JID:IJA AID:8250 /FLA [m3G; v1.242; Prn:29/08/2018; 15:50] P.1 (1-17)

International Journal of Approximate Reasoning ••• (••••) •••–•••

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

Partial meet pseudo-contractions ☆

Yuri David Santos [a], Vinícius Bitencourt Matos [b], Márcio Moretto Ribeiro [c], Renata Wassermann [b]

[a] Faculty of Philosophy & Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen, Groningen, The Netherlands
[b] Department of Computer Science, University of São Paulo (USP), São Paulo, Brazil
[c] School of Arts, Sciences and Humanities (EACH), University of São Paulo (USP), São Paulo, Brazil

ARTICLE INFO

Article history:
Received 14 February 2018
Received in revised form 21 August 2018
Accepted 24 August 2018
Available online xxxx

Keywords:
???

ABSTRACT

In the AGM paradigm for belief revision, epistemic states are represented by logically closed sets of sentences, the so-called belief sets. An alternative approach uses belief bases, arbitrary sets of sentences. Both approaches have their problems when it comes to contraction operations. Belief bases are more expressive, but, at the same time, they present a serious syntax dependence.

Between those two extremes lie a whole gamut of operations called pseudo-contractions, some of which may be interesting alternatives to the classical ones, providing a good balance between syntax dependence and expressivity.

In this paper we explore some very natural and general constructions for pseudo-contractions, showing some of their properties and giving their axiomatic characterizations. We also illustrate possible practical scenarios where they can be employed.

1. Introduction

Rational agents, be they humans or machines, must have some representation of their knowledge or belief system. *Belief revision* wants to understand how these agents should change these representations when they are faced with new information. As Gärdenfors has claimed [11, preface], there is little use in knowing how to represent knowledge if we do not know how to change these representations. The problem is not trivial, as new information may contradict what the agent previously believed, and it is not always clear how it can be accommodated.

In the most widely accepted theory of belief revision, known as AGM due to the initials of the authors of [1], the belief state of an agent is represented by a set of formulas closed under logical consequence. This representation has allowed for very elegant results, linking mathematical constructions to properties defining rational outcomes. These properties, the *AGM rationality postulates*, describe very intuitive properties, such as the *inclusion* postulate for contraction, which states that when the agent is giving up a belief, no new information should be added.

Nevertheless, the logical elegance of the AGM theory brings with it some undesirable effects when applied in practice: logical closure is computationally complex for any reasonably expressive language; in most logics, inconsistency (even if

☆ This paper is part of the Virtual special issue on Defeasible and Ampliative Reasoning, Edited by Ivan Varzinczak, Richard Booth, Giovanni Casini, Szymon Klarman and Gilles Richard.

E-mail addresses: y.david.santos@rug.nl (Y.D. Santos), vinicius.matos@usp.br (V.B. Matos), marciomr@usp.br (M.M. Ribeiro).

https://doi.org/10.1016/j.ijar.2018.08.006
0888-613X/

momentary) will lead to trivialization; there is no distinction between beliefs that are explicit, or explicitly acquired, and beliefs that were implicitly derived.

In further studies (such as [15]), a generalization of the AGM theory was proposed wherein belief states are represented by arbitrary sets of logical sentences, not necessarily closed, called *belief bases*. Belief bases have a major advantage over closed belief sets: they are more expressive, as many different bases can be representations for one and the same belief state. This is particularly useful in the case of an inconsistent belief state. With logically closed sets, typically an inconsistency will clutter the set with all sentences of the language, making any two inconsistent belief states indistinguishable.

The same mathematical constructions defined for the AGM theory can be applied for belief bases. But here, the inclusion postulate, so intuitive for closed sets, turns out to be too restrictive. Without adding new elements it is not possible to weaken any formula, which could be enough to perform the contraction in consonance with the *success* postulate (i.e., effectively removing the input sentence). This means that more than what is necessary will be removed, considering that we want a *minimal change* (regarded as one of the base principles of belief revision). This problem can also be viewed as the problem of *syntactic dependence*, i.e., for two bases representing the same belief state, the codification of the base (the individual formulas it contains) can lead to different contractions. This phenomenon does not exist in the context of belief sets, since they represent a belief state in the *knowledge level* [27], ignoring syntactic variations.

Thus, on the one hand, we have operations on belief states represented by logically closed sets, which are completely independent of the syntactical form of formulas and, on the other hand, we have base operations where the outcome is completely dependent on the syntax. What we are looking for in this work is a middle ground, where some syntactical distinctions may be blurred. Throughout the paper, we present different examples of what this means. Typically, we are interested in weakening formulas, giving up parts of beliefs. As a motivation, consider the following example:

Example 1. [17] Suppose I believe, for good and independent reasons, that *Andy is son of the mayor* (a) and *Bob is son of the mayor* (b). Then I hear the mayor say: "I certainly have nothing against our youth studying abroad. My only son did it for three years". I then have to retract $a \wedge b$ from my base $B = \{a, b\}$. But it is reasonable to retain a belief that either Andy or Bob is the son of the mayor, i.e., the result of the contraction should be $\{a \vee b\}$.

Hansson called these operations *pseudo-contractions* [17].

Pseudo-contractions are base change operations that violate the inclusion postulate. The goal of this paper is to investigate a particular construction for pseudo-contractions. We first argue for the benefits of this investigation, then we demonstrate some of the properties and relationships of this operation with other pseudo-contractions in the literature, such as the ones proposed by Nebel [26] and by Ribeiro and Wassermann [29]. It is important to note that in departing from the inclusion postulate for bases, we are not advocating for the arbitrary inclusion of formulas, but only for those that would be allowed in the corresponding closed set.

The original motivation for this research program came from databases and ontologies, where the knowledge base is a simplified representation for a more complex set of beliefs.

Example 2. Imagine that we have the knowledge base below:

$$KB = \{\text{likesCold} \wedge \neg\text{flies}, \text{isBird}, \text{isBird} \rightarrow \text{flies}\}.$$

Here, $\{\text{flies}, \neg\text{flies}\} \subseteq Cn(KB)$, that is, KB is inconsistent. It is easy to spot the inconsistency in this case, but one can easily come up with other examples where the chain of reasoning leading to flies is much more complicated, which would be more realistic. It is not uncommon for real life ontologies (KBs) to contain some inconsistencies. This, of course, does not mean that the users of that KB want to treat it as any other inconsistent KB, such as $\{\text{flies}, \neg\text{flies}\}$ or $\mathfrak{L}$ (the whole language) itself: they still find the information in KB meaningful, although strictly speaking $\varphi \in Cn(KB)$ for any φ. So, on the one hand, we want the syntactic form of KB to be taken into account, not only its logical closure ($Cn(KB)$). The bits of information likesCold and isBird, for example, seem to be part of the KB, whereas ¬likesCold and ¬isBird do not (although from a purely logical point of view, which is probably not the user's point of view, ¬likesCold and ¬isBird do follow from the explosion caused by the inconsistency). On the other hand, if we consider KB to be completely in the syntactical realm, a contraction of ¬flies in this case could be accomplished by removing likesCold ∧ ¬flies, but then we also lose likesCold ($\text{likesCold} \notin Cn(KB - \neg\text{flies})$). Sure, there might be cases where we want likesCold to go with ¬flies: maybe they were entered together and therefore should be deleted together. It is plausible, however, that in some cases the users of this KB (and this is application/domain dependent) would expect likesCold to survive the removal of ¬flies, making $\{\text{likesCold} \wedge \neg\text{flies}, \text{isBird}, \text{isBird} \rightarrow \text{flies}\} - \neg\text{flies} = \{\text{likesCold}, \neg\text{flies}, \text{isBird}, \text{isBird} \rightarrow \text{flies}\} - \neg\text{flies}$. This example shows that an intermediate degree of syntax dependence (not zero but not complete) might be desired.

The above example supports the idea that KBs must neither be treated as belief sets (deductively closed) nor simply as mere sets of formulas, where only the information explicitly represented matters. We come back to ontology engineering as an example of application in Section 5.

The rest of this paper is organized as follows. The subjects of belief revision and pseudo-contractions are approached in Sections 2 and 3, respectively. In Section 4, we introduce a very general pseudo-contraction, show some of its properties