Distributed Artificial Intelligence: Multi-Agent Systems

Benjamin Smith
[redacted]

Contents

Keywords
Distributed; Artificial Intelligence; Multi-Agent Systems; Distributed Computing;

Abstract
The following paper is a synthesis of related work in distributed artificial intelligence, multi-agent systems and distributed problem solving. The design for these systems are not standardized and communication conflicts occur between agents. The software design issue being observed is the methods for interaction and organization of a networked system of intelligent agents. A solution to the issues surrounding these methods could be a centralized design in conjunction with the distributed artificial intelligence design.

1 Introduction

Artificial intelligence (AI) has become ubiquitous in society. The use for AI extends through individual and a multitude of uses. The analysis of datasets to determine weather predictions, classifying target domains through computer vision, natural language processing, part-of-speech tagging, and so on are only a few examples of AI in society. In conjunction with distributed computing, AI tasks are able to decrease latency by distributing their workload across multiple compute nodes. To handle giant datasets, high-efficient data centers are dedicated to this computation and analysis. Though, AI do not have to exist in one geographical location.

Instead of the idea of delegating tasks through a network, as in distributed computing methodology, the question arises about how autonomous individual AI entities that operate over a network, *distributed artificial intelligence*(DAI), function as a single system with these tasks. This concept has been referred to as system of *intelligent agents* and is defined as *multi-agent system*(MAS). MAS and distributed AI have become synonyms for one another, though they do have their own distinct attributes (2). For this paper, DAI and MAS will be considered as the same and used interchangeably since that is the accepted trend in the field (2). The design of these systems are not standardized (3), which can be observed by the issues of such a unique system: commitment, coordination, negotiation, cooperation, distributed problem solving, centralized versus distributive (1). This paper investigates the software design issues of DAI and MAS, specifically the software design toward interaction and organization of such a system. This paper then introduces the concept of centralized design in response to the inconsistent design choices for agents. The setting considered is a system involving a group of agents under the DAI.

2 Distributed Artificial Intelligence: Multi-Agent Systems

The concept of distributed computing is to delegate workable task units to compute nodes that will process the task and intuitively return a result. The youthful subfield of AI, DAI, is built from this same process of task delegation (6). However, these tasks are delegated to what is referred to as *agents* instead of compute nodes. These agents are autonomous AI entities capable of learning and performing actions in response to events (6). A key characteristic of DAI systems is that the agents are capable of learning from one another, thus being autonomous AI entities. Analysis into how these agents learn can lead to more efficient DAI architecture design (5). This same analysis is also a relevant topic in cognitive sciences and artificial intelligence (1).

It is interesting to note that DAI are also to suggested to be an *intelligent system* (1) and it is proposed that such a system can simulate human reasoning, knowledge, and expertise for a given task (1).

2.1 Cognitive Science of an Agent

Cognitive science in terms of MAS entities involve the decision processes of a singular agent in relation to being in a group of other distinct agents. It was proposed that agents can be used to analyze the social interaction of one agent with other agents and that could be used to model intuitions about reasoning, knowledge, and planning of autonomous AI (1). It can be proposed that observing an agents methods for reasoning about other agents actions could be used to also reason about other environmentally non-intelligent processes (1). This proposition stands on the application of dynamic behaviour since if agents were static entities, their behaviour would be easily predictable. The intentional stance[1] of an agent is one taken from a high-level view of their purpose and so is considered an intentional system (6).

2.1.1 Reasoning and Knowledge

Reasoning in terms of an intelligent agent are the ability to change its own beliefs and actions, as

[1]Term coined by philosopher Daniel Dennett for the level of abstraction in which we view the behaviour of an entity in terms of mental properties

well as the same for other agents (1). The idea of reasoning and action is considered the agents commitment to the task. Most of the work done in this topic is focused on the understanding the knowledge that's needed to take an action, and how that knowledge is received through communication (1). Such knowledge is required to take an action to an event and through the agents ability to learn, new knowledge is gained from that action (1).
Agents run into problems when modal propositional logic is used with logical omniscience (all valid formulae and consequences are derived by the agent) and common knowledge (the agent knows a fact is true and they also know that all other agents know) (1). To tackle these issues of knowledge, the concept of defeasible knowledge is introduced that is implemented by the following

- Foundation theory: beliefs are sustained by explicitly justifications and are dropped if the justifications change
- Coherence theory: beliefs persist until challenged (1)

The concept of a decision space is proposed goals that are bound to their beliefs by preconditions and current conditions (1). Where the decision space are the set of hypothesis to be played by the learner agent on the data drawn from the distribution/task provided.

The last topic of reasoning is that of reasoning about other agents actions. An agent must be able to reason about other agents actions in order to recognize their own potential actions (1). Through communication, an agent has a casual understanding of actions taken by another agent and can construct feasible conclusions on a response over the set of all possible actions (1). However, this method is unable to infer why one conclusion is more suitable over another (1) and so a method to successfully weight these responses should be given. This plan recognition model uses defeasible reasoning and direct argumentation ascription of belief and used to support communication between agents (1).

2.1.2 Assessment

The basic question of DAI is when and what can an agent do and is solved by assessing the distribution situation of its environment which can involve many other agents (1). This assessment involves inquiring, abstracting, and organizing information of the environment. If the information correlates with the agents expectations, then appropriate actions can be executed. If not, then new expectations may be required. (1)

2.1.3 Organization

A DAI requires knowledge to create agents that can reconfigure their actions and interactions as the context or environment change (1). The structure of a DAI is the pattern of information, the control of relationships between agents, and the distribution of problem solving capabilities among them (1). With this established, the notion of a *global plan* is introduced. With this in mind, the following conditions must be met to ensure success at problem solving:

- Coverage: any necessary portion of the overall problem must be within the problem solving capabilities of at least one agent
- Connectivity: agents must act in a manner that permits the covered activities to be developed and integrated into an overall solution
- Capability: coverage and connectivity must be achievable within the communication resource limitations, as well as the reliability specifications of the group (1)

Where the structure of the DAI will provide the necessary information to perform the underlying mechanisms of the conditions. An organization of a DAI is defined by the embedded beliefs, commitments, and intentions of the agents within (1).

DAI, intuitively, must align their actions to the global plan goal and form negotiation and cooperation tactics to achieve this goal. The following system goals aid agents in their attempt to find solutions for the global goal.

- Increase the solution creation rate by forming subsolutions in parallel
- Minimize time agents must wait for results from other agents by coordinating activity
- Improve overall problem solving by permitting agents to exchange predictive information
- Assign important tasks to multiple agents to increase the chance that a solution will be found
- Improve the use of physical resources by permitting agents to exchange tasks

- Improve the use of individual agent expertise by allowing agents to exchange goals, constraints, partial solutions, and knowledge
- Reduce the amount of unnecessary duplication of effort by allowing agents to recognize and avoid redundant activities
- Have agents verify a solution with each other through rederivation of their expertise and knowledge
- Increase variety of solutions by allowing agents to form local solutions without being overly influenced by other agents
- Reduce communication resources by specifying what messages are permitted to be exchanged (3)

By following these goals, DAI agents can maximize effective organization across multiple dimensions (3)

2.1.4 Communication

Communication over a network can be accomplished a multitude of ways. DAI systems are known to utilize no communication, primitive communication, plan and information passing, blackboard exchange, message passing, and high-level communication to name a few (1). The act of communication is usually to determine how they can/will help each other. This process of negotiation and cooperation occurs throughout the DAI.
When no communication exists between agents, an agent only infers the other agents plans (1). The use of game-theory is implemented by use of payoff matrices that contain agent pay-off for an interaction (1). A step above that is the primitive communication method where agents communicate by use of signals. These signals are usually fixed interpretations that is known by all agents in the DAI (1). Blackboard exchange utilizes a shared memory in which agents write messages, post partial results, and find information (1). Message passing involves typical message passing systems such as serialized JSON or OpenMPI solutions. Finally, excluding plan and information passing, high-level communication refers to the understanding of intention based on beliefs, facts, and previous knowledge (1).

As mentioned, plan and information passing is a method of communication used throughout a DAI. This involves a total plan being communicated via two way transaction between two distinct agents. The total plan acceptance of an agent is based on a first-in-first-out (FIFO) scheme (1). However, a few issues arise concerning passing total plans. One issue is that passing total plans is computationally expensive (1) to process and also in regards to network bandwidth. Another issue is that the total plan is not guaranteed over any medium of communication (1). Finally, the state of the system may change while a plan is being communicated and becomes inapplicable to the new system state (1). Therefore, although the previous sections discussed plan passing, total plan passing is not a best strategy and another mode of communication to deliver and receive plans should be put in place.

The ability of agents to negotiate is to identify potential interactions through communication with other agents or by reasoning about the current state of the environment and the intentions of other agents in the system (3). Modifying the agents intention or the other agents intentions in order to prevent harmful situations by creating cooperative solutions (**?**). This can be accomplished through the above plan sharing and negotiation on task importance and solution creativity. These descriptions capture human interaction where foundations of DAI have been established (3).

3 Distributed Problem Solving

Up to this point, a *collaborative reasoning* system has been discussed involving agents solving the same problem collaboratively (5). In context of a distributed computing scenario, now a *distributed problem solving* system is considered as an improvement on design efficiency of task completion (5). In the same context as distributed computing and where DAI get their foundation of methodology from (2), tasks are delegated and assigned to agents in the distributed problem solving (DPS) case. The tasks are computed asynchronously and each autonomous agent will plan their actions accordingly (5). In the end, the partial solutions will be synthesized into an overall complete solution. DPS, by some, is considered a subset of MAS (2) and intuitively inherit the characteristics of the agents. There are four steps to the distributed problem solving method:

- Problem decomposition
- Task assignment
- Local problem solving
- Solution synthesis (5)

Which essentially follow along with typical distributed computing methodology and so increases performance via parallelization of task analysis. In regard to task assignment and solving, a DPS is a MAS under a few assumptions:

- Benevolence Assumption: agents want to help each other whenever possible
- Common Goal Assumption: motivation for benevolence is that agents operate under a common goals (2)

The benevolence assumption does not guarantee cooperation and coherent coordination due to difficulties timing and local perspectives (2). As well with the common goal assumption, the same issues from the benevolence assumption hold in terms of globally coherent tasks (2). Since agents are able to communicate with one another, there is the risk that they may inundate other agents with superfluous information (2). What is worse is that agents may cross purposes and communicate with another agent causing a distraction and side-track towards unimportant goals (2). This can be caused by a degree of redundancy of knowledge in an agents system (5).

4 Centralized Multi-Agent Design

In centralized multi-agent planning one agent acts as the central pool that collects the other agents plans and assesses them to determine conflicts - thus eliminating redundant work in the system. The central agent also determines if limited resources are available for any of the participating agents. Finally, a safety analysis is completed in order to determine if any potential interactions may lead to conflicts. Conflicting plans are not dismissed, but instead placed in a critical plan dialog based on priority. The plans and critical plans are then communicated back to the other agents to implement. (1)

However, in this situation there is no central agent with a global view (1) and so a proposition for a more centralized design is introduced. So a concept of a centralized design that manages the global view of the network of agents is introduced. Agents retain information of the (common) global goal but the centralized designer would act accordingly to make sure each agent is benevolent and work as a whole (2). Having a centralized designer to micro-manage communication, plan orientation, agent organization, and domain knowledge priority could potentially eliminate some of the issues mentioned before with respect to the topic.

A centralized designer could also allow for a non-benevolent environment for agents to come and go free, in an "open society" setting (2). These agents would be bound by the "laws" of the centralized designer and as long as those laws are obeyed, the agent may operate within the "society" (2). This could lead for more efficient results from independent agents acting on their own interests and beliefs (2).

5 Conclusion

DAI, a subdomain of AI, are becoming more prevalent like static AI solutions and can offer a deeper insight into what it means to communicate and reason as an AI (1). Through reasoning and prediction of intent, intelligent agents are able to take benevolent action that aids other agents with respect to their common goal. The process of negotiation and cooperation about how to manage plans and actions between agents are accomplished based on the limited resources available, timing, and concerns for conflict on those actions. Through the organization and assessment of the environmental situation, the structure for the DAI is defined. That structure lays way to issues involving communication and planning where plans cannot be executed due to conflict or accidental distraction from agents whose knowledge is different than that of agent being communicated with. Attempts of having a centralized agent in charge of plan management have been made, but as efficient as they are at solving plan conflicts, they do not hold the global view and can only operate as their notion of the common goal exists - which could lay to bias plan and critical plan management. To approach this issue, a centralized entity or centralized designer that operates on the global view and is able to manage and micromanage agents plans and communication respectively is suggested.

As agents can be observed in terms of human reasoning - so can the overall system that they operate with in. With this approach, the consideration of a centralized governing entity is proposed to manage the global view in which the intelligent agents operate. There is research involving a democratized DAI that involve agents being able to cast votes to other agents (4). An issue with this is that agents must have built up a "reputation" in order to vote and could so be swayed by a highly reputed that actions are taken toward a bias (4). Having a centralized entity manage and prevent such issues gives a fair opportunity for all agents to participant in the DAI environment and global goal.

6 Future Work

The focus of this paper was on the overall structure of DAI/MAS systems and how they operate. What was not discussed was the issues surrounding the interoperability of agents (4) (1). The problems surrounding this issue involve the challenge of coordinating agent specific goals within the global goal among MAS systems (4). This is a more narrow view inside DAI systems whereas the larger view was taken in this paper, and so intuitively the next step of investigation in DAI structures.

The concept of plan passing was introduced and issues surrounding total plan passing were raised. Research into appropriate and efficient solutions to communicate plans should follow.

References

[1] Chaib-Draa, B., Moulin, B., Mandiau, R., and Millot, P. Trends in distributed artificial Intelligence Artificial Intelligence Review 6, pp 35-66, 1992

[2] Durfee. E and Rosenschein, J. Distributed problem solving and multi-agent systems: comparisons and examples AAAI Technical Report WS-94-02, 1994

[3] Durfee. E., Lesser, V., and Corkill, D. Trends in cooperative distributed problem solving IEEE Transactions on Knowledge and Data Engineering, 1(1):63–83, 1995

[4] Monte, G. and Goertzel, B. Distributed, decentralized, and democratized artificial intelligence Technological Forecasting and Social Change 141, 2019

[5] Shaw, M. and Sikora, R A distributed problem-solving approach to rule induction: learning in distributed artificial intelligence systems Technical Report, CMU-RI-TR-90-28, The Robotics Institute, Carnegie Mellon University, 1990

[6] Wooldridge, M. and Jennings, N. Agent theories, architecture, and languages: a survey In: Intelligent Agents. ATAL 1994. pp 1-39