Artificial Intelligence: a concise introduction

Vicenç Torra

School of Informatics
University of Skövde
SE-54128 Skövde, Sweden
e-mail: vtorra@his.se

Chapter 5. Distributed artificial intelligence

Artificial intelligence has benefited from parallel machines and computers with multiple cores. Algorithms and methods have been developed to exploit them. For example, there are parallel algorithms for search problems (i.e., best-first heuristic search, A*). See e.g. [3]. Distributed problem solving approaches also exist on a similar basis. That is, problems are decomposed into simpler tasks that can be performed in parallel. Consider for example, the cases in [5] about distributed planning.

– Centralized planning for distributed plans. We have a single description of the objectives of the plan, the plan is built in a centralized way, but then some actions will be executed in parallel.
– Distributed planning for centralized plans. We still have a single description of the objectives of the plan, but the plan will be developed by means of the cooperation of different *planning agents*. As [5] states "the overall problem-formulation task can be thought of as being decomposed and distributed among various planning specialists, each of which might then proceed to generate its portion of the plan".

Following [1,8] we call this subarea as parallel AI, or massively parallel AI.

We can say that parallel AI focuses on having a single and relatively well defined problem and then developing algorithms that exploit parallelism to solve it. Parallelism is either used in the computation of the solution or on the implementation of the solution. Parallel AI in this sense can be embedded in any intelligence system. We can say that Deep Blue and Watson used these techniques.

Decentralized AI follows a different perspective. In this case, we have a set of agents (or autonomous software) each one with its own goals and information, and solutions have to be found globally by these agents. In order to find a solution, agents need to interact and cooperate. E.g., agents need to exchange information with the other agents and reach agreements.

Decentralized optimization under uncertainty is an example of problem that can be solved within the framework of decentralized AI. The overall goal is to solve an optimization problem among different partners (e.g. a supply chain problem involving different companies) however the interests of each partner is not communicated to a central authority. The problem has to be solved through the interaction of these partners, each partner having partial knowledge of others' interests, and all seasoned with doses of uncertainty (e.g., rise production, delivery times and benefits can only be estimated).

We could write here multiagent systems (MAS) instead of decentralized AI. Nevertheless, there is discussion on the relationship between MAS and AI. In particular, there is discussion on whether MAS is just AI. Nevertheless, we will discuss multiagent systems as the usual way to implement decentralized AI in Section 1. An outline of the discussion on the relationship between MAS and AI is given in Section 1.1.

Types of AI. We have distinguished the following types of AI.

- **Centralized AI.** A single agent needs to solve a problem, and it solves it by means of a single-threaded algorithm.
- **Parallel AI.** A single agent needs to solve a problem. Parallelism is either used in the computation of the solution or on the implementation of the solution.
- **Decentralized AI (and multiagent systems).** A set of agents, each with its own local goals have to solve (or not) a common problem. Each agent can only access to local information.

1 Multiagent systems

Agents are autonomous pieces of software. This means that they make decisions and act based on the information they have available and their own goals. The concept of agent is sufficiently broad so that we can use it to represent multiple devices. That is, we can see from the agent perspective all type of *devices* from thermostats (whose goal is to maintain a given temperature with the actions of heating and cooling) to humans.

In artificial intelligence we focus on intelligent agents, where we expect that agents have an intelligent behavior. We will not discuss here our previous discussion on what is intelligence and what is artificial intelligence. In short, intelligent agents will have the characteristics of AI systems.

Intelligent agents are expected to have goals and that these goals drive their actions. In order to achieve them they need to interact. This interaction is in the form of communication, cooperation, coordination, and negotiation.

In this description we have repeated several times that agents have goals. We will discuss goals for agents in more detail in Section 1.2.

Main characteristics of agents. The two main characteristics of intelligent agents are:

- Autonomy. This means that decisions are made on the basis of own goals and interests.

– Interaction. This appears in the form of communication, cooperation, coordination, and negotiation.

When designing a multiagent system, we need to design agents so that they have an intelligent behavior (this corresponds to agent design) and the society where these agents interact (this corresponds to society design).

1.1 Multiagent systems and related areas

Wooldridge in his book [10] underlines that the area of multiagent systems is highly multi-disciplinary and discusses the relationship between multiagent systems and a few other areas. In particular, he mentions distributed and concurrent systems, artificial intelligence, economics and game theory, and social science. He also underlines what makes MAS different from these disciplines and areas.

From the point of view of AI, the most relevant discussion is of course what makes MAS not all just AI. Two points are made. The first one is summarized using Etzioni's quotation "Intelligent agents are ninety-nine percent computer science and one percent AI" [6]. It is argued that most agents in multi-agent systems do not need all capabilities of AI systems. The second is that AI "has largely ignored the *social* aspects of agency".

So, putting in the other way round, we can state that MAS is the area of AI that (besides of implementing decentralized AI) focuses on the social aspects of agents. That is, how agents communicate, cooperate, coordinate and negotiate.

This links with our initial discussion (see Section 4.7 **??** in the introduction of this course) on situatedness and the need of social embedding. Recall that we pointed out that H. M. Collins [4] underlined the need of socialization for achieving intelligence.

> Putting all this together, we can say with confidence that if we can't train a computer without a body to act like a socialised human, giving it the ability to move around in the world encountering the same physical situations is not going to solve the problem. On the other hand, if we can find out what is involved in the sort of socialising process undergone by a Madeleine – let us call it "socialisability" – we may be able to apply it to an immobile box. (Collins 1996 [4], p. 105)

With respect to the other areas mentioned. The main difference between MAS and distributed and concurrent systems is that agents are autonomous (and intelligent!) and syncronization needs will appear as a consequence of their behavior and reasoning (and not previously known). Game theory is used in MAS but while in economics game theory is mainly used in a descriptive way and without much emphasis on how they are computed, MAS focus on computation issues.

1.2 On the goals and related terms

Agents are expected to achieve their goals. For this, we need them to settle goals somehow. Once the goals are settled, the agent will use them to decide which actions have to be applied using e.g. planning or other search methods.

It is usual in multiagent systems to use the term intentions as rough synonym of goals. Intentions are defined in [10] as the "state of affairs that an agent has chosen and committed to", and the following characteristics are given for them.

- Intentions lead to action. When actions are decided by means of planning, intentions drive planning. The agent wants to achieve the intention, and when a plan is not suitable for achieving it, other plans will be designed.
- Intentions persists. Intentions are not changed at the first sign of trouble. Agents maintain them as long as there is no important change of their beliefs (e.g., a previous intention is removed because the agent is now convinced that the intention will not be achieved).
- Intentions constrain future deliberations. All actions that are inconsistent with the intentions (that make the intentions unreachable or more difficult to achieve) are discarded.
- There is relationship between the beliefs and the intentions of any agent. First, agents can consider that the intentions will be achieved and use them for further reasoning about future events. Believing that one cannot achieve a state and maintaining it as an intention is not rational.

There is a term related to intention: desire. For desires consistency is not a requirement. On the contrary, it is expected that intentions are consistent. Intentions have stronger influence on the actions of the agent, and they are desires to which the agent is committed. So, the agent may select from the desires which to commit.

In this setting, beliefs represent agent's knowledge. In multiagent systems it is usual to use belief, because the knowledge of the agent is not necessarily true.

Among the architectures for building agents, there is one that precisely takes its name from these terms: BDI, which stands for belief-desire-intention. See e.g. [2,10].

1.3 Interactions: Cooperation and conflict

Agents need to interact, and this causes situations of cooperation and conflict between them. Game theory studies mathematical models for these interactions. Game theory distinguishes two areas, cooperative and non-cooperative game theory.

In non-cooperative game theory we distinguish between static and dynamic games.

Static games (also known as simultaneous decision games) are those games (problems) in which agents made decisions one at a time with no knowledge on the decision made by the other agents. The prisoner's dilemma is an example of such a game. Concepts as dominance, Nash equilibria and zero-sum games are central in game theory.

Example 1. (Prisoner's dilemma) Two people are imprisoned for a serious offence. They are isolated and cannot communicate. No serious evidence exist and without confession they will be charged 1 year on a lesser charge. If one (say, A) betrays the other (say B), who keeps silent, A will be set free and B will be charged 3 years. If both betray, both are charged 2 years.

Dynamic games is when decisions are made sequentially and previous decisions are known by the agent when making a new one.

1.4 Agreements through negotiation

Two agents can make agreements when they can reach a situation in which both get some benefit from the previous state. Negotiation [7,9] corresponds to the techniques for reaching an agreement.

Four components are considered in negotiation. They are the following ones. See [10] for details.

- A negotiation set. The space of proposals that agents can make.
- A protocol. It defines the legal proposals that agents can make. It is a function of the negotiation history.
- Strategies for each agent. They permit to compute the new proposals of the agents.
- A rule to know when the negotiation is finished and the outcome of the negotiation.

Among the techniques for negotiation we find auctions. They have a clear protocol to allocate goods to agents. There are several types of auctions that differ on the protocol.

Another approach to negotiation is based on game theory.

A third one is argumentation. In argumentation agents try to reach the agreement by means of a discussion of their position. Logic is one of the tools used in argumentation.

1.5 Communication

Communication between agents is needed

- Blackboard systems.
- Message passing. Speech acts, FIPA-ACL, JADE (implements FIPA-ACL).

Final discussion

References

1. Bond, A. H., Gasser, L. (1988) A survey of distributed artificial intelligence, `http://www.exso.com/nsurvo.pdf`

2. Bratman, M. E., Israel, D. J., Pollack, M. E. (1998) Plans and resource-bounded practical reasoning, Computational intelligence 4 349-355.
3. Burns, E., Lemons, S., Ruml, W., Zhou, R. (2010) Best-first heuristic search for multicore machines, Journal of Artificial Intelligence Research 39 689-743.
4. Collins, H. W. (1996) Embedded or embodied? a review of Hubert Dreyfus' What computers still can't do, Artificial Intelligence 80 99-117.
5. Durfee, E. H. (2001) Distributed problem solving and planning, Proc. ACAI 2001, LNAI 286 118-149.
6. Etzioni, O. (1997) Moving up the information food chain: Deploying softbots on the world wide web, AI Magazine 18:2 11-18. See also Proc. 13th AAAI 4-8.
7. Jennings, N., Faratin, P., Lomuscio, A., Parsons, S., Wooldridge, M., Sierra, C. (2001) Automated negotiation: prospects, methods and challenges, Group Decis. Negot. 10:199-215.
8. Kitano, H., Hendler, J., Higuchi, T., Moldovan, D., Waltz, D. (1991) Massively parallel artificial intelligence, Proc. IJCAI 1991.
9. Lopes, F., Wooldridge, M., Novais, A. Q. (2008) Negotiation among autonomous computational agents: principles, analysis and challenges, Artif. Intel. Rev. 29 1-44.
10. Wooldridge, M. (2002) An introduction to multiagent systems, John Wiley and Sons.