Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)
Doctoral Consortium Track

On the Learnability of Knowledge in Multi-Agent Logics

Ionela G Mocanu
University of Edinburgh
i.g.mocanu@ed.ac.uk

Abstract

Since knowledge engineering is an inherently challenging and somewhat unbounded task, machine learning has been widely proposed as an alternative. In real world scenarios, we often need to explicitly model multiple agents, where intelligent agents act towards achieving goals either by coordinating with the other agents or by overseeing the opponents moves, if in a competitive context. We consider the knowledge acquisition problem where agents have knowledge about the world and other agents and then acquire new knowledge (both about the world as well as other agents) in service of answering queries. We propose a model of implicit learning, or more generally, learning to reason, which bypasses the intractable step of producing an explicit representation of the learned knowledge. We show that polynomial-time learnability results can be obtained when limited to knowledge bases and observations consisting of conjunctions of modal literals.

1 Introduction

Knowledge engineering is an inherently challenging and somewhat unbounded task, and in that sense, machine learning has been widely proposed as an alternative. But most machine learning systems focus on inferring representations with respect to an underlying environment that assumes a single agent. However, in many real-world scenarios, we need to explicitly model multiple agents, where intelligent agents act towards achieving goals either by coordinating with the other agents or by overseeing the opponents moves if in a competitive context. In this sense, several sophisticated formal logics have been proposed for modelling such contexts, from areas such as philosophy, knowledge representation and game theory.

Epistemic logics have been widely studied to reason about multi-agent systems (MAS), often in combination with temporal modalities. The typical language extends propositional logic by adding n modalities K_i representing the knowledge of agent i, B_i representing beliefs of agents, as well as modalities representing various other mental states (desires, intentions, awareness) and/or the flow of time. Modal propositional logic is widely recognised as a specification language for a range of applications such as robotics, games and distributed systems, in the multi-agent-based systems such as air traffic control for example [Belardinelli and Lomuscio, 2009]. While highly expressive, reasoning about other agents' beliefs in complex scenarios comes at a computational cost, which is the case of the K_n logic where entailment checking is PSPACE-complete. In that sense, [Lakemeyer and Lespérance, 2012] has proposed a knowledge base that is in a proper epistemic knowledge base (PEKB, sets of modal literals) form and if encoded into prime implicate normal form, entailment can be decided in polynomial time. However, this computational speed-up comes with few limitations for example the knowledge base is restricted to a specific form, which becomes computationally costly as the modal depth increases. This form of knowledge might be useful for certain applications however it does not handle an important epistemic notion: knowing whether, which is encoded as $K_i\phi \vee K_i\neg\phi$. That is, it does not cover any form of incomplete knowledge or horn clauses, therefore very limited forms of inference can be gained using this approach. Disjunctions have shown to constitute the computational cliff when it comes to reasoning. That is because disjunctions have two main applications of use: one is to express rules such as horn clauses ($\phi \rightarrow \psi$), which are used for chaining in reasoning, and the second application which to express incomplete knowledge, such as the knowing whether the sentence above holds.

We propose a multi-agent logic of belief that performs reasoning with partial information in a multi-agent context. We design a multi-agent subjective language that adopts a perspectival view and reasons from the point of view of a single root agent. The agent is an omniscient agent and reasons about the beliefs of the other agents, therefore the focus of this approach is on beliefs as opposed to knowledge. This will distinguish between what is true in the real world and what the agents know or believe about the world. For example, the beliefs of agent A about the world may differ from agent B's knowledge and what agent A believes B to know may differ from what B actually believes. In the context of multi-agent systems, we are interested in reasoning about the entailment of an input query about the other agents with respect to a background theory that contains the beliefs of agents in the application domain.

Motivated by the trade-off of a computationally tractable deductive reasoning system and one that is limited from a practical point of view, [Liu *et al.*, 2013] proposed a logic that lies between these extremes; they proposed a logic of belief. This language proves to be tractable in the propositional case and deals with disjunctions in a controlled manner. The two forms of disjunctions are dealt with separately within the logic: the horn clauses are expressed as explicit belief formulas and are closed under unit propagation, whereas the incomplete knowledge is reasoned about by splitting into cases. The level of case splitting tolerated is bound by a level index k, with the meaning that the higher the level k, the more resources are required in order to derive the implicit beliefs from the explicit ones.

2 Multi-Agent Subjective Language

Intelligent agents need knowledge about the world in order to reach good decisions. Knowledge is contained within agents' minds in the form of sentences in a knowledge representation language that is stored in a knowledge base.

Knowledge base (KB): a set of sentences, is the central component of a knowledge-based agent. Each sentence is expressed in a language called a knowledge representation language and represents some assertion about the world. Our proposed reasoning service takes a perspectival point of view, meaning that it reasons from the viewpoint of a single root agent. This approach was proposed by [Muise *et al.*, 2015], which, despite the inability to handle disjunctive beliefs, it allows nested beliefs (of bounded depth) and non-homogeneous agents. It assumes a finite set of agents, each with (incomplete and possibly incorrect) beliefs about the world. In our setting, we perform reasoning for the designated agent, the root agent symbolically represented as $\star$. The novelty in our approach is that we combine the computational property of a reasoning service based on Subjective Logic with the expressiveness advantages of reasoning as another agent in a multi-agent context. We will assume that the root agent's view of the world differs from what a particular agent i believes. The state of the world represents the mental model of a particular agent, an omniscient agent, that perceives an environment that includes all other agents, so every reasoned formula is in this sense epistemic (it does not change the state of the world, different from ontic).

3 Example

Consider for example a card game, in a more simplified version, consisting of only four cards: clubs, diamond, heart and spades, and two agents A and B. Now let's suppose that A draws one card from the set. Agent A now knows her card, for example, spades, but she does not know what card does agent B pick. Using multi-agent epistemic logic we can model this type of knowledge into logic. For example, we can infer that agent A knows that B knows his own card, but not A's card.

4 Syntax

We assume a finite set of agents $\mathcal{A} = \{\star, A, B...\}$, a countable set of propositional variables $\mathcal{P} = \{p, q, ...\}$ and the standard logical connectives $\neg, \wedge, \vee$. Multi-agent subjective language $\mathcal{MSL}$ is the propositional epistemic logic whose atom formulas are beliefs atoms of the form $B_k^A\phi$, read as ϕ is believed by agent A at level k, where B_k is a modal operator for any $k > 0$ and ϕ is either another belief formula of depth $d-1$ or formulas from the language of $\mathcal{P}$. The modal operator B_k for belief and B_k^A, together with the set of propositional variables form the infinite set of (modal and objective) atoms: $\{p, q, .., B_k^A, B_i^B B_j^A q, B_0^\star, ...\}$ with arbitrary nestings of maximum bound d. $\mathcal{MSL}$ is the least set of expressions: a. if $\phi \in \mathcal{MSL}$, $k \geq 0$ then $B_k^A\phi \in \mathcal{MSL}$ called belief atom of agent A at level k; and b. if $\phi, \psi \in \mathcal{MSL}$, then $\neg\phi$, $(\phi \vee \psi) \in \mathcal{MSL}$. All predicates must occur inside the modal operator. In contrast with the previous work on limited belief [Liu *et al.*, 2013], modalities can be nested up to a bound degree d. We contrast this form of knowledge with the standard possible-world logic of belief which has a disadvantage of logical omniscience because all agents are aware of their possible worlds.

5 Progress and Limitations

The contribution of this project is an efficient framework for reasoning in a multi-agent logic. Various strategies of obtaining this framework have been analysed and, of course, each comes with its limitations and trade-offs. In particular, we have the expressiveness-efficiency trade-off. In a subjective multi-agent logic, we have limited the range of possible applications by shifting the reasoning from the perspective of the root agent only. However, by integrating it with the limited belief reasoning system [Liu *et al.*, 2004], we would obtain a tractable reasoning system. The success criteria of this project would be a sound and complete reasoning algorithm integrated for multi-agent modal logic, for which we can show polynomial guarantees. This project is at a theoretical stage where the formal framework of this logic is yet under development.

References

[Belardinelli and Lomuscio, 2009] Francesco Belardinelli and Alessio Lomuscio. Quantified epistemic logics for reasoning about knowledge in multi-agent systems. *Artif. Intell.*, 173(9-10):982–1013, 2009.

[Lakemeyer and Lespérance, 2012] Gerhard Lakemeyer and Yves Lespérance. Efficient reasoning in multiagent epistemic logics. In *ECAI 2012 - Montpellier, France, August 27-31 , 2012*, pages 498–503, 2012.

[Liu *et al.*, 2004] Yongmei Liu, Gerhard Lakemeyer, and Hector J. Levesque. A logic of limited belief for reasoning with disjunctive information. In *KR*, page 587–597, 2004.

[Liu *et al.*, 2013] Guohua Liu, Tomi Janhunen, and Ilkka Niemelä. Introducing real variables and integer objective functions to answer set programming. In *KDPD 2013, September 11-13, 2013*, pages 118–135, 2013.

[Muise *et al.*, 2015] C. Muise, V. Belle, P. Felli, S. McIlraith, T. Miller, A. Pearce, and L. Sonenberg. Planning over multi-agent epistemic states: A classical planning approach. In *Proc. AAAI*, 2015.