7-29-2013

Learning the Meaning of the Vervet Alarm Calls using a Cognitive and Computational Model

Nisrine Ait Khayi-Enyinda

LEARNING THE MEANING OF THE VERVET ALARM CALLS USING A
COGNITIVE AND COMPUTATIONAL MODEL

by

Nisrine Ait Khayi-Enyinda

A Thesis

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Master of Science

Major: Computer Science

The University of Memphis

August 2013

ACKNOWLEDGMENTS

ABSTRACT

This thesis explains how the infant vervet, *Chlorocebus pygerthrus*, learns the meaning of vervet alarm calls using the Learning Intelligent Distribution Agent's (LIDA) perceptual learning mechanism. We consider an approach of multiple meanings which correspond to a feeling-based meaning, an action-based meaning, and a referential meaning. The first part of simulations was performed to test the learning of the meaning of these alarm calls while the infant is attached physically to the mother. The second part of simulations was performed to study the infant's understanding of these alarm calls while the infant is detached physically from the mother. The results show that a LIDA-based agent is capable to learn such multiple meanings. The agent learned in sequence the feeling-based meaning, the action-based meaning, and the referential meaning. The LIDA agent achieved a good performance of understanding. This was verified by checking the correct escape action after hearing a specific alarm call.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

Many researchers in the language evolution field have agreed on the existence of an early form of communication preceding human language (Bickerton 1990, 1996; Wray, 1998) known as protolanguage. According to Bickerton (1990), a protolanguage is a simple form of communication involving little structure, emerging from primate vocalizations by means of evolutionary pressures, perhaps eventually leading to a full-fledged human language. Bickerton also states that infantile human speech and protolanguage share common mechanisms and characteristics, such as a limited vocabulary. Chomsky and colleagues (1965) were among the few language theorists claiming that human language is entirely different from animal communication.

Because animal communication is a product of biological phenomena and the gradual evolution of processes involving neurobiology (Loula, Gudwin, Ribeiro & Queiroz, 2010b), modeling non-human primate communication may give insight into solving the problem of human language understanding. Oller and colleagues (2005) claim non-human primate communication systems belong to the fixed signals category. However, Campbell's monkey alarm calls contradict this claim (Lemasson, Gautier & Hausberger, 2003; Lemasson, Hausberger, & Zuberbühler, 2005). Oller and colleagues also suggests that natural selection couples a fixed signal to a function that is not modifiable in the individual.   For example, a primate call serving as an alarm cannot be reassigned as a courtship signal. There are a limited number of these functions, such as threat, greeting, contact, affiliation, invitation, etc. Fixed signals also appear in very early stages of human infant vocalization. Thus, modeling and implementing fixed signals could be a starting step toward modeling and implementing human language in a

cognitive architecture. For this purpose, we take the vervet alarm call system as a case study of animal communication, which belongs to fixed signals serving as a warning function against dangerous predators. Additionally, vervet alarm calls comprise a well-studied case among the primate communications.

Vervet monkeys are indigenous to Southern and East Africa. They are semiarboreal, inhabiting savanna, riverine woodlands, coastal forests and mountains in groups of up to 30 members. Field studies (including the play-back experiments done by Seyfarth and colleagues), revealed the existence of distinct vervet alarm calls (Seyfarth, Cheney, & Marler, 1980). These calls are acoustically distinct, and are used in different contexts. In this work, we focus on those serving a warning function of danger from predators. They are used by adults to warn the rest of the group of dangerous predators in the vicinity. Cheney and colleagues claimed that vervet alarm calls incorporate both reference to an object, as well as a disposition to behave toward that object in a particular way (Cheney & Seyfarth, 1997). They refer to a particular sort of immediate danger, and they function to designate particular classes of predators. In fact, vervet juveniles emit an eagle call for avian predators, a leopard call for the terrestrial predators, and a snake call for serpentine like objects. Each alarm call typically triggers a specific escape behavior into a location safe from a specific predator. Vervet adults climb to the top of trees in response to a leopard call, run to the bushes when an eagle call is sounded, and stand bipedally and look down and scan the area upon hearing a snake call. An important result of these experiments is that the vervet infants and juveniles often produce alarm calls in the wrong context. In fact, infants give eagle alarm calls to a very broad class of visual stimuli found in the air above (e.g., birds, failing leaves, etc.), leopard calls to

various terrestrial mammals, and snake calls to long and thin objects. Through time and experience, they gradually use the correct alarm calls, and they respond appropriately to each of them (Dorothy, Cheney, & Seyfarth, 1997; Zangenehpour, Ghazanfar, Lewkowicz & Zatorre, 2002). This provides direct evidence that vervet infants learn the meaning of these alarm calls.

In linguistics, a meaning represents the information conveyed by a sender in its message to the receiver, modified by any inference the receiver makes as a function of the current context. Controversy has permeated the debate about the meaning of vervet alarm calls. John Smith described vervet alarm calls as "referring to different escape actions," while the psychologist John Marshall (Dorothy, Cheney, & Seyfarth,1990) has averred on the basis of plausibility that vervet alarm calls refer to the predator type rather than the fearful emotions aroused by predators. To analyze the meaning process, several approaches have been used. Franklin introduced (1995) the quadratic understanding concept. In Franklin' words: "A system's understanding of a concept, or of collection of concepts, seems to vary with the complexity of its connections from the given concepts to other knowledge. Roughly, the more connections, the more understanding" (p.348).

According to this concept, each vervet alarm call can have multiple meanings, thus multiple connections. One connection is established from an alarm call to the corresponding predator, another one to the escape action, and another one to the fear feeling. Another meaning analysis process in biological and artificial systems is the semiotics of Charles S. Peirce. According to him signs can be classified to three classes: icons, indexes and symbols. These classes (icon, index and symbol) reflect the relationship between the sign and its object (Peirce, 1998). Icons look like their objects

(e.g. diagrams, portraits). Indexes are influenced directly by their objects (e.g. thermometer). Symbols hold an agreement-based relationship with their corresponding objects (e.g. alphanumeric). Hence, symbolic signs are established to convey various purposes in internal and external world. Human vocal communications are a well-known example of symbol sign systems. According to the Peircian classification, alarm calls operate in a specific way even in the absence of their referents. Thus, each alarm call is a symbol of a predator class. In this work, we studied multiples meanings of an alarm call including the referential meanings a la Peirce.

In recent decades, the use of computer simulations has increased in the language evolution field (Cangelosi, & Parisi, 2001). Computer simulation is a useful tool for studying language as a complex system (e.g., Steels, 1997), which has properties such as the emergence of the language, as well as simple signaling systems. In fact, linguistic behaviors emerge through the interaction between diverse components of the complex system, their neural, cognitive, communication abilities and their physical environment. In this work, we focused on the emergence of simple communication systems in an animal context using vervet monkeys. Using computer simulations has the benefit of testing the internal validity of theories by studying language or protolanguage as a complex system, and concluding how ecological factors, such as agents' spatial organization can influence the evolution of language and communication. However, a drawback of this approach is the simplifying assumptions required to decrease the computational cost, and the arbitrariness of some details. This can have an impact on the realism of the experiments as well as the results. In this work, we adopted a two-dimensional grid-based simulation composed of a main cognitive agent labeled INFANT

that learns the meaning of vervet alarm calls through interacting with other autonomous agents in a highly predatory environment. Further details about the environment, design, and implementation of the simulation will be provided later.

Another issue faced in research on communication evolution is related to the symbol grounding problem (Harnad & Glenberg, 1990; Robertson, 2000), namely, how the meaning of vocal symbols is acquired. In this work, we use a computational and cognitive model known as the Learning Intelligent Distributed Agent (LIDA) model, a cognitive architecture that controls autonomous software agents "living" in complex and dynamic environments. The major principle guiding LIDA is that every autonomous agent, be it human, animal, or artificial (e.g., software agent or robot), must frequently and continually sense its environment, interpret what it senses, and then act (Franklin & Graesser, 1997). LIDA is a hybrid system of cognition (Franklin et al., 2012), which blends various features from connectionist models and symbolic processing, with all symbols being grounded in the physical world in the sense of Brooks and Stein (Barsalou, 1999; R. Brooks & Stein, 1994). LIDA has various modules for perception, working memory, declarative memory, emotions, semantic memory, episodic memory, action selection, and conscious-like behavior. Despite the cognitive richness of the LIDA model that makes the realization of multiple human and primate tasks feasible, LIDA has been criticized as focusing on low level intelligence tasks such as object recognition, and lacking high level cognitive functions such as language understanding (Duch, Oentaryo & Pasquier, 2008). Our main contribution is beginning to overcome this gap by modeling vervet alarm calls. Accomplishing such work is a first step toward solving the human language understanding problem. Using the various LIDA cognitive modules, the

INFANT learner agent, controlled by LIDA, links the vocal symbols (vervet alarm calls) with external objects of its environment (predators), corresponding escape actions, and fear feeling. We assume all the objects and categories are grounded in its Perceptual Associative Memory.

This thesis is organized as follows: Chapter 2 explains the LIDA model and its cognitive cycle. Chapter 3 describes the LIDA-based perceptual learning mechanism. Chapter 4 briefly highlights the LIDA computational framework, especially the modules used in our simulation implementation. It then describes the design and the implementation of the two-dimensional grid environment. Finally, it explains the design and the implementation of the LIDA agent. Chapter 5 describes the experiments, their results and their interpretation. Finally, chapter 6 summarizes our work, describes our findings, and introduces some future directions.

2 The LIDA model and its cognitive cycle

The LIDA model is a systems-level, conceptual model that covers a large portion of human cognition while implementing some ideas of Global Workspace Theory (GWT) (Baars, 1988, 1997). Many pre-conscious processes are implemented by various codelets, which are small pieces of code, each running independently. These are specialized for some simple tasks, and often play the role of a daemon watching for an appropriate condition under which to act. These codelets operate asynchronously, independently of other processes in LIDA. There are several codelets classes in LIDA. One class is called structure building codelets (SBC) which are hypothesized to perform a number of central functions playing a role in the learning process. Each SBC can be seen as a daemon that is triggered when a specific pattern is matched in the Workspace. The SBC then responds by modifying existing structures or constructing new ones. A task example of a SBC is adding a new link between nodes. Another example could be creating a new instantiation of a node, or a new node for a category, object, event, feeling etc.

A SBC is implemented as a data structure. Each such codelet has the following attributes:

1. ***Base-level activation*** measures the usefulness of the SBC and it is modified by selectionist learning.

2. ***Context*** is the node structure or pattern that SBC is "looking for".

3. A***ction*** specifies what the codelet does when activated. It is typically short and performs a simple task.

In LIDA, there are several SBC types. We focus more on SBCs that add new referential and causality links in the Workspace's Current Situational Model CSM, from

the alarm calls nodes to their corresponding referents nodes, escape action nodes, and nodes representing fear feelings.

The LIDA model and its ensuing architecture are grounded in the LIDA cognitive cycle. The cognitive cycle (as described in Figure 1) is based on the fact that every autonomous agent (Franklin & Graesser, 1997) continually senses its environment, understands its current situation, and then selects an appropriate response (action). The agent's "life" can be regarded as consisting of a continual sequence of these cognitive cycles. Each cycle comprises three main phases of understanding, attending, and acting.
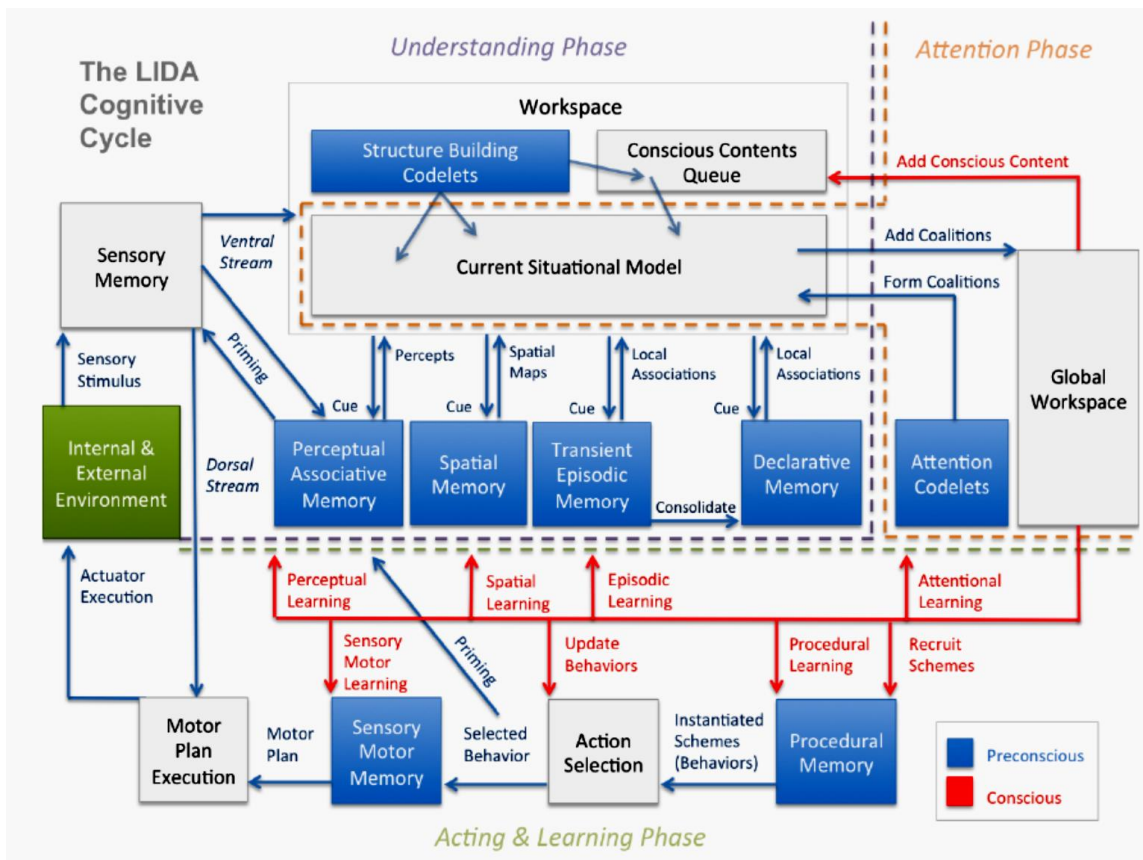
Figure 1. The LIDA cognitive cycle consisting of: 1) Understanding phase 2) Attention phase 3) Acting and learning phase.

Just as atoms have inner structure, the LIDA model hypothesizes a rich inner structure for its cognitive cycles (Baars & Franklin, 2003; Franklin, Baars, Ramamurthy, & Ventura, 2005; see Figure 1). What follows is a brief description of each phase of the cognitive cycle.

The understanding phase is initiated after receiving a sensory stimulus, which activates low level feature detectors that pre-process the received data, and add an initial meaning to it. The preprocessed data is sent directly to the Workspace or to the Perceptual Associative Memory (also called recognition memory) where higher level entities, such as objects, feelings, events, categories, relations etc. are recognized. The entities (nodes or links) in this long-term perceptual memory, whose activations rise over a threshold, form the current percept. This resulting percept is moved asynchronously to the preconscious Workspace. Here, a preconscious model of the agent's current situation, labeled the Current Situational Model (CSM), is updated. This percept and items from the Current Situational Model cue both Transient Episodic Memory and Declarative Memory (autobiographical and semantic) producing local associations from these short-term and long-term episodic memories. These local associations are combined with the percept to update the Current Situational Model. This process typically requires the SBCs, which have the role of monitoring the Workspace to fulfill their specified tasks as described previously. This newly updated model constitutes the agent's best understanding of its current situation within its world.

The attention phase starts when the attention codelets bring portions of the Workspace content to the Global Workspace by forming coalitions. All attention codelets are tasked with finding in the CSM structures matching their own content of concern.

Each attention codelet has the following properties (Faghihi, McCall, & Franklin, 2011): 1) *base level activation* which measures the codelet's usefulness in bringing information to consciousness, and is modulated through learning; 2) *concern:* the content, whose presence in CSM causes the codelet to act; and 3) *current activation* which reflects the saliency (e.g., novelty, urgency etc.) of its concern. A competition for consciousness among the formed coalitions, takes place in the Global Workspace in order to select the most salient and relevant coalition. The winning coalition is broadcasted globally. This causes the initiation of the acting and learning phase.

     The acting and learning phase involves multiple and parallel learning processes of the broadcasted conscious content as described above in Figure 1. The possible learning processes include:

- ***Perceptual learning*** occurs through learning new entities and relationships, and reinforcing old ones in LIDA's Perceptual Associative Memory once the conscious broadcast reaches the Perceptual Associative Memory.

- ***Procedural learning*** occurs through adding new action schemes, with their contexts and expected results, into the procedural memory. Old schemes can be reinforced.

- ***Episodic learning*** occurs through encoding new broadcasted events in the Transient Episodic Memory. When the consciousness mechanism broadcasts, its contents are encoded into Transient Episodic Memory (TEM), and may be later consolidated into LIDA's long-term Declarative Memory (DM) which stores the knowledge and facts as well as autobiographical memories.

- *Attentional learning* occurs through adding new attentional codelets or reinforcing the base-level activation of existing ones.

Procedural Memory is one of the primary recipients of this conscious broadcast. It stores templates of possible actions including their contexts and possible results, as well as an activation value that measures for each template the likelihood that a selected action within its context produces the expected result. Templates whose contexts match with the contents of the conscious broadcast, instantiate instances of themselves with their variables specified to the current situation. These instantiations are passed to the action selection mechanism, which chooses a single action from one of them. The chosen action then goes to Sensory Motor Memory, where it is executed by an appropriate algorithm called a motor plan. The action taken affects the environment, completing the cycle.

3   Learning the meaning of vervet alarm

Multiple Meaning of Vervet Alarm Calls

Multiple meanings of a concept refer to its multiple connections to other knowledge. The more connections, the more understanding. According to this approach to meaning assessment, multiple relationships should be built in the vervet mind from an alarm call to other concepts. Field experiments (Seyfart et al., 1980) revealed the occurrence of various events while the infant learns the meaning of various alarm calls. One event is the vocalizing of alarm calls by adult vervets. Another event is their executing specific escape actions into locations safe from predators.   Alarm calls also trigger some fearful reactions in the adult vervets such as body shaking and fearful face expressions. These events can be translated into two distinct causality relationships. The first one is between each alarm call and its corresponding escape action, and the second one is between each alarm call and the fear feeling. In addition, field studies revealed the referential functionality of vervet alarm calling system. Vervet alarm calls provide vervet listeners with sufficient contextual information to enable them to respond suitably to particular alarm calls as though they had direct evidence of the presence of the predator. This is implicit evidence that the referential relationship between each alarm call type and its corresponding predator class is already learned in the adult vervet's mind. Figures 2, 3, 4 describe various meanings of vervet alarm calls.
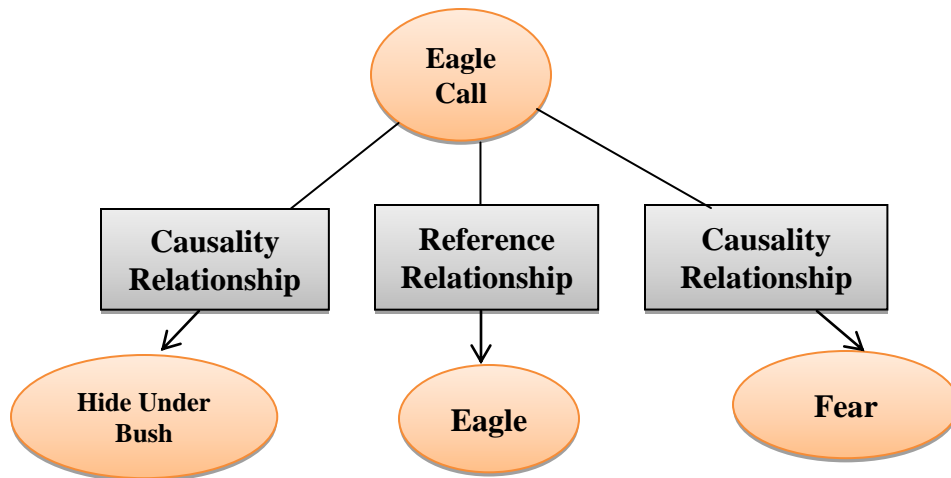
Figure 2. Eagle call meaning: consist of an eagle as a referent; fear and
hiding under bush as results of hearing an eagle call
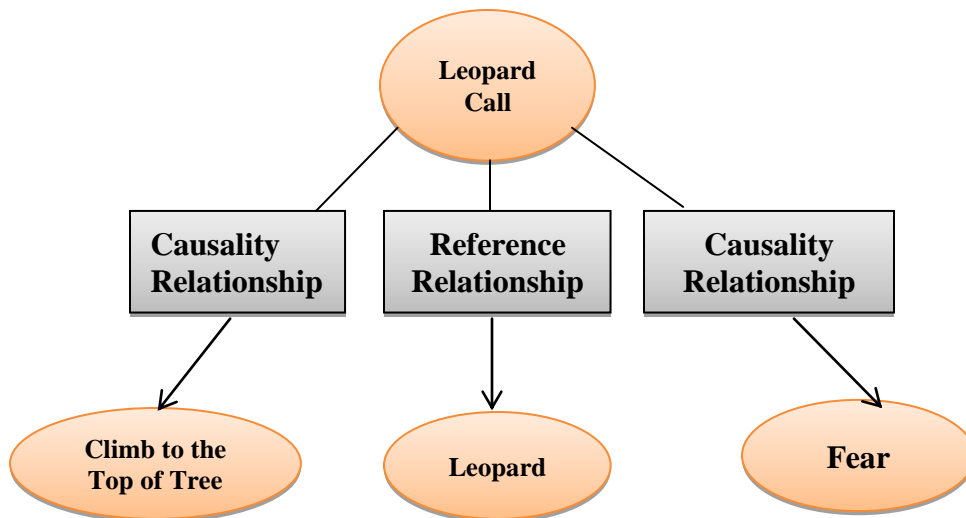
Figure 3. Leopard call meaning: consists of a leopard as a referent;
fear and climbing tree as results of hearing a leopard call
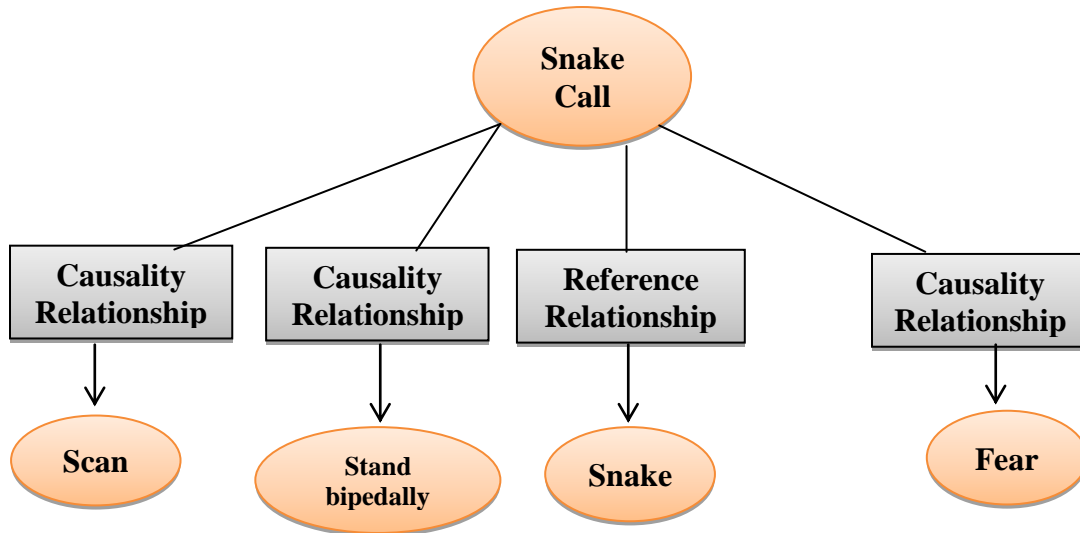
Figure 4. Snake call meaning: Consists of a snake as a referent; fear, standing bipedally, and scanning as a result of hearing a snake call

LIDA-based perceptual learning mechanism

Perceptual Associative Memory (PAM) is implemented in the LIDA architecture as a slipnet, a semantic net[1] with passing activation (Hofstadter & Mitchell, 1994). Perceptual learning in the LIDA model occurs with consciousness. It has two modes: *instructionlist mode*, which creates a new item for the first time in PAM with an initial amount of base level activation, and *selectionnist* mode which strengthens an existent item by reinforcing its base level activation using a sigmoid function (Edelman 1987).

Learning the meaning of vervet alarm calls occurs when new referential and causality relationships from the vervet alarm calls to the predators, escape actions, and fear feelings are established in the vervet's mind. In the LIDA terminology, we talk about adding a new referential link from each alarm call instance node to the corresponding

---

[1] Labels play no functional role in the LIDA cognitive architecture

predator instance node, and other causality links to the corresponding escape action and the vervet's fear instance node in the Workspace. These pre-conscious operations are implemented by Structure Building Codelets (SBC) which are classified into three categories: 1) Referential-Meaning Codelets 2) Fear-Meaning Codelets, and 3) Action-Meaning Codelets. Next, we describe the functionality of each codelet's class.

Referential-Meaning Codelets

In LIDA, nodes and links with high activation (above a threshold) are instantiated in the preconscious Workspace, and they point to their corresponding root nodes and links in PAM.

Referential-Meaning Codelets add a new referential link in the Workspace's Current Situational Model (CSM) from an alarm call node's instance to its corresponding predator node's instance. Seyfarth and colleagues (1980) pointed out that juveniles sometime produce alarm calls in a wrong context. They utter eagle calls upon spotting any instance of the avian category (e.g., falling tree leaves, birds etc.), utter leopard calls upon detecting a terrestrial animal (e.g., Cheetah), and utter snake calls upon noticing any serpentine object.

This is direct evidence that vervet infants understand that an eagle call, snake call, and leopard call refers respectively to an avian category, serpentine like category, and terrestrial category. In other words, the vervet's brain avoids any association between a specific call and an instance outside of the corresponding category. For example, the infant's brain may associate in an early stage of learning an eagle call with a crow, but not with a lion because it doesn't belong to the avian category. Similar logic is applied for snake and leopard calls. Based on these experimental observations, referential

meaning codelets (responsible for adding relationships from alarm calls to their corresponding predators' classes) are inborn in the vervet's mind. Through experience, the infants learn to refine the external referent category to be more specific. In fact, their brains reinforce the correct associations from the alarm call to the corresponding predator, and the inappropriate associations decay. As a result, the eagle calls are associated with eagles only, leopard calls are associated with leopards, and snakes calls with snakes.

The following is a description of the functionality of each meaning codelet's category.

Eagle Call Referential-Meaning Codelet

If an eagle call node and an object of avian category exist in the CSM, the Eagle-Call Referential Meaning Codelet adds a referential link from the eagle call node to the avian instance node with a total activation of 1.0, which helps to bring the learned link to consciousness.

Leopard Call Referential-Meaning Codelet

If a leopard call instance and an object of terrestrial category exist in the CSM, the Leopard Call Referential Meaning Codelet adds a referential link from the leopard call node to the terrestrial instance node with a total activation 1.0

Snake Call Referential-Meaning Codelet

If a snake call instance and an object of snake like category exist in the CSM, the Snake Call Referential Meaning Codelet adds a referential link from the snake call node to the serpentine-like instance node with a total activation 1.0.

Action Meaning Codelet

An Action-Meaning Codelet adds a causality link in the CSM from an alarm call node to its corresponding escape action node. The context of this codelet is an alarm call and an escape action. This is a generic codelet, which acts if its context is matched in the CSM.

Fear Meaning Codelet

In the LIDA model, emotions are considered as feelings with cognitive content, such as being angry at a specific person, the shame at saying an inappropriate thing, etc. Franklin and Graesser (1997) state that every autonomous agent must be equipped with primitive motivations that motivate its selection of actions, in order to form its own agenda. Such motivations may sometime be causal or in the form of productions rules *(if condition)* in an artificial agent. In the LIDA model, these motivations are implemented by feelings (Franklin & Ramamurthy, 2006). Vervet agents use fear as a primary motivation to select the appropriate escape action upon hearing an alarm call.

We consider a Fear-Meaning Codelet, whose task is adding a causality link in the CSM from an alarm call node to the fear node. In another work context, a generic Emotion-Meaning Codelet or a Causality-Meaning Codelet can be employed. The context of the Fear-Meaning Codelet is presence of the alarm call node and the node representing self-fear. This may result in associating, in an early stage, the fear with non-threatening objects (e.g. , tree, bush) perceived simultaneously with an alarm call. Thus, the infant's perception of the non-threatening object, in a further time, may trigger its fear feeling. As the infant grows older, the meaningful relationships are reinforced and the insignificant ones decay. All the newly created referential and causality links in the CSM

are not learned into PAM, unless they succeed in being brought to consciousness by attentional codelets. These referential and causality links are broadcast and added to the PAM node structure with a specific value of base level activation. If an existing link is broadcast in a later cognitive cycle, its base level activation is reinforced. The learning of the meaning of each alarm call may take several cognitive cycles to be accomplished.

The implementation of the base level activation of the learned link can be done using a sigmoid function, which defines the behavior of the base level activation of the newly learned links. The sigmoid function is defined as follows:

$$f(x) = \frac{1}{1 + \frac{1}{\exp(a \times x + c)}}$$

Where:

- $f(x)$: the new base level activation of the link or node in PAM.

- x : the current base level activation of the link or node in PAM.

- *a and c* : are real numbers for linear parameterization. Their default values are 1.0 and 0.0, respectively.

Another important concept that affects the base level activation's behavior of each PAM's element is the decay concept. All elements in PAM decay over time. The decay rate follows an inverse sigmoidal of the current value of the base level activation. The higher the base level activation of an item, the slower its decay rate (Scott, 2006).

3   LIDA Framework & Simulation Design and implementation

LIDA Framework

The LIDA framework (Snaider, McCall, & Franklin, 2011) is a generic computational implementation of the various modules and components of the LIDA cognitive model, using the Java programming language. The framework is easily customizable for specific domains and problems. This customization can be done through the LIDA configuration file (an XML file) which allows the developer, at a low level, to configure several parameters such as decaying strategies and base level activations. Another feature of the framework is the specification of the XML file; the developer does not need to implement the entire agent in Java; he can just define much of it using this file.

A masterpiece of the LIDA framework is the task manager, which schedules and executes all the tasks of the application such as recognition tasks, attentional codelets tasks, structure building codelet tasks, etc. The task manager organizes all the tasks in a task queue to schedule the LIDA tasks for execution. Each position in the task queue represents a discrete instant in simulation time, which we call a tick. A tick is considered as a time unit, and its duration can be configured by the developer in milliseconds. This mechanism allows the simulation experiments to be run in various modes: -slow mode-step-by-step mode – different speeds.

The framework is implemented with an object oriented approach. Thus, while implementing the LIDA agent, which represents the vervet infant, we call the generic classes of each module needed, and we override functions to implement specific tasks.

Next, we explain the design and implementation of the LIDA agent modules and how these modules are related to the LIDA framework.

ALife Environment Design and Implementation

Artificial life attempts to understand the essential general properties of living systems by synthesizing life-like behavior in software, hardware, and biochemistry. The use of these approaches demonstrates the capability of explaining various aspects of language including the evolution of signaling systems, the grounding of symbols, and the evolution of meanings (Bedau, 2003; Kirby, 2002). Taking advantage of this approach, we designed and implemented a two-dimensional ALife environment to test learning the meanings of vervet alarm calls. The environment consists of a grid of cells, populated with a LIDA-based autonomous agent labeled INFANT, and other agents (mother agent (MAMA), vervet agents (VERVET), and predators) which are controlled by simple rules in the form of "*If condition–Then action*". The agents' control is consistent with Nagal's assumptions (1974). In his words: "*Learning "what it is like" to be an animal of a certain sort means learning how that animal goes about deciding where to go next and what to do next*".

The agents make continuous navigational decisions to minimize the risk of being attacked by predators and maintain a state of good health. To attain these survival goals, the agents perform various actions, such as escaping into locations safe from predators (e.g., climbing to the top of trees, or hiding in a bush), vocalizing various alarm calls, and foraging for food. The INFANT depends on the MAMA agent during the first simulation part, where we assumed a physical attachment between them. A band of vervets, that at a given time, occupies only a small region in the wild, was simulated. Hence, the agents' (MAMA, VERVET, INFANT) vision system consists of a line of sight; the agent can see

the objects located in every cell in the line. There is an exception for the predators' vision system: an eagle can't see a vervet agent hidden in a bush. Conversely, the agent's hearing system is extended to every cell in the environment. In other words, the agent in the ALife environment is able to perceive the sound regardless of its location. In fact, the sound spreads quickly in the small region occupied at a given time by a small group of vervets.

The ALife environment is generic and flexible. The grid size, the agent's vision, and hearing systems can be adjusted depending on the nature of the experiments. Also, additional features can be added as needed. Therefore, it is an effective research and computational tool to test various theories.

What follows is a description of the ALife environment, its structure, the objects that populate it, including their properties and methods. We describe also the available actions for vervets during both stages of the simulation.

Environment

The environment is a square grid which consists of cells. The size of the environment can be adjusted to fit the needs of any experiment. In the wild, a size of vervet band ranges from ten to thirty monkeys. If we consider a Poisson distribution to describe the population distribution of such troops, the median size is fifteen vervets. Hence, the grid environment was composed of 20x20 cells, which is consistent with a region occupied by this number of vervets during the daytime. Figure 5 represents the two-dimensional ALife grid environment.

Figure 5. Two-dimension ALife grid environment

The main Java class that implements the environment is called "ALifeWorld". This class has the main properties:

1. Cells: an array of cells

2. Width

3. Height

4. Objects: objects that populate the environment

5. Actions: Actions available for the agents

Cell Object: Cell is a class that defines a specific location in the environment. It has the following attributes:

1. Occupancy : Integer variable whose value is the sum of the objects' sizes in a cell. Each object has a size as an attribute and configured in the XML specification file of the agent.

2. Capacity : Integer variable that determines the maximum number of objects that can be populated in a cell based on their sizes.

3. Objects : Set of objects that occupy the cell and whose summed sizes is less than the cell's capacity.

4. Sound : String variable that has three possible values:

"eagle call"

"snake call"

"leopard call"

5. X-coordinate : Integer variable that can vary in [0, width]

6. Y-coordinate : Integer variable that can vary in [0, height]

ALifeObject: A Java class defines each element comprising the cell. There are two types: 1) animated objects that move including the mother agent, vervet agent, juvenile agent, and predator agents; and 2) non- animated objects such as trees and bushes.

1. id

2. name

3. size

4. isHidden

5. ClimbedTree

6. IsTrembling

7. Direction

8. hasStood

9. health

"isHidden" is a Boolean variable set to true when the agent is hidden under a bush.

"*ClimedTree"* is a Boolean variable set to true when the agent performs the action "climb

tree". "*IsTrembling"* is a Boolean variable set to true when the agent is feeling fear.

"*hasStood"* is also a Boolean variable set to true when the agent stands bipedally.

Direction defines where the agent is heading. It has four possible values:

1. North

2. South

3. East

4. West

"Health" is implemented as a volatile double variable that can be modified by multiple

threads. Its range is in the interval [0.0, 1.0]. A detailed explanation will be provided

later in the agent's health recognition task. These attributes are all applicable for the

animated agents only.

UpdateState function

This function is called every time point (tick) in the simulation to update the environment. Every agent can perform its own actions through overriding this function.

World Operations

In the LIDA model, there are several types of actions: internal actions (e.g., imagining an action), external actions that involve a muscle movement (e.g., run climb a tree, hide under a bush, etc.), and implicit actions such as see and hear.

In this simulation, the agents in the ALife environment have interactive abilities. They perform various actions: move, attack, climb to the top of trees, hide under bush, and vocalize diverse alarm calls. The escape actions are specific to preys, while attacking is performed by predators only. During the physical attachment between the mother and the infant, the effect of the actions executed by the MAMA includes the INFANT. All the actions override a main function labeled *"performOperation"* which takes as the following parameters: *"ALifeWorld world"* (the current environment), *"ALifeObject" subject* (agent that preforms the action), *and "ALifeObject [] objects"* (the set of agents on which a selected action is performed.

The following is a brief description of the main actions available for agents:

Move

Based on its current direction, the agent moves to the next cell. For example, if the agent occupies the Cell (x, y) and its current direction is north, taking Move action results in placing the agent in the Cell (x, y-1). This results in decreasing the agent's health by a small amount. Health decrease is viewed as an energetic cost of the agent's movement.

MamaMove

Based on its current direction, the *"MamaMove"* action changes the current cell of the MAMA and the INFANT simultaneously during their physical attachment.

VocalizeAlarmCall

1. VocalizeEagleCall

2. VocalizeSnakeCall

3. VocalizeLeopardCall

Vocalizing an alarm call by a VERVET or MAMA, results in changing the internal variable sound of the current cell. At a specific tick of the simulation, vocalizing eagle call, snake call, and leopard call changes the sound value to "eagle call", "snake call" and "leopard call", respectively. At the following tick, the sound is spread to the surrounding cells of the grid environment. Therefore, every sound attribute of each cell is set to the specific alarm call.

HideUnderBush

VERVET and MAMA agents hide in the bush upon hearing an eagle call. If an agent performs this action, its internal Boolean variable, *"isHidden",* is set to true when arriving to a cell that contains a bush object. Additionally, the agent stops feeling fear caused by hearing the eagle call. This is translated by setting the variable, *"IsTrembling",* to false. If the agent moves from the cell that contains a bush object, its Boolean internal variable (*"isHidden"*) is turned to false.

ClimbTree

VERVET and MAMA agents climb a tree upon hearing a leopard call. When a VERVET or MAMA agent performs this action, its internal Boolean variable, *"ClimbedTree",* is set to true when arriving to a cell that contains a tree. Moreover, the

vervet agent stops feeling fear caused by hearing the leopard call. This results in setting the variable, "is*Trembling",* to false. When the agent climbs down from the tree (moves from the cell that has a tree object), its internal Boolean variable, *"ClimbedTree",* is set to false.

Stand Bipedally

A VERVET or MAMA agent stands bibedally after hearing a snake call. Performing this action results in turning the internal Boolean variable, *"hasStood",* to true. If there is no snake call event in the environment, the *"hasStood"* variable is updated to false.

Attack

1. EagleAttack

2. SnakeAttack

3. LeopardAttack

Attack actions have three parameters: the environment, the predator which performs the action, and the prey agent that is attacked by the predator. Field experiments show that successful predator attacks, which occur rarely in real wild life, most often result in the immediate death of vervets. Failed attacks produce serious injuries. In this work, we simulate the common case, so predator attacks result in decreasing the health amount of the attacked agent instead of its death.

There are three types of attacks: First, an eagle performs an *"EagleAttack"* action if it sees a VERVET, MAMA, or INFANT agent unless they are hidden in the bush. Second, a leopard performs a *"LeopardAttack"* action if it sees a VERVET, MAMA, or INFANT agent unless they are on the top of a tree. This is because the leopard can't

ascend to the small branches at the top of a tree. Lastly, a snake agent performs a

"*SnakeAttack*" action if it sees a VERVET, MAMA, or INFANT agent unless they flee.

Turn

1. TurnRight

2. TurnLeft

3. TurnAround

Performing turning actions results in changing the agent's direction.

SeeObjectsInCell

As mentioned previously, "see" is considered as an implicit action in the LIDA

model. A vervet agent detects all *animated* and *non-animated* objects in its current cell.

There is another see action that allows the vervet agent to detect objects in all cells along

its line of sight.

EatOperation

A vervet agent grabs a food in its current cell. The execution of this action results

in increasing the health of the agent

LIDA Agent Design and Implementation

We consider two experimental categories: 1) learning the meaning of the vervet

alarm calls while the INFANT's movement is confined to a few actions. This is because

it is attached physically to the mother agent; and 2) Assessing the INFANT's

understanding of the meaning of these alarm calls where the primary assumption is the

detachment of the INFANT from the MAMA agent. Therefore, we expand its procedural

memory to contain new schemes and actions (e.g., climb tree, hide under bush, and stand

bipedally). Next, we present the design and the implementation of the INFANT's modules.

Sensory Memory

This module is implemented as a class in the LIDA framework. In our implementation, we create a new Java class proper to the INFANT that extends the Java generic class *SensoryMemoryImpl* in the framework. This class inherits all the functions of the generic class and we implement the following additional sensors:

- *Sound sensors* allow the agent to detect sound in the environment, more specifically the alarm calls produced by other vervets. The INFANT is able to sense the sound regardless of its location in the environment because the sound is scattered in all cells of the grid environment.

- *Infant-Mother sensors* allow the INFANT to sense emotions and feelings from the mother agent, such as fear, especially for the first simulation stage where the main assumption is physical attachment.

- *OriginCellObjects sensors* allow the INFANT to recognize all the objects in its cell. A detailed description of the environment, including the cell, will be described later.

- *NextCellObjects sensors* allow the INFANT to recognize all the objects in every cell in its line of sight.

- *Health sensors* allow the INFANT to sense its health. The health system's agent is implemented as a double variable in the [0.0, 1.0] interval.

Perceptual Associative Memory (PAM)

PAM Design. Perceptual Associative Memory (PAM) is implemented as a modified slipnet (Hofstadter & Mitchel, 1994). It allows the agent to distinguish,

- 29 -

and identify external and internal information. Oliphant has defined animal

communication as follows (Oliphant, 1997):

An act of communication is a causal chain of events, whereby one individual, the sender,

exhibits a behavior in response to a particular situation, and a second individual, the

receiver, responds to this behavior. Such an interaction is communicative if it involves

manipulation on the part of the sender and exploitation on the part of the receiver (p.).

Following this definition, the vervet infants acquire the meaning of distinct alarm

calls from observing the following events: 1) Detection of the predator in the

environment; 2) Hearing alarm calls; and 3) Escape actions into safe locations. The event

is considered as the primary representation in the LIDA agent PAM design. In the LIDA

model, event-based representations draw inspiration from research on thematic roles

(McCall, Franklin, & Friedlander, 2010). Events are represented as nodes with thematic

role links binding to Agent, Object, Location, Feelings and other node types. This

representation is consistent with Carlson's definition of thematic roles in events

representations (Carlson, 1998). In his words: *"The basic idea that there is a smallish,*

finite number of distinct roles with names like "Agent," "Instrument" , "Goal",

"Patient", "Location", and so forth that have direct semantic import..."

We assume that the INFANT has already learned to recognize the events involved

in this simulation (detection of predators, hearing alarm calls and escape actions). This

can be realized in several cognitive cycles.

As mentioned previously, an event is represented in PAM as a node that has

multiple thematic role links that lead to it from multiple nodes which play roles in the

event. An event node is activated in PAM based on the amount of activation received

from its children nodes and passed through the thematic role links. For this purpose, we consider a new implementation of the propagation task in PAM. This task serves to excite the link's sink (in this case sink is the event node) based on the link's new activation. If this puts the sink over its percept threshold, then both link and sink will be sent as a percept.

The mathematic equation of the excitation is as follows:

$$\text{Excitation of Sink} = \text{excitation amount} * \text{Base level activation of link}$$

The excitation of the sink is the amount of activation passed to the event node by each thematic role link which has a specific excitation amount. The base level activation of each link reflects the weight of each role in the event.

In the initialization of PAM parameters, we set the base level activation of each thematic role link associated with an event, based on the significance of each thematic role in the event. As mentioned previously, there are three events types:

1- *Detection of a predator*. We generate three events, of this type, in PAM:

   1) I see an eagle; 2) I see a leopard; 3) I see a snake.

2- *Hearing alarm calls :* We generate three events, of this type, in PAM:

   1) I hear an eagle call; 2) I hear a leopard call; 3) I hear a snake call

3- *Escape actions.* We generate three events, of this type, in PAM:

   1) Mother agent hides under bush;

   2) Mother agent climbs to the top of the tree.

   3) Mother agent scans the area.

In the first stage of the simulation with the physical attachment assumption, it's more likely that the infant recognizes the actions performed by its mother, due to its body position. In the second stage of the simulation with the de-attachment assumption, the INFANT can recognize additional events where the corresponding actions are performed by the other vervets. For example, a VERVET climbs to the top of the tree event as well as MAMA climbs the top of tree event. Now we break down the agent PAM to the following events:

Detection Predator Events

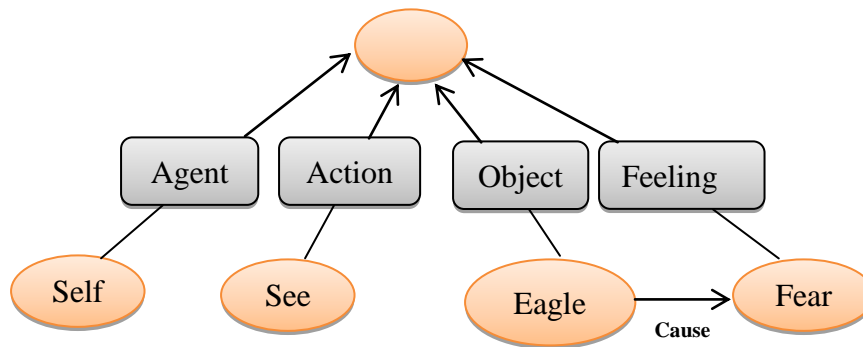Figures 6, 7, 8 describe the events of seeing various predators.


Figure 6. Event Representation in PAM "I see an eagle"

Figure 7. Event "I see Leopard**"**
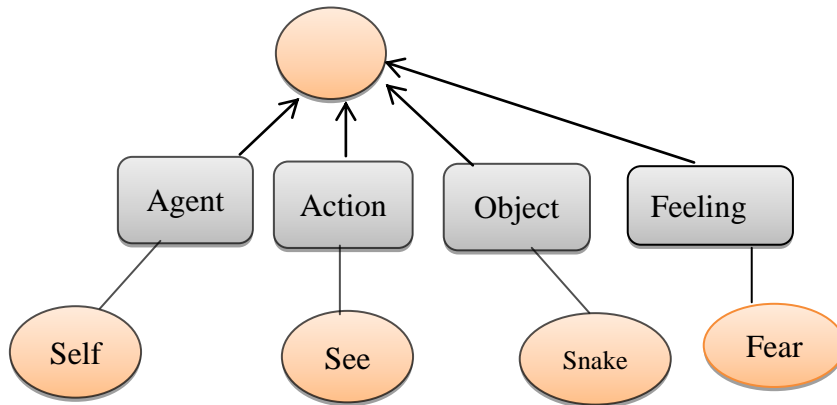
Figure 8. Event "I see snake"

The three events of seeing an eagle, leopard, or snake, share these three thematic roles:

1. The agent thematic role is the agent itself. It is represented by the self-node. The LIDA model supports a self-system composed of three components: the ProtoSelf, the Minimal (Core) Self and the Extended Self (Ramamurthy & Franklin, 2011; Gallagher, 2004). The self-node in this event belongs to the self-as-experiencer (the experiencing self). The LIDA agent uses an object feature detector to detect any object in its current cell (animated or non-animated). The self-node is considered as any other object; hence it is always activated in the Current Situational Model.

   Gallup and colleagues (Gallup, Jr., Anderson, & Shillito, 2002) conducted mirror test experiments to answer the question "can animals recognize themselves in a mirror?" The results of the mirror test show the inability of some monkeys to recognize themselves in a mirror but some monkeys have   this capability (Waal,

M., Freeman, & Hall, 2005; Rajalael, Reininger, Lancaster, & Populin, 2010). These findings don't contradict our assumption that the vervet infant recognizes itself in reality, in spite of its inability to recognize itself in front of a mirror.

2. The action thematic role is attached to the see node which is considered an implicit action in the LIDA model. This node is also activated continuously in the LIDA agent CSM.

3. The object thematic role is attached to the predator node: eagle, snake, and leopard nodes.

The main dissimilarity among the three events is the feeling thematic role, which is attached to seeing an eagle event and seeing a snake. Several researchers in psychology performed various studies and experiments to answer the question "is the fear of specific predators innate or does it involve learning?" Seyfarth and colleagues (1980) indicate that the infant vervets emit fewer snake alarm calls to snakes than vervet adults, and make more snake alarm calls to inappropriate objects. This was implicit evidence that the infant's brain is prepared to learn the fear of snakes very quickly. We assume then that the fear feeling is part of seeing the snake event. On the other hand, Worden (1996) claimed that the vervets are born with an innate fear of birds. According to him; infant vervets innately produce eagle calls in the presence of birds. By observing their adult peers' reactions, such as facial experiences or body reaction, they reinforce fear for only dangerous birds such as eagles. This justifies the fear feeling thematic role in "I see eagle" event. Lastly, for the fear of leopards, we assume that the LIDA agent acquires this fear from sensing it from the mother agent. This is attained computationally through the mother-infant sensors.

Now, we describe the representation of the events perceived by the LIDA agent during its physical attachment with the mother.

Hearing Alarm Calls Events

Figures 9, 10 and 11 describe the events of hearing various alarm calls.

Figure 9. Event representation in PAM "I hear eagle call"
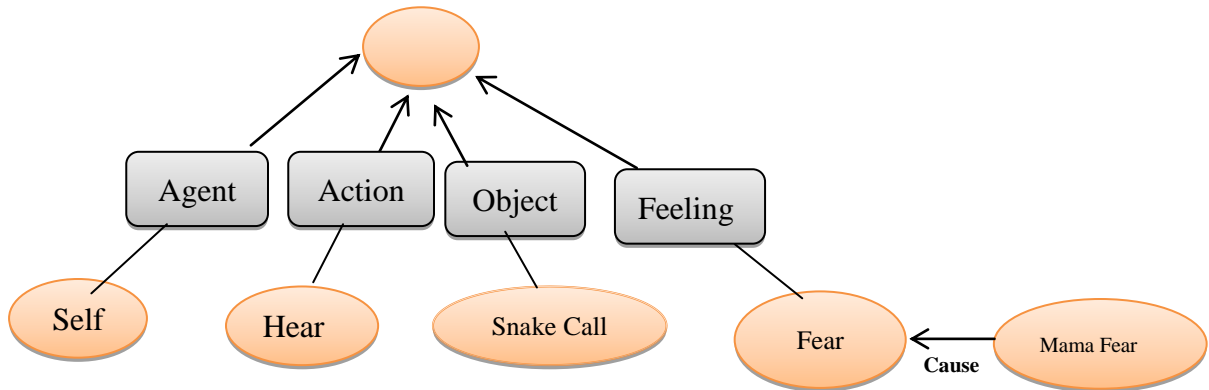

Figure 10. Event representation in PAM "I hear snake call"
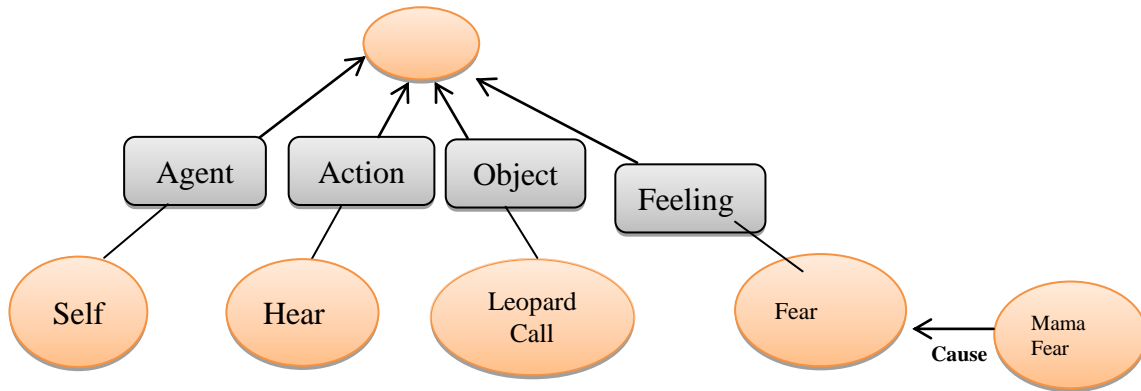
Figure 11. *Event* representation in PAM "I hear leopard call"

The three events of hearing the three alarm calls share these thematic roles:

1. The agent thematic role is the agent itself. It is represented by the self-node as explained previously.

2. The action thematic role is hearing. This action node is activated in CSM upon hearing any alarm or sound in general call.

3. The object thematic role link leads from each alarm call node, which is an acoustic node.

4. The feeling thematic role link leads from the fear feeling node. The physical attachment of the vervet infant to its mother in the first stage, allows it to sense the mother's fear directly after hearing each alarm call by the means of the infant-mother sensors. Many psychological studies have shown that the emotional bond between the infant (human or animal) and its mother (or caregiver) contributes to the infant's experience of diverse feelings and emotions including fear. This justifies the innate causality link in PAM, from the mother's fear to the LIDA agent's fear. (Harlow & Harlow, 1969)

Escape Actions Events

Figures 12, 13, 14 and 15 describe the events of escape actions executed by the MAMA.
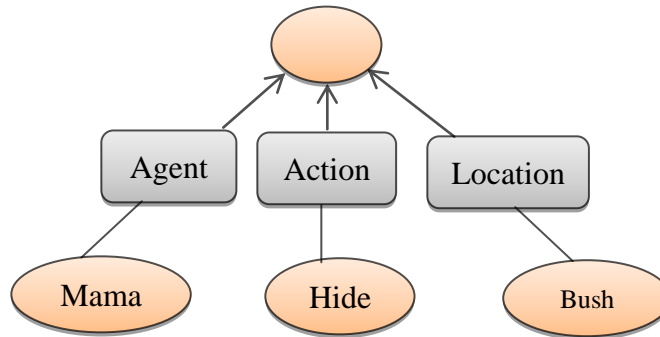
Figure 12. Event representation
"Mama hides under a bush"

Figure 13. Event representation
"MAMA climbs top of tree"
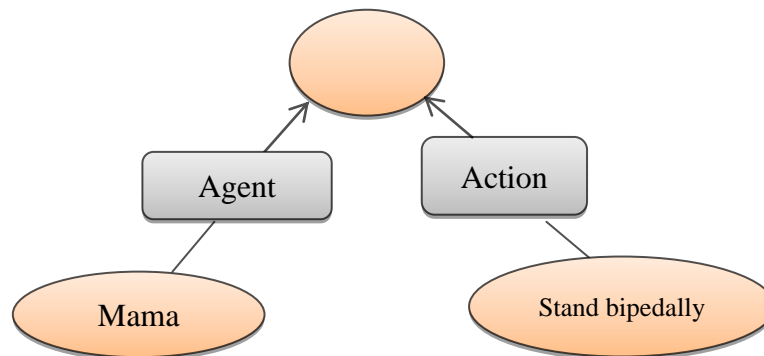
Figure 14. Event representation "Mama stands bipedally"

Figure 15. Event representations "Mama scans the area"

Agent thematic link leads from the mother node. Because the INFANT is attached physically to the MAMA agent in the first stage, it is more likely to recognize an escape action performed by the mother agent than by any other VERVET agent.

1. Action thematic role. Each alarm call elicits different escape actions. The action thematic role link leads from: 1) hide action node in the event "MAMA hides under the bush "; 2) climb action node in the event " MAMA climbs tree"; and 3) stands bipedally action node in the event "MAMA stands bipedally".

2. Location thematic role, whose link is attached to: 1) the bush node in "MAMA hides under bush event" and 2) tree node in "MAMA climbs tree" event.

Recognition Tasks

Features detectors in LIDA represent the main mechanism for executing recognition tasks. They descend on the incoming sensation in sensory memory. Those that find features (bits of meaning, single chunks) relevant to their specialty activate appropriate nodes in Perceptual Associative Memory (Franklin, Baars, Ramamurthy & Ventura, 2005). Four categories are used: 1) Object Features Detectors; 2) Mother Fear

Feature Detectors; 3) Alarm Call Feature Detectors; 4) Health Detector; 5) Action

Feature Detectors (hide under bushes, climb trees, and stand bipedally). Listed below is a

further description of the functionality of each category.

Object Feature Detector

The function of the Object Feature Detector is recognizing objects visually in every

cell of the line of the sight of the LIDA agent. There are two types of visual objects: 1)

animated such as mother agent, vervet agent, juvenile agent etc.; and 2) *non-animated*

such as trees and bushes. All objects are detected by using the same object feature

detector algorithm. In addition, we configure the object associated with its feature

detector in the specification XML file which specifies much of the architecture of the

software agent. A supplementary function of this detector is allowing the INFANT to

recognize itself by adding a self-node in its Perceptual Associative Memory (PAM). The

self-node is considered an *animated* object just as vervets. Its feature detector is

configured in the primary XML file by adding the self- node in the object feature detector

configuration as an object.

Mother Fear Feature Detector

The main assumption of the first stage of the simulation is the physical attachment

of the LIDA agent to the mother agent. Therefore, the INFANT learns the fear of entities

and events (e.g., snake, leopard, alarm calls) by sensing the mother's fear which is

computationally implemented as a Boolean variable labeled "*isTrembling*". Every

VERVET quivers when seeing a predator or hearing an alarm call. This is implemented

computationally through Tremble action which consists of setting the variable value of

"*isTrembling*" to true. When the VERVET escapes to a location safe from predators after

hearing an alarm call or detecting a predator, the *"isTrembling"* variable is then set to false. Ultimately the variable, *"isTrembling",* is primarily used to enable the INFANT to recognize the mother's fear.

Alarm Call Feature Detector

The recognition of vervet alarm calls is one of the most important tasks for the LIDA agent. The sound is computationally implemented as a string variable, *sound,* associated with each cell of the two-dimension grid environment. The *sound* value changes when VERVET agents perform distinct vocalization actions associated with a particular predator. The variable *sound* of each cell is updated to the following values: 1) *"eaglecall"* when *VocalizeEagleCall* action is performed; 2) "leopardcall" when *"VocalizeLeopardCall"* action is performed and 3) *"snakecall"* when *"VocalizeSnakeCall"* action is performed. A VERVET vocalizes an alarm call if no one else did it while spotting a predator in the vicinity.

Health Feature Detector

The agent's health is an internal real variable. The INFANT loses an amount of its health if it experiences a dangerous event such as being attacked by a predator. In the wild, it is rare that vervets are killed when being attacked by predators. It is vital that the agent maintains a good health during the simulation's iterations. Thus, we boost the INFANT's health by the nursing action or eating food action during both simulation stages of physical attachment and detachment. This feature detector activates three nodes in PAM, depending on the health value. If the health value is greater than 0.66, the good Health node is activated in PAM. If the value is greater than 0.33, fair Health node is

activated. Finally, if the health value drops under 0.33 the bad health node is activated.

During all experiments, we tried to maintain a fair health for the INFANT.

Action Feature Detector

Action feature detectors allow the LIDA agent to recognize the actions performed by the other agents. The observer is not merely contemplating the action of the other agent, it is attempting to understand or predict the outcome of the action it observes. Actions of other agents convey valuable information for learning skills or engaging in communication. Perceiving the escape actions performed by adult vervets plays a role in learning such actions. The representation of the observed escape actions in the LIDA agent's PAM allows the infant to learn the action meaning of various alarm calls by building causal relationships from such calls to their corresponding escape actions. We implement computationally the recognition of each escape action using the following *Boolean* variables: "*isHidden*" for hiding under bush, "*hasClimbed*" for climbing to the top of trees, and *"hasStood"* for standing bipedally. Now we describe the Structure Building Codelets module of the LIDA agent.

Structure Building Codelets

A LIDA Structure Building Codelet (SBC) is a small process (or daemon) that performs specific tasks in the Workspace, such as modifying existing structures in the CSM, or adding new structures (e.g., nodes, link etc.). A SBC operates asynchronously and independently of other processes in LIDA. Each SBC is triggered when a specific representation is present in the Workspace. As a data structure, the SBC has a base-level activation, a context, and an algorithm. As explained previously, the base-level activation measures the usefulness of the codelet, and is modified by selectionist learning. The

context is the node structure or pattern that SBC is "looking for" in the Workspace. The action or algorithm specifies what the codelet does when activated.

We implemented three Structure Building Codelets' categories: 1) Referential SBC; 2) Action-Meaning SBC; and 3) Fear-Meaning SBC. For this purpose, we create a new Java class that inherits the *StructureBuildingCodeletImpl* class from the LIDA framework. Next, we override the *runThisFrameworkTask()* function, by implementing the job of each SBC. Lastly, we configured each SBC in the primary XML file of the LIDA agent and in the Factory XML file. Table 1 summarizes the Structure Building Codelets (SBC) used in this work:

Table 1

Structure Building Codelets and their descriptions

SBC Name	SBC task description (in CSM)
EagleCallReferential Codelet	- Add a referential link from eagle call node to avian nodes
LeopardCallReferential Codelet	- Add a referential link from leopard call node to terrestrial nodes
SnakeCallReferential Codelet	- Add a referential link from snake call node to serpentine like nodes
Fear-Meaning Codelet	- Add a causality link from fear node to all nodes that present non-threatening objects and that are instantiated in the CSM.
Action-Meaning Codelet	- Add a causality link from alarm call node to action escape node that is instantiated in the CSM.

Three different Referential SBCs were implemented for several reasons. The EagleCallReferential Codelet and SnakeCallRefrential Codelet are hardwired in the

infant's vervet mind. In fact, field experiments revealed the tendency of vervet infants and juveniles to produce eagle calls and snake calls when seeing, respectively, an avian instance and serpentine like instance. However, the LeopardCallReferential Codelet is not hardwired in the vervet's mind.

Procedural Memory

LIDA's procedural memory initiates the process of deciding what to do next. It's implemented using a scheme net data structure which is a directed graph whose nodes are called schemes. This is similar to Drescher's schema mechanism but with many fewer parameters (Drescher, 1991). A scheme has a context, an action, a result, and a base-level activation. In the first simulation stage, the primary assumption of the LIDA agent is the physical attachment to the mother agent. Consequently, the LIDA agent performs few actions such as turning left, turning right, and turning around. The LIDA agent selects these actions when hearing an alarm call or sensing the mother's fear. The infant tries to search for more cues to understand these perceived salient events and feelings. Table 2 summarizes the LIDA agent's actions in the first part of the simulation:

Table 2

LIDA Agent's actions during its physical attachment with the mother agent

Action Name	Description
Turn Left	- Instruct the agent to attempt to rotate left
Turn Right	- Instruct the agent to attempt to rotate right
Turn Around	- Instruct the agent to attempt to rotate to the opposite current direction

During the second stage of the simulation, the LIDA agent de-attaches from the mother agent. Thus, we expand the procedural memory to contain additional schemes such as moving, hiding under bush, climbing a tree, and standing. Table 3 summarizes the LIDA agent actions.

Table 3

LIDA Agent's schemes during its physical de-attachment with the mother agent

Action Name		Description
Turn Left	-	Instruct the agent to attempt to rotate left
Turn Right	-	Instruct the agent to attempt to rotate right
Turn Around (Scan)	-	Instruct the agent to attempt to rotate to the opposite current direction
Move	-	Instruct agent to move to the next cell
Eat	-	Instruct agent to grasp the food in the current cell
Hide under bush	-	Instruct agent to move to a cell where there is a bush
Climb tree	-	Instruct agent to move to a cell where there is a tree to climb it
Stand bipedally	-	Instruct agent to stand bipedally

4 Experiments & Results

Learning the meanings of vervet alarm calls was tested by running numerous simulations using a two-dimensional grid-based simulation and a LIDA-based agent labeled INFANT, implemented using the LIDA Framework's modules (Snaider et al., 2011). In each simulation, various objects (*animated* and *non-animated*) were placed randomly. There are fourteen VERVET agents that learned the meanings of alarm calls during their infancy. In addition, the ALife environment (Figure 16) was populated with other *animated* agents such as MAMA agent, eagles, leopards, snakes and *non-animated* objects such as trees, bushes and food. The number of the *animated* and *non-animated* objects was defined in a way to be consistent with a population of a band of vervets that occupies a small region in the wild during a specific time.

Figure 16. Vervet-ALife environment (Left panel of the GUI). The right panels of the GUI describe the cognitive components of the LIDA agent: - PAM Graphs that allow us to visualize the content of PAM in form of a directed graph with nodes and links. This content remains static until adding the new learned links from alarm calls to the corresponding predators, escape actions and fear feeling. The CSM model is dynamic and shows the entities perceived by the LIDA agent at each tick of the simulation. Global Workspace describes the broadcast content, broadcast trigger [2], activation of coalitions [3], broadcast count and broadcast time (in ticks). The bottom panel shows the task queue where multiple tasks are scheduled to run at a point in time (tick) in the simulation. There is also a logger that displays some other information such as the agents 'action.

---

[2] Trigger is a computational technique used to start competition between the   coalitions in the Global Workspace.

[3] A coalition is a *node structure* that has nodes and links.

The *animated* agents are controlled by simple productions rules in the form of "*if condition else action"*. As explained previously, LIDA-based perceptual learning consists of implementing and reinforcing the base level activation of new entities and existing ones respectively. Hence, we recorded the base level activation of the newly learned links.

The simulations were divided into two main stages. First, we carried out experiments to test the learning of the multiple meanings of the vervet alarm calls while the INFANT is attached physically to the MAMA agent. Secondly, we performed a set of experiments to determine whether the INFANT understands the meaning of these alarm calls by evaluating the correctness of its escape actions upon perceiving an alarm call. The main assumption in this second stage is the physical de-attachment of the INFANT from the MAMA. The performance of the INFANT's understanding of the alarm calls was assessed.

Before performing the simulations, we carried out a pilot study in order to make sure the LIDA agent works. Such pilot studies have been done in the past by our CCRG colleagues for diverse research purposes related to the LIDA cognitive architecture. But since we are studying and testing the LIDA based- perceptual learning for the first time, we conducted new preliminary experiments to tune the LIDA parameters, particularly those related to the perceptual learning such as the Structure Building Codelets tick parameter. Once we found the suitable values, we kept them fixed for the rest of the simulations. One main element of the LIDA framework is the task manager, which controls the execution of various LIDA-tasks (Snaider et al., 2011). This task manager

maintains a task queue that schedules the LIDA tasks for execution. Each position in the task queue represents a discrete instant in simulation time, which we call a tick. Ticks are numbered along the simulation, for example tick 1, tick 2, and so on. Each task is scheduled to be executed at a specific tick. So, a single LIDA-task scheduled for a tick is queued in some position t. All tasks scheduled for a particular tick are executed before the task manager advances to the next tick. The order of execution of tasks is not defined a priori. All the tasks at the scheduled tick are executed in a random order. In addition, there is a parameter labeled tick duration which represents milliseconds. In our simulations, we consider one tick duration equal to 10 milliseconds.

What follows is a summary of the tuned values of the LIDA model's internal parameters. Some of them were retrieved from our CCRG colleague's research work (Madl, Baars & Franklin, 2011).

Table 4

Suitable values of the main internal parameters of the LIDA model

Parameter	Value [Ticks]	Value [MS]
1. Sensory Memory ticks	2	20
2. Feature Detectors ticks	3	30
3. Attention Codelets ticks	3	30
4. Structure Building Codelets ticks	3	30
5. Scheme Selection ticks	11	110

Ticks are numbered along the simulation, for example tick 1, tick 2 etc. The above parameters present the frequency of executing a specific task. For example, the feature detectors look at the content of the sensory memory every 3 ticks of the

simulation (at 3 ticks, 6 ticks, 9 ticks etc.). The process of looking at the content of the sensory memory or other modules of LIDA occurs synchronously, and the execution of tasks is asynchronous. The following is a description of each LIDA parameter.

1. The Sensory Memory ticks parameter indicates how often the LIDA agent's sensors operate to sense external data from the environment or its internal data.

2. The Feature Detectors ticks parameter indicates how often the feature detectors look at the content of the sensory memory in order to find specific patterns. If they find the searched pattern, they activate the corresponding entity in the Perceptual Associative Memory (PAM)

3. The Attention Codelets ticks internal parameter indicates how often attention codelets look at the content of the Current Situational Model in order to find relevant portions that match with their concern. Once they find their matched concern, they start to act by forming coalitions and bringing them to the Global Workspace to compete for consciousness.

4. The Structure Building Codelets ticks parameter governs how often the Structure Building Codelets look at the content of the Workspace. Once they find their matched context, they start to operate according to their task.

5. The Scheme Selection ticks parameter governs how often an action is selected from the procedural memory depending on the broadcast conscious content.

Part I: Testing Learning of Meanings of Alarm Calls

Simulations were conducted using each predator type (leopard, eagle, and snake) individually in order to test learning the meanings of the corresponding vervet alarm call.

In all simulations, the ALife grid environment is composed of the INFANT agent, the

MAMA agent, fourteen VERVET agents, trees, and bushes. The INFANT and MAMA

are placed in the same cell in order to comply with the assumption of the physical

attachment between them. This is done via setting a configuration file labeled

"*objects.proprety*" which allows the developers to adjust the attributes of each object

such as its location in the environment, its size, its icon etc. (Snaider, McCall & Franklin,

2011).

In the LIDA model, perceptual learning consists of reinforcing the base level

activation of links and nodes in the Perceptual Associative Memory. Hence, the base

level activations of the newly learned referential and causal links that correspond to the

meanings of vervet alarm calls are recorded. The next table describes the learned links

whose base level activations are recorded in the executed simulations.

Table 5

The learned links that represent the multiple meanings

Source	Sink
Eagle call node	Fear node
Eagle call node	Hide under bush node
Eagle call node	Eagle node
Leopard call node	Fear node
Leopard call node	Climb tree node
Leopard call node	Leopard node
Leopard node	Fear node
Snake call node	Fear node
Snake call node	Stand node
Snake call node	Snake node

The visualization of the occurrence and the progress of learning the meanings of vervet alarm calls is realized by plotting the base level activation of the learned links (leading from an alarm call to the corresponding predator, escape action and the fear feeling) at the time of the broadcast (in ticks) of each link.

As mentioned previously, we adopt a multiple-meanings assessment approach. Each alarm call has three types of meanings: a reference-based meaning, an action-based meaning, and a feeling-based meaning. We studied the temporal order of learning each type of meaning, in order to check whether the INFANT's mind learns, as it is expected to happen in the wild; first the fear meaning, followed by the action meaning, followed by the referential meaning. In fact, the body position of the INFANT (Picture 1) permits him to perceive the mother's fear feeling quickly, followed by the mother's escape actions and finally seeing predators. This order was expected to affect the temporal order of learning the multiple meanings.

Another datum collected from the simulations is the length of time required for learning each type of meaning. It was calculated as the difference between the first broadcast time (in ticks) of a learned link and the broadcast time when the learning is saturated. Although, the sigmoid function approaches 1.0 asymptotically, we assume practically that the learning stops at 0.9999. A comparison of timespan of learning was done between the three types of meanings. The following is a visual representation of the results.

Results and Discussion

The results in figures 17, 18 and 19 show the capacity of the INFANT agent, controlled by the LIDA cognitive architecture, to learn the relationships leading from the

eagle call to the fear feeling, hiding under a bush, and the eagle predator, respectively. Each simulation was performed using the same series of 35 randomly generated environments. The INFANT learned the fear-based meaning at an average point of time equal to 372252.9524 (in ticks). Second, the meaning associated with hiding under bush at an average point of time equal to 781230.9 (in ticks). Lastly, the reference-based meaning related to eagle was learned. As shown in Figure 1, the base level activation of each learned link is reinforced at each broadcast using a sigmoid function.
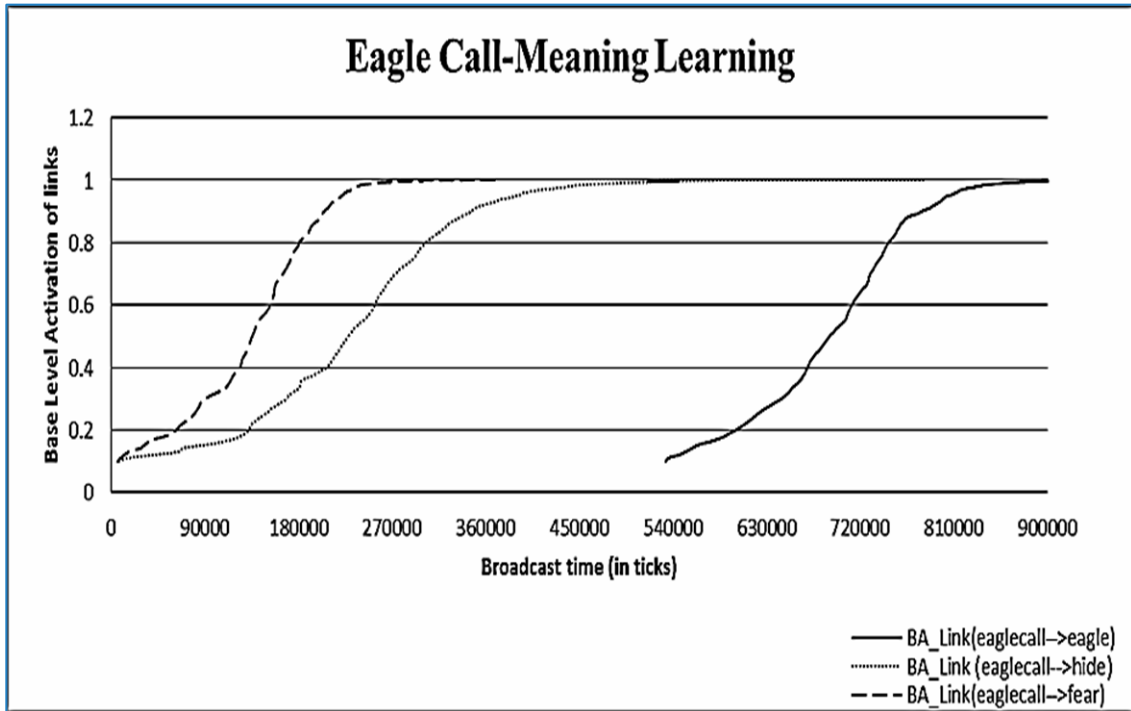

Figure 17. Base-level activations of learned links from eagle call node to the eagle node, the hide under the bush node, and the fear node at each broadcast time.

Figure18 shows the behavior of the sigmoid function used in learning the three links leading from eagle call node to eagle node, fear node, and hide under bush node. The x axis represents the number of times that a learned link comes to the consciousness

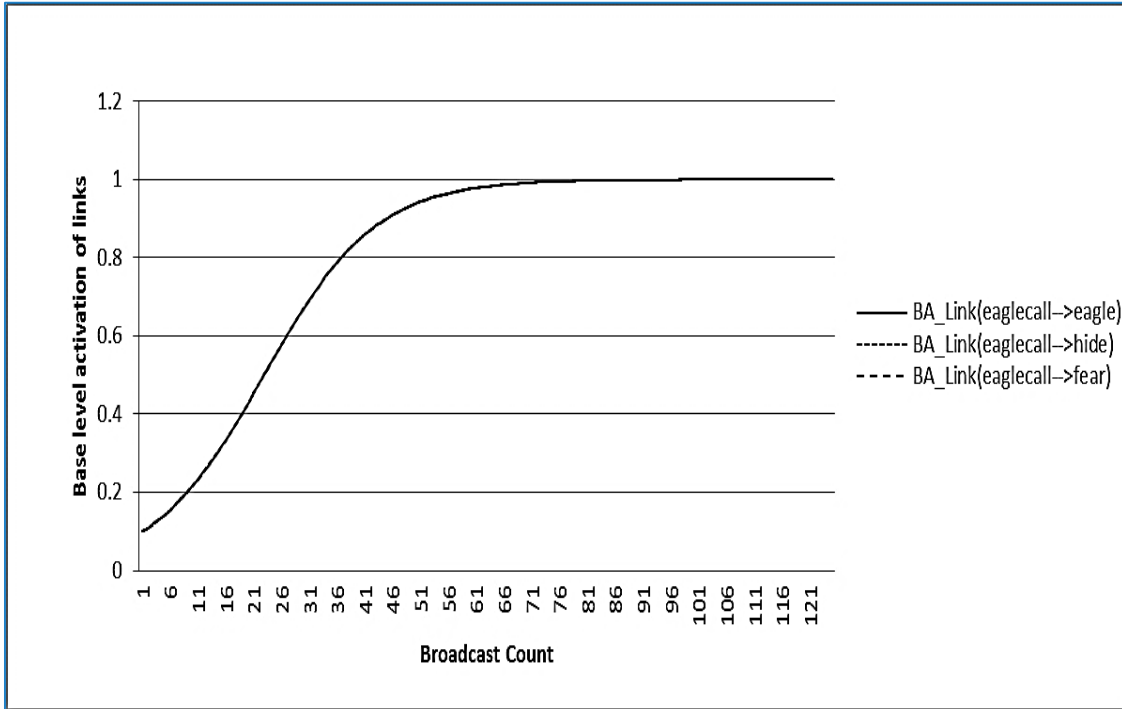and the y axis represents the progress of the base level activation of each learned link at each broadcast.

Figure 18. Base level activation behavior of learned links from eagle call node to eagle node hides under bush node, and fear node along the occurrences of these links in the consciousness.

Figures 18, 21 and 24 are similar because we used a sigmoid function to excite the three learned links with the same amount of 0.1. The obtained curves are in line with the expected "S" shape of a sigmoid function and they reveal a progression from small starts that accelerate and approach a maximum over time.
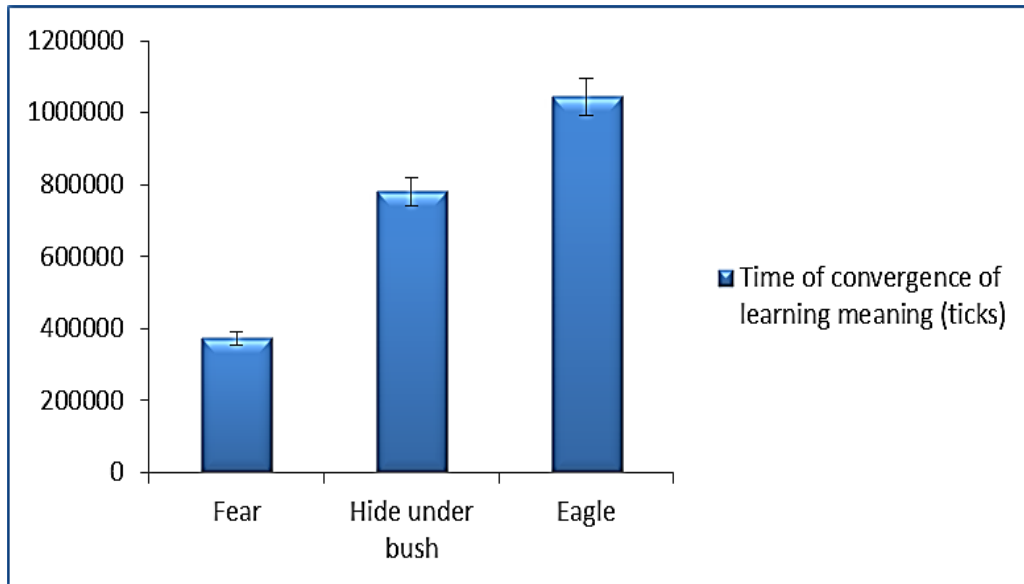
Figure 19. Time of convergence of learning each meaning of eagle call. This time corresponds to a point of time in the simulation when the base level activation of a learned link approaches the maximum. This reflects the temporal order of complete learning of various meanings of eagle calls.

Figures 20, 21 and 22 describe the results of learning the meaning of snake call.

Figure 20.   Base level activations of learned links from snake call node to snake node, stand node, and fear node consecutively, at each broadcast time.
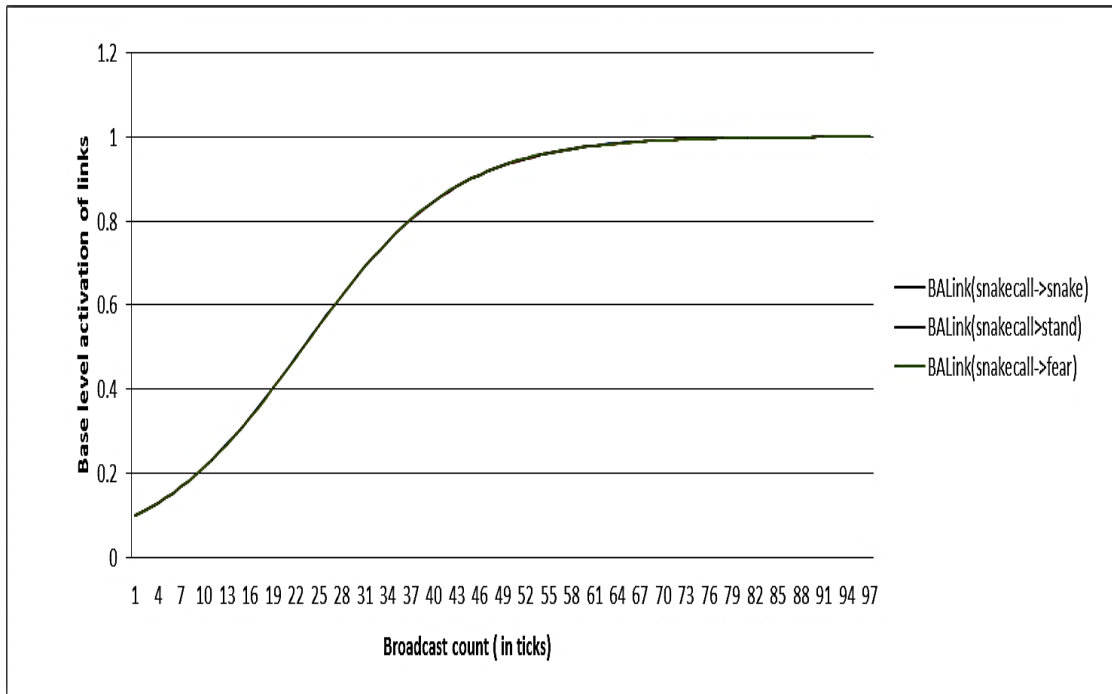
Figure 21. Base level activations of learned links from snake call node to snake node, stand node, and fear node along the occurrences of these links in the consciousness.

The results show that the INFANT learned the relationships leading from the snake call to the fear feeling, standing bipedally, and the snake respectively. Each simulation was performed using the same series of 35 randomly generated environments. The INFANT learned in sequence, the fear-based meaning at an average point of time equal to 149547.619 (in ticks), the action-based meaning associated with standing bipedally at an average point of time equal to 183071.5 (in ticks) and lastly, the reference-based meaning related to the snake predator.
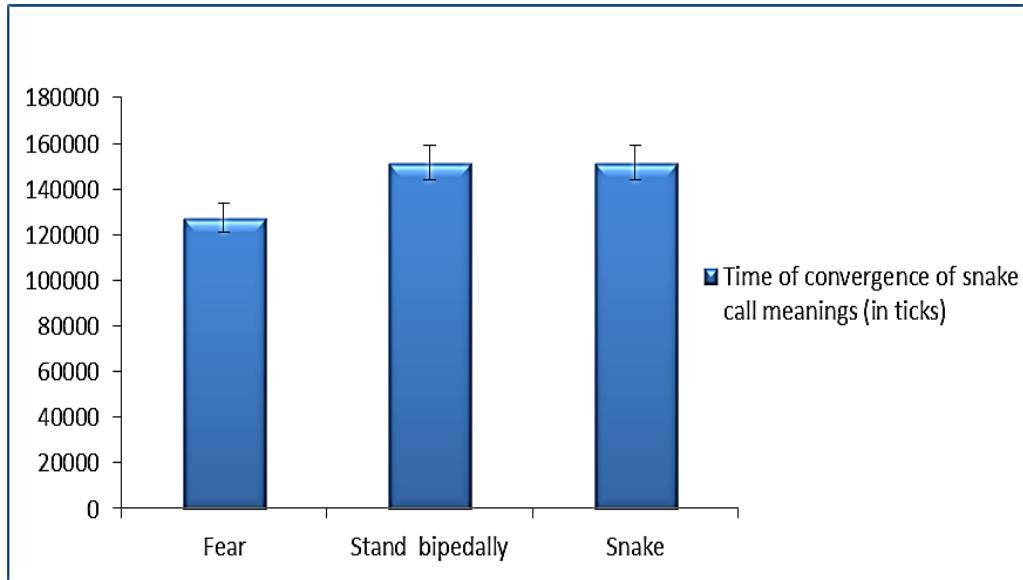
Figure22. Time of convergence of learning each meaning of a snake call.

Figures 23, 24 and 25 describe the results of learning the meaning of leopard call.
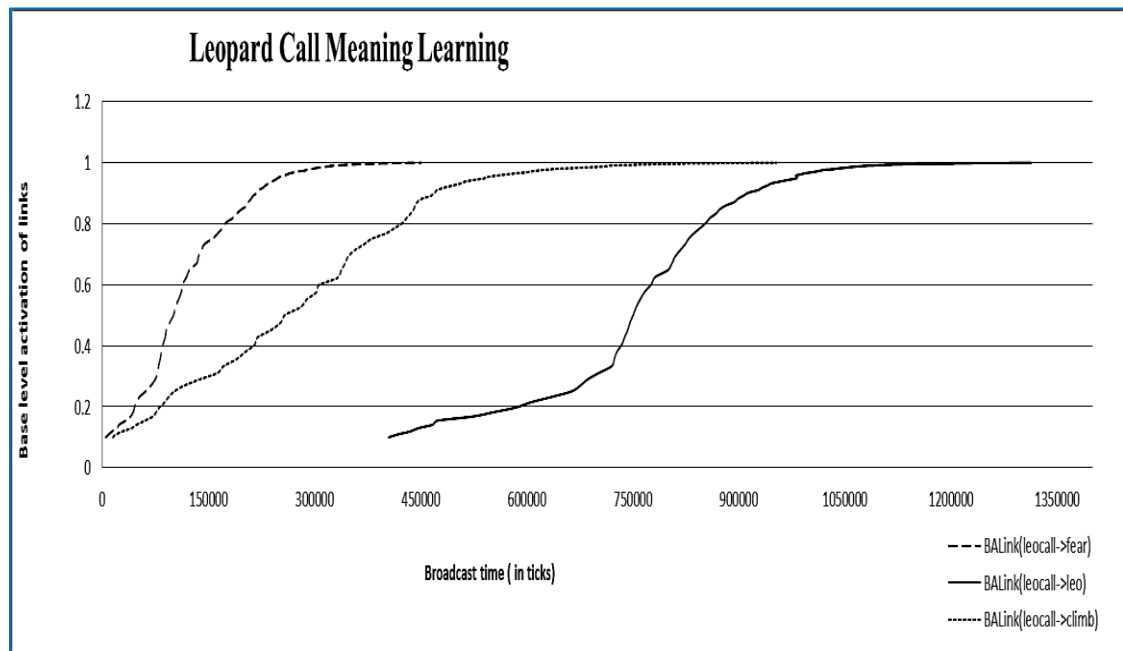

Figure 23. Base level activations of learned links from leopard call node to leopard node, climb tree node, and fear node at each broadcast time.
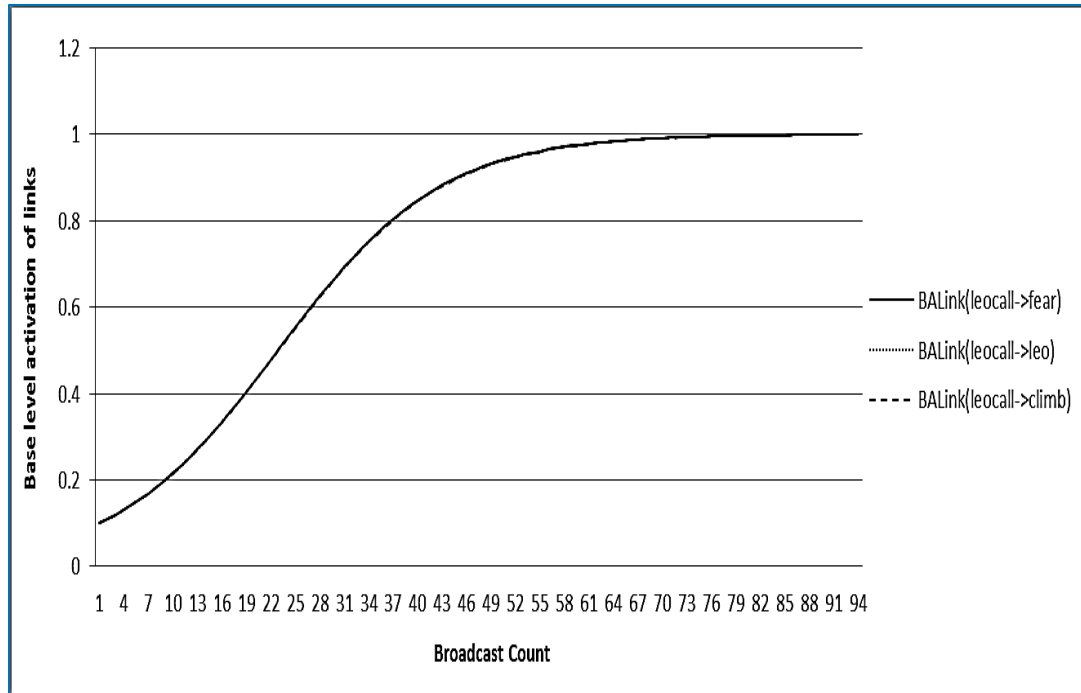
Figure24. Base level activations of learned links from leopard call node to leopard node, climb tree node, and fear node along the occurrences of these links in the consciousness.

Another set of results show that the INFANT's brain added new relationships leading from the leopard call to the fear feeling, climbing a tree, and the leopard, respectively. As mentioned previously, each simulation was performed using the same series of 35 randomly generated environments. The INFANT learned the various meanings of the leopard call in the following temporal order: The fear-based meaning at an average point of time equal to 446335.0476 (in ticks), the action-based meaning associated with climbing a tree, and lastly the reference-based meaning related to the leopard predator.

Figure 25. Time of convergence of learning each meaning of the leopard call

The results of learning the meaning of the three distinct vervet alarm calls differed as follow: In the eagle call and leopard call meanings results (Figure 17 and Figure 23), the vervet's mind quickly associated the eagle call with fear and hiding under bush and the leopard call with fear and climbing a tree. It takes much more time to associate the eagle call and leopard call with the eagle and the leopard, respectively. These results are consistent with what is expected to be learned in the wild. In fact, most of the time, the vervet infant is held by his mother. Hence, the INFANT is able to feel his mother's fear, perceive her hiding under bush, and climbing the tree faster than seeing the eagle and the leopard. However, the result of learning the snake call meaning (Figure 20) showed that the INFANT's mind associated the snake call with the fear, standing bipedally, and snake within a short time interval. In fact, the vervet infant is able to see the snake quickly in spite of being held by the mother most of the time.

Figure 26. Base level activation of learned link leading from leopard node to fear node at each broadcast time.

As explained before, the INFANT is not born with a fear of leopards. Our results show that the INFANT's brain added a causal relationship from the leopard to the fear feeling. This causal association is learned at a late point of time (in ticks) in the simulations. In the wild, the vervet infant is held by the mother most of time. Hence, it's expected that the INFANT will rarely spot the leopard.

In summary, the results illustrate the capacity of the INFANT agent, controlled by the LIDA cognitive architecture, to associate each alarm call with its multiple meanings in the following temporal order: The fear feeling, the equivalent escape action, and the corresponding predator class. This temporal order is in line with the primary assumption of the physical attachment between the MAMA and the INFANT. During infancy, the body position (picture 1) of the infant allows him to perceive the mother's fear and her escape action faster than spotting a predator in the vicinity.

***Picture* 1**. Body position of a vervet infant

Modeling the vervet alarm calls using the LIDA model can be viewed as the first step toward modeling human language understanding. Although, the human vocal system is more complex than the vervet vocal system, this work can be a foundation for learning the meanings of spoken words, especially that we adopted a multi-meanings assessment approach.

Experiment II: Evaluating the Understanding of Vervet Alarm Calls

The main advantage of modeling the learning of the meaning of vervet alarm calls using the computational and cognitive model LIDA is the ability to check the learning through looking at the implemented base level activations of the learned links in the Perceptual Associative Memory. In the wild, upon de-attaching from their mother, the vervet infants become capable of escaping appropriately upon hearing an alarm call. In order to simulate the reality of vervet monkeys, several experiments were performed to evaluate the INFANT's understanding of the meanings of alarm calls by gauging the correctness of the escape actions executed by the INFANT upon hearing an alarm call. The Perceptual Associative Memory of the INFANT in this stage comprises the newly

learned links leading from each alarm call node to the fear feeling node, the appropriate escape action, and the corresponding predator node.

The INFANT agent is de-attached physically from the MAMA agent in this stage of the simulation. Several schemes were added to the procedural memory of the INFANT to express this de-attachment. The schemes associated with the escape actions are set as follows:

Table 6

Schemes of the INFANT

Context	Action
Eagle node	Hide under Bush
Leopard node, fear node	Climb Tree
Snake node, fear node	Stand Bipedally

Other schemes were added to the procedural memory of the INFANT such as grabbing food and fleeing. The INFANT was assumed to know how to escape properly upon spotting a predator in the vicinity. Then, it was expected that the INFANT selects the suitable action upon hearing an alarm call. This understanding occurs by the means of the referential and causal relationships established during the first stage of the simulation. The INFANT's performance was calculated as the ratio of the number of correct escape actions of the INFANT after perceiving an alarm call to the total number of the actions of the INFANT including incorrect actions or no actions after perceiving an alarm call.

$$Performance = \frac{Number\ of\ correct\ escape\ actions}{Total\ Number\ of\ escape\ actions}$$

Table 7

Performance of understanding the meaning of various alarm calls

Alarm Call	Mean of performance
Eagle Call	0.760710864
Leopard Call	0.505581234
Snake Call	0.574280061

The INFANT has seven available actions in the ALife grid environment (e.g. hide under bush, climb a tree, stand bipedally etc.). Hence, the probability that the INFANT takes a random action is 0.125. The results show that the INFANT was able to escape correctly upon hearing an alarm call with an average performance. This is a good result in comparison with a random result. Additional procedural learning or tuning is needed to improve its overall performance.

5 CONCLUSION

We studied the vocal alarm calling system of vervet monkeys using a causation mechanism in order to propose an explanation of how the vervet mind learns the meanings of such communicative signals. For this purpose, a two-dimension simulation was designed and implemented using an ALife grid environment populated with an INFANT agent controlled by the LIDA cognitive architecture, and a MAMA agent, and other VERVET and predator agents controlled by production rules. Simulations were split into two categories: 1- The first stage of simulations were based on the assumption of the physical attachment between the INFANT and the MAMA, and aimed to test the convergence of learning the multiple meanings of distinct alarm calls; 2-The second part of the simulations were based on the assumption of the later de-attachment of the INFANT from the MAMA, and were done in order to check the comprehension of the alarm calls.

This work provides a research contribution in the following directions:

- A novel multiple meanings approach was adopted to study the meanings of vervet alarm calls. Three meaning types were considered successively: a feeling-based meaning, an action-based meaning, and a reference-based meaning. Approaching vervet alarm calls with multiple meanings can give us a fundamental insight on modeling human words which convey multiple meanings as well.

- Successful modeling of the meanings of vervet alarm calls using the LIDA cognitive architecture represents a first step toward realizing the goal of

language processing in LIDA, which is one of the important and complex high-level cognitive functions.

- The performed study was a good validation of the LIDA-based perceptual learning mechanism, particularly in learning relationships. The results, and especially the temporal order of learning the meanings of each alarm call, were consistent with the reality of vervet monkeys in the wild.

- The two-dimension grid ALife environment used in this study showed the importance of computational simulations in studying the convergence of meanings of simple communicative acts such as vervet calls, and it may also be an efficient tool in studying more complex vocal systems that have syntax, grammar etc.

REFERENCES

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

Baars, B.J. (1997). *In The Theater of Consciousness*. New York, NY: Oxford University Press.

Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Sciences,* 7(4), 166-172.

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

Bedau, M. A. (2003). Artificial life: organization, adaptation, and complexity from the bottomup. *Trends in Cognitive Sciences*, 7, 505-512.

Bickerton, D. (1990). *Language & species*. Chicago: University of Chicago Press.

Bickerton, D. (1996). *Language and human behavior*. London: UCL Press.

Brooks, R. A., & Stein, L.A. (1994). Building Brains For Bodies. *Autonomous Robots,* 1(1), 7-25.

Brown, R. S. (2006). *Domain Independent Perception in the LIDA Cognitive Software Agent*. (Unpublished master's thesis). University of Memphis, Memphis, TN.

Cangelosi, A., & Parisi, D. (2002). Computer simulation: A new scientific approach to the study of language evolution. In Cangelosi, A. and Parisi, D., editors, *Simulating the Evolution of Language,* chapter 1, pages 3–28. Springer, London.

Carlson, G. N. (1998). Thematic Roles and the Individuation of Events. Rothstein, S. (ed.) *Events and Grammar*. Dordrecht/Boston/London: Kluwer Academic Publishers, pp. 35-51.

Cheney, D.L., & Seyfarth, R.M. (1997). *Why animals don't have language*. The Tanner Lectures on Human Values, delivered at Cambridge University.

Chomsky, N. (1965). *Aspects of a Theory of Syntax.* Cambridge: MIT Press.

Cheney, D. L., & Seyfarth, R. M. (1990). *How monkeys see the world: inside the mind of another species*. Chicago: University of Chicago Press.

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

Duch, W., Oentaryo, R.J., Pasquier, M. (2008). Cognitive architectures: Where do we go from here? *In Proceeding of the 2008 Conference on Artificial General Intelligence 2008: Proceedings of the First AGI Conference, Pei Wang, Ben Goertzel, and Stan Franklin (Eds.).* IOS Press, Amsterdam, The Netherlands, The Netherlands, 122-136.

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

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

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

Franklin, S. (1995). *Artificial minds*. Cambridge, Mass.: MIT Press.

Franklin, S. (2005). Cognitive robots: Perceptual associative memory and learning. *In Proceedings of the 14th Annual International Workshop on Robot and Human Interactive Communication* (RO-MAN 2005) (pp. 427–433).

Franklin, S., & Graesser, A. (1997). Is it an Agent, or Just a Program? Taxonomy for Autonomous Agents. *Paper presented at the Proceedings of the Workshop on Intelligent Agents III, Agent Theories, Architectures, and Languages* 21-35.

Franklin, S., & F. Patterson Jr. G (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., & 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., Baars, B. J., Ramamurthy, U., & Ventura, M. (2005). The Role of Consciousness in Memory. *Brains, Minds and Media*, 1, 1–38.

Gallagher, S. (2000). Philosophical conceptions of the self: implications for cognitive science. *Trends Cogn Sci.,* 4, pp. 14– 21.

Gallup, G. G., Jr., Anderson, J. L., & Shillito, D. P. (2002). *The mirror test*. In M. Bekoff, C. Allen, & G. M. Burghardt (Eds.), the cognitive animal: Empirical and theoretical perspectives on animal cognition (pp. 325–333). University of Chicago Press.

Glenberg, A. M., & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language,* 43(3), 379–401.

Harlow, H.F. & Harlow, M. K. (1969). Effects of various mother-infant relationships on rhesus monkey behaviors. In B.M. Foss (Ed.) *Determinants of infant behavior* (Vol.4).London: Methuen.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1-3), 335-346.

Hofstadter, D. R., &Mitchell, M. (1994). An overview of the Copycat project. In K. Holyoak & J. Barnden (eds), *Advances in Connectionist and Neural Computation Theory, Volume 2: Analogical Connections.* Norwood, NJ: Ablex Publishing Corporation.

Kirby, S. (2002). Natural Language From Artificial Life. *Artificial Life,* 8(2), 185-215.

Lemasson, A., Gautier, J. P., Hausberger, M. (2003). Vocal similarities and social bonds in Campbell's monkey (Cercopithecus campbelli*). Comptes Rendus Biologies*, 326, pp. 1185- 1193.

Lemasson, A., Hausberger, M., Zuberbühler, K. (2005). Socially meaningful vocal plasticity in Campbell's monkeys. *J Comp Psychol* 119: 220–229.

Loula, A., Gudwin, R., Ribeiro, S., & Queiroz, J. (2010b). On Building Meaning: a biologically-inspired experiment on symbol-based communication. In Cutsuridis, V., A. Hussain, A.K., Barros, I., Aleksander, L., Smith, & L., Chrisley, editors (Eds) . *Brain Inspired Cognitive Systems,* pp. 77-94. USA: Springer.

Madl, T., Baars, B. J., & Franklin, S. (2011). The Timing of the Cognitive Cycle. *PLOS ONE,* 6(4), e14803.

McCall, R., Franklin, S., & Friedlander, D. (2010). Grounded Event-Based and Modal Representations for Objects, Relations, Beliefs, Etc. *Paper presented at the FLAIRS-23,* Daytona Beach, FL.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review,* 83, 435-4.

Oliphant, M. (1997). *Formal approaches to innate and learned communication: Laying the foundation for language*. (Unpublished doctoral dissertation). University of California, San Diego.

Oller, D. K., & Griebel, U. (2004). *Evolution of communication systems a comparative approach*. Cambridge, Mass.: MIT Press.

Peirce, C.S. (1998). *The Essential Peirce Selected Philosophical Writings, Volume 1* (1893-1913), Peirce Edition Project, eds., Indiana University Press, Bloomington and Indianapolis, IN.

Rajala, A. Z., Reininger, K. R., Lancaster, K. M., & Populin, L. C. (2010). Rhesus monkeys (Macaca mulatta) do recognize themselves in the mirror: implications for the evolution of self-recognition. *PLOS ONE*, 5(9), e12865.

Ramamurthy, U., Franklin, S. (2011). Self-System in a model of Cognition. *In Proceedings of Machine Consciousness Symposium at the Artificial Intelligence and Simulation of Behavior Convention* (AISB'11), University of York, UK, p 51-54.

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.

Snaider, J., McCall, R., & Franklin, S. (2011). The LIDA Framework as a General Tool for AGI. *The Fourth Conference on Artificial General Intelligence (Springer lecture notes in artificial intelligence),* Mountain View, California, USA.

Steels, L. (1997). The Synthetic Modeling of Language Origins. *Evolution of Communication,* 1(1), 1-34.

Waal, F. B., c, M., Freeman, C. A., & Hall, M. J. (2005). The monkey in the mirror: hardly a stranger. *Proceedings of the National Academy of Sciences of the United States of America*, 102(32), 11140-11147.

Worden, R.P. (1996). Primate Social Intelligence. *Cognitive Science: A Multidisciplinary Journal*, 20(4), 579-616.

Wray, A. (1998). *Protolanguage as a holistic system for social interaction*. Language and Communication, 18.47–67.

Zangenehpour S., Ghazanfar A. A., Lewkowicz D. J., Zatorre R. J. (2009). Heterochrony and cross-species intersensory matching by infant vervet monkeys. *PLOS ONE* 4, e4302. Doi: 10.1371/journal.pone.0004302.