Ontology Development with 4-Valued Implication Connectives

Cristian Cocos and Wendy MacCaull

Centre for Logic and Information, St. Francis Xavier University
Antigonish, Nova Scotia, Canada
{ccocos, wmaccaul}@stfx.ca

Abstract

We examine bilattice-based logic and extract a self-standing compendium of theoretical results that are of interest to ontology developers in a paraconsistency-tolerant context. This logical framework boasts at least two major paraconsistent consequence relations, one of which is non-monotonic. These consequence relations are at the core of RaDON, a reasoner for paraconsistent ontologies. Applying RaDON requires the intermediate step of translating propositional connectives into 4-valued logic correspondents. The existing translation guidelines are insufficient, and we strive to improve them.

1. Introduction and Motivation

Ontologists looking for a way to diversify their ontology development toolkit will surely be tempted by the promises of a reasoning device that can accommodate contradictory knowledge with grace. Contradictory information is unavoidable in large knowledge bases built through knowledge base merger; as such, finding a way to avoid collapse in the face of inconsistency emerges as a worthy research theme. Fortunately, such efforts have been underway for some time—to our knowledge, at least since Belnap's writings in the '70s [Be92].

The goal of isolating contradictory portions of a theory or a knowledge base has spawned the burgeoning field of Paraconsistent Logics. Within this cluster of research preoccupations, the one area that looks to us to be the most promising is the four-valued logics route (henceforth dubbed the "Belnap-Avron style" ([Be92], [AA96])): certain carefully-defined logical frameworks whose propositions can take more values than the classical {true, false} have the interesting property that a contradiction $(p \& \neg p)$ does not necessarily imply every other proposition: contradictions, it is often said, are not "explosive," or do not entail "explosion."

We undertake, in the following, a presentation of some of the four-valued logic theoretical background, relying heavily on Avron's pioneering work. The main element of novelty is that we have adopted the stance of a working ontologist, whose interest is less in exploring ramifications and subtleties of primarily theoretical interest, but more on providing a reasonably complete foundation from which to derive concrete procedural principles, useful for the day-to-day activity of ontology and knowledge base development. We hence tend to view our contribution as being situated at the junction between pure theory and practice, even though at times we have attempted to fill some gaps in the theory itself. Sections two to five below are thus the result of distilling, filtering, and synthesizing an impressive amount of theoretical facts.

Sections six to eight deal with the application of these principles to ontology merging and classification. We present a particular paraconsistent reasoner, possibly the most complete one, and reflect on what we perceive as one of the more difficult stages of the approach: the translational phase, whereby natural language constructs (propositional connectives, to be more exact) are to be regimented into four-valued logical canons. We note that this translation is to be undertaken by human operators, and that the existing four-valued logic rendering guidelines leave a certain latitude that could generate discomfort for those ontologists devoted to the ideal of a one and only accurate translation. We hence propose a few more such rules, which, in our opinion and experience, should help towards restricting the translational leeway.

2. Bilattice-Based Logics: Syntactic Framework

The core of the Belnap-Avron style of paraconsistent logics consists in the notion of a logical bilattice. We will assume, in the following, that the notion of a lattice requires no further explanation. Just as the name suggests, a *bi*lattice is essentially a set endowed with two lattice structures:

DEFINITION 2.1 ([AA96]): A *bilattice* is a structure $\mathcal{B} = (B, \leq_t, \leq_k, \neg)$, such that B is a set of cardinal at least two, $(B, \leq_t)$ and $(B, \leq_k)$ are lattices, and $\neg$ is a unary operation on B that obeys the following laws:

1. If $a \leq_k b$, than $\neg a \geq_k \neg b$,

2. If $a \leq_t b$, than $\neg a \geq_t \neg b$,
3. $\neg\neg a = a$.

The two partial ordering relations ($\leq_k$ and $\leq_t$) reflect intuitively the differences in the amount of knowledge and, respectively, the measure of truth exhibited by the elements of B. Let f ("false") and t ("true") be the least and, respectively, the greatest elements of $(B, \leq_t)$, and $\perp$ and $\top$ of $(B, \leq_k)$. In the latter order, $\perp$ and $\top$ stand, intuitively, for lack of information (no knowledge) and over-determination. The meet and join of the $(B, \leq_t)$ lattice will be denoted by & and $\vee$ respectively, whereas for $(B, \leq_k)$ by $\otimes$ and $\oplus$.

In the standard case of Belnap's *FOUR* bilattice, B has only four elements: the classical truth values of t (truth) and f (false), plus $\perp$ and $\top$.

The concept of a *designated* truth value has been introduced in order to capture the notions of 'valid formula' and 'consequence relation' in many-valued logics in a manner analogous to classical logic. In Belnap's terms, a designated value represents a truth value known to be "at least true," or "told true" [Be92]. A generalization thereof in bilattice context is the notion of *bifilter*:

DEFINITION 2.2 ([AA96]): A *bifilter* of/on a bilattice $\mathcal{B} = (B, \leq_t, \leq_k)$ is a nonempty proper subset of B ($\mathcal{F} \subset B$, $\mathcal{F} \neq B$) such that:

$a \,\&\, b \in \mathcal{F}$ iff $a \in \mathcal{F}$ and $b \in \mathcal{F}$;

$a \otimes b \in \mathcal{F}$ iff $a \in \mathcal{F}$ and $b \in \mathcal{F}$.

A bifilter is called *prime* if it also satisfies:

$a \vee b \in \mathcal{F}$ iff $a \in \mathcal{F}$ or $b \in \mathcal{F}$;

$a \oplus b \in \mathcal{F}$ iff $a \in \mathcal{F}$ or $b \in \mathcal{F}$.

With this, a *logical bilattice* is a pair ($\mathcal{B}$,$\mathcal{F}$) in which $\mathcal{B}$ is a bilattice and $\mathcal{F}$ is a prime bifilter on $\mathcal{B}$.

A standard example of a logical bilattice is Belnap's *FOUR*, whose only bifilter is $\mathcal{F} = \{\top, t\}$.

Given a logical bilattice ($\mathcal{B}$,$\mathcal{F}$), one defines the syntax of a language based on bilattices as follows:

DEFINITION 2.4 ([AA96]): The language *BL* (bilattice-based language) is the standard propositional language over $\{\&, \vee, \neg, \otimes, \oplus\}$. *BL*(4) is *BL* enriched with the propositional constants $\{f, t, \perp, \top\}$. *BL*($\mathcal{B}$) is *BL* enriched with a propositional constant for every element of *B*.

Finally, the bilattice-based language most interesting to us contains one further "logical constant" $\supset$ defined as:

$$a \supset b \triangleq \begin{cases} b \text{ if } a \in \mathcal{F} \\ t \text{ if } a \notin \mathcal{F} \end{cases}.$$

With this, $BL_{\supset}$, $BL_{\supset}(4)$ and $BL_{\supset}(\mathcal{B})$ are the corresponding language extensions. (Note that throughout this definition we have used the same symbols for propositional connectives and propositional constants as for the bilattice operations and elements of the lattice; this procedure will be employed for the remainder of the article where there is no danger of confusing the semantic layer with the syntactic level.)

3. Consequence Relations: General Theory

So far we have a language, the bilattice-based language. In order to obtain a logic, we need a way to capture the idea of a sentence being related to others by way of consequence. This resonates with the seldom stated "consensus" according to which "Logic … [is] the science of consequence relations" [Av91]. A characterization of what counts as a consequence relation, however, can follow different paths according to the extent to which one wishes to accommodate the various efforts encountered under the "Logic" moniker. Here are the most usual constraints encountered in the literature ([Av91], [Av08], [An08] etc.):

DEFINITION 2.5: Let $\mathcal{L}$ be a propositional language. Let Greek letters such as φ, θ, ψ denote $\mathcal{L}$'s formulas, and T, T', S, S' stand for finite sets of formulas from $\mathcal{L}$. Let $\triangleright$ be the consequence relation of standard bivalent propositional logic. A binary relation $\vDash$ between sets of formulas of $\mathcal{L}$ and formulas of $\mathcal{L}$ is said to obey:

i. *reflexivity* (*identity*): if $\varphi \in T$ then $T \vDash \varphi$;
ii. *monotonicity* (*dilution*, *weakening*): if $T \vDash \varphi$ and $T \subseteq T'$ then $T' \vDash \varphi$;
iii. *supraclassicality*: if $T \triangleright \varphi$ then $T \vDash \varphi$;*
iv. *transitivity* (*cut*): if $T \vDash \psi$ and $T, \psi \vDash \varphi$ then $T \vDash \varphi$.

DEFINITION 2.6: With the same notations as above, a binary relation $\vDash$ between sets of formulas of $\mathcal{L}$ and *sets of* formulas of $\mathcal{L}$ is said to obey:†

i. *reflexivity*: if $T \cap S \neq \varnothing$ then $T \vDash S$;
ii. *monotonicity* (*dilution*, *weakening*): if $T \vDash S$, $T \subseteq T'$ and $S \subseteq S'$ then $T' \vDash S'$;
iii. *transitivity* (*cut*): if $T \vDash \psi, S$ and $T, \psi \vDash S$ then $T \vDash S$.

DEFINITION 2.7: Same as definition 2.6, only that T, T', S, S' stand for finite *multi*sets of formulas from $\mathcal{L}$.

DEFINITION 2.8 [Av91]: A *consequence relation* (CR) is any relation satisfying constraints i and iii of definition 2.7 above. A *Scott consequence relation* (*scr*) is any relation complying with definition 2.6. Finally, a *Tarskian*

* Here obviously $\mathcal{L}$ must contain the language of classic propositional logic.
† In the following we will assume that the reader is familiar with the Gentzen-style notation, esp. regarding the meaning of the presence of a *set* of formulas to the *right* of the turnstile symbol.

consequence relation (*tcr*) is a relation that obeys conditions i, ii, and iv of definition 2.5.

There are several variations on Avron's terminology: there are, for example, authors that require a classical consequence relation (tcr) to also observe condition iii of definition 2.5, though for the purposes of our study we will follow the canons of definition 2.8. Probably the most interesting fact to note is the absence of (the requirement of) monotonicity in the definition of a CR. This is, as a matter of fact, something that will prove useful in the context of bilattice-based logics. This is also something destined to accommodate non-monotonic logics.

4. Bilattice-Based Logics: Paraconsistent Consequence Relations

Given a logical bilattice ($\mathcal{B}$,$\mathcal{F}$), and the corresponding bilattice-based language $BL_{\supset}(\mathcal{B})$, the natural semantics of the language emerges in a manner analogous to classical propositional logic. A valuation will be a function from sentences of $BL_{\supset}(\mathcal{B})$ to B defined in the usual manner. The following relations are of crucial importance for us:

DEFINITION 2.9: Let ($\mathcal{B}$,$\mathcal{F}$) be a logical bilattice, and Γ, Δ finite sets of formulas from $BL_{\supset}(\mathcal{B})$. (a) We write $\Gamma \vDash_{(\mathcal{B},\mathcal{F})} \Delta$ iff each valuation that takes all elements of Γ into $\mathcal{F}$, takes at least one element of Δ into $\mathcal{F}$. (b) If Γ and Δ contain no propositional constants (i.e. $\Gamma \cup \Delta \subseteq BL_{\supset}$), we define the relation $\Gamma \vDash \Delta$ which holds iff $\Gamma \vDash_{(\mathcal{B},\mathcal{F})} \Delta$ for every logical bilattice ($\mathcal{B}$,$\mathcal{F}$). (c) Finally, if Γ and Δ contain only sentences from $BL_{\supset}(4)$ (that is, contain as propositional constants only elements from $\{t, f, \top, \bot\}$), we write $\Gamma \vDash_4 \Delta$ iff $\Gamma \vDash_{(\mathcal{B},\mathcal{F})} \Delta$ for every logical bilattice ($\mathcal{B}$,$\mathcal{F}$).

The next step would be to prove that these relations are, in fact, consequence relations in either of the senses mentioned above (CR, tcr or scr). Interestingly enough, we have been unable to locate in Avron's writings a statement devoted to proclaiming (much less *proving*) this unambiguously, though we have no reason to suspect that this is anything but a very pardonable oversight. For the sake of a clean record, we will fill this gap in the following. In order to do that, we will have to employ the resources of associated deduction mechanisms. Our focus will be on the latter two relations, as, at least from our point of view, $\vDash_{(\mathcal{B},\mathcal{F})}$ is of strictly theoretical interest. And while both $\vDash$ and $\vDash_4$ admit valid formulas—and thus Hilbert-style proof systems—the deduction machineries we will use are of the Gentzen-type sequent calculi variety, and have been detailed in [AA96]—which is why we will not provide any details about them here. Here are the relevant properties of $\vDash$ and $\vDash_4$:

PROPOSITION 2.10 [AA96]: Both $\vDash$ and $\vDash_4$ are monotonous and paraconsistent (i.e. $p, \neg p \nvDash q$ and $p, \neg p \nvDash_4 q$), and have cut free, sound, and complete Gentzen-type proof systems ($GBL_{\supset}$ and $GBL_{\supset}(4)$)—that is, with the notations of definition 2.9, $\Gamma \vDash \Delta$ iff $\Gamma \vdash_{GBL_{\supset}} \Delta$ and $\Gamma \vDash_4 \Delta$ iff $\Gamma \vdash_{GBL_{\supset}(4)} \Delta$, and both $\vdash_{GBL_{\supset}}$ and $\vdash_{GBL_{\supset}(4)}$ are transitive in the sense of definition 2.6 above.

COROLLARY 2.11: $\vDash$ and $\vDash_4$ are scrs.

Proof.

Proposition 2.10 gives monotonicity and transitivity (the latter via the fact that $\vDash$ is coextensive with $\vdash_{GBL_{\supset}}$, and $\vDash_4$ is coextensive with $\vdash_{GBL_{\supset}(4)}$). On the other hand, it is quite obvious from the way there were defined, that both $\vDash$ and $\vDash_4$ are reflexive. The conditions of definition 2.6 have thus been fulfilled. ■

Finally, two further relations that we will need are derived from the above $\vDash$ and $\vDash_4$. [AA96] defines the symmetric versions $\vDash^s$ and $\vDash_4^s$ of $\vDash$ and $\vDash_4$ respectively, by imposing certain conditions on the latter two. Given that these definitions have no impact for us, we will just state the relevant properties of $\vDash^s$ and $\vDash_4^s$:

PROPOSITION 2.13 [AA96]: $\vDash^s$ and $\vDash_4^s$ are:

a) non-monotonic CRs;

b) stronger than their non-symmetric correspondents (that is, if $\Gamma \vDash^s \psi$, then $\Gamma \vDash \psi$, and if $\Gamma \vDash_4^s \psi$, then $\Gamma \vDash_4 \psi$).

We can now proceed to cover the corresponding implication connectives.

5. Implication Connectives: General Theory and Particular Choices

It is impossible to tell what makes certain authors dub certain propositional connectives as "implications." This is, as a matter of fact, one of the main reasons why Avron has felt compelled to isolate a class of connectives that "certainly deserve" to be called implication connectives. An outline of this class of connectives is given in [Av91]:

DEFINITION 2.14 [Av91]: Given a CR $\Vdash$, we call a binary connective $\twoheadrightarrow$ an *internal implication* relative to (or for) $\Vdash$ if $\Gamma, \psi \Vdash \varphi$ iff $\Gamma \Vdash \psi \twoheadrightarrow \varphi$.

Our hope is, now, that the connective $\supset$ defined above is an internal implication for $\vDash$ (which, of course, would make it an internal implication for $\vDash_4$ as well). This is, indeed, the case, although, again, we have not been able to locate a statement to this effect in Avron's papers that deal with bilattice-based logics. So let us fill this gap here:

PROPOSITION 2.15: The propositional connective $\supset$ introduced above in definition 2.4 is an internal implication for $\vDash$.

Proof.

[AA96] proves the following properties of $GBL_{\supset}$ (and $GBL_{\supset}(4)$): (a) $\Gamma, \psi \vdash_{GBL_{\supset}} \varphi$, then $\Gamma \vdash_{GBL_{\supset}} \psi \supset \varphi$; (b)

$\psi, \psi \supset \varphi \vdash_{GBL_{\supset}} \varphi$. Given the soundness of $GBL_{\supset}$ (proposition 2.10 above), (a) translates into "$\Gamma, \psi \vDash \varphi$, then $\Gamma \vDash \psi \supset \varphi$," which is half of our proof. The other direction is also easy to prove—though not quite as trivial—and also proceeds via the syntactic route. In virtue of $GBL_{\supset}$'s soundness, what we have to prove in fact is that if $\Gamma \vdash_{GBL_{\supset}} \psi \supset \varphi$, then $\Gamma, \psi \vdash_{GBL_{\supset}} \varphi$. To show this, it suffices to apply cut elimination (transitivity of $\vdash_{GBL_{\supset}}$) to (b) and $\Gamma \vdash_{GBL_{\supset}} \psi \supset \varphi$: we get $\Gamma, \psi \vdash_{GBL_{\supset}} \varphi$, which in virtue of $GBL_{\supset}$'s soundness, translates into "$\Gamma, \psi \vDash \varphi$." ■

Given that $\supset$ has now gained recognition as a legitimate implication connective, the properties we used in the previous proof according to which (a) if $\Gamma, \psi \vdash_{GBL_{\supset}} \varphi$, then $\Gamma \vdash_{GBL_{\supset}} \psi \supset \varphi$ and (b) $\psi, \psi \supset \varphi \vdash_{GBL_{\supset}} \varphi$ can be recast as:

PROPOSITION 2.16 [AA69]: In $GBL_{\supset}$ deduction theorem and *modus ponens* hold for $\supset$.

The implication connective $\supset$ will be henceforth dubbed *internal implication*. This implication, unfortunately, suffers from two major shortcomings: two sentences that "imply" each other (i.e. we have that $\vDash \varphi \supset \psi$ and $\vDash \varphi \supset \psi$ are true) cannot be deemed equivalent (i.e. inter-substitutable). The second 'drawback' is that it contradicts our intuitions vis-à-vis implication connectives in yet another way: the premise of a true implication with false conclusion is not necessarily false. As a consequence, one further connective has been introduced:

DEFINITION 2.16 [AA96]: (*strong implication*)
$\varphi \to \psi \triangleq (\varphi \supset \psi) \& (\neg \psi \supset \neg \varphi)$.

This, turns out to be an internal implication for $\vDash^{s}$ (and $\vDash_4^{s}$):

PROPOSITION 2.17 [AA96]: $\Gamma, \psi \vDash^{s} \varphi$ iff $\Gamma \vDash^{s} \psi \to \varphi$ and $\Gamma, \psi \vDash_4^{s} \varphi$ iff $\Gamma \vDash_4^{s} \psi \to \varphi$.

Finally we have:

DEFINITION 2.18: (*material "implication"*)
$\varphi \rightsquigarrow \psi \triangleq \neg \varphi \vee \psi$.

Note that, as discussed above, this last "implication" connective bears this name only in virtue of its association with the material implication connective of classical propositional logic. It has absolutely no claims of representing an implication connective in a bilattice-based logic context in any rigorous way. The three connectives thus defined nevertheless *do* overlap classical material implication if we are to confine ourselves to the standard propositional *fragment* of $BL_{\supset}(4)$.

The theoretical background has been finally laid. What remains now is to give a short presentation of the use of these three connectives in the context of ontology development efforts.

6. Reasoning with Paraconsistent Ontologies: ParOWL/RaDON

The practical value of 4-valued paraconsistent logics resides in their capacity to isolate contradictions, so that this phenomenon—impossible to avoid in real world data processing and knowledge base development—does not end up infecting the rest of the data. This is the main motivation behind Belnap's pioneering efforts ([Be92]), and has been hailed as one of the most promising venues for dealing with large heterogeneous databases.

Our interest lies with a certain type of data representation, namely OWL-based ontologies. Under the auspices of the Semantic Web, ontologies have flourished rapidly, such that, as of this writing, there are literally thousands of such knowledge repositories devoted to most knowledge fields. Built against a Description Logic background, OWL-based ontologies can benefit from the existence of consistency-checking and "classifying" devices ("reasoners"). ParOWL ([Ma07]) is one such reasoner that has been developed precisely with paraconsistent ontologies in mind.

The proliferation of ontologies has also led to the existence of competing (or, at least, overlapping) ontologies. It is, in this respect, quite natural to attempt at consolidating the data by merging competing/overlapping ontologies. And as it happens, the merger process often results in contradictory data, which is where paraconsistent reasoners' value comes into play.

Before expounding on a particular case of ontology merging effort, a few words about ParOWL in its latest incarnation (RaDON) are in order.

One of the features that recommend RaDON (nowadays part of the NeOn toolkit[‡]) is that it allows for simultaneous usage of the three different inclusions studied above. Since OWL only recognizes one type of inclusion/implication, an inconsistent ontology would have to be partitioned into three ontologies, each containing only one type of implication. The splitting process will have to be performed more or less manually by the ontologist and the domain expert. The three files plus one .owl file containing instance-level statements constitute the input to the paraconsistent reasoner.

RaDON has been designed so as to handle the latest OWL version (OWL version 2), based on the $\mathcal{SROIQ}$ variant of Description Logic, whereas the old ParOWL was tailored to the much less expressive $\mathcal{ALC}$ DL flavor. Consequently, the supporting RaDON paper ([Ma09]) devotes significant space to the developing of a four-valued semantics for $\mathcal{SROIQ}$, which is where we refer the reader in need of further details about the innards of the RaDON reasoner.

‡ http://www.neon-toolkit.org

7. Merging Ontologies

A concrete ontology merging effort involving ParOWL has made the subject of [Im09]. The merger consists of several steps, one of which requires

> A[n] … opinion from the domain experts about different inclusions [that] can be used … . Before invoking the reasoner, in the 4th step of the merging process, the new system provides an interface to capture and classify the type of class inclusions and produces four OWL files for ParOWL.

The significant detail is that domain experts are, in the ideal scenario, supposed to analyze those (and only those) class inclusions and implications in the inconsistent ontology that are responsible for this status, and separate them into three .owl files. Leaving aside the fact that most of the times it is really hard to pinpoint the exact culprits in an inconsistency situation, the experts are expected to discern between the three types of implications based on a set of application rules. These rules will preoccupy us in the following.

8. Implication Connectives: Application Criteria

It is our feeling that the guidelines for the application of the three "major notions of implication used in the literature" could use some improvement.

For a quick review of the guidelines, let us start with [AA96], where these implication connectives were initially introduced.

Avron works out a concrete instance that supposedly illustrates their paradigm use. He builds a toy knowledge base made up of three "if … then …" laws:

1. If tweety is a bird, then tweety flies;
2. If tweety is a penguin, then tweety is a bird;
3. If tweety is a penguin, then tweety does not fly.

The knowledge base is rounded up with:

4. tweety is a bird.

It is easy to see that, if the three "if … then …" connectors are to be captured my means of the same connector as the material implication of standard bivalent logic, the knowledge base collapses in inconsistency. By couching the above using the resources of 4-valued logic and $BL_{\supset}$, one will avoid such an outcome: The three laws above will have to be represented via the three implication connectives described above: strong ($\rightarrow$), internal ($\supset$) and material ($\leadsto$). The problem, as we see it, is which one goes where. And why.

Avron's guidelines rely on the "strength" of the implication: represent each "law" according to (as far as we can tell) whether they allow for counterexamples. Number 3 above, for example, is supposed to be represented as a strong implication, as there is no (known) exception to it:

$$penguin(tweety) \rightarrow \neg flyingAnimal(tweety)$$

The second proposition certainly has no exceptions, though due to the fact that "penguins are not typical birds, … they should not inherit all the properties we expect birds to have." Avron thus suggests that this should be modeled via an internal implication:

$$penguin(tweety) \supset bird(tweety)$$

Leaving aside the questionable individuation principle that transpires from the fact that a bird might not have all the characteristics that a bird has, we certainly expect some further clarification as to how the formal definition of an internal implication leads to its application in such cases of atypical birds. Finally, the first proposition, which is patently false in a classical world, gets to be captured via material implication, as it "states only a default feature of birds. … [Therefore it] does not cause automatic inference of flying abilities just from the fact that something is a bird." While this is certainly beyond controversy (given that $\leadsto$ does not satisfy *modus ponens*), it is less clear for someone looking to apply $\leadsto$ how strong the connection between penguin and flyingAnimal should be in order to be represented by material implication.

In [Ma07] strong inclusion is advertised as something that is appropriate for modeling of "universal truth," though the authors stop short of defining that unambiguously. The best [Ma07] suggestion might be that "Internal inclusion should be used whenever it is important to infer the consequent even if the antecedent may be contradictory," though this is still lacking in terms of usability in real world ontology development situations, due to, among others, the presence of the "important" qualifier. The scenario of the material implication benefits from a larger discussion, which can be summarized as follows: an implication should be represented as a material implication if, among others, one does not need (the 4-valued version of) *modus ponens*, nor does one want contradictory information propagated backwards. The latter stems from what the authors dub "the weak form of contraposition reasoning," which, presumably, looks like this:

$a \leadsto b$ holds (i.e. is designated)

b doesn't hold (i.e. is not designated)

$\therefore a$ is told false (i.e. false or $\top$)

Finally, another hint comes via the authors' insistence that should one use "only one kind of inclusion," one should lend preference to strong inclusion, as this "should serve the ontology engineer's original intention most closely." We are, however, left wondering what this "original intention" might be, as this does not transpire from the paper with sufficient clarity.

It is, again, important to stress at this point that we are regarding the issue from the point of view of ontology engineers engaged in the actual development of real-world ontologies, hence working out clear employment criteria for the three implications is crucial for the success of our enterprise. The above hints do leave room for improvement: they, in our opinion, allow too much of a leeway when it comes to making a useful decision in picking between the three in real-world cases.

It is also important to clarify what it means to make a decision in a concrete physical situation based on a four-valued proposition. From the semantics presented above, a good guess might be that the designated values represented the new "true," hence should one, for example, be confronted with the sentence "This house is on fire" and told that its truth value is $\top$, one should proceed as if the house were indeed on fire. Let us turn to [Ma07] again for clarifications. The authors recommend that internal implication be used "whenever it is important to infer the consequent even if the antecedent may be contradictory." What does it mean "to infer" though? Inferring the consequent could either mean asserting it as true, or giving it truth value $\top$—certainly nothing else. A quick look at the truth table (see below Figure 1) makes it clear that if by asserting the implication one means that the implication is true, a contradictory ($\top$) antecedent would make the consequent also true, whereas if the compound proposition is taken to have a designated truth value, a contradictory antecedent would make the consequent designated as well. It is thus quite conceivable that natural language terms like "infer," "assert" etc. can be legitimately used in either way in a 4-valued context—as long, of course, as the use stays consistent throughout—though for the remainder, we will adopt the stance that such terms are indicative of designated values.

In order to produce further guidelines for the real-life application of the three types of implication connectives, we will attempt to extract interesting and useful rules of inference by carefully inspecting the truth table for the respective connectives. We will borrow the master truth table from [Ma07], where $\rightsquigarrow$, $\supset$ and $\rightarrow$ stand for $\alpha \rightsquigarrow \beta$, $\alpha \supset \beta$ and $\alpha \rightarrow \beta$ respectively:

α	f	f	f	f	t	t	t	t	$\top$	$\top$	$\top$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$
β	f	t	$\top$	$\bot$	f	t	$\top$	$\bot$	f	t	$\top$	$\bot$	f	t	$\top$	$\bot$
$\rightsquigarrow$	t	t	t	t	f	t	$\top$	$\bot$	$\top$	t	$\top$	t	$\bot$	t	t	$\bot$
$\supset$	t	t	t	t	f	t	$\top$	$\bot$	f	t	$\top$	$\bot$	t	t	t	t
$\rightarrow$	t	t	t	t	f	t	f	$\bot$	f	t	$\top$	$\bot$	$\bot$	t	$\bot$	t

Figure 1: Master truth table ([Ma07])

To facilitate our search, we will re-draw the implication part of the table by grouping the elementary truth values into groups of possible relevance as follows: $D = \{t, \top\}$ represents the designated "truth value", $D' = \{f, \bot\}$ is its complementary ("no evidence of being true"), $kf = \{f, \top\}$ stands for "told (or known) false," and $kf' = \{t, \bot\}$ its complementary ("no evidence of being false"). With these, (parts of) the master table can be rendered in, among others, the ways illustrated in figures 2 and 3 below.

α	$\alpha \rightsquigarrow \beta$	$\alpha \supset \beta$	$\alpha \rightarrow \beta$	$\beta_{\rightsquigarrow}$	$\beta_{\supset}$	$\beta_{\rightarrow}$
t				D	D	t
f				$D \cup D'$	$D \cup D'$	$D \cup D'$
$\top$				$D \cup D'$	D	D
$\bot$	D			D	$D \cup D'$	kf'
D				$D \cup D'$	D	D
kf				$D \cup D'$	$D \cup D'$	$D \cup D'$
kf'				D	$D \cup D'$	kf'

Figure 2: β table

These tables show the truth value of the consequent (last three columns (β, resp. α)) in an argument where the premises are the first column (α, resp. β) plus either of columns two to four (one of the three implications). Obviously, the $D \cup D' = \{t, f, \top, \bot\}$ cells cannot furnish significant facts, so we would have to look elsewhere for interesting correlations.

The α table allow us to recover two major facts inventoried by [Ma07], namely "weak contraposition" of material inclusion (see above), and contraposition of strong inclusion:

$\alpha \rightarrow \beta$ holds (D)

β is told false (kf)

$\therefore \alpha$ is told false (kf).

β	$\alpha \rightsquigarrow \beta$	$\alpha \supset \beta$	$\alpha \rightarrow \beta$	$\alpha_{\rightsquigarrow}$	$\alpha_{\supset}$	$\alpha_{\rightarrow}$
D'				kf	D'	D'
t				$D \cup D'$	$D \cup D'$	$D \cup D'$
f				kf	D'	f
$\top$	D			$D \cup D'$	$D \cup D'$	kf
$\bot$				kf	D'	D'
kf				$D \cup D'$	$D \cup D'$	kf
kf'				$D \cup D'$	$D \cup D'$	$D \cup D'$

Figure 3: α table

Aside from these two, we regard the following three argument classes as having potential in the context of knowledge base development:

1	2	3
$\alpha \rightsquigarrow \beta$ (D) α (kf') $\therefore \beta$ (D)	$\alpha \rightarrow \beta$ (D) α (kf') $\therefore \beta$ (kf')	$\alpha \supset \beta$ (D) β (D') $\therefore \alpha$ (D')

In the first case, having no evidence of being false leads to the asserting of the consequent of the material implication, which thus can be construed as an implication connective very sensitive to truth.

The second argument says, in essence, that no information about the falsity of the antecedent of a strong implication leads to no information about the falsity of the consequent, which is what one would expect from an implication connective that has been designed to transmit (forward) only truth/designated values. This rule thus rounds off the image of strong implication as the implication to use when it is important to not "produce" information out of nothing—which is quite the opposite of the previous rule.

The third argument might be interpreted as a way to "refute" the antecedent of an internal implication, given the non-designated status of the consequent, though care should be taken not to confuse this with contraposition, which does not hold for this connective.

9. Summary and Conclusions

Our endeavor was initially intended as a simple and quick attempt at extracting a minimum of theoretical knowledge required in order to proceed at applying paraconsistent techniques to the development of knowledge bases (and ontologies in particular). We have, nevertheless, found ourselves compelled to escalate it into a rather lengthy presentation of some of the theoretical background of bilattice-based logic, primarily due to the fact that Avron's writings on the subject—our main source of inspiration—are not only markedly abstract and vast, but also rather far from a minimally friendly exposure destined to appeal to the casual applied logician in search of a quick many-valued logic fix. We have consequently undertaken to (among others) harmonize the notations and terminology, select the relevant facts and fill some of the more glaring gaps, all in the hope that this new presentation might be closer to the needs of the actual working ontologist.

While this may primarily be regarded as a work of bilattice-based logic synthesis, we have also ventured to present some applied part in the guise of paraconsistent reasoners that utilize the theoretical principles expounded. In this respect, our attention has focused on an essential step involving domain expert and/or ontology engineer input, whereby compound natural language sentences comprising the "if … then …" propositional connective are to be translated in bilattice-based language; this translation, we remarked, is insufficiently supervised/guided, hence leaves a certain interpretational latitude whose outcome could ultimately affect the accuracy of the resulting knowledge base or ontology. The cure, as we saw it, was to improve the translation guidelines—and this by subjecting the corresponding propositional constructs of the bilattice-based language to further analysis from various angles. In our view, and based on our ontology development experience, our resulting proposals, while far from revolutionary, have the potential to make a difference. Nevertheless, we[§] see the continuation of such efforts as a promising ontology development research venue.

References

[AA00] Arieli, O., Avron A., 2000. "Bilattices and Paraconsistency" in *Frontiers in Paraconsistent Logic* (edited by D. Batens, C. Mortensen, G. Priest, and J. Van-Bengedem), 11-28, Research Studies Press.

[AA94] Arieli, O., Avron A., 1994. "Logical Bilattices and Inconsistent Data," *Proc. 9th IEEE Annual Symp. On Logic in Computer Science*, IEEE Press, 468-476.

[AA96] Arieli, O., Avron A., 1996. "Reasoning with logical bilattices," *Journal of Logic, Language, and Information*, Vol.5, No.1, pages 25-63.

[AA98] Arieli, O., Avron A., 1998. "The Value of Four Values," *Artifical Intelligence* 102, 97-141.

[An08] Antonelli, G.A., 2008. "Non-monotonic Logic," *The Stanford Encyclopedia of Philosophy (Winter 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2008/entries/logic-nonmonotonic/>.

[Av10] Avron, A. 2010. "Tonk – A Full Mathematical Solution," to appear in Ruth Manor Festchrift.

[Av91] Avron, A., 1991. "Simple Consequence Relations," *Information and Computation* 92, 105-139

[Av91b] Avron, A., 1991. "Natural 3-valued Logics—Characterization and Proof Theory," *Journal of Symbolic Logic* 56, 276-294.

[Be09] Beall, J.C., Restall, G., 2009. "Logical Consequence," *The Stanford Encyclopedia of Philosophy (Fall 2009 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2009/entries/logical-consequence/>.

[Be92] Anderson, A.R., Belnap, N.D., Dunn, J.M. 1992. *Entailment: the Logic of Relevance and Necessity*, volume 2, Princeton University Press.

[Im09] Imam, F., MacCaull, W., 2009. "Integrating Healthcare Ontologies: An Inconsistency Tolerant Approach and Case Study," *Business Process Management Workshops*, Lecture Notes in Business Information Processing, Springer Berlin Heidelberg, V. 17 373-384.

[Ma07] Ma, Y., Hitzler, P., Lin, Z., 2007. "Algorithms for Paraconsistent Reasoning with OWL," in: Enrico Franconi, Michael Kifer, Wolfgang May (eds.), *The Semantic Web: Research and Applications. Proceedings of the 4th European Semantic Web Conference*, ESWC2007, Innsbruck, Austria, June 2007. Springer Lecture Notes in Computer Science 4519, pp. 399-413.

[Ma09] Ma, Y., Hitzler, P., 2009. "Paraconsistent Reasoning for OWL 2," in: Axel Polleres, Terrance Swift (Eds.), *Web Reasoning and Rule Systems, Third International Conference*, RR 2009, Chantilly, VA, USA, October 2009, Proceedings. Lecture Notes in Computer Science Vol. 5837, Springer, pp. 197-211.

[§] As well as others—e.g. [Ma07], [Ma09].