Multi-valued logics

Matthew L. Ginsberg*

THE LOGIC GROUP
KNOWLEDGE SYSTEMS LABORATORY
Department of Computer Science
Stanford University
Stanford, California 94305

ABSTRACT

A great deal of recent theoretical work in inference has involved extending classical logic in some way. I argue that these extensions share two properties: firstly, the formal addition of truth values encoding intermediate levels of validity between true (i.e., valid) and false (i.e., invalid) and, secondly, the addition of truth values encoding intermediate levels of certainty between true or false on the one hand (complete information) and unknown (no information) on the other. Each of these properties can be described by associating lattice structures to the collection of truth values involved; this observation lead us to describe a general framework of which both default logics and truth maintenance systems are special cases.

1 Introduction

There has been increasing interest in AI generally in inference methods which are extensions of the description provided by first order logic. Circumscription [9], default logic [10] and probabilistic inference schemes such as that discussed in [7] are examples.

Research in truth maintenance systems [4] has involved recording information concerning not only the truth or falsity of a given conclusion, but also justifications for that truth or falsity. This is useful in providing explanations, and also in the revision of inferences drawn using non-monotonic inference rules. Assumption-based truth maintenance systems [3] provide an interesting extension of this idea, taking the truth value of a given proposition to be the set of contexts in which it will hold.

My intention in this paper is to show that these different approaches can be subsumed under a uniform framework. Hopefully, such a framework will lead to a greater understanding of the natures of the individual approaches. In addition, an implementation of the general approach should facilitate the implementation of any of the individual approaches mentioned earlier, in addition to combinations of them (such as probabilistic truth maintenance systems) or new ones yet to be devised.

The ideas presented in this paper should not be taken as supporting any specific multi-valued logic, but as supporting a multi-valued approach to inference generally. The specific logic selected in any given application can be expected to depend upon the domain being explored.

2 A motivating example

Let me motivate the approach I am proposing with a rather tired example. Suppose that Tweety is a bird, and that birds fly by default.

*This work supported by the Office of Naval Research.

Any of the standard formalizations of default reasoning (such as [9] or [10]) will allow us to conclude that Tweety can fly; suppose that we do so, adding this conclusion to our knowledge base. Only now do we learn that Tweety is in fact a penguin.

The difficulty is that this new fact is in contradiction with the information just added to our knowledge base. Having incorporated the fact the Tweety can fly into this knowledge base, we are unable to withdraw it gracefully.

Truth maintenance systems [4] provide a way around this difficulty. The idea is to mark a statement not as merely "true" or "false", but as true or false *for a reason.* Thus Tweety's flying may depend on Tweety's being a penguin *not* being in the knowledge base; having recorded this, it is straightforward to adjust our knowledge base to record the consequences of the new information.

The truth maintenance approach, however, provides us with a great deal more power than is needed to solve this particular problem. We drew a default conclusion which was subsequently overturned by the arrival of new information. Surely we should be able to deal with this without recording the justification for the inference involved; it should be necessary merely to record the fact that the conclusion never achieved more than default status.

In this particular example, we would like to be able to label the conclusion that Tweety can fly not as true, but as true by default. The default value explicitly admits to the possibility of new information overturning the tentative conclusion it represents.

3 Truth values

3.1 Lattices

This approach is not a new one. There is an extensive literature discussing the ramifications of choosing the truth value assigned to a given statement from a continuum of possibilities instead of simply the two-point set $\{t, f\}$. Typical examples are a suggestion of Scott's in 1982 [12] and one of Sandewall's in 1985 [11].

Scott notices that we can partially order statements by their truth or falsity, and looks at this as corresponding to an assignment to these statements of truth values chosen from some set L which is partially ordered by some relation $\leq_t$ (the reason for the subscript will be apparent shortly).

He goes on to note that if we can associate to the partial order $\leq_t$ greatest lower bound and least upper bound operations, the set L is what is known to mathematicians as a *lattice* [8]. Essentially, a lattice is a triple $\{L, \wedge, \vee\}$ where $\wedge$ and $\vee$ are binary operations from $L \times L$ to L which are *idempotent, commutative and associative*:

$$a \wedge a = a \vee a = a$$

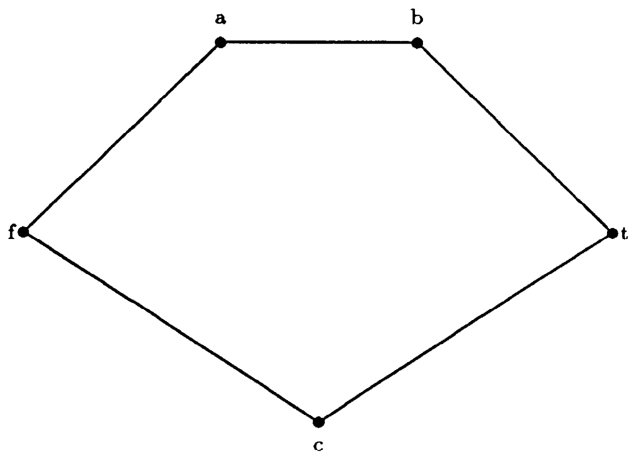

Figure 1: A lattice

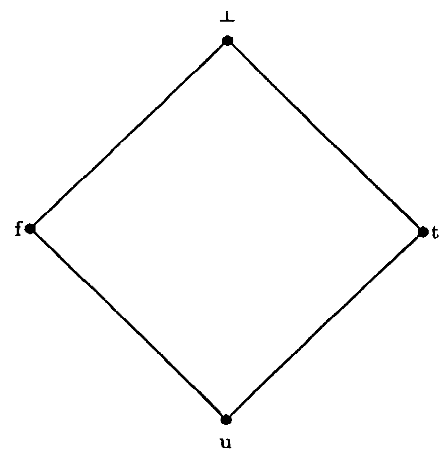

Figure 3: The smallest non-trivial bilattice

Figure 2: The two-point lattice

$$a \wedge b = b \wedge a; \quad a \vee b = b \vee a$$

$$(a \wedge b) \wedge c = a \wedge (b \wedge c); \quad (a \vee b) \vee c = a \vee (b \vee c).$$

In terms of the partial order mentioned earlier, we have $a \wedge b = \text{glb}(a,b)$ and $a \vee b = \text{lub}(a,b)$. $\wedge$ is called the *meet* operation of the lattice; $\vee$ is called the *join*. We also require that if $a \leq b$, then $a \wedge b = \text{glb}(a,b) = a$ and $a \vee b = \text{lub}(a,b) = b$. This is captured by the *absorption identities*:

$$a \wedge (a \vee b) = a; \quad a \vee (a \wedge b) = a.$$

Lattices can be represented graphically. Given such a representation, we will take the view that $p \leq_t q$ if a path can be drawn on the graph from p to q which moves uniformly from left to right on the page.

In the lattice in figure 1, f is the minimal element of the lattice, and t is the maximal element. We also have $a \leq_t b$; a and c are incomparable since there is no unidirectional path connecting them.

Up to isomorphism, there is a unique two-point lattice, shown in figure 2. The truth values in first order logic are chosen from this lattice; all we are saying here is that $f \leq_t t$; "true" is more true than "false".

3.2 Uncertainty

Sandewall's proposal, although also based on lattices, is a different one. Instead of ordering truth values based on truth or falsity, he orders them based on the completeness of the information they represent. Specifically, Sandewall suggests that the truth values be subsets of the unit interval $[0,1]$, the truth value indicating that the probability of the statement in question is known to lie somewhere in the associated probability interval. This proposal also appears in [7] and [5].

The lattice operation used is that of set inclusion. Thus true, corresponding to the singleton set $\{1\}$, is *incomparable* to false, which corresponds to the singleton $\{0\}$. (And each is in turn incomparable with any other point probability, such as $\{0.4\}$.) Instead, the inclusion of one truth value in another relates to our acquiring more information about the statement in question. The minimal element of the lattice is the full unit interval $[0,1]$; the fact that the probability of some statement lies in this interval contains no real information at all.

This is in sharp contrast with knowing, for example, that the probability of the statement in question is .5. If the probability of a coin's coming up heads is .5, the coin is fair; if nothing is known about the probability, it may well not be.

It is clear that the partial order corresponding to Sandewall's notion is conceptually separate from that in Scott's construction. To capture it, we introduce a second partial order $\leq_k$ onto our lattice of truth values, interpreting $p \leq_k q$ to mean loosely that the evidence underlying an assignment of the truth value p is subsumed by the evidence underlying an assignment of the truth value q. Informally, more is known about a statement whose truth value is q than is known about one whose truth value is p.

Since f and t in the two-point lattice corresponding to first order logic should be incomparable with respect to this second partial order, there is no way to introduce this second lattice structure onto the lattice in figure 2. Instead, we need to introduce two additional truth values $\text{glb}_k(t,f)$ and $\text{lub}_k(t,f)$, as shown in figure 3. Just as $p \leq_t q$ if p is to the left of q in a graphical representation, we will adopt the convention that $p \leq_k q$ if p is *below* q on the page.

The two new values are given by u (unknown) and $\perp$ (contradictory). The latter indicates a truth value subsuming both true *and* false; this truth value will be assigned to a given statement just in case it is possible to prove it true using one method and false using another.

We will denote the two lattice operations corresponding to $\leq_k$ by $+$ (lub_k) and $\cdot$ (glb_k) respectively. In general, we define a *bilattice* to be a quintuple $(B, \wedge, \vee, \cdot, +)$ such that:

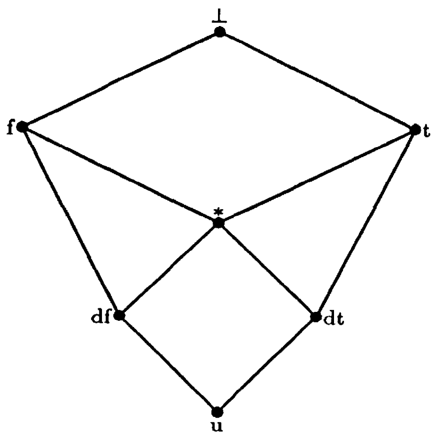

Figure 4: D, the bilattice for default logic

1. $(B, \wedge, \vee)$ and $(B, \cdot, +)$ are both lattices, and
2. Each operation respects the lattice relations in the alternate lattice. For example, we require that if $p \leq_k q$ and $r \leq_k s$, then $p \wedge r \leq_k q \wedge s$. Equivalently, $\wedge$ must be a lattice homomorphism from the product lattice $(B \times B, \cdot, +)$ into $(B, \cdot, +)$ (and similarly for $\vee$, $\cdot$ and $+$).

Just as figure 2 depicts the smallest non-trivial lattice, figure 3 depicts the smallest bilattice which is non-trivial in each lattice direction. Belnap [1,2] has considered the possibility of selecting truth values from this bilattice.

Another bilattice is shown in figure 4; this is the bilattice of truth values in default logic. In addition to the old values of t, f, u and $\perp$, a sentence can also be labelled as dt (true by default) or df (false by default). The additional value $* = dt + df$ labels statements which are both true *and* false by default. This is of course distinct from u (indicating that no information at all is available) or $\perp$ (indicating the presence of a proven contradiction). We will discuss this bilattice in greater detail in a subsequent section.

Before proceeding, however, note that this bilattice shares the elements t, f, $\perp$ and u with the previous one. In fact, *any* bilattice will have four distinguished elements, corresponding to the maximal and minimal elements under the two partial orders. We will denote these distinguished elements in this fashion throughout the paper.

4 Logical operations

In order to apply these ideas, it is insufficient merely to give a framework in which to describe the truth values associated to the sentences of our language. We must also be able to perform inference using these truth values. We now turn to the issue of describing logical operations in a multi-valued setting.

4.1 Extensions and logical connectives

Let L be the set of all well-formed formulae in our language. We will define a *truth function* to be any mapping

$$\phi : L \rightarrow B,$$

corresponding to an assignment of some truth value in the bilattice B to each formula in L.

In first order logic, consistency is defined for truth functions ϕ that are models, so that for each well-formed formula p, $\phi(p) = t$ or $\phi(p) = f$. We will continue to use this definition in the case of multi-valued logics, calling ϕ a *model* if ϕ maps L into the two-point set $\{t, f\}$.

If ϕ and ψ are two truth functions with $\phi(p) \leq_k \psi(p)$ for all $p \in L$, we will write $\phi \leq_k \psi$ and say that ψ is an *extension* of ϕ. If the inequality is strict for at least one $p \in L$, we will write $\phi <_k \psi$ and say that the extension is *proper*. If ψ is a model, we will say that it is a *complete* extension of ϕ.

Informally, an extension of a truth function is what is obtained upon the acquisition of more information about some sentence or sentences in L. The extension will be proper if and only if the new information was not already implicit in the existing truth values.

The usual logical operators of negation, conjunction, disjunction and implication can be described in terms of natural operations on the bilattice structure of our truth values. Conjunction and disjunction are the most easily described, since they are essentially captured by the lattice operators $\wedge$ and $\vee$. In order for a model to be consistent, we therefore require:

$$\phi(p \wedge q) = \phi(p) \wedge \phi(q) \tag{1}$$

$$\phi(p \vee q) = \phi(p) \vee \phi(q). \tag{2}$$

Negation is rather different. Clearly we want to have in general that $\phi(\neg p) \leq_t \phi(\neg q)$ if and only if $\phi(p) \geq_t \phi(q)$. Somewhat less transparent is that we should have $\phi(\neg p) \leq_k \phi(\neg q)$ if and only if $\phi(p) \leq_k \phi(q)$: if we know less about p than about q, we also know less about the negation of p than about that of q. Additionally, we require that $\phi(\neg\neg p) = \phi(p)$.

This leads us to define negation in terms of a map $\neg$ from B to itself such that:

1. $\neg$ is a bilattice isomorphism from $(B, \wedge, \vee, \cdot, +)$ to $(B, \vee, \wedge, \cdot, +)$, and
2. $\neg^2 = 1$.

Note that in the first condition, we have reversed the order of $\wedge$ and $\vee$ between the two bilattices while retaining the order of $\cdot$ and $+$. This corresponds to the observations of the previous paragraph.

For a model, we require:

$$\phi(\neg p) = \neg\phi(p). \tag{3}$$

We handle implication by retaining the usual identification

$$(p \rightarrow q) \equiv (\neg p \vee q).$$

This gives us

$$\phi(p \rightarrow q) = \neg\phi(p) \vee \phi(q). \tag{4}$$

We deal with quantification by noting that $\forall x.p \rightarrow p^t_x$, where t is substitutable for x in p and p^t_x is the result of replacing some (but not necessarily all) of the occurrences of x in p with t. This leads us to assume:

$$\phi(\forall x.p) = \text{glb}_t\{\phi(p^t_x) | t \text{ is substitutable for } x \text{ in } p\}. \tag{5}$$

The existential operator is similar:

$$\phi(\exists x.p) = \text{lub}_t\{\phi(p^t_x) | t \text{ is substitutable for } x \text{ in } p\}. \tag{6}$$

In general, we will call a truth function ϕ *consistent* if it has a complete extension satisfying (1)–(6).

Here are some predicate calculus examples:

x	$\phi(x)$	$\psi(x)$	$\phi(x)$	$\psi(x)$	$\phi(x)$	$\psi(x)$
A	t	t	f	f	f	f
B	f	f	u	t (or f)	u	?
$A \wedge B$	f	f	u	f	t	t

In the first two cases, ψ is a consistent complete extension of ϕ. Since ϕ in the third case has no consistent complete extension, it is itself inconsistent.

Suppose that ϕ is consistent, and let $\{\psi_i\}$ be the set of its consistent complete extensions. We define $\overline{\phi}$ to be the greatest lower bound of the ψ_i:

$$\overline{\phi} = \text{glb}_k\{\psi|\psi \text{ is a consistent complete extension of } \phi\}.$$

In the following two examples, ϕ has two consistent complete extensions given by ϕ_1 and ϕ_2, and $\overline{\phi}$ is the greatest lower bound of these.

	$\phi(x)$	$\phi_1(x)$	$\phi_2(x)$	$\overline{\phi}(x)$
A	t	t	t	t
B	u	t	f	u
$A \vee B$	u	t	t	t

	$\phi(x)$	$\phi_1(x)$	$\phi_2(x)$	$\overline{\phi}(x)$
A	u	t	f	u
$\neg A$	u	f	t	u
$A \vee \neg A$	u	t	t	t

The above construction is closely related to the usual notion of logical inference. In fact, if we denote by ϕ_p the truth function given by

$$\phi_p(q) = \begin{cases} t, & \text{if } q = p; \\ u, & \text{otherwise,} \end{cases}$$

we have:

Theorem 4.1 *$p \models q$ if and only if $\overline{\phi_p} \geq_k \phi_q$.*

Proof. All proofs can be found in [6].

If p is consistent, ϕ_p is the k-minimal truth function in which p is true; the point of the theorem is that q will be true in $\overline{\phi_p}$ if and only if $p \models q$.

4.2 Closure

It might seem that $\overline{\phi}$ is a natural choice for the closure of a truth function in general, but it suffers the drawback of having $\overline{\phi}(p) \geq_k \text{glb}_k\{t, f\}$ for all p. As our bilattice of truth values becomes more complex, such a closure will be insensitive to some of the information contained in ϕ. In the default bilattice D, for example, we have $\overline{\phi}(p) \geq_k *$. Contrast this with theorem 5.1, where the closure of ϕ can also take the values dt, df or u.

The general construction is somewhat more involved; the reader is referred to [6] for details. If we denote the closure of some truth function ϕ by cl(ϕ), the key features of the construction are the following:

1. It can be described completely in terms of the bilattice structure of the truth values.
2. Logical inference always "adds" information to a truth function, so that $\phi \leq_k \text{cl}(\phi)$ in all cases.
3. The construction is non-monotonic, so that it is possible to have $\phi <_k \psi$ without $\text{cl}(\phi) \leq_k \text{cl}(\psi)$. An example of this is given in the next section.

The final remark above refers only to a portion of what is generally referred to as "non-monotonic" behavior. Consider a truth function with $\phi(p) = dt$ but $\text{cl}(\phi)(p) = f$, for example; here inference is behaving "non-monotonically" in the sense that $\phi(p) >_t u$ but $\text{cl}(\phi)(p) <_t u$. It is behaving monotonically, however, in that $\phi(p) <_k \text{cl}(\phi)(p)$. It turns out that the computational difficulties which plague non-monotonic inference systems arise principally as a result of the potential non-monotonicity in the k sense; loosely speaking, k-monotonicity is enough to guarantee that we can maintain our knowledge base using updates. There are therefore substantial practical advantages to be gained by recognizing situations where it can be demonstrated that the closure operation is k-monotonic. Details are in [6].

5 Examples

Let me end by very briefly describing Reiter's default reasoning and truth maintenance in terms of this sort of construction. The second of these is extremely straightforward, essentially requiring us merely to identify those statements that support some fixed one. Default reasoning is a bit more intricate, since the philosophy underlying Reiter's approach is very different from that of the one we have been presenting.

5.1 Default logic

Reiter defines default reasoning in terms of a *default theory* (R, T) where T is a collection of first order sentences, and R is a collection of *defaults*, each of the form

$$\frac{\alpha : \beta_1, \ldots, \beta_m}{w},$$

indicating that if α holds and all of the $\beta's$ are possible, then w holds. If every default rule is of the form

$$\frac{\alpha : w}{w},$$

so that we infer w from α in the absence of information to the contrary, the default theory is called *normal.*

Reiter goes on to define an *extension* of a default theory (R, T), and shows that these extensions correspond to the collections of facts derivable from such a theory. Since there may be conflicting default rules, it is possible that a given default theory have more than one extension.

The bilattice for default logic appeared in figure 4.

Theorem 5.1 *Let (R, T) be a normal default theory, with default rules*

$$\frac{\alpha_i : w_i}{w_i},$$

where i indexes the elements in R. Associate to (R, T) a truth function ϕ given by

$$\phi(p) = \begin{cases} t, & \text{if } p \in T; \\ dt, & \text{if } p \text{ is } \alpha_i \rightarrow w_i \text{ for some } i \text{ but } p \notin T; \\ u, & \text{otherwise.} \end{cases}$$

Then:

$$\text{cl}(\phi)(p) = \begin{cases} t, & \text{iff } T \models p; \\ f, & \text{iff } T \models \neg p; \\ *, & \text{iff } p \text{ is true in some extensions of } (R,T) \text{ and false in others;} \\ dt, & \text{iff } T \not\models p \text{ but } p \text{ is true in some of the extensions and false in none;} \\ df, & \text{iff } T \not\models \neg p \text{ but } p \text{ is false in some of the extensions and true in none;} \\ u, & \text{iff } p \text{ is undecided in all extensions of } (R,T). \end{cases}$$

Theorem 5.2 *The closure operation is potentially non-monotonic.*

Suppose that we have two default rules, one of which indicates that birds can fly, and the other that flying things are stupid. Consider the following two truth functions:

p	$\phi(p)$	$\psi(p)$
bird(Tweety)	t	t
penguin(Tweety)	u	t
flies(Tweety)	u	u
dumb(Tweety)	u	u

Clearly $\phi <_k \psi$. But in light of the previous theorem, the closures of ϕ and ψ are given by:

p	$\phi(p)$	$\psi(p)$
bird(Tweety)	t	t
penguin(Tweety)	df	t
flies(Tweety)	dt	f
dumb(Tweety)	dt	u

We do not have $\text{cl}(\phi) \leq_k \text{cl}(\psi)$. The point is that the fact that Tweety is now known not to fly keeps the default rule about stupidity from firing. □

5.2 Truth maintenance

In a truth maintenance system, the truth values assigned to propositions contain information concerning the reasons for their truth or falsity. We can capture this using a multi-valued logic in which the truth values consist of pairs $[\ a\ .\ b\ ]$ where a and b are respectively justifications for the truth and falsity of the statement in question. We can assume that these justifications are themselves in disjunctive normal form, consisting of a list of parallel conjunctive justifications.

An example will make this clearer. Suppose that p is the statement $q \vee (r \wedge s)$. Then if q, r and s are all in the knowledge base, the truth value of p will be

$$[\ (\{q\}\{r,s\})\ .\ \text{nil}\ ]$$

Either q or the $\{r,s\}$ pair provides independent justification for p, and there is no justification for $\neg p$. We will assume in general that if the truth value of p is $[\ a\ .\ b\ ]$, either a or b is empty; in other words, that either p or $\neg p$ is unjustified.

Given two justifications j_1 and j_2 expressed in disjunctive normal form, we write $j_1 \leq j_2$ if every conjunctive subclause in j_1 contains some subclause in j_2 as a subset:

$$(a_1 \ldots a_n) \leq (b_1 \ldots b_m)$$

if for each a_i, there is some b_j with $b_j \subseteq a_i$. It is not hard to see that the empty justification (containing no information) is a minimal element under this partial order, while the justification $(\{\})$ consisting of a single empty conjunct (a justification needing no premises) is maximal.

If $j_1 \leq j_2$, we now define:

$$[\ j_1\ .\ \text{nil}\ ] \leq_t [\ j_2\ .\ \text{nil}\ ] \qquad [\ j_1\ .\ \text{nil}\ ] \leq_k [\ j_2\ .\ \text{nil}\ ] \qquad (7)$$

and

$$[\ \text{nil}\ .\ j_1\ ] \geq_t [\ \text{nil}\ .\ j_2\ ] \qquad [\ \text{nil}\ .\ j_1\ ] \leq_k [\ \text{nil}\ .\ j_2\ ] \qquad (8)$$

The k-join $t \cdot f$ is of course $\perp$ as usual. Note the sense of the first inequality in (8).

The analog to theorem 5.1 is now:

Theorem 5.3 *Let $p_1,\ldots,p_n$ be possible assumptions in our knowledge base; and suppose that ϕ is given by*

$$\phi(q) = \begin{cases} [\ (\{p_i\})\ .\ \text{nil}\ ], & \text{if } q = p_i \text{ for some } i; \\ u, & \text{otherwise.} \end{cases}$$

Then if $q_1,\ldots,q_m$ form a subset of the p_i's and x is an arbitrary sentence,

$$[\ (\{q_1,\ldots,q_m\})\ .\ \text{nil}\ ] \leq_k \text{cl}(\phi)(x)$$

if and only if the q_j's form a justification of x.

6 Future work

It is painfully clear that the work presented in this paper only scratches the surface of the approach being discussed.

Both theoretical and engineering issues need to be explored. There are many other non-standard approaches to inference; can they be captured in this framework? Circumscription and probabilistic schemes seem especially important candidates.

Equally important is an implementation of the ideas we have discussed. Ideally, a general-purpose inference engine can be constructed which accepts as input four functions giving the two glb and two lub operations in the bilattice, and which then performs suitable multi-valued inference. The key issue is the determination of what price must be paid in terms of efficiency for the increased generality of the approach we are proposing.

Work in each of these areas is currently under way at Stanford.

REFERENCES

[1] N. D. Belnap. How a computer should think. In *Proceedings of the Oxford International Symposium on Contemporary Aspects of Philosophy*, pages 30–56, 1975.

[2] N. D. Belnap. A useful four-valued logic. In G. Epstein and J. Dumm, editors, *Modern Uses of Multiple-valued Logic*, pages 8–37, D. Reidel Publishing Company, Boston, 1977.

[3] J. de Kleer. An assumption-based truth maintenance system. *Artificial Intelligence*, 28:127–162, 1986.

[4] J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12:231–272, 1979.

[5] M. L. Ginsberg. *Analyzing Incomplete Information.* Technical Report 84-17, KSL, Stanford University, 1984.

[6] M. L. Ginsberg. *Multi-valued Logics.* Technical Report 86-29, KSL, Stanford University, 1986.

[7] M. L. Ginsberg. Non-monotonic reasoning using Dempster's rule. In *Proceedings of the American Association for Artificial Intelligence*, pages 126–129, 1984.

[8] G. Grätzer. *General Lattice Theory.* Birkhäuser Verlag, Basel, 1978.

[9] J. McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 89–116, 1986.

[10] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[11] E. Sandewall. A functional approach to non-monotonic logic. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 100–106, 1985.

[12] D. S. Scott. Some ordered sets in computer science. In I. Rival, editor, *Ordered Sets*, pages 677–718, D. Reidel Publishing Company, Boston, 1982.