Being consistent about inconsistency:

Toward the rational fusing of inconsistent propositional logic bases

Didier Dubois and Henri Prade

Abstract. This short note discusses preliminary steps in the way to a general picture of inconsistency handling in an information fusion perspective, by relating some natural approaches.

1. Introduction

The fusion of pieces of information coming from different sources and the handling of inconsistency have motivated works in logic [2, 11] as well as in artificial intelligence (see, e.g. [7]). The issues are then the appropriate representation of the information provided by each source, and the handling of conflicts between sources, and possibly inside sources. However, it is usually assumed that the information provided by each source is consistent. In the following short discussion, we allow for the inconsistency of information provided by individual sources in a propositional format, and show it leads to interesting pending questions.

2. Background on the handling of possibly inconsistent propositional knowledge bases

By a knowledge base K here, we mean a finite set of propositional formulas $K = \{\varphi_i | i = 1, \cdots, m\}$. Then, given two knowledge bases K and $K' = \{\psi_j | j = 1, \cdots, n\}$, one can define their conjunctive and disjunctive combinations respectively in a syntactic manner as

$$K \otimes K' = \{\varphi_i \wedge \psi_j \mid i = 1, \cdots, m, \text{ and } j = 1, \cdots, n\}$$
$$K \oplus K' = \{\varphi_i \vee \psi_j \mid i = 1, \cdots, m, \text{ and } j = 1, \cdots, n\}$$

Note that $K \otimes K$ and $K \oplus K$ have the same semantic content as K. Moreover, $K \otimes K'$ has the same set of models (possibly empty) as the mere union $K \cup K'$ of K and K'. When K and K' are consistent, but not $K \otimes K'$, classical inference becomes trivial, and the disjunctive combination $K \oplus K'$ is of interest. It tacitly assumes that one source is right but we do not know which source [8]. This idea may be refined by considering the maximal consistent bases of $K \cup K'$ and then combining them disjunctively. Or yet it has been proposed for a long time to compute the intersection of the sets of consequences of each maximal consistent bases of $K \cup K'$ (it was first proposed by Rescher and Manor already in 1970 [12]). In that case we are back to the problem of non-trivial inference from a single inconsistent knowledge base.

So consider the case when a single base K is inconsistent. Let $S_1, \cdots, S_k$ be the maximal consistent subbases of K, and let $C(S)$ denote the deductive closure of a consistent propositional base S. It is easy to see that $C(S_1 \oplus \cdots \oplus S_k) \subseteq C(S_1) \cap \cdots \cap C(S_k)$. Note that in order to obtain an equality rather this inclusion, one should replace $\oplus$ by a richer disjunction, where we not only perform disjunctions of the form $\psi^1 \vee \cdots \vee \psi^k$ with $\psi^i \in S_i$ for $i = 1, \cdots, k$, but *all* the disjunctions of this kind where now ψ^i is any conjunction $\psi_1^i \wedge \cdots \wedge \psi_j^i$ of propositions taken in S_i (where j ranges between 1 and $card(S_i)$).

Several other approaches exist to obviate the explosive nature of classical inference.

In [4, 5], it has been proposed to associate an inconsistent propositional logic base $K = \{\varphi_i | i = 1, \cdots, m\}$ with its so-called "paraconsistent completion" K^o, defined as $K^o = \{(\varphi_i, 1, x_i) \mid i = 1, \cdots, m\}$ where $x_i = 1$ if $\exists A \subseteq K, A$ consistent, $A \vdash \neg\varphi_i$, and where $x_i = 0$ otherwise.

This method partitions K^o into the inconsistency-free part K_{fr}^o and the paraconsistent part K_{par}^o of K^o, respectively defined by

$$K_{fr}^o = \{\varphi_i \mid (\varphi_i, 1, 0) \in K^o, i = 1, \cdots, m\}$$

$$K_{par}^o = \{\varphi_i \mid (\varphi_i, 1, 1) \in K^o, i = 1, \cdots, m\}$$

However, reducing K to K_{fr}^o, which is consistent, is quite conservative, since while every consequence from K_{fr}^o is also a consequence of all maximal consistent subbases of K, the converse is false, as shown by the following counter-example [4]:

Example
Let $K = \{\alpha, \neg\alpha \vee \neg\beta, \beta, \neg\alpha \vee \gamma, \neg\beta \vee \gamma\}$. Then it can be verified that

- K has three maximal consistent subbases $\{\neg\alpha \vee \neg\beta, \beta, \neg\alpha \vee \gamma, \neg\beta \vee \gamma\}$, $\{\alpha, \beta, \neg\alpha \vee \gamma, \neg\beta \vee \gamma\}$, $\{\alpha, \neg\alpha \vee \neg\beta, \neg\alpha \vee \gamma, \neg\beta \vee \gamma\}$;
- $K_{fr}^o = \{\neg\alpha \vee \gamma, \neg\beta \vee \gamma\}$ is the intersection of these three bases (this is always true);
- K_{fr}^o does not entail γ;
- γ is entailed by each maximal consistent subbase of K.

Besides, the "paraconsistent completion" can be extended to layered knowledge bases in the setting of possibilistic logic [5, 9], taking into account the levels of certainty of the pieces of information.

A bolder inference mechanism that provides a set of consequences larger than $C(S_1) \cap \cdots \cap C(S_k)$ is provided by the so-called *argued inference*, which is defined as follows: ψ is an argued consequence of a (possibly inconsistent) base K if there is a consistent subset of K that entails ψ and none that entails $\neg\psi$. A set of two argued consequences is always consistent, but it is no longer the case for larger sets of argued consequences [4]. This contrasts with the fact that $C(S_1) \cap \cdots \cap C(S_k)$ is clearly consistent.

These approaches can be encompassed by extending the paraconsistent completion $\mathcal{L}^o$ to the whole language $\mathcal{L}$ of K, in the spirit of the proposal by Arieli [1]: $\forall\phi \in \mathcal{L}$:

- $\phi \in \mathcal{L}_T$ if and only if there is a consistent subset of K that entails ϕ and none that entails $\neg\phi$; and we can write $(\phi, 1, 0) \in \mathcal{L}^o$.
- $\phi \in \mathcal{L}_F$ if and only if there is a consistent subset of K that entails $\neg\phi$ and none that entails ϕ; and we can write $(\phi, 0, 1) \in \mathcal{L}^o$.
- $\phi \in \mathcal{L}_U$ if and only if there is no consistent subset of K that entails ϕ nor any that entails $\neg\phi$; and we can write $(\phi, 0, 0) \in \mathcal{L}^o$.
- $\phi \in \mathcal{L}_I$ if and only if there is a consistent subset of K that entails ϕ and another one that entails $\neg\phi$; and we can write $(\phi, 1, 1) \in \mathcal{L}^o$.

In the above definition one can restrict to maximal consistent subbases of K. These four sets of formulas $\mathcal{L}_T, \mathcal{L}_F, \mathcal{L}_U, \mathcal{L}_I$ partition the language. It can be checked that $K^o \subset \mathcal{L}^o$ (in particular, $K^o_{fr} \subset \mathcal{L}^o_T$ and $K^o_{par} \subset \mathcal{L}^o_I$), and that $\mathcal{L}_T$ is the set of argued consequences. One can view the four annotations by pairs of Boolean values as akin to Belnap [2] epistemic truth-values, TRUE, FALSE, NONE and BOTH respectively. However, Belnap logic comes down to computing epistemic statuses of atomic propositions based in information from various sources, then obtaining the epistemic status of other formulas via truth-tables extending the usual ones to four values.

The paraconsistent completion and its extension also stem from possibilistic logic [9], since it comes down to computing degrees of necessity $N_i(\phi)$, and $N_i(\neg\phi)$ whereby $N_i(\phi) = 1$ if and only if $S_i \vdash \phi$; otherwise $N_i(\phi) = 0$. Necessity measures N_i are $\wedge$-decomposable functions ($N_i(\phi \wedge \psi) = \min(N_i(\phi), N_i(\psi))$) akin to KD modalities. The extended paraconsistent completion can be formally defined as

$$\mathcal{L}^o = \{(\phi, \max_{i=1}^k N_i(\phi), \max_{i=1}^k N_i(\neg\phi)) : \phi \in \mathcal{L}\}.$$

Note that $\max_{i=1}^k N_i$ is just a general monotonic Boolean set-function (and conversely any such set function is of this form for some integer k), and the annotated formulas can be encoded in a non-regular modal logic [10].

It is easy to see that for the cautious approach using intersection, it holds that $C(S_1) \cap \cdots \cap C(S_k) = \{\phi : \min_{i=1}^k N_i(\phi) = 1\}$. The companion completion of $\mathcal{L}^o$ of the form $\mathcal{L}^{cautious} = \{(\phi, \min_{i=1}^k N_i(\phi), \min_{i=1}^k N_i(\neg\phi)) : \phi \in \mathcal{L}\}$ only

contains triples of the form $(\phi, 1, 0), (\phi, 0, 1), (\phi, 0, 0)$, with $(\phi, 1, 0) \in \mathcal{L}^{cautious}$ if and only if $(\neg\phi, 0, 1) \in \mathcal{L}^{cautious}$. The set $\mathcal{L}_T^{cautious} = \{\phi : (\phi, 1, 0) \in \mathcal{L}^{cautious}\}$ is consistent, deductively closed and equal to $C(S_1) \cap \cdots \cap C(S_k)$. Indeed, $\min_{i=1}^{k} N_i$ is again a necessity measure.

3. Toward fusing inconsistent knowledge bases

In this section, we turn to the case when the source bases are inconsistent already.

3.1. A misleadingly simple problem

Let us consider the simple case of the fusion of the two atomic knowledge bases $K = \{p, \neg q, q\}$ and $K' = \{p, \neg p, q\}$. There are at least four ways of envisaging their fusion:

- A first idea is to merely concatenate compute $K \cup K' = \{p, \neg p, q, \neg q\}$; we get a state of contradiction both about p and about q, which is not appealing, as one may feel that some information has been lost in the process (the same result would obtain by merging $\{p, \neg p\}$ and $\{q, \neg q\}$).
- However, we can observe that all the maximal consistent sub-bases of K (resp. K') entail p (resp. q). So if we apply an inconsistency-tolerant inference to each of K and K' prior to merging, we may feel to be entitled to derive $p \wedge q$ in the end.
- Alternatively, consider $K \oplus K' = \{p, p \vee \neg p, p \vee q, p \vee \neg q, \neg p \vee \neg q, \neg q \vee q, \neg p \vee q, q\}$. It may be reduced to $(K \oplus K')^d = \{p, \neg p \vee \neg q, q\}$ if we delete each formula which is subsumed by another one. Although $(K \oplus K')^d$ is inconsistent, it is syntactically different from $K \cup K'$, and now all consistent subbases of $(K \oplus K')^d$ only entail $p \vee q$.
- Finally, note that p (resp. q) is inconsistency-free in K (resp. K'). This may suggest to view K as a tentative representation of a situation where the associated source asserts p, and states it has contradictory information about q, in the sense that $(q, 1, 1) \in K_{par}^o$, which we shall denote by $\perp_q$. Similarly, K' corresponds to a situation where the associated source asserts q, and states it has contradictory information about p, namely $\perp_p$. It may be considered reasonable to assume that $p \vee \perp_p = p$ (since $p \vee (\neg p \wedge p) = p$, and it agrees with the truth-table of disjunction for Belnap logic), we get

$$\{p, \perp_q\} \oplus \{q, \perp_p\} = \{p, q, p \vee q, \perp_p \vee \perp_q\},$$

which may be reduced to $\{p, q, \perp_p \vee \perp_q\}$. It expresses that p and q can be asserted, and that we have contradictory information about p or about q.

The latter approach, that uses a reification of local inconsistencies, does not sound counterintuitive; it is just an optimistic way of fusing the information coming from the two sources stating respectively that p is true and information about q is conflicting (paraconsistent) and that q is true and information about p is conflicting. It is less optimistic than the second approach which throws away much information to restore consistency of source information prior to merging.
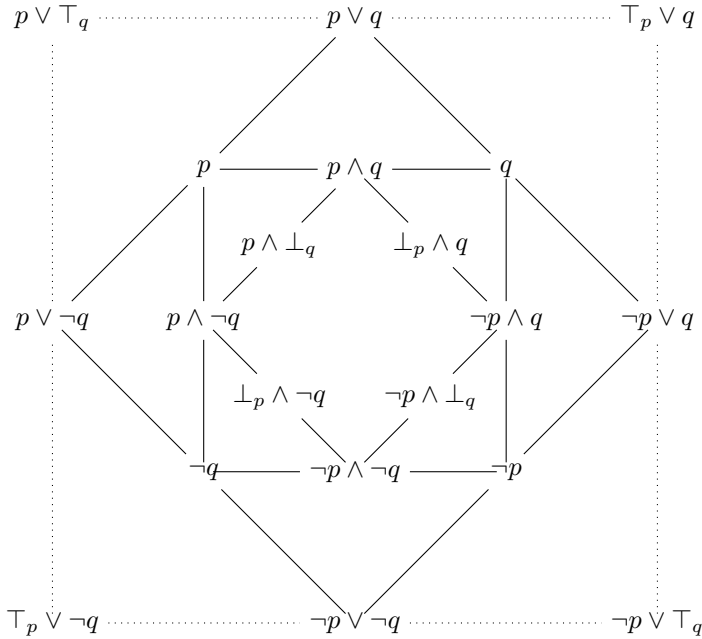
Figure 1. Conjunctions and disjunctions in squares

The basic elements of the calculus implicitly underlying this latter view may be nicely pictured on Figure 1, where three squares in solid lines (which are not squares of opposition, at least in the classical sense) are displayed. The middle of the side of a square is always the conjunct of its edges, and the vertices of the isosceles triangles are the disjunction of the edges of their opposite side.

3.2. Elements for a discussion

Two basic issues, in order to make sense of the problem from an artificial intelligence point of view, are certainly i) to have a proper representation of what has to be fused (which requires a clear understanding of what is supposed to be stated); and ii) what are the principle underlying the merging process or the conjoint exploitation of the knowledge bases reflecting the sources of information.

As it may be expected, one observes that if information is modelled by an inconsistent set of atomic formulas, one may be in one of the four states of information p, $\neg p$, $\perp_p$, and $\top_p$ regarding a propositional variable p. Here the symbol $\top_p$ would encode a local tautology $p \vee \neg p$ in a symmetric way as does $\perp_p$ for the contradiction. One may then be tempted to complete the three squares by vertices of the form $\top_p \vee q$, etc. (see the external square in dotted lines). But it is hard to defend the idea that $(\top_p \vee q) \wedge (\top_q \vee p) = p \vee q$, as, strictly speaking, it is just a tautology. However, just as the inconsistency tolerant approaches try to confine

the contradiction to a local feature not destroying knowledge bases, one could similarly, as a topic of further thought, try to have a local approach to tautologies that would avoid the opposite form of trivialization (nothing follows from a tautology).

The possibility of defining four statuses for formulas with respect to an inconsistent knowledge bases may suggest to try either a multiple-valued logic approach (as in [2]), or a modal logic approach (as in [3]), or an argumentative approach (as in [1]).

The multiple-valued logic view may be considered misleading as one may argue that truth-tables are a clumsy way of handling states of information attached to atomic formulas [6]. In fact, the symbol p in the above attempt is ambiguous when put along with $\perp_p$, as writing p in the knowledge base here refers to *the fact of consistently asserting* p, without its being contradicted by other assertions: in the annotation framework outlined in the previous section, it really stands for $(p, 1, 0)$, while $\perp_p$ stands for $(p, 1, 1)$.

If we understand the four states of formulas induced by an inconsistent knowledge base K as encoded by the paraconsistent completion of the language, one may see that it may lead us to questioning the edges of the second inner square (even though they look safe), namely we can doubt whether, in an inconsistent framework, the conjunction of p and q is $p \wedge q$ at all. Indeed if instead of p, $\neg p$, $\perp_p$, and $\top_p$, we write $(p, 1, 0)$, $(p, 0, 1)$, $(p, 1, 1)$, and $(p, 0, 0) \in \mathcal{L}^o$, then it is easy to check that $(p, 1, 0) \in \mathcal{L}^o$, $(q, 1, 0) \in \mathcal{L}^o$ do not imply $(p \wedge q, 1, 0) \in \mathcal{L}^o$, if we work with an inconsistent knowledge base [6]. This is as well at odds with Belnap logic, which assumes the four-valued logic truth-tables extend the ones of Boolean logic. So the proposed set of squares should not be taken for granted under all approaches to inconsistency.

4. Conclusion

The above discussion only (tentative) aim is to draw attention on basic difficulties in the problem of handling, and more particularly fusing and inferring from propositional bases in the presence of inconsistency, whereby the syntactic framework of propositional logic seems to be ambiguous and insufficient to express what is going on.

Clearly, even if many works have been published on inconsistency and what to do with it, a number of questions remain, and one of the ambitions of this note is to make preliminary steps in the way to a general picture, relating some natural attempts. The only other purpose is to seize the occasion of this Festschrift for Jean-Yves Béziau to heartily wish him a happy birthday.

References

[1] O. Arieli. Conflict-tolerant semantics for argumentation frameworks. Proc. 13th Euro. Conf. Logics in Artificial Intelligence (JELIA'12), (L. Fariñas del Cerro, A. Herzig, J. Mengin, eds.), Toulouse, Sept. 26-28, Springer, LNCS, Vol. 7519, 28-40, 2012.

[2] N. D. Belnap. A useful four-valued logic. In: (J. M. Dunn and G. Epstein, editors), Modern Uses of Multiple-Valued Logic, pages 7-37. D. Reidel, Dordrecht, 1977.

[3] J.-Y. Béziau. A new four-valued approach to modal logic. Logique & Analyse, 213, 109-121, 2011.

[4] S. Benferhat, D. Dubois, and H. Prade. Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study – Part 1: The flat case. Studia Logica, 58, 17-45, 1997.

[5] S. Benferhat, D. Dubois, and H. Prade. Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study – Part 2: The prioritized case. In: (E. Orlowska, editor), Logic at work, Physica-Verlag, Heidelberg, 473-511, 1998.

[6] D. Dubois. On ignorance and contradiction considered as truth-values. Logic Journal of the IGPL, Oxford University Press, 16 (2), 195-216, 2008.

[7] D. Dubois, P. Everaere, S. Konieczny, and O. Papini. Approches de la révision et de la fusion d'informations. In: (P. Marquis, O. Papini, H. Prade, editors), Panorama de l'Intelligence Artificielle. Volume 1: Représentation des Connaissances et Formalisation des Raisonnements, pages 321-361, Cépaduès Éditions, Toulouse, 2014.

[8] D. Dubois, W. Liu, J. Ma, and H. Prade. Toward a general framework for information fusion. Proc. 10th Int. Conf. on Modeling Decisions for Artificial Intelligence (MDAI'13), Barcelona, (V. Torra, Y. Narukawa, G. Navarro-Arribas, D. Megías, eds.), Nov. 20-22, Springer, LNCS 8234, 37-48, 2013.

[9] D. Dubois, H. Prade, Possibilistic logic – An overview. In: Handbook of the History of Logic, Vol. 9: Computational Logic, (D. M. Gabbay, J. Siekmann, J. Woods, eds.), Elsevier, 197-255, 2014.

[10] D. Dubois, H. Prade, A. Rico. Qualitative capacities as imprecise possibilities. Proc. 12th Europ. Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'13), Utrecht, July 8-10, (L. C. van der Gaag, ed.) Springer, LNCS 7958, 169-180, 2013.

[11] N. Rescher, and R. Brandom. The Logic of Inconsistency. Basil Blackwell, Oxford, 1980.

[12] N. Rescher and R. Manor. On inference from inconsistent premises. Theory and Decision, 1:179219, 1970.

Didier Dubois and Henri Prade

Institut de Recherche en Informatique de Toulouse (IRIT), Université Paul Sabatier, 31062 Toulouse Cedex 09, France

e-mail: `dubois@irit.fr, prade@irit.fr`