

A Cognitive Architecture for Agent-Based Artificial Life Simulation

Ronaldo Vieira^[0000-0002-7109-0897], Bruno Dembogurski^[0000-0002-6719-0537],
Leandro Alvim and Filipe Braidá

Department of Computer Science
Universidade Federal Rural do Rio de Janeiro
Nova Iguaçu, Brazil
ronaldo.vieira@ufrj.br, {brunodembogurski, alvim.lgm, filipebraidá}@gmail.com

Abstract. The ability to simulate living beings that behave in a credible way is a fundamental aspect in digital games. This is due to its interdisciplinary characteristic, that brings together different fields of knowledge to better understand biological life and its processes. In this context, the design of an intelligent agent is a hard task as it involves a complex system, which has several interconnected components. In this work a virtual mind architecture for intelligent agents is proposed, where it simulates the cognitive processes of an actual brain, in this case attention and memory, in order to reproduce behaviors similar to those of actual living beings. A prototype is then proposed, where the architecture is applied on agents that represent virtual animals in a semantic-modeled ecosystem, and conduct a proof-of-concept experiment with it to demonstrate its effectiveness. In this experiment, the behavior of the virtual animals were consistent with reality, thus, validating the architecture's ability to simulate living beings.

Keywords: artificial intelligence, intelligent agents, cognitive architecture

1 Introduction

The simulation of realistic virtual living beings has many applications, ranging from use in non-player characters (NPCs) in games to crowd and traffic simulation. In games, realism is often a crucial factor to the entertainment of the players, whereas in scientific simulations, realism is needed in order for them to be considered a valid simulation of real life.

Ultima Online [18], one of the earliest massively multiplayer online games released, simulated an ecosystem with various species of virtual animals, where they interacted with each other and with the players. *The Elder Scrolls V: Skyrim* [5], in turn, presented NPCs with daily routines that react to unexpected scenarios, such as the sudden presence of a dragon nearby. The simulation of evacuation of people on accident scenarios also needs that the virtual people behave similar to their real counterparts.

The efforts made on this purpose often employ the concept of intelligent agents, from artificial intelligence. In these, the agents carry a virtual mind model that resemble the functioning of an actual brain, providing them with realistic behavior. In this work, the

problem of simulating virtual living beings that behave in a credible way is addressed with the proposal of a virtual mind architecture for intelligent agents.

The proposed architecture has the goal to be generic enough to be usable in a variety of applications and flexible enough to allow problem-specific extensions. It also should be sufficiently complex in order to generate complex behavior while sufficiently simple in order for its use to be viable on digital games or other real time applications, where it is often needed to simulate the behavior of hundred or thousands of characters many times per second.

The main contribution of this paper is to introduce a virtual mind architecture that fulfills the aforementioned requirements while using a simple high-level abstraction of the reasoning workflow of a brain that is both computationally cheap and easy to understand and predict. Also in this paper, the proposed architecture is applied in the context of simulating a small ecosystem, where its operation is exemplified.

In Sect. 2, the concept of intelligent agents is presented. Next, in Sect. 3, concepts about the cognitive processes of the brain that are essential to the understanding of the proposed architecture are explained. After that, the proposed architecture, its components and their interaction are described in Sect. 4. A proof-of-concept prototype that uses the architecture is then depicted on Sect. 5, and the experiment is made in Sect. 6. Conclusions and future works are discussed on Sect. 7.

2 Intelligent Agents

In a generic way, an agent can be defined as an autonomous entity that perceives its environment through sensors and act on it through actuators [23]. With this definition, many things can be considered agents: a human being perceives its environment through its sight, hearing, smell, taste and other sensors, while act on it with its hands, feet and voice; a smartphone perceives its environment with the aid of its proximity, luminosity and touch sensors and act through its speakers and screens. However, the concept of agents is explored in depth on the Artificial Intelligence field, and is often used as means to develop complex systems, simulations and games.

Besides sensors and actuators, this definition of agents also have a program, that determines how the agent select its actions based on the inputs captured by its senses. An agent is said to be intelligent when its program can maximize its performance on the tasks it was developed for through its perception of the environment. To demonstrate how the definition of an agent can vary, a reading of the work by Franklin and Graesser [10] is suggested, where a taxonomy for autonomous agents is presented and, also, the many agent definitions used in the literature.

The term autonomous agent has been an intense discussion topic since the early 90's [16], due to its ability to act in complex and changing multiagent environments. The main core of an autonomous agent is its architecture, which comprehends the description of its features (modules and parts) and how they interact with each other. Thus, since its first appearance, numerous architectures have been proposed in order to create better and smarter agents, each presenting different features or expanding the previous ones. For a thorough explanation of the research field and its main threads refer to [16].

Several works make use of intelligent agents to simulate the behavior of human beings or other living beings. [1] and [9] use them in the context of simulating characters in a narrative, while [3] apply them to build realistic bots for the game *Unreal Tournament 2004*, and [17] applies them in the medical context, as virtual patients and clinical advisors. [4] creates a agent system in order to explain the cognitive processes of a human brain. As can be noted, intelligent agents can be applied to solve a variety of different problems.

The next section explain some of the cognitive processes of human brains on which the proposed mind architecture is based.

3 Cognitive Processes

Cognition can be defined as “the mental action or process of acquiring knowledge and understanding through thought, experience, and the senses” [20]. It encompasses processes such as knowledge, attention, memory, reasoning and problem solving. In this section, concepts of the human memory and attention that are relevant to the proposed architecture are discussed.

3.1 Working memory

On neuroscientific literature, three types of memory are often accepted to exist: long-term memory, short-term memory and working memory [7]. When new memories are created, they are short-term ones. If any of them is deemed important by the brain (by either conscious or subconscious decision), it gets stored as a long-term memory.

The working memory can be considered a temporary memory where information is kept for immediate use by the brain’s cognitive processes. Its existence allows the brain to manipulate information: when comparing two objects, for instance, it is necessary that a representation of both of them be stored in the working memory so that the comparison can be made.

Some cognitive architectures [1] and [4] make use of both long-term and short-term storing. They are out of the scope of this work. Their future addition, however, is noted in Sect. 7.2.

3.2 Attention

Though the human field of vision seem to be wide, our eyes can only perceive with precision, in fact, within a few centimeters radius from the point where we are focusing. The human hearing often has gaps that are filled with what our brain sees fit [24]. However, even with those and other known limitations of our sensory systems, the brain is constantly bombarded with information from our surroundings. It was necessary during evolution the development of a mechanism to decide which information was relevant and which was not. That mechanism is called attentional filter [12].

The attentional filter’s job is to allow only the most relevant information to enter consciousness (and, consequently, the working memory). After driving for a while, for example, the attentional filter will eventually prevent the engine noise from entering

consciousness. However, if the engine starts producing an unfamiliar noise – possibly due to some mechanism needing attention –, this information is promptly brought to consciousness. That is the first principle used by the attentional filter: change. The other principle, importance, can be exemplified by when someone in a full bus starts paying attention to a specific conversation once something important, such as their own name, was cited.

4 Proposed Architecture

This section explains in details the proposed architecture, each of its components and the interaction between them.

There are several mind architectures in the literature, AuRA (Autonomous Robot Architecture) [2], The Soar Cognitive Architecture [11], IRMA [6], BDI (Belief-Desire-Intention) architecture model [22], FAtiMA [8] and ORIENT [13], just to name a few. The latter is built on top of FAtiMA (which uses the OCC - cognitive theory of emotions [19]) plus motivational and learning components from the PSI theory [21]. In this paper, however, the agents behaviors are derived from a much simpler workflow, where a motivation system is abstracted as their needs and, combined with their perception of the environment, results in dynamic realistic agents whose behavior can still be directed/predicted.

The virtual mind model proposed here simulates in a high level of abstraction both the architecture and cognitive flow of an actual brain, in order to allow intelligent agents carrying this artificial mind, once disposed in an semantic-modeled environment, to display behaviors that are similar to those of actual living beings.

Despite all the technology and scientific knowledge available nowadays, it is still impossible to perfectly simulate a working brain. Thus, this work's objective is not to try such feat, but to find a complexity level for the architecture that is both high enough for it to be able to generate credible behavior and low enough to make its use on real time applications – such as games – viable.

An abstraction of the senses, needs and goals of a living being, along with some cognitive processes (perception, memory, attention and reasoning) and a semantic-modeled environment or world are the components of the proposed virtual mind architecture.

4.1 Semantic representation

For the virtual mind be able to reason, it needs information about the surroundings of the agent. Thus, the architecture requires a semantic-modeled world, that is, an environment where each perceivable entity has semantic attributes that represent its meaning. For instance, an object can have attributes for its size, color, what it is made of, whether it is alive, *etc.* The agents are not an exception – as any other entity in the world, they also must have these semantic information.

The semantic information of an entity can be represented with any data structure able to hold key-value entries. However, a frame system [15] is recommended, since it allows for complex semantic modeling with inheritance and commonsense reasoning [25].

4.2 The agent

Different from the static objects in the world, the agents can move and interact with near objects and agents. That is possible because they have senses which perceive their surroundings, needs that guide their goals and actions, and a mind where the actual reasoning happens. In the following sections, each of these components will be described in terms of its role in the architecture and its properties.

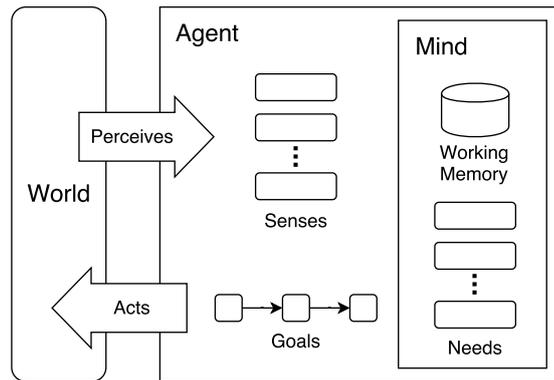


Fig. 1. Architecture of an agent.

4.3 Senses and perceptions

The role of the human senses, such as sight and hearing, is to obtain information from its environment and make them available to the brain so that it has sufficient inputs to decide which actions to take. Analogously, in this architecture, in order for an agent to interact with the world, it should have one or more senses. Here, a sense's job is to capture semantic information from the entities in the world that are within its range.

In a Sight sense, for instance, a field-of-view algorithm may be used as its range in order to simulate the sight of a living being. It may be able to perceive attributes like size, color and distance. These captured semantic information are made available to the agent's mind in the form of perceptions.

A perception is, basically, a subset of semantic information of an object or agent extracted by a sense, and can be represented with the same data structure. At the moment of its creation, some non-static attributes can also be filled by the sense, such as the distance attribute, which can only be determined at perception creation time.

4.4 Goals

In the other end of the agent's interaction with the world are the goals. While the senses are responsible for reading the current world state, the goals' role is to alter them. The

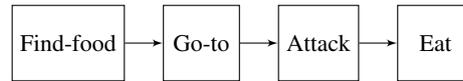


Fig. 2. A goal chain. The arrows denote prerequisite relationships.

agent is at all times executing a goal, and the choice of the most suitable goal to execute is byproduct of the reasoning cycle of the agent's mind. The goals of an agent in this work are directly mapped to actions in the environment, unlike [1] and [9], where goals and actions are hierarchically structured.

A goal should have the following properties:

- A set of **completion conditions**. A goal is considered as completed if and only if each of its completion conditions are satisfied. A Go-to goal, for instance, may have the unique completion condition of the distance between the agent and its target being less than a specified threshold.
- Optionally, a **previous goal** (or prerequisite goal). The previous goal should be necessarily completed in order for the current goal to be executed. If during the execution of an goal its prerequisite ceases to be completed, the agent will execute its prerequisite instead until it is completed again.
- Optionally, a **next goal**. The next goal is the next goal in the goal chain, that is, the goal that will be executed after the completion of the current one.
- Optionally, a **target**. It is represented by a perception that represents the target of the goal. An Attack goal, for instance, may require a target which will be attacked. Not all goals will require a target, though. It may be the case for a Sleep goal, for example. The goal's target is inherited from previous goals to next goals in a goal chain or manually defined at its creation. In figure 2, the Find-food goal may not require a target beforehand, but may set its target as any potential food source once it has found one so that the next goal – Go-to – will then inherit it as its target.
- A **execution cycle**. The execution cycle of a goal is an algorithm that is executed on each reasoning cycle of the agent. That algorithm will, often, perform actions on the goal's target, in the agent itself or alter the world's state. The execution cycle of a Eat goal, for instance, may have instructions to decrease the intensity of the agent's Hunger need and also reduce the weight attribute of the target entity.

The structure formed by goals and their next and previous goals is named goal chain. In the subsequent sections, the term goal will be used to refer to both single goals and goal chains. The management of a goal chain, that is, the choice of the next goal once the current is completed and the choice of the previous goal once it ceases to be completed is one of the steps of the agent's reasoning cycle.

4.5 Needs

According to Maslow [14], the human being has various needs to be fulfilled simultaneously, such as breathing, temperature, feeding, safety, health, shelter, being part of a group or community and happiness. Some of them have greater priorities than others,

forming a hierarchy of needs. A human being will not give much thought to its need to be respected if, say, their need to breathe is not being at least reasonably satisfied.

The needs of an agent are the central aspect in their decision-making process. Each need is bound to a goal that will satisfy it. The current goal of an agent in a given moment will always be the goal bound to the need with the most priority among its needs at that moment. The calculation of the priority of a need will be explained in detail on Sect. 4.7.

The need to feed or sleep, for example, are derived from a number of biological processes of an animal's body. However, their intensity increase over time can be represented in a simpler way without a significant loss of realism through a increase rate. That way, the needs should have, besides their bound goal and a intensity value, a increase rate. This rate define how quickly the intensity value of a need will increase over time.

Some needs, such as safety and shelter, may not possess the behavior of increasing their intensity over time. In that case, they can be assigned with a zero increase rate, causing their intensity value to be fixed. Furthermore, it is convenient to set their initial intensity value to a reasonable value.

At last, a need should be able to evaluate how relevant is a given perception. The relevancy evaluation of perceptions can be implemented as either a set of logic rules such as "if the perception source is a bigger animal, then return 1", a set of fuzzy rules such as "if the size attribute from the perception ranges between 1 and 2, then return the interpolated value between 0 and 1" or any other algorithm or equation.

The aforementioned needs hierarchy proposed by Maslow [14] raises a diverse list of needs of human beings that are usable in this architecture. However, it is important to note that any simulation aspect that fits in the description of a need can be represented as one.

In summary, needs should have the following properties:

- An **intensity value**. A value in the range $[0, 1]$ that represents how intense is the need in a given moment. A hungry agent, for instance, would have a value near 1 as the intensity of their Hunger need.
- An **increase rate**. A rate that defines how quickly the intensity of the need will increase. A need with an increase rate of 0.2 per minute would take exactly five minutes to go from fully satisfied (intensity 0) to critically neglected (intensity 1).
- A **bound goal**. A goal or goal chain that will be executed by the agent whenever the need is chosen as the one with the most priority by the reasoning cycle.
- A **relevancy evaluator**. An algorithm that, given a perception, returns its relevancy score to the need in the $[-1, 1]$ interval. Positive scores amplify the need's intensity, while negative scores inhibit it.

4.6 The mind

The mind of an agent translate perceptions obtained by its senses in goals, in order to satisfy its needs. It consists of a working memory and the reasoning cycle algorithm.

The working memory stores a limited set of perceptions and serves as a temporary memory, representing the current thoughts of the agent. It has a capacity limit that defines how many perceptions it is able to store simultaneously. In [1], the short-term

memory is in charge of this aspect. [3] also has a working memory module, where a number of different sensory inputs compete for attention. In this architecture, however, the perceptions in the working memory already represent the agent's attention.

The agent's attentional filter, in the reasoning cycle, guarantees that only the most relevant perceptions will enter the working memory. The reasoning cycle is the core of the architecture. It represents the agent's process of thought and decision-making through its limited perception of the world.

4.7 Reasoning Cycle

The first step in an agent's reasoning cycle is to obtain a new set of perceptions of the world through the agent's senses. Each sense is queried for its current set of perceptions, then all perceptions that share the same source are merged. The agent's working memory can already be filled with perceptions from earlier reasoning cycles, so those are also considered for the merging. At this point, each entity in the world is represented in the agent's mind by at most one perception.

Then, each need n of the agent will assign a relevancy score $r(p, n) \in [-1, 1]$ to each perception p using its relevancy evaluator. That way, each perception will be assigned a number of different relevancy scores, one for each of the agent's needs.

The working memory is then updated to carry only the k most overall relevant perceptions, where k is its capacity limit. The overall relevancy of a perception ($r(p)$) is just the sum of the absolute values of all the relevancy scores assigned to it, as shown by equation 1. A perception that got 0.8 and -0.5 as relevancy scores by the two needs of the agent, for instance, will have a overall relevancy score of 1.3. This step works as an attentional filter to the agent's mind, ensuring that only the most relevant entities perceived by its senses will be considered.

$$r(p) = \sum_n |r(p, n)| \quad (1)$$

With the working memory updated, the agent is ready to decide which of its needs demands more attention at the moment. That decision is made using equation 2, by choosing the need n that presents the highest priority value $p(n)$, given as a function of its intensity $i(n)$ and relevancy $r(n)$ values.

$$p(n) = i(n) * (1 + r(n)) \quad (2)$$

The relevancy of a need n , as shown in equation 3, is given as a function of each of the relevancy scores $r(p, n)$ assigned by it to the perceptions currently in the agent's working memory W . In this case, however, each relevancy score is penalized by a factor of the distance of the perception to the agent.

$$r(n) = \sum_{p \in W} \frac{r(p, n)}{1 + \ln(1 + d(p))} \quad (3)$$

It can be noted on equation 2 that when an agent does not have any perception stored in its working memory, $r(n)$ will be null, and the priority of the needs will be based

solely on its intensity. On the other hand, when two needs have the same intensity value, their priority values will vary based only on the current perceptions of the agent. A virtual human being whose Hunger and Thirst needs have the same intensity at a given moment, for example, would easily choose to satisfy their hunger should an apple be near them. Besides, it is guaranteed that any need with a non-null increase rate, if neglected long enough, will eventually have a high enough intensity value, causing it to be chosen by the agent.

In equation 3, in turn, can be noted that the primary factor to determine the relevancy of a need is how relevant to this need are the current perceptions in the agent's working memory. For a virtual animal, for example, a perception that represents a threatening animal would hold a positive relevancy score to its Safety need and, thereafter, would have a positive impact on the relevance of that need. Still, if many members of its own specie were near, generating negative relevancy scores and impacting negatively the Safety's relevance, the threatening animal would then not look so threatening.

The relevance of the perceptions are penalized proportionally to the natural logarithm of the distance of those perceptions to the agent. That way, a virtual human being with a Hunger need would prefer a near apple than an apple that is 50 meters away. The natural logarithm is used in order to reduce the penalizing impact of the distance – an apple that is 50 meters away is still reasonably more relevant to a Hunger need than no apple at all.

Once the need with the most priority is identified, if its bound goal is already being executed by the agent, then nothing is done. Otherwise, the agent's current goal is replaced with that need's bound goal.

The last step on the reasoning cycle consists in executing the execution cycle of the current goal. The goal management also happens: if the current goal is completed or if the previous one is not, the appropriate arrangements are made.

4.8 Extensions

Although the proposed architecture is based on high-level abstractions of cognitive processes of human beings and its design is focused to use on digital games, it can be adapted to other contexts, such as (i) transit simulations, with agents as vehicles in a road or city; (ii) crowd simulations, with agents as people; (iii) narratives, as in [1], with agents as characters in a story; (iv) simulation of ecosystems, as described in Sect. 5; and even (v) population-based metaheuristics, to solve combinatorial optimization problems.

5 Prototype

This section has the objective to present one of the possible applications of the architecture, provide a concrete example and allow for experiments to demonstrate its operation.

The chosen context for this prototype is the simulation of a small ecosystem, with the animals as agents and other static entities. For this ecosystem, two species of animals were chosen to represent respectively the prey and predator roles: rabbit and fox. Also, a grass entity was created to represent the prey's food and a rock entity to represent a object that is not important to any of the agents. The objective of this prototype is to use

the proposed virtual mind architecture in intelligent agents simulating virtual animals. For that, their senses, needs and goals were chosen accordingly to those of real animals in a ecosystem.

In the following subsections, the semantic model used, as well as the chosen senses, needs and goals are described.

5.1 Semantic model

A implementation of Minsky's frames [15] was used as the data structure to represent the semantic information of the prototype's world.

Each entity in the world (fox, rabbit, grass and rock) possess a instance frame containing its semantic attributes. That instance frame, in turn, inherits attributes from a generic frame in the semantic hierarchy. This way it is possible to infer that, for example, a rabbit has four legs, since it is a mammal, and the mammal generic frame it inherits from possess a "legs" attribute set to four.

5.2 Senses, needs and goals

The Sight and Hearing senses were created to represent – unsurprisingly – the sight and hearing of the virtual animals, respectively. Their perception ranges are depicted graphically by figure 3. The parameters and initial values of these senses as well as those of the needs and goals were chosen empirically, with values that provided a satisfactory level of realism while working in a similar way to the senses of the corresponding animals in nature.

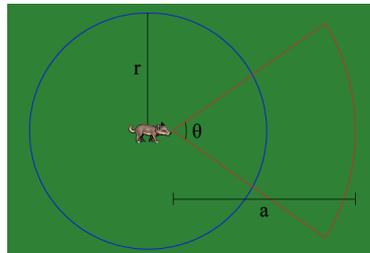


Fig. 3. Graphic representation of the range of the agent's sight (red) and hearing (blue).

Two needs were created for the agents: Hunger and Safety. They represent two of main aspects of the survival in a ecosystem: feeding and avoiding potential threats. Their bound goals were defined as the goal chains depicted by figure 4, respectively. The detailed description of each need and goal is omitted, but can be assumed to act abstractly as they would in real life.

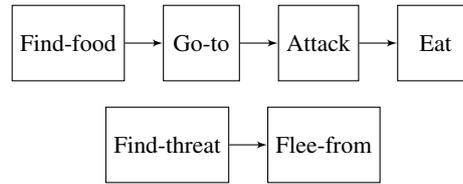


Fig. 4. Goal chains bound to the Hunger (on top) and Safety (on bottom) needs.

6 Experiment

The proposed experiment involves placing four entities reasonably close to each other in a world: a fox, a rabbit, a grass and a rock. It is expected, as the virtual mind architecture described on Sect. 4 is used with the accordingly developed components described on Sect. 5, that the agents display behaviors similar to reality. As the agent's reasoning cycle is executed dozens of times per second in the simulation, a single moment was chosen to present and explain the operation of the architecture and justify some of the behaviors shown by the agents.



Fig. 5. Initial setting and some key moments of the experiment.

The chosen moment happens right at the start of the simulation, when the rabbit will decide to flee from the fox. The following paragraphs will demonstrate their reasoning cycles in that moment of the simulation.

Initially, all entities are close to each other. As the first execution of the agent's reasoning cycle, their working memories will be first filled with perceptions from the world and their first goals will be defined.

The fox's reasoning cycle starts with the retrieval of its current perceptions from its senses. In figure 6 the perceptions retrieved from the Sight and Hearing senses are described. Though the agent does not know, it can be noted that these perceptions come, respectively, from the rabbit, grass and rock.

The rabbit's perception was assigned with a 1.0 relevancy score from the fox's Hunger need, since it fits its criteria for potential food sources. The fox, rabbit, grass and rock's perceptions will be referred to as, respectively, p_f , p_r , p_g e p_o .

The working memory of the agents was defined to hold up to four perceptions simultaneously. This way, all the three obtained perceptions will be included in the

Perception 1	Perception 2	Perception 3
(other attributes) distance: 4.896 r(Hunger): 1.0 r(Safety): 0.0	(other attributes) distance: 6.326 r(Hunger): 0.0 r(Safety): 0.0	(other attributes) distance: 10.191 r(Hunger): 0.0 r(Safety): 0.0

Fig. 6. Perceptions retrieved by the fox's senses. Each box represents a perception, where their semantic attributes and relevancy scores are listed.

fox's working memory. Should its capacity limit be reduced to one perception, only the rabbit's perception would be stored as its overall relevance value equals 1.0, while all the others have 0.0.

In this moment, the intensities of the Hunger and Safety needs are, respectively, 0.15 and 0.3. Their relevance is calculated using the perceptions currently stored in the working memory.

$$\begin{aligned}
r(\text{Hunger}) &= \frac{r(p_r, \text{Hunger})}{1 + \ln(d(p_r))} + \frac{r(p_g, \text{Hunger})}{1 + \ln(d(p_g))} + \frac{r(p_o, \text{Hunger})}{1 + \ln(d(p_o))} \\
&= \frac{1}{1 + \ln(5.896)} + \frac{0}{1 + \ln(7.326)} + \frac{0}{1 + \ln(11.191)} \\
&= 0.36
\end{aligned} \tag{4a}$$

$$\begin{aligned}
r(\text{Safety}) &= \frac{r(p_r, \text{Safety})}{1 + \ln(d(p_r))} + \frac{r(p_g, \text{Safety})}{1 + \ln(d(p_g))} + \frac{r(p_o, \text{Safety})}{1 + \ln(d(p_o))} \\
&= \frac{0}{1 + \ln(5.896)} + \frac{0}{1 + \ln(7.326)} + \frac{0}{1 + \ln(11.191)} \\
&= 0
\end{aligned} \tag{4b}$$

As shown by equations 4a and 4b, the relevancy scores respective to the Hunger and Safety needs are 0.36 and 0. With these values, is then possible to calculate the priority score of them. The equations 5a and 5b depict the necessary calculations.

$$\begin{aligned}
p(\text{Hunger}) &= i(\text{Hunger}) * (1 + r(\text{Hunger})) \\
&= 0.15 * (1 + 0.36) \\
&= 0.204
\end{aligned} \tag{5a}$$

$$\begin{aligned}
p(\text{Safety}) &= i(\text{Safety}) * (1 + r(\text{Safety})) \\
&= 0.3 * (1 + 0) \\
&= 0.3
\end{aligned} \tag{5b}$$

This way, it can be noted that the Safety's priority score (0.3) is higher than Hunger's (0.204), thus, Safety is the chosen need. Once having chosen a need, the agent's current goal will become the first goal of the goal chain bound to that need. In this case, the

bound goal chain is composed by the Find-threat and Flee-from goals, and the former will become the fox's current goal.

The goal management is the last step of the fox's reasoning cycle. The mind will check that the current goal's completion conditions are not satisfied and no previous goal is present, and will thus execute it. As part of the goal's execution cycle, it will check the perceptions currently on the agent's working memory for any potential threats. However, no such perception will be found. The fox will then keep searching for threats in the next reasoning cycles, until its Hunger (which intensity value increases with time) become the need with the highest priority score. And so ends the fox's reasoning cycle.

The rabbit's reasoning cycle starts retrieving its current perceptions from its Sight and Hearing senses. Figure 7 describes the perceptions obtained by them that refers to, respectively, the fox, grass and rock. It can be noted that the perception relative to the grass and fox received 1.0 as relevancy score from the Hunger and Safety needs, respectively.

Perception 1	Perception 2	Perception 3
(other attributes) distance: 4.896 r(Hunger): 0.0 r(Safety): 1.0	(other attributes) distance: 5.046 r(Hunger): 1.0 r(Safety): 0.0	(other attributes) distance: 6.352 r(Hunger): 0.0 r(Safety): 0.0

Fig. 7. Perceptions retrieved by the rabbit's senses. Each box represents a perception, where their semantic attributes and relevancy scores are listed.

The rabbit's attentional filter also permits all three perceptions to enter its working memory. It is then calculated which of its needs will be satisfied, following the same process as made by the fox.

$$\begin{aligned}
 r(\text{Hunger}) &= \frac{r(p_f, \text{Hunger})}{1 + \ln(d(p_f))} + \frac{r(p_g, \text{Hunger})}{1 + \ln(d(p_g))} + \frac{r(p_o, \text{Hunger})}{1 + \ln(d(p_o))} \\
 &= \frac{0}{1 + \ln(5.896)} + \frac{1}{1 + \ln(6.046)} + \frac{0}{1 + \ln(7.352)} \\
 &= 0.357
 \end{aligned} \tag{6a}$$

$$\begin{aligned}
 r(\text{Safety}) &= \frac{r(p_f, \text{Safety})}{1 + \ln(d(p_f))} + \frac{r(p_g, \text{Safety})}{1 + \ln(d(p_g))} + \frac{r(p_o, \text{Safety})}{1 + \ln(d(p_o))} \\
 &= \frac{1}{1 + \ln(5.896)} + \frac{0}{1 + \ln(6.046)} + \frac{0}{1 + \ln(7.352)} \\
 &= 0.36
 \end{aligned} \tag{6b}$$

As shown in the equations 6a and 6b, the relevancy scores respective to the Hunger and Safety needs are 0.357 e 0.36. Is then possible to calculate their priority scores,

which is made by equations 7a and 7b using the intensity values (0.15 for Hunger and 0.3 for Safety) and the obtained relevancy scores.

$$\begin{aligned} p(\text{Hunger}) &= i(\text{Hunger}) * (1 + r(\text{Hunger})) \\ &= 0.15 * (1 + 0.357) \\ &= 0.204 \end{aligned} \tag{7a}$$

$$\begin{aligned} p(\text{Safety}) &= i(\text{Safety}) * (1 + r(\text{Safety})) \\ &= 0.3 * (1 + 0.36) \\ &= 0.408 \end{aligned} \tag{7b}$$

With the obtained results, the rabbit will choose its Safety need (0.408) to be satisfied rather than its Hunger (0.204). For that reason, the goal chain formed by the Find-threat and Flee-from goals will be chosen as its current goal.

The rabbit's goal management step proceeds similarly to that of the fox. However, in the execution cycle of the Find-threat goal, a potential threat perception is found. That perception – that refers to the fox – is set as the target of the current goal so that, in the next reasoning cycle, the mind will check that the goal's completion conditions were satisfied and will move on to the next goal in the chain: Flee-from.

In the subsequent reasoning cycles, the Flee-from goal will move the rabbit away from the fox. Even though it possesses a wide field of view, by moving in the opposite direction, the rabbit will not be able to detect the fox with its Sight sense. However, it will still be able to know its location by the Hearing sense, and will keep running away until it reaches a sufficient distance.

After the chosen moment, the rabbit will run in the opposite direction of the fox until it feels safe enough: the distance from the fox will trigger the completion of its Flee-from goal. Eventually, the fox's Hunger need will become sufficiently intense, causing it to be chosen by the reasoning cycle. The fox, then, will chase the rabbit until it is close enough to attack. After feeding itself, the fox will have satiated its hunger and will be back to exploring the world.

As depicted in the chosen moment and in the rest of the simulation, the architecture was able to choose dynamically an appropriate behavior for the virtual animals in all situations through their reasoning cycles. The rabbit, having a threat and a food source at sight, decided to take care of its safety and flee. The fox, on the other hand, ignored the presence of the rabbit until it was hungry enough, and then chased it.

7 Conclusion

Simulating virtual living beings that behave in a realistic way is not a trivial problem. In order to perfectly represent a living being, it would be necessary to simulate the low level interaction between the chemical elements in its body, which is impossible by the currently available technology and scientific knowledge. That being said, the current approaches to this problem in the literature often make use of models with higher levels of abstraction where virtual minds emulate, in a simplified way, the functioning of real brains.

7.1 Proposal

In this work a cognitive architecture for intelligent agents was proposed, that is, an architecture that simulates the functioning of a brain and its cognitive processes, whose objective is to reproduce on agents behaviors similar to those of living beings, once they are disposed in a semantic-modeled virtual world.

The proposed architecture translates the concept of sensors of intelligent agents as living beings' senses. The job of a sense is to capture information about the surroundings of the agent in order to foment its reasoning. These information are passed to the agent's mind in the form of perceptions. Once inside the mind, its attentional filter ensures that only the most relevant perceptions enters the agent's working memory, and thus, be considered in its reasoning process. The mind then decides which of the agent's needs demands most attention taking into account its current intensity and the content of the working memory. The goal chain associated to that need is then executed by the agent. The goals encompass the concept of actuators of intelligent agents. This way, the agent perform actions that are consistent with its internal state and its perception of the external world.

A prototype that simulates virtual animals in a small ecosystem was developed in order to validate the architecture. In the experiment made, the two agents acted accordingly to what would be expected of actual animals pursuing their survival and, therefore, demonstrated the proposed architecture's ability to simulate living beings.

7.2 Research directions

As future improvements to the architecture, would be interesting to have (i) grouping of similar perceptions into one collective perception; (ii) occasional errors on the senses' detection; (iii) goals uncoupled from the needs; (iv) satisfaction of multiple needs simultaneously; and (v) a short-term and a long-term memory.

As for the developed prototype, some interesting features are (i) non-binary relevancy scores; (ii) implement reproduction as a need; (iii) learning by experience; and (iv) allow evolution and natural selection.

References

1. Araujo, S., Chaimowicz, L.: Atores Virtuais Autônomos para Sistemas Narrativos Interativos. Master's thesis, Universidade Federal de Minas Gerais (2009)
2. Arkin, R.C.: Integrating behavioral, perceptual, and world knowledge in reactive navigation. *Robotics and Autonomous Systems* 6(1), 105 – 122 (1990), <http://www.sciencedirect.com/science/article/pii/S0921889005800314>, designing Autonomous Agents
3. Arrabales, R., Ledezma, A., Sanchis, A.: Towards conscious-like behavior in computer game characters. In: *Computational Intelligence and Games, 2009. CIG 2009. IEEE Symposium on*. pp. 217–224. IEEE (2009)
4. Bach, J.: *Principles of synthetic intelligence PSI: an architecture of motivated cognition*, vol. 4. Oxford University Press (2009)
5. Bethesda Game Studios: *The Elder Scrolls V: Skyrim*. [PC CD-ROM] (2015)

6. Bratman, M.E., Israel, D.J., Pollack, M.E.: Plans and resource-bounded practical reasoning. Tech. Rep. 425, AI Center, SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025 (Sep 1988)
7. Cowan, N.: What are the differences between long-term, short-term, and working memory? *Progress in brain research* 169, 323–338 (2008)
8. Dias, J., Mascarenhas, S., Paiva, A.: FATiMA Modular: Towards an Agent Architecture with a Generic Appraisal Framework, pp. 44–56. Springer International Publishing, Cham (2014), https://doi.org/10.1007/978-3-319-12973-0_3
9. Franco, A.O., Maia, J.G., Neto, J.A., Gomes, F.A.: An interactive storytelling model for non-player characters on electronic rpgs. In: *Computer Games and Digital Entertainment (SBGames)*, 2015 14th Brazilian Symposium on. pp. 52–60. IEEE (2015)
10. Franklin, S., Graesser, A.: Is it an agent, or just a program?: A taxonomy for autonomous agents. In: *Proceedings of the Workshop on Intelligent Agents III, Agent Theories, Architectures, and Languages*. pp. 21–35. ECAI '96, Springer-Verlag, London, UK, UK (1997), <http://dl.acm.org/citation.cfm?id=648203.749270>
11. Laird, J.E.: *The Soar Cognitive Architecture*. The MIT Press (2012)
12. Levitin, D.J.: *The organized mind: Thinking straight in the age of information overload*. Penguin (2014)
13. Lim, M.Y., Dias, J., Aylett, R., Paiva, A.: Creating adaptive affective autonomous npcs. *Autonomous Agents and Multi-Agent Systems* 24(2), 287–311 (Mar 2012), <https://doi.org/10.1007/s10458-010-9161-2>
14. Maslow, A.H.: A theory of human motivation. *Psychological review* 50(4), 370 (1943)
15. Minsky, M.: A framework for representing knowledge. *The psychology of computer vision* 73, 211–277 (1975)
16. MÜLLER, J.P.: Architectures and applications of intelligent agents: A survey. *The Knowledge Engineering Review* 13(4), 353–380 (1999)
17. Nirenburg, S., McShane, M., Beale, S., Catizone, R.: A cognitive architecture for simulating bodies and minds. In: *AMIA Annual Symposium Proceedings*. vol. 2011, p. 905. American Medical Informatics Association (2011)
18. Origin Systems: *Ultima online*. [CD-ROM] (1997)
19. Ortony, A., Clore, G.L., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge University Press (1990)
20. Oxford Dictionaries: cognition - definition of cognition in english | oxford dictionaries. Available at <<https://en.oxforddictionaries.com/definition/cognition>> (2010), accessed on: 07/07/2017
21. Prisyakov, V.F., Prisyakova, L.M.: Mathematical modeling of emotions. *Cybernetics and Systems Analysis* 30(1), 142–149 (Jan 1994), <https://doi.org/10.1007/BF02366374>
22. Rao, A.S., Georgeff, M.P.: *Modeling rational agents within a bdi-architecture* (1991)
23. Russell, S.J., Norvig, P.: *Artificial Intelligence: A Modern Approach*. Pearson Education, 2 edn. (2003)
24. Warren, R.M., et al.: Perceptual restoration of missing speech sounds. *Science* 167(3917), 392–393 (1970)
25. Winston, P.H.: *Frames and commonsense*. In: *Artificial intelligence*, 3rd (repr. with corrections 1993) ed, chap. 10, pp. 209–230. Addison-Wesley (1993)