When can knowledge-based systems be called agents?

Jaime Simão SICHMAN*
Yves DEMAZEAU
Olivier BOISSIER

LIFIA — Institut IMAG
46, Avenue Félix Viallet
38031 — Grenoble — France

jaime@lifia.imag.fr
demazeau@lifia.imag.fr
boissier@lifia.imag.fr

Abstract

Distributed Artificial Intelligence (DAI) has become an established research domain in recent years. As opposed to classical AI approaches, where the metaphor of intelligence is based on an *individual human behaviour* and the emphasis is put on knowledge representation and inference methods, the metaphor used in DAI is based on *social behaviour*, and its emphasis is put on actions and interactions. An "intelligent" social behaviour can arise either in the case of coarse-grained "intelligent" members of the society (called cognitive agents) or in the case of fine-grained "unintelligent" members of the society (called reactive agents). Interests in this field may be divided into two major areas: *Distributed Problem Solving* (DPS) and *Multi-Agent Systems* (MAS). One can consider the DPS approach as one phase of the MAS approach. In either of these areas, the entities are called *agents*, and not knowledge-based systems as in classical AI approaches. The goal of this paper is to put together some ideas and definitions of the MAS field, and to give some answers to the question "When can knowledge-based systems be called agents?". To illustrate, we will use as an example a model of a cognitive agent being currently developed at the LIFIA which deals with conflicting data.

Keywords: Multi-Agent Systems, Distributed Artificial Intelligence, Internal Models of Cognitive Agents.

*On leave from Escola Politécnica da Universidade de São Paulo, Depto. Eng. Computação e Sistemas Digitais, Av. Prof. Luciano Gualberto, 158 - trav. 3, CEP: 05508, São Paulo, SP, Brazil, email: jssichma@brusp.bitnet

1 Distributed Artificial Intelligence

Distributed Artificial Intelligence (DAI) has become an established research domain in recent years. As opposed to classical AI approaches, where the metaphor of intelligence is based on an *individual human behaviour* and the emphasis is put on knowledge representation and inference methods, the metaphor used in DAI is based on *social behaviour*, and its emphasis is put on actions and interactions.

Such an approach is desirable when one has to solve large and complex problems, which require knowledge from several different domains, and which sometimes involve physically distributed data managing. From a computational point of view, the advantages of such an approach is the control of complexity, graceful degradation, support for evolution and replicability [CHAN81].

Interests in this field may be divided into two major areas: *Distributed Problem Solving* (DPS) and *Multi-Agent Systems* (MAS). In both cases, we use the expression *agent*, and not knowledge-based system or another expression to describe the entities which take part in the problem solving activities. Within the MAS area, we have also two different approaches, regarding the granularity of the agents: fine-grained, very simple and "unintelligent" *reactive agents* and coarse-grained, "intelligent" *cognitive agents*.

We will use a very simple example in order to illustrate our ideas all throughout this paper. The example is the following:

Yves wants to paint the Brazilian flag on the wall of his office, but he has neither knowledge nor skills to accomplish this task. Nevertheless, there are five people in the next room, with different capabilities: Jaime, who knows the colors of the Brazilian flag and its format, Olivier, who is a very good painter, Matthew, who is a great jazz singer, Philippe, who is a good artist and Patrick, who is also a good painter. All of them have the capabilities of working together in a same society, that is, they can perceive their environment (they can see), communicate with one another (they can speak the same language) and act (they can use their hands).

In the context of our example, Yves will play sometimes the role of the system's designer, and sometimes the role of the final user, but he does not take part in the problem solving activities. Therefore, he should not be considered as an agent.

In the following subsections, we will show the differences between a DPS approach and a MAS approach. We will also present some historical remarks and the main problems involved in DAI research. Section 2 presents the main concepts of MAS, notably those referring to cognitive agents. We will show then in section 3 the additional features that must be incorporated in a knowledge-based system so that one can call it an agent. Finally, we will illustrate these concepts in section 4, when we will present a model of a cognitive agent that is being currently developed at the LIFIA and which deals with conflicting data, followed by our conclusions.

1.1 A Distributed Problem Solving Approach

The main difference between a DPS approach and a MAS approach may be explained by the use of the term *autonomous*. In a DPS approach, the agents are designed *in order to solve some particular problem*. One can think of DPS as a kind of extension or evolution of the techniques used in distributed systems applied in a AI context. This means that a priori the agents cannot be used to solve any other similar problems. In other words, there is a *closed world assumption* regarding these agents. Schematically, the DPS approach might be represented in the following way:

(problem to be solved) $\Longrightarrow$ (design of the agents) $\Longrightarrow$ (problem solving) $\Longrightarrow$ (solution)

A DPS approach to solve our example would be the following:

Yves knows that Jaime is the only one who knows the colors and the format of the Brazilian flag. He asks him, therefore, to instruct Philippe how to draw the flag. Once the format is drawn, he asks Olivier and Patrick to start painting the flag, according to the colors Jaime has told them. Matthew does not take part in the work. When Yves returns, the wall is painted and he knows exactly who has done what and when.

The important points in this approach are the following:

- although working cooperatively, there is no need for the agents to represent *explicitly* what are the other's skills and goals. This is *implicitly* represented by the designer;
- the description and decomposition of the task is *in most cases fully decided by the designer*. This is an extreme case, but even if there is some dynamic task decomposition, the methods used are strongly dependent of the application domain;
- *in most cases, conflict will not arise* if the task has been properly divided by the designer. Once more, this is a critical case, but if conflict arises, it is strongly related with the application domain;
- even if they can communicate, there is *no need of a complex conversation* to achieve the goal;
- new agents *cannot be added dynamically* to the society. This means that such a system could not be considered as an open system [HEWI86].

1.2 A Multi-Agent System Approach

On the other hand, a MAS approach is concerned with the activity of *autonomous* agents in a multi-agent world. By autonomous, we mean that the agents have their own existence, despite of the prior existence of any problem to be solved. The agents coexist in a common environment and each one can collaborate with the others in order to achieve a common goal [DEMA90]. Schematically:

(existing agents) $\Longrightarrow$ (problem to be solved) $\Longrightarrow$ (problem solving) $\Longrightarrow$ (solution)

A MAS approach (using a cognitive agent model) to solve our example would be the following:

Yves presents the task to everyone in the room and goes away. Then, the people in the room begin to communicate in order to establish the division of the task, constraints, etc. Once decided, they start doing their work. From this point on, the process is similar to the DPS approach. When Yves returns, the wall is painted, but he has no idea about the execution of the task.

The important points in this approach (using a cognitive agent model) are the following:

- *the decomposition of the task is done by the agents*, and not by the user or designer. At the limit, there may be a dynamic reorganization, that is, the agents may decide that they should *change their behaviour* in order to better accomplish their task;

- the agents are *autonomous*, that is, that they can have their own local goals. In our example, Olivier may refuse to cooperate with the others if he thinks that it is more important for him to go home and take care of his daughter. Therefore, *conflicts usually arise*, due to the existence of local and global goals. Moreover, a *complex conversation* must be held in order to establish the role of each agent in the problem solving activities;

- if we consider that Matthew has some knowledge about American jazz groups and the others can sing and play some musical instruments, the same agents will be able to form a jazz group. Therefore, agents are able to *solve other problems*, if they have the required skills to accomplish it;

- *a MAS is an open system*, so any agent may enter or leave the society whenever it wants to. In our example, if Olivier leaves the society, another way must be found to accomplish the task. For instance, Patrick can be required to paint the whole flag alone. If a new agent enters the society, the agents must incorporate its capabilities and skills. This is accomplished by the maintenance of an *explicit representation of the other agents' skills and goals*;

- *the environment may change*, and the agents must incorporate these changes in their internal world model. In our example, if Yves asks again to paint a Brazilian flag after the task has been finished, Jaime may tell him that there is already a flag painted on his wall. Mobile robots which navigate in different and unknown rooms are another example of agents which must interact with a dynamic world.

From the former example, one can notice that the DPS approach may be considered as one phase of the MAS approach.

1.3 Historical Remarks

Historically, the first steps in DAI research were related to the DPS approach. Some results and tecniques from social psychology, social sciences and ethology were then incorporated into the domain, which has lead to the appearance of the MAS approach.

The first attempts to solve problems cooperatively can be found in the seventies [FERB91a]. One of the first of these attempts was the HEARSAY-II project [ERMA80], a speech understanding system which introduced the Blackboard model [ENGE88]. At the same time, some computational models, like Lenat's beings [LENA75] and Agha and Hewitt's actors [AGHA85], were developed in order to deal with problems such as shared resources, complex control and the arising of complex behaviours from very simple interactions.

Later, in the eighties, Konolige and Nilsson were the first researchers who pointed out the importance of explicitly representing in one agent the other's knowledge, in a planning context [KONO80]. Fox has discussed the use of human organizations models in cooperative work [FOX81] and Smith developed the contract net protocol [SMIT80], a negotiation protocol based on a market structure. By this time, object-oriented techniques [STEFb86] were becoming popular. Brooks designed his "reactive" robots, where the main idea was to use simple and small architectures to design an autonomous robot [BROO86]. Hewitt has introduced the concept of open systems [HEWI86], that is crucial for MAS research. Finally, some the first environments were developed to test and implement some of these ideas, like Corkill's DVMT [CORK83], Cardozo's DPSK [CARD87] and Gasser's MACE [GASS87].

1.4 Main Problems in Distributed Artificial Intelligence

The main problems in DAI may be basically divided into five different areas, either on a DPS or on a MAS approach [BOND88]:

- *description, decomposition and allocation of tasks:* how can one easily describe and decompose a complex task into subtasks, statically or dynamically, how are these subtasks to be allocated and in which order they should be performed;
- *interaction, languages and communication:* which primitives should a communication protocol present in order to express the semantic concepts arising in a cooperative work;
- *coordination, control and coherent behaviour:* how can one assure a coherent global behaviour in a set of agents, each one with its own skills and goals, how should the control of such a system be designed;
- *conflict and uncertainty:* as no agent has total information about its environment, how can the conflicts that may arise be solved, and how should uncertain, incomplete data be dealt with in order to guarantee coherent results;
- *programming languages and environments:* from a computational point of view, which are the programming languages that should be used in such systems, and what are the requirements of an environment to make testing these policies possible.

Obviously, some of these problem have been investigated by researchers in other areas, like distributed systems and logics, for a long time. One can consider that DAI research is one of the first attempts to put these problems together in the same context.

2 Multi-Agent Systems

As we have said before, research in MAS may be divided into *reactive* and *cognitive* agents. It is out of our scope to discuss in detail reactive agents in this paper. In the following subsections, we will present a definition for the term agent, describe briefly the differences between reactive and cognitive agents and present some features of cognitive agents.

2.1 What is an agent?

Currently, one cannot find in the literature a universal and unique definition of the term agent. We will adopt Ferber's definition [FERB91a], which is very general:

An agent is a **real** *or* **virtual** *entity which is emerged in an* **environment** *where it can take some actions, which is able to* **perceive** *and* **represent** *partially this environment, which is able to* **communicate** *with the other agents and which possesses an* **autonomous behaviour** *that is a consequence of its observations, its knowledge and its interactions with the other agents.*

2.2 Reactive Agents

Reactive agents [FERB91a] are based on *biological/ethological organization models*, such as a colony of ants. The main idea is that a complex "intelligent" behaviour may emerge from a collection of very simple entities. A single ant does not seem to be very intelligent, but a colony of ants acts as an intelligent entity, for instance when searching for food. In a certain way, this approach is similar to the connectionist approach, although it is not possible to change the behaviour of a neural network once it is trained, and reactive agents do change when their environment does so.

A reactive agent behaves in a *stimulus-response mode*, that is, *it does not have a memory of the actions carried out in the past nor any prevision of the actions to be taken in the future.* Everything it knows about the actions and behaviours of the other members of the society is perceived by changes in the environment. For instance, the ants usually follow a chemical trace left by the others when searching for food. Therefore, *there is no high-level communication model nor an explicit representation of the environment or of the members of the society.* Societies of reactive agents have usually *a great number of members.*

Some recent research in this area are Eco Problem Solving [FERB91b], Coordination Patterns [DEMA91] and Situated Actions [ROSE86].

2.3 Cognitive Agents

Cognitive agents are based on *social organization models*, in the sense of human societies (groups, hierarchies, markets) [FERB91a]. Within the agents, *there is an explicit representation of the environment and of the members of the society.* They can *reason about the actions taken in the past and plan the actions to be taken in the future.* From a society point of view, it is generally composed of *a small number of members.*

We will now describe some concepts and ideas related to a cognitive agent. We will refer to our example to illustrate them, and we will link these concepts with the problems in DAI research described in the previous section.

2.3.1 Possible Social Behaviours

One can classify the possible behaviours of such agents in an environment according to two criteria [DEMA90]:

- *locality of the task:* global or local. A global task is a task which concerns all the agents, and a local task concerns only one agent;
- *capability of performing the task:* able or unable. An agent is able to perform a task if he has the required skills to accomplish it.

According to these criteria, the following behaviours may occur:

- *cohabitation:* an agent can solve its own local task and does so. One can note that an agent is not obliged to cooperate with the others simply because they are emerged in the same environment. In our example, Matthew can start singing while the others are involved in drawing and painting the Brazilian flag;
- *cooperation:* an agent cannot solve its own local task (either it lacks sufficient skills or it cannot achieve the performance required), and so asks the others to cooperate with it. Regarding our example, if Yves had asked only Olivier to paint the Brazilian flag, Olivier could have asked Patrick to help him in order to finish the task more quickly;
- *collaboration:* an agent can solve a global task alone. If there are many agents, an election mechanism must be carried out. In our example, if Yves does not designate anyone to draw the flag, Philippe may suggest that he has the skill to accomplish it;

- *distribution:* some global goals must be accomplished by a collective action. This is our example presented when we have described the MAS approach. Neither of the agents can achieve the global task by itself, so the task must be divided and allocated according to some criteria.

2.3.2 Internal and External Descriptions

Referring to the problem of decomposing and allocating a task in a MAS approach, one question must be addressed: *how the agents know what are the capabilities and/or skills of the other members of the society?* This question is crucial for a cooperative/distributed behaviour.
The answer is that *every agent must have an external representation of all agents of the society.* We adopt the term external description to refer to a compact representation of the capabilities of each agent. For instance, in our example, an external description of Jaime could be:

```
Name: Jaime
Knows: Brazilian flag format
Knows: Brazilian flag colors
Committed: drawing a Brazilian flag in Yves' office
```

We are not going to discuss the notion of commitment in this paper. The basic idea is that *static* information (like knowledge) and *dynamic* information (like a commitment in a problem solving activity) are present in this description. On the other hand, an *internal description* shows how the agent behaves internally, that is, how Jaime represents and manipulates his knowledge about the Brazilian flag. This kind of description may be explicitly represented or not. In the case it is, it enables the agent to reason about its own internal behaviour. An interesting point is that two different agents may have the same external description, even if they are organized internally in a different manner. In our example, if Olivier is left-handed and Patrick right-handed, they actually paint in a different way, but their external description may be the same. The scope of this external description is an open research topic, and it is closely related with the communication protocols, to be described next.

2.3.3 Communication Protocols

As we have said in the beginning of this paper, DAI research stresses the importance of actions and interactions between agents and the environment. We can divide an agent's interface to the world into *perception* and *communication.* Perception means an intrinsic capability of an agent towards its environment, like seeing. One can think of a camera in a mobile robot. Communication means changing data and knowledge between agents.
As this communication is very complex, communication protocols ought to be defined in order to express the semantic concepts involved in cooperative problem solving. One proposal adopted by many researchers is the use of speech acts [CAMP90]. Briefly, this theory proposes that each utterance has to be viewed as a kind of action, similarly to the actions carried out on the environment. Therefore, each speech act carries some propositional content plus an illocutionary force, like demanding information, requesting some order, warning etc. As an example, the statements "Paint this flag" and "Could you paint this flag" have the same propositional content, but the first looks like an order, and the second like a request. One open research topic is to what extent these concepts are to be present in the primitives or in the content of the message itself.

Another interesting idea which is being currently used in the field is the concept of *law-governed systems* [MINS89]. Its main idea is that protocols for distributed systems must be enforced, in order to avoid their violation. So, all the communication is governed by laws that have a predefined sequence of rules. As an example, the rules of the road may be considered as a law that everyone knows. When a driver uses his turn signal before turning right, he does not actually communicate with the others directly, but the others are aware of his intention to turn to the right because everyone knows the rules of the road.

2.3.4 The Control of an Agent

As we have seen so far, an agent can carry out the following activities:

- *perceive its environment*;
- *communicate with other agents*;
- *plan how to achieve a task*;
- *execute some problem solving activity.*

It is clear that a complex control policy must be present in order to achieve a coherent global behaviour. The control may be *centralized* or *decentralized.* In our example, a centralized control could be a situation when Jaime, who has the knowledge about the problem, would tell everybody what to do and when. In some sense, this policy is similar to some highly hierarchical human organizations. On the other hand, a decentralized control means that everyone in the society may help in establishing the role that the agents may play. In our example, it could happen that Olivier and Patrick ask the others what the colors of the flag are and Philippe asks what he is supposed to draw.

Like the definition of the term agent, a unique control policy does not yet exists. Moreover, there are two different dimensions of control: *agent control* and *society control.* Agent control addresses how an agent should internally organize its activities. On the other hand, society control addresses how to organize the set of agents and how to control their interactions. Boissier [BOIS92] tries to clarify these concepts, when he investigates how the agent and social control could be hierarchically splitted, when applying a MAS approach to computer vision problems. Levi's work [LEVI89] is another attempt to formalise some concepts in the agent control. From a Software Engineering point of view, Shoham [SHOH92] proposes a framework for what he calls agent-oriented programming, defining the internal states and the cycle of a generic agent interpreter. Ishida [ISHI92] addresses social control and proposes a model for organization-centered problem solving, where some control structures based on feedback between problem solving activities and organization self-design are defined.

2.3.5 Some existing systems

Altough MAS is a relatively recent field, there are already a significant number of systems which have been developed using its techniques. We can cite works in natural language processing [STEFa92], intelligent buildings and learning tutors [COEL92]. From an industrial point of view, the Esprit II project ARCHON (Architecture for Co-operative Heterogeneous On-Line Systems) [OLIV91] is trying to integrate in a MAS perspective some pre-existing knowledge-based systems so that they can cooperate in a problem solving activity. The researchers have developed a high-level layer where they represent externally the other agents and achieve communication, using a Blackboard model.

3 From Knowledge-Based Systems to Agents

At this point, we are already able to distinguish the features one has to incorporate in a knowledge based-system so that one can call it an agent (in a cognitive sense). These features are shown in figure 1. We adopt here the term knowledge-based system (KBS) in the sense of a classical expert-system, that is, a system with an inference engine and domain specific knowledge represented in a proper formalism, like production rules. Moreover, *we consider a knowledge-based system as an isolated problem solver.*

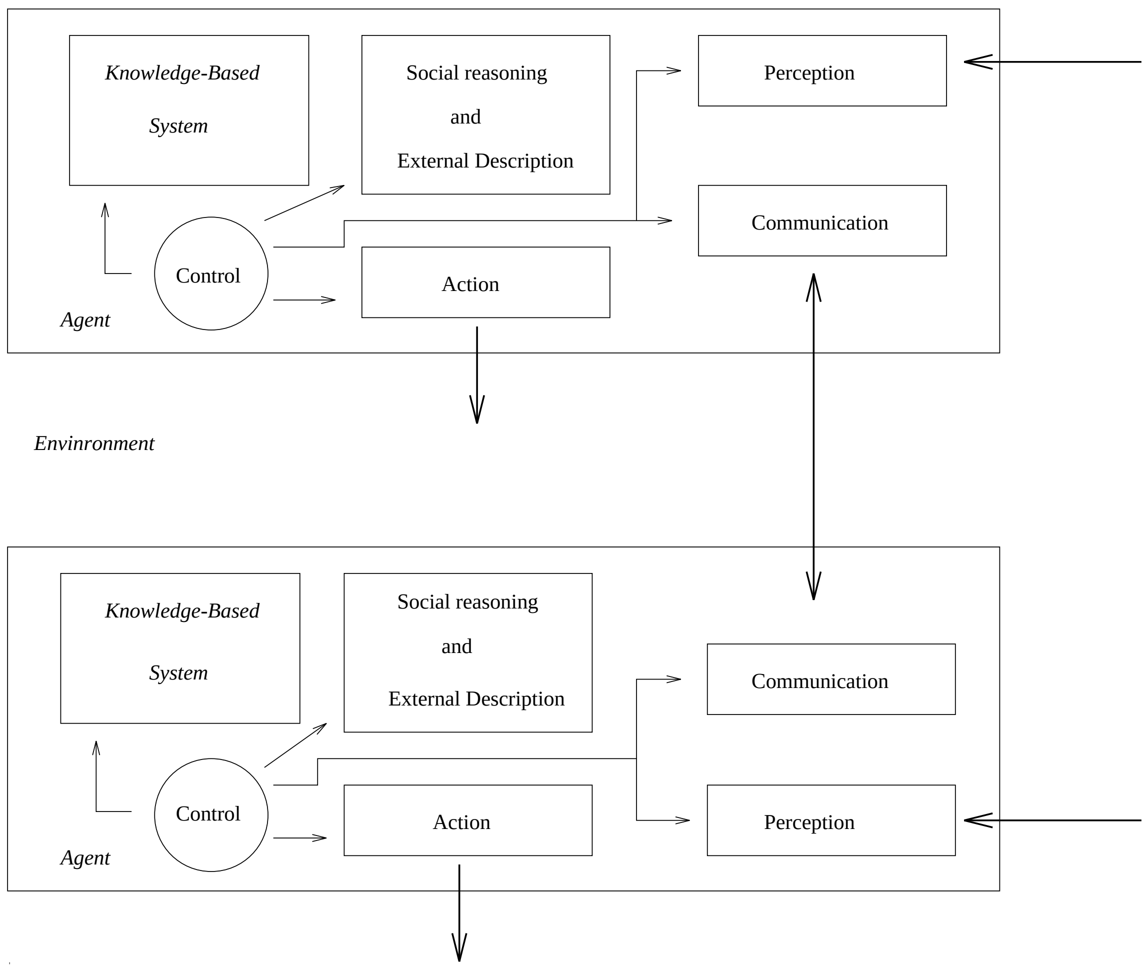

Figure 1: From Knowledge-Based Systems to Agents

Basically, the following additional features must be present in a KBS so that one can call it an agent:

- *perception capabilities:* an agent must perceive its environment, in the sense that it must interpret changes in the latter. This is not only related to a human-machine interface, as it is usual in KBS. The environment may change as a consequence of the actions of the other agents;
- *communication capabilities:* an agent must be able to communicate with the others. This is not usual in KBS. Moreover, the term communication is employed here in a broader sense, meaning that the agents can change not only data, but pieces of knowledge and partial plans as well;
- *action capabilities:* an agent must be able to act in his environment, as a result of its problem solving activities. Once more, this feature has a broader sense than merely a human-machine interface;

- *social reasoning:* an agent must be able to reason about the other agents' activities. This is accomplished by an internal representation of the other members of the society and the related reasoning mechanisms;

- *multi-layered control structure:* an agent must decide when to perceive, communicate, plan and act. So, a complex control structure must be provided in order to make possible the activation of these different features.

We are not employing the term KBS to design any implementation technique. At this level, our social reasoning module may be implemented using production rules and an inference engine. What we intend to clarify is the difference of meaning of these two expressions, *agent* and *KBS*, when applied in a model level.

4 An Example of a Cognitive Agent

The model of a cognitive agent [DEMA90] presented in figure 2 is a very generic one. Its purpose is not to specify what an agent architecture should be, but to illustrate how its behaviour can be described.

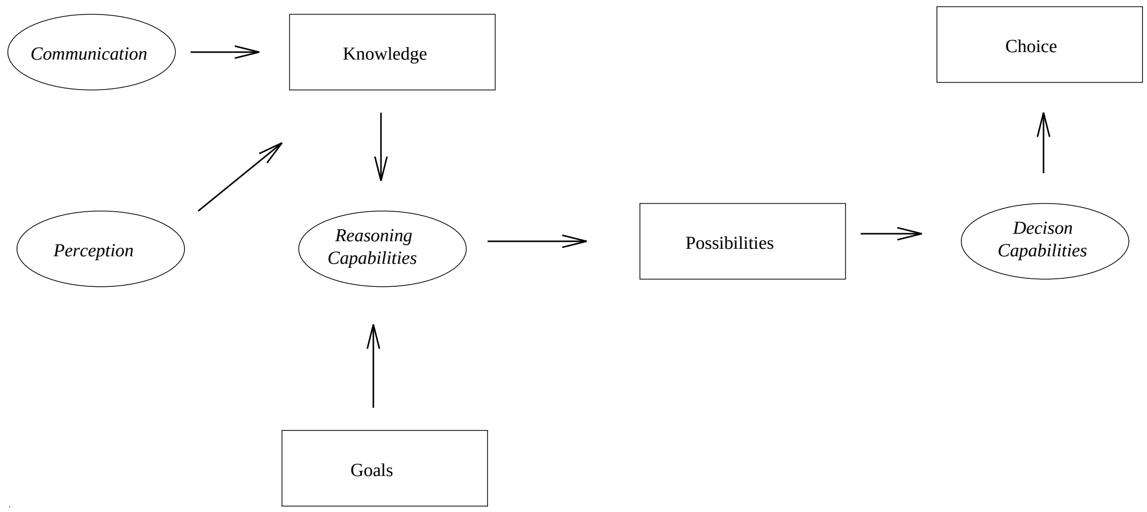

Figure 2: A Cognitive Agent Model

We will call *knowledge* the agent's representations of the world and of the problem to be solved. This knowledge may be acquired through *communication* or *perception*. We also represent the *goals* an agent has to accomplish. These goals may be encoded in the algorithms or dynamically acquired, by communication with the other agents or by observing the changes in the environment and reasoning about the other agents' behaviour. Once a goal is established, and taking into account its current knowledge, an agent uses its *reasoning capabilities* in order to plan how it could accomplish this goal. There might be more than one solution to this process, we will call them *possibilities*. Finally, a *choice* of the best possibility is made by using its *decison capabilities*.

At the LIFIA, we are currently involved in the development of a multi-agent environment to test the behaviour of a community of such agents. This environment is based on the DPSK+P system [CARD91], an object-oriented distributed problem solving kernel, which has a graphical monitoring tool [SICH91]. We have already developed a high-level interagent communication protocol [BERT92] which is being implemented. The next step is to refine this generic model in order to investigate some particular aspects of control, conflict detection and resolution, and possible enhancements in the communication protocol.

One original approach which is being addressed by our team is a dual representation of knowledge acquired *by perception* and *by communication*. Some researchers, like Galliers [GALL89] [GALL90] and Gaspar [GASP90], have pointed out that conflicts may have a positive role in a MAS, and propose some mechanisms of belief revision in order to avoid incoherence. This process may be considered as a type of *learning* method, as an agent may change and evolve its current knowledge about the environment and the other members of the society. As an agent has multiple sources of knowledge acquisition, conflict may arise and must be dealt with. We are particularly interested in studying how to represent and incorporate knowledge acquired from different modalities, such as perception and communication.

Our thesis is that *not every conflict that arises must be resolved.* Regarding our example, we will show two different situations of conflict:

Situation 1: While painting the Brazilian flag, Olivier tells Patrick that he thinks the USA is going to win the next Olympic Games. Patrick disagrees, and tells Olivier that he thinks that Russia is going to win.

Situation 2: While painting the Brazilian flag, Olivier tells Patrick that the inner circle should be painted green, and Patrick disagrees, telling him that he is sure that it should be painted blue.

Considering the problem solving activities, it is obvious that the first situation does not pose any problem regarding the cooperative task. On the other hand, the second situation is closely related to the task beeing executed, and must be resolved. Some possible solutions could be:

- *Patrick tries to convince Olivier that the right colour is blue;*
- *Patrick asks Jaime to show Olivier that he has misunderstood some information;*
- *if there is a model of the flag, Patrick may ask Olivier to look at this model.*

On the other hand, as information perceived and communicated are from different modalities, it should be interesting to represent and treat them separately. The originality of our approach resides in this dual representation. In many and most circumstances, humans tend to trust more in perceived than in communicated information. If someone tells a person that the green wall in front of him is red, he would not believe (and actually would incorporate internally that his interlocuter is color blind). Nevertheless, there are some situations where this does not happen. If someone is working at his office and everyone in the same floor starts running and telling that the other side of the building is burning, probably this person will trust more in this communicated information (actually, he is not able to see the fire) and will start to run.

Our proposal is to organize internally the knowledge acquired by an agent in two different knowledge bases: a communication knowledge base (K_c) and a perception knowledge base (K_p). Each knowledge base must be coherent internally, but the set of all knowledge an agent possesses must not. So, statements representing "I think the wall is green" and "My neighbour thinks that the wall is red" could be represented at the same agent, each one on a different base.

Let us use P to designate perceptive knowledge and C to designate communicated knowledge acquired by an agent. We want to investigate the following topics:

- how to detect and model new perceived and communicated knowledge (generation of P and C);
- detection, evaluation and resolving conflicts in the integration of P in K_p;
- detection, evaluation and resolving conflicts in the integration of C in K_c;
- detection, evaluation and managing conflicts between K_p and K_c.

Obviously, conflicts may arise within K_c, when information is communicated by different agents. For simplicity, we are not addressing this point at the present moment, but we think that we will be able to extend this dual representation in the future, in order to investigate this topic. We are currently studying some logic formalisms, as modal logics and belief logics, in order to select the one that better suits our model.

5 Conclusions

We have presented in this paper some ideas and concepts currently used in Distributed Artificial Intelligence, notably referring to cognitive agents in a Multi-Agent System approach. We have also presented what additional features a knowledge-based system must incorporate so that one can call it an agent. Finally, we have presented a model of a cognitive agent which is being developed at the LIFIA, addressing the duality of communicated and perceived knowledge in a Multi-Agent System.

6 Acknowledgments

The first author's current research is being supported by EPUSP and FAPESP - Fundação de Amparo à Pesquisa do Estado de São Paulo, Grant Number 91/1943-5. We would like to thank Matthew Turk for his useful comments during the revision of this paper.

References

[AGHA85] AGHA, G.; HEWITT, C. Concurrent programming using ACTORs: exploiting large-scale parallelism. In: BOND, A. H.; GASSER, L., ed. **Readings in distributed artificial intelligence**. San Mateo, Morgan Kaufmann Publishers Inc., 1988.

[BERT92] BERTHET, S.; DEMAZEAU, Y.; BOISSIER, O. Knowing each other better. In: INTERNATIONAL WORKSHOP ON DISTRIBUTED ARTIFICIAL INTELLIGENCE, 11, Glenn Arbor, Michigan, 1992. **Proceedings**. p. 23-41.

[BOIS92] BOISSIER, O. **Le problème du contrôle dans un système de vision integré**. Grenoble, 1992. Thèse de Doctorat, Institut National Polytechnique de Grenoble. Octobre, 1992. /In french/

[BOND88] BOND, A. H.; GASSER, L., ed. **Readings in distributed artificial intelligence**. San Mateo, Morgan Kaufmann Publishers Inc., 1988. Chapter 1, p. 3-35.

[BROO86] BROOKS, R. A. A robust layered control system for a mobile robot, **IEEE Journal of Robotics and Automation**, v. 2, n. 1, p. 14-23, March 1986.

[CAMP90] CAMPBELL, J. A.; D'INVERNO, M. P. Knowledge interchange protocols. In: DEMAZEAU, Y.; MULLER, J. P., ed. **Decentralized A. I.**. Amsterdam, Elsevier Science Publishers B. V., 1990. p. 63-80.

[CARD91] CARDOZO, E. et al. **DPSK+P**: a distributed problem solving kernel for object-oriented applications. Internal report. ITA, São José dos Campos, 1991.

[CARD87] CARDOZO, E. **DPSK**: a kernel for distributed problem solving. Pittsburgh, 1987. 144 p. PHD Thesis - CAED, Carnegie Mellon University.

[CHAN81] CHANDRASEKARAN, B. Natural and social system metaphors for distributed problem solving: introduction to the issue. **IEEE Transactions on Systems, Man and Cybernetics**, v. 11, n. 1, p. 1-5, Jan. 1981.

[COEL92] COELHO, H. et al. Experiments on achieving communication in communities of autonomous agents. In: INTERNATIONAL WORKING CONFERENCE ON DECISION SUPPORT SYSTEMS: EXPERIMENTS AND EXPECTATIONS, Fontainbleau, France, 1992. **Proceedings**.

[CORK83] CORKILL, D. D.; LESSER, V. R. The Distributed Vehicle Monitoring Testbed : a tool for investigating distributed problem solving networks. **AI Magazine**, v. 4, n. 3, p. 15-33, Fall 1983.

[DEMA91] DEMAZEAU, Y. Coordination patterns in multi-agent worlds. In: IEE COLOQUIUM ON INTELLIGENT AGENTS, Savoy Place, London, 1991. **Proceedings**.

[DEMA90] DEMAZEAU, Y.; MULLER, J. P. Decentralized artificial intelligence. In: DEMAZEAU, Y.; MULLER, J. P., ed. **Decentralized A. I.**. Amsterdam, Elsevier Science Publishers B. V., 1990. p. 3-13.

[ENGE88] ENGELMORE, R.; MORGAN, T., ed. **Blackboard systems**. Reading, Addison-Wesley, 1988.

[ERMA80] ERMAN, L. D. et al. The Hearsay-II speech-understanding system: integrating knowledge to resolve uncertainty. **Computing Surveys**, v. 12, n. 2, p. 213-253, June 1980.

[FERB91a] FERBER, J.; GASSER, L. **Intelligence Artificielle Distribuée**. In: INTERNATIONAL WORKSHOP ON EXPERT SYSTEMS & THEIR APPLICATIONS, 11., Avignon, France, 1991. Cours no. 9.

[FERB91b] FERBER, J.; JACOPIN, E. The framework of eco-problem solving. In: DEMAZEAU, Y.; MULLER, J. P., ed. **Decentralized A. I. 2**. Amsterdam, Elsevier Science Publishers B. V., 1991. p. 181-93.

[FOX81] FOX, M. S. An organizational view of distributed systems. **IEEE Transactions on Systems, Man, and Cybernetics**, v.11, n.1, p.70-80, Jan. 1981.

[GALL90] GALLIERS, J. R. Modelling autonomous belief revision in dialogue. In: DEMAZEAU, Y.; MULLER, J. P., ed. **Decentralized A. I. 2**. Amsterdam, Elsevier Science Publishers B. V., 1991. p. 231-43.

[GALL89] GALLIERS, J. R. The positive role of conflict in cooperative multi-agent systems. In: DEMAZEAU, Y.; MULLER, J. P., ed. **Decentralized A. I.**. Amsterdam, Elsevier Science Publishers B. V., 1990. p. 33-46.

[GASP90] GASPAR, G. Communication and belief changes in a society of agents: towards a formal model of an autonomous agent. In: DEMAZEAU, Y.; MULLER, J. P., ed. **Decentralized A. I. 2**. Amsterdam, Elsevier Science Publishers B. V., 1991. p. 245-55.

[GASS87] GASSER, L. et al. MACE: a flexible testbed for distributed AI research. In: HUNHS, M. N. ed. **Distributed Artificial Intelligence**. Los Altos, Morgan Kaufmann Publishers Inc., 1987.

[HEWI86] HEWITT, C. Offices are open systems. In: BOND, A. H.; GASSER, L., ed. **Readings in distributed artificial intelligence**. San Mateo, Morgan Kaufmann Publishers Inc., 1988. p. 321-9.

[ISHI92] ISHIDA, T. The Tower of Babel: towards organization-centered problem solving. In: INTERNATIONAL WORKSHOP ON DISTRIBUTED ARTIFICIAL INTELLIGENCE, 11, Glenn Arbor, Michigan, 1992. **Proceedings**. p. 141-53.

[KONO80] KONOLIGE, K.; NILSSON, N. J. Multi-agent planning systems. In: NATIONAL CONFERENCE OF THE AMERICAN ASSOCIATION FOR ARTIFICIAL INTELLIGENCE, 1., Stanford, 1980. **Proceedings**. Menlo Park, AAAI, 1980. p. 138-41.

[LENA75] LENAT, D. B. Beings: knowledge as interacting experts. In: BOND, A. H.; GASSER, L., ed. **Readings in distributed artificial intelligence**. San Mateo, Morgan Kaufmann Publishers Inc., 1988. p. 161-8.

[LEVI89] LEVI, P. Architectures of Individual and distributed autonomous agents. In: KANADE, T. et al. ed. **Intelligent Autonomous Systems 2**, v. 1, p. 315-24. The IAS Foundation, 1990.

[MINS89] MINSKY, N. H. **The imposition of protocols over open distributed systems**. Internal repport, LCSR-TR-154, Rutgers University, 1989.

[OLIV91] OLIVEIRA, E.; QIEGANG, L. Towards a generic monitor for cooperation. In: NATIONAL CONFERENCE OF THE AMERICAN ASSOCIATION FOR ARTIFICIAL INTELLIGENCE, 10., Anhein, 1991. **Proceedings**.

[ROSE86] ROSENSCHEIN, S. J.; KAELBLING, L. P. The synthesis of digital machines with provable epistemic properties. In: HALPERN, J. Y. ed. **Theoretical Aspects of Reasoning about Knowledge**. Palo Alto, Morgan Kauffman, 1986. p. 83-98.

[SHOH92] SHOHAM, Y. Agent oriented programming. In: INTERNATIONAL WORKSHOP ON DISTRIBUTED ARTIFICIAL INTELLIGENCE, 11, Glenn Arbor, Michigan, 1992. **Proceedings**. p. 345-53.

[SICH91] SICHMAN, J. S. **Uma ferramenta de monitoração para um núcleo de resolução distribuída de problemas orientado a objetos**. São Paulo, 1991. 166 p. MS Thesis - Escola Politécnica, Universidade de São Paulo. /In portuguese/

[SMIT80] SMITH, R. G. The contract net protocol: high-level communication and control in a distributed problem solver. **IEEE Transactions on Computers**, v. 29, n. 12, p. 1104-1113, Dec. 1980.

[STEFa92] STEFANINI, M. H. et al. Talisman: a multi-agent system governed by linguistic laws for natural language processing. In: INTERNATIONAL CONFERENCE ON COMPUTATIONAL LINGUISTICS (COLING), 14., Nantes, France, 1992. **Proceedings**.

[STEFb86] STEFIK, M.; BOBROW, D. G. Object-oriented programming: themes and variations. **The AI Magazine**, v.6, n.4, p.40-62, Jan. 1986.