in PAM nodes. These include nodes representing cues and targets and their position. Specifically there are nodes for the colors *green*, *pink*, and *black* as well as for a small *rectangle* and the positions *left* and *right*. Because of the simplicity of the perceptual algorithm used we employ the following technique in order to simulate the timing of more realistic perception: when patterns are detected in the sensory stimuli, the corresponding nodes

Behavior Label	Activation	Context	Action	Adding Result	Deleting Result
left, pressLeft	0.5667	Nodes (rectangle[8],black[14],left[4]) Links ()	action.Left	Nodes () Links ()	Nodes () Links...
right, pressRight	0.2667	Nodes (right[6],rectangle[8],black[14]) Links ()	action.Right	Nodes () Links ()	Nodes () Links...

Tick at Selection	Selection Count	Action
4850	0	Left

Figure 3. A screenshot of the LIDA-based attention agent running. The above panel contains internal contents of Action Selection. The top table in this panel displays candidate behaviors.

are only excited with a small amount of activation. Thus detectors would need to excite the node multiple times to sufficiently activate the node. Combined with a detector operation frequency of five (see Table 1) we are able to simulate a longer time for perceptual recognition.

The agent has a default attention codelet, which creates a coalition from the most highly activated Workspace content when it runs.

The agent performs both instructionalist and selectionist learning with each conscious broadcast. Whenever the default attention codelet was responsible for creating the winning coalition, a new attention codelet is learned (instructionalist) with content of concern equal to that of the broadcast. If a non-default attention codelet is responsible for a winning coalition its base-level activation is reinforced.

The default attention codelet’s base-level activation was not reinforced. The agent has two built-in schemes for responding “left” and “right”. The first scheme has the context *left, black* and the action “respond left” while the other has the context *right, black* and the action “respond right”. The agent’s Action Selection periodically chooses the action of the most highly activated scheme, which, in this case, will be *left* or *right* (Figure 3).

In addition to replicating the aforementioned attention experiment, this agent implementation allows us to tune the internal parameters of the LIDA model and suggests additional parameters. We started with those parameters and parameter values suggested by Madl et al. (2011) based on neuroscientific data and LIDA’s cognitive cycle hypothesis. This hypothesis states that a LIDA agent’s reaction time is in accordance with the duration of a cognitive cycle and takes approximately 260–390ms i.e. comparable to a human’s. An initial phase of perception (stimulus recognition) lasts 80–100ms from stimulus onset under optimal conditions. It is followed by a conscious broadcast 200–280ms after stimulus onset, and an action selection phase 60–110ms from the start of the conscious phase (Madl, et al., 2011). The attention agent was initially given the internal parameters identified by Madl et al. However, since it included additional faculties, such as a more detailed Action Selection module, new parameters were introduced and tuned for this work.

Before the parameters are presented it is necessary to discuss the parameter format. Processes in LIDA agents are implemented by “tasks” which have an execution frequency. Each task’s frequency is controlled by its *ticksPerRun* parameter. That is, how frequently in *ticks* (simulated milliseconds) the task is executed. A *ticksPerRun* of 1 implies the task is run at every time step, i.e., as frequently as possible.

Table 1 shows a comparison between the parameters used in our experiment and those used by Madl et al. The *sensing frequency* parameter controls how often the agent senses the environment, updating the contents of its Sensory Memory. The *feature detector frequency* parameter (FDP) controls how often each feature detector runs. Our results are comparable with the Madl’s for the FDP because there was a difference in the amount of excitation that nodes received from feature detectors. Perhaps a better approach for future work is to require feature detectors to run every tick. Then the parameter for feature detectors could be in “activation per tick” units like decay rates. The *attention codelet frequency* parameter (ACP) controls how often attention codelets are run. The difference between our ACP and Madl’s is due to the fact that our task required multiple broadcasts while Madl’s experiment only considered simple reactions requiring only one conscious broadcast. For our task the broadcast occurred too infrequently using a parameter value of 200. We chose a value of 100 which is within the range of the believed 7-10hz rate of consciousness (Doesburg et al., 2009). The *NoBroadcastOccuring trigger* parameter controls the maximum time that can pass before a competition for consciousness is held. The *Action selection frequency* parameter governs the frequency of the process that leads to the selection of an action.

	Parameter Name	Reaction Time Agent	Attention Agent
1	Sensing frequency	20 ticks	20 ticks
2	Feature detector frequency	30 ticks	5 ticks
4	Attention codelet frequency	200 ticks	100 ticks
5	NoBroadcastOccuring trigger frequency	200 ticks	200 ticks
6	Action selection frequency	100 ticks	100 ticks
7	Node decay rate	0.006 / tick	0.01 / tick
8	Coalition, scheme, and behavior decay rate	n/a	0.01 / tick
9	Default attention codelet base-level activation	n/a	0.5
10	Action selection candidate threshold	n/a	0.9

Table 1 Comparison of internal parameters for Reaction Time Agent (Madl et al. 2011) and Attention Agent

Decay is an important component of LIDA agents as the activations of nodes, links, codelets, coalitions, schemes,

and behaviors all decay. The *node decay rate* parameter determines how quickly the activation of nodes representing objects is decayed. The *coalition, scheme, and behavior decay rate* parameter governs how much activation is decayed away from coalitions, schemes, and behaviors every time step. These two decay rate parameters are equivalent³ (Table 1) and thus there is effectively only one decay rate parameter for this agent. The *default attention codelet base-level activation* specifies the general importance of the default attention codelet. Its value is lower than that of newly learned attention codelets, which are learned with a base-level activation of 1.0. The *action selection candidate threshold* is a threshold for an action to be selected. All parameters not mentioned in Table were given values e.g. 1 or 0 so that they would have no effect.

Discussion of Attention Agent Experiment. The attention agent was run in a series of 100 trials including both congruent and incongruent trials. The average reaction time was 360ms for all congruent trials and 380ms for all incongruent trials. Our results for congruent and incongruent trials are in accordance with the results of Van Bockstaele's and others (Del Cul et al., 2007; Fize et al., 2011; Gaillard et al., 2009). Indeed, we found that the reaction times were significantly different for the congruent and incongruent trials.

In this experiment the agent is alert to its task *al la* Posner since its sole sensory input is the visual image of the computer screen (Michael I Posner, et al., 2004). This experiment is also a good example of how Posner's "orienting" is implemented in LIDA (Michael I Posner, et al., 2004). For example suppose that during the first trial the cue and the target appeared in the same location: left. When the agent becomes conscious of the target appearing in the left it undergoes attentional learning. The attention codelet concerned with left would have its base-level activation reinforced. In the following trial, because of the learning from the previous trial, it will now be easier for a target appearing on the left to come to consciousness.

In this congruent case, two factors play an important role in action selection. The first factor relates to the intersection between the conscious broadcast and the contexts of the two schemes in Procedural Memory. Since the cue and target both appeared on the left, the scheme containing the action "press left" will have significantly more activation than the scheme containing the action "press right" because its context will better match the recent contents of consciousness.

The second factor relates to the fact that the conscious broadcast is sent directly to both Procedural Memory *and* Action Selection. When a conscious broadcast arrives in Action Selection and a behavior is present that behavior's context is updated and its activation may change as a result. Suppose during a congruent trial that the scheme for "press left" is instantiated and added to Action Selection as a result of several conscious broadcasts of the cue appearing on the left. Then, later, when a strong broadcast of the target arrives in Action Selection, the "press left" behavior is updated receiving sufficient new activation from that broadcast to go over the action selection threshold. Note that the arrival of a conscious broadcast carrying a target carries enough activation by itself to win the action selection competition and trigger action selection.

In the incongruent case, where, e.g., the cue appears on the right and target on the left, initial conscious broadcasts are related to the cue appearing the right and instantiate the scheme with action "press right". Thus, the scheme is instantiated in Action Selection albeit without sufficient activation for selection because the cue was not black. When the conscious broadcast of the target arrives, it excites this "incorrect" behavior but, again, the intersection between broadcast and context is insufficient to activate the behavior above selection threshold. In Procedural Memory, the scheme with action "press left" strongly matches this broadcast, is instantiated as a result, and enters Action Selection with enough activation to be selected right away. Since the instantiation time for a scheme is 20ms, the selection in the incongruent case requires an additional 20ms as compared to the congruent case. Recall that in the congruent case the "correct" behavior was already instantiated in the Action Selection at the time of the broadcast of the target.

The congruent case is a good example of how LIDA explains executive attention. The conscious broadcast of the cue (attention to the cue) primed the Action Selection part of LIDA allowing it to respond more quickly to a target appearing on the same side.

CONCLUSION

We present a conceptual and computational model of attention and attentional learning for intelligent software

³ Given the simplicity of this agent's perception we only cite one node decay parameter; however, for agents with more complex perception multiple such parameters likely having different values would be necessary [32].

agents in the context of the LIDA cognitive model. We explain how the LIDA model implements several aspects of attention including those identified [7] by Posner: 1.) alerting: “maintaining an alert state”; 2.) orienting: “focusing our senses on the information we want; and 3.) executive attention: “the ability to manage attention towards goals and planning”.

Attentional learning is defined as the learning of to what to attend (Estes, 1993; Vidnyánszky, et al., 2003). This form of learning has not received its deserved attention in cognitive architecture research (see the BICA Table (Samsonovich, 2010)), primarily since so few architectures contain explicit attention mechanisms. Attentional learning in the LIDA model occurs with each conscious broadcast, during which two modes of attentional learning may occur. In *selectionist learning*, the attention codelet that brought the winning coalition to consciousness has its base-level activation reinforced. In *instructionalist learning*, a brand-new attention codelet is created with a specific content of concern.

We presented a LIDA-based agent implementation, which performs an attention manipulation task. The results of the agent’s performance on this task are comparable with human data.

This work represents a first step toward implementing the full range of cognitive faculties associated with attention. The present theory can be considered as a base implementation of attention and attentional learning in a cognitive architecture. Future work will involve replicating more experiments (e.g., causal learning, episodic learning) with LIDA agents using the same parameters we obtained in this study and comparing the results with human data.

Our future work would be to explore how attention in the LIDA model controls and regulates a variety of processes, such as monitoring a task’s procedures, detecting errors, resolving conflicts and modifying a task’s procedures when necessary.

ACKNOWLEDGMENT

We would like to thank the FQRNT for its financial support and all the members of CCRG lab who have contributed to the development of the LIDA’s computational model.

REFERENCES

- Anderson, J. R., & Lebiere, C. (2003). The Newell test for a theory of cognition. *Behavioral and Brain Sciences*, 5(26), 587-640.
- Baars, B. J. (1988). A cognitive theory of consciousness. *Cambridge: Cambridge University Press*.
- Baars, B. J. (1997). In the Theater of Consciousness: The Workspace of the Mind. *Oxford: Oxford University Press*.
- Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Science* (6), 47-52.
- Bliss, J. P. (2003). An investigation of alarm related accidents and incidents in aviation. *International Journal of Aviation Psychology*, 13(3), 249-268.
- Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1), 14-23.
- Chun, M. M., & Jiang, Y. (1999). Top-down attentional guidance based on implicit learning of visual covaria-tion. *Psychological Science*, 10, 360-365.
- Connor, C. E., Egeth, H. E., & Yantis, S. (2004). Visual Attention: Bottom-Up Versus Top-Down. *Current Biology*, 14, R850–R852.
- Damasio, A. R. (1999). The Feeling of What Happens: Body and Emotion in the Making of Consciousness. *New York: Harcourt Inc*.
- Dayan, P., Kakade, S., & Montague, P. R. (2000). Learning and selective attention. *Nature America*, 1218-1223.
- de Vega, M., Glenberg, A., & Graesser, A. (2008). Symbols and Embodiment: Debates on meaning and cognition. *Oxford: Oxford University Press*.
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* 79, 1-37.
- Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS Biol* 5: e260.
- Doesburg, S., Green, J., McDonald, J., & Ward, L. (2009). Rhythms of consciousness: binocular rivalry reveals large-scale oscillatory network dynamics mediating visual perception. *PloS one* 4: e6142.
- Dolan, R. J. (2002). Emotion, Cognition, and Behavior. *Science Vol. 298*, 1191-1194

- Drescher, G. L. (1988). Learning from Experience Without Prior Knowledge in a Complicated World. *Proceedings of the AAAI Symposium on Parallel Models*. AAAI Press.
- Drescher, G. L. (1991). *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. Cambridge, MA: MIT Press.
- Edelman, G. M. (1987). *Neural Darwinism*. New York: Basic Books.
- Estes, W. K. (1993). *Classification and cognition*. Oxford: Oxford University Press.
- Fan, J., McCandliss, B. D., Fossella, b. J., Flombaum, J. I., & Posner, M. I. (2005). The activation of attentional networks. *NeuroImage* 26, 471 – 479.
- Fize, D., Cauchoix, M., & Fabre-Thorpe, M. (2011). Humans and monkeys share visual representations. *publmed.gov*.
- Franklin, S. (2000). Deliberation and Voluntary Action in “Conscious” Software Agents. *Neural Network World*, 10, 505-521.
- Franklin, S. (2005). Evolutionary Pressures and a Stable World for Animals and Robots: A Commentary on Merker. *Consciousness and Cognition*, 14, 115-118.
- Franklin, S., Baars, B. J., Ramamurthy, U., & Ventura, M. (2005). The Role of Consciousness in Memory. *Brains, Minds and Media*, Vol.1, bmm150 (urn:nbn:de:0009-3-1505).
- Franklin, S., Madl, T., D'Mello, S. K., & Snider, J. (in review). LIDA: A Systems-level Architecture for Cognition, Emotion, and Learning.
- Franklin, S., & Patterson, F. G. J. (2006). The Lida Architecture: Adding New Modes of Learning to an Intelligent, Autonomous, Software Agent. *Integrated Design and Process Technology, IDPT-2006, San Diego, CA, Society for Design and Process Science*.
- Franklin, S., & Ramamurthy, U. (2006). Motivations, Values and Emotions: 3 sides of the same coin. *Proceedings of the Sixth International Workshop on Epigenetic Robotics, Paris, France, 128*, 41-48.
- Franklin, S., & Ramamurthy, U. (2006). Motivations, values and emotions: three sides of the same coin. *In Proceedings of the Sixth International Workshop on Epigenetic Robotics, Paris, France, September 2006, Lund University Cognitive Studies*, 128, 41-48.
- Gaillard, R., Dehaene, S., Adam, C., Cle´menceau, S., & Hasboun, D. (2009). Converging intracranial markers of conscious access. *PLoS Biol* 7.
- Gelman, R. S. (1969). Conservation acquisition: a problem of learning to attend to relevant attributes. *Journal of Experimental Child Psychology*, 7, 167–187.
- Haerich, P. (1994). Startle Reflex Modification: Effects of Attention Vary With Emotional Valence. *Psychological Science*, 407-410.
- Häkkinen, M. (2010). *Why Alarms Fail: A Cognitive Explanatory Model*. University of Jyväskylä.
- Harnad, S. (1990). The Symbol Grounding Problem. *Physica D* 42, 335-346.
- Hintzman, D. L. (2011). Research Strategy in the Study of Memory: Fads, Fallacies, and the Search for the ,
 ÅúCoordinates of Truth, Åù. *Perspectives on Psychological Science*, 6(3), 253-271.
- Hofstadter, D., R., & Mitchell, M. (1994). The Copycat Project: A model of mental fluidity and analogy-making. *In advances in connectionist and neural computation theory, Vol.2: logical connections, ed. K.J. Holyoak, and J.A. Barnden. Norwood N.J.: Ablex*.
- Iordanova, M. D., McNally, G. P., & Westbrook, R. F. (2006). Opioid Receptors in the Nucleus Accumbens Regulate Attentional Learning in the Blocking Paradigm. *Journal of Neuroscience*, 26, 4036-4045.
- Jackson, J. V. (1987). Idea for a mind. *SIGART Newsletter*, 23-26.
- James, W. (1890). *The Principles of Psychology*. Cambridge, MA: Harvard University Press.
- Kanwisher, N. (2001). Neural events and perceptual awareness. *Cognition*, 79:89(79), 89-113.
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology* 4(4):, 219-227.
- Kruschke, J. K. (2010). Models of attentional learning. *In: E. M. Pothos and A. Wills (eds.), Models of Learning and Categorization, Cambridge University Press*.
- Lachman, R., Tatsuoka, M., & Bonk, W. J. (1961). Human behavior during the tsunami of May. *Science*, 133, 1405-1409.
- Langley, P., Laird, J. E., & Rogers, S. (2008). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*. <http://csl.stanford.edu/research/ongoing/icarus/>.
- Leach, P. (1997). *Your Baby and Child: From Birth to Age Five*. New York: Alfred A. Knopf.
- LeDoux, J. E. (2000). Emotion circuits in the brain. *Annual Review of Neuroscience*, 23, 155-184.
- Lewin, K. (1951). Field theory in social science: selected theoretical papers. *New York: Harper & Row*.

- Liddell, B. J., Brown, K. J., Kemp, A. H., Barton, M. J., Das, P., Peduto, A., et al. (2005). A direct brainstem–amygdala–cortical alarm system for subliminal signals of fear. *NeuroImage* 24, 235–243.
- Madl, T., Baars, B. J., & Franklin, S. (2011). The Timing of the Cognitive Cycle. *PLoS ONE*.
- Maes, P. (1989). How to do the right thing. *Connection Science*, 1, 291–323.
- Mateo, J. M. (2010). Alarm calls elicit predator-specific physiological responses. *Biol. Lett.*, 6, 623–625.
- Merker, B. (2005). The liabilities of mobility: A selection pressure for the transition to consciousness in animal evolution. *Consciousness and Cognition*, 14, 89–114.
- Miller, E., & Cohen, J. (2001). An integrative theory of prefrontal cortex function. *Annu Rev Neurosci*, 24 (1), 167–202.
- Miller, S. M., & Fu, W.-T. (2007). The role of temporal sequence learning in guiding visual attention allocation. *Proceedings of the human factors and ergonomics society 51st annual meeting*, 1368–1372.
- Minsky, M. (1986). *Society of Mind*. New York: Simon and Schuster.
- Negatu, A., & Franklin, S. (2002). An action selection mechanism for 'conscious' software agents. *Cognitive Science Quarterly*, 2, 363–386.
- Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press.
- Ogawa, H., & Yagi, A. (2002). The Implicit Processing in Multiple Object Tracking. *Technical Report on Attention and Cognition*, 10, 1–4.
- Ornstein, R. (1986). *Multimind*. Boston: Houghton Mifflin.
- Phelps, E. A. (2006). Emotion and Cognition: Insights from studies of the human amygdala. *Annual Review of Psychology*, 57, 27–53.
- Posner, M. I. (2004). *Cognitive Neuroscience of Attention*: The Guilford Press.
- Posner, M. I. (2009). Michael Posner on the Anatomy of attentional networks – a historical perspective. *GoCognitive*.
- Posner, M. I., & Fan, J. (2004). ATTENTION AS AN ORGAN SYSTEM. In J. R. Pomerantz & M. C. Crair (Eds.), *Topics in Integrative Neuroscience: From Cells to Cognition*. Cambridge UK: Cambridge University Press.
- Posner, M. I., & Rothbart, M. K. (1998). Attention, self-regulation and consciousness. *Phil. Trans. R. Soc. Lond. B* 353 (1377), 1915–1927.
- Posner, M. I., Rothbart, M. K., & Sheese, B. E. (2007). The anterior cingulate gyrus and the mechanism of self-regulation. *Cognitive, Affective, & Behavioral Neuroscience*, 7 (4), 391–395.
- Posner, M. I., Snyder, C. R. R., & Davidson, D. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, 109, 160–174.
- Purves, D., Brannon, E., Cabeza, R., Huettel, S. A., LaBar, K., Platt, M., et al. (2008). *Principles of cognitive neuroscience*. Sunderland, Massachusetts: Sinauer Associates.
- Ratey, J. J. (2001). *A User's Guide to the Brain: Perception, Attention, and the Four Theaters of the Brain*. Vintage Books.
- Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61(2), 168–185.
- Rolls, E. T. (2000). Neurophysiology and functions of the primate amygdala, and the neural basis of emotion. In *The Amygdala: a Functional Analysis* (ed. Aggleton, J. P.). Oxford Univ. Press, Oxford, UK, 447–478.
- Samsonovich, A. V. (2010). Toward a Unified Catalog of Implemented Cognitive Architectures. *Proceeding of the 2010 Conference on Biologically Inspired Cognitive Architectures*, 195–244.
- Sarter, M., Givens, B., & Bruno, J. P. (2001). The cognitive neuroscience of sustained attention: where top-down meets bottom-up. *Brain Research Reviews* 35, 146–160.
- Schneider, W., & Chein, J. M. (2003). Controlled & automatic processing: behavior, theory, and biological mechanisms. *Cognitive Science Quarterly*, 27, 525–559.
- Shanahan, M. (2010). *Embodiment and the Inner Life*. Oxford: Oxford University Press.
- Sloman, A. (1987). Motives Mechanisms Emotions. in *Emotion and Cognition* 1(3), 217–234.
- Sloman, A. (1998). Damasio, Descartes, Alarms and Meta-management. *Proceedings Symposium on Cognitive Agents: Modeling Human Cognition*. San Diego: IEEE.
- Squire, L. R., & Kandel, E. R. (2000). *Memory: From Mind to Molecules*: W. H. Freeman.
- Thorne, G. C., & Alice, T. (2003). What is Attention? , June 2006, from www.cdl.org/resource-library/articles/attention2.php?
- Van Bockstaele, B., Verschuere, B., Houwer, J. D., & Crombez, G. (2010). On the costs and benefits of directing attention towards or away from threat-related stimuli: A classical conditioning experiment. *Behaviour Research and Therapy*, 48, 692–697.

- Vidnyánszky, Z., & Sohn, W. (2003). Attentional learning: learning to bias sensory competition. *Journal of Vision*, 3(174a).
- Xiao, Y., Seagull, F. J., Nieves-Khouw, F., Barczak, N., & Perkins, S. (2004). Organizational-historical analysis of the failure-to-respond-to-alarms problems. *IEEE TRANSACTIONS on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 34(6).
- Yoshida, H., & Smith, L. B. (2003). Known and novel noun extensions: Attention at two levels of abstraction. *Child Development*, 76, 564-577.
- Ziemke, T., Zlatev, J., & Frank, R. M. (2007). *Body, Language and Mind: Volume 1: Embodiment*: Mouton de Gruyter.