DEFAULT REASONING USING MONOTONIC LOGIC: A Modest Proposal

Jane Terry Nutter

Department of Computer Science, Tulane University

ABSTRACT

This paper presents a simple extension of first order predicate logic to include a default operator. Rules of inference governing the operator are specified, and a model theory for interpreting sentences involving default operators is developed based on standard Tarskian semantics. The resulting system is trivially sound. It is argued that (a) this logic provides an adequate basis for default reasoning in A.I. systems, and (b) unlike most logics proposed for this purpose, it retains the virtues of standard first order logic, including both monotonicity and simplicity.

I INTRODUCTORY COMMENTS

Reasoning from incomplete information and from default generalizations follows patterns which standard first order predicate logic does not reflect. The most striking deviation involves making inferences whose conclusions may be counterindicated by further information which does not explicitly contradict anything previously known.

In standard logics, if a set of premises entails a conclusion, any larger set containing all those premises also entails that conclusion. Logics with this property are called monotonic. The above departure from standard logic's reasoning patterns has led researchers to adopt and develop non-monotonic logics for use in A.I. systems (see e.g. McDermott and Doyle, 1980; McDermott, 1982; Reiter, 1980; Aronson *et al.*, 1980; Duda *et al.*, 1978; and others).

Critics of non-monotonic logics have pointed out technical weaknesses (see e.g. Davis, 1980), and Israel (1980) argues persuasively that its supporters confuse logic with complex judgments of kinds logic cannot deal with. But unless some alternative appears for dealing with default reasoning, non-monotonic logic must continue to appeal.

I have argued elsewhere that (1) default reasoning only appears non-monotonic if we fail to distinguish warranted assertions from warranted assumptions (Nutter, 1982), and (2) any logic which adequately models default reasoning must be monotonic (Nutter, 1983).

If this is accepted, a conservative approach promises more than do radical ones, not because it is "safer", but because it simultaneously preserves what is of value in standard logic -- its simplicity, clarity, and adequacy in domains of complete information -- and allows distinguishing warranted assumptions from warranted assertions (generalizations from universals, default consequences from genuines inferences, etc.). This paper presents an almost trivial extension of first order predicate logic which distinguishes generalizations and the resulting presumptions from universals and their consequents and permits reasoning from default generalizations.

II DEFAULTS AS PROPOSITIONAL OPERATORS

The logic we are aiming for must differ from standard logic in three ways. First, it must distinguish "guarded" from "unguarded" propositions (presumptions from the related assertions). Second, it must allow for appropriate inferences involving guarded propositions, with "inheritence" of the guard. Third, it's semantics must allow as compatible a guarded proposition and the proposition's unguarded negation.

So our first task is to find an appropriate representation for this guard. It cannot be a new connective in the ordinary truth functional sense, the most familiar kind of logical constant, because default generalizations are not truth functions of their component propositions. Nor is the guard a quantifier: that would reduce defaults to some kind of statistical operators. In fact, no standard category of logical constant seems quite appropriate. Neither can we use a non-logical constant, since the meaning of a non-logical constant cannot affect the course of reasoning in a logic.

So we need a new kind of logical operator, which operates on propositions and affects the course of reasoning. There is already a unary operator on formulas of that general kind: negation. The default operator belongs to a new class not because it binds formulas, but because it does so without being truth-functional. The resulting proposition is obviously related to its component in an interesting way, but its truth value is not a function of the component's.

This paper describes how to extend standard first order predicate logic to a logic appropriate for default reasoning. This involves four tasks. (1) We must characterize the *grammar* of the extended language. That is, we must give rules for determining whether a particular string of symbols represents a proposition in the extended language. (2) We must explain how to extend the *deductive system* of the original logic so that all and only desirable inferences will be legal. (3) We must describe the *semantics* for the language: that is, we must say what the interpretations of the language are, how truth values on the interpretations are determined, and how logical entailment is defined. Finally (4) we should investigate at least briefly the *metalogic* of the extended system. That is, we should assure ourselves that the new extended system is sound: it never permits false conclusions to be derived from true premises. We now take up these tasks.

III GRAMMAR

The first component of a formal logic is its grammar, which determines the *language* of the logic, which may be defined as either the set of formulas or the set of propositions which are considered well formed. Formulas are distinguished from propositions by containing open variable occurrences, that is, placeholders which could *either* be replaced by specific individual constants -- names -- *or* be bound by an existential or a universal quantifier -- turned into an equivalent of either "something" or "everything" -- but which at present occur unbound. In this section, we present the rules of grammar for the language of our logic.

The technical discussions in the sections below will presuppose familiarity with the usual ways that formal logics are described, and will presuppose a complete description of a first order predicate logic. The informal descriptions should convey an overview for readers who lack this familiarity and motivate some of the less obvious decisions embodied in the technical discussions.

A. Informal Description

The grammar holds relatively few surprises. In simple terms, wherever in English one might reasonably prefix a clause with "There is reason to believe that" or some similar construction, the logic will allow its equivalent of the clause to be governed by our default operator *p* (for "*p*resumably").

The only real decision involves whether to let the operator govern formulas or only propositions. While its use governing formulas is probably eliminable, it seems to me both innocuous and well motivated. While "Birds generally fly," which becomes $\forall\alpha(\mathrm{Bird}(\alpha) \Rightarrow p\ \mathrm{Flies}(\alpha))$ in the formal logic, probably comes to the same thing as "In general birds fly" $[p\ \forall\alpha\ (\mathrm{Bird}(\alpha)) \Rightarrow \mathrm{Flies}(\alpha))]$ in the long run, both are said. Furthermore, "If anything is a bird, then it has wings and presumably flies" must be considerably reformulated before it can be equivalently expressed without having *p* bind a formula instead of a sentence. Hence *p* may directly govern either a formula or a sentence.

B. Technical Presentation

Let *L* be a standard first order predicate calculus without equality and without function symbols, with a Jaskowski-style natural deduction system (Jaskowski, 1934) containing explicit introduction and elimination rules for the standard connectives and quantifiers and a small set of "bookkeeping" rules. (The precise formal system is unimportant.) The language of *L* is L, the logic with defaults which we will build from *L* is *LD*, and its language is LD. The default operator is written *p*.

Throughout the discussion below, Roman capitals vary over formulas and propositions of L, while Greek letters Φ and Ψ vary over formulas and propositions of LD, and α varies of variables of LD.

Def. The set of formulas of LD is the smallest set satisfying the following clauses:

(a) All atomic formulas of L are formulas of LD;

(b) For any formula Φ of LD, $p\Phi$ is a formula of LD;

(c) For any formulas Φ,Ψ of LD, the following are all formulas of LD: $\sim\Phi$; $\Phi \vee \Psi$; $\Phi \wedge \Psi$; $\Phi \Rightarrow \Psi$; $\Phi \equiv \Psi$.

(d) For any formula Φ of LD in which α occurs open, the following are also formulas of LD: $\forall\alpha\ \Phi$; $\exists\alpha\ \Phi$.

For any formula Φ of LD, Φ is a proposition of LD if and only if Φ contains no open occurrences of variables.

IV DEDUCTIVE SYSTEM

As commented above, the deductive system is a standard natural deduction system, with the usual five connectives ("not", written "$\sim$", "and", written "$\wedge$", "or", written "$\vee$", "if...then...", written "$\Rightarrow$", and "if and only if", written "$\equiv$") and both quantifiers ("for all", written "$\forall$", and "there is", written "$\exists$"). This will be the least altered portion of the logic.

A. Informal Description

The guarded status of default generalizations must be inherited through inferences, and from a guarded proposition, it must be possible to infer a guarded version of any conclusion which could be inferred from the unguarded version of the premise. That is, if you could infer from "Roger flies" that "If my dog sees Roger, he will find him fascinating," then from "Presumably Roger flies" you may infer "Presumably, if my dog sees Roger, he will find him fascinating," and so on. However, it generally takes more than one premise to warrant an inference. How many of the premises may be generalizations? There are several ways to go on this.

We could allow only one guarded premise per inference. But suppose we have "In general, if A then B," and we have "There is reason to believe that A." We would normally feel free to infer "There is reason to believe that B." But this inference has two guarded premises, and without further information, there is no way to avoid using both guarded premises in a single inference and still get the conclusion.

Alternatively, we could "string out" the guards, putting as many *p*s on the front of the conclusion as there were guarded premises contributing to it. Primae facie, this might seem to give not only a guard, but also a sort of certainty measure. Unfortunately, it's a very poor measure for anything interesting. Here's why.

Default generalizations are not statistical. They don't necessarily "accumulate". There are several reasons for this. First, they embody judgments of typicallity. As such, related default generalizations will frequently stand or fall together: if the instance is typical, they all hold, while if the instance is atypical, each is likely to fail. Certainly this phenomenon cannot be relied upon, but it is common enough to upset the measure. (For more on this and related points, see Nutter, 1982 and Nutter, 1983b.)

Second, a single guarded proposition might be used five or six times in the course of a long derivation. Is the result any less certain for having that proposition enter many times rather than once? By the time a *chain* of inferences have taken place, the number of *p*s on the front may no longer represent *even* the number of distinct generalizations involved in reaching it.

Third, there is no reason to suppose that all the original generalizations are equally reliable. Five extremely strong generalizations probably provide better support than one relatively weak one. If we could assign weights reflecting reliability -- i.e. probabilities -- we would be dealing with statistical generalizations and not with defaults anyhow.

Furthermore, from "Generally, if A than B" and "There is reason to believe that A," we don't conclude "There is reason to believe that there is reason to believe that B." We conclude that there is reason to believe B. Stringing out *p*s does not reflect any practice present in English usage. If that practice played a role in the kind of reasoning we are trying to model, it seems reasonable to suppose that language would reflect it. I have therefore decided to "inherit" the guard if any or all premises involved it, without regard to how many. This is the force of the first rule of inference below.

Even with the rule in that form, it would be possible to generate long strings of *p*s. The second rule says that if there are more than one consecutive qualifications, they may be collapsed into a single one.

B. Technical Development

We retain all the standard rules of inference in L (generalizing them to allow for sentences of LD and not just of L), and add the following two rules.

RpI: Suppose that $\Psi = \{\psi^1, \ldots, \psi^n\} \subseteq$ LD, $\Psi' = \{\psi^{1\prime}, \ldots, \psi^{n\prime}\} \subseteq$ LD, $\phi \in$ LD and for all i, $1 \le i \le n$, $\psi i' = \psi i$ or $\psi i' = p\psi i$. Suppose further that some rule of inference warrants inferring ϕ from Ψ. Then from Ψ' you may infer $p\phi$.

RpE: From $pp\phi$ you may infer $p\phi$.

Given $\Phi \subseteq$ LD, $\psi \in$ LD, we say that ψ *is provable from* Φ (or Φ *proves* ψ, written $\Phi \vdash \psi$) provided that ψ can be derived from premises in Φ following the rules of inference in the usual and obvious way.

V SEMANTICS

This is the portion of the logic which deals with issues of "meaning" and truth. The sense of meaning involved is extremely primitive. For the non-logical terms, it reflects only the function of words as referring in some way to objects, completely neglecting not only *how* they refer, but even *to what* they refer in ordinary English. Formal semantics are *formal*: they deal with what can be said about inheritance of truth conditions on the basis of logical form alone, which throws out the overwhelming majority of what we would ordinarily mean when we talk about meaning.

The semantics given below represents the most complete departure from standard logic in *LD*. The standard core of the system retains the familiar semantic properties, but across the portion of LD which is not in L, the logic is three-valued. This can be represented in many ways; I have chosen to represent the truth values as non-empty subsets (instead of elements) of the set {t,f} of familiar truth values. Intuitively, {t} means "only true", and occurs where in ordinary logic one would expect t; similarly {f} means "only false". The "newcomer" is {t,f}, which can best be thought of as "well, yes and no," or both true *and* false. Taken in this way, *LD* provides an intriguing face for lovers of radical departures: it remains sound while violating the traditional law of non-contradiction.

A. Evaluating Truth of Presumptions

What does it *mean* for a default generalization to be true? It does *not* mean that the statement inside the *p* is true. It does not even mean that there is no reason to believe that it is false: $p\phi \wedge p \sim\phi$ means something like, "There is reason to believe that ϕ is true, but there is also reason to believe that it is false." That sentence is consistent if unhelpful.

Part of what we want our default operator to mean which cannot be modeled by formal semantics: neither causality nor typicality is a formal semantic concept. But we can state certain constraints on how our *p* operator works.

Any true proposition is evidence for itself. (If we *knew* that ϕ is true, then we would certainly be willing to suppose it. The logic cannot distinguish the truth value of $p\phi$ depending on whether we know it or not.) Hence if ϕ is true, then $p\phi$ should be *at least* true. If ϕ and ψ are both only true, then clearly their conjunction should be only true as well; if either is only false, their conjunction should be only false. Otherwise, the conjunction should be both true and false. Similar comments work for the other connectives.

There is a restriction on our interpretations which is not obvious from an intuitive standpoint. It arises from the fact that for any ϕ, $\phi \vDash p\phi$. Suppose we have $\phi \Rightarrow \psi$, and also $p\phi$. Then we can infer $p\psi$. For this to be sound, it must be the case that whenever $p\phi$ and $\phi \Rightarrow \psi$ are both true, so is $p\psi$.

But consider an interpretation on which ϕ is false. Then $\phi \Rightarrow \psi$ is true regardless of the truth value of ψ (or anything else about ψ, for that matter; this is one of the so-called paradoxes of implication of first order logic). Suppose further that the interpretation makes $p\phi$ true (as well as false). This interpretation *must* make every proposition of the form $p\psi$ at least true, in order to support the inference of $p\psi$ from $\phi \Rightarrow \psi$ and $p\phi$. Now from what we said above, if ψ is true, then $p\psi$ is already at least true. But $p\psi$ must *also* be at least true for false ψ.

Hence for any interpretation, if even one false proposition has the property that there is reason to believe it on that interpretation, then all false propositions have that property on that interpretation (though the presumption forms of the true propositions may be only true, so this does not trivialize these interpretations.) This fact has its basis in the ordinary truth functional definition of "if ... then ...". Relevance logic can prevent the undesired deductions, but so far a denotational model theory which distinguishes relevance logic from standard logic is lacking.

B. Interpretations

An interpretation i = (U,e) of *LD* is defined from ordinary Tarskian interpretations of *L* by much the obvious sort of extension given the comments above, except that we alter the treatment of ordinary propositions slightly. In particular, where before t represented true and f false, we will now use {t} and {f} respectively. That is, for all propositions ϕ, $e(\phi) \subseteq \{t,f\}$; as one would expect, for A $\in$ L, e(A) = {t} or e(A) = {f}. It works out that for $\phi \in$ LD $e(\phi) = \{f\}$ or $e(\phi) = \{t\}$ or $e(\phi) = \{t,f\}$ (i.e. the presumption is either true, false, or both), but it will never happen that $e(\phi) = \Lambda$.

In particular, having defined the evaluation function over atomic propositions of LD in the usual way (except for the substitution of {f} and {t} for f and t), we use the following induction clauses:

1. For all $\phi \in$ LD, if ϕ has the form *p*A for A $\in$ L, then $e(A) \subseteq e(\phi)$.
2. For all $\phi \in$ LD, if ϕ has the form $pp\psi$ for some $\psi \in$ LD, then $e(\phi) = e(p\psi)$.
3. For all ϕ, $\psi \in$ LD, we have

$$e(\sim\phi) = \{f\} \text{ if } e(\phi)=\{t\};\ \{t\} \text{ if } e(\phi)=\{f\};\ \{t,f\} \text{ if } e(\phi)=\{t,f\}.$$

$$e(\phi \wedge \psi) = \{f\} \text{ if } t \notin e(\phi) \text{ or } t \notin e(\psi);\ \{t\} \text{ if } f \notin e(\phi) \text{ and } f \notin e(\psi);\ \{t,f\} \text{ otherwise.}$$

$$e(\phi \vee \psi) = \{f\} \text{ if } t \notin e(\phi) \text{ and } t \notin e(\psi);\ \{t\} \text{ if } f \notin e(\phi) \text{ or } f \notin e(\psi);\ \{t,f\} \text{ otherwise.}$$

$$e(\phi \Rightarrow \psi) = \{f\} \text{ if } f \notin e(\phi) \text{ and } t \notin e(\psi);\ \{t\} \text{ if } t \notin e(\phi) \text{ or } f \notin e(\psi);\ \{t,f\} \text{ otherwise.}$$

$$e(\phi \equiv \psi) = \{f\} \text{ if } e(\phi) \cap e(\psi) = \Lambda;\ \{t\} \text{ if } e(\phi) = e(\psi) = \{t\} \text{ or } e(\phi) = e(\psi) = \{f\};\ \{t,f\} \text{ otherwise.}$$

$$
\begin{aligned}
e(\forall\alpha\ \Phi) = \{f\} &\text{ if there is a } u \in U \text{ such that } e(\Phi{:}u/\alpha) = \{f\};\\
\{t\} &\text{ if for all } u \in U,\ e(\Phi{:}u/\alpha) = \{t\};\\
\{t,f\} &\text{ otherwise.}
\end{aligned}
$$

$$
\begin{aligned}
e(\exists\alpha\ \Phi) = \{f\} &\text{ if for all } u \in U,\ e(\Phi{:}u/\alpha) = \{f\};\\
\{t\} &\text{ if there is a } u \in U \text{ such that } e(\Phi{:}u/\alpha) = \{t\};\\
\{t,f\} &\text{ otherwise.}
\end{aligned}
$$

4. If there is a $\Phi \in LD$ such that $e(\Phi) = \{f\}$ but $t \in e(p\Phi)$, then for all $\Psi \in LD$, $f \in e(p\Psi)$.

(Notational remark: $e(\Phi{:}u/\alpha)$ means the result of evaluating Φ, treating α as if it were a constant and $e(\alpha) = u$.)

The standard equivalences follow trivially from these definitions. Now we define satisfaction as follows: for all $i = (U,e)$ an interpretation of *LD*, and for all $\Phi \in LD$, i satisfies Φ (written $i \models \Phi$) if and only if $t \in e(\Phi)$.

VI METALOGIC

As has been indicated from the outset, the characteristic we are most concerned with for logics in inference system implementations is soundness. For this logic, the following theorem follows trivially from the soundness of standard first order predicate logic:

Theorem The logic *LD* is sound, i.e. for all $\Psi \subseteq LD$, $\Phi \in LD$, if $\Psi \vdash \Phi$ then $\Psi \models \Phi$.

VII CONCLUDING REMARKS

A. Issues of relevance

As noted above, the apparently paradoxical features which result from adopting a truth-functional view of implication can be avoided by using a deductive system based on relevance logic. Such a deductive system will permit inferences which are a proper subset of those allowed under the traditional view presented here. Martins (1983) has developed a variant of relevance logic for belief revision, which has been implemented in SNePS (Shapiro, 1979). An inference system including default operators as described in this paper is being implemented on this belief revision system. Limitations of space prohibit fuller discussion here; for details, see Nutter (1973b).

B. Significance of the logic

It is important not to *over*estimate what a logic for default reasoning can do once we have one. As Israel (1980) points out, arriving at useful generalizations and resolving conflicts among conclusions of default reasoning both exceed the scope of logic. To those who wonder what help the system will be once it has deduced "$p\Phi \wedge p{\sim}\Phi$", the answer is, none at all. *If* the non-logical portion of your system contains heuristics saying that conflicting evidence indicates a need for further investigation, and *if* your system further contains some subsystem for investigating, *then* the ability to deduce statements of the kind above can be used to trigger investigation; but this goes beyond the logic alone.

It does not follow that such a logic is useless. I have already argued that implementing a logic of the kind outlined here allows implementing useful salience rules in natural language generation, for instance. Furthermore, this kind of inference will play an intimate role in any system which investigates failed expectations, conflicting expectations, and the like. Only it will not *constitute* such a system: instead it is a tool for the system to use.

The logic presented here makes no pretense to philosophical depth. As a logic, it is trivial; as a philosophical view of generalization, it is perhaps the shallowest possible. But it can be used to form non-trivial systems modeling deep views. From one point of view, that is what logic is for.

VIII ACKNOWLEDGEMENTS

Many thanks to Stuart Shapiro and to the SNePS Research Group for their many helpful comments and suggestions. This research was carried out in the Department of Computer Science, State University of New York at Buffalo, Amherst, New York.

REFERENCES

[1] Aronson, A. R., Jacobs, B. E., and Minker, J., "A note on fuzzy deduction." *JACM* 27:4 (1980) 599-603.

[2] Davis, M., "The mathematics of non-monotonic reasoning." *Artif. Intell.* 13:1-2 (1980) 73-80.

[3] Duda, R. O., Hart, P. E., Nillson, N. J., and Sutherland, G. L., "Semantic network representations in rule-based inference systems" In *Pattern - Directed Inference Systems*, D. A. Waterman and F. Hayes-Roth, editors, Academic Press (New York, 1978) pp. 203-223.

[4] Israel, D.J., "What's wrong with non-monotonic logic?" *Proc. AAAI-80*. Palo Alto, California, August, 1980, pp. 99-101.

[5] Jaskowski, S., "On the rules of supposition in formal logic." *Studia Logica* 1 (1934).

[6] Martins, J. P., "Reasoning in Multiple Belief Spaces," Technical Report 203, State University of New York at Buffalo, Amherst, New York, June 1983.

[7] McDermott, D. V. and Doyle, J., "Non-monotonic logic I." *Artif. Intell.* 13:1-2 (1980) 41-72.

[8] McDermott, D., "Non-monotonic logic II." *JACM* 29:1 (1982) 33-57.

[9] Nutter, J. T., "Defaults revisited, or "Tell me if you're guessing." *Proc. Cog. Sci. 4*. Ann Arbor, Michigan, August, 1982, pp. 67-69.

[10] Nutter, J. T., "What else is wrong with non-monotonic logics? Representational and informational shortcomings." *Proc. Cog. Sci. 5* Rochester, New York, May, 1983.

[11] Nutter, J. T., "Default reasoning in A.I. systems" (1983b) draft.

[12] Reiter, R., "A logic for default reasoning." *Artif. Intell.* 13:1-2 (1980) 81-132.

[13] Shapiro, S., "The SNePS semantic network processing system" In *Associative Networks*, N. V. Findler, ed., Academic Press (New York, 1979) pp. 791-796.