Artificial Intelligence 218 (2015) 1–22

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Belief revision in Horn theories ☆

James P. Delgrande [a], Pavlos Peppas [b,c,*]

[a] School of Computing Science, Simon Fraser University, Burnaby, B.C., V5A 1S6, Canada
[b] Centre for Quantum Computation and Intelligent Systems, FEIT, University of Technology, Sydney, NSW 2007, Australia
[c] Dept of Business Administration, University of Patras, Patras 265 00, Greece

ARTICLE INFO

Article history:
Received 1 May 2013
Received in revised form 24 July 2014
Accepted 27 August 2014
Available online 10 September 2014

Keywords:
Knowledge representation
Belief revision
Horn logic

ABSTRACT

This paper investigates belief revision where the underlying logic is that governing Horn clauses. We show that classical (AGM) belief revision doesn't immediately generalise to the Horn case. In particular, a standard construction based on a total preorder over possible worlds may violate the accepted (AGM) postulates. Conversely, in the obvious extension to the AGM approach, Horn revision functions are not captured by total preorders over possible worlds. We address these difficulties by introducing two modifications to the AGM approach. First, the semantic construction is restricted to "well behaved" orderings, what we call *Horn compliant* orderings. Second, the revision postulates are augmented by an additional postulate. Both restrictions are redundant in the AGM approach, but not in the Horn case. In a representation result we show that the class of revision functions captured by Horn compliant total preorders over possible worlds is precisely that given by the (extended) set of Horn revision postulates. Further, we show that Horn revision is compatible with work in iterated revision and work concerning relevance in revision. We also consider specific revision operators. Arguably this work is interesting for several reasons. It extends AGM revision to inferentially-weaker Horn theories; hence it sheds light on the theoretical underpinnings of belief change, as well as generalising the AGM paradigm. Thus, this work is relevant to revision in areas that employ Horn clauses, such as deductive databases and logic programming, as well as areas in which inference is weaker than classical logic, such as in description logic.

1. Introduction

The area of belief change studies how an agent may modify its beliefs given new information about its environment. The best-known approach to belief change is the AGM paradigm [1,16], named after the original developers. This work focussed on belief *revision*, in which new information is incorporated into an agent's belief corpus, as well as belief *contraction*, in which an agent may reduce its set of beliefs. The AGM approach addresses belief change at an abstract level, in which an agent's beliefs are characterised by *belief sets* or deductively closed sets of sentences, and where the underlying logic includes classical propositional logic. In this approach to revision, a set of rationality postulates is given which arguably any revision function should satisfy. As well, a semantic construction of revision functions has been given, in terms of a total

☆ This article is an extension and elaboration of the work published in [9].

* Corresponding author.
E-mail addresses: jim@cs.sfu.ca (J.P. Delgrande), pavlos@upatras.gr (P. Peppas).

http://dx.doi.org/10.1016/j.artint.2014.08.006
0004-3702/

preorder over possible worlds, called a *faithful ranking* [18]. These syntactic and semantic approaches have been shown to capture the same set of revision functions.

In this paper we address AGM-style belief revision in the language of *Horn clauses*, where a Horn clause can be expressed as a rule of the form $a_1 \wedge a_2 \wedge \cdots \wedge a_n \rightarrow a$ for $n \geq 0$, and where a, a_i $(1 \leq i \leq n)$ are atoms, and a is an atom or the constant falsum $\bot$. (Thus, expressed in conjunctive normal form, a Horn clause is a clause with at most one positive literal.) In our approach, an agent's beliefs are represented by a Horn clause belief set, and the input is a Horn formula, consisting of a conjunction of Horn clauses. It proves to be the case that AGM-style belief revision doesn't transfer directly to Horn knowledge bases. On the one hand, in the Horn case the AGM postulate set is unsound with respect to faithful rankings over possible worlds. On the other hand, given a Horn revision function that satisfies the AGM postulates, there may be no corresponding faithful ranking that captures the revision function or, alternately, there may be several faithful rankings that capture the function.

Nonetheless, we show that the AGM approach can be extended to the Horn case naturally and satisfactorily. On the semantic side, we impose a "well-behaved" condition on faithful rankings, expressing that a ranking must be coherent with respect to Horn revision. We call such rankings *Horn compliant*. On the syntactic, postulational, side, we add a postulate to the standard suite of AGM postulates. Interestingly, in the AGM approach this additional postulate is redundant, in that it follows as a theorem from the other AGM postulates. In the Horn case, in which the language is less expressive than in the classical case, this postulate is independent of the others. Given these adjustments to the AGM approach, we then prove a representation result, showing that the class of Horn revision functions conforming to the extended postulate set is the same as those capturable by Horn compliant faithful rankings. Moreover, we prove Horn revision, as modelled herein, is consistent with Darwiche and Pearl postulates for iterated revision [8] and with Parikh's postulate for relevance [22]. A final contribution of our work is the development of two specific Horn revision operators, called *basic Horn revision* and *canonical Horn revision*, with *polynomial* time complexity ($O(n)$ and $O(n^2 \log n)$ respectively).

This topic is interesting for several reasons. It sheds light on the theory of belief change, in that it weakens the assumption that the underlying logic contains propositional logic. In doing so, it shows that the AGM approach is more generally applicable than perhaps originally believed. That is, our results provide a *broadening* of the AGM approach to include Horn reasoning, and not just a modification of the AGM approach to accommodate Horn reasoning. Horn clauses are a very useful restriction of classical logic, and have found widespread application in artificial intelligence and database theory. As well, results here may also be relevant to belief change in description logics, a topic that has also received recent attention.

The next section gives basic notation and definitions used in the paper. The third section introduces belief change and related work that has been carried out in belief change in Horn clause reasoning. This is followed by a discussion of issues that arise in Horn clause belief revision (Section 4). Section 5 develops the approach, and in particular presents the representation result for Horn formula revision. Section 6 discusses iterated Horn revision, while Section 7 discusses relevance and Horn revision, and Section 8 introduces specific Horn revision operators and examines their computational complexity. The paper concludes with a discussion of future work and a brief conclusion.

2. Formal preliminaries

We introduce here the terminology that we will use in the rest of the paper. $\mathcal{P} = \{a, b, c, \ldots\}$ is a finite set of propositional variables. $\mathcal{L}_H$ denotes the Horn formula language over $\mathcal{P} \cup \{\bot\}$. That is, $\mathcal{L}_H$ is the least set given by:

1. $a_1 \wedge a_2 \wedge \cdots \wedge a_n \rightarrow a$ is a *Horn clause*, where $n \geq 0$ and $a, a_i \in \mathcal{P} \cup \{\bot\}$ for $1 \leq i \leq n$.
   If $n = 0$ then $\rightarrow a$ is also written a, and is a *fact*.
2. If ϕ is a Horn clause then (ϕ) is a *Horn formula*.
3. If ϕ and ψ are Horn formulas then so is $(\phi \wedge \psi)$.

In our approach, we deal exclusively with Horn formulas, and so *formula* will refer to a Horn formula; the only exception is when we discuss background work, in which case the context is clear. Formulas are denoted by lower case Greek letters; sets of formulas are denoted by upper case Greek letters. Parentheses are required in 2 above to distinguish, for example $(p \wedge q \rightarrow r)$ from $(p) \wedge (q \rightarrow r)$. Nonetheless we freely drop parentheses when the meaning is clear.

An *interpretation* w is a subset of $\mathcal{P}$, where $a \in w$ means that a is assigned *true* in w and $a \notin w$ means that a is assigned *false* in w. Occasionally we will explicitly list negated atoms in an interpretation; for example for $\mathcal{P} = \{p, q\}$ the interpretation $\{p\}$ will sometimes be written $\{p, \neg q\}$ or more briefly $p\bar{q}$. The symbol $\bot$ is always assigned *false*. $\mathcal{M}$ is the set of *interpretations* or *(possible) worlds* (we use these terms interchangeably). Sentences of $\mathcal{L}_H$ are *true* or *false* in an interpretation according to the standard rules in propositional logic. Truth of ϕ in w is denoted $w \models \phi$. As well, for $W \subseteq \mathcal{M}$, $W \models \phi$ iff for every $w \in W$, $w \models \phi$. For formula ϕ, $[\phi]$ is the set of models of ϕ. For set of worlds W, $t_H(W)$ denotes the set of formulas satisfied by all worlds in W, i.e.

$$t_H(W) = \{\phi \in \mathcal{L}_H \mid m \models \phi \text{ for every } m \in W\}.$$

Note that this means that $t_H(\emptyset) = \mathcal{L}_H$.

$\Gamma \vdash \phi$ iff ϕ is derivable from the set of formulas Γ. Again, members of Γ and ϕ are Horn, and $\vdash$ is defined in terms of Horn formulas; see [12] for details. $\psi \vdash \phi$ is an abbreviation for $\{\psi\} \vdash \phi$, and $\psi \equiv \phi$ is logical equivalence, i.e. $\psi \vdash \phi$