

Learning Probabilistic Ontologies with Distributed Parameter Learning

Giuseppe Cota¹, Riccardo Zese¹, Elena Bellodi¹, Fabrizio Riguzzi², and
Evelina Lamma¹

¹ Dipartimento di Ingegneria – University of Ferrara
Via Saragat 1, I-44122, Ferrara, Italy

² Dipartimento di Matematica e Informatica – University of Ferrara
Via Saragat 1, I-44122, Ferrara, Italy
`[giuseppe.cota,riccardo.zese,elena.bellodi,evelina.lamma,
fabrizio.riguzzi]@unife.it`

Abstract. We consider the problem of learning both the structure and the parameters of Probabilistic Description Logics under DISPONTE. DISPONTE (“DISTRIBUTION Semantics for Probabilistic ONTOlogiEs”) adapts the distribution semantics for Probabilistic Logic Programming to Description Logics. The system LEAP for “LEARNING Probabilistic description logics” learns both the structure and the parameters of DISPONTE knowledge bases (KBs) by exploiting the algorithms CELOE and EDGE. The former stands for “Class Expression Learning for Ontology Engineering” and it is used to generate good candidate axioms to add to the KB, while the latter learns the probabilistic parameters and evaluates the KB. EDGE for “Em over bDds for description loGics paramEter learning” is an algorithm for learning the parameters of probabilistic ontologies from data. In order to contain the computational cost, a distributed version of EDGE called EDGE^{MR} was developed. EDGE^{MR} exploits the MapReduce (MR) strategy by means of the Message Passing Interface. In this paper we propose the system LEAP^{MR}. It is a re-engineered version of LEAP which is able to use distributed parallel parameter learning algorithms such as EDGE^{MR}.

Keywords: Probabilistic Description Logics, Structure Learning, Parameter Learning, MapReduce, Message Passing Interface.

1 Introduction

In real world domains the information is often uncertain, hence it is of foremost importance to model uncertainty in representations of the world, including Description Logics (DLs).

In [10, 16, 6] the authors studied the use of probabilistic DLs and various approaches for representing uncertainty in DLs.

Moreover, some works have started to appear about learning the probabilistic parameters or the whole structure of probabilistic ontologies. These are motivated, on one hand, from the fact that specifying the values of the probabilities is

a difficult task for humans and data is usually available that could be leveraged for tuning them, and, on the other hand, from the fact that in some domains there exist poor-structured knowledge bases which could be improved [10, 9].

In Probabilistic Logic Programming (PLP) various proposals for representing uncertainty have been presented. One of the most successful approaches is the distribution semantics [15]. In [3, 14, 11] the authors proposed an approach to represent probabilistic axioms in DLs called DISPONTE (“DIstribution Semantics for Probabilistic ONTologiEs”), which adapts the distribution semantics for Probabilistic Logic Programming to DLs.

LEAP [13] for “LEArning Probabilistic description logics” is an algorithm for learning the structure and the parameters of probabilistic DLs following DISPONTE. It combines the learning system CELOE [8] with EDGE [12]. The former, CELOE (“Class Expression Learning for Ontology Engineering”), provides a method to build new (subsumption) axioms that can be added to the KB, while the latter is used to learn the parameters of these probabilistic axioms.

EDGE stands for “Em over bDds for description lOGics paramETER learning” and learns the parameters of a probabilistic theory. This algorithm is rather expensive from a computational point of view. Therefore, in order to reduce EDGE running time, we developed EDGE^{MR} [5]. It represents a distributed implementation of EDGE and uses a simple MapReduce approach based on the Message Passing Interface (MPI).

In this paper we present an evolution of LEAP called LEAP^{MR} which adapts the LEAP algorithm to use EDGE^{MR}. In addition, due to a software re-engineering effort, it was possible to remove the RMI module used by LEAP. To the best of our knowledge there are no other algorithms that perform distributed structure learning of probabilistic DLs.

The paper is structured as follows. Section 2 introduces Description Logics and summarizes DISPONTE. Sections 3 and 4 briefly describe the EDGE and EDGE^{MR} algorithms. Section 5 presents LEAP^{MR}. Finally, Section 7 draws conclusions.

2 Description Logics and DISPONTE

Description Logics (DLs) are a family of logic based knowledge representation formalisms which are of particular interest for representing ontologies and for the Semantic Web. For an extensive introduction to DLs we refer to [1, 2].

While DLs are a fragment of first order logic, they are usually represented using a syntax based on concepts and roles. A concept corresponds to a set of individuals while a role corresponds to a set of couples of individuals of the domain.

A query over a KB is usually an axiom for which we want to test the entailment from the KB. The entailment test may be reduced to checking the unsatisfiability of a concept in the KB, i.e., the emptiness of the concept.

DISPONTE [3] (“DIstribution Semantics for Probabilistic ONTologiEs”) applies the distribution semantics to probabilistic ontologies [15]. In DISPONTE

a *probabilistic knowledge base* \mathcal{K} is a set of certain and probabilistic axioms. *Certain axioms* take the form of regular DL axioms. *Probabilistic axioms* take the form $p :: E$, where p is a real number in $[0, 1]$ and E is a DL axiom. A DISPONTE KB defines a distribution over DL KBs called *worlds* assuming that the axioms are independent. Each world w is obtained by including every certain axiom plus a subset of chosen probabilistic axioms.

3 Parameter Learning for Probabilistic DLs

EDGE [12] is a parameter learning algorithm which adapts the algorithm EM-BLEM [4], developed for learning the parameters for probabilistic logic programs, to the case of probabilistic DLs under DISPONTE. Inspired by [7], it performs an Expectation-Maximization cycle over Binary Decision Diagrams (BDDs).

EDGE performs supervised parameter learning. It takes as input a DISPONTE KB and a number of positive and negative examples that represent the queries in the form of concept membership axioms, i.e., in the form $a : C$ for an individual a and a class C .

First, EDGE generates, for each query, the BDD encoding its explanations using BUNDLE [14]. Then, EDGE starts the EM cycle in which the steps of Expectation and Maximization are iterated until a local maximum of the log-likelihood (LL) of the examples is reached. The LL of the examples is guaranteed to increase at each iteration. EDGE stops when the difference between the LL of the current iteration and that of the previous one drops below a threshold ϵ or when this difference is below a fraction δ of the previous LL . Finally, EDGE returns the reached LL and the new probabilities for the probabilistic axioms.

EDGE is written in Java, hence it is highly portable. For further information about EDGE please refer to [12].

4 Distributed Parameter Learning for Probabilistic DLs

In this section we briefly describe a parallel version of EDGE that exploits the MapReduce approach in order to compute the parameters. We called this algorithm EDGE^{MR} [5].

Like most MapReduce frameworks, EDGE^{MR} 's architecture follows a master-slave model. The communication between the master and the slaves is done by means of the Message Passing Interface (MPI).

In a distributed context, the performances depend on the scheduling strategy. In order to evaluate different methods, we developed two scheduling strategies: *single-step scheduling* and *dynamic scheduling*. These are used during the queries computation phase.

Single-step Scheduling if N is the number of the slaves, the master divides the total number of queries into $N + 1$ chunks, i.e. the number of slaves plus the master. Then the master begins to compute its queries while, for the other chunks of queries, the master starts a thread for sending each chunk

to the corresponding slave. After the master has terminated dealing with its queries, it waits for the results from the slaves. When the slowest slave returns its results to the master, EDGE^{MR} proceeds to the EM cycle.

Dynamic Scheduling is more flexible and adaptive than single-step scheduling. At first, each machine is assigned a fixed-size chunk of queries in order. Then, if the master ends handling its chunk it just takes the next one, instead, if a slave ends handling its chunk, it asks the master for another one and the master replies by sending a new chunk of queries to the slave. During this phase the master runs a thread listener that waits for the slaves' requests of new chunks and for each request the listener starts a new thread that sends a chunk to the slave which has done the request (to improve the performances this is done through a thread pool). When all the queries are evaluated, EDGE^{MR} starts the EM cycle.

Experimental results conducted in [5] show that dynamic scheduling has usually better performances than single-step.

5 Structure Learning with Distributed Parameter Learning

LEAP^{MR} is an evolution of the LEAP system [13]. While the latter exploits EDGE, the first was adapted to be able to perform EDGE^{MR} . Moreover, after a process of software re-engineering it was possible to remove the RMI communication module used by LEAP and therefore reduce some communication overhead.

It performs structure and parameter learning of probabilistic ontologies under DISPONTE by exploiting: (1) CELOE [8] for the structure, and (2) EDGE^{MR} (Section 4) for the parameters.

CELOE [8] was implemented in Java and belongs to the open-source framework DL-Learner³. Let us consider a knowledge base \mathcal{K} and a concept name **Target** whose formal description, i.e. class description, we want to learn. It learns a set of n class expressions C_i ($1 \leq i \leq n$) from a set of positive and negative examples. Let $\mathcal{K}' = \mathcal{K} \cup \{C\}$ where \mathcal{K} is the background knowledge, we say that a concept C covers an example e if $\mathcal{K}' \models e$. The class expressions found are sorted according to a heuristic. Such expressions can be used to generate candidate axioms of the form $C_i \sqsubseteq \text{Target}$.

In order to learn an ontology, LEAP^{MR} first searches for good candidate probabilistic subsumption axioms by means of CELOE, then it performs a greedy search in the space of theories using EDGE^{MR} to evaluate the theories using the log-likelihood as heuristic.

LEAP^{MR} takes as input the knowledge base \mathcal{K} and a set of examples, then generates a set of candidate axioms by exploiting CELOE. A first execution of EDGE^{MR} is applied to \mathcal{K} to compute the initial value of the parameters and of the LL . Then LEAP^{MR} adds to \mathcal{K} one probabilistic subsumption axiom

³ <http://dl-learner.org/>

generated from CELOE. After each addition, EDGE^{MR} is performed on the extended KB to compute the LL of the data and the parameters. If the LL is better than the current best, the new axiom is kept in the knowledge base and the parameter of the probabilistic axiom are updated, otherwise the learned axiom is removed from the ontology and the previous parameters are restored. The final theory is obtained from the union of the initial ontology and the probabilistic axioms learned.

6 Experiments

In order to test how much the exploitation of EDGE^{MR} can improve the performances of LEAP^{MR} , we did a preliminary test where we considered the Moral⁴ KB which qualitatively simulates moral reasoning. It contains 202 individuals and 4710 axioms (22 axioms are probabilistic).

We performed the experiments on a cluster of 64-bit Linux machines with 8-cores Intel Haswell 2.40 GHz CPUs and 2 GB (max) memory allotted to Java per node. We allotted 1, 3, 5, 9 and 17 nodes, where the execution with 1 node corresponds to the execution of LEAP, while for the other configurations we used the dynamic scheduling with chunks containing 3 queries. For each experiment 2 candidate probabilistic axioms are generated by using CELOE and a maximum of 3 explanations per query was set for EDGE^{MR} . Table 1 shows the speedup obtained as a function of the number of machines (nodes). The speedup is the ratio of the running time of 1 worker to the one of n workers. We can note that the speedup is significant even if it is sublinear, showing that a certain amount of overhead (the resources, and thereby the time, spent for the MPI communications) is present.

Dataset	N. of Nodes			
	3	5	9	9
Moral	2.3	3.6	6.5	11.0

Table 1. Speedup of LEAP^{MR} relative to LEAP for Moral KB.

7 Conclusions

The paper presents the algorithm LEAP^{MR} for learning the structure of probabilistic description logics under DISPONTE. LEAP^{MR} performs EDGE^{MR} which is a MapReduce implementation of EDGE, exploiting modern computing infrastructures for performing distributed parameter learning.

We are currently working for distributing both the structure and the parameter learning of probabilistic knowledge bases by exploiting EDGE^{MR} also during

⁴ <https://archive.ics.uci.edu/ml/datasets/Moral+Reasoner>

the building of the class expressions. In particular we would like to distribute the scoring function used to evaluate the obtained refinements.

References

1. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, New York, NY, USA (2003)
2. Baader, F., Horrocks, I., Sattler, U.: *Description Logics*, chap. 3, pp. 135–179. Elsevier Science (2008)
3. Bellodi, E., Lamma, E., Riguzzi, F., Albani, S.: A distribution semantics for probabilistic ontologies. In: *International Workshop on Uncertainty Reasoning for the Semantic Web*. CEUR Workshop Proceedings, vol. 778, pp. 75–86. Sun SITE Central Europe (2011)
4. Bellodi, E., Riguzzi, F.: Expectation Maximization over Binary Decision Diagrams for probabilistic logic programs. *Intell. Data Anal.* 17(2), 343–363 (2013)
5. Cota, G., Zese, R., Bellodi, E., Lamma, E., Riguzzi, F.: Distributed parameter learning for probabilistic ontologies (2015), to appear
6. Fleischhacker, D., Völker, J.: Inductive learning of disjointness axioms. In: *On the Move to Meaningful Internet Systems: OTM 2011*, pp. 680–697. Springer (2011)
7. Ishihata, M., Kameya, Y., Sato, T., Minato, S.: Propositionalizing the EM algorithm by BDDs. In: *Late Breaking Papers of the International Conference on Inductive Logic Programming*. pp. 44–49 (2008)
8. Lehmann, J., Auer, S., Bühmann, L., Tramp, S.: Class expression learning for ontology engineering. *J. Web Semant.* 9(1), 71–81 (2011)
9. Minervini, P., d’Amato, C., Fanizzi, N.: Learning probabilistic description logic concepts: Under different assumptions on missing knowledge. In: *Proceedings of the 27th Annual ACM Symposium on Applied Computing*. pp. 378–383. ACM (2012)
10. Ochoa-Luna, J.E., Revoredo, K., Cozman, F.G.: Learning probabilistic description logics: A framework and algorithms. In: *Advances in Artificial Intelligence*, pp. 28–39. Springer (2011)
11. Riguzzi, F., Bellodi, E., Lamma, E., Zese, R.: Epistemic and statistical probabilistic ontologies. In: *Uncertainty Reasoning for the Semantic Web*. CEUR Workshop Proceedings, vol. 900, pp. 3–14. Sun SITE Central Europe (2012)
12. Riguzzi, F., Bellodi, E., Lamma, E., Zese, R.: Parameter learning for probabilistic ontologies. In: *RR 2013. LNCS*, vol. 7994, pp. 265–270. Springer Berlin Heidelberg (2013)
13. Riguzzi, F., Bellodi, E., Lamma, E., Zese, R., Cota, G.: Learning probabilistic description logics. In: *Uncertainty Reasoning for the Semantic Web III*, pp. 63–78. LNCS, Springer International Publishing (2014)
14. Riguzzi, F., Lamma, E., Bellodi, E., Zese, R.: BUNDLE: A reasoner for probabilistic ontologies. In: *RR 2013. LNCS*, vol. 7994, pp. 183–197. Springer Berlin Heidelberg (2013)
15. Sato, T.: A statistical learning method for logic programs with distribution semantics. In: *Proceedings of the 12th International Conference on Logic Programming*. pp. 715–729. MIT Press (1995)
16. Völker, J., Niepert, M.: Statistical schema induction. In: *The Semantic Web: Research and Applications*, pp. 124–138. Springer Berlin Heidelberg (2011)