Journal of Applied Logic 13 (2015) 414–440

Contents lists available at ScienceDirect

Journal of Applied Logic

www.elsevier.com/locate/jal

A concept for the evolution of relational probabilistic belief states and the computation of their changes under optimum entropy semantics

Nico Potyka [a,*], Christoph Beierle [a,**], Gabriele Kern-Isberner [b,**]

[a] *Faculty of Mathematics and Computer Science, University of Hagen, 58084 Hagen, Germany*
[b] *Department of Computer Science, University of Technology Dortmund, 44227 Dortmund, Germany*

A R T I C L E I N F O

Article history:
Received 11 February 2014
Available online 8 January 2015

Keywords:
Probabilistic logic
Conditional logic
Relational conditional
First-order knowledge base
Maximum entropy
Minimum relative entropy
Belief state
Belief change
Iterative belief changes
Computation of belief changes

A B S T R A C T

Coping with uncertain knowledge and changing beliefs is essential for reasoning in dynamic environments. We generalize an approach to adjust probabilistic belief states by use of the relative entropy in a propositional setting to relational languages, leading to a concept for the evolution of relational probabilistic belief states. As a second contribution of this paper, we present a method to compute the corresponding belief state changes by considering a dual problem and present first application and experimental results. The computed belief state usually factorizes and we explain how the factorization can be exploited to store the belief state more compactly and to simplify its computation by exploiting equivalences of worlds. Finally, we present results on the computational complexity of determining equivalence classes.

1. Introduction

When autonomous agents act in dynamic environments, they have to deal with information that is uncertain and subject to changes. Over the years, different approaches have been developed to deal with both problems [1,18,29,14,24]. In this work, we will consider a framework based on probabilistic conditional logics [31,26]. As an illustration and running example, we assume an application scenario where an agent has to watch some pets. Sometimes the pets attack each other and our agent has to separate them. The agent's knowledge base might contain *deterministic conditionals* like $(Bird(X) \wedge Dog(X))[0]$ expressing that a pet cannot be both a bird and a dog. There may also be *non-deterministic conditionals* like

* Principal corresponding author.
** Corresponding author.
E-mail addresses: nico.potyka@fernuni-hagen.de (N. Potyka), christoph.beierle@fernuni-hagen.de (C. Beierle), gabriele.kern-isberner@cs.uni-dortmund.de (G. Kern-Isberner).

http://dx.doi.org/10.1016/j.jal.2015.01.001
1570-8683/

$(Attacks(X,Y) \mid GT(Y,X))[0.1]$ expressing that the (subjective) probability that an animal attacks a greater one is 10%, where the binary predicate $Attacks(X,Y)$ stands for X *attacks* Y and $GT(X,Y)$ expresses X *is greater than* Y. When we receive further information, e.g., we come to know that the pet is definitely a dog, or we read about changes in the general behaviour of small dogs in a magazine, we have to adapt our beliefs to this new information, requiring a concept for the evolution of belief states.

A probabilistic conditional semantics defines which probability distributions satisfy a probabilistic conditional. However, interpreting probabilistic conditionals built from open formulas with variables is not an easy task as there is no established understanding of such formulas. In any case, semantics for first-order probabilistic conditionals should coincide with the usual propositional conditional semantics which is defined via conditional probabilities in those cases where the conditional does not contain any variables. So, we will first explore the propositional case in a bit more depth before addressing first-order, or relational issues, respectively. When a (propositional) probabilistic conditional knowledge base $\mathcal{R}$ is given, we are often interested in a best distribution satisfying the conditionals in $\mathcal{R}$ that we can use for inferences. An appropriate selection criterion is the principle of maximum entropy [31] that, intuitively speaking, states that we should select the most uniform one among all satisfying distributions. To adapt the distribution to new knowledge, we will consider the distribution that minimizes the relative entropy to the prior distribution [19]. Even though relative entropy is not a metric, it satisfies some geometric properties similar to the Euclidean distance and can be regarded as an information-theoretic distance measure [3].

Following this idea, our agent's knowledge state consists of a knowledge base $\mathcal{R}$ reflecting her explicit knowledge and a probability distribution $\mathcal{P}$ corresponding to background knowledge. The unique distribution that satisfies the conditionals in $\mathcal{R}$ and minimizes relative entropy with respect to $\mathcal{P}$ reflects her epistemic state. According to [20], we distinguish between two belief change operations. *Revision* deals with new information in a static world. That is, prior explicit knowledge remains valid, even though our epistemic state may change. For instance, we might learn about the pet *bully* that it is a dog. Then $\mathcal{P}(Dog(bully))$ should become 1, but the prior explicit knowledge in $\mathcal{R}$ should remain valid. *Update* deals with new information in a dynamic world. Therefore, the new knowledge might be in conflict with the prior explicit knowledge. For example, if we observe that small pets are getting more aggressive, e.g., due to a newly added ingredient to their food, we should increase our belief that an animal attacks a greater animal. In this case, the explicit knowledge can be subject to change as well. This is in line with the distinction that is usually made in belief change theory. However, the classical AGM-theory [1] is much too weak to be able to handle such advanced change operations, as it can deal neither with probabilities nor with conditionals. In [20] it is shown how both operations can be implemented for a propositional probabilistic language by use of the relative entropy, and in [2] a corresponding conceptual agent model providing a series of powerful belief management operations is developed. The MECore system [9] implements these ideas and allows the user, e.g., to define an initial knowledge base, to apply both belief change operations, or to query the current epistemic state.

Note, however, that all approaches mentioned so far only deal with propositional logic and do not cover conditionals with variables like $(Attacks(X,Y) \mid GT(Y,X))[0.1]$. In this paper, we generalize the belief change operations from [20] to relational languages. Our results hold for a full class of conditional semantics and in particular for the relational *grounding* and *aggregating* semantics [11,22]. Basically, all belief change operations can be reduced to the core functionality of minimizing relative entropy with respect to a prior distribution.

Besides providing a concept for changing relational probabilistic belief states, the second major contribution of this paper is an alternative approach to compute such belief changes. We consider the Wolfe dual of relative entropy minimization [12], which yields an unconstrained convex optimization problem, and solve it with L-BFGS [40]. For entropy maximization under aggregating semantics, a significant performance gain compared to a recent iterative scaling implementation [7] is obtained.