OU

A functional approach to non-monotonic logic[1]

Erik Sandewall
Department of Computer and Information Science, Linköping University, Linköping, Sweden

Received May 3, 1985
Revision accepted August 1, 1985

Axiom sets and their extensions are viewed as functions from the set of formulas in the language to a set of four truth values, t, f, u for undefined, and k for contradiction. Such functions form a lattice with "contains less information" as the partial order $\sqsubseteq$, and "combination of several sources of knowledge" as the least-upper-bound operation $\sqcup$. Inference rules are expressed as binary relations between such functions. We show that the usual criterium on fixpoints, namely, to be minimal, does not apply correctly in the case of non-monotonic inference rules. A stronger concept, approachable fixpoints, is introduced and proven to be sufficient for the existence of a derivation of the fixpoint. In addition, the usefulness of our approach is demonstrated by concise proofs for some previously known results about normal default rules.

Les ensembles d'axiomes et leurs extensions sont considérés comme des applications de l'ensemble des formules du langage, vers un ensemble de quatre valeurs de vérité t, f, u pour l'indéfini, et k pour la contradiction. Ces applications forment un ensemble ordonné avec "contient moins d'information" symbolisé par l'ordre partiel $\sqsubseteq$, et "combinaison de différentes sources de connaissance" symbolisé par l'opération d'union $\sqcup$. Les règles d'inférence sont exprimées par des relations binaires entre ces fonctions. Nous montrons que le critère habituel pour les invariants, c'est-à-dire qu'ils sont minimaux, ne s'applique pas correctement dans le cas des règles d'inférence non-monotones. Un concept plus fort, compatible avec les invariants, est introduit et l'on prouve qu'il est suffisant pour qu'existe une dérivation de l'invariant. De plus, la pertinence de notre approche est démontrée grâce à des démonstrations concises venant de résultats déjà connus à propos des règles de défaut normales.

[Traduit par la revue]

Comput. Intell. 1, 80–87 (1985)

1. Introduction and overview

Non-monotonic logic may be studied in terms of either non-monotonic inference rules (Reiter 1980; Goodwin 1984) or non-monotonic operators in the language such as the *Unless* operator (Sandewall 1972; McDermott and Doyle 1980). In this paper we pursue the former approach.

The concept of fixpoints is central to the study of non-monotonic logic: For a given set v of propositions and a given set of rules, we are looking for an extension, i.e., a set v' of propositions which contains v as a subset, and which is a fixpoint of the set of rules. Fixpoints are also used in the denotational semantics approach to the theory of programming languages (Scott 1970; see also e.g., Manna 1974; Stoy 1977; Blikle 1981). There, the recursive definition of a function is viewed as a functional, i.e., an operator on partial functions, and the function is viewed as the fixpoint of the same functional.

In this paper we propose that the functional approach that is taken in denotational semantics can be adapted and serves conveniently for the study of non-monotonic logic. This is attractive since logical inference is often viewed as a high-level form of computation, and since computational inference often needs to be non-monotonic.

Using the functional approach, we present a result regarding criteria on the desirable fixpoints in the case of non-monotonic inference rules. In the monotonic case, the criterium of being *a minimal* fixpoint (i.e., no other fixpoint is "smaller") is sufficient for eliminating fixpoints that contain spurious information not warranted by the given facts and inference rules. Also, there is of course just one minimal fixpoint, which is then *the least* fixpoint. It is well known that in the case of non-monotonic rules, there is in general no single least fixpoint. But in addition, the criterium of being minimal is not sufficient for eliminating spurious fixpoints. There may be minimal fixpoints which have the given set of propositions as a subset, but which still cannot be reached or approached (in the sense of a limit) by any derivation using the given set of rules. We introduce the concept of an *approachable* fixpoint, which is stronger than the concept of a minimal fixpoint, and which is proven to be a sufficient condition for the existence of a derivation that reaches or approaches the fixpoint.

Besides this result, the usefulness of the functional approach is also demonstrated through greatly simplified proofs of some of Reiter's (1980) results for normal default theories.

The following formal machinery is used. We start from two domains: a domain L whose elements are called *formulas* and a domain J of *truth values*. V is the domain of *valuations*, i.e., continuous functions from L to J.

A set of axioms is seen as a valuation that maps some formulas (the axioms) to t (for true) and "all" other formulas to u (for undefined). (Exception is made for the top element of the domain L.)

A valuation may be thought of as a partial information state. If valuations were used in a treatment of knowledge and belief, then the state of mind of one agent could be appropriately represented as a valuation, since the agent may believe some propositions to be true and some others to be false and does not know about yet others. One such information state "contains less information" than another one, if the only way they can differ is that the former state assigns undefined to a proposition whereas the latter state assigns true or false. That relation of "containing less information" is clearly a partial order and will be represented by the symbol $\sqsubseteq$.

Logical deduction may then be seen as a process that goes from an initial information state to successive other states that contain "more" information in the sense of $\sqsubseteq$. A set of inference rules corresponds therefore to a binary relation of valuations, i.e., a subset of $V \times V$, which will be called inference

[1]This research was sponsored by the Swedish Board of Technical Development.

relation. An inference relation holds between two valuations v and v' if $v \sqsubseteq v'$ and the augmentation of information from v to v' results from a one-step application of an inference rule. A *derivation* using an inference relation F is a sequence of valuations,

$$v_0, v_1, \ldots$$

where

$$\langle v_i, v_{i+1} \rangle \in F$$

for each $i \geq 0$.

Inference relations can be used for expressing both proofs and semantics, provided that there are syntactic functions and predicates on L which characterize the abstract syntax of the language. This includes predicates which indicate whether a formula is a conjunction, a disjunction, an implication, atomic, etc., as well as functions, e.g., for composing the conjunction of two other formulas. The conventions for calculating the truth value of a propositional expression may then be seen as an inference relation F where $F(v, v')$, e.g., in the case where

$$v(a) = t$$

$$v(b) = t$$

$$v(a \wedge b) = u$$

$$v'(a \wedge b) = t$$

$$v'(x) = v(x) \qquad \text{for all other formulas } x$$

Here $a \wedge b$ refers of course to the formula obtained by composing the formula a, the conjunction operator, and the formula b.

The concepts and results of conventional logic can easily be rephrased along these lines, allowing a uniform treatment of proofs and models. In the case of non-monotonic logic, there is however a particular advantage with using this approach. A non-monotonic rule

$$a, \textit{Unless } (b) \Rightarrow c$$

can now be seen as an inference relation F which allows $F(v, v')$, e.g., in the case where

$$v(a) = t$$

$$v(b) = u$$

$$v(c) = u$$

$$v'(b) = f$$

$$v'(c) = t$$

$$v'(x) = v(x) \qquad \text{otherwise}$$

In other words, one derivation step using F will change the truth values of the two formulas b and c at the same time. This is different from the viewpoint in ordinary logic, where the intuition is that each formula or proposition has "its" truth value, so that rules of inference may contribute additional information about "the" truth of a proposition. In non-monotonic logic, we must be prepared to recognize multiple extensions of the given axioms, or multiple fixpoints of an inference relation. It therefore makes sense to correlate assignments of truth values in the way just described.

By the usual notions of non-monotonic logic, we should be

Instead of the usual type of definition such as "a valuation is inconsistent if it assigns the truth value t both to a proposition and to its negation," it is convenient to use the four-valued logic proposed by Belnap (1977). Besides the truth values t, f, and u, we also allow the truth value k for "contradiction." A valuation is consistent if it does not assign the truth value k to any proposition.

One can think of these truth values as the ones assigned by committees (cf. Borgida and Imielinski 1984): If some members of the committee assign the value t to a proposition and the others assign the value u, then the committee assigns t, but if some members assign t and others assign f, then the committee assigns the value k to the proposition. "Committees" are a convenient metaphor whenever there are several parallel sources of knowledge, such as when several inference rules are being used.

Let us now go back to the example of a non-monotonic rule above. The inference relation F should allow $F(v, v')$ also in the case where

$$v(a) = v'(a) = t$$

$$v(b) = u$$

$$v(c) = f$$

$$v'(b) = f$$

$$v'(c) = k$$

True and false add up to a contradiction, as in the case for the committee. In this case, we rightfully obtain a contradiction: c had previously been assumed or proven to be false, and now we derived that it is true. Therefore v' (or one of its successors) may well be a fixpoint of the inference relation, but it will not be a consistent fixpoint.

However, in line with the idea that F should change the truth values of both b and c at the same time, we should also have $F(v, v')$ in the following case:

$$v(a) = v'(a) = t$$

$$v(b) = t$$

$$v(c) = u$$

$$v'(b) = k$$

$$v'(c) = t$$

But in this case, v is really a valuation (or in other words, an information state) where the non-monotonic rule should not be applied at all, since the non-monotonic antecedent, "b is not known to be true," was not satisfied. Formally speaking, in the last example we want v to be approved as a consistent fixpoint for the inference relation, and therefore $F(v, v')$ should not hold for any v' different from v and particularly not for an inconsistent v'. In the case before that, however, we do wish $F(v, v')$ to hold for the inference relation F, so that the contradiction is made explicit.

We resolve this matter by using two related inference relations, usually written F and G, and it is the consistent fixpoints of G that we are interested in. Thus in the last case of the example, $G(v, v')$ would not apply, whereas $F(v, v')$ applies. The inference relations G and F are in general related so that $G \subseteq F$, and the difference occurs in cases like the one just discussed, where G rejects using a rule if it means that a non-monotonic antecedent (a formula given as the argument to *Unless*) then obtains the truth value k.

The reason for using the inference relation F at all is that it has simpler mathematical properties, so it is convenient to prove results for F first, and then to transfer them in a "guarded" way to the relation G.

By using inference relations in this way, we can abstract away the detailed mechanics of applying inference rules. We identify the properties of inference relations if they have been constructed from ordinary inference rules: monotonicity, compactness, and so on. We then carry out the proofs in terms of those more abstract properties.

The present paper only reports results on fixpoints for inference relations that express non-monotonic inference rules. However, we believe that the notions of valuations and inference relations may also be useful in other ways. The possibility of using valuations as arguments for operators expressing knowledge and belief has already been touched upon, and of course operators expressing information transfer ("tells that," "asks whether") belong to the same general category.

Also, if a valuation is seen as an information state during a deductive process, then the choice of a deduction strategy may be seen as a transformation from a given inference rleation to another related one, for example a transformation from F to a subset of F or to a subset of F^*. The relationship between the original and the derived inference relation may be characterized either by using logic on the meta level or by a relational algebra.

This has been an outline of the key ideas of the paper. We now proceed to the systematic treatment.

2. Monotonicity and fixpoints for relations

The domain J contains the elements $\{u, t, f, k\}$ as already discussed, with the partial order $\sqsubseteq$ described by the following figure:

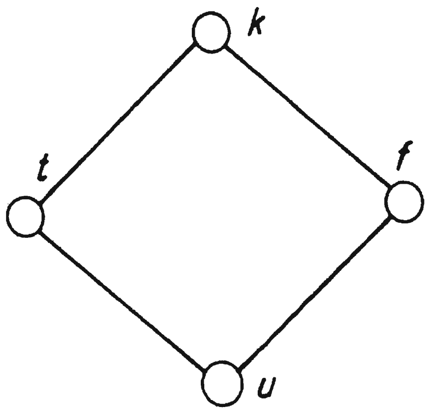

(The reader who is not familiar with Scott's domain theory may skip the next two paragraphs and proceed using the intuitions introduced in the previous section.)

The domain J is therefore a flat lattice with u as the bottom element and k as the top element. The flat domain L is called a *language* and its elements are called *formulas*, with lb as the bottom element and lt as the top element.

A *valuation* is a function v from L to J which satisfies

$$v(lb) = u$$

$$v(lt) = k$$

which in particular guarantees that valuations are continuous (and therefore monotone) functions. A valuation is *consistent* iff no formula other than lt is mapped to k.

Valuations form a lattice with the partial order $\sqsubseteq$ defined in the following way:

$$v \sqsubseteq v' \text{ iff } (\forall x) v(x) \sqsubseteq v'(x)$$

We say then that v' is an *extension* of v, and that v' is *above* v (with respect to $\sqsupseteq$). A valuation v is *finite* iff $v(l) = u$ except for a finite number of formulas l.

An *inference relation* is a binary relation on V, i.e., a subset of $V \times V$. The operation $\cup$ and the relation $\subseteq$ are therefore defined on inference relations. The partial order $\sqsubseteq$ could also be extended to inference relations, but that will not be needed in this paper.

An inference relation F is *conservative* iff

$$F(v, v') \rightarrow v \sqsubseteq v'$$

which can now be written

$$F \subseteq \sqsubseteq$$

A *chain* is a sequence of valuations where

$$v_0 \sqsubseteq v_1 \sqsubseteq \ldots$$

It is well known that each such chain has a least upper bound (l.u.b.) in a lattice. A *derivation* from v to v' using a conservative inference relation F is a chain where

$$v = v_0$$

$$F(v_i, v_{i+1}) \quad \text{for all } i \geq 0$$

v' is the l.u.b. of the chain

Notice that we say a derivation — to v' — even in the case of an infinite sequence where v' is never reached, just approached as a limit.

An inference relation F is *monotonic* iff

$$v \sqsubseteq v' \wedge F(v, y) \rightarrow (\exists y')(F(v', y') \wedge y \sqsubseteq y')$$

It is *linear* iff

$$F(v, y) \rightarrow F(v \sqcup z, y \sqcup z)$$

Clearly every linear inference relation is also monotonic.

We need an existential quantifier rather than a universal quantifier in the definition of monotonicity, since for the given v there may be several y such that $F(v, y)$, each of which expressing one of the conclusions that can be immediately added from v. If now v' is another valuation which contains "more" information than v, we wish that each of the steps that can be taken from v to one of the y should have a counterpart from v' to a correponding y'.

A valuation v' is a *fixpoint* of an inference relation F iff

$$F(v', v'') \rightarrow v' = v''$$

Thus in particular v' is a fixpoint if F does not allow any "successor" v''. For a given valuation v and inference relation F, we shall be interested in fixpoints of F above v, i.e., fixpoints of F which are $\sqsupseteq v$.

A fixpoint of F above v is *minimal* iff no "smaller" (by $\sqsubseteq$) fixpoint exists for the same F and v.

What has been described so far uses some of the tools of denotational semantics, but in a different fashion than usual. The differences are dictated by our desire to deal with logic using these tools. The reason for that, again is the wish to consider non-monotonic deduction. Anyway, the obvious properties of the monotonic case follow easily, in particular:

Proposition 1. For a conservative, monotonic inference relation, (a) there is a unique least fixpoint above each v; and (b) the l.u.b. of any derivation from v is $\sqsubseteq$ the least fixpoint.

Proof. Let the inference relation be called F.
(a) Consider a set $\{v_i\}$ of fixpoints of F above v. We wish to prove that $v' = \sqcap\, v_i$ is also a fixpoint of F. Suppose

$$F(v', z)$$

By monotonicity, for each v_i there exists a valuation z_i such that

$$F(v_i, z_i)$$

$$z \sqsubseteq z_i$$

Since each v_i is a fixpoint, we have

$$v_i = z_i$$

and therefore

$$z \sqsubseteq v' (= \sqcap\, z_i)$$

and since F is conservative,

$$v' = z \qquad []$$

(b) Let $v_0, v_1, \ldots$ be a derivation using F, and let v^* be the least fixpoint of F above v_0. By induction, it is easily proven that

$$v_i \sqsubseteq v^*$$

for each i. []

The following concept is also of interest in the non-monotonic case: A valuation v is *maximally consistent* w.r.t. a deduction F iff it is consistent and

$$F(v, v') \rightarrow v = v' \vee [v' \text{ is inconsistent}]$$

3. Linearity and compactness

There are two significant properties, linearity and compactness, for inference relations formed using a number of rules. This section studies consequences of these properties. Linearity was defined in the previous section.

Proposition 2. If F is a conservative, linear inference relation, and there exists a derivation from v to v' using F, then for every y such that

$$v \sqsubseteq y \sqsubseteq v'$$

there exists a derivation from y to v' using F.

Proof. Let the derivation from v to v' be

$$v_0, v_1, \ldots$$

where $v_0 = v$. The sequence

$$v_0 \sqcup y, v_1 \sqcup y, \ldots$$

is a derivation using F, and by the assumptions we have

$$v_0 \sqcup y = y$$

and the l.u.b. of the sequence is v'. []

Proposition 3. If F is a conservative, linear inference relation and there are derivations from v to v' and from v to v'' using F, then there is a derivation from v to $v' \sqcup v''$ using F.

Proof. Let the derivations of v' and v'' be

$$v'_0, v'_1, \ldots$$

$$v''_0, v''_1, \ldots$$

where $v'_0 = v''_0 = v$. Since F is linear, the sequence

$$v'_0, v'_1,$$

$$v'_1 \sqcup v''_1, v'_2 \sqcup v''_1,$$

$$v'_2 \sqcup v''_2, v'_3 \sqcup v''_2, \ldots$$

is a derivation from v using F, and it has $v' \sqcup v''$ as its l.u.b. []

The proof of Proposition 3 generalizes easily:

Proposition 4. If F is a conservative, linear inference relation and there are derivations using F from v to each member of a non-empty (possibly infinite) set W of valuations, then there is a derivation from v to the l.u.b. of W.

Proof (outline). Use the same technique as in the previous proof, but select some ordering of the members of W (presumed denumerable), and construct the new derivation in a triangular fashion, so that it may use the first j terms of the derivation for the first member of W, the first $j-1$ terms of the derivation for the second member, etc. []

The other important property is compactness. An inference relation F is *compact* iff whenever $F(v, v')$ there exist some finite valuations y, y' such that

$$y \sqsubseteq v \text{ (which of course means } v = v \sqcup y)$$

$$v' = v \sqcup y'$$

We shall see in a moment that the inference relations obtained from "ordinary" rules of inference are compact, but let us first identify a consequence of compactness.

Proposition 5. If F is a conservative, linear, and compact inference relation and there are derivations using F from v to v' and from v' to v'', then there is also one from v to v''.

Proof. Let the derivation from v to v' be

$$v_0, v_1, \ldots$$

and let the derivation from v' to v'' be

$$v'_0, v'_1, \ldots$$

where for each $i \geq 0$ there are some finite y_i, z_i such that

$$F(y_i, z_i)$$

$$y_i \sqsubseteq v'_i$$

$$v'_{i+1} = v'_i \sqcup z_i$$

Construct now a sequence of valuations as follows:

$$v_0, v_1, \ldots, v_{n_0},$$

$$v_{n_0} \sqcup z_0, v_{n_0+1} \sqcup z_0, \ldots, v_{n_1} \sqcup z_0,$$

$$v_{n_1} \sqcup z_0 \sqcup z_1, \ldots$$

$$\ldots$$

where the n_i are selected so that

$$n_i \leq n_{i+1}$$

$$y_i \sqsubseteq v_{n_i}$$

for $i \geq 0$. By compactness such a sequence must exist, and by the linearity the sequence is a derivation using F, and its l.u.b. is clearly v''. []

We can now proceed to introducing the counterparts of inference rules in our system.

A *kernel* is a pair $\langle v, v' \rangle$, where v and v' are finite valuations, and $v \sqsubseteq v'$.

Kernels may be used for expressing how the truth value of a composite expression follows from the truth value of its component(s), or vice versa. One such example was given in the introductory section. For another example, the rule that if a is true then $\neg a$ is false, is expressed by the kernel $\langle v, v' \rangle$ where

$$v(a) = v'(a) = t$$

$$v(\neg a) = u$$

$$v'(\neg a) = f$$

The *direct realization* of a kernel $\langle v, v' \rangle$ is the inference relation formed as

$$\{\langle v \sqcup y, v' \sqcup y \rangle \mid y \in V\}$$

In other words, the direct realization of a kernel $\langle v, v' \rangle$ is the set of all possible pairs $\langle y, y' \rangle$ such that $v \sqsubseteq y$ and $y' = y \sqcup v'$. Each such pair characterizes a derivation step that is allowed by the kernel, i.e., if the preconditions v are satisfied in y the conclusions v' may be accumulated to y giving y'.

The direct realization of a set of kernels is defined to be the union (using $\cup$) of the direct realizations of the individual kernels. This has the effect that from each valuation v there are several successors, corresponding to the choices of which derivation step to take. From this definition it follows:

Proposition 6. The direct realization of a set of kernels is conservative, linear, and compact.

4. Non-monotonic rules

We shall now characterize those inference relations which correspond to (what we intuitively think of as) a set of non-monotonic (NM) inference rules.

Following Goodwin (1984) approximately, an NM rule is a triple $\langle M, N, C \rangle$ of finite sets of formulas, where M is the *monotonic antecedents*, N is the *non-monotonic antecedents*, and C is the *consequents*.

The idea is that if each member of M is true and each member of N is false or undefined, then each member of C can be inferred to be true. At the same time, the assumption is made that all members of N are false. In practice, we will usually be interested in infinite sets of NM rules, corresponding to the set of substitution instances of what is maybe intuitively thought of as a single rule.

Each of M, N, and C may be the empty set. If N is empty we have a monotonic rule. If C is empty we have what Reiter '1980) called a normal default rule.

The kernel that *corresponds to* an NM rule $\langle M, N, C \rangle$ is the pair $\langle v, v' \rangle$, where

$$v(m) = v'(m) = t \quad \text{for all } m \text{ in } M$$

$$v'(n) = f \quad \text{for all } n \text{ in } N$$

$$v'(c) = t \quad \text{for all } c \text{ in } C$$

and all other values are u.

Thus non-monotonic rules differ from the monotonic ones, partly by causing *several* formulas to change their truth value as one inference step is performed. A possible objection against this way of dealing with non-monotonic antecedents is that the resulting valuation should differentiate explicitly between that information which has been obtained as a consequent and that which was "merely" assumed in order to be able to apply the rule, i.e., the assignment to the non-monotonic antecedent(s). We, however, view that as a book-keeping issue which need not concern the formal treatment of the inference relation as such.

The direct realization of an NM rule is the direct realization of its corresponding kernel.

We let H be that inference relation which performs the obvious deductions of conventional, propositional logic. For example, if a valuation v satisfies

$$v(a \wedge b) = t$$

then the valuation $H^*(v)$, which is the least fixpoint of H above v (Proposition 1), satisfies

$$H^*(v)(a) = H^*(v)(b) = t$$

(unless a contradiction occurs in which case $H^*(l) = k$ for all l). We work presently on a more specific definition and analysis of this inference relation H.

We want the direct realization of a set of NM rules to be the inference relation which has as subsets the realizations of each of the rules, but which is also able to do trivial derivations of the truth values. We therefore formally define the *direct realization* of a *set* of NM rules as $H \cup$ (the union of the direct realization of each of the rules). Clearly the direct realization of a set of NM rules is linear, conservative, and compact.

The *restricted realization* of an NM rule $\langle M, N, C \rangle$ is a subset of the direct realization of the same rule. It is obtained by excluding all those pairs $\langle v, v' \rangle$ where $v'(n) = k$ for some $n \in N$. Notice that pairs $\langle v, v' \rangle$ where $v'(c) = k$ for some $c \in C$ are not excluded, unless some $v'(n)$ is also k. The restricted realization of a set of NM rules is obtained as the union of H and the restricted realization of each of the NM rules. The restricted realization is conservative and compact but not monotonic (and therefore not linear).

An example may be useful at this point. For the following examples, we assume that the language consists of the formulas $\{a, b, c, \ldots\}$. A valuation will be written as $[x, y, \ldots]$ meaning the valuation v where $v(a) = x$, $v(b) = y$, etc. If either of the sets in a rule is a singleton, then the curly brackets around it will be omitted, and the empty set will be written as a dash. Thus $\langle a, -, c \rangle$ is an example of a rule, meaning the same as $\langle \{a\}, \{\}, \{c\} \rangle$.

Example 1. Suppose we have the following rules:

$$\langle -, a, b \rangle$$

$$\langle b, -, c \rangle$$

Informally, this says: b holds unless a is known to be true; if b then c. The restricted realization G of this set of rules satisfies

$G([u, u, u], [f, t, u])$ (this uses the rule "b unless a")

$G([f, t, u], [f, t, t])$ (this uses the rule "if b then c")

and $[f, t, t]$ is also a fixpoint for G over $[u, u, u]$.

If we start instead from a valuation where a is known to be true, e.g., $v = [t, u, u]$, there is no valuation v' such that $G(v, v')$. The direct realization F of the same set of rules holds of course for the same argument pairs as G, but also

$$F([t, u, u], [k, t, u])$$

since the direct realization will proceed even if it introduces a contradiction. The valuation $[t, u, u]$ is therefore a fixpoint for G, but not for F. However, it is a maximally consistent extension of $[t, u, u]$ w.r.t. F.

5. Correct extensions, approachability

Throughout this section, we assume that v is a valuation and R is a set of NM rules whose direct realization is F and whose restricted realization is G.

A valuation v' is termed a *correct extension* of v w.r.t. R iff:
(1) v' is a fixpoint of G above v (meaning in particular that $v \sqsubseteq v'$)
(2) v' is consistent
(3) there is some derivation from v to v' using G.
(The adjective "correct" is used since in the terminology used here, the phrase "v' is an extension of v" means simply that $v' \sqsupseteq v$.) The notion of correct extensions expresses stringently what are the desirable fixpoints for given v and G. If v' is not a fixpoint then some additional derivation steps remain to be performed. If v' is inconsistent it is for one of the following reasons:
(a) v is inconsistent
(b) the set of NM rules implies an inconsistence (e.g., if $v(a) = t$ and one of the rules is $\langle -, -, \neg a \rangle$)
(c) some non-monotonic antecedent was assumed to be f at the beginning of the derivation, and later in the derivation the assumption was invalidated either by using an NM rule with the same proposition as a consequent or by using the inference relation H.

Finally, if there is no derivation from v to v' then v' is "unfounded," like a nonminimal fixpoint in the monotonic case.

The third condition in the definition corresponds to Goodwin's requirement of well foundedness. The above formulation is however problematic in that it refers to the existence of a derivation. By contrast, the minimality requirement on a fixpoint need not refer to derivations; it just states that no "smaller" fixpoint exists. Similarly, we would like to have a static condition, which guarantees the existence of a derivation, instead of having to prove its existence whenever needed. The remainder of this section will give such a result.

Although the definition of correct extensions uses G, the inference relation F provides a partial characterization of them and will be used as a tool.

Proposition 7. Each consistent fixpoint of G above v is a maximally consistent extension of v w.r.t. F.

Proof. The proof follows easily from the definition of G from F, since

$$F(v, v') \wedge \neg G(v, v')$$

implies that v' is inconsistent.

The converse does not hold, i.e., there are maximally consistent extensions of v w.r.t. F which are not fixpoints of G over v. This happens in those cases where a truth maintenance system has to shift IN nodes to OUT status. Consider

Example 2. Suppose we have the following NM rules:

$$\langle -, a, b \rangle$$

$$\langle -, -, a \rangle$$

Then $[f, t]$ is a maximally consistent extension of $[u, u]$ w.r.t. F, but $G([f, t], [k, t])$.

Not every set of NM rules has a correct extension:

Example 3. Consider the NM rules

$$\langle -, a, b \rangle$$

$$\langle b, -, a \rangle$$

Then of course

$$G([u, u], [f, t])$$
$$G([f, t], [k, t])$$

and $[u, u]$ does not have any correct extension w.r.t. these NM rules.

Some of the propositions in the earlier sections can now be extended to apply to the realizations of NM rules. The basic idea is to first use those results for F, and then to transfer the result to G by introducing a consistency requirement.

Proposition 2A. If v' is a correct extension of v w.r.t. R, then for every y such that

$$v \sqsubseteq y \sqsubseteq v'$$

there exists a derivation from y to v' using G.

Proof. According to Proposition 2 this holds for F. However, it follows from the definition of G that if $F(z, z')$ and z' is consistent, then $G(z, z')$. Since v' is consistent, so must all intermediate steps in the derivation from y be, because G is conservative. Therefore we have a derivation from y to v' using G. []

Corollary ("minimality of extensions" (Reiter 1980)). If v' and v'' are correct extensions of v w.r.t. R, then

$$v' \sqsubseteq v'' \rightarrow v' = v''$$

Proposition 3A. If there are derivations from v to v' and from v to v'' using G, and $v' \sqcup v''$ is consistent, then there is a derivation from v to $v' \sqcup v''$ using G.

Proof. The proof follows directly from Proposition 3.

Corollary. If v' and v'' are distinct correct extensions of v w.r.t. R, then $v' \sqcup v''$ is inconsistent.

This corollary subsumes Reiter's (1980) "orthogonality of extension" theorem. Proposition 4 of course extends similarly.

Let us return now to the issue of the third criterion in the definition of a correct extension. This requirement cannot be omitted, since that would allow fixpoints for which there is no support. For example, the valuation $[t, u, u]$ is a consistent fixpoint of G in Example 1 above. In the monotonic case, such fixpoints are eliminated by the requirement to be minimal, but that requirement is not sufficient here since $[t, u, u]$ is indeed a minimal fixpoint in the example (no "smaller" fixpoint exists). However, we can substitute instead of the third requirement another one which is similar in spirit to minimality, as follows.

A fixpoint v' of G above v is called *approachable* from v iff

$$v \sqsubseteq y \sqsubset v' \rightarrow (\exists y')(y \sqsubset y' \sqsubseteq v' \wedge G(y, y'))$$

Intuitively, this says that whenever we are on the path from v to v', there is some step allowed by G that will take us closer to v'.

This concept is a strengthening of the concept of least fixpoint, since the definition directly implies:

Proposition 8. If v' is an approachable fixpoint of G above v, then it is minimal.

It is easily seen that the least fixpoint $[t, u, u]$ in Example 1 is not approachable. This approachability condition can replace the third condition in the definition of the correct extension, and this property can be deduced from the following proposition.

Proposition 9. If v' is a consistent fixpoint of G above v, and v' is approachable from v, then it is a correct extension of v w.r.t. R.

Proof. v' immediately satisfies the first two conditions for being a correct extension. It remains to show that there is a derivation from v to v' using G.

Suppose this were not the case, i.e., every derivation from v whose members are $\sqsubseteq v'$ (the existence of at least one such chain is guaranteed by the definition of "approachable") has a l.u.b. $y \sqsubset v'$. Let v'' be the l.u.b. of all such y. By Proposition 4 there is a derivation from v to v''. since now $v'' \sqsubset v'$, consider the y' whose existence is guaranteed by the approachability, such that

$$G(v'', y')$$

$$v'' \sqsubset y' \sqsubseteq v'$$

The derivation step from v'' to y' must have been obtained using the extension of a kernel $\langle z, z' \rangle$ such that $z \sqsubseteq v''$, and z' is not $\sqsubseteq v''$. By compactness, in any derivation,

$$z_0, z_1, \ldots$$

of v'', there must be some element z_n such that $z \sqsubseteq z_n$. But then there is also a derivation of $z_n \sqcup z'$, which is not $\sqsubseteq v''$. This is a contradiction. []

In this way we have obtained the desired counterpart of the fixpoint criterion of the monotonic case.

6. Normal default rules

Reiter (1980) introduced the concept of *normal default rules*. He showed that every normal theory has an extension (in our terms a correct extension) and proved semi-monotonicity, i.e., larger set of normal default rules has a larger extension. His results, which are given with fairly complicated proofs, can now be obtained more easily from the material presented above.

A *normal default rule* is an NM rule of the form $\langle M, \{n\}, \{\} \rangle$.

If v is a valuation then v^* is the (obviously unique) least fixpoint of H above v. Clearly the * operation is monotonic. The valuation v is *fully consistent* iff v^* is consistent.

A valuation v is *saturated* w.r.t. a formula a iff $v(a) = v^*(a)$.

Let G be the restricted realization of a set of normal default rules. A derivation using G is *cautious* iff in each derivation step $\langle v, v' \rangle$ that uses a rule $\langle M, \{n\}, \{\} \rangle$, v is saturated w.r.t. n. The idea is that in a cautious derivation it is *not* possible to have the following scenario. Let the formula c be $\neg a$. Start from the valuation $[u, u, f]$, i.e., $\neg a$ is false but the conclusion that a is true has not been drawn. Use the extension of the NM rule $\langle -, a, b \rangle$ to derive $[f, t, f]$. After that, use the inference relation H which administrates simple truth-value calculations to derive $[k, t, f]$. The cases that we exclude by being cautious in this sense are the ones where an implementation would have to backtrack or (in a truth-maintenance system) shift propositions from IN to OUT status, because a non-monotonic antecedent, which was temporarily accepted since no proof had been bound so far, had to be retracted later when a proof was found.

Since H does derivation steps according to propositional logic, one can derive a valuation that is saturated w.r.t. n in a finite number of steps. If follows:

Proposition 10. Let v be a consistent valuation, and let G be the restricted realization of a set of normal default rules. Each step in a cautious derivation from v using G is fully consistent.

Proof. We first prove that any step is consistent. By the definition of restricted realizations, a derivation step using a normal default rule cannot introduce k into the valuation. Suppose a derivation step according to H in a cautious derivation does go from a consistent to an inconsistent valuation. Let n be the non-monotonic antecedent of the most recent derivation step, v' to v'', that uses an NM rule. We must have $v'(n) = u$, $v''(n) = f$, and $v'(l) = v''(l)$ for all other formulas l. (By the definition of restricted realizations, we could not have had $v'(n) = t$, $v''(n) = k$.) But by a familiar result of propositional logic, if there was a derivation using H from v'' to a contradiction, there must have been a derivation using H from v' to a valuation y where $y(n) = t$, which means v' was not properly saturated. Contradiction.

The full consistency follows easily, which completes the proof. []

Corollary. If v is a consistent valuation and there is a cautious derivation from v to v' using G, then v' is fully consistent.

Proposition 11. If G is the restricted realization of a set of normal default rules and there is a derivation from a valuation v to a fully consistent valuation y using G, then there is a cautious derivation using G from v to some y' such that $y \sqsubseteq y' \sqsubseteq y^*$.

Proof. We prove by induction on the derivation of y, using the monotonicity of F, that there is a derivation using F from v to some such y' (just saturate sufficiently before applying each rule). But since y^* is consistent, so is y', and the derivation using F is also a derivation using G. []

The main results are now:

Proposition 12. If v is a fully consistent valuation and G is the restricted realization of a set of normal default rules, then v has a correct extension w.r.t. G.

Proof. Let W be the set of l.u.b. of cautious derivations from v using G. By the corollary of Proposition 10, each member of W is fully consistent. A subset $Y \subseteq W$ is called *maximally consistent* iff the l.u.b. of its members is fully consistent, but the l.u.b. of every strict superset of Y which is still a subset of W, is not fully consistent. Consider some maximally consistent Y (it is clear that some exist, although maybe $Y = W$). Let y be the l.u.b. of Y. Clearly y is a fixpoint of H, i.e., $y = y^*$. According to Proposition 4 there is a derivation from v to y, and by Proposition 11 there is then a cautious derivation from v to some y' which is between y and y^*, so $y' = y$. However there cannot be any derivation from y using G, except the trivial derivation consisting only of y, since otherwise Y would not be maximal. Therefore y is a fixpoint for G and a correct extension. []

Propositon 13 (semi-monotonicity). Let v be a fully consistent valuation, $R' \subseteq R''$ be two sets of normal default rules, and $G' \subseteq G''$ be the restricted realizations of R' and R''. If v' is a correct extension of v w.r.t. G', then there exists some correct extension v'' of v w.r.t. G'' for which $v' \sqsubseteq v''$.

Proof. Let y be a correct extension of v' w.r.t. G''. Since using G'' there is a derivation from v to v' and a derivation from v' to y, there is also a derivation from v to y (Proposition 5). Let W be the set of l.u.b. of cautious derivations using G'' from v to, or passing through y. Proceed as in the previous proof. []

Reiter's proofs for the last two results are considerably more involved. The functional approach taken in this paper makes it possible to reason more abstractly, and therefore more concisely.

7. Conclusion

The functional view of logic focuses on inference relations, i.e., binary relations in a space of possible "information states" each of which represents partial knowledge in an agent or represents a state in the process performed by an inference machine. In the present paper, the information states are functions from propositions to truth values in a four-valued logic. Other kinds of information states are also conceivable. We have demonstrated the utility of this approach for studying non-monotonic inference rules.

The following are some promising directions for continued work:

(1) To characterize inference strategies in terms of transformations or criteria on inference relations.

(2) In particular, since non-monotonic deduction may be seen as a process that explores the tree of possible derivations from a given initial valuation, and which discards all branches where an inconsistent valuation has been reached, we would like to characterize "reason maintenance" methods that take us from an inconsistent valuation in a derivation, "sideways" to a consistent valuation along some other derivation from the same initial valuation.

(3) To consider information states with other truth-value domains, for example, for probability values or with information states which have an entirely different structure than that of mapping propositions to truth values. We would like to find such types of information states for which the same general properties apply (e.g., compactness) as were studied in this paper.

Acknowledgements

I wish to thank Jacek Leszczylowski for his constructive comments on an earlier version of this paper, and Brian Mayoh who brought Belnap's paper to my attention. The word "conservative" was chosen to conform to the terminology in Brown (1985) which uses the term "belief conserving" rules of inference.

BELNAP, N. D. 1977. A useful four-valued logic. *In* Modern uses of multiple-valued logic. *Edited by* J. M. Dunn and G. Epstein. Dordrecht.

BLIKLE, A. 1981. Notes on the mathematical semantics of programming languages. Mathematics Department, Linköping University, Linköping, Sweden. Report LiTH-MAT-R-81-19.

BROWN, A. 1985. Modal propositional semantics for reason maintenance systems. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, 1985, pp. 178–184.

BORGIDA, A., and IMIELINSKI, T. 1984. Decision making in committees—A framework for dealing with inconsistency and non-monotonicity. Proceedings of the American Association for Artificial Intelligence Non-Monotonic Reasoning Workshop.

GOODWIN, J. 1984. WATSON: a dependency directed inference system. Proceedings of the American Association for Artificial Intelligence Non-Monotonic Reasoning Workshop.

MANNA, Z. 1974. Mathematical theory of computation. McGraw-Hill Book Company, New York, NY.

MCDERMOTT, D., and DOYLE, J. 1980. Non-monotonic logic I. Artificial Intelligence, **13**, pp. 41–72.

REITER, R. 1980. A logic for default reasoning. Artificial Intelligence, **13**, pp. 81–132.

SANDEWALL, E. 1972. An approach to the frame problem, and its implementation. *In* Machine intelligence 7. *Edited by* B. Meltzer, and D. Michie. Wiley, New York, NY, pp. 195–204.

SCOTT, D. 1970. Outline of a mathematical theory of computation. Proceedings of the 4th Annual Princeton Conference on Information Sciences and Systems, Princeton, NJ, pp. 169–176.

STOY, J. 1977. Denotational semantics: the Scott-Strachey approach to programming language theory. The MIT Press, Cambridge, MA.