# PrimEmo: A Neural Implementation of Survival Circuits Supporting Primitive Emotions

**Luc Caspar**[1] and **Roger K. Moore**[2]

**Abstract.** Affective and cognitive sciences are both fields interested in the inner workings of the brain. While affective science focuses on the concept of emotions and how they are produced, cognitive science considers the brain as a whole, treating it as a system made of a multitude of independent subsystems. Since its inception, affective science has produced a plethora of theories and models each trying to solve the mysteries surrounding the definition and origin of emotions in the brain. Cognitive science on the other hand has had a rather chaotic history shifting its focus every decade, before finally being influenced by computer science and adopting the point of view of the brain as the most elaborate computational device. In this point of view, the gray matter residing within our skull is reduced to a system that takes sensory information on its inputs, processes it and outputs more data or actions. Each field has, more or less, evolved independently from one another so far, but both are now facing fundamental problems. On the one hand, the concept of emotions has yet to be completely defined and modeled. On the other, cognitive science is still trying to produce architectures imbuing artificial agents with human-level intelligence. This paper introduces a neural architecture based on the "*survival circuits*" framework (LeDoux, 2012) supporting primitive emotions and providing survival skills to artificial agents. Side-stepping the problem of defining the concept of emotions, the suggested neural structure focuses on identifying and modeling parts of the brain involved in survival functions (defense, thermo-regulation, maintenance of energy, nutritional supplies, reproduction and fluid balance). The neural implementation of this system provides a proto-brain for groups of artificial agents trying to survive in a dynamic virtual environment. By comparison with a hard-coded control logic (Scheutz, 2004), our architecture allows for a more equitable sharing of the resources and a longer life expectancy. It is our belief that, in the long term, the system suggested in this paper could become a robust basis upon which more elaborate cognitive architectures, such as ACT-R or SOAR, could be built, hence moving one step closer to endowing artificial agents with human-level intelligence.

## 1 INTRODUCTION

Ever since humans have walked the earth it seems they have been fascinated by emotions, how they shape our everyday life, our inner thoughts and our relationship with each other. Due to their subjective, nature everyone has his own thoughts and theory on the matter. However, even though so many minds are at work trying to solve this puzzle, we are nowhere near a complete understanding of "*What is an emotion?*" [15].

[1] University of Sheffield, United Kingdom, email: LCaspar1@sheffield.ac.uk
[2] University of Sheffield, United Kingdom, email: r.k.moore@sheffield.ac.uk

Since the time of Aristotle [2] in ancient Greece, emotions (called "*passions*" at the time) have been separated from reason, if not opposed to one another, in the human mind. Passion, according to Aristotle, was the fire that burns from the inside and makes people act irrationally or on impulse. While reason was seen as the cold blooded, carefully calculating and logical part of the mind, that one uses to solve problems and manage one's own life. This seemingly natural separation of the mind in two halves, passion against reason, would endure and lead to the creation of the affective and cognitive sciences, but that would not happen until the birth of René Descartes and William James, respectively.

In 1637, Descartes opened the gate of cognitive science with his now famous "*Cogito, Ergo sum*" in "*Discourse on the Method*" [7], sparking the interest of many a scientist on the subject of the human brain and its functioning. Ironically, 12 years later he would also write what is now considered to be the first theory on emotions in his "*Passions of the Soul*" [8]. While the behaviorists took over the field of cognitive science [5, 21], it would not be before James' paper titled "*What is an Emotion?*" [15] that the field of affective science would really see the light of day. Each field went its separate way, the behaviorists dismissing the concept of emotions since it was thought to be part of the conscious mind, impossible to quantify or measure and thus of no concern to them [21, 5]. Affective science meanwhile began its age-old debate over the definition of emotions [4].

Many cognitive architectures have since been developed, such as ACT-R [35], SOAR [18] and EPIC [16]. Although, the ACT-R and EPIC theories are devoid of any emotions and SOAR, even if possessing an emotion module, makes only limited use of the information output. Similarly, a plethora of theories and models of emotions are now available, but as Grandjean, Sander and Scherer [29] noticed there are maybe too many of them with no connections in-between and no ways to compare them. The problem has grown since the appearance of service robots, focusing the attention on emotions recognition and expression to ease human-robot interaction, distancing the whole field from the fundamental emotion riddle. There is clearly a need for a different approach if the field is to move forward again, but where should one begin to look for a solution?

Sloman [34] and LeDoux [22], both suggested a similar solution: acknowledging the fact that the concept of emotions is still blurry at the moment and rather concentrate on the functions fulfilled by emotions. In what LeDoux calls a theory of "*Survival Circuits*" (detailed in section 2), emotions are seen as a system that evolved in every living creatures as a better survival strategy. Animals that had emotions were better at detecting dangers in their surroundings and so more efficient in avoiding them. The same is true of happiness that would usually drive animals to repeat actions that made them happy in the first place.

It is our opinion that most cognitive architectures, be it SOAR, ACT-R or EPIC, are lacking a robust emotion basis upon which to build a truly intelligent system (see DAC [38] and LIDA [12] for different attempts at solving this problem). This paper then introduces a new neural architecture in section 2, which focuses on the innate circuits and the survival functions of emotions. The goal, discussed in section 4, is to show that artificial agents based on this architecture are better at adapting to and surviving in dynamic environments than other cognitive agents. Finally, section 5 highlights the steps that remain to build a complete cognitive architecture endowed with emotions.

## 2 METHOD

### 2.1 Survival circuits

The idea of hard wired structures residing in the older part of the brain and related to emotions is not entirely new. Its origin could be traced back to Darwin [6], who first suggested the existence of innate circuits common to all humankind and even part of the animal kingdom, enabling us to feel and express a certain set of emotions in similar fashions. According to this theory of basic emotions, it is the reason why even if one was to travel to the most remote regions of the world, a simple smile would suffice for everybody to understand that you are in fact happy. The term "*Basic emotions*" was coined for those universally understood emotions, while the rest are usually referred to as "*Complex emotions*", since one's culture is embedded into and required for their expression. Scientists at the time were quick to adopt the concept and each set out to define the set of basic emotions shared by all and more importantly to describe at the same time, which were the related circuits in the brain. Two of the most notable researchers engaged in this furious race are Paul Ekman, whose work is still the basis for many a project to this day and who developed a standard set of six basic emotions: fear, anger, happiness, sadness, disgust and surprise [9]. And Robert Plutchik, most famous for the graphical representation of his theory, the aptly named "*Wheel of emotions*" (see Figure 1), which described basic emotions as being similar to primary colors and complex emotions were obtained by combining those primary colors together in particular ways [27].

However appealing the theory of basic emotions might be, it has been heavily challenged in the recent years. Most prominent among all the criticisms, was the fact that sets of different sizes and containing different labels were put forth by each research project working on the subject [4, 5, 21, 20]. Furthermore, scientists came to realize that some of the labels used in the nomenclature could describe two similar, but still different states of arousal [22]. Another important issue, pointed out by the opponents of the basic emotion concept, was the fact that there is a discrepancy between the set of basic emotions listed by the main theory, based on brain research in animals, and the sets obtained by research projects using human subjects, including Ekman's work. The advent of imaging technologies, such as functional Magnetic Resonance Imaging (fMRI) and Positron Emission Tomography (PET) scans, that allowed to directly observe the brain's activation patterns while patients were feeling and expressing emotions, brought the debate over the veracity of the basic emotions theory to an end. Even though, the spatial resolution of this type of technology is still lacking, the evidence pointed toward the fact that patterns of activations did not have dedicated circuits, as previously believed. Nowadays, the theory of Basic emotions has lost most of its supporters, but some parts of its core concepts can still be found in the work of Orthony [10], Alexandrov [10] and others.

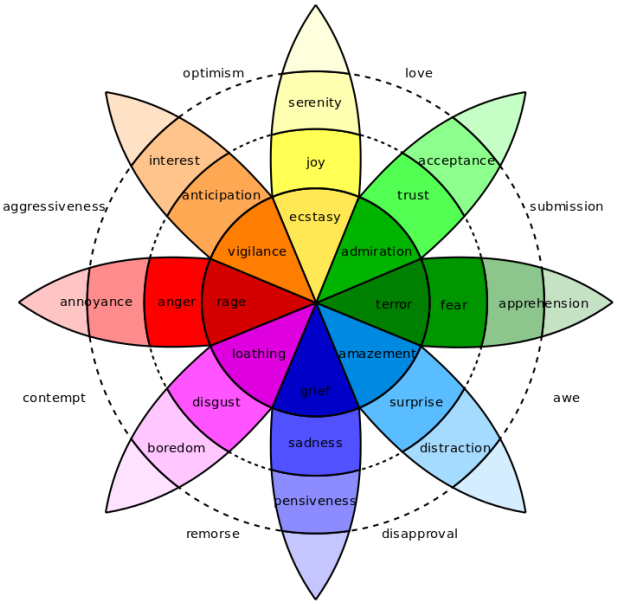


**Figure 1.** The wheel of emotions as a visual representation of Plutchik's theory on basic emotions and their interactions. Source: Wikimedia commons, under public domain licensing.

The survival circuits framework introduced by LeDoux in 2012 [22] relies on the same fundamental idea of innate circuits nested at the base of the brain. However, the key difference here is that those circuits are not dedicated to the production of any emotion or feeling. Survival circuits have their origin in the single cell organisms that populated the earth billions of years ago. Those early life forms, even though very simple, were able to detect and avoid noxious chemical components and steer themselves toward the more nutritious ones. Since then, evolution has run its course, but LeDoux argues that the blueprint for the basis of the vertebrate brain is pretty similar across species; parts of the human brain can also be found in a rat brain and vice versa. Obviously, there are some differences in size and shape, but the same primitive mechanisms that help us survive are present in the rat's, as well as, the human's brain. Thus, the primary function of survival circuits is to ensure the survival of the individual by monitoring the state of the internal state and detecting relevant changes in the environment. Each circuit can be conceived as a control system tasked to maintain one of the individual's basic needs within a comfortable range and initiate a reaction if any stimulus were to unsettle the established balance.

The similarity that Darwin observed in the expression of emotions across culture and species are reduced to the fact that the set of innate triggers and reactions defining each survival circuit has been inherited from a common ancestor. Similarly, the existence of complex emotions influenced by culture can be explained by the ability of survival circuits to create new associations between insignificant stimuli and innate reactions through the process called conditioning [26].

However, if survival circuits are not dedicated to producing emotions, how can they constitute a basis for a model of emotions? According to LeDoux, the triggering of an innate circuit will simultaneously send signals for the release of hormones, in both the brain and body, and to sub-cortical and motor areas to react accordingly. The

sensory feedback from the reaction coupled with the knowledge that an innate detector has been triggered, the formation of new memories and the heightened state of activity in the brain, are all gathered into a global organismic state that can be processed further by higher-level cognitive functions. The job of these functions will be to categorize and finally assign an emotion label to the global organismic state [22].

## 2.2 PrimEmo: A neural implementation of the survival circuits

Relying on LeDouxs framework of survival circuits, this section investigates a plausible model for such mechanisms, before implementing it in populations of artificial neurons as the main control process for primitive robots.

### *2.2.1 The hypothalamus*

The hypothalamus in the mammalian brain is part of a larger complex called the diencephalon along with the thalamus and the pituitary gland. It is symmetrically distributed around the third ventricle and its anterior boundary, laying over the optic chiasm, is defined by the lamina terminalis and the anterior commissure, while its posterior border is occupied by the mammillary bodies [30, 11, 37, 17, 3].

In the vertebrate brain the hypothalamus is in control of the Autonomic Nervous System (ANS) [30, 11, 37, 17, 3] responsible for the well-being of the body's internal state. For a long time it was believed that the hypothalamus contained different command centers, for aggression, energy, thermoregulation, fluid balance and reproduction, that could realize fully fledged behaviors. However, recent discoveries suggest that, on the contrary, the hypothalamus only influences the likelihood for a behavior to be executed by sending signals to widespread areas of the forebrain, cortex and toward the pituitary gland for the release of hormones. Its functioning is therefore similar to a control system, processing sensory information and adjusting the release of hormones in the brain and blood to maintain the body's homeostasis [30].

Given its role as a control system and its innate propensity to detect stimuli relevant to the self-preservation of the body, we consider the hypothalamus as being one such survival circuit.

### *2.2.2 The amygdala*

In the vertebrates brain, the amygdala refers to a small almond shaped collection of nuclei that lies in the temporal lobes pole below and medial to the cortex on each side [11, 37, 17, 3, 40].

In the field of affective science, the amygdala has always been the center of attention, for it was believed to be the main constituent of the limbic system introduced by Paul MacLean [23]. As the concept of the limbic system survived throughout the years, so did the idea of the amygdala being the production center for emotions [5, 22, 3]. Results from studies on the expression of fear in animals concluded that the amygdala was not the only part involved in the process as many already suspected, but that it was essential to the elicitation of negative emotions, such as fear, anger and disgust [28, 21, 19, 5]. However, with the advent of new technologies, especially the fMRI, researchers discovered that the amygdala was involved in processing positive emotions as well [39].

As of recently, some scientists have begun to suspect that the amygdala might only play an indirect role in the production of emotions and that it would be better to think of it as an innate detector processing external sensory information[22]. Akin to the hypothalamus the amygdala can be seen as a control system (see Figure 2) responsible for monitoring the external environment and reacting to stimuli relevant to the individual's survival. As a complement to the hypothalamus, the amygdala is the second survival circuit in this model.

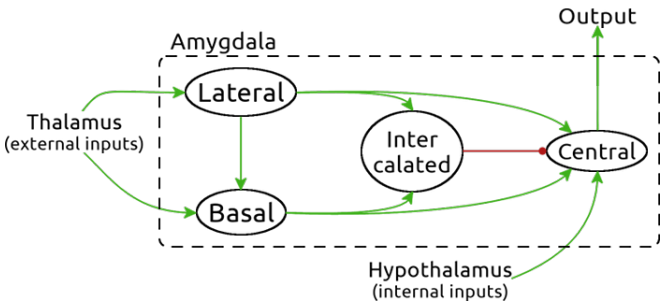


**Figure 2.** Diagram outlining the inner workings of the amygdala. The lateral and basal nuclei act as the main inputs, propagating information to both the intercalated neurons and the central nucleus. The central nucleus, the main ouput, combines information from the hypothalamus with data processed inside the amygdala. The intercalated neurons are probably the most important part of the amygdala-thalamus feedback loop, as they prevent the central nucleus from over-activating by gradually inhibiting its activity, hence implementing a simple "*habituation*" mechanism.

### *2.2.3 The thalamus*

The thalamus is the second part of the diencephalon. Roughly egg-shaped, it sits in the gap between the posterior wall of the interventricular foramen and the subarachnoid space below the fornix and the splenium of the corpus callosum and above the pineal gland and tectum [32, 33, 11, 37, 17, 3]. The thalamus is usually referred to as the gateway to the brain, as it is believed that any external sensory data must first go through the thalamus on its way to the cortex or other parts of the brain [32, 33]. For a long time, the role of the thalamus was reduced to a mere proxy that redirects any incoming information toward its corresponding processing center in the neocortex [33, 36]. However, recent studies have revised this view and found that, through bidirectional connections with wide spread areas of the brain, feedback loops could be defined, where the thalamus would first send sensory data to a sub-system of the brain. The sub-system would then process the information further to try and make sense of it. Finally, a signal would be fed back to the thalamus, modulating the flow of sensory information entering the sub-system.

As the thalamus' core function is not directly related to the survival of the individual, but rather to the restriction on the amount of information entering the brain, its inclusion might seem pointless. Nevertheless, the thalamus is a necessary constituent for the amygdala to correctly perform its role as a control system.

## 2.3 Experimental setup

To assess the performance of our architecture in terms of survival capabilities, we replicated a "resource foraging task" experiment suggested by Scheutz [31] and compared the subsequents results to those originally obtained by Scheutz.

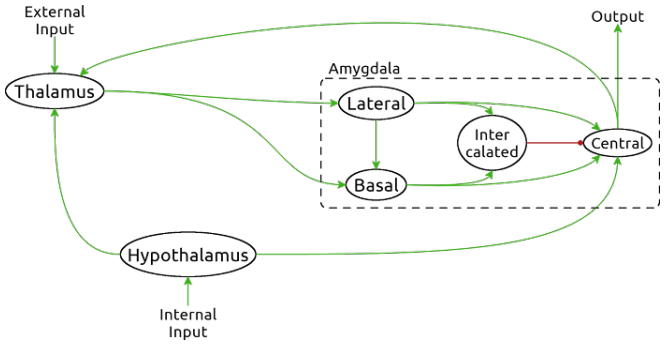


**Figure 3.** A complete view of the brain areas included in the PrimEmo model and their interactions.

The main simulation environment is quite simple and consists in a virtual unlimited two-dimensional floor, on which energy resources, obstacles and robots are placed. Obstacles are static columns placed randomly inside a predetermined area of $8 \times 8$ meters at the beginning of the game and will remain in place until the end. On the contrary energy resources, which are also randomly scattered inside the area, will be depleted by the foraging robots and disappear. However, no agents should ever lack energy, for resources are continually replaced. Each source contains a maximum of 400 units of energy. In the initialization phase 25 units of both obstacles and energy sources will be placed randomly in the area. Furthermore, even though resources are allowed to appear randomly for each turn, the number of resources that can be on the map at one time has been limited to 50 units for reasons of computational power.

The robot used in this simulation to represent an agent is the E-puck [25]. It comes equipped with 8 infra-red distance sensors, allowing it to detect objects in every directions, and an rgb camera. It has been selected for its small size and its availability, since this experiment, if implemented with real agents, involves groups of up to 50 robots. In the frame of the simulation, an agent can perform one out of five possible actions:

- Move: selected by default if nothing appears in the robot's immediate surroundings. Moving around expands energy at a rate proportional to the robot's linear speed. More over, if a robot is below the critical level of energy (set to 400 units) its maximum speed will be limited to $\frac{1}{7}^{th}$ of its original value.
- Forage: when the agent is on top of an energy source it will consume energy at a rate of 100 units per turn. There is no limit to how much energy an agent can consume.
- Fight & Flee: when two or more agents meet, they enter a battle mode, each agent has to choose whether it will fight or run away. Knowing that if an agent decides to stay and fight it will consume 50 units of energy, while running away depletes its energy level by an amount equivalent to moving around at full speed for 10-20 turns. Each agent will make the decision following an algorithm based on its "*action tendency*", a coefficient between 0 and 1 indicative of the probability the agent will fight.
- Procreate: If an agent manages to survive for more than 250 turns and has more than 2200 units of energy left, it will then automatically procreate. This means that a new agent will appear at a random location on inside the $8 \times 8$ meters area. The act of procreation costs 2000 units to the agent.

In the initialization phase, a group of 25 robots are scattered inside the area, each with an initial energy level of 2000 units. For the same reason as the sources, there cannot be more than 50 robots at one time on the map. If an agent fulfills the procreation criteria, but the robot population has already reached its maximum, the agent will be allowed to choose another action to perform and no energy is consumed for the procreation.

When moving both the speed and direction of the robot are computed according to the following equation:

$$D = \sum_n g_r \times resource(n) + \sum_m g_a \times agent(m) + \sum_k g_t \times obstacle(k) \quad (1)$$

Where $g_r$ is a coefficient indicative of the agent's need in energy, $g_a$ represents how much a robot wants to enter a battle (if positive the robot will be moving toward others, while a negative value means that the robot will actively avoid others), and $g_t$ correspond to obstacle avoidance (similarly, a positive value means the robot will try as much as possible to run into obstacle, while a negative value forces it to avoid obstacles). $resource(n)$, $agent(m)$ and $obstacle(k)$ are the scaled vector from the agent's current position to the $n^{th}$ resource, $m^{th}$ nearby agent (excluding itself) and $k^{th}$ obstacle, respectively.

Finally, at the end of a turn each agent loses 1 unit of energy to processing power, urging them to move around and forage for energy as much as possible, since an energy level of 0 means death.

## 2.4 Hard coded control scheme

In Scheutz's experiment four types of agent were defined: asocial non-adaptive, social non-adaptive, asocial adaptive and emotional (social and adaptive). When a social agent finds itself in a conflict, it fights only if its action tendency is the highest, otherwise it runs away. Thus, battle between social agents are resolved after one turn only. Asocial agents, in the same situation, do not take the action tendency of others into consideration, but use their own action tendency as a fighting probability. Therefore, a conflict between asocial robots can last until one of them decides to run away or dies.

The capacity of an agent to adjust its action tendency makes the difference between adaptive and non-adaptive individuals. On the one hand, if a robot is non-adaptive then its action tendency is fixed and chosen at random upon initialization, following a Gaussian distribution of mean 0.5 and a standard deviation of 0.125. Adaptive agents, on the other hand, adjust their action tendency depending on whether they lost or won the last fight. A losing agent increases its action tendency, thus increasing its probability of winning the next battle, while a winning agent decreases its action tendency.

The adaptation rule (AR), as described by Scheutz [31], defines the amount by which action tendency should be increased or decreased. If $r$ is the basic action tendency (randomly initialized at the beginning of the simulation) and $m$ is the current action tendency, then:

- After a loss:

$$AR^+(m) = \begin{cases} m + \frac{(1-m)}{2} & \text{if } m \geq r \\ 2 \times m & \text{if } m \leq \frac{r}{2} \\ r + (2 \times m - r) \times \frac{(1-r)}{2 \times r^2} & \text{otherwise.} \end{cases} \quad (2)$$

- After winning:

$$AR^-(m) = \begin{cases} 2 \times m - 1 & \text{if } m \geq r + \frac{1-r}{2} \\ \frac{m}{2} & \text{if } m \leq r \\ \frac{r}{2} + \frac{r \times (m-r)}{(1-r)^3} & \text{otherwise.} \end{cases} \quad (3)$$

Emotional agents will have the added possibility of setting a non-zero value for their fighting tendency $g_a$, by following the equation:

$$g_a = 100 \times actiontendency - 50 \quad (4)$$

Regardless of the agents type, the value for the other gains are fixed to $g_r = 20$ and $g_t = -20$.

## 2.5 Control based on PrimEmo

This second control scheme used to investigate the performances of the PrimEmo architecture is split into two main components. On the one hand, the model of survival circuits detailed in the previous sub-section (see Figure 3) has been implemented with the LEABRA framework in the Emergent software [1]. Each part of the brain being realized by a different set of layers of artificial neurons. To drive its activation and act as a sort of proto-brain for each artificial agent, the network takes as input:

- the sensory data coming from all eight of the proximity sensors;
- the maximum value of the action tendency sampled over all nearby robots, if any;
- the agent's energy level;
- and a "*fear*" value that is proportional to the distance between the agent and the closest obstacle on the map.

Due to the algorithm behind both the activation and learning processes in the LEABRA framework, an additional set of inputs is fed back to the network for learning purposes. Since the architecture is still in its primitive form and not imbued with the concept of "*good*" and "*bad*", the learning signal is made of four numbers, in the range $[0, 1]$ representing a correct value for each of the $g_r$, $g_a$, $g_t$ and $action\_tendency$ for the present situation. Once all inputs and learning signals have been sent to the network, it alternates between activation phases and learning phases, until it reaches an equilibrium. At the end of the cycle, the activation values on the network's output layer represent the updated coefficients that influences the robot's behavior for the next turn.

A python script acts as both an interface between the proto-brain and simulated robot, and as a pre-motor area sending commands for the next action to perform, to the actuators. In its role as an interface, the script keeps track of the agent's state, that is:

- the level of energy;
- the agent's position in a 2D plane;
- the agent's velocity along the same axes;
- the maximum value for the action tendency, sampled over all nearby robots;
- and the distances between the agent and the closest obstacle and energy source.

With this information, the script is capable of computing the "*fear*" value, as well as the correct values for the four coefficients. All required inputs are then formatted in an appropriate structure and sent to the brain.

Once the network has reached equilibrium, the script retrieves the activations of the output layer and uses them to determine the next action to perform, following the process described in sub-section 2.3. At the end of each turn it sends commands to the simulated environment, to either move the agent, fight, flee, forage or procreate. In case an agent finds itself in a battle situation, similarly to Scheutz's emotional agents, it only fights when its action tendency is higher than the action tendencies of all other agents involved in the battle. Otherwise it flees for a random number of turns chosen in the range $[10, 20]$.

In contrast to Scheutz's control scheme, where some coefficients were static and some bound by equations, agents in this group are free to set all their coefficients. Moreover, since they are learning over time, it is possible for them to optimize their strategy for survival.

# 3 RESULTS

The results presented in this section were obtained after running two series of 40 simulations, each simulation lasting 150,000 cycles and initialized with 25 robots, 25 obstacles and 25 energy sources. The first set of simulations were run using the "*emotional*" agents defined by Scheutz, since those were the most efficient of all the different types present in the experiment [31]. While agents for the second set of simulations were controlled by our neural implementation of the survival circuits. In both cases the simulation was executed in a Webots [24] virtual environment.

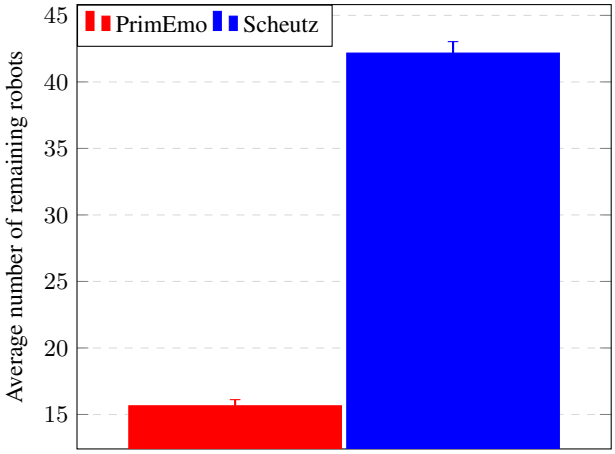


**Figure 4.** Diagram comparing the average number of robots still alive at the end of the simulation.

Figures 4 and 5 are concerned with the average number of remaining robots and life expectancy (expressed in number of simulation cycles) of a virtual agent for both control schemes. These results are the most important ones since they provide a way to quantify the performance of each scheme in regard to the general survival goal.

Seeing that the results obtained during the simulation are not as clear cut as one would have expected (see section 4 for more details), the diagrams and plot of figures 6, 7 and 8 will help in analyzing the difference in strategies adopted by the groups of robots for each control scheme.

The figures 6 and 7 represent the number of cycles during which a robot executes the "*move*" and "*fight*" actions respectively. Knowing that before the "*fight*" action becomes available to a robot, it will have to meet with at least one of its peers.

Finally, the level of energy for each agent was extracted for each simulation cycle and the average difference between the robot with the lowest and the highest levels of energy has been plotted in Figure 8. This difference can be seen as a measure of how equitable the repartition of energy resources is within a group of robots. Additionally, it will help shed light onto the reason why fewer agents

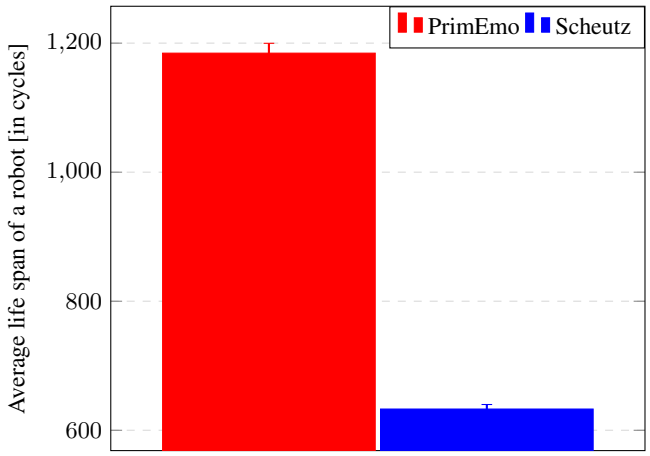


**Figure 5.** Diagram comparing the average life span, in cycles, of a virtual agent.

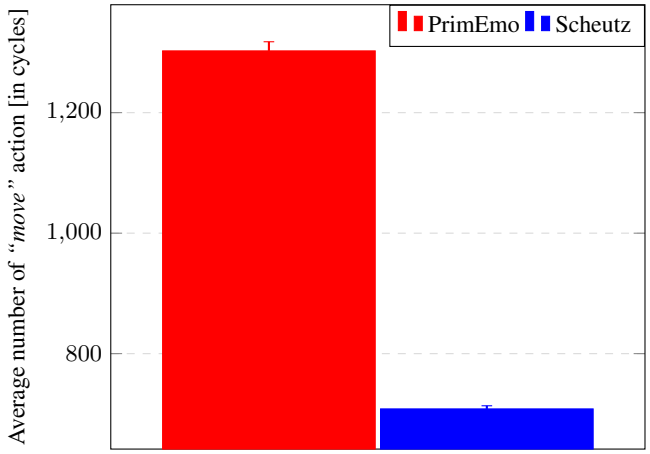


**Figure 6.** Diagram comparing the average number of cycles a robot spends moving in the environment.

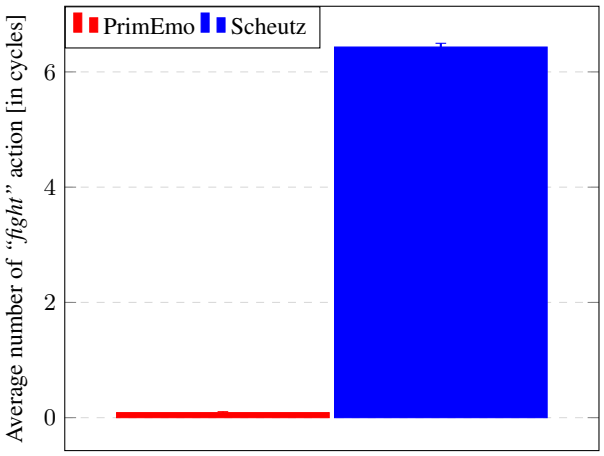


**Figure 7.** Diagram comparing the average number of cycles a robot executes the "*fight*" action.

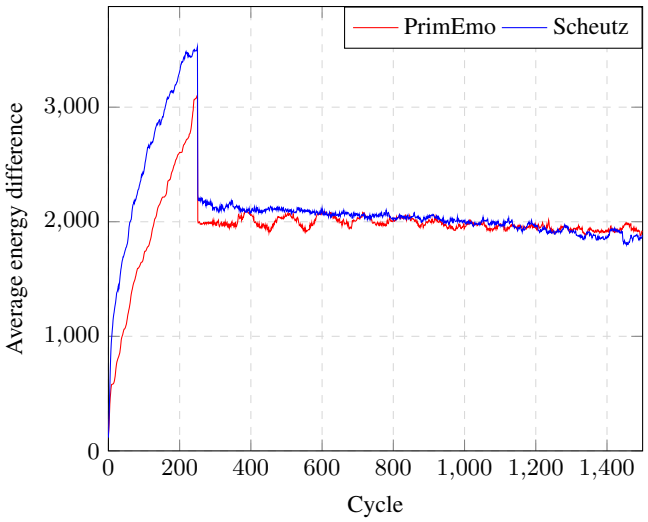


**Figure 8.** Diagram comparing the average difference between the robot with the highest and the robot with the lowest energy levels.

controlled by the PrimEmo architecture remain at the end of a simulation, when compared to Scheutz's control scheme.

## 4 DISCUSSION

The results presented in the previous section might, at first glance, appear odd since neither control schemes seems to have the upper hand in the race for better survival (see Figures 4 and 5). The virtual agents using the PrimEmo architecture as brains were indeed unable to increase or even maintain the original population number until the end of the simulation, but they did live almost twice as long as their counterparts controlled by Scheutz's algorithm. While the contrary is true for Scheutz's group of robots. However, by putting the first two diagrams into a broader perspective encompassing data on the number of cycles dedicated to moving and fighting, and the differences in energy within a group, an interesting picture can be drawn and it becomes possible to explain parts of those seemingly odd results.

The overall question could be rephrased as follows: if robots are able to live longer, why is it that so few of them remain in the end? The answer comes in two parts, the first has to do with the strategy for sharing energy resources amongst a population. By analyzing the average differences in energy between the robot with the highest and the one with the lowest levels (see Figure 8), it is possible to infer that, on the one hand, the group using the PrimEmo brains implement a more equitable strategy, since the difference in energy always remain inferior to that of Scheutz's group. As a consequence, robots at the low end of the energy spectrum have an easy access to energy resources and are allowed to live longer, hence the extended average life span (Figure 5). At the other end of the spectrum, robots with high energy levels have to move to avoid their peers (Figures 6 and 7) and find sources to forage, thus expanding a lot of their precious energy into movement rather than saving it for procreation. This results in a better quality of life for the overall population. However, and here is where the second part of the answer intervenes, because of the high energy requirements for procreation, it also means that

dying agents are not necessarily replaced and the population slowly decreases.

On the other hand, the strategy implemented in Scheutz's algorithms tends to share energy resources more equally. Therefore, robots with low energy will often have to fight for their right to forage a source (see Figure 7), and robots with high levels spend less time avoiding others (see Figure 6). This results in a bigger gap between high and low energy robots, with low energy robots dying quickly (see Figure 5) and being replaced as quickly by the offspring of high level robots, which are easily able to meet the procreation requirements. Moreover, since Scheutz's algorithm has been optimized for this experiment, the rate at which procreation happens is higher than the death rate resulting in an overall population increase (see Figure 4).

## 5 FUTURE WORK

At the moment the PrimEmo architecture is still a work in progress, as evidenced by the fact that the structure cannot even differentiate "*good*" from "*bad*". This means that primitive emotions supported by the network are really only states of arousal triggered by stimuli endangering the agent's homeostasis.

To complete the PrimEmo architecture and support full-fledged primitive emotions, two more models would have to be incorporated. First, would be the Primary Value Learned Value (PVLV) model, introduced by Hazy, Frank and O'Reilly [14], describing the process by which the dopaminergic system, along with the ventral striatum, amygdala and hypothalamus, is able to predict the reward associated with an action not yet performed and how this same system is capable of learning from its own mistakes. The PVLV would endow the PrimEmo architecture with a sense of "*right*" and "*wrong*", and with a robust learning mechanism based on reinforcement. The second model of interest for this project, is the Pre-frontal cortex Basal ganglia Working Memory (PBWM) structure, suggested by Hazy, Frank and O'Reilly [13] to explain the gating mechanism implemented by the basal ganglia in the brain.The PrimEmo architecture, would greatly benefit from the PBWM model, since it is capable not only to manage working memory, to a certain degree, but also to make concrete decisions about what action to perform, rather than simply influencing the decision process from afar.

On the long term, and once the PrimEmo architecture has been completed and thoroughly tested, it is our belief that it could be used as a replacement for the core module of the famous ACT-R system. In ACT-R the role of the main module, is to consider all stored production rules and choose the one best suited, given the situation. This role could easily be filled in by both the PBWM and PVLV models. Although, the PrimEmo architecture could offer more than a simple replacement for the production module, since it is able to keep track of the state of both the agent's internal and external environments, manage its own working memory and maintain its homeostasis, the data produced by the PrimEmo module could also be used to influence the functioning of all the other modules in the system. The modified PBWM and PVLV part of the structure could be used as part of the update mechanism for both the intentional and declarative modules, along with their associated buffers. The more primitive part of the architecture, presented in this paper, can also imbue the commands sent to the manual module with a sense of urgency. However, our belief that the PrimEmo architecture could replace, if not enhance, the core module of the ACT-R system, do not stem from only its capabilities. The constraint, imposed by the ACT-R theory, that modules can only communicate with the central hub (and even then only through buffers), but not with each others, would ultimately force PrimEmo into the central position, since it needs to be able to monitor the global state of the system and to talk to every part of it.

Be that as it may, the integration of the PrimEmo structure into the ACT-R system still remains theoretical for the moment. As mentioned in the previous section, the first part of the PrimEmo architecture did not clearly out perform Scheutz's emotional robots[31]. Thus, the network presented in this paper still needs to be further refined, before the project can move on to the next stage and make use of the PVLV and PBWM models. For this reason, in the immediate future more experimentation will be done to fully understand the agents' behavioral patterns and pinpoint the parts that need to be improved upon.

## REFERENCES


[1] Brad Aisa, Brian Mingus, and Randy O'Reilly, 'The Emergent neural modeling system', *Neural Networks*, **21**(8), 1146–1152, (2008).
[2] Aristotle, *The Nicomachean Ethics*, Penguin Books, 2004.
[3] Mark F. Bear, Barry W. Connors, and Michael A. Paradisio, *Neuroscience: Exploring the brain*, Lippincott Williams & Wilkins, 2007.
[4] Tim Dalgleish, 'The emotional brain.', *Nature reviews. Neuroscience*, **5**(7), 583–9, (jul 2004).
[5] Antonio Damasio, *Descartes' error: Emotion, reason and the human brain*, 2008.
[6] Charles Darwin, *The expression of the emotions in man and animals*, 1872.
[7] René Descartes, 'Discourse on the Method', *Philosophy*, **1**, xii, 249, (1637).
[8] René Descartes, *The Passions of the Soul*, Hackett Publishing, 1649.
[9] Paul Ekman. Basic emotions, 1999.
[10] JM Fellous and MA Arbib, *Who needs emotions?: The brain meets the robot*, 2005.
[11] A. N. Felten, D. L., and Shetty, *Netter's Atlas of Neuroscience*, 2010.
[12] Stan Franklin, Tamas Madl, Sidney D'Mello, and Javier Snaider, 'LIDA: A Systems-level Architecture for Cognition, Emotion, and Learning', *IEEE Transactions on Autonomous Mental Development*, **6**(1), 19–41, (mar 2014).
[13] T. E. Hazy, M. J. Frank, and R. C. O'Reilly, 'Banishing the homunculus: Making working memory work', *Neuroscience*, **139**(1), 105–118, (2006).
[14] Thomas E. Hazy, Michael J. Frank, and Randall C. O'Reilly, 'Neural mechanisms of acquired phasic dopamine responses in learning', *Neuroscience and Biobehavioral Reviews*, **34**(5), 701–720, (2010).
[15] William James, 'What is an Emotion?', *Mind*, **9**, 188–205, (1884).
[16] David E. Kieras. The EPIC Architecture, 2002.
[17] J A Kiernan and N Rajakumar, *Barr's The human nervous system*, Wolters Kluwer Health/Lippincott Williams & Wilkins, 2014.
[18] John E. Laird, *The Soar Cognitive Architecture*, volume 171, The MIT Press, apr 2012.
[19] J E LeDoux, 'Emotion circuits in the brain.', *Annual review of neuroscience*, **23**, 155–84, (jan 2000).
[20] Joseph E LeDoux, *Synaptic Self: How Our Brains Become Who We Are*, Penguin Publishing Group, 2003.
[21] Joseph E LeDoux. The emotional brain, fear, and the amygdala, 2003.
[22] Joseph E LeDoux, 'Rethinking the emotional brain.', *Neuron*, **73**(4), 653–76, (feb 2012).
[23] Paul D. MacLean, 'The triune brain, emotion and scientific bias', *The neurosciences: Second study program*, 336–349, (1970).
[24] Olivier Michel, 'Webots TM : Professional Mobile Robot Simulation', *International Journal of Advanced Robotic Systems*, **1**(1), 39–42, (2004).
[25] Francesco Mondada, Michael Bonani, Xavier Raemy, James Pugh, Christopher Cianci, Adam Klaptocz, Jean-christophe Zufferey, Dario Floreano, and Alcherio Martinoli. The e-puck , a Robot Designed for Education in Engineering, 2006.
[26] Ivan Petrovich Pavlov, *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*, Oxford University Press: Humphrey Milford, 1927.
[27] Robert Plutchik, *The Emotions: Facts and Theories, and a New Model*, 1962.

[28] Edmund T. Rolls, *Emotion Explained*, 2007.
[29] David Sander, Didier Grandjean, and Klaus R. Scherer, 'A systems approach to appraisal mechanisms in emotion.', *Neural networks : the official journal of the International Neural Network Society*, **18**(4), 317–52, (may 2005).
[30] Clifford B. Saper and Bradford B. Lowell, 'The hypothalamus', *Current Biology*, **24**(23), R1111–R1116, (dec 2014).
[31] Scheutz Matthias, *Useful roles of emotions in artificial agents: A case study from artificial life*, Ph.D. dissertation, Notre Dame, 2004.
[32] S Murray Sherman and R W Guillery, 'The role of the thalamus in the flow of information to the cortex.', *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, **357**(1428), 1695–708, (dec 2002).
[33] S. Murray Sherman and R. W. Guillery, *Exploring the Thalamus and Its Role in Cortical Function*, 2005.
[34] Aaron Sloman, 'Beyond shallow models of emotion', *Cognitive Processing: International Quarterly of . . .* , **2**, 177–198, (2001).
[35] Niels Taatgen, John R Anderson, Christian Lebiere, and Niels Taatgen, 'The Atomic Components of Thought', *Trends in Cognitive Sciences*, **3**(2), 82, (jan 1998).
[36] Brian B Theyel, Daniel A Llano, and S Murray Sherman, 'The corticothalamocortical circuit drives higher-order cortex in the mouse', *Nature reviews. Neuroscience*, **13**(1), 84–89, (2010).
[37] Arthur W. Toga, *Brain mapping : an encyclopedic reference*, volume 1-3, 2015.
[38] Paul Verschure and Armin Duff Eds, *Distributed Adaptive Control: Theory and Practice*, Barcelona, Spain, 2014.
[39] Tor D Wager, Jian Kang, Timothy D Johnson, Thomas E Nichols, Ajay B Satpute, and Lisa Feldman Barrett, 'A Bayesian Model of Category-Specific Emotional Brain Responses.', *PLoS computational biology*, **11**(4), e1004066, (apr 2015).
[40] Paul J. Whalen and Elizabeth A. Phelps, *The Human Amygdala*, Guilford Press, 2009.