Notes on Bilattices

Melvin Fitting

April 15, 2018

1 Pre-bilattices

This is probably more than you need to know, but the material is rather straightforward. If you are unfamiliar with the notion of a lattice, a good example to keep in mind is all the subsets of some set, with $\subseteq$ as the ordering relation. All lattices I consider will have tops and bottoms—largest and smallest elements. (For the set example, the top is the entire set, and the bottom is the empty set.) To keep terminology simple, in everything that follows the term *lattice* means lattice with a top and a bottom.

If you are not familiar with the terminology, in a lattice the greatest lower bound and the least upper bound of two-element sets is required to exist (and hence also for any finite set). The greatest-lower-bound operation is usually called *meet* and the least-upper-bound operation is called *join*. Lattice meet and join operations are always commutative and associative. Requiring a top and a bottom amounts to saying there is an element bigger than all others and an element smaller than all others.

Definition 1.1 A *pre-bilattice* is a structure $\mathcal{B} = \langle \mathcal{B}, \leq_t, \leq_k \rangle$ in which $\mathcal{B}$ is a non-empty set and $\leq_t$ and $\leq_k$ are partial orderings each giving $\mathcal{B}$ the structure of a lattice.

Think of the members of $\mathcal{B}$ as pieces of information that are act as truth values in some generalized sense. It is to deal with this dual role, information and degree of truth, that we have the two ordering relations.

The ordering $\leq_k$ should be thought of as ranking "degree of information". Thus if $x \leq_k y$, y gives us at least as much information as x (and possibly more). I suppose this really should be written as $\leq_i$, using i for information instead of k for knowledge, but k has become standard. The meet and join operations for $\leq_k$ are denoted $\otimes$ and $\oplus$. The $\otimes$ operation is called *consensus*: $x \otimes y$ is the most information that x and y agree on. The $\oplus$ operation is called *gullability*—a person who is gullable will believe anything. Then $x \oplus y$ should be thought of as combining the information in x with that in y, without worrying about whether the pieces fit together or not. The bottom in the $\leq_k$ ordering is denoted by $\bot$ and the top by $\top$. Think of $\bot$ as representing the state of complete ignorance—no information. Likewise $\top$ represents full information, possibly including inconsistencies.

The relation $\leq_t$ is an ordering on the "degree of truth." The bottom in this ordering will be denoted by *false* and the top by *true*. Thus *false* $\leq_t x \leq_t$ *true* for any $x \in \mathcal{B}$. The meet and join operations for $\leq_t$ will be denoted by $\wedge$ and $\vee$. It is easy to check that when restricted to *false* and *true*, these obey the usual truth-table rules. It is also easy to check that when restricted to *false*, $\bot$ and *true* they obey the rules of Kleene's strong three-valued logic (this works equally well if we restrict to *false*, $\top$ and *true*, but it is better to think of this as a version of Priest's logic LP).

In a lattice, meets and joins of finite sets must exist. What is called a completeness assumption extends this to infinite sets as well. Completeness is needed to adequately interpret quantifiers. Here is the bilattice version of completeness.

Definition 1.2 A pre-bilattice $\langle \mathcal{B}, \leq_t, \leq_k \rangle$ is *complete* if all meets and joins exist, with respect to both orderings. I'll denote infinitary meet and join with respect to $\leq_t$ by $\bigwedge$ and $\bigvee$, and by $\prod$ and $\sum$ for the $\leq_k$ ordering.

1.1 Examples

Suppose we have a certain group of people, $\mathcal{P}$, whose opinions we value. If we ask these people about the status of a sentence X, some will call it true, some false. But also, some may decline to express an opinion, and some may be uncertain enough to say they have reasons for calling it both true and false. We can, then, assign X a kind of generalized truth value, $\langle P, N \rangle$, where P is the set of people in $\mathcal{P}$ who say X is true and N is the set who say it is false. As just noted, we do not require that $P \cup N = \mathcal{P}$, nor that $P \cap N = \emptyset$.

Orderings can be introduced into our people-based structure: set $\langle P_1, N_1 \rangle \leq_k \langle P_2, N_2 \rangle$ if $P_1 \subseteq P_2$ and $N_1 \subseteq N_2$, and set $\langle P_1, N_1 \rangle \leq_t \langle P_2, N_2 \rangle$ if $P_1 \subseteq P_2$ and $N_2 \subseteq N_1$ (note the reversal here). Thus, information goes up if more people express a positive or negative opinion, and truth goes up if people drop negative opinions or add positive ones. This gives us the structure of a pre-bilattice. In it, for example, $\langle P_1, N_1 \rangle \wedge \langle P_2, N_2 \rangle = \langle P_1 \cap P_2, N_1 \cup N_2 \rangle$, and $\langle P_1, N_1 \rangle \otimes \langle P_2, N_2 \rangle = \langle P_1 \cap P_2, N_1 \cap N_2 \rangle$. Reflection should convince you that these are quite natural operations. Also, $\bot = \langle \emptyset, \emptyset \rangle$, $\top = \langle \mathcal{P}, \mathcal{P} \rangle$, $\mathit{false} = \langle \emptyset, \mathcal{P} \rangle$, and $\mathit{true} = \langle \mathcal{P}, \emptyset \rangle$. You should reflect on these too.

As another example, consider a "fuzzy" truth value space, in which truth values are pairs $\langle p, n \rangle$ of real numbers in the interval $[0, 1]$, where p is "degree of belief," and n is "degree of doubt." Appropriate orderings for this example are $\langle p_1, n_1 \rangle \leq_k \langle p_2, n_2 \rangle$ if $p_1 \leq p_2$ and $n_1 \leq n_2$; and $\langle p_1, n_1 \rangle \leq_t \langle p_2, n_2 \rangle$ if $p_1 \leq p_2$ and $n_2 \leq n_1$.

The two examples above can be combined if we consider a collection of people, each of whom has "fuzzy" opinions. I won't follow up on this—you probably get the general idea.

Figure 1 shows the simplest non-trivial example of a pre-bilattice: only the four extreme elements exist and are distinct. It can be thought of as a special case of the people pre-bilattice above, in which there is only one person. This is a fundamental example, and originated before bilattices as such arose—it is the four-valued logic due to Belnap, [1, 2], and will be called $\mathcal{FOUR}$ here. Think of the left-right direction as characterizing the $\leq_t$ ordering: a move to the right is an increase. The meet operation for the $\leq_t$ ordering, $\wedge$, is then characterized by: $x \wedge y$ is the rightmost thing that is left of both x and y. The join operation, $\vee$ is dual to this. In a similar way the up-down direction characterizes the $\leq_k$ ordering: a move up is an increase in information. $x \otimes y$ is the uppermost thing below both x and y, and $\oplus$ is dual. Spatial conventions like these will be used throughout.

Figure 2 shows a pre-bilattice in which subtler distinctions can be registered. As is also the case with $\mathcal{FOUR}$, $\bot$ represents a state of complete ignorance, and $\top$ one of information overload—solid evidence has been supplied both for and against some proposition. Likewise *false* represents the situation in which we have convincing evidence against some proposition, and no evidence in its favor, while *true* is just the opposite. But in Figure 2 there are two more states. Think of *fd* as a state in which we have no evidence in favor of a proposition, but we have some weak evidence against—read *fd* as "false with doubts." Think of *td* likewise as "true with doubts."

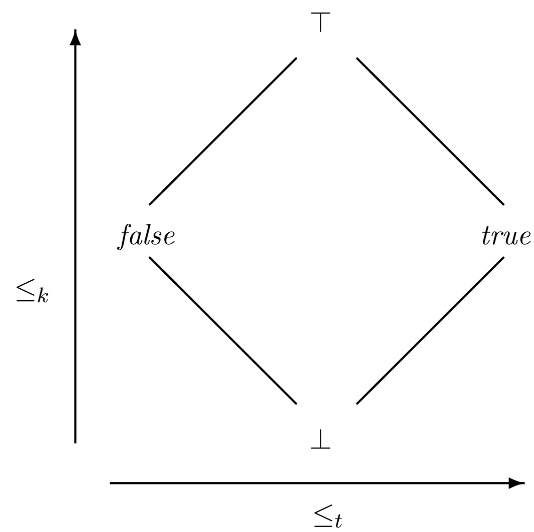

Figure 1: The Bilattice $\mathcal{FOUR}$

1.2 Bilattices

A pre-bilattice has two orderings, with no postulated connections between them. I'll reserve the term *bilattice* for pre-bilattices where there are useful connections between orderings. Ginsberg's original definition of bilattice postulated a connection through a negation operation. Here I will use stronger notions that also trace back to Ginsberg, [8].

Definition 1.3 A pre-bilattice $\langle \mathcal{B}, \leq_t, \leq_k \rangle$ is:

1. *an interlaced bilattice* if each of the operations $\wedge$, $\vee$, $\otimes$, and $\oplus$ is monotone with respect to both orderings (the interlacing conditions);

2. *an infinitarily interlaced bilattice* if it is complete and all four infinitary meet and join operations are monotone with respect to both orderings;

3. *a distributive bilattice* if all 12 distributive laws connecting $\wedge$, $\vee$, $\otimes$, and $\oplus$ are valid;

4. *an infinitarily distributive bilattice* if it is complete and infinitary, as well as finitary, distributive laws are valid. Examples of infinitary distributive laws are: $a \wedge \sum_i b_i = \sum_i (a \wedge b_i)$, and $a \otimes \bigwedge_i b_i = \bigwedge_i (a \otimes b_i)$.

A lattice is called distributive if it satisfies distributive laws; for example, a pre-bilattice is a lattice with respect to the $\leq_k$ ordering, and this lattice is distributive if $x \otimes (y \oplus z) = (x \otimes y) \oplus (x \otimes z)$ and $x \oplus (y \otimes z) = (x \oplus y) \otimes (x \oplus z)$ holds. Saying a pre-bilattice is distributive requires that we have distributive lattices with respect to both orderings and, in addition, we have "mixed" distributive laws, such as $x \otimes (y \vee z) = (x \otimes y) \vee (x \otimes z)$. All examples from Section 1.1 are distributive bilattices, and infinitary distributivity is satisfied as well.

In a lattice, meet and join operations are always monotone with respect to the lattice ordering. Thus we always have that $x_1 \leq_t y_1$ and $x_2 \leq_t y_2$ implies $(x_1 \wedge y_1) \leq_t (x_2 \wedge y_2)$. Saying we have an interlaced bilattice adds to this the requirement that monotonicity also work "across" orderings;

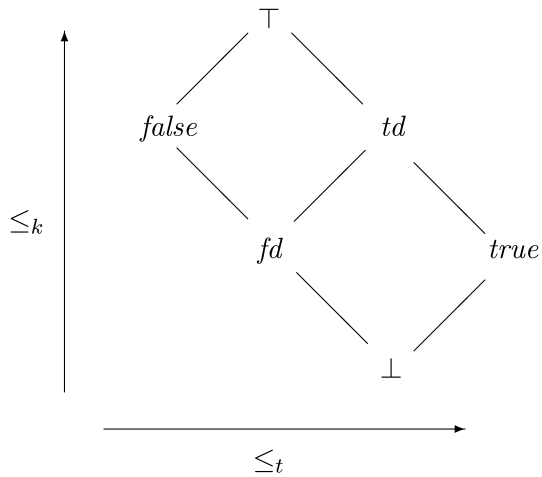

Figure 2: A Six-Valued Bilattice

for example $x_1 \leq_k y_1$ and $x_2 \leq_k y_2$ implies $(x_1 \wedge y_1) \leq_k (x_2 \wedge y_2)$. It is not hard to show that every (infinitarily) distributive bilattice is also (infinitarily) interlaced, hence the examples of Section 1.1 satisfy the interlacing conditions.

1.3 Negation and Conflation

Some bilattices have natural symmetries, and these can be used to characterize interesting subsystems.

Definition 1.4 A bilattice has a *negation* operation if there is a mapping, $\neg$, that reverses the $\leq_t$ ordering, leaves unchanged the $\leq_k$ ordering, and $\neg\neg x = x$. Likewise a bilattice has a *conflation* operation if there is a mapping, $-$, that reverses the $\leq_k$ ordering, leaves unchanged the $\leq_t$ ordering, and $- - x = x$. If a bilattice has both operations, they *commute* if $-\neg x = \neg - x$ for all x.

In the people example of Section 1.1, there are natural notions of negation and conflation. Take $\neg\langle P, N\rangle$ to be $\langle N, P\rangle$—the roles of for and against are switched. And take $-\langle P, N\rangle$ to be $\langle \mathcal{P} - N, \mathcal{P} - P\rangle$, where $\mathcal{P}$ is the set of people. This amounts to a kind of switching to a default position—the people who affirm under a conflation are the people who originally did not deny, for instance. The "fuzzy" example has a similarly defined negation and conflation—I'll leave their formulation to you. For both examples, negation and conflation commute.

In the example of Figure 1, there is a negation operation under which $\neg$*true* = *false*, $\neg$*false* = *true*, and $\bot$ and $\top$ are left unchanged. There is also a conflation under which $-\bot = \top$, $-\top = \bot$ and *true* and *false* are left unchanged. In this example negation and conflation commute. In any bilattice, if a negation or conflation exists the behavior on the extreme elements $\bot$, $\top$, *false*, and *true* will be as it is in $\mathcal{FOUR}$.

The example of Figure 2 does not have either a negation or a conflation. One might, for instance, try introducing a negation by adding to the usual conditions for the extreme elements the requirement that $\neg$*td* = *fd* and $\neg$*fd* = *td*, but this will not work. We have *fd* $\leq_k$ *false* and negation is required not to affect the $\leq_k$ ordering, so we should have *td* $\leq_k$ *true*, but in fact we have the opposite. There is a deeper reason for the lack of conflation and negation in this example that will become clear in the next section.

Definition 1.5 Suppose $\mathcal{B}$ is a bilattice with a conflation operation. Call $x \in \mathcal{B}$ *exact* if $x = -x$ and *consistent* if $x \leq_k -x$.

In the bilattice example involving people, Section 1.1, the exact values are those $\langle P, N \rangle$ where N is the complement of P—everyone expresses an unambiguous opinion. The consistent values are those where $P \cap N = \emptyset$, that is, people may be undecided, but they are never contradictory. In the bilattice $\mathcal{FOUR}$ of Figure 1, the exact members are $\{\mathit{false}, \mathit{true}\}$, the classical truth values, and the consistent ones are $\{\mathit{false}, \bot, \mathit{true}\}$, which behave like the values of Kleene's strong three-valued logic, with respect to $\neg$, $\wedge$, and $\vee$. This phenomenon, in fact, is not uncommon. The exact part of a complete bilattice with commuting conflation and negation is always closed under $\neg$, $\wedge$, and $\vee$, and similarly for the consistent part. In addition, the consistent part will always be closed under the infinitary version of $\otimes$, and under the infinitary version of $\oplus$ when applied to a directed set. It is essentially these conditions that were used in [4] for the special case of Kleene's strong three-valued logic, but in fact they obtain much more generally.

2 Constructing Bilattices

There are several ways of constructing bilattices that also provide some intuition concerning them. Only one is discussed here. It traces back to [9] with extensions of mine, though underlying ideas actually go back somewhat earlier.

Suppose we have notions of positive and negative evidence. For instance, positive evidence for a mathematical conjecture might consist of plausibility arguments, computer experiments, almost correct proofs, and so on. Actual proofs would be best possible, of course. Negative evidence might also consist of various informal arguments, with counter-examples as best possible. Let us say we have a way of ranking evidence—this piece is better than that. More formally, say we have two lattices, $\mathcal{L}_1 = \langle \mathbf{L}_1, \leq_1 \rangle$ and $\mathcal{L}_2 = \langle \mathbf{L}_2, \leq_2 \rangle$, where members of $\mathbf{L}_1$ are things that can serve as positive evidence, with $\leq_1$ as a comparison relation, and similarly for $\mathbf{L}_2$ as pieces of negative evidence. The lattices need not be the same.

Definition 2.1 A *bilattice product* $\mathcal{L}_1 \odot \mathcal{L}_2$ is the structure $\langle \mathbf{L}_1 \times \mathbf{L}_2, \leq_t, \leq_k \rangle$ where:

1. $\langle x_1, x_2 \rangle \leq_t \langle y_1, y_2 \rangle$ if $x_1 \leq_1 y_1$ and $y_2 \leq_2 x_2$
2. $\langle x_1, x_2 \rangle \leq_k \langle y_1, y_2 \rangle$ if $x_1 \leq_1 y_1$ and $x_2 \leq_2 y_2$

Think of a member $\langle x, y \rangle$ of $\mathbf{L}_1 \times \mathbf{L}_2$ as encoding evidence about some assertion: evidence for, x, and evidence against, y. Then an increase in information amounts to saying evidence in general goes up. An increase in truth says evidence for increases while evidence against decreases. Earlier examples concerning people and "fuzzyness" are both special cases of this construction.

It is straightforward to show that $\mathcal{L}_1 \odot \mathcal{L}_2$ is always an interlaced bilattice, and is complete if both $\mathcal{L}_1$ and $\mathcal{L}_2$ are complete as lattices. And further, if both $\mathcal{L}_2$ and $\mathcal{L}_2$ are distributive lattices, $\mathcal{L}_1 \odot \mathcal{L}_2$ will be a distributive bilattice.

If $\mathcal{L}_1 = \mathcal{L}_2$ then a negation operaton can be introduced into $\mathcal{L}_1 \odot \mathcal{L}_2$. Set $\neg\langle x, y \rangle = \langle y, x \rangle$. That is, negation switches the roles of positive and negative evidence. Next, suppose $\mathcal{L}_1 = \mathcal{L}_2 = \mathcal{L}$ has what is called a *de Morgen complement* operation, an operation that maps x to $\overline{x}$ such that $x \leq y$ implies $\overline{y} \leq \overline{x}$, and $\overline{\overline{x}} = x$. Then a conflation operation can also be introduced into the bilattice product: set $-\langle x, y \rangle = \langle \overline{y}, \overline{x} \rangle$. Defined these ways, negation and conflation will commute.

The machinery just set forth for constructing various kinds of bilattices is completely general. That is, every distributive bilattice is isomorphic to $\mathcal{L}_1 \odot \mathcal{L}_2$ for some distributive lattices $\mathcal{L}_1$ and $\mathcal{L}_2$, and similarly for the other cases. Proof can be found for the various parts of this family of results in [9, 5, 6, 7]. In a way the result is a descendant of the Polarities Theorem of Dunn, [3].

Consider the familiar lattice $\mathcal{B}$ whose carrier is $\{false, true\}$, with $false < true$. This is a distributive lattice for which the operation $\overline{false} = true$ and $\overline{true} = false$ is a de Morgen complement. Then $\mathcal{B} \odot \mathcal{B}$ is a distributive bilattice with a negation and a conflation. It is, in fact, isomorphically the bilattice $\mathcal{FOUR}$ of Figure 1. Further, let $\mathcal{C}$ be the lattice whose carrier is $\{0, \frac{1}{2}, 1\}$, ordered numerically. Then $\mathcal{B} \odot \mathcal{C}$ is isomorphically the bilattice of Figure 2. Since $\mathcal{B}$ and $\mathcal{C}$ are different, there is no negation or conflation. But since both are distributive lattices, $\mathcal{B} \odot \mathcal{C}$ is a distributive bilattice.

3 Kleene's strong three-valued logic generalized

Kleene's strong three-valued logic has probably been the most popular of the various logics used in Kripke's approach to self-reference. In this logic, for instance, $true \vee \bot = true$, informally because if we get enough further information about the second component of the disjunction to assign it a classical truth value, whether we find that component to be *true* or *false* we would still evaluate the disjunction to *true*, so extra information is not really relevant—we can assign *true* right now. Because of the nature of bilattices, this is the simplest partial logic for us to generalize. Let v be a valuation, an assignment of values in $\mathcal{B}$ to *atomic* sentences. This extends to a mapping on *all* sentences. I'll denote this extension by v^s; the superscript is for "strong".

1. $v^s(T(t)) = v(T(t))$.
2. $v^s(X \wedge Y) = v^s(X) \wedge v^s(Y)$
3. $v^s(X \vee Y) = v^s(X) \vee v^s(Y)$
4. $v^s(\neg X) = \neg v^s(X)$

In item 2 the occurrence of $\wedge$ on the left is syntactic—it is a symbol of the language; the occurrence of $\wedge$ on the right denotes the meet operation of $\mathcal{B}$ with respect to the $\leq_t$ ordering. Similar remarks apply to 3 and 4 as well.

4 Kleene's weak three-valued logic generalized

Kleene's weak three-valued logic assigns a value of $\bot$ to any compound formula in which some part has been assigned $\bot$. Thus, for instance, $true \vee \bot = \bot$, which is a different outcome than we get in Kleene's strong three-valued logic. The weak logic too can be generalized to the bilattice setting—[6] proposes an approach, but here I follow a different one. For motivation, consider once again the bilattice example based on people, from the beginning of Section 1.1. Suppose we have two bilattice values, $A = \langle P_1, N_1 \rangle$ and $B = \langle P_2, N_2 \rangle$, where the P_i and N_i are sets of people, those expressing opinions for, and against, respectively. Of course $A \wedge B$ was defined earlier, but suppose we want to 'cut this down' by only considering people who have actually expressed an opinion on both propositions A and B. As far as A is concerned, $A \oplus \neg A = \langle P_1 \cup N_1, P_1 \cup N_1 \rangle$, and taking the consensus, $\otimes$, of this with an arbitrary member of the people bilattice does, indeed, cut things down to those who have expressed an opinion concerning A. Similarly for B. To keep notational clutter down, suppose I write $\|X\|$ for $X \oplus \neg X$, so what we want for a 'cut down' conjunction is

$(A \wedge B) \otimes \|A\| \otimes \|B\|$. We can do a similar thing with disjunction, of course. Negation is somewhat simpler since $\neg A \otimes \|A\| = \neg A$, so we can avoid extra complication in this case.

This suggests we define the following operators for any complete bilattice with negation. The superscript w is for "weak," and in fact, confined to the consistent part of the bilattice $\mathcal{FOUR}$, they are the connectives of Kleene's weak three-valued logic.

1. $X \wedge^w Y = (X \wedge Y) \otimes \|X\| \otimes \|Y\|$
2. $X \vee^w Y = (X \vee Y) \otimes \|X\| \otimes \|Y\|$

Once again let v be a valuation in $\mathcal{B}$; I'll extend it to a mapping v^w on all sentences as follows.

1. $v^w(T(t)) = v(T(t))$.
2. $v^w(X \wedge Y) = v^w(X) \wedge^w v^w(Y)$
3. $v^w(X \vee Y) = v^w(X) \vee^w v^w(Y)$
4. $v^w(\neg X) = \neg v^w(X)$

References

[1] N. D. Belnap, Jr. How a computer should think. In G. Ryle, editor, *Contemporary Aspects of Philosophy*, pages 30–55. Oriel Press Ltd, Stocksfield, 1977.

[2] N. D. Belnap, Jr. A useful four-valued logic. In J. M. Dunn and G. Epstein, editors, *Modern Uses of Multiple-Valued Logic*, pages 8–37. D. Reidel, Dordrecht, 1977.

[3] J. M. Dunn. Relevance logic and entailment. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, Volume III*, chapter 3, pages 117–224. D. Reidel, 1986.

[4] M. C. Fitting. Notes on the mathematical aspects of Kripke's theory of truth. *Notre Dame Journal of Formal Logic*, 27:75–88, 1986.

[5] M. C. Fitting. Bilattices in logic programming. In G. Epstein, editor, *The Twentieth International Symposium on Multiple-Valued Logic*, pages 238–246. IEEE, 1990.

[6] M. C. Fitting. Kleene's three-valued logics and their children. *Fundamenta Informaticae*, 20:113–131, 1994.

[7] G. Gargov. Knowledge, uncertainty and ignorance in logic: bilattices and beyond. *Journal of Applied Non-Classical Logics*, 9, 1999.

[8] M. Ginsberg. Multi-valued logics. In *Proceedings of AAAI-86, Fifth National Conference on Artificial Intelligence*, pages 243–247, Los Altos, 1986. Morgan Kaufman.

[9] M. Ginsberg. Multivalued logics: a uniform approach to reasoning in AI. *Computational Intelligence*, 4:256–316, 1988.