ABOUT THE LIMITATIONS OF LOGIC-BASED APPROACHES TO THE FORMALISATION OF BELIEF FUSION

Eric Grégoire

CRIL CNRS & IRCICA,
rue Jean Souvraz SP18 F-62307 Lens France gregoire@cril.univ-artois.fr

ABSTRACT

In this paper, it is shown that logical approaches to the formalisation of belief fusion suffer from several drawbacks. A first one is due to the fact that the various epistemological roles of involved beliefs are not taken into account. As a result, the set-theoretic union operator is not necessary adequate for fusing beliefs even when such an operator delivers a consistent set of beliefs and when no preference can rank order beliefs. A second problem is linked to an inherent property of deduction and of most inference systems that conducts weaker information to be hidden by stronger ones. Finally, it is claimed that some postulates –especially the syntax-independence ones- founding many logic-based belief fusion and revision approaches are themselves questionable.

1. INTRODUCTION

The controversial debate about the possible roles of logic in the representation of knowledge and formalisation of reasoning is not new in the artificial intelligence community. It can be traced back to M.L. Minsky [14] who emphasised that standard logic, lacking a form of non-monotonicity, is not an adequate tool for knowledge representation and reasoning. This led to the emergence of a (still) very active domain of research about non-monotonic logics, which are intended to represent common-sense knowledge and reasoning. The actual success of this research domain was itself the object of controversial discussions in the late 80s' (see e.g. [7,11]). One should also remember the confrontation of procedural vs. declarative approaches to knowledge representation and reasoning in the 80s', and many discussions about the actual role of logic in that respect: should it be a competence or a specification model vs. an implementation tool (see e.g. [16])? Does it embody unavoidable principles and rules that should be obeyed by any rational –artificial or human- agents (see e.g. [15])? Should it be extended to cover other forms of reasoning outside the mathematical field?

In this paper, we do not actually take position in this still open and controversial debate. Instead, we want to shed some light on some current limitations of logical approaches to knowledge representation and reasoning that did not attract much attention so far. More precisely, we address the specific situation where beliefs[1] of several agents are to be fused. We show that usual logic-based approaches to belief fusion suffer from several drawbacks. The first one is due to the fact that the various epistemological roles of involved beliefs are not taken into account. As a result, the set-theoretic union operator is not necessary adequate for fusing beliefs even when such an operator delivers a consistent set of beliefs and when no preference does rank order beliefs. The second problem is linked to an inherent property of deduction (and of many forms of inference) that conducts weaker information to be hidden by stronger ones. We also claim that some rationality postulates adopted by most logic-based studies about belief revision and belief fusion are highly questionable. In this last aspect, we agree with Friedman's and Halpern's point of view [4]. More specifically, we claim here that syntax-independence is often not an adequate postulate.

Indeed, a very active field of research in symbolic artificial intelligence concerns the logical formalisation of the dynamics of beliefs of rational agents. These last two decades, many authors have concentrated on how a rational agent should revise his (her) beliefs in the presence of inconsistency, and how the beliefs of several agents should be merged (see e.g. [2] or [3] for a survey). Most authors dealing with the logic-based formalisation of belief revision and belief fusion have based their works on the following consideration: when the (multi-)set of beliefs of an agent becomes inconsistent, then his (her) beliefs should be revised in order to restore consistency. Similarly, when the beliefs of several agents should be fused, then it is often proposed to take

[1] In this paper, we do not make any difference between *belief* and *knowledge* since our discussion does not require us to make such a difference. Likewise, we shall use the expression belief base and knowledge base, indiscriminately. However, we claim that actual fusion operators should take this distinction into account.

the set[2]-theoretic union of the beliefs unless this set becomes inconsistent. Otherwise, some operation should be done to recover consistency.

Let us show how such an approach misses several important points.

2. THE SET-THEORETIC UNION OPERATOR IS NOT ALWAYS APPROPRIATE TO THE FORMALISATION OF BELIEFS

To motivate this claim, let us take a very simple example. Assume that we have a first belief base containing the formula "When the switch is on and the lamp-bulb is ok, then the lights should be on", that is to be merged with a second one containing the formula "When the switch is on and the switch is ok, then the lights are on". Assume that both bases are consistent and do not contradict one another. Then, according to most authors, the merged bases should be their set-theoretical union. Indeed, no logical conflict arises and no preference or priority ranks order beliefs. Thus, from the fused beliefs, we can infer that the lights are on when the switch is on, provided that at least *one* of the conditions asserting that the lamps is ok and that the switch is ok is true. We shall thus be able to infer that lights are on even when the lamp bulb is broken, provided that the switch is both ok and on. This is clearly an unwanted conclusion. Actually, the intended meaning of the beliefs in the initial bases expresses *necessary* conditions for the lights to be on. In this specific case, taking their set-theoretic union transforms them as mere *sufficient* conditions for the lights to be on. Clearly, the fusion should have transformed the formula as "When the switch is on and when the lamp-bulb is ok and when the switch is ok then the lights are on", which cannot be obtained by a set-theoretic union.

Actually, most current logic-based approaches lacks an ability to take the various epistemological statuses of the involved beliefs and knowledge into account. In the example, the status of necessary conditions for the lights to be on is not preserved when additional conditions are introduced in the knowledge base.

To some extent, this problem is also linked to the monotonicity property of standard logic. Whenever a conclusion (e.g. lamps-on) can be inferred, no matter additional information comes in, this conclusion can still be deduced. However, it is important to emphasise that such a problem *still* exists in nonmonotonic logics.

For example, assume that we represent the above example using McCarthy's abnormality propositions Ab_i [13], which represent abnormal conditions that are minimised.

$KB_1 = \{\text{Switch-on} \wedge \neg Ab_1 \Rightarrow \text{lights-on},$
$\quad \text{lamp-bulb-ko} \Rightarrow Ab_1\}$

$KB_2 = \{\text{Switch-on} \wedge \neg Ab_2 \Rightarrow \text{lights-on},$
$\quad \text{switch-ko} \Rightarrow Ab_2\}$

Taking the set-theoretic union of the bases and inferring conclusions under the models minimizing the abnormality propositions Ab_i yields the same problematic conclusions that lights are on, when e.g. we also have the additional information switch-on and lamp-bulb-ko.

In [6], a pre-processing procedure is discussed, which prevents such a problem to occur under the above circumscriptive framework, to a given extent.

However, let us stress that a universal logic-based approach to belief fusion that takes care of the various epistemological statuses of the involved knowledge is still far from being available.

3. WEAKER INFORMATION IS HIDDEN IN DEDUCTION AND MOST LOGIC-BASED INFERENCE PROCESSES

When several belief bases are to be fused, another important drawback of logic-based approaches is that weaker information is hidden in deduction and in most logic-based inference processes.

Let us illustrate this by means of a very simple example again. Assume that we are given a first belief base containing a formula representing the information "When the switch is on then the lights are on" and a second one about the same physical device, containing the logically weaker but more precise information "When the switch is on and when the switch is not broken then the lamps are on". Clearly these two pieces of information are mutually consistent. Putting them together inside a same base, will allow us to infer that the lights are on when the switch is on, even when it is also believed that the switch is broken! Once again, because of the monotonicity property of standard logic, whenever we have that a conclusion can be inferred from a set of premises, no additional piece of information will allow us to retract such an inference. In the above example, taking both the second piece of information and the piece of information that the switch is broken as the additional one, will not allow us to retract the conclusion that the lights are on when the switch is on when we also know at the same time that the switch is on and the rule that when the switch is on then the lights are on.

Although this feature is again linked to the monotonicity property of standard logic, switching to current nonmonotonic logics does solve the problem. Indeed, such a phenomenon even appears within a single *consistent* belief base for which the issue of retracting conclusions because inconsistency occurs is not relevant.

For example, assume again that the above example is represented using McCarthy's abnormality propositions [13].

[2] Or multi-set. For clarity of presentation, we shall use the word "set", only.

KB_1 = {switch-on $\Rightarrow$ lights-on, switch-on}
KB_2 = {switch-on $\wedge \neg Ab_1 \Rightarrow$ lights-on,
switch-ko $\Rightarrow Ab_1$}

In such a case, under a minimal semantics minimizing abnormality predicates Ab_i, we shall infer lights-on from $KB_1 \cup KB_2 \cup$ {switch-ko}. Indeed, such a non-monotonic logic (as most of them) allows one to derive a superset of the conclusions that are classically entailed. Switching to such a nonmonotonic logic will thus not solve the problem.

Actually, such a problem might concern any deductive conclusion from a set of premises.

Let us for instance consider $KB_1 = \{A \Rightarrow B, A \wedge \neg Ab_1 \Rightarrow B\}$. Clearly, $A \Rightarrow B$ subsumes $A \wedge \neg Ab_1 \Rightarrow B$ and $\neg Ab_1$ is thus not required to be *true* for B to be inferred from A. Actually, detecting such situations might require us to consider the whole contents of the knowledge bases. As an example, we can see that $KB_1 = \{A \Rightarrow C, C \Rightarrow B, A \wedge \neg Ab_1 \Rightarrow B\}$ leads to an identical problem.

In [Grégoire 03], some a priori checks preventing such a problem to occur are discussed, under this circumscriptive-oriented framework, at least to some extent.

In the general case, whenever we need to merge two belief base KB_1 and KB_2, we should check whether some information that can be inferred from KB_1 is not subsumed by KB_2. However, such a process is often computationally intractable. Indeed, in the case that both knowledge bases are in CNF, and that f and f' are clauses, we have that f strictly subsumes f' when $f \subset f'$. Assume that KB_1 entails f. In order to check whether f is not hidden by more stronger information from KB_2, we need to check whether there exists a subset f' of f such that f' is entailed by KB_2. Clearly, such a problem is NP-hard in the general case.

Actually, devising a reasoning agent that avoids subsuming weaker information requires us to drop a fundamental property of logic-based inference systems.

4. ABOUT RATIONALITY POSTULATES

Many approaches to belief fusion are based on AGM-like [1] rationality postulates that state the properties that should be obeyed by the fusion operator (see e.g. the seminal work by [8]).

One such a postulate is “syntax-independence” which claims that replacing a piece of information by a logically equivalent one should not influence the fusion process. Actually, the debate about whether we should adopt a syntax-dependent or a syntax-free representation of beliefs can be traced backed to [9] and [10]. The latter author stressed how the semantic-approach is too coarse-grained, while a purely syntactical one is too fine-grained for representational purpose. In the following, we claim that a syntax-independence postulate is even more questionable when fusing beliefs is under consideration.

Syntax-independence means that splitting or merging formulas into logically equivalent ones will not affect the fusion process. In many respects, such a requirement is wrong. For example, when a knowledge base is already an accumulation of information from several sources, the duplication of information is sometimes a natural way of enforcing the assertion of the *true* nature of this latter information. Accordingly, dropping (or weakening) this information in a fusion process can be less acceptable than dropping another piece of information that is just asserted once. Similarly, splitting a conjunctive formulas into its conjuncts might exhibit a different *actual* meaning. Indeed, it can be claimed that when a user expresses the formula $a \wedge b \wedge c$, he (she) might want to stress that the three facts a, b and c are crossly-related *true*, which is not necessarily the case when he (she) introduces the three different facts independently in the knowledge base.

Other syntax-related properties can be important in a fusion process [5]. For example, one could want to take the length of the involved formulas into account to define the resulting fused knowledge, especially in case of contradictory information. In some applications, we accept to falsify basic facts whereas we are less likely to falsify longer formulas. For example, in the definition of the first expert systems shells *à la* OPS-5, some wired heuristics in the reasoning control system were implementing such a preference. The motivation was twofold. On the one hand, a long formula can represent a more permanent piece of knowledge representing e.g. an uncontroversial rule whereas a single literal is a tentative conclusion or observation that can be contradicted. Once again, such a difference between various forms of knowledge depends on a epistemological-based classification of beliefs, which is still to be taken into account logic-based belief fusion approaches.

In the same vein, it can be argued that the longer a formula is –the length of formula being a clearly syntactical concept-, the more specific it can be; a more specific piece of information could be preferred in case of conflicting information to be fused because it is more focused and less subject to possible exceptions. In other domains, the opposite view can be adopted. Facts are *true observations* while rules can be hypothetical and be defeated.

5. CONCLUSIONS

In this position paper, we have presented some limitations of current logic-based approaches to knowledge and belief fusion that attracted less attention so far. Actually, we have stressed on two different kinds of problems. The first one is due to the fact that most logic-based belief fusion systems do not differentiate between the various possible epistemological roles of the involved beliefs. The second one is linked to the inherent nature of deduction and logic-based inference systems that subsume logically weaker information, whereas weaker information can exhibit more precise information than does a stronger one.

ACKNOWLEDGEMENTS

This paper has been supported by the IRCICA, and the EC under a FEDER program.

REFERENCES

[1] Alchourrón C.E., Gärdenfors P. and Makinson D., “On the logic of theory change: partial meet functions for contraction and revision”, Journ. Of Symbolic Logic, vol. 50, pp. 510-530, 1985.

[2] Bloch I. and Hunter A. (guest eds.), “Fusion: general concepts and characteristics”, Int. Journ. of Intelligent Systems, vol. 16, pp. 1107-1134, 2000.

[3] Cholvy L., “Reasoning about merging information”, Handbook of Defeasible Reasoning and Uncertainty Systems”, vol. 3, pp. 133-263, 1998.

[4] Friedman N. and Halpern J.Y., “Belief revision: a Critique”, Proc. of KR’96, pp. 421-431, 1996.

[5] Grégoire E., « Syntax and semantics in knowledge fusion: a mixed approach », *Proc. of the 6th Int. Conf. on Sensor Fusion: Architectures, Algorithms and Applications*, B.V. Dasarathy (ed.), Orlando, pp. 60-64, april 2002.

[6] Grégoire E., « Fusing Cooperative Technical-Specification Knowledge Components », *Int. Journ. on Artificial Intelligence Tools (IJAIT)*, vol. 12, n°3, pp. 265-278, 2003.

[7] Hanks S. and McDermott D., “Default reasoning, nonmonotonic logic and the frame problem”, Proc. of AAAI-86, pp. 328-333, 1986.

[8] Konieczny S. and Pino Pérez R., “Merging with integrity constraints”, Proc. of ECSQARU’99, LNCS 1638, pp. 233-244, 1999.

[9] Konolige K., A deduction model of belief, Research Notes in Artificial Intelligence, Pitman, London, 1986.

[10] Levesque, H.J., “A logic of implicit and explicit belief”, Proc. of AAAI-84, pp. 198-202, 1984.

[11] Levesque, H.J. (ed.), “Taking issue/forum: A critique of pure reason”, Computational Intelligence, 3(3), 1987.

[12] Lehmann D., “Belief revision, revisited”, Proc. of IJCAI’95, pp. 1534-1540, 1995.

[13] McCarthy J., “Applications of circumscription to formalizing common-sense knowledge”, Artificial Intelligence, vol. 28, pp. 89-116, 1986.

[14] Minsky M.L., “A framework for representing knowledge” Winston P. (ed.), The Psychology of Computer Vision, McGraw-Hill, New York, pp. 34-57, 1974.

[15] Moore R.C., “The role of logic in knowledge representation and common-sense reasoning”, Proc. of AAAI-82, pp. 428-433.

[16] Newell A., “The knowledge level”, AI Magazine, 2, pp. 1-20, 1980.

[17] Thayse A. (ed.), “From Standard Logic to Logic Programming”, Wiley, London, 1988.