International Journal of Approximate Reasoning 82 (2017) 39–55

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

# Scalable learning and inference in Markov logic networks

Zhengya Sun *, Yangyang Zhao, Zhuoyu Wei, Wensheng Zhang, Jue Wang

*Institute of Automation, Chinese Academy of Sciences, China*

ARTICLE INFO

*Article history:*
Received 10 March 2016
Received in revised form 2 September 2016
Accepted 2 December 2016
Available online 8 December 2016

*Keywords:*
Markov logic networks
Structure learning
Probabilistic inference
Large scale machine learning

ABSTRACT

Markov logic networks (MLNs) have emerged as a powerful representation that incorporates first-order logic and probabilistic graphical models. They have shown very good results in many problem domains. However, current implementations of MLNs do not scale well due to the large search space and the intractable clause groundings, which is preventing their widespread adoption. In this paper, we propose a general framework named Ground Network Sampling (GNS) for scaling up MLN learning and inference. GNS offers a new instantiation perspective by encoding ground substitutions as simple paths in the Herbrand universe, which uses the interactions existing among the objects to constrain the search space. To further make this search tractable for large scale problems, GNS integrates random walks and subgraph pattern mining, gradually building up a representative subset of simple paths. When inference is concerned, a template network is introduced to quickly locate promising paths that can ground given logical statements. The resulting sampled paths are then transformed into ground clauses, which can be used for clause creation and probabilistic inference. The experiments on several real-world datasets demonstrate that our approach offers better scalability while maintaining comparable or better predictive performance compared to state-of-the-art MLN techniques.

© 2016 Elsevier Inc. All rights reserved.

## 1. Introduction

Markov logic networks (MLNs) have gained wide attention in recent years due to their ability to unify the strengths of first-order logic and probabilistic graphical models [1,2]. Recent advances in MLNs have revealed their superiority in dealing with real world data, such as biochemistry [3,4], computer vision [5,6] and text [7–9]. Typically, an MLN attaches weights to first-order clauses and views these as templates for Markov network features. On the one hand, learning MLN structure from data contributes to capturing uncertain dependencies underlying the relational data in the form of clauses together with their weights [10–12]. On the other hand, inference in MLNs allows us to reason probabilistically about complex relationships according to some uncertain logical structure [13–15].

The learning and inference tasks in MLNs, although important, can be computationally prohibitive with the continuous increase in data size. Inference is crucial in structure learning, as they are intrinsically related to the problem of counting true groundings [16]. Besides, structure learning not only has a super-exponentially large search space, but also produces an overwhelming number of groundings during clause evaluation [17]. In brief, the exhaustive search and full grounding may fail to finish in a reasonable amount of time, which severely limits the size of domains they can be applied to.

* Corresponding author.
*E-mail address:* zhengya.sun@ia.ac.cn (Z. Sun).

http://dx.doi.org/10.1016/j.ijar.2016.12.003
0888-613X/© 2016 Elsevier Inc. All rights reserved.

The search problem has motivated the quest for relational pathfinding [18] that exploits the training data to constrain the space of candidates, by finding paths of true ground atoms linked via their arguments and generalizing them into first-order clauses. As each path corresponds to a conjunction that has at least one support in the data, this allows the search to focus on regions with promising clauses. However, pathfinding typically requires exhaustive search over an exponential number of paths. Hence, bottom-up structure learning [11] only uses relational pathfinding to produce very short paths, with the resulting short clauses greedily joined into longer ones. Hypergraph lifting [19] and structural motif learning [20] on the other hand employ some form of structural clustering to induce a more compact representation, and facilitate efficient relational pathfinding for clause discovery. Unfortunately, clustering the observed data can itself be quite expensive, especially in domains involving a huge number of objects and complex relations among them, due to the combinatorics involved.

The intractability of counting how often the grounded clauses are satisfied promotes numerous studies on approximate inference algorithms [1,21,22]. Among them, Markov Chain Monte Carlo (MCMC) [23] sampling is arguably the most popular one, which approximates the target probability by generating a sequence of samples from a well-designed Markov chain. Despite its successes, the applicability of MCMC inference in MLNs is severely limited by the exponential cost of grounding all the logical statements. Several lifted inference algorithms improve upon MCMC by grouping random variables that are symmetrical give the first-order structure, and then sampling over the high-level representation [24–26]. However, they tend to lose their advantage in the presence of evidence, as evidence breaks symmetries. Another way to avoid full grounding is to perform lazy inference which achieves both memory and time savings by utilizing the relational sparseness, i.e., only a small fraction of ground atoms are true, and the great majority of clauses are trivially satisfied [14,27]. As lazy inference algorithms only ground atoms and clauses as needed, they incur the overhead in checking whether an atom or clause is in memory, retrieving and grounding the clauses activated by an atom. They may have increased inference time when ground atoms are updated frequently, along with the relevant clauses activated correspondingly.

Our current study aims at addressing the above problems in a unified and scalable manner. Based on MCMC simulation, we propose a novel MLN learning and inference framework named Ground Network Sampling (GNS) which encodes ground substitutions as simple paths in the Herbrand universe, i.e., a sequence of interactions among distinct objects.

GNS starts by converting the unary atoms (i.e., an atom that involves only one argument) into binary ones, so that the already-present ground atoms can all be treated as interactions between their arguments. Intuitively, the objects interact when they appear in a ground atom. The resulting argument pairs that share common objects are then concatenated in order of increasing length, hence forming a partial order graph (POG). Every node in the POG corresponds to an effective path which may describe whether or not the initial and terminal objects interact with each other. Subsequently, GNS performs a random walk on the POG with a prescribed transition probability matrix, and returns every target node it visits. On the one hand, GNS creates ground clauses from the uniformly sampled paths, and variabilizes them for discriminate weight learning. On the other hand, when inference is concerned, GNS adapts the uniform sampler to quickly locate promising paths that can ground the given logical statements. Based on the resulting ground network, GNS returns the probability estimations for query predicates. The experimental results on several real-world datasets validate the efficiency and efficacy of our proposed approach compared to the state-of-the-arts.

The main contributions of this paper are summarized as follows:

- We offer a new instantiation perspective by encoding the ground substitutions as simple paths in the Herbrand universe, and thus is suitable for obtaining representative ground networks for learning and inference in MLNs, where full grounding is infeasible or very expensive.
- We combine the benefits of random walks and subgraph pattern mining, which avoids exploring the entire POG and can speed up the search significantly. We also discuss the convergence rate to the stationary distribution. The output set consists of the simple paths that correspond to the non-trivial groundings of the discriminative clauses.
- We ensure efficient inference by leveraging the template network to locate promising paths that can ground the given clauses instead of all possible clauses. This amounts to performing a constrained random walk through the POG, with the resulting samples being used to estimate the probabilities of the query predicates.

The remainder of this paper is organized as follows. Section 2 provides the background and some related work on MLNs. Section 3 presents our scalable learning and inference approach, and the experimental results are reported in Section 4. Section 5 summarizes the conclusion of this work.

## 2. Markov logic networks

This section provides some background on representation, structure learning, discriminative weight learning, and conditional inference for Markov logic networks. Let us briefly review some basic terminology which will be used throughout this paper. Constants starting with an uppercase letter (e.g. Bob) represent objects in the domain of interest. A lowercase letter (e.g. $a$) indicates a variable ranging over the objects. Predicates represent relations among objects or attributes of objects. An atom is a predicate symbol applied to a tuple of arguments, which may be variables or constants. A ground atom is an atom all of whose arguments are constants. Variables and constants are often typed. A positive literal is an atom, and a

Download English Version:

https://daneshyari.com/en/article/4945309

Download Persian Version:

https://daneshyari.com/article/4945309

Daneshyari.com