

# Towards an Architecture for Emotional BDI Agents

David Pereira, Eugénio Oliveira and Nelma Moreira

Technical Report Series: DCC-2005-09



Departamento de Ciência de Computadores – Faculdade de Ciências

&

Laboratório de Inteligência Artificial e Ciência de Computadores

---

Universidade do Porto

Rua do Campo Alegre, 823 4150 Porto, Portugal

Tel: +351+2+6078830 – Fax: +351+2+6003654

<http://www.ncc.up.pt/fcup/DCC/Pubs/treports.html>

# Towards an Architecture for Emotional BDI Agents

David Pereira and Nelma Moreira  
DCC-FCUP & LIACC  
{dpereira,nam}@ncc.up.pt

Eugénio Oliveira  
NIAD&R-FEUP & LIACC  
eco@fe.up.pt

January 11, 2006

## Abstract

In this report we present the Emotional-BDI architecture, an extension to the BDI architecture supporting Artificial Emotions and including internal representations for agent's Capabilities and Resources. The architecture we present here, is conceptual, defining which components should exist so that Emotional- BDI agents can use Effective Capabilities as well as Effective Resources in order to better cope with highly dynamic environments.

**Keywords:** BDI Agents, Artificial Emotions

## 1 Introduction

The BDI architecture is one of the most well known and studied software agents' architectures. The main reasons are, most notably, the architecture's widely accepted philosophical roots [BIP88, Den87], its logical frameworks for modelling and reasoning about BDI agents [GPP<sup>+</sup>98] and a considerable set of software systems which employ the architecture's concepts [GR96, Geo].

More recently, the attention of the Agent Community started to focus on a new (general) kind of agency model: Emotional Agents, which use results about the beneficial aspects of human emotions [Dam94] and apply artificial versions of them in Rational Agents.

In this report we present the result of a research work aimed at extending the classic BDI architecture with Artificial Emotions: the conceptual Emotional-BDI architecture. These Artificial Emotions are metaphors for the emotions that can be computationally implemented. As a support for these Artificial Emotions, the architecture is equipped with an internal representation of Resources and Capabilities, making the agents "Self-Aware". The main idea is to use Artificial Emotions to help the BDI architecture to incorporate a more accurate model of Practical Reasoning [BIP88] through the interconnection of the mechanisms which are charged of managing the Emotional State of an agent, the Resources and Capabilities, and all the mechanisms which compose the original BDI architecture. This way, the kind of processing done in the architecture we propose is the same as the BDI architecture, except for the fact that there can be applied different kinds of algorithms for performing the same tasks, and these algorithms will be chosen depending on the agents Emotional State, Resources and Capabilities.

This report is organised as follows. In Section 2 we will present, from our perspective, the positive and negative aspects of the BDI architecture and draw some conclusions about its real usability in today's problem-solving reality. In Section 3 we establish a connection between the problems detected and the solutions provided by the functional roles of emotions which deal with Resources and Capabilities management. In Section 4 we will start by introducing all the concepts which we consider as prerequisites to the development of a conceptual Emotional-BDI architecture, present this architecture and describe its new components, the interactions between all the architecture's components in terms of information flux and processing, and some topics about implementation. Finally, in Section 5 we will draw some conclusions about the work to be presented in this report and also point out the paths for future research.

## 2 Motivation

The BDI architecture has shown to be a very successful one, as it can be seen from the number of software systems which implemented it [Geo, GR96, RG95]. However, we believe that the architecture can still be more effective so that it can face the new generation of problems characterised by high levels of unpredictability, complexity and dynamics. Some opinions about the necessity of this enhancement process were already put forward by [GPP<sup>+</sup>98] and, therefore, we followed those clues to conduct our own research.

### 2.1 Pros

First, let us review the main reasons for the BDI architecture's success in the Agents Community. These can be divided in three different classes [GPP<sup>+</sup>98]:

**Strong philosophical roots:** The BDI Architecture was created by Bratman et al. [BIP88] with the goal of being able to establish a good balance between reactive and deliberative behaviour, while still being a pure deliberative architecture. For that, and in addition to the philosophical concept of Intentional System [Den87] that underlies the own notion of agent, the architecture uses the concept of Resource Bounded Practical Reasoning which is a variant of the classical reasoning directed towards the execution of actions, but where Intentions have a central role in driving means-ends reasoning, constraining deliberation and influence the Belief base of an agent [BIP88, Wei99]. Both philosophical concepts are widely accepted and also provide a natural way of describing the behaviour which agents should exhibit only through the use of the following mental states: Beliefs, Desires and Intentions. Therefore, the development of BDI agents is feasible either by computer programming experts or experts of the domain where the agents will act.

**Elegant logical formalisms:** Another attractive aspect of the BDI architecture is the set of logical frameworks exclusively developed to reason about BDI agents, and which also provide an important guidance towards the correct development of software programs which make use of the architecture. The most well know are the Rao and Georgeff's **BDI<sub>CTL</sub>\*** framework and the **KARO** framework of Woek et al. For a more detailed overview about these two logical frameworks refer to [vdHW03].

**Software implementations:** The BDI architecture has been applied to software systems on the realms of academic and real-world software systems. In the realm of academic research, programming languages were developed which embody the BDI model and thus diminish the amount of *ad hoc* coding, such as the PRS, dMARS, AgentSpeak and 3APL [DAM]. Also in the realm of academic research there exists a considerable set of agent programming frameworks which provide a set of tools for an easier and consistent development of BDI agents. Well known and used frameworks are JACK [DAM], Jadex BDI Agent System [LB05] and Jason [BH05]. On the realm of real-word applications, the BDI architecture was applied with a great amount of success. The most important application of this architecture is Georgeff's project for diagnose faults in the reaction control system of the Shuttle Discovery [Geo], although there are other examples [GR96, RG95] of the validity of this architecture in software implementations

### 2.2 Cons

Lets now analyse the main problems associated with the usage of the BDI architecture for the development of agent-based systems.

**Lack of information about resource bounds:** The BDI architecture uses only specific roles of Intentions [Wei99] to control the problem of acting under resource bounds. In our point of view this does not seems to be enough, since today's problems are characterised by a continuously growing of complexity and unpredictability, under severe resource bounds. We believe that the BDI model lacks an explicit internal representation of the means which an agent can count on

in order to decide which is the best way of acting on its environment without unnecessarily compromising future actions and also its overall performance.

**The problem of agents reconsideration:** This problem is the consequence of the relation that exists between an environment's change rate and the frequency and amount of reconsideration (deliberation) which BDI agents which inhabit this environment do. Since the frequency of reconsideration carried out by BDI agents must be defined in advance, these may miss important changes that occur in the environment or may carry out unfruitful reconsideration, in the case of not occurring significant changes on the environment. Therefore, we argue that both the deliberation processes and the instruments used in them should be dynamic and as much adaptive to environment's changes as possible.

**Lack of other human-like mental states:** The description of entities and the prediction of their behaviour under the rules of Dennet's Intentional Stance [Den87] is not limited to the usage of Beliefs, Desires and Intentions. Despite the fact that these mental states are flexible enough to model BDI agents for acting in a wide range of scenarios, there are some cases where the usage of other mental states would be appropriate [Nor04]. The usage of other mental states would provide to computer scientists and domain specific experts, respectively, new computational structures to implement software agents and new concepts for developing more refined models of BDI agents.

### 2.3 Comparing the pros and cons

From the above results we have concluded that the BDI architecture is still being a valuable concept to have in account when the development of software systems which requires Rational Agents is the case. However, this architecture suffers from problems which are far from being ignored and urge to be solved. In our opinion the first attempt to be made should fall upon the additional usage of accurate and specialised components to deal with the explicit resources which agents have, both for their usage in reconsideration processes and in all the functions which are part of the architecture's processing cycle [Wei99]. Wooldridge et al. also reached a similar conclusion some years ago [GPP<sup>+</sup>98], when they have proposed the development of specialised heuristics for dealing with the problems encountered in the BDI architecture. The components we are proposing are surely elements of a set composed by such heuristics.

## 3 Artificial Emotions, Agent's means and the BDI architecture

The usage of Artificial Emotions in the BDI architecture is not new [PT96]. However, the idea of using them as mechanisms for controlling the means of the agents for acting upon their environment is new, at least for what we are aware.

From the set of the three fragilities we have found in the BDI architecture, only one is not fully dependent on the addition of new concepts to the architecture: the *lack of other human-like mental states* can be solved by using the same approach used for deciding that Beliefs, Desires and Intentions would form the base of the BDI architecture. In other words, we can use the Intentional Stance [Den87] and add commonsense definitions of new mental states, such as emotions, and make them influence the BDI architecture through the commonsense understanding of the way they affect positively the reasoning performed by humans. For instance, we can define a new concept such as Fear which is something like an informational data structure which reports situations which an agent should avoid.

The problems of the *lack of information about resource bounds* and the *problem of agents reconsideration* requires a more refined approach, in terms of the usage of Artificial Emotions, mapping the way they are activated and which kind of plans of action they trigger. Consider, for instance, the problem of the amount and frequency of reconsideration done by an agent when it faces a threat. It seems acceptable that, instead of reactively going away, it should apply the maximum of its resources

and allocate them to the best solution-search algorithms it possesses, thus finding the best solution for escaping the threat. In fact, this is the role of driving attention and self-awareness updating that emotions play in humans. If an agent acts upon rules similar to the previous one, the reconsideration process may become dynamic (both on the amount of the means used and on the frequency of its execution) and adaptive, thus making the agent respond better to the changes occurring in the environment.

For the application of such clues that emotions present in humans give to us, on the perspective of a correctly usage of the available means, we have also to have a clear representation of these kind of means. Therefore, this leads us to the problem of the lack of information about resource bounds. This kind of information influences the Emotional State of the agent, but also is influenced by it, since the emotional state affects the way perception is achieved. The architecture we present tries to link all these concepts by giving an abstract idea of how they should be defined and how they should interact. Therefore, we will present only a conceptual architecture, leaving more low-level details to future implementations which will be conducted for analysing with more detail the validity of our approach, relatively to other architectures.

## 4 The Emotional-BDI architecture

We will now describe the proposed extension of the BDI architecture for supporting Artificial Emotions. We will start by introducing the basic concepts for such extension, followed by the conceptual architecture itself and also an abstract interpreter which maps the architecture's processing cycle.

### 4.1 Effective Resources and Effective Capabilities

The need for an explicit internal representation of the means that an agent has to execute upon its environment has been stress as the basis for the extension we are proposing. We argue that the concepts of Resource and Capability should form the basis of such internal representation.

**Capabilities:** These are abstract plans of action which the agent has available to act upon its environment. In terms of agent conception, a set of Capabilities can be seen as a dynamic Plan Library, where some of its plans are marked as being impossible to be considered either temporarily or permanently. Padgham and Lambrix introduced the idea of Capabilities in the BDI architecture [PL05]. Although they consider as Capability both the *ability* and *opportunity* and therefore an agent does not depend on any other resource to engage their execution, they enforce the idea that a Capability can be considered as a Plan which may not always be available.

**Resources:** These are means that turn Capabilities into plans of action which can be performed by the agent in its environment. Resources can be either physical (CPU time, disk space, available memory, etc.) or virtual (energy sources on a virtual world, other agents, etc), or both. They are also the source of the dynamic associated with the Capabilities, since the availability or unavailability of Capabilities depends on the availability of the required Resources.

Agents are not omniscient entities, thus only have information about a limited part of their environment and about itself. In particular this also applies to their Resources and Capabilities. Thus, what an agent can count on is its Effective Resources, which are a subset of all the available Resources which the agent is aware of. Capabilities, when instantiated to the required Effective Resources become Effective Capabilities.

The process of both Resources and Capabilities become effective is done by the following function which we decided to call *Effective Capabilities and Effective Resources revision function*, or *EC-ER-rf*. This function's behaviour can be described as follows:

- $EC-ER-rf : \mathbf{Cap} \times \mathbf{Res} \times \mathbf{Percept} \times \langle \mathbf{Cap} \times \mathbf{Res} \rangle \rightarrow \langle \mathbf{Cap} \times \mathbf{Res} \rangle$

This function revises and updates the Resources and Capabilities which in fact are available for an agent to use. The information that carries data needed for this process comes from the

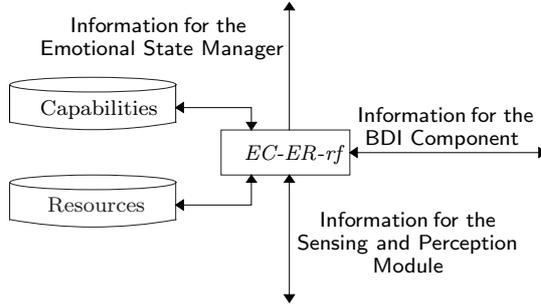


Figure 1: The Effective Capabilities Manager module.

data perceived by the agent and also from the various functions of the BDI architecture, which report the amount of Resources consumed, and which of the Capabilities were used. The result of this evaluation is forwarded to the mechanism responsible for the Emotional State update, so an agent become adapted to the new reality of the means it has.

A schema of the entities which forms this component of our architecture is presented in Figure 1.

## 4.2 Sensing and Perception

In order to interpret the meaning of the stimuli that occurs in an environment, and that is captured by the agents sensorial machinery which is inhabiting that environment, we need to equip these agents with mechanisms which have rules that binds these stimuli to concepts which "make sense" in these agents reasoning procedures.

For the architecture we are presenting, we propose a *Sensing and Perception Module*, a mechanism capable of obtaining the desired information from stimuli provided by an environment. This module we propose is composed of two sub-mechanisms responsible for performing different but complementing tasks: a *sensing filter* which deals directly with the stimuli captured by sensors of the agent; a *perception filter*, which attributes meaning to previous filtered data. In more detail we have:

***sensing filter***: this filter is responsible for extracting information directly from the stimuli captured by an agent's sensorial machinery. For that, it is connected to a repository of information extraction rules which can be directly applied to the information provided by the sensors. The rules which are applied are dynamically chosen, depending on the Effective Capabilities and on the Emotional State of the agent. Therefore, the kind of information extracted for further processing will emulated the role of driving attention that real emotions play in human reasoning.

***perception filter***: this filter is responsible for giving meaning to data previously processed by the *sensing filter*. In order to attribute a semantics for chunks of data, this mechanism uses a repository of *semantic association rules*. The result of the application of this filter is the production of concepts which can be forwarded for both *belief revision processes* of the agent, and also for *Effective Resources and Effective Capabilities revision*. Also the kind of rules that are available for being applied depends on the Effective Capabilities and Emotional State of the agent.

An illustrative example of how meaning is attributed to stimuli is presented in Figure 2. The full *Sensing and Perception Module* is depicted in Figure 3.

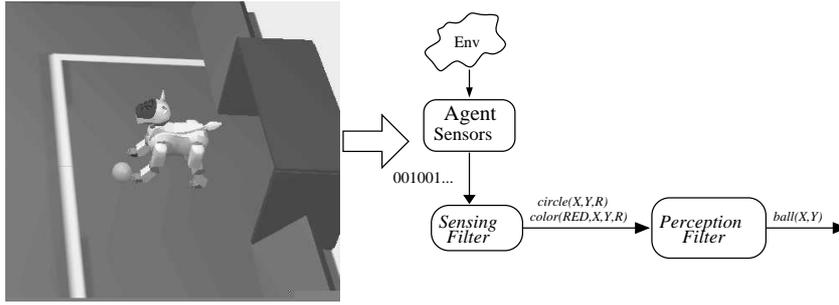


Figure 2: Identification of a stimuli as a ball of a robotic soccer match.

### 4.3 The Emotional State Manager

The Emotional State Manager is the component which is responsible by controlling all the Resources and Capabilities usage on all the information processing phases of the architecture we are proposing. It should then be composed by an internal structure which exhibits an adaptive behaviour, depending on the changes that occur in the environments where agents are standing. Thus, in order to avoid falling into a problem near to the problem of reconsideration detected in the BDI architecture, we do not present any fixed structure which this Emotional State Manager should incorporate. Instead, we only propose the following conduits that we consider to be fundamental for this kind of component:

- it should base itself on a well defined set of Artificial Emotions which relates efficiently the kind of tasks the agent has to perform, and the way changes on the environment affect the internal state of the agent. For instance, if the primary goal of an agent is survival, then Fear should be present in the Emotional Set so the agent can be alerted when its survival changes diminish considerably.
- there should exist various information extraction functions for each of the Artificial Emotions which are elements of the set of Artificial Emotions. This dues to the fact that there are various sources of stimuli which change the same emotion in different ways. For instance, for a Fireman which is fighting a fire, the Fear should be elicited both by fear proximity and changes on the wind, which may change the fire's direction.
- there should exist a decay rate for each of the Artificial Emotions on the set of Artificial Emotions, but always depending on the state of the emotional eliciting source. If the Fireman continuously notices fire near to him, he will continue in a Fear emotional state. However, if it escapes the fire, the Fear rate will diminish with time.

An interesting model for Emotional State managing can be seen on the work of Oliveira and Sarmiento [OS03].

### 4.4 The conceptual architecture

The Emotional-BDI architecture we are proposing is an extended version of the classic BDI architecture with the addition of three new components – the *Sensing and Perception Manager*, the *Effective Capabilities and Effective Resources revision function* and the *Emotional State Manager* – and the definition of interactions between only these new components and between these new components and the old ones. Since we have already introduced the new modules of this architecture, we can present the overall conceptual Emotional-BDI architecture, as is shown in 4.

The interactions between the new components have already been described in Sections 4.1, 4.2 and 4.3. We need now to describe the interaction between these new components and the ones which are present in the original BDI architecture.

The new components of the BDI component of our architecture have the same name as the original ones plus an apostrophe concatenated at the end of their name. We denote the Beliefs

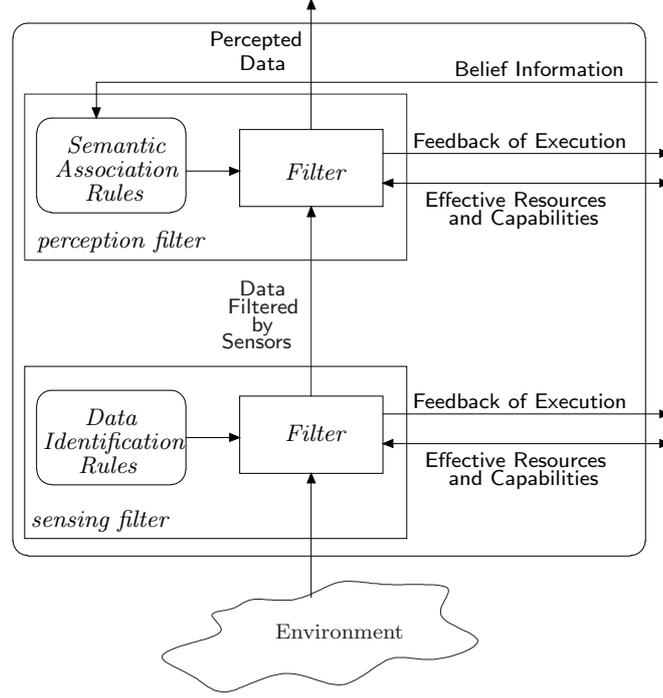


Figure 3: The Sense and Perception manager.

by **Bel**, the Desires by **Des**, the Intentions by **Int**, the Effective Resources by **ERes**, the Effective Capabilities by **ECap**, the information perceived by **Percept**, and finally the Emotional State by **EState**, respectively. These new components are the following:

- $brf' : \text{Bel} \times \text{Percept} \times \text{ECap} \times \text{ERes} \times \text{EState} \rightarrow \text{Bel}$

This function updates the Beliefs of the agent based on already existing Beliefs and on the new information perceived by the *Sensing and Perception Module*. The computational resources spent and the revision algorithms used are defined by an evaluation made on the base of the Effective Resources, Effective Capabilities and the Emotional State of the agent.

- $gen - options' : \text{Bel} \times \text{Int} \times \text{ECap} \times \text{ERes} \times \text{EState} \rightarrow \text{Des}$

This function is responsible by means-ends reasoning, the process of recursively elaborating hierarchical plans which defines progressively more specific Intentions, until these become satisfiable by the execution of actions. This function is also controlled by both the Effective Resources/Capabilities and the Emotional State which enables the use of distinct algorithms, making it adaptive to the changes of the environment.

- $filter' : \text{Bel} \times \text{Des} \times \text{Int} \times \text{ECap} \times \text{ERes} \times \text{EState} \rightarrow \text{Int}$

This function updates the Intentions based on already existing Intentions, Beliefs and Desires. Through it, Intentions should be dropped if they were already achieved or if they will never be achievable, and should also retain the ones which are considered to be fruitful in the future. Once again, distinct kinds of algorithms may be used, depending on the Emotional State of the agent.

- $execute' : \text{Int} \times \text{ECap} \times \text{ERes} \times \text{EState} \rightarrow \text{Action}$

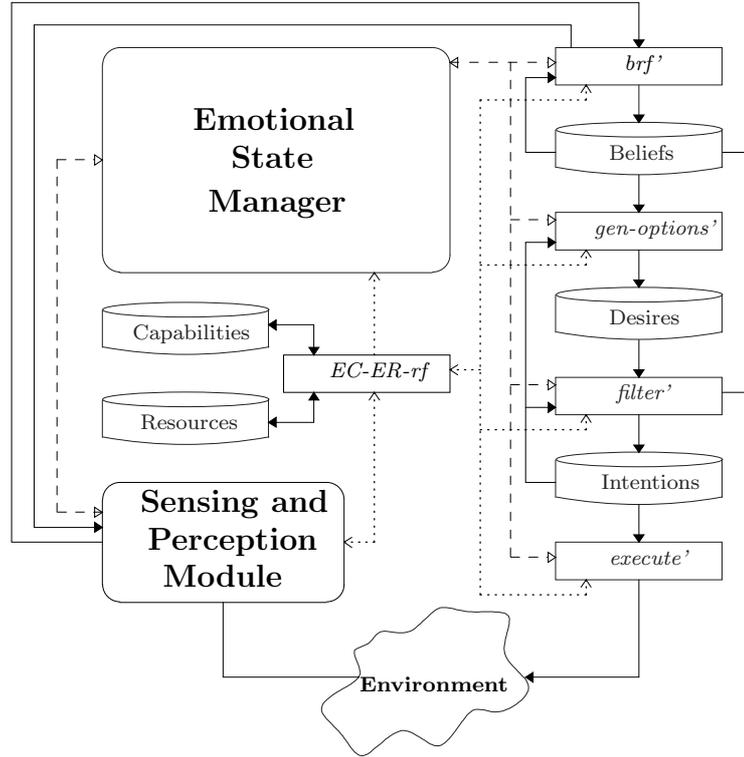


Figure 4: The conceptual Emotional-BDI architecture.

This function selects an Intention to be fulfilled by the direct execution of an action. The action and the Intention selection can be a dynamic process, depending on the available Effective Resources, Effective Capabilities and the Emotional State.

Besides having the functionality that was described, each of the above functions also informs about the subset of Effective Resources and Effective Capabilities used so that the *EC-ER-rf* function can update correctly what is left to continue to be used. These functions also inform the Emotional State Manager about their overall performance, in order to update the Emotional State of the agent. These analyses can contain, for instance, the number of solutions obtained, the time taken to obtain them, which algorithms were used, if errors were detected, etc.

Based on the schematic view of the conceptual Emotional-BDI architecture and on the properties of each of the above functions, we can also present an abstract interpreter for our architecture. As happens with the schematic view, this interpreter is an extension of the original one [Wei99].

The abstract interpreter we have just presented (Figure 5) is very similar to the original one, except that the various components are initialised separately and all the functions include as input the Effective Capabilities, the Effective Resources and the Emotional State of the agent. Moreover, this interpreter enforces the idea that our architecture works as one component only, as stated in [Dam94].

## 5 Conclusions and future work

In this report we have presented the Emotional-BDI architecture, an extension to the original BDI architecture with the addition of an internal representation about what means an agent can count

```

Emotional-BDI-interpreter
initialize-emotional-state-manager(ES);
initialize-sensing-perception-module(ES);
initialize-EC-ER-rv(C,R);
initialize-bdi-state();
repeat
  options := option-generator(sensor-input,B,D,I,EC,ER,ES);
  selected-options := deliberate(options,B,D,I,EC,ER,ES);
  execute(I,EC,ER,ES);
  get-new-external-events(sensing-perception-filter,EC,ER,ES);
  drop-successful-attitudes(B,G,I,EC,ER,E);
  drop-impossible-attitudes(B,G,I,EC,ER,ES);
end repeat

```

Figure 5: The Emotional-BDI abstract interpreter.

on, and an Emotional State capable of controlling the usage of these means in order for the agent to perform adaptively to its environment. We believe that both the components fill the gaps which we argued that exist in the BDI architecture.

Our future work will be on the implementation of Emotional-BDI agents in dynamic and complex environments like PyroSim [OS03], but also in static environments like for instance a simple maze with energy sources and obstacles, possible of being developed, for instance, in freeBots<sup>1</sup>.

Parallel to the implementation issues we have just referred, we will try to extend the existing logical frameworks to support the concepts present in the Emotional-BDI architecture (and the implementation issues themselves) with the intent of avoiding unpleasant *ad hoc* coding techniques and to have the possibility to verify properties of our architecture.

## References

- [BH05] Rafael H. Bordini and Jomi F. Hübner. Jason – a java-based agentspeak interpreter used with saci for multi-agent distribution over the net. <http://jason.sourceforge.net/>, March 2005.
- [BIP88] Michael E. Bratman, David Israel, and Martha E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.
- [DAM] Daniela Demergasso Davide Ancona and Viviana Mascardi. A survey on languages for programming bdi-style agents.
- [Dam94] Antonio R. Damasio. *Descartes' Error: Emotion, Reason and the Human Brain*. Grosset/Putnam, 1994.
- [Den87] D. C. Dennett. *The Intentional Stance*. 1987.
- [Geo] Michael Georgeff. The mind of the machine. <http://www.nasa.gov>.
- [GPP<sup>+</sup>98] Michael P. Georgeff, Barney Pell, Martha E. Pollack, Milind Tambe, and Michael Wooldridge. The belief-desire-intention model of agency. In Jörg P. Müller, Munindar P. Singh, and Anand S. Rao, editors, *ATAL*, volume 1555 of *Lecture Notes in Computer Science*, pages 1–10. Springer, 1998.
- [GR96] Michael P. Georgeff and Anand S. Rao. A profile of the australian artificial intelligence institute. *IEEE Expert*, 11(6):89–92, 1996.
- [LB05] Alexander Pokahr Lars Braubach. Jadex bdi agent system. <http://vsis-www.informatik.uni-hamburg.de/projects/jadex/>, May 2005.

<sup>1</sup>freeBots is a robot simulation framework that is being developed at DCC-FC & LIACC, by Luís Damas.

- [Nor04] Emma Norling. Folk psychology for human modelling: Extending the bdi paradigm. In *AAMAS*, pages 202–209. IEEE Computer Society, 2004.
- [OS03] Eugenio Oliveira and Luís Sarmento. Emotional advantage for adaptability and autonomy. In *AAMAS*, pages 305–312. ACM, 2003.
- [PL05] Lin Padgham and Patrick Lambrix. Formalisations of capabilities for bdi-agents. *Autonomous Agents and Multi-Agent Systems*, (10):249–271, 2005.
- [PT96] Lin Padgham and Guy Taylor. A system for modelling agents having emotion and personality. In Lawrence Cavedon, Anand S. Rao, and Wayne Wobcke, editors, *PRICAI Workshop on Intelligent Agent Systems*, volume 1209 of *Lecture Notes in Computer Science*, pages 59–71. Springer, 1996.
- [RG95] A. S. Rao and M. P. Georgeff. BDI-agents: from theory to practice. In *Proceedings of the First Intl. Conference on Multiagent Systems*, San Francisco, 1995.
- [vdHW03] Wiebe van der Hoek and Michael Wooldridge. Towards a logic of rational agency. *Logic Journal of the IGPL*, 11(2):135–159, 2003.
- [Wei99] Gerhard Weiss, editor. *Multiagent systems : a modern approach to distributed artificial intelligence*. MIT Press, Cambridge, Mass., 1999.