On Paraconsistent Belief Revision in LP (Extended Abstract)

Nicolas Schwind[1] , **Sébastien Konieczny**[2] , **Ramón Pino Pérez**[2]
[1]National Institute of Advanced Industrial Science and Technology, Tokyo, Japan
[2]CRIL - CNRS, Université d'Artois, Lens, France
nicolas-schwind@aist.go.jp, {konieczny,pinoperez}@cril.fr

Abstract

Belief revision aims at incorporating a new piece of information into the beliefs of an agent. This process can be modeled by the standard Katsuno and Mendelzon's (KM) rationality postulates. However, these postulates suppose a classical setting and do not govern the case where (some parts of) the beliefs of the agent are not consistent. In this work we discuss how to adapt these postulates when the underlying logic is Priest's Logic of Paradox, in order to model a rational change, while being a conservative extension of KM belief revision. This paper is a summary of (Schwind, Konieczny, and Pino Pérez 2022).

1 Introduction

Belief revision accommodates into an agent's beliefs φ a new, reliable, piece of evidence μ, where both φ and μ are represented as propositional formulae. A revision operator $\circ$ is expected to satisfy a set of properties called KM postulates (denoted by (R1-R6)) (Alchourrón, Gärdenfors, and Makinson 1985; Katsuno and Mendelzon 1991). These postulates translate the three core principles of belief change: primacy of update, consistency, and minimality of change. Every revision operator can be constructed by associating each φ with a (plausibility) total preorder $\preceq_\varphi$ over worlds: when revising φ by μ, one takes as a result the models of μ that are minimal in $\preceq_\varphi$.

An issue with the KM framework is that it does not govern the case of inconsistent inputs, whether it be the agent's beliefs φ or the new information μ. So a solution is to use a paraconsistent logic, define the revision to be the conjunction, that allows one to derive sensible conclusions. But this is not a belief *change* operation. As an example, consider that the current beliefs of the agents are $p \wedge \neg p \wedge q$ and that the new piece of information is $\neg q$. In this case, it is arguable to consider that the variables p and q are independent, and that the expected result could be $p \wedge \neg p \wedge \neg q$, i.e., one may want to keep holding the conflict on p, while changing the beliefs on the truth value of q. Simply speaking, we want to do more than just conjunction.

For this purpose, we rephrase the KM postulates in a particular paraconsistent setting: Priest's Logic of Paradox (Priest 1979) (LP for short), a three-valued logic, with the third value meaning "inconsistent" ("both true and false"), that allows to isolate inconsistencies in the concerned propositional variables. In LP, a world (which we call LP world), is a three-valued interpretation. An LP world is an LP model of a formula α if it makes α "true" or "both". The LP entailment relation is defined in terms of inclusion between sets of LP models, i.e., $\alpha \models_{LP} \beta$ iff $[\![\alpha]\!] \subseteq [\![\beta]\!]$, where $[\![\alpha]\!]$ denotes the set of LP models of α. And $\alpha \equiv_{LP} \beta$ denotes the case where $[\![\alpha]\!] = [\![\beta]\!]$.

We discuss how to adapt the KM postulates in order to model a rational change, while being a conservative extension of KM belief revision. This requires in particular to adequately adapt the definition of belief *expansion*, since its direct translation is not adequate for non classical settings. We provide a representation theorem for this class of revision operators in terms of plausibility preorders on LP worlds (faithful assignments). And we define a whole family of distance-based operators that generalize Dalal revision in this setting.

2 Proposal

Representative LP models

First, it is useful to see that all LP worlds can be partially ordered, i.e., $\omega \preceq_{LP} \omega'$ if whenever ω' associates a variable with a classical value "true" of "false", ω associates that variable with the same classical value. Thus $\omega \preceq_{LP} \omega'$ could be read as ω' is "less classical" than ω. In particular, the LP worlds that are minimal with respect to $\preceq_{LP}$ form the set of all classical worlds. A crucial observation is that if an LP world ω is an LP model of a formula α, then all LP worlds ω' such that $\omega \preceq_{LP} \omega'$ are also LP models of α. This means that the only "meaningful" LP models of a formula α are the set of its minimal elements w.r.t. $\preceq_{LP}$, denoted by $[\![\alpha]\!]_\star$. We call the set $[\![\alpha]\!]_\star$ the *representative* set of LP models of α, and we can show that $[\![\alpha]\!] = [\![\beta]\!]$ if and only if $[\![\alpha]\!]_\star = [\![\beta]\!]_\star$.

Belief Expansion in LP

The notion of representative set allows one to formalize the notion of *expansion* in the LP setting, an operation on which belief revision relies. The expansion of φ by μ, denoted by $\varphi + \mu$, simply consists in "adding" μ into φ. In particular, $\varphi + \mu$ is a formula that does not question what could be already derived from φ, and it particular it may be an inconsistent formula (in the classical sense). In the classical case, expansion corresponds to the conjunction $\varphi \wedge \mu$. But while in the classical case, all classical worlds have the same status, this is not the case in LP: as we have seen before, some LP worlds are more "important" than others to characterize the LP models of a formula. So in LP, we rather focus on the *representative* set of LP models of φ to perform the selection: we define the LP-expansion $\varphi +_{LP} \mu$ as a formula whose representative set is characterized by $[\![\varphi +_{LP} \mu]\!]_\star = [\![\varphi]\!]_\star \cap [\![\mu]\!]$. That is, one selects the representative LP models of φ that are LP models of μ. When this set is non-empty, we say that the expansion is conclusive. Interestingly, in the same way conjunction is the only operator satisfying the Gärdenfors expansion postulates in the classical setting (Gärdenfors 1988), we show that $+_{LP}$ as defined above is the only operator that satisfies (an adaptation of) the expansion postulates to the LP setting.

LP Revision

Based on our LP-expansion operator, we propose the following set of postulates for LP revision:

(LP1) $\varphi \circ \mu \models_{LP} \mu$

(LP2) If $\varphi +_{LP} \mu$ is conclusive, then $\varphi \circ \mu \equiv_{LP} \varphi +_{LP} \mu$

(LP4) If $\varphi \equiv_{LP} \varphi'$ and $\mu \equiv_{LP} \mu'$, then $\varphi \circ \mu \equiv_{LP} \varphi' \circ \mu'$

(LP5) $(\varphi \circ \mu) +_{LP} \mu' \models_{LP} \varphi \circ (\mu \wedge \mu')$

(LP6) If $(\varphi \circ \mu) +_{LP} \mu'$ is conclusive, then $\varphi \circ (\mu \wedge \mu') \models_{LP} (\varphi \circ \mu) +_{LP} \mu'$

These postulates are similar to the original KM ones, except that we use the LP entailment in place of the classical entailment, and we use the LP-expansion instead of the classical expansion (i.e., instead of the conjunction). Thus, each postulate (LPi) above is a direct translation of the original KM postulate (Ri). Noteworthy, (LP3) does not appear in the list. Indeed, the KM postulate (R3) says that the revised result $\varphi \circ \mu$ should be consistent whenever the new information μ is consistent. When interpreting the notion of consistency in terms of non-emptiness of set of models, a direct adaptation of this postulate to LP would not make sense anymore: since every formula has a non-empty set of (representative) LP models, it is trivially satisfied. However, we discuss some possible adaptation of (R3), and show that the set of all postulates (including the adaptation of (R3)) can be viewed as a conservative extension of the KM postulates in LP (see (Schwind, Konieczny, and Pino Pérez 2022) for details).

Representation Result

Similarly to the classical case, every LP revision operator can be characterized in terms of an *LP faithful assignment*, i.e., by associating each φ with a total preorder over LP worlds $\preceq_\varphi$. This time, the first level of each total preorder φ corresponds to the *representative* set of LP models of φ. To revise φ by a formula μ, one takes as a result a formula $\varphi \circ \mu$ whose LP models are the LP-closure of the LP models of μ that are minimal in $\preceq_\varphi$, where the LP-closure of a set of LP worlds S, denoted by $Cl(S)$, is defined by $Cl(S) = \{\omega \mid \exists \omega' \in S, \omega' \preceq_{LP} \omega\}$ (accordingly for every set of LP worlds S, there is always a formula α such that $[\![\alpha]\!] = Cl(S)$).

Theorem 1. *An operator $\circ$ is an LP revision operator (i.e., it satisfies (LP1-LP6)) if and only if there is an LP faithful assignment $\varphi \mapsto \preceq_\varphi$ associating every formula with a total preorder over LP worlds such that for all formulae φ, μ, $[\![\varphi \circ \mu]\!] = Cl(min([\![\mu]\!], \preceq_\varphi))$.*

Distance-Based LP Revision

Lastly, we introduce a class of LP revision operators that are based on a *distance* between LP worlds. In the classical case, an interesting example of distance is the Hamming distance between worlds, which defines the well-known Dalal revision operator. We propose to extend that distance to the LP case. In the classical case, Hamming distance consists in counting the number of differences, of "changes", between two classical worlds. But in the three-valued LP setting, we have three values, and the "change" between "true" and "both" may not be the same as the change between "true" and "false". So to be as general as possible, we consider a distance which counts the number of changes between two LP worlds, where this change is given by a distance d_b between two truth values: $d(\omega, \omega') = \sum_{x_i} d_b(\omega(x_i), \omega'(x_i))$. For instance, the choice of a value for $d_b(true, both)$ represents, for the underlying agent, the cost of change from the value "true" to "both" for any variable. We show that under very natural conditions on d_b, the only value that matters in the definition of such a distance-based operator is the cost of change C from a classical truth value ("true" or "false") to the value "both". This means that overall, to define a distance-based operator, one has a one-dimensional choice space which corresponds to whether one wants to model a revision behavior that is *reluctant* to change (a high value for C), or *inclined* to change (a low value for C).

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.

Gärdenfors, P. 1988. *Knowledge in flux*. MIT Press.

Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52:263–294.

Priest, G. 1979. The logic of paradox. *Journal of Philosophical Logic* 8(1):219–241.

Schwind, N.; Konieczny, S.; and Pino Pérez, R. 2022. On paraconsistent belief revision in LP. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI'22)*.