Polytechnics University of Bucharest
Faculty of Computers and Automation
Department of Computer Science

A Framework for Developing Embodied Intelligent Agents with Affective Behavior

Research Report #1

Affective Computing for Artificial Agents

Author:
Valentin Lungu, B.Sc.

Scientific adviser:
Adina Magda Florea, Ph.D.

Bucharest
September 2011

Abstract

This paper proposes an architecture based on rule-based systems that uses a forward chaining and a backward chaining inference engine, a truth maintenance system and emotion simulation to achieve reasoning, fast decision-making intelligent artificial characters.

An agent needs to be able to accomplish its goals. Hierarchical goal decomposition is a powerful tool, allowing the agent to represent and solve complex problems. A backward chaining inference engine is best at breaking down goals into sub goals.

Agents in a dynamic environment where multiple aspects of the world are currently changing must be able to infer new knowledge about the world. Also, said agents should also be able to act in uncertain conditions (conditions of uncertain knowledge). A forward chaining inference engine is used to infer knowledge about the world that is not strictly goal-related, and a truth maintenance system is used to handle conflicting knowledge and maintain a consistent set of beliefs about the world.

Emotion integration is necessary in generating complex believable behavior, making the agent decision-making process less predictable and more realistic as well as generating actions in time comparable to human reaction time.

The study of what is now called *emotional intelligence* has revealed yet another aspect of human intelligence. Emotions were shown to have great influence on many of human activities, including decision-making, planning, communication, and behavior.

This information also provides new ideas to researchers in the fields of affective computing and artificial life about how emotion simulation can be used in order to improve artificial agent behavior.

Emotion theories attribute several effects of emotions on cognitive, decision and team-working capabilities (such as directing/focusing attention, motivation, perception, behavior towards objects, events and other agents.

This paper attempts to analyze some important aspects of emotions of human emotions and model them in order to create more believable artificial characters and more effective agents and multi-agent systems. This paper mainly deals with the aspect of memory and emotional influences on memory, perception, and reasoning.

We study various emotion theories, emotion computational simulation models, and, drawing from their wisdom, we propose a model of artificial

agents, which attempts to simulate the effects of human-like emotions on the cognitive and reactive abilities of artificial agents.

This paper also intends to show that emotion simulation does not only serve to improve human-computer interaction and behavior analysis, but may also be used in order to improve artificial agent and multi-agent system performance and effectiveness.

We describe an artificial agent emotion-driven architecture that not only attempts to provide complex believable behavior and representation for virtual characters, but that attempts to improve agent performance and effectiveness by mimicking human emotion mechanics such as motivation, attention narrowing and the effects of emotion on memory.

The work has been funded by the Sectoral Operational Programme Human Resources Development 2007-2013 of the Romanian Ministry of Labour, Family and Social Protection through the Financial Agreement POSDRU/88/1.5/S/61178.

Contents

List of Figures

Chapter 1

Introduction

Autonomous artificial intelligent entities are required to be able to exist in a complex environment and respond to relevant changes in the environment, deliberate about the selection and application of operators and be able to pursue and accomplish goals. This paper proposes a rule-based system that involves both forward and backward chaining and a rule maintenance system in order to provide an architecture for the development of fast decision-making agents. In the following sections, the features of this framework, the reasons for which they are needed and the way they interface with each other will be elaborated. Last, but not least, we discuss the value of integrating emotions with our agents, and a possible way to do so.

Several beneficial effects on human cognition, decision-making and social interaction have been attributed to emotions, while several studies and projects in the field of artificial intelligence have shown the positive effects of artificial emotion simulation in human-computer interaction and simulation applications. This paper presents various emotion theories and artificial emotion simulation models, and proposes a model of emotions that tries to reproduce the human-like effects on cognitive structures, processes, patterns and functions in artificial agents through emotion simulation and emotion based memory management, leading thus to more believable agents, while also improving agent performance. The paper is structured as follows:

Chapter 3 overviews some significant human emotion theories (from the standpoint of artificial intelligence) and discusses how they may be used to enhance virtual characters and multi-agent systems.

Chapter 2.1 discusses the effects of emotions on human memory that may be replicated/simulated in order to improve agent memory storage, pattern-matching/lookup speed and attention focus/narrowing which may help in designing resource constrained agents.

Chapter 2.2 presents related work in affective computing, namely computational emotional models and their main features.

Chapter 3 covers our proposed model for emotion representation and simulation and the ways it influences/interacts with knowledge, cognition, perception and the decision process, while chapter 4 presents the framework architecture.

Last, but not least, chapter 5 presents conclusions and future work.

Autonomous artificial entities are required to be able to exist in a complex environment and respond to relevant changes in the environment, reason about the selection and application of operators and actions and be able to pursue and accomplish goals. This type of agent has applications in various fields, including character behavior in virtual environments, as used in e-learning applications and serious games.

Emotion integration is necessary in such agents in order to generate complex believable behavior. It has been shown that emotions influence human cognition, interaction, decision-making and memory management in a beneficial and meaningful way. Several studies in the field of artificial intelligence have also been conducted ([35, 5, 15]) and have shown the necessity of emotion simulation in artificial intelligence, especially when human-computer interaction is involved.

This paper presents and artificial emotion simulation model and agent architecture (chapters 3, 4) that implements various human-inspired mechanics, such as emotions (section chapter 3.1), knowledge indexing (section 3.2), attention narrowing and memory management (sections 3.3 and 3.2) in an effort to provide artificial characters in virtual environments with believable behavior, as well as improve the functioning of artificial agent cognitive structures and procedures, improving agent performance.

Chapter 2

Inspiration

This chapter presents the theoretical psychology basis upon which our framework is built. We start out by presenting how emotions are perceived and assessed, and then present the effects that emotions have on human memory, learning and perception, thereby influencing both short and long-term behavior.

Further on, we move to related work in the field of emotion simulation for artificial agents, lessons learned from these and how our model expands on them.

2.1 Psychological Theory

The following section passes through the main ideas of several theories on human emotions [10], presents their main ideas and effects on physical and cognitive processes and makes references to their application regarding artificial agent design.

Increased interest in *emotional intelligence* has revealed yet another aspect of human intelligence. It has been shown that emotions have a major impact on the performance of many of our everyday tasks, including decision-making, planning, communication, and behavior [44].

Emotion is linked with mood, temperament, personality and disposition, and motivation. It is thought that emotions act as catalysts in brain activity, that increase or decrease the brain's activity level, direct our attention and behavior and inform us of the importance of events around us [5].

An important distinction is between the emotion and the results of the emotion, (behaviors, emotional expressions). Often, a certain behavior is triggered as a direct result of an emotional state (crying, fight or flight). If it is possible to experience a certain emotion without the associated behavior, then the behavior is not essential to the emotion.

Another means of distinguishing emotions concerns their occurrence in time. Some emotions occur over a short period of time (surprise), while others are more long lasting (love). Longer lasting emotions could be viewed as a predisposition to experience a certain emotion regarding a certain object or person (this may be implemented as a slow rate of emotion decay associated with a certain object or agent). A distinction is made between emotion episodes and emotional dispositions. An emotional disposition means the subject is generally inclined to experience certain emotions (e.g. an irritable person is disposed to feel irritation more easily than others).

Some theories place emotions within a more general category of *affective states* where affective states can also include emotion-related phenomena such as pleasure and pain, motivational states, moods, dispositions and traits.

2.1.1 Emotion Theories

Some theories argue that cognitive activity is necessary for emotions to occur. This means to imply the fact that emotions are about something (an object) or possess intentionality. The cognitive activity in question (judgment, evaluation) may be conscious or not.

Lazarus' theory [27] states that an emotion is a three stage disturbance that occur in the following order: 1. cognitive appraisal - the subject assesses the event cognitively. 2. physiological change - biological changes occur as a reaction to cognitive appraisal (increased heart rate, adrenal response). 3. action - the individual feels the emotion and chooses how to react. Lazarus accentuates that the type and intensity of emotion are controlled through cognitive processes. These processes influence the forming of the emotional reaction by altering the relationship between the person and the environment. Our model uses a similar sequence to that described by Lazarus in order to cognitively appraise events and choose an action: 1. triggering event. 2. cognitive assessment of event and associated emotional state 3. modify emotional state and physiological response (if any) 4. an action or series of actions are inserted into an action queue (be it an agenda or a plan, as used by forward or backward chaining inference engines respectively) according to the exhibited emotion and its intensity.

As shown in this section, emotion simulation can be used in agent and multi-agent systems either in order to replace several existing systems that perform the same functions that could be performed solely based on emotion simulation or new systems that agents and multi-agent systems should have.

Perceptual theory

This theory argues that conceptually based cognition is unnecessary for emotions to be about something. It states that bodily changes themselves perceive the meaningful content of the emotion. According to perceptual theory, emotions become similar to other senses, such as vision or touch, and provide information about the subject and the environment.

Cannon-Bard theory

The Cannon-Bard theory is a psychological theory which suggests that people feel emotions first and then act upon them. These actions include changes in muscular tension, perspiration, etc. The theory is based on the premise that one reacts to an event and experiences the associated emotion concurrently. The subject first perceives the stimulus, then experiences the associated emotion. Subsequently, the perception of this associated emotion influences the subject's reaction to the stimulus. The theory states that a subject is able to react to a stimulus only after perceiving the experience itself and the associated emotion.

Two-factor theory

Another cognitive theory is the Singer-Schachter theory. This theory is based on a study that shows that subjects can have different emotional reactions despite being in the same physiological state. During the study, experiments were performed where the subjects were placed in the same physiological state with an injection of adrenaline and observed to express either anger or amusement depending on whether a planted *peer* exhibited the same emotion, thus proving that a combination of cognitive appraisal and internal bodily state determine emotional response. According to the theory, *cognitions are used to interpret the meaning of physiological reactions to outside events*. This study shows that emotions are also a social mechanism and prompts research into using emotions as a useful tool in multi-agent systems.

2.1.2 Emotions and Memory

Emotions have a big impact on memory. Many studies have shown that the most vivid memories tend to be of emotionally charged events, which are recalled more often and with greater detail than neutral events. This can be linked to human evolution, when survival depended on behavioral patterns developed through a process of trial and error that was reinforced through life and death decisions. Some theories state that this process of learning became embedded in what is known as *instinct*.

A similar system may be used to decide the importance and relevance of events and knowledge and manage memory elements in artificial agents.

Emotions have positive effects on cognitive processes, such as increasing the likelihood of memory consolidation (the process of creating a permanent record). Also, over time, the quality of memories decreases, however memories involving arousing stimuli remain the same or improve [48]. A similar system may be used to organize and manage an artificial agents memory elements, improving memory storage efficiency, knowledge recall and rule firing speed. This mechanic is detailed in sections 3.2 and 3.3, and its effects on the decision making process are presented in section 3.5.

2.1.3 Emotions and Perception

Increase in arousal levels leads to *attention narrowing*, a decrease in the range of stimuli from the environment to which the subject is sensitive. This theory states that attention will be focused on arousing details, so that information regarding the source of the emotional arousal is memorized while peripheral details may not be [48].

Emotional events are also more likely to be processed when attention is limited, which implies the prioritized processing of emotional information [26]. Such a system of prioritized processing would be useful in resource-constrained agents or agents that exist in fast-changing dynamic environments (or environments that have a high event density ratio over short periods of time). The mechanics of such a system are discussed in chapter 3.3.

2.2 Computational Models

We begin this chapter by presenting two popular emotion models with different goals and approaches in emotion simulation, and end by presenting the findings of the PETEEI study that shows the importance of learning and evolving emotions.

As shown in this section (2.2), there are several key aspects to take into consideration when designing a memory and emotion architecture, and we believe that most are covered by our model, presented in the following section.

2.2.1 Ortony, Clore and Collins

The OCC (Ortony Clore & Collins) model has become the standard model for the synthesis of emotion. The main idea of this model is that emotions are valenced reactions to events, agents or objects (the particular nature being determined by contextual information). This means that the intensity of a particular emotion depends on objects, agents and events present in the environment [35]. In this model, emotions are reactions to perceptions and their interpretation is dependent on the triggering event, for example: an agent may be pleased or displeased by the consequences of an event, it may approve or disapprove of the actions of another agent or like or dislike an object. Another differentiation taken into account within the model is the target of an event or action and the cause of the event/action: events can have consequences for others or for oneself and an acting agent can be another or oneself. These consequences (whether for another or oneself) can be divided into desirable or undesirable and relevant or irrelevant.

In the OCC model, emotion intensity is determined by three intensity variables: *desirability* - reaction towards events and is evaluated with regard to goals, *praiseworthiness* relates to actions of other agents (and self), evaluated with regard to standards, and *appealingness* depends on reaction to objects and its value is determined with regard to attitudes.

The authors also define a set of intensity variables: sense of reality, proximity, unexpectedness and arousal are the global variables that influence all three emotion categories. There is also a threshold value for each emotion, below which an emotion is not subjectively felt. Once an emotion is felt however, it influences the subject in meaningful ways, such as focusing attention, increasing the prominence of events in memory, and influencing judgments.

As shown, according to the OCC model, *behavior is a response to an emotional state in conjunction with a particular initiating event* [35]. As in the

OCC model, in our model, emotions are valenced reactions to events, agents or objects determined by contextual information; in fact we will be using a simplified version of the OCC emotion evaluation process similar to the one used in [15] in early validation of our model.

2.2.2 Connectionist Model (Soar)

In this model, emotions are regarded as subconscious signals and evaluations that inform, modify, and receive feedback from a variety of sources including higher cognitive processes and the sensoriomotor system. Because the project focuses on decision-making, the model emphasizes those aspects of emotion that influence higher cognition [5].

The model uses two factors, *clarity* and *confusion*, in order to provide an agent with an assessment of how well it can deal with the current situation. Confusion is a sign that the agent's current knowledge, rule set and abilities are inadequate to handle the current situation, while clarity comes when the agent's internal model best represents events in the world. Since confusion is a sign of danger it is painful and to be avoided, while clarity, indicating the safety of good decisions, is pleasurable and sought-after.

Pleasure/pain, arousal and clarity/confusion form the basis of the emotional system. Instead of creating separate systems for fear, anger, etc., the authors assume that humans attach emotional labels to various configurations of these factors. For example, fear comes from the anticipation of pain. Anxiety is similar to fear, except the level of arousal is lower [5]. The advantage of a generalized system is that it does not require specialized processing.

Arousal is the main way that the emotional system interacts with the cognitive system. Memory and attention are most affected by changes in arousal. Highly aroused people tend to rely on well-learned knowledge and habits even if they have more relevant knowledge available.

In the Soar implementation of this system, the Soar rules that comprise have special conditions such that different kinds of rules only fire at differing levels of arousal. For example, highly cognitive rules will not fire at high levels of arousal, while more purely emotional rules may only fire at such levels [5]. This allows for a very general approach to emotions.

The Soar model emphasizes the effects of emotions on cognition and decision-making, effects considered important and also present in our own model (3.5, 2.1.2, 2.1.3).

2.2.3 Evolving Emotions (PETEEI)

Emotion is a complex process often linked with many other processes, the most important of which is learning: memory and experience shape the dynamic nature of the emotional process. *PETEEI is a general model for simulating emotions in agents, with a particular emphasis on incorporating various learning mechanisms so that it can produce emotions according to its own experience* [15]. The model was also developed with the capability to recognize and react to the various moods and emotional states of the user.

In order for an agent to simulate a believable emotional experience, it must adapt its emotional process based on its own experience. The model uses Q-learning with the reward treated as a probability distribution to learn about events and uses a heuristic approach to define a pattern and to further define the probability of an action, a1, to occur given that an action, a0, has occurred [15], thus learning about the user's patterns of actions. External feedback was used to learn what actions are pleasing or displeasing to the user. The application also uses accumulators to associate objects to emotions (Pavlovian conditioning). The accumulators are incremented by the repetition and intensity of the object-emotion occurrence.

An evaluation involving twenty-one subjects indicated that simulating the dynamic emotional process through learning provides a significantly more believable agent [15]. The PETEEI study focuses on the importance of learning agents and evolving emotions. Our memory and emotion architecture is built primarily with these two aspects in mind (chapter 4).

2.2.4 Meyer

In a position paper, Meyer also argues that a model inspired by human emotions can provide several reasonable and useful benefits [[54]]:

- reduce and control the nedeterminism involved in an agent's decision making process
- acheive flexible colaboration and cooperation between agents
- build better human-agent interfaces in hybrid human-agent multi agent systems

The paper aragues that, as long as many branches of computer science require heuristics to manage algorithm complexities, a model of affect based on the human emotion model and its perceived purpose would provide useful heuristics (as it serves the same purpose in humans).

Although many simplified implementations of the OCC model have been made, there has been no attempt at formalizing the complete logical structure of the model [[53]]. Meyer et al. attempt to put together a formal description of the OCC model in order to remove a number of ambiguities that stand in the way of its implementation.

Further work started based on the OCC model, and involved its formalization in order to remove a number of ambiguities that stand in the way of its full implementation. Starting from a qualitative formalization of the well-known model, Meyer et al. turned their attention to the quantitative aspects of emotions and investigated how these could be incorporated into the qualitative model [[51]].

The work is currently ongoing, *issues such a computational complexity and the possible need to empirically determine parameter settings have been taken into account when developing this formalization, but have not yet been explicitly dealt with* [[51]].

Chapter 3

Theoretical Model

In this section we first present our emotion representation scheme and its behavior, and then elaborate how emotions are used within and influence/-manage knowledge recall, the knowledge base hierarchy, attention span and the decision process. The model is based on the Cannon-Bard theory of emotion (section 2.1.1), cognitive appraisal is necessary for emotion synthesis and emotions are first felt and then acted upon. The model also implements the presented effects of emotion on memory (sections 2.1.3, 2.1.2) and may be expanded according to the perceptual theory that emotions also provide information about the subject and the environment.

3.1 Emotion Simulation

Analogous to the way primary colors combine, primary emotions could blend to form the full spectrum of human emotional experience. For example interpersonal anger and disgust could blend to form contempt.

We have chosen to represent / encode emotions as a series of weights, each describing the influence of one of the basic emotions intensity on the agents current emotional state (the eight basic emotions are those shown in Fig. 3.2, Plutchiks Wheel of Emotion: joy, trust, fear, surprise, sadness, disgust, anger and anticipation). The sum of emotion weights in our representation always equals one. The enforcement of this condition when certain emotions are modified by external stimuli allows gradual shifts from one dominant emotion to another, thereby avoiding sudden emotion changes (such as going from extreme anger (rage) to joy; the renormalization would increase joy and decrease anger in order to keep the sum of emotion weights equal to one;

Figure 3.1: Basic Emotions

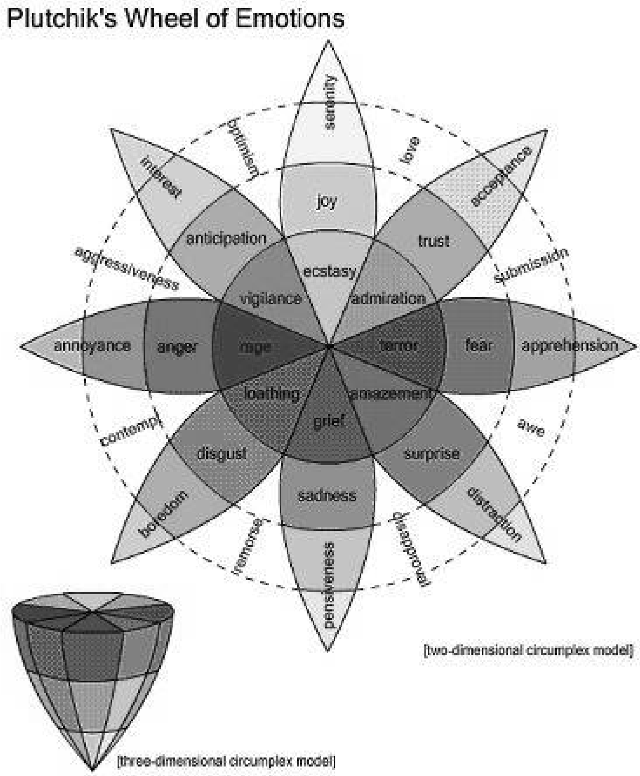

Figure 3.2: Plutchik's Wheel of Emotions

however, the dominant emotion would not become joy unless the stimulus was extremely powerful). The dominant (main) emotion is the emotion or combination of two adjacent emotions with the greatest weight / combined weights. Predisposition towards simple or complex emotions is left as the subject of future studies. For now, we will assume that characters / humans prefer simpler emotions as opposed to more complex ones.

Initially the agent is in a neutral, indifferent state, called the baseline state, where all emotion weights are equal to each other and equal to 0.125 (1/8). This representation was chosen because it provides an emotional state to relate to and because it is simpler to compute the effects of events on the emotional state as opposed to a percentage approach where it is possible for the sum of percentages to be lower than 100%. Emotion intensity (either one of the eight basic emotions or one of the eight compound emotions shown in Fig. 3.2 is calculated relative to this baseline, using the following formula:

$$\frac{\sum E_r}{n_r}.\varphi,$$

where E_r is the value of the weight that belongs to each basic emotion that contributes to the emotion being evaluated minus the baseline weight, as modified by internal and external events, based on the agent's preferences and expectations, as in the OCC model [35]), n_r is the number of relevant emotions and φ is a normalization factor that scales the result into the [0, 1] range ($\varphi = \frac{1}{1-0.125} = 1.142857$). The emotion with the highest intensity according to the above formula is chosen as the dominant (main) emotional state. If no clear emotion is dominant (most weights are within ε of each other), the character is considered to be in the baseline state.

Increased interest in emotions has revealed another aspect of human intelligence, showing that emotions have a major impact on the performance of many human tasks, including decision-making, planning, communication and behavior ([10]).

We intend to use our emotion simulation model in order to achieve the same effects that emotions have on human cognition, in artificial agents: catalyst in brain activity, motivation, direct attention and behavior, establish importance of events around us, therefore we follow a human cognitive emotion theory.

According to Lazarus ([27]), an emotion is a three-stage process:

1. cognitive appraisal *the subject assesses the event cognitively*

2. physiological change *biological changes occur as a reaction to cognitive appraisal (increased heart rate, adrenal response)*

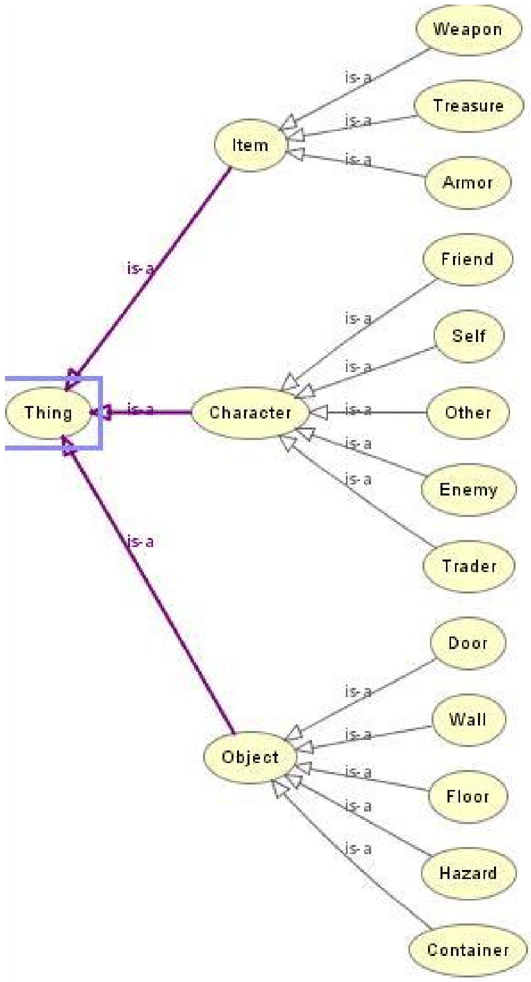

Figure 3.3: Knowledge Representation

3.2 Knowledge Representation

In order to enable the artificial character to learn about its environment and user it will be equipped with a perception module that will allow the agent to perceive external events (environment changes, other agents actions) as well as internal events (such as adding new knowledge or initiating an action) as well the associated main emotion and its intensity (the emotion that the agent is feeling in regard to the event) in chronological order. In our model, associated emotions reinforce some rules and facts, while suppressing others. This is how agent learning is achieved.

The agent's knowledge is organized in an inheritance-based hierarchical manner (Fig. 3.3) according to patterns and characteristics of items and other agents that it determines and considers important (the agent builds its ontology automatically); all objects have a common root node, similar to a top-level concept, which expands into subcategories which in turn branch out into new subcategories, based on observed common features. Individuals are represented as leaves in the knowledge tree. An individual object / agent (or a subcategory), may be absorbed into its parent category (while retaining information about which category an individual knowledge object belongs to) based on its frequency of use, time of last access and emotional response intensity (agents are more likely to learn and retain for longer periods of time and sometimes improve knowledge on subjects with a strong emotional load, in order to prune out useless, unimportant or infrequent knowledge while still retaining some information on how to react to stimuli and behave towards other agents (based on the characteristics of the group it belongs to). The agent will also be able to repeat knowledge it believes important in the long-term to itself, in order to increase its access frequency (similar to human learning).

The occurrence of events and knowledge in general are represented as lists of symbols in a manner similar to CLIPS or LISP.

The artificial agent will be able to learn about its environment (and other agents) through the emotions and emotion intensities associated to its knowledge items (facts and rules).

The agent's knowledge is organized in an inheritance-based hierarchical manner called an ontology. We are using an ontology because we want to be able to share common understanding of the structure of information among people or software agents and separate domain knowledge from operational knowledge.

Learning about the environment, human users and other agents is key to providing believable behavior in artificial characters ([15]). Each class and individual in the ontology has an associated emotion vector and intensity, which allows the agent to learn how to interact with a certain individual, as well as a class of individuals, as these associated emotions are propagated up the ontology hierarchy.

Rules and knowledge used in certain situations receive a fraction of the agent's current emotional state (this also travels upward through the inference chain adn ontology, also modifying rules and superclasses in order to help the agent cope with new, unknown individuals. For example, if the agent mainly encounters hostile characters (with whom it associates fear/anger) the emotion associated with the Character:Other class in the ontology, which is an average of all the individuals in that class, will be skewed in order to reflect this, and the agent will act accordingly with newly encountered agents).

The emotion e associated with an ontological class is

$$e = \frac{\sum\limits_{n_{subclasses}} e_{subclass}}{n_{subclasses}} \tag{3.5}$$

3.3 Perception and Memory Management

In fast-changing environments an agent may not have enough resources in order to react to every event that occurs, therefore a way of prioritizing important / significant events. This may be achieved through attention narrowing (section 2.1.3). Events are stored into an event queue and sorted according to emotional arousal associated with the event, ensuring more important events are processed first. Emotion preference is a domain-specific problem, for example fear (usually signaling danger) may be considered more important than disgust, and therefore should be processed with a higher priority. Events will be queued according to emotion intensity. Let's assume for example that the agent is running a labyrinth, and it does so in a turn-based manner, and has a limited number of actions each turn. It will not waste its time searching for traps if there are more serious threats, such as hostile creatures, around.

It has been determined that humans are more likely to remember bits of knowledge in similar emotional contexts that they learned or have used said knowledge in the past. In order to replicate this mechanic in artificial agents, we propose a system of indexing knowledge based on dominant emotion (including secondary emotions) and emotion intensity, similar to a hash-table, because there is a high probability for the same emotional context to be triggered in similar situations which may be solved by similar types of reasoning. For this purpose, both knowledge and rules (in rule-based systems for example) will have corresponding emotion types and intensities which will be updated every time the rule / knowledge item is used. These will be indexed in a table according to main emotion and sorted according to emotion intensity. When knowledge is required, the corresponding memory elements to the states main emotion and intensity are retrieved first and then the search is expanded around this point in the index table.

Emotions have a big impact on memory. Many studies have shown that the most vivid memories tend to be of emotionally charged events, which are recalled more often and with greater detail than neutral events. This can be linked to human evolution, when survival depended on behavioral patterns developed through a process of trial and error that was reinforced through life and death decisions. Some theories state that this process of learning became embedded in what is known as *instinct*.

A similar system may be used to decide the importance and relevance of events and knowledge and manage memory elements in artificial agents. This process is described in sections.

It has been discovered in humans that an increase in arousal levels leads to attention narrowing ([48]). Attention narrowing is a decrease in the range of stimuli from the environment that the subject is sensitive to, in other words, perception will be focused on more arousing events and details first and these will be processed sooner than non-arousing details and events [48]. This means that highly emotional events are more likely to be processed when attention is limited, leading experts to attest that emotional information is processed according to priority ([26]).

In fast-changing environments an agent may not have enough resources in order to react to every event that occurs, therefore a way of prioritizing important / significant events. This may be achieved through attention narrowing. Events are stored into an event queue and sorted according to emotional arousal associated with the event, ensuring more important events are processed first. Emotion preference is a domain-specific problem, for example fear (usually signaling danger) may be considered more important than disgust, and therefore should be processed with a higher priority. Events will be queued according to emotion intensity. Let's assume for example that the agent running the labyrinth does so in a turn-based manner, and has a limited number of actions each turn. It will not waste its time searching for traps if there are more serious threats, such as hostile creatures, around.

Humans are more likely to remember bits of informationand exhibit similar behavior and reasoning as exhibited in past similar emotional contexts. This can be used as a way of remembering context-appropriate facts and rules in artificial agents and may speed up the reasoning process as it does in humans.

We have devised a way of indexing facts and rules based on dominant emotion and emotion intensity similar to a hash-table. Both facts and rules have an associated emotion and intensity, which will be updated by a fraction of the agent's emotional state every time they are used to infer new data. These will be indexed in a table according to main emotion and sorted according to emotion intensity value (each cell in the table is a fact/rule).

At each new inference stage, a relevant rules relevant facts list are created, starting from the table cell that most closely matches the agent's current state and expanding outwards. These lists are then used to create an operator activation list, which will be triggered in the order that they are inferred.

The main differences between this model and other emotion simulation approaches are the perception and memory management mechanics that al-

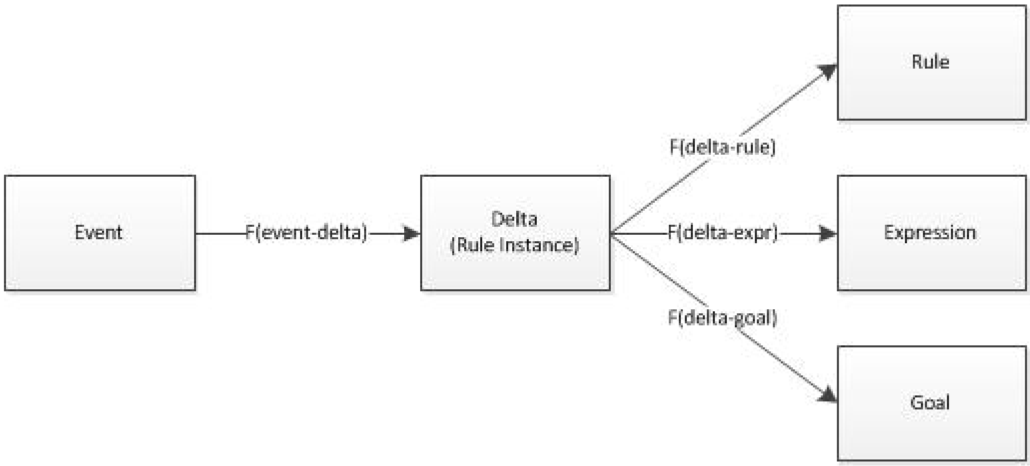

Figure 3.4: Learning through Emotion Influence

low emotions to indirectly influence the agent's decision and planning processes. Another difference from other approaches is that in this model, emotions serve a purpose, and are not merelely a goal in and of themselves.

3.4 Learning

Our model uses emotion intensity as the main means to learn by. Associated emotion intensity weighs each element in the model (rules, facts) and reinforces some while supressing others, however, by using the similarity function to weigh elements instead of the emotion intensity itself, the elements being enforced / suppressed change dynamically depending on the agent's emotion state. This means that knowledge relevant to the agent should be able to access knowledge more relevant to its state faster.

We want the agent to adapt its behavior to a dynamic environment, so whenever an event changes the agent's (emotional) state, we influence all inferences (and knowledge structures that took part in those inferences) so that their associated emotion matches a little bit more with the agent's current state, increasing the likelihood that and agent will trigger certain rules in similar emotional contexts.

As shown in Fig. 3.4, the agent learns based on feedback from the environment and other agents. When the agent perceives an event, it assumes that it is a consequence of all rules triggered that turn, therefore the emotion associated with the perceived event influences all rule instances triggered that inference turn, which in turn influence the rule that created them, the goal that the rule helps achieve (if any) and all expressions part of the rule, by a different factor F for each influence type.

3.5 Decision Process

Emotions and Decision-Making Although it is clear that emotions sometimes impede deliberative decision-making, one school of thought affirms that emotions provide a way of coding and compacting experience to enhance fast response selection. In evolutionary terms, it is better to respond immediately to a threat than take the time to rationally consider the best course of action.

Another use for emotion simulation in multi-agent systems is generating complex believable behavior, important in simulation environments and human computer interaction.

Future work involves exploring and integrating both schools of thought into the architecture.

We are not proposing a specific mechanic that influences the decision making process, however, knowledge retrieval and event handling are influenced through attention narrowing and knowledge indexing, thereby determining rule order in the agents agenda or plan. It should be clear that the agent adapts its behavior according to the dominant emotion and its context; for example, if its current state is fear, it will make efforts (introduce operators in the agenda or formulate plans, depending on the particularities of the agent) in order to eliminate the threat or get away from this situation. It will also make efforts in the future (if able) to avoid circumstances that led to the current state. Another example, if something causes the agent joy, the agent will be more likely to fire operators or formulate plans that lead to the joyous occurrence.

3.6 Visual Representation

The weight vector representation discussed in the previous section is very well suited to graphical representation, as we can have one graphical representation for each base emotion, and use the emotion vector weights to interpolate them. This resulting representation may in turn be interpolated with animation sequences such as speaking or breathing. We may also vary activity levels by speeding up or slowing down the animation according to emotion intensity. The baseline state yields the agents base movement speed. This is increased or decreased according to emotion intensity.

An alternative graphical representation would be to assign a color to each base emotion, as shown in Fig. 3.1 and blend them according to emotion weights; however, this may become confusing, as there are only three primary colors.

Chapter 4

Architecture

The proposed agent architecture is presented in Fig. 4.1.

Facts and rules are fed into the inference engine according to their associated emotion and its intensity, by way of knowledge indexing, as described in sections 3.3, 3.2.

The inference engine uses two inference methods (forward and backward chaining) running in parallel and produce rule activations along with an associated emotion vector, which are added to an agenda and a plan respectively ([29]). These are then merged into a conflict set and sorted according to associated emotion intensity and then trigerred.

The architecture also uses a truth maintenance system ([29]) so that the agent maintains a consistent world view, each new inferred fact has a justification and we are able to propagate associated emotions up the inference chain.

We are not proposing a specific mechanic that influences the decision making process, however, knowledge retrieval and event handling are influenced through attention narrowing and knowledge indexing, thereby determining rule order in the agents agenda or plan. It should be clear that the agent adapts its behavior according to the dominant emotion and its context; for example, if its current state is fear, it will make efforts (introduce operators in the agenda or formulate plans, depending on the particularities of the agent) in order to eliminate the threat or get away from this situation. It will also make efforts in the future (if able) to avoid circumstances that led to the current state. Another example, if something causes the agent joy, the agent will be more likely to fire operators or formulate plans that lead to the joyous occurrence.

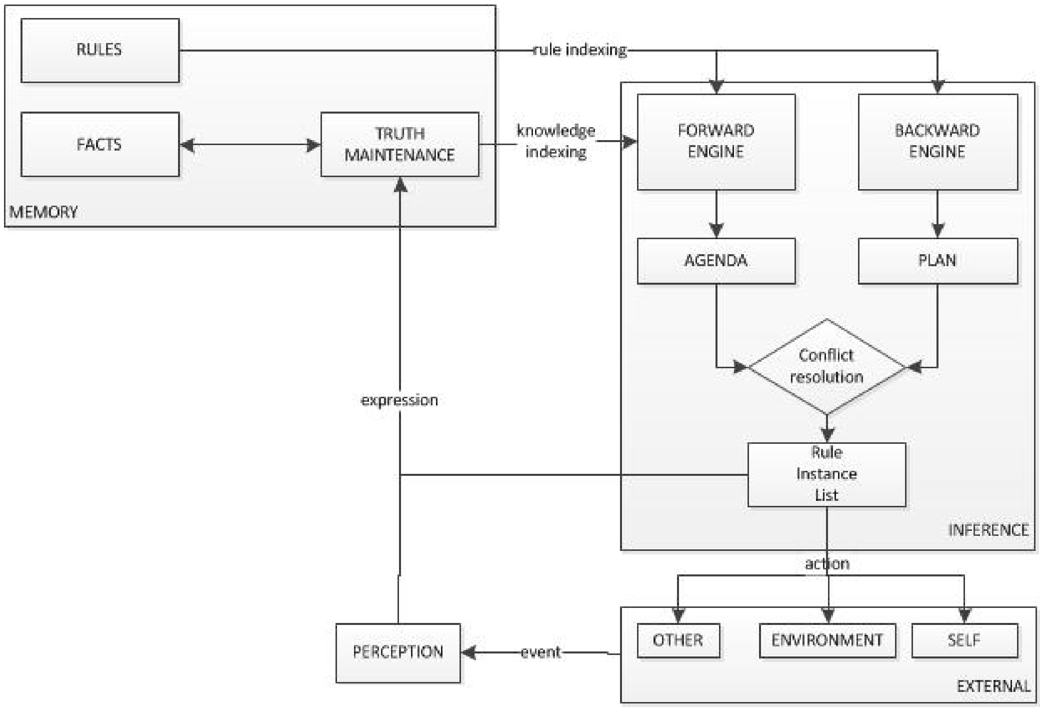

Figure 4.1: Architecture

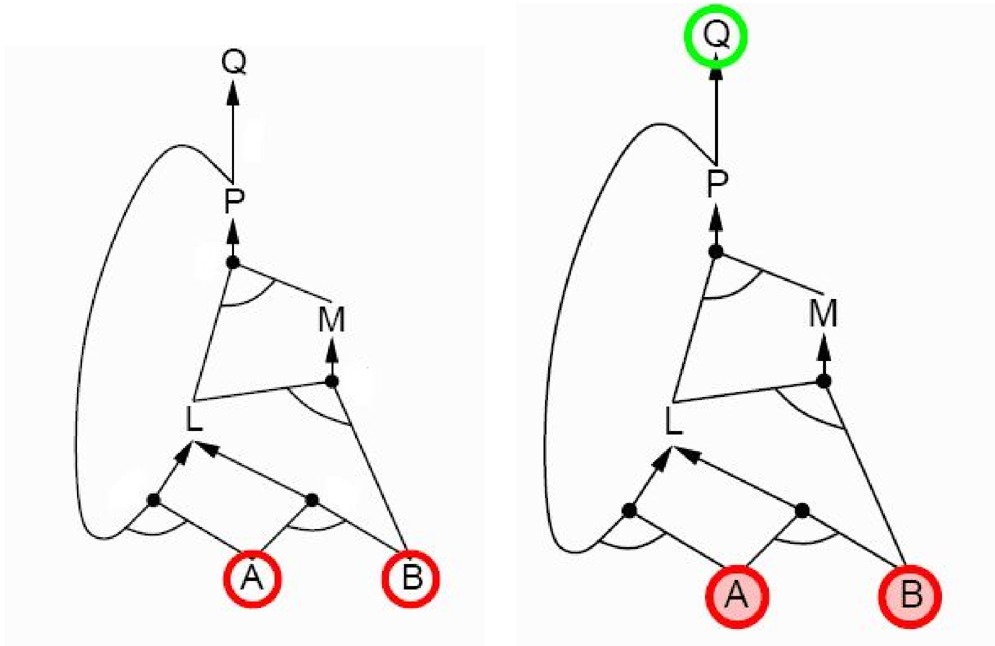

(a) Forward chaining

(b) Backward chaining

Figure 4.2: Inference methods used

4.1 Inference Engine

Forward Chaining

Forward chaining is a data-driven inference method. The engine hungrily awaits new knowledge in order to apply rules that match existing data.

Forward chaining starts with the available data and uses inference rules to extract more data (from an end user for example) until a goal is reached. An inference engine using forward chaining searches the inference rules until it finds one where the antecedent (If clause) is known to be true. When found it can conclude, or infer, the consequent, resulting in the addition of new information to its data.

Forward chaining inference engines will iterate through this process until a goal is reached. This seems as an ineffective way to go about solving a problem, especially when we have a large number of rules and data and only a few paths that lead to goal states. Backward chaining seems to be the better problem solver, and indeed it is; however, backward chaining does have a weakness, it never infers data that is not explicitly goal-related, even if somewhere down the inference path it may lead to goal-related data, or worse, conflicting data. In essence, we are using the forward chaining part of our inference engine to gather knowledge about the complex, dynamic environment, where multiple aspects of the world may be changing at any given moment. When run, a forward-chaining algorithm presents a conflict set of rules to apply (all rules that match current knowledge). Ways of deciding on application order will be discussed in a further section. At times, our forward-chaining inference engine may infer contradictory data; we need a way of detecting and dealing with such situations. The best approach is a truth maintenance system. Also, given that rule-based systems may have a large number of rules working with a large knowledge base, we need a fast pattern-matching algorithm (pattern-matching may be as much as 80

Backward Chaining

Backward chaining is a lazy inference method. It only does as much work as it has to. Backward chaining starts with a list of goals (or a hypothesis) and works backwards from the consequent to the antecedent to see if there is data available that will support any of these consequents. An inference engine using backward chaining would search the inference rules until it finds one which has a consequent that matches a desired goal. If the antecedent of that rule is not known to be true, then it is added to the list of goals (in order for one's goal to be confirmed one must also provide data that confirms this new rule), making it ideal at building an at least partially (we may have to accomplish disjoint goals, as well as having to accomplish several goals concurrently with no specified order in order to confirm a rule) ordered plan to follow in order to achieve the agents goals, dynamically decomposing goals into subgoals recursively until the agent selects among primitive operators that perform actions in the world. The backward chaining part of our inference engine accomplishes our agents goal-oriented planning, resulting in the composition of a plan, a sequence of partially-ordered goals to pursue in the future.

4.2 Truth Maintenance

A truth maintenance system is used to maintain a consistent set of beliefs about the world. It is based on the following principles:

each action in the problem-solving process has an associated justification

when a contradiction is obtained, find the minimal set of assumptions that generated the contradiction. Select an element from this set and defeat it. The justification for the contradiction is no longer valid and the contradiction is removed.

propagate the effects of adding a justification and of eliminating a belief (keep consistency)

A truth maintenance system keeps beliefs as a network of two types of nodes: beliefs and justifications. A belief is an expression which can be true or false. A belief node contains a label (IN if the belief is considered true, OUT otherwise), a list of justifications for the node, and a list of the justifications of which the node is part of (and the label it must have in order for that justification to be valid). Justification nodes contain the following information: the inference type of a justification (in our case premise (always IN) or modus ponens (inferred)), a list of nodes the justification justifies and a list of nodes (and necessary associated values) that participated in the inference. This representation allows the propagation of consequences through the belief network of an agent. A truth maintenance system is easily integrated with the RETE pattern-matching algorithm (nodes that are IN will be considered true, and present in the RETE network, and nodes that are OUT will be considered false and will not be present in the RETE network; nodes that change state from OUT to IN will be added and nodes that change state from IN to OUT will be retracted from the RETE network).

As can be seen, the RETE network and a TMS are easy to interface with each other, since they operate on different areas of a rule system (working memory vs. rule base).

Truth maintenance systems are not usually used alongside backward chaining inference engines because the inference engine rarely changes a nodes state; however, should this occur, the change is fed into the TMS by the inference engine and the changes propagated through the network, just as in the forward chaining case.

4.3 Grammar

We have designed a language in order to specify agent knowledge and rules. This chapter describes its grammar. The basic element of the language is the expression.

```
expression :
        ((IN| assert) | (OUT| retract)) symbol
          symbol symbol
```

Each expression, regardless whether it appears in the knowledge base or a rule, represents a fact. An expression is made up of a modifier that states the desired state of the fact in the architecture's truth maintenance system (IN or OUT) and a series of three symbols which need to represent an ontology fact (subject property object). Each symbol may be a variable, a reference, a value, a condition or an operator.

```
symbol :
        variable | reference | value |
          condition | operator
```

Further on, these symbols are reduced to basic literals:

```
variable :
        '?'identifier
reference :
        identifier
value :
        Int | Float | String
operator :
        ( + | - | * | / ) symbol
condition :
        ( < | > | <= | >= | <> ) symbol
```

Knowledge base facts may not contain variables, conditions or operators. The knowledge base structure is specified using the **class**, **property** and **individual** keywords.

A class specification is made up of the *class* keyword, possibly a list of classes that the class extends and a list of (property name, value / reference) pairs that specify restrictions.

```
class :
        'class' id ('extends' id (',' id)* )? (
            id ':' (value | reference) )*
```

Properties are defined by using the *property* keyword, a property name and a pair of classes that are the property's range and domain.

```
property :
        'property' id ('extends' id)? ':'
          domain '->' range
```

Last but not least, individuals in the agent's knowledge representation scheme are specified by the *individual* keyword and an identifier, followed by a list of (property, value) pairs which translate to expression of the form ((modifier) identifier property value).

```
individual :
        'individual' id ':' id ((IN|OUT) id '='
          (value | reference))*
```

Another key part of the reasoning engine are rules. Rules are specified through the use of the *rule* keyword and an identifier and are made up of a series of expressions known as the antecedent, which means these expressions have to be evaluated as true for the rule to be triggered and another series of expressions named the consequent, which represent facts that become true when the rule is triggered.

```
rule :
        'rule' id ('repeatable')? (expression)+
          '==>' ((expression | action))+
```

A rule's consequent may also contain actions

```
action :
        'action' id symbol* expression*
```

which are operators that the agent can apply to the environment or other agents (specified as a list of parameters) and that have consequences that happen over time and are not always guaranteed (the list of expressions that make up an action are the *expected* consequences of an action).

The driving force of the architecture are *goals*. Each goal is made up of one expression which the agent will spend energy and resources to fulfill. Goals also have an attached salience value that represents their importance.

```
goal :
        'goal' identifier (salience Int)?
          expression+
```

Last but not least, definitions for the building blocks of our grammar: values and identifiers.

```
Int :
        ('-')?('0'..'9')+
Float :
        ('0'..'9')+ '.' ('0'..'9')+
String :
        ('a'..'z'|'A'..'Z')+

identifier :
        String (String | Int | '_' | '-')*
```

In the following section we present a source code input example used to test inference and planning.

4.3.1 Input example

```
self = Fritz

actions {
        action MoveTo ?orc          // action header
                IN ?self InRange ?orc    // expr / consequence
}

classes {
        class Creature
        class Sound
        class Species
        class Color
}

properties {
        property MakesSound : Creature -> Sound
        property Eats : Creature -> Creature
        property HasSpecies : Creature -> Species
        property HasColor : Creature -> Color

        property HasHealth : Creature -> Int

        property distance : Creature -> Int
}

individuals {
        individual Croak : Sound
        individual Chirp : Sound

        individual Fly : Creature
        individual Fritz : Creature
```

```
            IN MakesSound = Croak
            IN Eats = Fly
            IN HasHealth = 10
            IN distance = 1

      individual Canary : Species
      individual Frog : Species

      individual Green : Color
      individual Yellow : Color
}

goals {
      goal FritzColor
            IN Fritz HasColor ?x

      goal IncreaseHealth
            IN ?self HasHealth +>=5
}

rules {
      rule Frog
            IN ?x MakesSound Croak
            IN ?x Eats Fly
            ==>
            IN ?x HasSpecies Frog

      rule FrogGreen
            IN ?x HasSpecies Frog
            ==>
            IN ?x HasColor Green

      rule Canary
            IN ?x MakesSound Chirp
            ==>
            IN ?x HasSpecies Canary

      rule CanaryYellow
            IN ?x HasSpecies Canary
            ==>
            IN ?x HasColor Yellow

      rule FritzChangeColor
            IN Fritz HasSpecies Frog
            IN Fritz HasColor Green
            ==>
            OUT Fritz HasSpecies Frog
            IN Fritz HasSpecies Canary
}
```

Chapter 5

Conclusion and Future Work

5.1 Conclusion

One of the goals of this work is to demonstrate how emotions can affect agent attention focus, performance, learning, memory management and decision making and how the effects can vary according to the personality of the agent. Another is to provide a flexible and scalable emotion representation model that reproduces the effects of emotion observed in humans and that allows for a general dynamic approach to emotion in artificial agents.

One of the goals of this work is to develop an emotion simulation model and agent architecture that provides artificial characters in virtual environments with believable behavior in order to enhance virtual environment experiences for human users.

Another goal, as presented in this paper, is to use artificial emotion simulation as motivation for agent behavior, as well as improve reasoning capabilities and memory management by simulating mechanics tied to emotions found in humans (see section 3).

The main contributions of this model are that emotion is not only tied to the agent's learning and reasoning process, but it is a key component, providing agents with the underlayer for the simulated emotion, making the reasoning process more akin to humans, achieving complex, believable behavior for virtual agents. Another important contribution of this model is the improvement of agent performance and effectiveness through emotion simulation as it is currently believed emotions do in humans.

5.2 Future Work

Future work includes mixed emotion resolution (i.e. when emotions on different sides of Plutchik's Wheel result as dominant) and exploring whether an explicit way in which the emotion model interacts with the emotion process is necessary or not.

There are at least two directions in which the model can be extended, both exploring the multi-agent capabilities of the model. First of all, the model can be extended to be used as a trust and reputation model among agents, based on the reciprocal altruism model that is believed to be used by humans (based on emotions). Another direction for future research, still in the field of multi-agent systems, would be to extend the model so that it supports the theory of reciprocal altruism (where an agent acts in a manner that temporarily reduces its fitness while increasing another agent's fitness, with the expectation that the other agent will act in a similar manner at a later time).

Another possible modification that may improve the system would be to get rid of the ontology (and thus the restrictions imposed by it), relying on user interpretation of symbols to give meaning to expressions. This would allow us to increase language complexity, allowing for expressions of any length and rules to reference and call other rules.

Bibliography

[1] M. Ben Ari. *Principles of concurrent and distributed programming.* Prentice Hall, N.Y., 1990.

[2] Christoph Bartneck. Integrating the occ model of emotions in embodied characters. *Design*, pages 1–5, 2002.

[3] J. Bates. The role of emotion in believable agents. In *Commun*, pages 122–125. ACM 37, 1994.

[4] Ronald Brachman and Hector Levesque. *Knowledge Representation and Reasoning.* Morgan Kaufman, 2004.

[5] Eric Chown, Randolph M. Jones, and Amy E. Henninger. An architecture for emotional decision-making agents. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 352–353, Bologna, Italy, July 2002.

[6] A.D. Craig. Interoception: The sense of the physiological condition of the body. *Current Opinion in Neurobiology*, 13:500–505, 2003.

[7] Michael David and Chen Sande. Serious games: Games that educate, train, and inform, 2006.

[8] Sara de Freitas and Steve Jarvis. A framework for developing serious games to meet learner needs. In *Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*, 2006.

[9] Sara de Freitas and Steve Jarvis. Serious games engaging training solutions: A research and development project for supporting training needs. *British Journal of Educational Technology*, 38(3):523–525, May 2007.

[10] Ronald de Sousa. Emotion. Online, available at http://plato.stanford.edu/entries/emotion/, September 2010. Stanford Encyclopedia of Philosophy.

[11] A. S. Drigas, L. G. Koukianakis, and J. G. Glentzes. Knowledge engineering and a virtual lab for hellenic cultural heritage. In *Proceedings of the 5th WSEAS Int. Conf. on Artificial Intelligence*, 2006.

[12] Oliver Duvel and Stefan Zerbst. *3D Game Engine Programming*. Course Technology PTR, 2004.

[13] David Eberly. *The 3D Game Engine Architecture: Engineering Real-time Applications with Wild Magic*. ?, 2005.

[14] David Eberly. *3D Game Engine Design: A Practical Approach to Real-Time Computer Graphics*. ?, 2006.

[15] Magy Seif El-Nasr, Thomas R. Ioerger, and John Yen. Peteei: a pet with evolving emotional intelligence. In *Proceedings of the third annual conference on Autonomous Agents*, pages 9–15, Seattle, Washington, United States, April 1999.

[16] Wolfgang f. Engel. *Shader X2 Introductions and Tutorials with DirectX 9*. Wordware Publishing Inc., 2004.

[17] Randima Fernando. *GPU Gems: Programming Techniques, Tips and Tricks for Real-Time Graphics*. Pearson Education, 2004.

[18] Adina Magda Florea and Adrian George Boangiu. *Bazele logice ale inteligentei artificiale*. Universitatea Politehnica Bucuresti, 1994.

[19] M.P. Georgeff, B. Pell, M.E. Pollack, , M. Tambe, and M. Wooldridge. The belief-desire-intention model of agency. In *Proceedings of the 5th international Workshop on intelligent Agents V, Agent theories, Architectures, and Languages*, 1999.

[20] Jason Gregory. *Game Engine Architecture*. A.K. Peters, 2009.

[21] A.E. Henninger, R.M. Jones, and E. Chown. Behaviors that emerge from emotion and cognition: implementation and evaluation of a symbolic-connectionist architecture. In *Proceedings of the Second international Joint Conference on Autonomous Agents and Multiagent Systems*, pages 321–328, Melbourne, Australia, 2003. ACM, New York, NY.

[22] Amy E. Henninger, Randolph M. Jones, and Eric Chown. Behaviors that emerge from emotion and cognition: A first evaluation. In *Interservice/Industry Training Systems and Education Conference(I/ITSEC)*, 2002.

[23] Obe Hostetter. Video games - the necessity of incorporating video games as part of constructivist learning. James Madison University, 2002. Department of Educational Technology.

[24] HUMAINE Network of Excellence. Humaine emotion annotation and representation language (earl). Online. available at http://emotion-research.net/projects/humaine/earl/index _html/ ?searchterm=emotion%20markup%20language.

[25] James Kennedy and Russell C. Eberhart. *Swarm Intelligence*. Morgan Kaufmann, 2001.

[26] E.A. Kensinger. Remembering emotional experiences: The contribution of valence and arousal. *Reviews in the Neurosciences*, 15:241–251, 2004.

[27] R.S. Lazarus. *Emotion and Adaptation.* Oxford University Press, New York, NY, 1991.

[28] Jill Fain Lehman, John Laird, and Paul Rosenbloom. A gentle introduction to soar, an architecture for human cognition. In *S. Sternberg and D. Scarborough (Eds).* MIT Press, 1996. Invitation to Cognitive Science.

[29] Valentin Lungu. Rule-based system for emotional decision-making agents. In *Sisteme distribuite ISSN 2067-5259*, Suceava, 2009. online.

[30] Valentin Lungu. Artificial emotion simulation model. In *Agents for Complex Systems Workshop*, Timisoara, 2010. SYNASC 2010.

[31] Xiao Mao and Zheng Li. Implementing emotion-based user-aware e-learning. In *CHI 2009, Spotlight on Works in Progress*, 2009.

[32] D. Moura and E. Oliveira. Fighting fire with agents: an agent coordination model for simulated firefighting. In *Proceedings of the 2007 Spring Simulation Multiconference*, pages 71–78, San Diego, CA, 2007. Society for Computer Simulation International. Volume 2.

[33] D. Nardi and R. Brachman. An introduction to description logics. Online. available at http://www.inf.unibz.it/ franconi/dl/course/dlhb/dlhb-01.pdf.

[34] E. Oliveira and L. Sarmento. Emotional advantage for adaptability and autonomy. In *Proceedings of the Second international Joint Conference on Autonomous Agents and Multiagent Systems*, pages 305–312, Melbourne, Australia, 2003. ACM, New York, NY.

[35] A. Ortony, G. Clore, and A. Collins. *The cognitive structure of emotions.* Cambridge University Press, Cambridge, 1988.

[36] H. Van Dyke Parunak, R. Bisson, S. Brueckner, R. Matthews, and J. Sauter. A model of emotions for situated agents. pages 993–995, New York, NY, 2006. ACM.

[37] D. Pereira, E. Oliveira, and N. Moreira. Formal modelling of emotions in bdi agents. *Computational Logic in Multi-Agent Systems: 8th international Workshop, CLIMA VIII*, 2008.

[38] David Pereira, Eugnio Oliveira, and Nelma Moreira. Modelling emotional bdi agents. In *Workshop on Formal Approaches to Multi-Agent Systems (FAMAS 2006)*, 2006.

[39] David Pereira, Eugnio Oliveira, Nelma Moreira, and Lus Sarmento. Towards an architecture for emotional bdi agents. In *EPIA 05: Proceedings of 12th Portuguese Conference on Artificial Intelligence*, pages 40–47. Springer, 2005.

[40] Matt Pharr. *GPU Gems 2: Programming Techniques for High-Performance Graphics and General-Purpose Computation.* Pearson Education, 2005.

[41] David Poole, Alain Mackworth, and Randy Goebel. *Computational Intelligence: a Logical Approach.* Oxford University Press, 1998.

[42] J. Rickel and W. Lewis Johnson. Steve: an animated pedagogical agent for procedural training in virtual environments.

[43] Stuart J. Russell and Peter Norvig. *Artificial Intelligence – A Modern Approach.* Prentice Hall, Englewood Cliffs, New Jersey 07632, 1995.

[44] K. Scherer. What are emotions and how can they be measured? *Social Science Information*, 44(4):695–729, 2005.

[45] S. Schiaffino, A. Amandia, I. Gasparini, and M. Pimenta. Personalization in e-learning: the adaptive system vs. the intelligent agent approches.

[46] M. Schrder, H. Pirker, and M. Lamolle. First suggestions for an emotion annotation and representation language. In *Proceedings of LREC'06 Workshop on Corpora for Research on Emotion and Affect*, pages 88–92, Genoa, Italy, 2006.

[47] M. Schroder, L. Devillers, K. Karpouzis, J. Martin, C. Pelachaud, C. Peter, H. Pirker, J. Schuller, J. Tao, and I. Wilson. What should a generic emotion markup language be able to represent? In *Second International Conference on Affective Computing andIntelligent Interaction: Lecture Notes in Computer Science*. Springer, 2007.

[48] T. Sharot and E.A. Phelps. How arousal modulates memory: Disentangling the effects of attention and retention. *Cognitive, Affective & Behavioral Neuroscience*, 4:294–306, 2004.

[49] M.P. Singh. Agent communication languages: Rethinking the principles. *Computer 31*, 12:40–47, 1998.

[50] Bas R. Steunebrink, Mehdi Dastani, and John-Jules Ch. Meyer. The occ model revisited.

[51] Bas R. Steunebrink, Mehdi Dastani, and John-Jules Ch. Meyer. Towards a quantitative model of emotions for intelligent agents.

[52] Bas R. Steunebrink, Mehdi Dastani, and John-Jules Ch. Meyer. A logic of emotions for intelligent agents. In *Proceedings of the 22nd national conference on Artificial intelligence - Volume 1*, pages 142–147. AAAI Press, 2007.

[53] Bas R. Steunebrink, Mehdi Dastani, and John-Jules Ch. Meyer. The occ model revisited. In *Proceedings of the 4th Workshop on Emotion and Computing – Current Research and Future Impact*, 2009.

[54] B.R. Steunebrink, M.M. Dastani, and J-J.Ch. Meyer. Emotions as heuristics in multi-agent systems. In *Proceedings of the 1st Workshop on Emotion and Computing - Current Research and Future Impact*, pages 15–18, 2006.

[55] Gerard Tel. *Introduction to Distributed Algorithms.* Cambridge University Press, 1994.

[56] W3C Incubator Group. Elements of an emotionml 1.0. Online, 2008. available at http://www.w3.org/2005/Incubator/emotion/XGR-emotionml/.

[57] Michael Wooldridge and Nicholas R. Jennings. Pitfalls of agent-oriented development. In *Proceedings of the Second International Conference on Autonomous Agents*, pages 385 – 391. ACM Press, 1998.

[58] Nick Yee. The labor of fun: How video games blur the boundaries of work and play. *Games and Culture*, 1:68–71, 2006.