New Developments in the LIDA Model

By Stan Franklin, Steve Strain, Sean Kugele, Tamas Madl, Nisrine Ait Khayi & Kevin Ryan

Abstract

Research on our LIDA Model of how minds work is now well into its second decade. Here we briefly summarize new developments concerning research on the model. One of these is the appearance of a tutorial article that offers a gentle way towards understanding the model. Another describes an outsider's view of LIDA's contributions. The list continues with descriptions of various stages of research on action execution, spatial memory, motivation, self, multi-cyclic processes, language, brain rhythms, mental imagery, and a computational framework.

Introduction

Systems level cognitive models are intended to model minds, which we take here to be control structures (Franklin, 1995, p.412) for autonomous agents (Franklin & Graesser, 1997). The LIDA (Learning Intelligent Decision[1] Agent) systems level cognitive model is intended to model human minds, some animal minds, and some artificial minds, be they software agents or robots. LIDA is a conceptual and partly computational model that serves to implement and flesh out a number of psychological theories (Baddeley, 1993; Barsalou, 1999; Conway, 2001; Ericsson & Kintsch, 1995; Glenberg, 1997; Minsky, 1985; Sloman, 1999), in particular the Global Workspace Theory of Baars (1988). Hence any LIDA agent, that is any agent whose control structure is based on the LIDA Model, is at least functionally conscious (Franklin, 2003). Research on LIDA has entered its second decade (Franklin & Patterson, 2006). This note is intended to summarize some of the newer developments of the LIDA Model.

The LIDA Tutorial

The LIDA Model is quite complex consisting of numerous independently and asynchronously operating modules (see Figure 1 below). It has been described in more than fifty published papers, presenting a considerable challenge to any would be student of the model. Thus the recent appearance of a LIDA tutorial paper (Franklin et al., 2016) summarizing the contents of these earlier papers, as well as new material, is a significant new LIDA development. The tutorial reduces the fifty some odd papers into only fifty some odd pages of text and figures.

---

[1] For historical reasons, this word was previously "distribution". It has been recently changed to better capture important aspects of the model in its name.

AI: Its Nature and Future

In 2016, Oxford University Press published philosopher/cognitive scientist Margaret Boden's *AI: Its Nature and Future* (2016), which pays considerable attention to our LIDA Model.

Pointing out that LIDA "…arises from a unified, systems-level theory of cognition…", Boden goes on to speak of LIDA as being "…deeply informed by cognitive psychology, having been developed for scientific, not technological, purposes." and "…designed to take into account a wide variety of well-known psychological phenomena, and a wide range of experimental evidence …" She says that "Integrating highly diverse experimental evidence…" LIDA is used "…to explore theories in cognitive psychology and neuroscience." She also says that "…the *philosophical* significance of LIDA, for instance, is that it specifies an organized set of virtual machines that shows how the diverse aspects of (functional) consciousness are possible." And, Boden points out that the LIDA Model speaks to the "binding" problem, to the frame problem, and avoids any central executive.

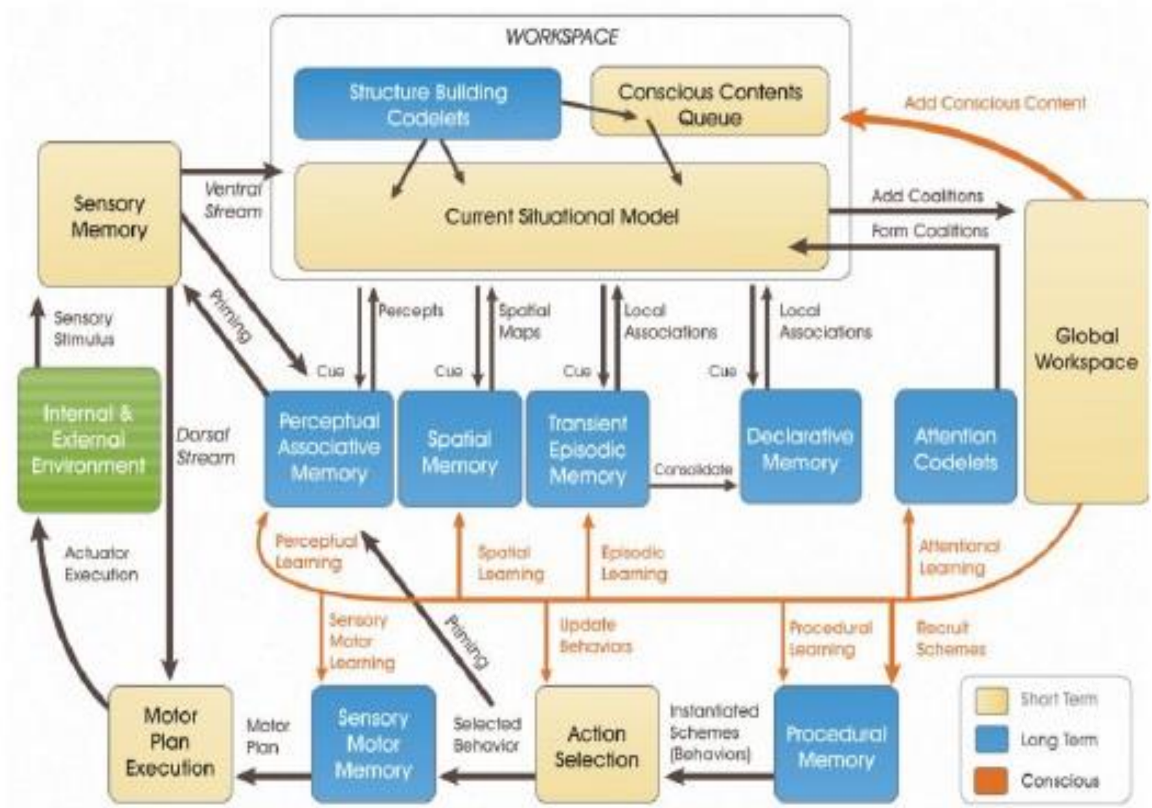

Figure 1: The LIDA Cognitive Cycle

Action Execution

The LIDA Model attempts to model minds generally, providing an architecture for the control structure of any number of different LIDA-based agents. Thus, the LIDA Model in its general form must remain uncommitted to particular mechanisms or specifications for senses, actions, and environments. Each of its many independent and asynchronous modules, mentioned above, must allow for implementation so as to serve various agents with a variety of senses, actions and environments.

Two of LIDA's most recently developed modules are devoted to action execution, which is concerned with creating a motor plan for a selected goal-directed behavior, and executing it. A motor plan template transforms a selected behavior into a sequence of executable actions. The Sensory Motor Memory (see Figure 1 above) learns and remembers motor plan templates (Dong & Franklin, 2014). Based on the subsumption architecture (Brooks, 1986), our LIDA agent testing this module adds analogs of the visual system's dorsal and ventral streams to the Model. Given an appropriate motor plan for the selected behavior, the Motor Plan Execution module instantiates a suitable motor plan, and executes it (Dong & Franklin, 2015b). Together the two modules allow a LIDA-based agent to execute a selected action, quite important for any autonomous agent.

We have also introduced a new type of sensorimotor learning to the LIDA Model (Dong & Franklin, 2015a). Using reinforcement learning, it stores and updates the rewards of pairs of data, motor commands and their contexts, allowing the agent to output effective commands based on its reward history. As is all learning in LIDA, this sensorimotor learning is cued by the agent's conscious content. A dynamic learning rate controls the effect of the newly arriving reward. The mechanism controlling the learning rate is inspired by the memory of errors hypothesis from neuroscience (Herzfeld, Vaswani, Marko, & Shadmehr, 2014). Our computer simulations indicate that using such a dynamic learning rate improves movement performance.

Spatial Memory

In any cognitive system, memory is most generally defined as the encoding, storing and recovery of information of some sort. The storage can be over various time scales. Cognitive modelers, and cognitive scientists in general, tend to divide the memory pie in many different ways. The LIDA Model has separate, asynchronous, modules for memory systems of diverse informational types. (In Figure 1, the modules for various long-term memory systems are dark colored.) Much earlier research was devoted to Perceptual Associative Memory, Transient Episodic Memory, Declarative Memory, and Procedural Memory. (In all these cases, there is much left to be done.) Recent work on Sensory Motor Memory was discussed in the preceding section.

Over the past couple of years we have begun to think seriously about how best to represent data in Spatial Memory, representations of spatial information concerning objects in the agent's environment, and its location within it. We picture long-term Spatial Memory as consisting of hierarchies of cognitive maps, each representing the size, shape and location of objects, and the directions and distances between them. In addition to long-term spatial memory, LIDA's working memory may contain one or few cognitive map segments and facilitate planning and updating. Inspired by place and grid cells involved in spatial representations in mammalian brains, cognitive map representations in LIDA also consist of hierarchical grids of place nodes, which can be associated with percepts and events. We have implemented prototype mechanisms for probabilistic cue integration and error correction, to mitigate the problems associated with accumulating errors from noisy sensors (see section on uncertainty below). So far we have only experimented with how human agents mentally represent such cognitive maps of neighborhoods (Madl, Franklin, Chen, Montaldi, & Trappl, 2016; Madl, Franklin, Chen, Trappl, & Montaldi, 2016).

Motivation

Every autonomous agent, be it human, animal or artificial, must act in pursuit of its own agenda (Franklin & Graesser, 1997). To produce that agenda requires motivation. Actions in the LIDA Model are motivated by feelings, including emotions, that is feelings with cognitive content (Johnston, 1999). An early paper lays this out, and relates feelings in this context to both values and utility (Franklin & Ramamurthy, 2006). More recent work fleshes out just how feelings play a major role in motivating the choice of actions (McCall, 2014; McCall, Franklin, Faghihi, & Snaider, submitted). Feelings arise in Sensory Memory (see Figure 1), are recognized in Perceptual Associative Memory, and become part of the percept that updates the Current Situational Model. There they arouse structure building codelets to produce various options advocating possible responses to the feeling (in accordance with appraisal theories of emotion (Franklin, et al., 2016)). The most salient of these likely wins the competition for consciousness in the Global Workspace, and is broadcast, in particular to Procedural Memory. There schemes proposing specific actions to implement the broadcast option are instantiated and forwarded to Action Selection, where a single action is selected as a response to the original feeling. Thus feelings act as motivators.

Self

Any systems-level cognitive model such as our LIDA Model that aspires to model consciousness, must attempt to account for the notion of self with its multiple aspects. We have made one attempt at describing how a number of different “selves” could be constructed within the LIDA Model (Ramamurthy & Franklin, 2011). These include the minimal (or core) self with its three sub-selves, self as subject, self as experiencer and self as agent. The sub-selves of the extended self are comprised of the autobiographical self, the self-concept, the volitional (or executive) self, and the narrative self.

More recently we have begun to augment this account by combining these constructs with key elements of Shaun Gallagher’s pattern theory of self (2013), namely his meta-theoretical list of aspects. These include minimal embodied aspects, minimal experiential aspects, affective aspects, intersubjective aspects, psychological/cognitive aspects, narrative aspects, extended aspects, and situated aspects. We explore the use of the various aspects of this pattern theory of self in producing each of the various selves within the LIDA Model. The three types of minimal self are all implemented using only minimal embodied aspects and minimal experiential aspects. All of these can be created within the current LIDA Model. The four types of extended self will require all eight aspects in the list. Some of these will require additional processes to be added to the LIDA Model.

This use of pattern theory is helping us to clarify various theoretical issues with including various "selves" in the LIDA Model, as well as open questions such as the relationships between different sub-selves. Using pattern theory also can enable us to set benchmarks for testing for the presence of various types of self in different LIDA-based agents.

Cyclic to Multicyclic Processes

The LIDA Model begins its fleshing out of Global Workspace Theory by postulating a cognitive cycle (see Figure 1 for a detailed diagram), which we view as a cognitive atom from which more complex cognitive processes are constructed. A LIDA agent spends its "life" in a continual, cascading (overlapping) sequence of such cognitive cycles, each sensing and understanding the agent's current situation, and choosing and executing an appropriate response. Such cycles occur five to ten times a second in humans (Madl, Baars, & Franklin, 2011). The first decade or more of our research was devoted to trying to understand what happens during a single cognitive cycle, taking in humans 200 to 500 ms. Now, having at least a partial overall discernment of the activity of a single cycle, we feel emboldened to turn some of our attention to more complex multi-cyclic processes such as planning, reasoning, and deliberation.

Language

LIDA has been criticized for focusing on low intelligence tasks, and lacking high cognitive functions such as language understanding[2] (Duch, Oentaryo, & Pasquier, 2008). To overcome this gap, and initiate language processing in the LIDA architecture, learning the meaning of the vervet monkey alarm calls was simulated. Field studies (Seyfarth, Cheney, & Marler, 1980) revealed the existence of three distinct alarm calls. Each call is emitted to warn the rest of the group of the danger from a predator in the vicinity. Upon hearing a particular alarm call, vervet monkeys typically escape into safe locations in a manner appropriate to the predator type signaled by that alarm. A LIDA based agent that learns the meaning of these alarm calls has been developed (Khayi-Enyinda, 2013). LIDA's perceptual learning mechanism was implemented to associate each alarm call with three distinct meanings: an action based meaning, a feeling based meaning and a referential based meaning. This multiple meaning assessment approach aligns with our

[2] LIDA'S predecessor, IDA, communicated with sailors in pseudo-natural language, pseudo in the sense that it was hand-crafted rather than learned, and that it only dealt with a very narrow domain.

ultimate goal of modeling human words that must convey multiple meanings. A manuscript describing this research has been submitted, reviewed, revised, and re-submitted (Ait Khayi & Franklin, 2018 to appear).

LIDA's Hypothesis Regarding Brain Rhythms

Marr proposed three levels of analysis for cognitive modeling—the computational, the representational/algorithmic, and the implementational (1982). As a general model of minds, LIDA's core concepts possess an applicability that spans many possible domains and implementations. Accordingly, LIDA's primary area of interest lies within Marr's computational and algorithmic levels. However, many classes of biological mind fall within LIDA's purview, and modeling biological minds from the perspective of the LIDA Model requires careful attention to the available evidence and the competing theories regarding the way in which brains effect control structures for behavior in humans and certain non-human animals.

A helpful metaphor may be found in the example problem of reverse engineering a software program. The primary goal is to uncover the algorithms that carry out the software's computations, but this might require, or at least be facilitated by investigation of the operations carried out in the hardware during the program's execution. We frequently assert that LIDA is a model of minds rather than brains. Having said that, we find that understanding those biological minds of interest to LIDA requires close and frequent reference to the way brains carry out computations. In practice, this has meant examination of biological minds at the implementation level as well as the algorithmic and computational levels.

While neuroscience manifests a solid theoretical consensus regarding the basic tenets of neuroanatomy and neuronal physiology, considerable controversy continues to pervade investigations into the cognitive aspects of neural function. The vast proliferation of evidence resulting from recent decades' technological advances have thus far failed to converge on a consensual framework for understanding the neural basis of cognition. Nonetheless, LIDA's perspective on biological minds currently commits to a particular collection of theoretical proposals situated squarely within the broader controversy. While a detailed treatment of these proposals lies outside the scope of the present discussion, we give a brief overview as follows.

The Cognitive Cycle Hypothesis and the Global Workspace Theory (GWT) of Consciousness form the backbone of the LIDA Model. GWT, originally a psychological theory (Baars, 1988), was recently updated (Baars, Franklin, & Ramsøy, 2013) into a neuropsychological theory known as Dynamic Global Workspace Theory (dGWT). Per dGWT, a global workspace is "a dynamic capacity for binding and propagation of neural signals over multiple task-related networks, a kind of neuronal cloud computing" (Baars, et al., 2013, p. 1). Per LIDA's Cognitive Cycle Hypothesis, the global workspace produces a quasiperiodic broadcast of unitary and internally consistent cognitive content that mediates an autonomous agent's action selection and learning, and, over time, comprises the agent's stream of consciousness.

The theoretical proposals of Freeman's Neurodynamics (Freeman, 2012; Freeman & Kozma, 2010) provide the framework most harmonious with LIDA's central hypotheses. Within this framework, a cognitive cycle comprises the emergence of a self-organized pattern of neurodynamic activity. LIDA's Rhythms Hypothesis postulates that the content of a cycle's broadcast from the global workspace manifests in experimentally observable brain rhythms as gamma (30-80 Hz) frequency activity scaffolded within a slow-wave structure of approximately theta (4-6 Hz) frequency that tracks the rhythm of successive broadcasts. Elaboration of this hypothesis within the framework of Freeman's neurodynamical theory is quite complex and is the subject of a publication currently under preparation.

Mental Imagery, Preconscious Simulation, and Grounded Cognition

Most humans report the ability to have sensory-like experiences in the absence of external stimuli. They describe experiences such as "having a song stuck in our heads" or "listening to our inner voices" or "seeing with our mind's eye". In the literature cited below, these phenomena are referred to as "mental imagery". Many experiments have been performed that suggest mental imagery facilitates, and may be critical for, a broad range of mental activities including prediction (Moulton & Kosslyn, 2009), problem solving (Qin & Simon, 1992; Shaver, Pierson, & Lang, 1975), mental rehearsal (Driskell, Copper, & Moran, 1994; Keller, 2012), and language comprehension (Bergen, Lindsay, Matlock, & Narayanan, 2007; Zwaan, Stanfield, & Yaxley, 2002). Cognitive models are needed to help explain the processes that underlie mental imagery. We have begun to leverage the LIDA model to gain insight into how the fundamental capabilities needed for mental imagery could be realized in artificial minds, and to apply these insights toward the construction of software agents that utilize mental imagery to their advantage.

Mental imagery is by definition a conscious process; however, there is an intriguing possibility that the same mechanisms underlying mental imagery also support preconscious cognitive processes, and enable grounded (embodied) cognition. The psychologist and cognitive scientist Lawrence Barsalou defines "simulation" as the "re-enactment of perceptual, motor, and introspective states acquired during experience with the world, body, and mind", and hypothesizes that

> [simulation] is not necessarily conscious but may also be unconscious, probably being unconscious even more often than conscious. Unconscious [simulations] may occur frequently during perception, memory, conceptualization, comprehension and reasoning, along with conscious [simulations]. When [simulations] reach awareness, they can be viewed as constituting mental imagery... (Barsalou, 2009)

It is a goal of our research program to explore the possibility of a unified set of mechanisms supporting mental imagery, preconscious simulation, and grounded cognition. The LIDA model provides an ideal foundation for exploring these topics, as it is one of the few biologically inspired cognitive architectures that attempts to model functional consciousness, and is firmly committed to grounded cognition (Franklin, Strain, McCall, & Baars, 2013).

Representing and computing with uncertainty in LIDA

Cognition must deal with large amounts of uncertainty, due to a partially observable environment, erroneous sensors, noisy neural computation, and limited cognitive resources. There is increasing evidence for probabilistic mechanisms in brains (Chater, Tenenbaum, & Yuille, 2006; Clark, 2013; Knill & Pouget, 2004). We have recently started exploring probabilistic computation for LIDA, as of now for the specific purpose of dealing with spatial uncertainty and complexity in navigation (Madl, Franklin, Chen, Montaldi, et al., 2016; Madl, Franklin, Chen, Trappl, et al., 2016). Work is underway to augment LIDA's representations (inspired by Barsalou's perceptual symbols and simulators (Barsalou, 1999)) with a representation and computation mechanism accounting both for the uncertainty in various domains, as well as approximately optimal inference given cognitive, time and memory limitations.

LIDA Framework in Python

In 2011, Snaider et al. presented the "LIDA Framework", a software framework written in the Java programming language that aims to simplify the process of developing LIDA agents. The LIDA Framework implements much of the low-level functionality that is needed to create a LIDA software agent, and provides default implementations for many of the LIDA modules. As a result, simple agents can often be created with a modest level of effort by leveraging "out of the box" functionality.

Inspired by the success of the LIDA Framework, a sister project is underway to implement a software framework in the Python programming language, which we refer to as lidapy. One of lidapy's primary goals has been to facilitate the creation of LIDA agents that are situated in complex and "real world" environments, with the eventual goal of supporting LIDA agents in a robotics context. Towards this end, lidapy has been designed from the ground up to integrate with the Robot Operating System (Quigley et al., 2009), a framework developed by the Open Source Robotics Foundation (OSRF) that was specifically designed to support large-scale software development in the robotics domain.

References

Ait Khayi, N., & Franklin, S. (2018 to appear). Initiating language in LIDA: learning the meaning of vervet alarm calls. *Biologically Inspired Cognitive Architectures*.
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.
Baddeley, A. D. (1993). Working memory and conscious awareness. In A. Collins, S. Gathercole, Martin A. Conway & P. Morris (Eds.), *Theories of memory* (pp. 11–28). Howe: Erlbaum.
Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577–609.
Barsalou, L. W. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 364*(1521), 1281-1289.
Bergen, B. K., Lindsay, S., Matlock, T., & Narayanan, S. (2007). Spatial and linguistic aspects of visual imagery in sentence comprehension. *Cognitive science, 31*(5), 733-764.
Boden, M. A. (2016). *AI: Its Nature and Future*. Oxford, UK: Oxford University Press.
Brooks, R. (1986). A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation, RA-2*(1), 14-23.
Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in cognitive sciences, 10*(7), 287-291.
Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36*(03), 181-204.
Conway, Martin A. (2001). Sensory–perceptual episodic memory and its context: autobiographical memory. *Philos. Trans. R. Soc. Lond B., 356*, 1375–1384.
Dong, D., & Franklin, S. (2014). *Sensory Motor System: Modeling the process of action execution.* Paper presented at the Proceedings of the 36th Annual Conference of the Cognitive Science Society.
Dong, D., & Franklin, S. (2015a). Modeling sensorimotor learning in LIDA using a dynamic learning rate. *Biologically Inspired Cognitive Architectures, 14*, 1-9.
Dong, D., & Franklin, S. (2015b). A New Action Execution Module for the Learning Intelligent Distribution Agent (LIDA): The Sensory Motor System. *Cognitive Computation*. doi: 10.1007/s12559-015-9322-3.
Driskell, J. E., Copper, C., & Moran, A. (1994). Does mental practice enhance performance? : American Psychological Association.
Duch, W., Oentaryo, R., & Pasquier, M. (2008). Cognitive Architectures: Where Do We Go From Here? In P. Wang, B. Goertzel & S. Franklin (Eds.), *Artificial general Intelligence, 2008: proceedings of the First AGI Conference* (pp. 122-137).
Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102*, 211–245.
Franklin, S. (1995). *Artificial Minds*. Cambridge, Ma: MIT Press.

Franklin, S. (2003). IDA: A Conscious Artifact? *Journal of Consciousness Studies, 10*, 47–66.

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., Madl, T., Strain, S., Faghihi, U., Dong, D., Kugele, S., . . . Chen, S. (2016). A LIDA cognitive model tutorial. *Biologically Inspired Cognitive Architectures*, 105-130. doi: 10.1016/j.bica.2016.04.003

Franklin, S., & Patterson, F. G. J. (2006). The LIDA Architecture: Adding New Modes of Learning to an Intelligent, Autonomous, Software Agent *IDPT-2006 Proceedings (Integrated Design and Process Technology)*: Society for Design and Process Science.

Franklin, S., & Ramamurthy, U. (2006). Motivations, Values and Emotions: Three sides of the same coin *Proceedings of the Sixth International Workshop on Epigenetic Robotics* (Vol. 128, pp. 41–48). Paris, France: Lund University Cognitive Studies.

Franklin, S., Strain, S., McCall, R., & Baars, B. (2013). Conceptual Commitments of the LIDA Model of Cognition. *Journal of Artificial General Intelligence, 4*(2), 1-22. doi: 10.2478/jagi-2013-0002

Freeman, W. (2012). *Neurodynamics: an exploration in mesoscopic brain dynamics*: Springer Science & Business Media.

Freeman, W. J., & Kozma, R. (2010). Freeman's mass action. *Scholarpedia, 5*(1), 8040.

Gallagher, S. (2013). A pattern theory of self. *Frontiers in Human Neuroscience, 7*.

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

Herzfeld, D. J., Vaswani, P. A., Marko, M. K., & Shadmehr, R. (2014). A memory of errors in sensorimotor learning. *Science, 345*(6202), 1349-1353.

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

Keller, P. E. (2012). Mental imagery in music performance: underlying mechanisms and potential benefits. *Annals of the New York Academy of Sciences, 1252*(1), 206-213.

Khayi-Enyinda, N. A. (2013). *Learning the meaning of the vervet alarm calls using a cognitive and computational model.* Master of Science, University of Memphis, Memphis.

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences, 27*(12), 712-719.

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., Chen, K., Montaldi, D., & Trappl, R. (2016). Towards real-world capable spatial memory in the LIDA cognitive architecture. *Biologically Inspired Cognitive Architectures, 16*, 87-104. doi: 10.1016/j.bica.2016.02.001

Madl, T., Franklin, S., Chen, K., Trappl, R., & Montaldi, D. (2016). Exploring the structure of spatial representations. *PLoS ONE*.

Marr, D. C. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman.

McCall, R. (2014). *Fundamental Motivation and Perception for a Systems-Level Cognitive Architecture.* PhD Thesis, University of Memphis, Memphis, TN USA.
McCall, R. J., Franklin, S., Faghihi, U., & Snaider, J. (submitted). Artificial Motivation for Cognitive Software Agents.
Minsky, M. (1985). *The Society of Mind*. New York: Simon and Schuster.
Moulton, S. T., & Kosslyn, S. M. (2009). Imagining predictions: mental imagery as mental emulation. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 364*(1521), 1273-1280.
Qin, Y., & Simon, H. A. (1992). *Imagery and mental models in problem solving.* Paper presented at the Proc. AAAI Symposium on Reasoning with Diagrammatic Representations., Stanford, CA.
Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., . . . Ng, A. Y. (2009). *ROS: an open-source Robot Operating System.* Paper presented at the ICRA workshop on open source software.
Ramamurthy, U., & Franklin, S. (2011). *Self System in a model of Cognition.* Paper presented at the Machine Consciousness Symposium at the Artificial Intelligence and Simulation of Behavior Convention (AISB'11, University of York, UK.
Seyfarth, R., Cheney, D., & Marler, P. (1980). Monkey responses to Three Different Alarm Calls: Evidence of Predator Classification and Semantic Communication. *Science, 210*(4471), 801-803.
Shaver, P., Pierson, L., & Lang, S. (1975). Converging evidence for the functional significance of imagery in problem solving. *Cognition, 3*(4), 359-375.
Sloman, A. (1999). What Sort of Architecture is Required for a Human-like Agent? In M. Wooldridge & A. S. Rao (Eds.), *Foundations of Rational Agency* (pp. 35–52). Dordrecht, Netherlands: Kluwer Academic Publishers.
Snaider, J., McCall, R., & Franklin, S. (2011). *The LIDA Framework as a General Tool for AGI*. Paper presented at the Artificial General Intelligence (AGI-11), Mountain View, CA.
Zwaan, R. A., Stanfield, R. A., & Yaxley, R. H. (2002). Language comprehenders mentally represent the shapes of objects. *Psychological science, 13*(2), 168-171.