Learnability with PAC Semantics for Multi-agent Beliefs

Extended Abstract

Ionela G. Mocanu
The University of Edinburgh
Edinburgh, United Kingdom
i.g.mocanu@ed.ac.uk

Vaishak Belle
The University of Edinburgh
Edinburgh, United Kingdom
vbelle@ed.ac.uk

Brendan Juba
Washington University in St. Louis
Missouri, United States
bjuba@wustl.edu

ABSTRACT

This work proposes a new technical foundation for demonstrating Probably Approximately Correct (PAC) learning with multiagent epistemic logics, using implicit learning to incorporate observations into the background knowledge. We explore the sample complexity and the circumstances in which the algorithm can be made efficient.

KEYWORDS

Multi-Agent Logics; Knowledge Acquisition; Only-Knowing

ACM Reference Format:
Ionela G. Mocanu, Vaishak Belle, and Brendan Juba. 2023. Learnability with PAC Semantics for Multi-agent Beliefs: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

Many AI applications nowadays model environments with multiple agents, where each agent acts using their own knowledge and beliefs to achieve goals either by coordinating with the other agents or by challenging an opponent's actions in a competitive context. Reasoning not just about the agent's world knowledge but also about other agents' mental state is referred to as *epistemic reasoning*, for which a variety of modal logics have been developed [4]. While a number of sophisticated formal logics have been proposed, they do not, to a large extent, address the problem of *knowledge acquisition.* A very recent line of work initiated the idea of an *implicit knowledge base* [7] constructed from observations, via so-called PAC-Semantics [12]. Intuitively, in implicit learning, we efficiently check entailment from all of the possible sentences of the logic that could be observed to hold with high probability, without explicitly representing each of these sentences. The implicit approach avoids the construction of an explicit hypothesis but still allows us to reason about queries against noisy observations. This can be contrasted to the popular approaches of *inductive logic programming* [11] and *statistical relational learning* [3, 5]. Since the development of this technique, learning with the PAC-semantics has been extended to certain fragments of first-order logic, e.g., [1]. Here, we continue this line of work and extend the problem of implicit learnability to epistemic multi-agent logic. We provide in fact, concrete results about sample complexity and correctness, as well as polynomial time guarantees under certain assumptions. We leverage some recent results on the so-called *Representation Theorem* explored for single-agent and multi-agent epistemic logics which reduces epistemic reasoning to propositional logic. In addition, space-treelike resolution or bounded-width resolution can be used to further achieve tractability for propositional logic.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. …$ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
…$15.00

2 PRELIMINARIES

Syntax Let $\mathcal{L}_n$ be the propositional language with formulas from the set of propositions P. Let $\mathcal{OL}_n$ be the epistemic language with two additional modal operators. First, K_i: $K_i\alpha$ is to be read as *"agent i knows α"*, where i ranges from the finite set of agents $Ag = \{A, ..., B\}$. Second, $O_i\alpha$ is to be read as *"all that agent i knows is α"* to express that a formula is *all* that is known. Borrowing from the dynamic version of $\mathcal{OL}_n$ [2], we introduce a dynamic operator $[\rho]$, such that $[\rho]\alpha$ is understood as formula α is *true* after receiving the information ρ.

Semantics The semantics is provided in terms of possible worlds and *k-structures* [2]. The k-structure uses sets of worlds at different levels, the idea being that the number of levels (so-called i-depth) corresponds to the number of alternating modalities in the formula. A world $w \in \mathcal{W}$ is a function from the set P to $\{0, 1\}$. We denote the set of worlds by $\mathcal{W}$ and the set of k-structures for an agent A by e^k_A. Therefore, with two agents $\{A, B\}$, models are triples (e^k_A, e^j_B, w), where e^k_A is a k-structure for A, e^j_B is a j-structure for B and w is a model. We write $e^k, e^j_B, w \models \alpha$ provided the alternating modalities for A are $\leq k$, and the alternating modalities for $B \leq j$. The language in general allows for arbitrary nestings of epistemic operators.

Sensing The observations received are represented via the action modality $[\rho]$, where ρ is an action standing for propositional conjunction, interpreted, say, as reading from a sensor. We model the sensors using the Sensing Theorem [2] and integrate them into the knowledge base. The sensing theorem establishes that $[\rho]K_A\alpha \equiv o \wedge K_A(o \rightarrow [\rho]\alpha)$, where $obs_A(\rho) = o$. By means of the Regression Theorem, we note that $[\rho]\alpha = \alpha$ because sensing actions are assumed to not affect truth in the real world. This is the essence of the Regression Theorem [2], where the application of an observational action in the context of a propositional formula yields the formula itself because sensing does not affect truth in the real world.

3 REASONING

The reasoning problem is as follows: given an epistemic knowledge base and a set of noisy partial observations received through the sensors, we would like to decide the entailment of a given query α with respect to the knowledge base and the implicitly learned hypothesis capturing the observations. In order to perform reasoning in the current setup, we are required to resolve two challenges: first, how can observations be incorporated into the knowledge, and second, how can the entailment of the query with respect to the background knowledge as well as the hypothesis be evaluated? It is important to appreciate that the second challenge deserves great attention because we are dealing with noisy observations. Roughly speaking, the way implicit learning works [7, 8] given a set of noisy observations is that the conjunction of the background knowledge together with each observation is used to check if the query formula logically follows. Suppose this happens for a high proportion of the observations. In that case, the query is accepted by the decision procedure which can be seen as implicitly including whatever formula might be captured by the high proportion of observations. So checking logical validity is an important computational component of the overall algorithm. The learning algorithm thus only runs in polynomial time if checking validity is in polynomial time.

In our setting, by way of the only-knowing modality, the reduction to propositional reasoning is achieved using the *Representation Theorem* denoted by the operator $|| \cdot ||$, first introduced by [2, 9]. Putting it together, suppose that ϕ, ψ are objective formulas and α is an epistemic formula that does not mention $\{O_i, []\}$ operators. Then $O_A(\phi \wedge O_B\psi) \models K_A\alpha$ iff $\models ||\alpha||_{\phi,\psi}$, where $||\alpha||_{\phi,\psi}$ is a propositional formula.

PAC-Semantics *PAC-Semantics* [12] is a weaker semantics (compared to classical entailment) for answering queries about background knowledge, where only a proportion of interpretations of the background knowledge are required to satisfy the query.

In the context of learning, the knowledge is represented on one hand by a collection of axioms about the world, and on the other, by a set of examples drawn from an (unknown) distribution. In this way, the algorithm uses both forms of knowledge to answer a given query, which one may not be able to answer using only the background knowledge or the standalone examples. We extend the definition of validity as follows:

Definition 1. [$(1-\epsilon)$-validity] Suppose we have a distribution $\mathcal{D}$ supported on $E_A^k \times E_B^j \times \mathcal{W}$ and α an epistemic formula $\alpha \in O\mathcal{L}_n$, that does not mention $\{O_i, []\}$ operators. Then it follows that α is $(1-\epsilon)$-valid iff $Pr_{(e_A^k, e_B^j, w) \in E_A^k \times E_B^j \times \mathcal{W}}[(e_A^k, e_B^j, w) \models \alpha] \geq 1-\epsilon$, where (e_A^k, e_B^j, w) is a model from $E_A^k \times E_B^j \times \mathcal{W}$. If $\epsilon = 0$, we say that α is perfectly valid, which then corresponds to classical validity.

As mentioned above, the key trick is to develop a decision procedure that checks the entailment of the query against the knowledge base and the observations. The proportion of times the query formula evaluates to *true* can be used as a reliable indicator of the formula's degree of validity, as guaranteed by Hoeffding's inequality [6]. The agent will have some knowledge, and will also be able to sense the world around them and receive readings describing the current state of the world. These readings are neither fully accessible nor are they exact. As a consequence, the observations can be noisy or inconsistent with each other, but they are always consistent with the knowledge base. Formally we have the notion of a *masking process* that randomly reveals only a few properties of the world [10]. These readings are conjunctions of propositional atoms and are drawn independently at random from some probability distribution $\boldsymbol{M}$ over $\mathcal{L}_n$ which is unknown to the agent. The masking process induces a probability distribution $\boldsymbol{M}(\mathcal{D})$ over observations ρ and this is aimed to model the root agent's sensors, which has two interpretations: the readings are absent due to a stochastic device failure, or the agent is unable to concurrently detect every aspect of the world. The reasoning problem of interest becomes deciding whether a query formula α is $(1-\epsilon)$-valid.

Formally, knowledge about the distribution $\mathcal{D}$ comes from the set of examples $\rho^{(c)} \in \mathcal{L}_n$. Additional knowledge comes from a collection of axioms, the knowledge base Σ. We assume here two agents A and B, but it can be generalized to multiple agents, from which agent A is the root agent. The background knowledge, for A and B respectively, is represented by $\Sigma, \Sigma' \in O\mathcal{L}_n$. The input query α is of the form $M\alpha'$, where M denotes a sequence of bounded modalities. And finally, we draw at least m partial observations which are propositional $\rho^{(1)}, \rho^{(2)}, ..., \rho^{(m)}$.

Algorithm 1: DecidePAC Implicit learning reduction

Input: Σ set of sentences from root agent A; input query α; partial observations: $\rho^{(1)}, \rho^{(2)}, ..., \rho^{(m)}$; parameters: $\epsilon, \gamma, \delta \in (0, 1)$.

Output: **Accept** if there exist formulas I witnessed *true* with probability at least $(1-\epsilon+\gamma)$ on $\boldsymbol{M}(\mathcal{D})$ such that $O_A(\Sigma) \models K_A(I \rightarrow \alpha)$
Reject if $O_A\Sigma \models K_A\alpha$ is not $(1-\epsilon-\gamma)$-valid under the distribution $\mathcal{D}$.

```
begin
    b ← ⌊ε × m⌋, FAILED ← 0.
    foreach c in m do
        if ρ^(c) ∧ O_A(Σ) ⊭ K_A(ρ^(c) → α) then
            Increment FAILED. if FAILED > b then
                return Reject
    return Accept
```

In implicit learning, the query α is answered from observations directly, without creating an explicit model. This is done by means of entailment: we repeatedly ask whether $O_A(\Sigma \wedge O_B\Sigma') \models [\rho^{(c)}]K_A\alpha$ for examples $\rho^{(c)} \in \boldsymbol{M}(\mathcal{D})$ where $c \in \{1, ..., m\}$. So this entailment checking with respect to each observation $\rho^{(c)}$ becomes our best approximation to $(1-\epsilon)$-validity. If at least $(1-\epsilon)$ of the examples entail the query α, the algorithm returns *Accept*. The estimation is more accurate the more samples we use. The concepts of accuracy and confidence are captured by the parameters $\gamma, \delta \in (0, 1)$, where γ represents the accuracy of the examples used and δ captures the confidence of the example received.

4 CONCLUSION

We have demonstrated new PAC-learning results with multi-agent epistemic logic, by leveraging the Representation Theorem to reduce modal reasoning to propositional reasoning. The algorithm is in principle applicable to any multi-agent logic, as long as a sound and complete procedure is used to evaluate epistemic queries against an epistemic knowledge base.

ACKNOWLEDGMENTS

We would like to thank our reviewers for their feedback and suggestions. Ionela Georgiana Mocanu was supported by the Engineering and Physical Sciences Research Council (EPSRC) Centre for Doctoral Training in Pervasive Parallelism (grant EP/L01503X/1) at the University of Edinburgh, School of Informatics. Vaishak Belle was partly supported by a Royal Society University Research Fellowship, UK, and partly supported by a grant from the UKRI Strategic Priorities Fund, UK to the UKRI Research Node on Trustworthy Autonomous Systems Governance and Regulation (EP/V026607/1, 2020–2024). Brendan Juba was supported by the NSF awards IIS-1908287 and IIS-1942336.

REFERENCES

[1] V. Belle and B. Juba. Implicitly learning to reason in first-order logic. *CoRR*, abs/1906.10106, 2019.
[2] V. Belle and G. Lakemeyer. Multiagent only knowing in dynamic systems. *Journal of Artificial Intelligence Research*, 49, 2014.
[3] L. De Raedt and K. Kersting. Statistical relational learning. In *Encyclopedia of Machine Learning*, pages 916–924. Springer, 2011.
[4] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
[5] L. Getoor and B. Taskar. Introduction to statistical relational learning (adaptive computation and machine learning). 2007.
[6] W. Hoeffding. *Probability Inequalities for sums of Bounded Random Variables*, pages 409–426. Springer New York, New York, NY, 1994.
[7] B. Juba. Learning implicitly in reasoning in pac-semantics. *CoRR*, abs/1209.0056, 2012.
[8] B. Juba. Implicit learning of common sense for reasoning. In *IJCAI*, pages 939–946, 2013.
[9] H. J. Levesque and G. Lakemeyer. *The Logic of Knowledge Bases*. MIT Press, 2000.
[10] L. Michael. Learning from partial observations. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence*, IJCAI'07, page 968–974, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.
[11] S. Muggleton and L. de Raedt. Inductive logic programming: Theory and methods. *The Journal of Logic Programming*, 19-20:629–679, 1994. Special Issue: Ten Years of Logic Programming.
[12] L. G. Valiant. Robust logics. *Artificial Intelligence*, 117(2):231–253, 2000.