Artificial Intelligence 229 (2015) 33–104

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

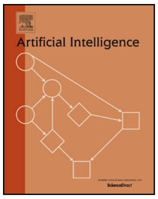

On updates of hybrid knowledge bases composed of ontologies and rules ☆

Martin Slota *, João Leite **, Theresa Swift *

NOVA LINCS & Departamento de Informática, Universidade Nova de Lisboa, Quinta da Torre, 2829-516 Caparica, Portugal

ARTICLE INFO

Article history:
Received 22 January 2014
Received in revised form 27 July 2015
Accepted 30 July 2015
Available online 4 August 2015

Keywords:
Hybrid knowledge bases
Ontologies
Description logics
Logic programs
Rules
Stable models
Answer sets
Belief change
Update

ABSTRACT

Throughout the last decade, two distinct knowledge representation paradigms have been standardised to capture rich metadata on the Web: *ontology languages* based on Classical Logic and *reasoning rules* based on Logic Programming. Both offer important features for knowledge representation and the interest in their integration has recently resulted in frameworks for *hybrid knowledge bases* that consist of an ontology and a rule component.

Instead of the usual static view of hybrid knowledge, in this paper we address its *dynamics* and in particular focus on *updates*. We develop two hybrid update semantics that fit the needs of particular use cases of hybrid knowledge and provide the expected results when used in specific application domains. The first semantics uses a given ontology update operator to update the ontology component of a hybrid knowledge base *in the presence of static rules*. Inspired by a realistic application, and based on a generalised notion of *splitting*, known from Logic Programming, the second semantics offers a way to *modularly combine* an ontology update operator with a rule update semantics. It can be used for performing updates of hybrid knowledge bases consisting of ontology and rule layers that share information through a rule-based interface.

Both of these developments constitute solutions to the problem of hybrid updates for restricted classes of hybrid knowledge bases. We examine their fundamental formal properties and show that despite the different ideas behind each of them, they are fully compatible with one another, i.e. when both are applicable, they lead to the same result.

1. Introduction

Recent standardisation efforts gave rise to widely accepted knowledge representation languages such as the Web Ontology Language (OWL)[1] and Rule Interchange Format (RIF),[2] based on Description Logics [11] and Logic Programming [22,42,61,74], respectively. This has fostered a large number of ontologies and rule bases with different levels of complexity and scale. Whereas description-logic based ontologies provide the logical underpinning of intelligent access and

---

☆ This is a combined, revised and extended version of the material presented in [83,84,91].

* Corresponding author.

** Primary corresponding author.

E-mail addresses: martin.slota@gmail.com (M. Slota), jleite@fct.unl.pt (J. Leite), theresasturn@gmail.com (T. Swift).

[1] http://www.w3.org/TR/owl-overview/.

[2] http://www.w3.org/2005/rules/wiki/RIF_Working_Group.

http://dx.doi.org/10.1016/j.artint.2015.07.008

information integration, rules are widely used to represent business policies, regulations and declarative guidelines about information.

Since both ontologies and rules offer important features for knowledge representation, considerable effort has been invested in identifying a unified hybrid knowledge framework where expressivity of both formalisms could be seamlessly combined. Over the years, work on hybrid knowledge bases has matured significantly and fundamental semantic as well as computational problems have been successfully addressed (see [25,52] for an overview). MKNF Knowledge Bases [77] form one of the most mature proposals. MKNF knowledge bases provide a tight integration of the two paradigms, allowing predicates to be defined concurrently in the ontology and using rules, while at the same time being faithful to the semantics both of their ontology component and their rules component. Further, they have thoroughly examined decidability and complexity properties including a tractable variant of the formalism based on the well-founded semantics [41,59].

While such formalisms make it possible to seamlessly combine rules and ontologies in a single unified framework, they do not take into account the *dynamic character* of application areas where they are to be used. More particularly, the essential support for keeping a hybrid knowledge base *up to date*, by incorporating new and possibly conflicting information, is still missing. Nonetheless, this topic has been extensively addressed in the context of both Description Logics and Logic Programs, when taken separately.

The formal underpinning of ontology updates originates in the ample area of belief change [3], more particularly in the principles and operators used to deal with change in action theories and relational databases with incomplete information [56,57,95,96]. An *update* is generally understood as a change operation that consists in bringing a knowledge base up to date when the *world described by it changes* [56]. More technically, update operators are based on the idea that the models of a knowledge base correspond to possible states of the represented world. When a change in the world needs to be recorded, inertia is applied to each of these possible states, making only the smallest necessary modifications to reflect the change, arriving at a new collection of possible states that represent the world after the update. Later, these ideas, and particularly Winslett's update operator [57,95], were applied to perform updates of Description Logic ontologies [15,27–29,36,58,73].

When updates were tackled in the context of Logic Programming, it was only natural to consider adapting the classical update postulates and operators to deal with them. However, this led to counterintuitive results because simply applying inertia to the current stable model, or even a richer semantic characterisation, such as SE-models [93], results in the loss of essential relationships between literals that are encoded in the rules [66,85,90]. Although state-of-the-art approaches to rule updates are guided by the same basic intuitions and aspirations as ontology updates, they build upon fundamentally different principles and methods.

Many are based on the *causal rejection principle* [6,7,37,66,78], while others employ syntactic transformations and other methods, such as abduction [80], forgetting [98], prioritisation [97], preferences [31], or dependencies on defeasible assumptions [82]. Inertia, instead of being applied to the current state, is applied to *rules*. Furthermore, rather than operating on the models of a logic program, rule update semantics refer to its syntactic structure: the individual rules and, in many cases, also the literals in heads and bodies of these rules. These properties render them seemingly irreconcilable with ontology updates where *state inertia* and *syntax-independence* are central.

In order to define *generic* hybrid update operators i.e., update operators that can deal with arbitrary hybrid knowledge bases, these apparently irreconcilable approaches to dealing with knowledge updates must first be integrated, both semantically and computationally. Despite the recent developments, which already provide a narrow bridge between ontology and rule updates [87–89], this problem is still far from having a suitable solution.

In Section 3 we take an important first step to defining a generic hybrid update operator by following a different approach. In particular, we show how the static semantics for MKNF knowledge bases can be adapted to allow for *updates of the ontology component* of a hybrid knowledge base while the rule component remains static. This encompasses practical applications of hybrid knowledge bases where the ontology contains highly dynamic information and rules represent business policies, preferences or behaviour that does not change or can be maintained manually, and can be overridden by ontology updates when necessary. For illustrative purposes, consider the following simple scenario:

Example 1 *(Electronic marketplace).* An electronic marketplace where agents sell and exchange items, resources and services needs to keep track of all the active users and sellers, available offers, product categorisation, etc. Though most of this information is kept within a Description Logic ontology, rules are used where the Closed World Assumption or reasoning with exceptions is needed. A small fragment of this hybrid knowledge looks as follows:

$$\mathsf{Seller} \sqsubseteq \mathsf{User} \tag{1}$$

$$\mathsf{ProspectiveSeller} \equiv \neg\mathsf{Seller} \sqcap \exists\mathsf{RecommendedBy}.\mathsf{Seller} \tag{2}$$

$$\neg\mathsf{User}(\mathbf{x}) \leftarrow {\sim}\mathsf{User}(\mathbf{x}). \tag{3}$$

$$\mathsf{PaysServiceFee}(\mathbf{x}) \leftarrow \mathsf{Seller}(\mathbf{x}), {\sim}\mathsf{Student}(\mathbf{x}). \tag{4}$$

Ontology axioms (1) and (2) express that every Seller in the marketplace is also its User and that every individual that is currently not a seller but has been recommended by some seller is considered a ProspectiveSeller (e.g. for marketing purposes). Rule (3) expresses that the concept User is interpreted under the Closed World Assumption, i.e. it is assumed