

Towards a Taxonomy of Decision Making Problems in Multi-Agent Systems

Christian Guttman*

School of Primary Health Care
Faculty of Medicine, Nursing and Health Sciences, Monash University
Notting Hill, 3168, VICTORIA, Australia
christian.guttman@gmail.com

Abstract. Taxonomies in the area of Multi-Agent Systems (MAS) classify problems according to the underlying principles and assumptions of the agents' abilities, rationality and interactions. A MAS typically consists of many autonomous agents that act in highly complex, open and uncertain domains. A taxonomy can be used to make an informed choice of an efficient algorithmic solution to a class of decision making problems, but due to the complexity of the agents' reasoning and modelling abilities, building such a taxonomy is difficult. This paper addresses this complexity by placing model representation, acquisition, use and refinement at the centre of our classification. We classify problems according to four agent modelling dimensions: model of self vs. model of others, learning vs. non-learning, individual vs. group input, and competition vs. collaboration. The main contributions are extensions of existing MAS taxonomies, a description of key principles and assumptions of agent modelling, and a framework that enables a choice for an adequate approach to a given MAS decision making problem.

1 Introduction

Coordination of activities in natural and engineered systems often require that agents make individual and joint decisions. Multi-Agent Systems (MAS) consists of autonomous agents that make their own decisions (and do not follow decisions made by others) and have their own beliefs (and do not rely on beliefs maintained by other agents) [1,2]. Choosing an adequate method for effective coordination requires a thorough understanding of the underlying assumptions and principles of how agents model their social surroundings and how these models are used in decision making.

An autonomous agent requires a model of its own behaviour and that of other agents to make informed decisions about taking its next action. Knowledgeable agents (i.e., agents that have models) can predict its own actions and those of other agents. A system with such agents is structured and predictable. This is opposed to a system with ignorant agents associated with chaotic behaviour [1,2]. Research has demonstrated that using agent models benefits agent coordination in a variety of scenarios [3,4,5]. However, despite the importance of agent models in coordination, previous taxonomies do not place a notable emphasis on an agent's ability to model agent behaviour for its decision making processes. Instead, taxonomies often organise the problem space

* This research was supported in part by Linkage Grant LP0774944 from the Australian Research Council.

considering macro features of decision making problems. For example, [6,7] centre the issue of how many agents are required to perform one or several tasks. [8] considers how to classify different types of joint activity problems, and [9] classifies based on heterogeneity, distribution, and autonomy. This research offers useful insights into MAS coordination, but it does not emphasise a central feature of agent systems: the complexity involved when an agent maintains models of itself and others to make decisions. A better understanding of the appropriateness of an approach to a problem requires a re-organisation of the problem space based on the underlying assumptions and principles of MAS.

This paper offers extensions to existing MAS taxonomies and advances the state of the art in artificial intelligence, and particularly MAS, as follows.

- **Organisation of the space of decision making problems.** Unlike previous taxonomies, we propose a classification structure of decision making problems in MAS that places the role of agent models at its centre.
- **Analysis of the problem space.** We identify critical assumptions that define four agent modelling dimensions across the space of problems. The location of a decision making problem in this space requires an analysis against these assumptions.
- **Prescriptive framework.** Our taxonomy enables an informed choice of an approach for a given problem class as we have a clear understanding of the underlying assumptions and principles of the role of agent models in coordination.
- **Identification of research opportunities.** This taxonomy is used to classify well-known approaches, some offer provable, others heuristical solutions to problems. Our taxonomy indicates underexplored types of decision making problems.

This research arranges decision making problems in MAS emphasising the use of models maintained by agents. This taxonomy is a first attempt to find an appropriate approach for a given problem, and offers a basis to develop a unified model.

Section 2 discusses related research. Section 3 defines the four agent modelling dimensions, offers a classification scheme and positions well known existing approaches. Section 4 discusses a possible unified approach. Section 5 concludes this paper.

2 Related Research

Section 2.1 discusses the use of agent models in of MAS coordination. Section 2.2 reviews related MAS taxonomies.

2.1 Role of Agent Models

[1] argue that the coordination of MAS will be chaotic if agents are not able to predict their own behaviour and that of others. Enabling agents to maintain models of the behaviour of other agents is an important requirement for the coordination of MAS [1]. Previous research has demonstrated that using agent models benefits agent coordination in predicting the decisions made by collaborators [4], matching students with tutors in collaborative support environments [5], and predicting the performance of soccer-agents in RoboCup [3]. Each initiative makes distinct assumptions that influence what to model (feature selection), how to model it (feature representation), and how to use models (model usage). Previous research has not adequately addressed the classification of MAS decision making based on the role of agent models.

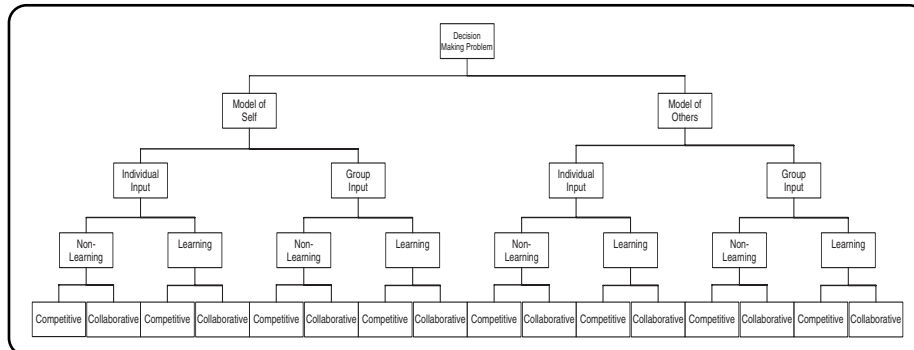


Fig. 1. A taxonomy on decision making problems in MAS

2.2 Multi-Agent System Taxonomies

Modelling the behaviour of agents is a crucial skill of an agent [1], but many MAS taxonomies do not place the role of agent models at the centre of the classification [9,6,7,8]. [7] has been widely used for task classifications and uses four categories.

	Task execution requires one agent	Task execution requires several agents
Separate task	Single Agent - Single Task (SA-ST)	Mult. Agents - Single Task (MA-ST)
Simultaneous tasks	Single Agent - Mult. Tasks (SA-MT)	Mult. Agents - Mult. Tasks (MA-MT)

This taxonomy demonstrates how a well structured classification can assist in understanding a complex problem space. As such, the taxonomy offers a useful starting point to understand fundamental types of decision making problems. However, this MAS taxonomy (as well as many others [6,7,8,9]) has a significant limitation as the classification is not based on models maintained by agents. Our taxonomy structures this space by placing the complexity of building agent models at the taxonomy's centre.

3 Four Agent Modelling Dimensions

The complex role of models maintained by agents is central to our MAS decision making taxonomy. We consider the role of agent models in being able to make particular types of decisions. This taxonomy enables positioning of a problem in the decision making space. As with other MAS taxonomies (e.g., [6,7]), our taxonomy identifies classes of decision making problems where provable optimal solutions exist while other classes can only be solved using heuristics. While this taxonomy is not exhaustive, it characterises the complexity of agent modelling in MAS decision making problems. Four axes describe the space of MAS decision making problem (defining 16 classes, Figure 1). We now discuss a rationale for each axis.

3.1 Dimension 1: Self Model versus Model of Others

Does an agent model only its own performance or also that of others? At one extreme, an agent has information that pertains to itself, for example, it estimates its own performance or it holds a value that represents the estimated pay-off for taking a particular




Fig. 2. Input offered by several agents (a) for each agent’s decision making process with consequences for itself, or (b) for group decision making process with consequences for entire group

action. Many approaches based on each agent’s knowledge of its own behaviour are *market-driven approaches*, because each agent is assumed to know its own performance best and accurately. In the Contract Net (CNET) protocol, a manager agent announces a task, each contractor assesses how well it performs the task and makes a bid, and the manager then assigns the task to the contractor which made the most adequate bid. The CNET protocol works best when the agents have accurate self-estimations, because the manager can rank the bids and select the highest bidder [10].

At the other extreme, an agent has information that pertains to other agents. Approaches where an agent requires information of other agents’ behaviour are referred to as *agent modelling approaches*. Agent models are used to predict decisions of collaborators [4], match students with tutors in collaborative support environments [5], and predict the performance of soccer-agents in RoboCup [3].

3.2 Dimension 2: Individual Input versus Group Input

The input of a decision making process is either derived from an individual, or a group (Figures 2(a) and 2(b)). At one end of this spectrum, an individual agent makes a decision and uses as input for the decision making process its own knowledge (Figure 2(a)). An example here is Multi Agent Reinforcement Learning (MARL), where multiple agents execute tasks individually and use reinforcement learning to coordinate their actions, taking into account various configurations (e.g., if agents can observe each others’ actions) [11,12,13]. MARL agents *do not jointly* select a team which then executes a given task, instead MARL is concerned with the internal coordination of a team.

At the other end of this spectrum, we have decision making processes that require the input of several agents (Figure 2(b)), which we refer to as *group decision making*. Voting is a preference aggregation procedure applied in situation where agents have conflicting preferences – agents “compete” as each aims to see its own preference implemented [14,15]. The aim of preference aggregation is to find a “collective decision” of what best reflects the “will” of the group.

3.3 Dimension 3: Learning versus Non-learning

We can divide this space by considering decision making problems that are made only once, or over multiple rounds where learning plays an important role [16]. In the latter case, each agent requires adequate processes to maintain models and refine them over time. For example, an agent can update its models whenever new information of its

own behaviour and that of agents is available. This information may be acquired from different sources. For example, [17]’s agents update their models using observations of other agent’s behaviour over several iterations. Agent models in [5] are refined using information derived from explicit communication. Further discussions on related topics of MAS learning can be found in [16]).

Many other MAS decisions do not consider multiple rounds and learning is therefore not required. The Contract Net (CNET) protocol is an example which considers only a single round before a decision is made [10]. In particular, the CNET protocol and many of its extensions describe how a contract is made after a single announcement of the task (i.e., in a single round). In these cases, managers and contractors are not required to be able to learn. Similarly, in many voting frameworks, a group makes a decision with little consideration to long term consequences [14,15].

3.4 Dimension 4: Collaboration versus Competition

An agent’s decision making style can range from being *collaborative* to being *competitive*. That is, an agent makes a decision intending to improve the welfare of a group or task (collaborative) or its own welfare (competitive).

Collaborative agents aim to maximise the welfare of the group, e.g., by finding the best allocation of a team to a task (i.e., an agent maximises the group or task utility before its own). These agents aim to offer a best “global solution” as opposed to competitive procedures that aim to find adequate trade-off’s between several parties. For example, [18]’s agents are collaborative, because each agent aims to detect and resolve problems that could jeopardise a successful completion of a mission.

A competitive agent exhibits behaviour that maximises its own utility – a behaviour also referred to as *self-rational or self-interested*. Agents exhibit self-rational behaviour in settings where resources are limited or different agents have opposing or conflicting beliefs. For example, in auctioning, a group decision is made based on the competitive bids of agents [19,20]. In these settings, an agent may only have information of its own evaluation of the auctioned item in question, but information by competing agents may be unreliable and its accuracy can not be trusted.

4 A Unified Approach to Distributed Decision Making?

Can we find a unified approach that represents and solves the problem classes defined in our taxonomy? Our taxonomy shows that there is a multitude of MAS approaches for decision making, and many are located at different ends of the dimensions as we discussed in Section 3. How can we build a computational model that unifies many, if not all approaches to decision making problems classified in our taxonomy?

[21] offers an initial framework that captures a wide variety of decision making problems located across the dimensions discussed in Section 3. [21] studies the refinements of allocations based on group decisions. We refer to this as Collective Iterative Allocation (CIA), because decisions are made together and allocations can be refined (and iterated over time). In CIA, agents model their own performance (as in CNET) and that of others. It allows for single round decision problems as well as for multiple rounds (where agents are able to learn). Different competition and collaboration approaches can be defined by the group decision policy. The CIA framework assumes that a decision is always made by a group (that is, the voting policy requires input from several

agents). One way to address this issue is to consider the conditions under which each agent should make its own decisions (e.g., as is done in MARL) or follow the decisions made by the group decision policy.

5 Conclusion

This paper discusses a taxonomy for MAS decision making problems. This paper offers extensions to existing taxonomies on decision making in MAS and makes four contributions. We propose a classification structure of decision making problems in MAS that places the role of agent models at its centre (this classification can be represented using a tree structure where classes are clearly separated). We identified critical assumptions which define four agent modelling dimensions in the space of problems. A location of a decision making problem in this space requires an analysis against these assumptions. Our taxonomy enables an informed choice of an approach for a given problem class as we have a clearer understanding of the underlying assumptions and principles of the role of agent models in coordination. Finally, this taxonomy can be used to identify research opportunities as it classifies well-known approaches. We also discussed that the CIA framework is a first step towards a unified approach for many decision making problems. A future research direction is to extend this framework to enable further unification. In future works, we aim to offer a comprehensive survey of existing approaches, as well as to continue the formalisation of the taxonomy discussed in this paper.

References

1. Bond, A.H., Gasser, L.: An analysis of problems and research in DAI. In: Bond, A.H., Gasser, L. (eds.) *Readings in Distributed Artificial Intelligence* (1988)
2. Wooldridge, M.: *Introduction to Multiagent Systems*. John Wiley & Sons, Inc., Chichester (2002)
3. Stone, P., Riley, P., Veloso, M.M.: Defining and using ideal teammate and opponent agent models. In: *Proceedings of the Innovative Applications of Artificial Intelligence Conference (IAAI)*, pp. 1040–1045 (2000)
4. Gmytrasiewicz, P.J., Durfee, E.H.: Rational communication in multi-agent environments. *Autonomous Agents and Multi-Agent Systems* 4(3), 233–272 (2001)
5. Vassileva, J., McCalla, G.I., Greer, J.E.: Multi-agent multi-user modeling in I-Help. *User Modeling and User-Adapted Interaction* 13(1–2), 179–210 (2003)
6. Dudek, G., Jenkin, M., Milios, E., Wilkes, D.: A taxonomy for multi-agent robotics. *Autonomous Robots* 3(4), 375–397 (1996)
7. Gerkey, B., Mataric, M.: Are (explicit) multi-robot coordination and multi-agent coordination really so different. In: *Proceedings of the AAAI Spring Symposium on Bridging the Multi-agent and Multi-robotic Research Gap*, pp. 1–3 (2004)
8. Klein, G., Feltovich, P., Bradshaw, J., Woods, D.: *Common ground and coordination in joint activity*. *Organizational Simulation* (2004)
9. Bird, S.: Toward a taxonomy of multi-agent systems. *International Journal of Man-Machine Studies* 39(4), 689–704 (1993)
10. Smith, R.G.: The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers* 29(12), 1104–1113 (1980)
11. Claus, C., Boutilier, C.: The dynamics of reinforcement learning in cooperative multiagent systems. In: *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI)*, pp. 746–752 (1998)

12. Shoham, Y., Powers, R., Grenager, T.: Multi-agent reinforcement learning: A critical survey. In: AAI Fall Symposium on Artificial Multi-Agent Learning (2004)
13. Sandholm, T.: Perspectives on Multiagent Learning. *Artificial Intelligence (Special Issue on Multiagent Learning)* 171, 382–391 (2007)
14. Arrow, K.J.: *Social choice and individual values*. J. Wiley, New York (1951)
15. Fishburn, P.: *The theory of social choice*. Princeton University Press, Princeton (1973)
16. Stone, P., Veloso, M.M.: Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots* 8(3), 345–383 (2000)
17. Suryadi, D., Gmytrasiewicz, P.J.: Learning models of other agents using influence diagrams. In: *Proceedings of the seventh International Conference on User Modeling (UM)*, Banff, Canada, pp. 223–232 (1999)
18. Tambe, M.: Towards flexible teamwork. *Journal of Artificial Intelligence Research* 7, 83–124 (1997)
19. Vickrey, W.: Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance* 16(1), 8–37 (1961)
20. Boutilier, C., Goldszmidt, M., Sabata, B.: Sequential auctions for the allocation of resources with complementarities. In: *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 527–523 (1999)
21. Guttman, C.: *Collective Iterative Allocation*. PhD thesis, Monash University (2008)