Formalizing Non-Monotonic Reasoning Systems

David W. Etherington
Department of Computer Science
University of British Columbia
Vancouver, B.C., Canada, V6T 1W5.

Technical Report 83-1

Abstract

In recent years, there has been considerable interest in non-monotonic reasoning systems. Unfortunately, formal rigor has not always kept pace with the enthusiastic propagation of new systems. Formalizing such systems may yield dividends in terms of both clarity and correctness. We show that Default Logic is a useful tool for the specification and description of non-monotonic systems, and present new results which enhance this usefulness.

Table of Contents

Table of Figures

Formalizing Non-Monotonic Reasoning Systems

David W. Etherington[1]
Department of Computer Science
University of British Columbia

Abstract

In recent years, there has been considerable interest in non-monotonic reasoning systems. Unfortunately, formal rigor has not always kept pace with the enthusiastic propagation of new systems. Formalizing such systems may yield dividends in terms of both clarity and correctness. We show that Default Logic is a useful tool for the specification and description of non-monotonic systems, and present new results which enhance this usefulness.

1. Introduction

Human common-sense reasoning appears to rely heavily upon the ability to use general rules subject to exceptions; what has been called prototypic or default information. Virtually none of the decisions one makes everyday are made with complete certainty. With little effort, an endless supply of more or less probable scenarios can be constructed which contraindicate any chosen course. Yet people are not paralyzed by indecision; they continue to act and to decide in spite of all this uncertainty.

AI researchers have placed great emphasis on this ability to act "rationally" in the absence of complete, definitive knowledge about situations. A variety of systems have been designed which address various facets of the problem. Unfortunately, many have been too informal to allow full understanding of their performance and evaluation of their correctness. Conversely, those formal systems which have been developed have typically lacked any computationally "realistic" inference mechanism. While one might have hoped for synergy, with a formal system providing a semantic foundation and an implementation providing "efficient" computation, this has been slow in coming.

[1] This research was supported in part by an I.W. Killam Predoctoral Fellowship and by an NSERC Postgraduate Scholarship.

We focus on one formal system, Default Logic [Reiter 1980], and show its usefulness in formalizing the semantics of more computationally oriented schemes. Default Logic appears to capture many of the intuitions underlying systems for reasoning with incomplete or prototypic information. This allows it to serve quite naturally as a system specification tool. We also present several new results which facilitate this application. These include a correct inference procedure and sufficient conditions for coherence of a default theory. These conditions considerably simplify the problem of proving a system coherent.

We illustrate the suitability of Default Logic for semantic specification by developing a semantics for the familiar inheritance hierarchies with exceptions. This semantics highlights the sources of many of the problems encountered by those trying to develop parallel inference algorithms for such structures. Properties of the underlying default theories suggest partial solutions to some of these.

There are a number of related open questions. Some of the more interesting problems and promising directions for future research are discussed briefly in the final sections.

2. An Introduction to Default Logic

In common-sense reasoning, conclusions are often based on both supporting evidence and the absence of contradictory evidence. Traditional logics cannot emulate this form of reasoning because they lack any means for considering the absence of knowledge. Default Logic was developed to address this shortcoming in a principled way, by augmenting a traditional first-order logic with a mechanism for predicating conclusions on the absence of specific knowledge.

A brief introduction to Default Logic follows. A detailed development can be found in [Reiter 1980]. The discussion is relatively self-contained, but a familiarity with first-order logic is assumed. (See [Mendelson 1964] for an introduction.) Readers already familiar with Default Logic may wish to proceed directly to the next section.

A *default theory*, (D,W), consists of a set of first-order formulae, W, and a set of defaults, D. A *default* is any expression of the form:

$$\frac{\alpha(x_1, \ldots, x_n) : \beta(x_1, \ldots, x_n)}{\gamma(x_1, \ldots, x_n)}.$$

$\alpha(x_1, \ldots, x_n)$, $\beta(x_1, \ldots, x_n)$, and $\gamma(x_1, \ldots, x_n)$ are all well-formed formulae whose free variables are among $x_1, \ldots, x_n$, and are called the *prerequisite*, *justification*, and *consequent* of the default, respectively.

Defaults serve as rules of inference or conjecture, augmenting those normally provided by first-order logic. Under certain conditions, they sanction inferences which could not be made in a strictly first-order framework. If the prerequisite of a default is known and its justification is "consistent" with what is known (i.e. its negation is not provable[2]), then its consequent may be inferred. The consequent's status is akin to that of a belief, subject to revision should the justification be denied at some point. Defaults are thus *non-monotonic* inference rules, since the addition of new information which denies the justification of a default may invalidate previously derived conclusions.

Since defaults allow reference to what is not provable in the determination of what is provable, the "theorems" of a default theory are not so easy to generate as are those of a first-order theory. What is provable both determines and is determined by what is not provable. To avoid this apparent circularity, the theorems of a default theory are defined by a fixed-point construction. A set, E, is an *extension* for a default theory, $\Delta = (D, W)$, if and only if it is a fixed-point of an operator which yields a minimal set containing W, closed under the provability relation, and containing the consequent of any default whose prerequisite is in E and whose justification is consistent with E. More formally, if S is a set of closed formulae and $\Gamma(S)$ is the smallest set satisfying the following three properties:

(1) $W \subseteq \Gamma(S)$

(2) $\Gamma(S) = Th(\Gamma(S))$[3]

(3) If $\frac{\alpha : \beta}{\omega} \in D$ and $\alpha \in \Gamma(S)$ and $\neg\beta \notin S$ then $\omega \in \Gamma(S)$.

then a set E is an extension for Δ if and only if $\Gamma(E) = E$, i.e. if and only if E is a fixed point of the operator Γ [Reiter 1980].

The defaults in D can be viewed as extending the first-order knowledge, contained in W, about an incompletely specified world. An extension is then interpreted as an acceptable set of beliefs one may hold about that world. Not every default theory has an extension, and some have more than one. The defaults select a restricted subset of the models of the underlying first-order theory. Thus any model for an extension of Δ will be a model for W; the converse is generally not true.

A simple example may help to illustrate. Consider the following default theory:

$W = \{ Block(A) \vee Block(B) \}$

---

[2] The reader is referred to [Reiter 1980] for a precise definition of "provability".

[3] $Th(X)$ is the closure of the set X under first-order provability. i.e. $\omega \in Th(X) \leftrightarrow X \vdash \omega$.

$$D = \{ \frac{: \neg Block(x)}{\neg Block(x)} \} .$$

While not strictly accurate, free variables in a default may be thought of as implicitly universally quantified, with scope covering the whole default. Thus, the above default may be loosely interpreted as saying: "If it is consistent to assume that something is not a block, do so." This theory has two extensions, E_1 and E_2:

$$E_1 = Th(\{ \neg Block(A), Block(B), Block(A) \vee Block(B) \})$$

$$E_2 = Th(\{ Block(A), \neg Block(B), Block(A) \vee Block(B) \})$$

Note that there is no extension containing both $\neg Block(A)$ and $\neg Block(B)$, since no model for W contains both. Neither is there an extension with both $Block(A)$ and $Block(B)$, since there is no support for concluding both.

3. Reasoning with Default Logic

Extensions play a fundamental role in Default Logic. An extension is a set of beliefs which are in some sense "justified" or "reasonable" in light of what is known about a world. Formally, extensions are attractive because they are both well-founded and complete: A formula enters an extension, E, only if it is in W, if it is provable from other formulae in E, or if it is the consequent of a default whose prerequisites are in E and whose justifications are not denied by E; furthermore, every formula which meets these requirements is in E. The first of these restrictions prevents extensions from containing spurious, unsupported beliefs. The second ensures that justified beliefs are not ignored. The restrictions are analogous to those which define the theorems of a first-order theory.

Since the individual extensions of a default theory are both well-founded and complete, it is quite natural to require any default inference system to restrict its conclusions to one common extension. If no extension of a theory contains a formula, then it is not in any acceptable set of beliefs associated with that theory. If conclusions are drawn from different extensions, they may be incompatible. Consider the blocks-world example from the previous section. In that example, both $\neg Block(A)$ and $\neg Block(B)$ are reasonable assumptions. They are drawn from different extensions, however, and concluding both leads to inconsistency.

Since reasonable conclusions must reside in an extension of the default theory under consideration, it is clearly important to know whether every theory has extensions. Simply put, the answer is no. For example, the theory:

$$W = \{ \}$$

$$D = \{ \frac{:A}{\neg A} \}$$

has no extension. Such theories are *incoherent;* they support no reasonable set of beliefs about the world. Beyond pointing out the existence of incoherent theories, the most useful answer would include a syntactic characterization of which theories have or do not have extensions. While no such characterization is known, there are sufficient conditions which guarantee extensions. We present three such conditions below, in order of increasing utility.

A theory, $(\{ \}, W)$, with no defaults has a unique extension, $Th(W)$, the logical closure of the underlying first-order theory. Of course, this is a trivial default theory. We mention it only to emphasize that, since Default Logic is a superset of first-order logic, the required results obtain for the area of overlap.

Distinguishing commonly encountered types of defaults leads to more enlightening results. Any default of the form:

$$\frac{\alpha : \beta}{\beta}$$

is said to be *normal.* Normal defaults are sufficient for knowledge representation and reasoning in many naturally occurring contexts. In fact, they can express any rule whose application is subject only to first-order prerequisites and the consistency of its conclusion with the rest of what is believed. Rules like:

"Assume a bird can fly unless you know otherwise.", or
"Assume a thing is not a block unless it is required to be."

translate easily into normal defaults:

$$\frac{Bird(x) : Can\text{-}fly(x)}{Can\text{-}fly(x)} \quad \text{and} \quad \frac{: \neg Block(x)}{\neg Block(x)} .$$

The consequent of a normal default is equivalent to its justification. Intuitively, this makes the default inapplicable where the consequent has been denied. Such defaults cannot introduce inconsistencies, they cannot refute the justifications of other, already applied, normal defaults, nor can they refute their own justifications. This gives rise to well-behaved theories. Any theory involving only normal defaults (a *normal theory*) must have at least one extension [Reiter 1980].

Their broad applicability and the guarantee of coherence makes normal defaults attractive for knowledge representation and reasoning. There are, however, some types of knowledge which normal defaults cannot represent. For example, Reiter and Criscuolo [1983] have noticed that defaults sometimes interact with one another, and that normal defaults cannot adequately constrain these interactions. One manifestation of this occurs when two defaults with distinct but not mutually exclusive prerequisites have contradictory consequents. In such circumstances it is not always clear which default should be applied. Commonsense reasoning usually prefers one of the competing defaults by virtue of its prerequisite being more specific. This preference cannot be enforced using only normal defaults. For example, assume we are given:

Typical adults are employed.
Typical high-school dropouts are adults.
Typical high-school dropouts are not employed.

This may be expressed by the following normal defaults:

$$\{ \frac{Adult(x) : Employed(x)}{Employed(x)}, \frac{Dropout(x) : Adult(x)}{Adult(x)}, \frac{Dropout(x) : \neg Employed(x)}{\neg Employed(x)} \}.$$

For a given a dropout, this theory can be seen to have two extensions which differ on her state of employment. Intuition dictates that we assume she is unemployed. Careful consideration shows that the conflict arises because typical dropouts are not *typical* adults; this atypicality should block the transitivity from *Dropout* through *Adult* to *Employed*. The first default incorporates no explicit reference to these exceptional circumstances which should block its application. One way to address this problem is to require that the case under consideration not be a known exceptional case. This requirement is then added to the justification. Thus the first default above becomes:

$$\frac{Adult(x) : Employed(x) \wedge \neg Dropout(x)}{Employed(x)},$$

which is not applicable to known dropouts.

Any default of the form:

$$\frac{\alpha : \beta \wedge \gamma}{\beta}$$

is said to be *semi-normal*. Semi-normal defaults differ from normal defaults by having non-tautologous conjuncts in their justifications which do not occur in their consequents. The assurances of well-behavedness associated with normal theories do not carry over to theories with semi-normal defaults. For example, the theory:

$$W = \{ \}$$
$$D = \{ \frac{:A \wedge \neg B}{A}, \frac{:B \wedge \neg C}{B}, \frac{:C \wedge \neg A}{C} \} \qquad (1)$$

has no extension. This appears to be a somewhat artificial example, inasmuch as we have been unable to find a natural situation which fits this pattern. Which semi-normal theories, then, are assured of extensions? Do all "natural" theories have extensions? Perhaps pathological examples are merely formal curiosities? We do not purport to answer these questions — partly because of the difficulty of delimiting the class of "natural" theories. There is, however, a large class of semi-normal theories which are coherent. We characterize this class, which appears to be sufficient for many common applications, in the next section.

4. Ordered Default Theories

There appears to be a unifying characteristic among default theories without extensions. Consider the theory:

$$W = \{\ \}$$

$$D = \{\frac{:A}{\neg A}\}$$

which has no extension. The only reasonable candidates are $\overline{E}_1 = Th(\{\ \})$ or $\overline{E}_2 = Th(\{\neg A\})$. A is consistent with $\overline{E}_1$, so to be an extension $\overline{E}_1$ must contain $\neg A$, which it does not. Similarly, A is inconsistent with $\overline{E}_2$, so $\overline{E}_2$ cannot contain $\neg A$. The problem is that the default's justification is denied by its consequent; not applying the default forces its application, and vice versa. Returning to the semi-normal theory (1), we see that applying any one default leaves one other applicable. Applying any two, however, results in the denial of the non-normal part of the justifications of at least one of them. Any set small enough to be an extension is too small; any set large enough is too large. This behaviour is characteristic of theories with no extension; the requirement that extensions be closed under the default rules forces the application of defaults whose consequents lead to the denial of justifications of other applied defaults.

The exact source of the problem can be further isolated by recalling that all normal theories have extensions. Since the justification and consequent of normal defaults are identical, no applicable default can refute the justifications of an already applied default: applied normal defaults have already asserted their justifications. This means that any normal default capable of refuting those justifications is inapplicable, since its justifications have already been refuted. It follows that that part of the justification which distinguishes non-normal defaults from normal defaults is integrally involved in making a theory incoherent. Restricting our attention to semi-normal default theories, we see that once a default has been applied, only those conjuncts of its justification not entailed by its consequent are susceptible to refutation by other defaults. These conjuncts play a key role in the discussion below.

The conflict between closure under defaults and consistency of justifications can occur only if some formula depends on the absence of another and at the same time may serve to support the inference of that formula. In the theory (1) above, for example, A depends on the absence of B, B on that of C, and C on that of A. Hence inferring A would block the inference of C, allowing the inference B, which would invalidate the inference of A, etc.

The examples presented so far have involved defaults in their simplest form:

$$\frac{\alpha : \beta_1 \wedge \cdots \wedge \beta_n}{\omega}$$

where α, ω and β_i are all literals (i.e. atomic formulae or negations of atomic formulae). The problem of determining dependencies is more complicated when α, ω and β_i are allowed to be arbitrary first-order formulae. For example, the consequent of a default may be an implication; applying that default would introduce new dependencies. The essential idea remains the same, however: determine whether the dependencies involve potentially unresolvable circularities. The following definitions outline a syntactic method for determining whether such circularities exist within a semi-normal theory.

Definition of $\ll$ and $\leqq\!\!\!\leqq$

Let $\Delta = (D,W)$ be a closed,[4] semi-normal default theory. Without loss of generality, assume all formulae are in clausal form. The partial relations, $\leqq\!\!\!\leqq$ and $\ll$, on *Literals* $\times$ *Literals*, are defined as follows:

(1) If $\alpha \in W$ then $\alpha = (\alpha_1 \vee \cdots \vee \alpha_n)$, for some $n \geq 1$.
For all $\alpha_i, \alpha_j \in \{\alpha_1, \ldots, \alpha_n\}$, if $\alpha_i \neq \alpha_j$ let $\neg\alpha_i \leqq\!\!\!\leqq \alpha_j$.

(2) If $\delta \in D$ then $\delta = \dfrac{\alpha : \beta \wedge \gamma}{\beta}$. Let $\alpha_1, \ldots \alpha_r$, $\beta_1, \ldots \beta_s$, and $\gamma_1, \ldots \gamma_t$ be the literals of the clausal forms of α, β, and γ, respectively. Then
(i) If $\alpha_i \in \{\alpha_1, \ldots, \alpha_r\}$ and $\beta_j \in \{\beta_1, \ldots, \beta_s\}$ let $\alpha_i \leqq\!\!\!\leqq \beta_j$.
(ii) If $\gamma_i \in \{\gamma_1, \ldots, \gamma_t\}$, $\beta_j \in \{\beta_1, \ldots, \beta_s\}$ and $\gamma_i \notin \{\beta_1, \ldots, \beta_s\}$ let $\neg\gamma_i \ll \beta_j$.
(iii) Also, $\beta = \beta_1 \wedge \cdots \wedge \beta_m$, for some $m \geq 1$.
For each $i \leq m$, $\beta_i = (\beta_{i,1} \vee \cdots \vee \beta_{i,m_i})$, where $m_i \geq 1$.
Thus if $\beta_{i,j}, \beta_{i,k} \in \{\beta_{1,1}, \ldots, \beta_{m,m_m}\}$ and $\beta_{i,j} \neq \beta_{i,k}$ let $\neg\beta_{i,j} \leqq\!\!\!\leqq \beta_{i,k}$.

(3) The expected transitivity relationships hold for $\ll$ and $\leqq\!\!\!\leqq$. i.e.
(i) If $\alpha \leqq\!\!\!\leqq \beta$ and $\beta \leqq\!\!\!\leqq \gamma$ then $\alpha \leqq\!\!\!\leqq \gamma$.
(ii) If $\alpha \ll \beta$ and $\beta \ll \gamma$ then $\alpha \ll \gamma$.
(iii) If $\alpha \ll \beta$ and $\beta \leqq\!\!\!\leqq \gamma$ or $\alpha \leqq\!\!\!\leqq \beta$ and $\beta \ll \gamma$ then $\alpha \ll \gamma$.

---

[4] The definition is readily extensible to open theories using a technique given in [Reiter 1980].

The definition is complex, but the intention is that $\alpha \underline{\ll} \beta$ or $\alpha \ll \beta$ if there is any way that α could figure in an inference of β in the theory as it stands. The intuition behind parts (1) and (2.iii) is that any disjunction of n literals can be interpreted as an implication of any one of those literals. e.g. $(\alpha_1 \vee \cdots \vee \alpha_n) \equiv [(\neg\alpha_1 \wedge \cdots \wedge \neg\alpha_{j-1} \wedge \neg\alpha_{j+1} \wedge \cdots \wedge \neg\alpha_n) \supset \alpha_j\ ]$. The special prominence we have alluded to for the conjuncts in a justification not entailed by the consequent is reflected in part (2.ii) by the use of the distinguished " $\ll$ " relation. The negation, $\neg\gamma_i$, occurs in part (2.ii) since it is not knowing $\neg\gamma_i$ which makes γ_i consistent.

Definition of Orderedness

A semi-normal default theory is said to be *ordered* if and only if there is no literal, α, such that $\alpha \ll \alpha$.

An ordered theory has no potentially unresolvable circular dependencies. The theory in example (1) is not ordered, since $B \ll A$, $C \ll B$, and $A \ll C$; hence $A \ll A$. The theory:

$$W = \{\ \}$$
$$D = \{\frac{: A \wedge \neg B}{A},\ \frac{: B \wedge \neg D}{B},\ \frac{: (C \supset D) \wedge \neg A}{(C \supset D)}\} \qquad (2)$$

is also not ordered. The defaults give rise to the following relationships:

$$\{B \ll A\},\quad \{D \ll B\},\quad \text{and}\quad \{C \underline{\ll} D,\ \ \neg D \underline{\ll} \neg C,\ \ A \ll \neg C,\ \ A \ll D\},$$

respectively. Hence $A \ll D \ll B \ll A$.

The significance of orderedness for semi-normal default theories is shown by Theorem 1.

Theorem 1

If a semi-normal default theory is ordered, then it has at least one extension. [5]

Normal theories are clearly ordered, since only non-normal defaults give rise to " $\ll$ " relationships. Thus the coherence of all normal theories is a corollary of Theorem 1. This is encouraging inasmuch as it suggests that orderedness is not merely a special purpose gimmick but that it subsumes an existing, widely applicable characterization.

It is important to notice that orderedness is only a sufficient condition for existence of extensions. Non-ordered theories have potentially unresolvable circularities but, for

---

[5] The proofs of this and all other results have been relegated to Appendix I for the sake of continuity.

one reason or another, these circularities do not always interfere. The theory (2) is not ordered, but it does have an extension: $Th(\{B, (C \supset D)\})$. The circularity would cause problems, however, if C were added to W: the resulting theory has no extensions. In other cases, two or more potential circularities may cancel each other out. At present, we do not know whether the given condition can be strengthened to one which is both necessary and sufficient for the coherence of semi-normal theories and yet is still decidable.

5. Constructing Extensions

Having delineated a large class of theories which have extensions, we turn to the problem of generating extensions. Reiter [1980] shows that extensions need not be recursively enumerable, and that it is not generally semi-decidable whether a formula is in any extension of a theory. Faced with such pessimism, further exploration might seem pointless. Still, there are tractable subcases.

Etherington [1982] has developed a procedure which can generate all the extensions of an arbitrary finite default theory.[6] The procedure centres on a relaxation style constraint propagation technique. Extensions are constructed by a series of successive approximations. Each approximation, H_j, is built up from the first-order components in W by applying defaults, one at a time. At each step, the default to be applied is chosen from those, not yet applied, whose prerequisites are "known" and whose justifications are consistent with both the previous approximation and the current state of the current approximation. When no more defaults are applicable, the procedure continues with the next approximation. If two successive approximations are the same, the procedure is said to *converge*.

The choice of which default to apply at each step of the inner loop may introduce a degree of non-determinism. Generality requires this non-determinism, however, since theories do not necessarily have unique extensions. Deterministic procedures can be constructed for theories which have unique extensions, or if full generality is not required.

In the presentation of the procedure, below, $CONSEQUENT(\frac{\alpha : \beta}{\gamma})$ is defined to be γ.

---

[6] A finite theory is one with only finitely many variables, constant symbols, predicate letters, and defaults. No function symbols are allowed, except of course the 0-ary function symbols, the constants. These restrictions make the universe of discourse (or Herbrand Universe) finite, ensuring only a finite number of closed instances of open defaults.

$H_0 \leftarrow W$; $j \leftarrow 0$;
repeat
  $j \leftarrow j + 1$; $h_0 \leftarrow W$; $GD_0 \leftarrow \{\ \}$; $i \leftarrow 0$;
  repeat
    $D_i \leftarrow \{ \frac{\alpha : \beta}{\gamma} \in D \mid (h_i \vdash \alpha), (h_i \not\vdash \neg\beta), (H_{j-1} \not\vdash \neg\beta) \}$;
    if $\neg null(D_i - GD_i)$ **then**
      choose δ **from** $(D_i - GD_i)$;
      $GD_{i+1} \leftarrow GD_i \cup \{\delta\}$;
      $h_{i+1} \leftarrow h_i \cup \{CONSEQUENT(\delta)\}$; **endif**;
    $i \leftarrow i + 1$;
  until $null(D_{i-1} - GD_{i-1})$;
  $H_j = h_{i-1}$
until $H_j = H_{j-1}$

To see how this procedure works, consider the theory:

$$W = \{A\}$$

$$D = \{\frac{A : B}{B}, \frac{A : C}{C}, \frac{B : D}{D}, \frac{B : \neg D \wedge \neg C}{\neg D}\},$$

which has the unique extension, $Th(\{A,B,C,D\})$. The procedure can generate any of the following sequences of approximations:

$H_0=\{A\}$		
$H_1=\{A,B,\neg D,C\}$	$H_0=\{A\}$	$H_0=\{A\}$
$H_2=\{A,B,C\}$	$H_1=\{A,C,B,D\}$	$H_1=\{A,B,C,D\}$
$H_3=\{A,B,D,C\}$	$H_2=H_1$	$H_2=H_1$
$H_4=H_3$		

(The formulae in each approximation are listed in the order in which they are derived.) In the first sequence of approximations, $\neg D$ occurs in H_1 because it can be inferred in h_2 *before* C is inferred in h_3.

Etherington [1982] proves:

There is a converging computation such that $H_n = H_{n-1}$ and $Th(H_n) = E$ if and only if E is an extension for the default theory (D, W).

In other words, the procedure can return every extension, and only extensions are returned. This result falls short in two respects: First, while the procedure can converge on every extension, there are appeals to ***non-provability***. In general, such tests are not computable, since arbitrary first-order formulae are involved. There are computable subcases, however. If the set

$$W \cup \{\alpha \mid \frac{\alpha : \beta}{\gamma} \in D\} \cup \{\beta \mid \frac{\alpha : \beta}{\gamma} \in D\}$$

belongs to a decision class for first-order provability, extensions are computable. Propositional and Monadic theories fall into this class, as do finite theories, provided W is also finite.

The second shortcoming is that some finite theories admit non-converging computations. The procedure may never terminate even though the theory has an extension and each step is computable. In such cases, the procedure cycles forever between two or more distinct H_j's. Fortunately this cyclic behaviour seems to be caused by features similar to those which make theories incoherent. We have characterized certain classes of ordered theories for which the procedure is more well-behaved.

Theorem 2 shows that one such class is the class of ordered, hierarchical theories. A default theory, $\Delta = (D, W)$, is *hierarchical* if it satisfies the following conditions:

(1) W contains only:
   a) Literals (i.e. Atomic formulae or their negations), and
   b) Disjuncts of the form $(\alpha \vee \beta)$ where α and β are literals.

(2) D contains only normal and semi-normal defaults of the form:

$$\frac{\alpha : \beta}{\beta} \quad \text{or} \quad \frac{\alpha : \beta \wedge \gamma_1 \wedge \cdots \wedge \gamma_n}{\beta}$$

where α, β, and γ_i are literals.

Theorem 2

For finite, ordered, hierarchical theories, the procedure given above always converges on an extension.

We will have more to say about hierarchical theories in the next section.

We conjecture that Theorem 2 can be generalized to apply to arbitrary ordered semi-normal theories, but we have no proof. The proof may require a more restrictive definition of D_i in the procedure, viz

$$D_i = \{\frac{\alpha : \beta}{\gamma} \in D \mid \alpha \in h_i, (h_i \cup H_{j-1}) \not\vdash \neg\beta\}$$

instead of

$$D_i = \{\frac{\alpha : \beta}{\gamma} \in D \mid \alpha \in h_i, h_i \not\vdash \neg\beta, H_{j-1} \not\vdash \neg\beta\}$$

but it can be shown that all the results of [Etherington 1982] and the current paper still hold for the stronger version, so this should present no problems.

For normal theories, an even stronger result can be proved:

Theorem 3

For finite normal theories, the procedure given above always converges on an extension immediately — i.e. $Th(H_1)$ *is always an extension.*

With these tools, we turn our attention to our stated goal: formalizing the semantics of non-monotonic inference systems.

6. Inheritance Hierarchies with Exceptions

To illustrate the suitability of Default Logic for formalizing the semantics of non-monotonic systems, we consider an application familiar to most of the AI community. Semantic networks have been widely adopted as a representational mechanism for AI. In such networks, "inference" is equated with inheritance of properties by nodes from their superiors. Recent work has considered the effects of allowing exceptions to inheritance within networks [Brachman 1982; Etherington and Reiter 1983; Fahlman 1979; Fahlman et al 1981; Touretzky 1982; Winograd 1980]. Such exceptions represent either explicit or implicit cancellation of the normal property inheritance which networks enjoy.

In the absence of exceptions, an inheritance hierarchy is a taxonomy organized by the usual IS-A relation, as in Figure 1.

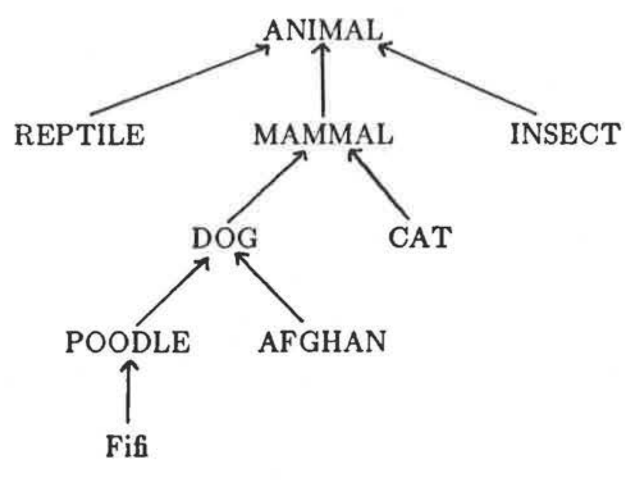

Figure 1 — Fragment of a taxonomy.

Schubert [1976] and Hayes [1977] have argued that inheritance hierarchies correspond quite naturally to certain theories of first-order logic. e.g.

$POODLE(Fifi)$
$\forall x.\ POODLE(x) \supset DOG(x)$
$\forall x.\ DOG(x) \supset MAMMAL(x)$
$\forall x.\ MAMMAL(x) \supset ANIMAL(x)$
etc.

Such a correspondence can be viewed as providing the semantics which "semantic" networks had previously lacked [Woods 1975]. The significant features of this semantics are these:

(1) Inheritance is a logical property of the representation. Given that $POODLE(Fifi)$, $MAMMAL(Fifi)$ is provable from the given formulae. Inheritance is the repeated application of modus ponens.
(2) The node labels of such a hierarchy are unary predicates: e.g. $DOG(*)$, $ANIMAL(*)$.
(3) No exceptions to inheritance are possible. If Fifi is a poodle, she must be an animal, regardless of any other properties she enjoys.

Unfortunately, this correspondence no longer applies when exceptions to inheritance are allowed. The logical properties of hierarchies change drastically when exceptions are permitted. For example, consider the following facts about elephants:

(1) Elephants are gray, except for albino elephants.
(2) All albino elephants are elephants.

Common-sense reasoning about "elephants" allows one, given an individual elephant not known to be an albino, to infer that she is gray. Subsequent discovery — perhaps by observation — that she is an albino elephant forces the retraction of the conclusion about her grayness. Thus, common-sense reasoning about exceptions is non-monotonic, in the sense that new information can invalidate previously derived facts. This non-monotonicity precludes the use of first-order representations, like those used for taxonomies, for formalizing hierarchies with exceptions.

We establish a correspondence between hierarchies with exceptions and hierarchical default theories. This correspondence provides a formal semantics and a notion of correct inference for such networks. As was the case for IS-A hierarchies, inheritance will emerge as a logical feature of the representation. Those properties $P_1, \cdots, P_n$ which an individual, b, inherits will be precisely those for which $P_1(b), \cdots, P_n(b)$ all belong to a common extension of the default theory. Should the theory have multiple extensions — an undesirable feature, as we shall see — then b may inherit different sets of properties depending on which extension is chosen.

To see how defaults might be used to represent hierarchies with exceptions, consider the elephant example, which can be represented by the default theory:

$W = \{\forall x.\ Albino\text{-}Elephant(x) \supset Elephant(x)\}$

$$D = \{ \frac{Elephant(x) : Gray(x) \wedge \neg Albino\text{-}Elephant(x)}{Gray(x)} \}.$$

It is easy to see that if we are told only *Elephant(Fred)* then, so far as we know, *Gray(Fred)* $\wedge$ ¬*Albino–Elephant(Fred)* is consistent; hence *Gray(Fred)* may be inferred. Given only *Albino–Elephant(Sue)* one can conclude *Elephant(Sue)* using first-order knowledge, but *Albino–Elephant(Sue)* "blocks" the application of the default, preventing the derivation of *Gray(Sue)*, as required.

We adopt a network representation with five link types. Other approaches to inheritance may omit one or more of these, but our formalism subsumes these. The five link types,[7] with their translations to default logic, are:

(1) Strict IS-A: A.———▶.B: A's are always B's. Since this is universally true, we identify it with the first order formula: $\forall x.\ A(x) \supset B(x)$.

(2) Strict ISN'T-A: A.┼┼┼┼┼┼▶.B: A's are never B's. Again, this is a universal statement, identified with: $\forall x.\ A(x) \supset \neg B(x)$.

(3) Default IS-A: A.———>.B: Normally A's are B's, but there may be exceptions. To provide for exceptions, we identify this with a default:

$$\frac{A(x) : B(x)}{B(x)}$$

(4) Default ISN'T-A: A.┼┼┼┼┼>.B: Normally A's are not B's, but exceptions are allowed. Identified with:

$$\frac{A(x) : \neg B(x)}{\neg B(x)}$$

(5) Exception: A.- - - - - ->
The exception link has no independent semantics; it serves only to make explicit the exceptions, if any, to the above default links. There must always be a default link at the head of an exception link; the exception then alters the semantics of that default link. There are two types of default links with exceptions; their graphical structures and translations are shown in Figure 2.

[7] Note that strict and default links are distinguished by solid and open arrowheads, respectively.

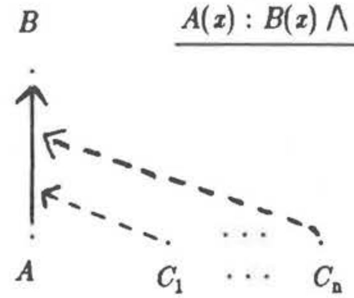

$$\frac{A(x) : B(x) \wedge \neg C_1(x) \wedge \cdots \wedge \neg C_n(x)}{B(x)}$$

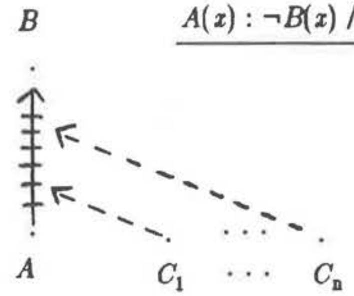

$$\frac{A(x) : \neg B(x) \wedge \neg C_1(x) \wedge \cdots \wedge \neg C_n(x)}{\neg B(x)}$$

Figure 2 — Links with exceptions.

We illustrate with an example from [Fahlman et al 1981].

Molluscs are normally shell-bearers.
Cephalopods must be Molluscs but normally are *not* shell-bearers.
Nautili must be Cephalopods and must be shell-bearers.

Our network representation of these facts is given in Figure 3.

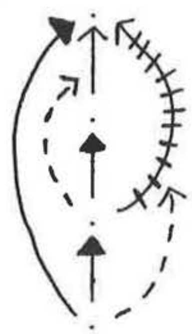

Figure 3 — Network representation of our knowledge about Molluscs.

The corresponding default theory is:

$$\{ \frac{M(x) : Sb(x) \mathbin{\&} \neg C(x)}{Sb(x)}, \ (x).C(x) \supset M(x), \ (x).N(x) \supset C(x),$$

$$\frac{C(x) : \neg Sb(x) \mathbin{\&} \neg N(x)}{\neg Sb(x)}, \ (x).N(x) \supset Sb(x) \}.$$

Given a particular Nautilus, this theory has a unique extension in which it is also a Cephalopod, a Mollusc, and a Shell-bearer. A Cephalopod not known to be a Nautilus will turn out to be a Mollusc with no shell.

It is instructive to compare our network representations with those of NETL [Fahlman et al 1981]. A basic difference is that in NETL there are no strict links; all IS-A and ISN'T-A links are potentially cancellable and hence are defaults. Moreover, NETL allows exception (*UNCANCEL) links only for ISN'T-A (*CANCEL) links. If we restrict the graph of Figure 3 to NETL-like links, we get Figure 4, which is essentially the graph given by Fahlman.

Figure 4 — NETL-like network representation of our knowledge about Molluscs.

This network corresponds to the theory:

$$\{\frac{M(x) : Sb(x)}{Sb(x)}, \frac{C(x) : M(x)}{M(x)}, \frac{N(x) : C(x)}{C(x)},$$

$$\frac{C(x) : \neg Sb(x) \,\&\, \neg N(x)}{\neg Sb(x)}, \frac{N(x) : Sb(x)}{Sb(x)}\}.$$

As before, a given Nautilus will also be a Cephalopod, a Mollusc, and a Shell-bearer. A Cephalopod not known to be a Nautilus, however, gives rise to *two* extensions, corresponding to an ambivalence about whether or not it has a shell. While counter-intuitive, this merely indicates that an exception to shell-bearing, namely being a Cephalopod, has not been explicitly represented in the network. Default Logic resolves the ambiguity by making the exception explicit, as in Figure 3. NETL, on the other hand, cannot make this exception explicit in the graphical representation, since it does not permit exception links to point to IS-A links.

How then does NETL conclude that a Cephalopod is not a Shell-bearer, without also concluding that it is a Shell-bearer? NETL resolves such ambiguities by means of an inference procedure which prefers shortest paths. Interpreted in terms of Default Logic, this "shortest path heuristic" is intended to favour one extension of the default theory. Thus, in the example above, the path from Cephalopod to ¬Shell-bearer is shorter than that to Shell-bearer so that, for NETL, the former wins. Unfortunately, this heuristic is not sufficient to replace the excluded exception type in all cases. Reiter and Criscuolo [1983] and Etherington [1982] show that it can lead to conclusions which are unintuitive or even invalid — i.e. not in any extension. Fahlman et al [1981] and Touretzky [1981, 1982] have also observed that shortest path algorithms can lead to anomalous conclusions. They describe attempts to restrict the form of networks to exclude structures which admit such problems. One effect of these restrictions is to permit only networks whose corresponding default theories have unique extensions.

An inference algorithm for network structures is correct only if it can be shown to derive conclusions all of which lie within a single extension of the underlying default theory. This criterion rules out shortest path inference for unrestricted networks. This is unfortunate, since shortest path inference has been popular for its relative efficiency and ease of implementation. Our results are not entirely pessimistic, however. Any network constructed using the five link-types given above corresponds to a hierarchical default theory. By insisting that any network constructed must correspond to an ordered theory, the coherence of a network knowledge representation system can be assured. For such systems, the procedure given in Section 5 is a correct and always converging inference algorithm.

It turns out that orderedness can be assured without reference to the full complexity of the mechanism described in Section 4. It is easy to see that any acyclic hierarchy gives rise to an ordered theory. The same is true if only the subgraph consisting of all IS-A links and explicit exceptions thereto has no cycles involving at least one exception link, or if there are no explicit exceptions to IS-A links.

In addition to pointing out the inadequacies of shortest path inferencing and to providing sufficient conditions for coherence and a correct inference mechanism, the semantics we have presented clarifies some of the outstanding problems in network inference. One of these, how to perform inferences in parallel, is considered in the next section.

7. Parallel Network Inference Algorithms

The computational complexity of inheritance problems, combined with some encouraging examples, has sparked interest in the possibility of determining inheritance in parallel. Fahlman [1979] has proposed a massively parallel machine architecture, NETL. NETL assigns one processor to each predicate in the knowledge base. "Inferencing" is performed by nodes passing "markers" to adjacent nodes in response to their own state and that of their immediate neighbours. Fahlman suggests that such architectures could achieve logarithmic speed improvements over traditional serial machines.

The formalization of inheritance hierarchies as default theories suggests, however, that there might be severe limitations to this approach. For example, correct inference requires that all conclusions share a common extension. For networks with more than one extension, inter-extension interference effects must be prevented. This seems impossible for a one pass parallel algorithm with purely local communication, especially in view of the inadequacies of the shortest path heuristic.

Even in knowledge bases with unique extensions, structures requiring an arbitrarily large radius of communication can be created. For example, the default theories corresponding to the networks in Figure 5 each have unique extensions. A network inference algorithm must reach F before propagating through B in the first network and

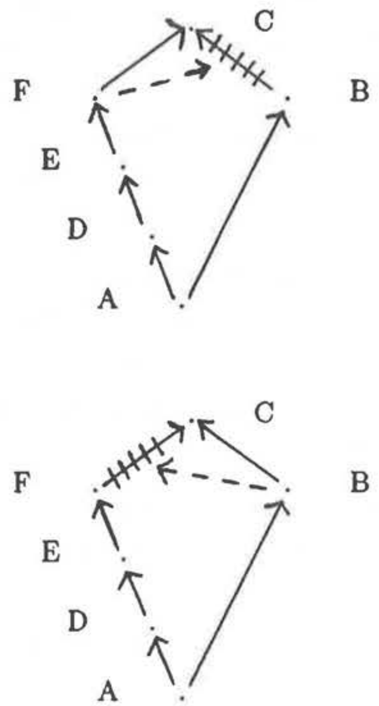

Figure 5 — Problems for local inheritance algorithms.

conversely in the second. The salient distinctions between the two networks are not local; hence they cannot be utilized to guide a purely local inference mechanism to the correct choices. Similar networks can be constructed which defeat marker-passing algorithms with any fixed radius.

This has prompted Touretzky [1981] to characterize a restricted class of network structures which admit parallel inferencing algorithms. In part, his restrictions appear to exclude networks whose corresponding default theory has more than one extension. Unfortunately, it is unclear how these restrictions affect the expressive power of the resulting networks. Moreover, Touretzky has observed that it is not possible to determine in parallel whether a network satisfies these restrictions.

Provided the network in question corresponds to an ordered theory, a form of limited parallelism can be achieved without sacrificing correctness. The key to this result lies in partitioning the network into subnetworks which are suitable for parallel processing. Essentially, each node in the network is numbered according to the number of exception links apon which it depends. This assigns each node to the lowest "level" possible while preserving the ordering amongst the nodes induced by the " $\ll$ " and " $\underline{\ll}$ " relations. Since the network is ordered, this can be done in parallel, in finite time proportional to the longest chain in the network. Processing then proceeds in k parallel steps, where k is the number of the highest level to which nodes were assigned. At step n, all links having exceptions which were asserted at step n–1 are disabled. The resulting sub-network, consisting of all remaining links impinging on nodes at levels less than or equal to n, is processed in parallel, ignoring exception links, with markers propagating from nodes asserted at step n–1. The "nodes asserted at level 0" are those in $Th(W)$. These correspond to the nodes for which the network is "activated". The result after step k is an extension.[8]

There are two caveats associated with this procedure: If both positive and negative markers reach a node in the same step, one must be chosen. Either choice will lead to an extension; we do not consider other ramifications of such choices here. Second, the algorithm assumes that all strict links propagate instantaneously. If this is not the case, each step in the algorithm must be followed by propagation along strict links, resolving conflicts as above. Note that conflicts are always resolved by changing assignments at the current level.

Provided that the inviolability of strict links is maintained, that default links are active only if their prerequisites are asserted and their justifications have not been denied, and that no node and its negation are asserted together (conflict resolution), *any* reasonable propagation algorithm (parallel or otherwise) may be used at each step.[9]

To illustrate the construction, we apply it to the moderately complex network of Figure 6. Rather than restrict ourselves to a particular parallel propagation algorithm at each step, we present a table showing all possibilities.

---

[8] This construction is that used in the proof of Theorem 1, where it is shown to yield an extension.

[9] To see this, it is necessary only to note that each step is, effectively, dealing with a normal theory. Arguments similar to those used in the proof of Theorem 3 can be used to show that the order of propagation is immaterial.

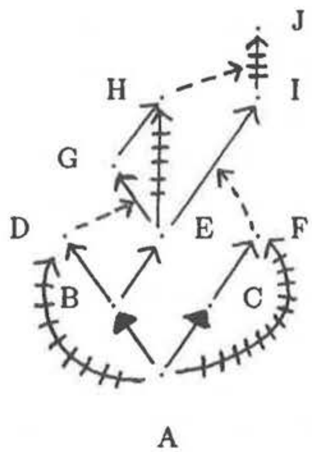

Figure 6— A multi-level inheritance graph.

The corresponding default theory, simplified to the propositional case and "activated" for A, is:

$$W = \{A, (A \supset B), (A \supset C)\}$$

$$D = \{\frac{A : \neg D}{\neg D}, \frac{A : \neg F}{\neg F}, \frac{B : D}{D}, \frac{C : F}{F}, \frac{B : E}{E}, \frac{E : \neg H}{\neg H},$$

$$\frac{E : G \wedge \neg D}{G}, \frac{G : H}{H}, \frac{E : I \wedge \neg F}{I},$$

$$\frac{I : \neg J \wedge \neg H}{\neg J}\}$$

The defaults above have been grouped according to the level to which their consequents are assigned (see Table 1). Table 2 shows the possibilities at each step; alternatives are

Level	Literals
1	A, B, C, D, ¬D, E, F, ¬F, ¬H
2	G, H, I
3	¬J

Table 1 — Levels of literals.

shown in separate columns, with major rows corresponding to steps in the algorithm.

Step 1	A, B, C, E, ¬H			
	D, F	D, ¬F	¬D, F	¬D, ¬F
Step 2		I	G	I, G
Step 3		¬J		¬J

Table 2 — Possible outcomes using different propagation schemes.

Thus the algorithm can, depending on the nature of the parallel marker propagation procedure, find:

$$E_0 = Th(W \cup \{A, B, C, E, \neg H, D, F\})$$
$$E_1 = Th(W \cup \{A, B, C, E, \neg H, D, \neg F, I, \neg J\})$$
$$E_2 = Th(W \cup \{A, B, C, E, \neg H, \neg D, F, G\})$$
$$E_3 = Th(W \cup \{A, B, C, E, \neg H, \neg D, \neg F, I, G, \neg J\})$$

all of which are extensions. Significantly, no choice of parallel procedure will enable the algorithm to find the theory's other four extensions:

$$E_4 = Th(W \cup \{A, B, C, E, H, D, F\})$$
$$E_5 = Th(W \cup \{A, B, C, E, H, D, \neg F, I\})$$
$$E_6 = Th(W \cup \{A, B, C, E, H, \neg D, F, G\})$$
$$E_7 = Th(W \cup \{A, B, C, E, H, \neg D, \neg F, G, I\})$$

because $\neg H$ is at level 1 and so can (and must) be inferred at step 1. H, being at level 2, is thus precluded before it can be inferred. We have not yet characterized the biases which this inability to find all extensions would induce in a reasoner.

Another potential problem with this approach stems from the fact that many network inference systems "prefer" one link-type over another (e.g. negation may override assertion). By breaking the network into sub-networks which are processed in turn, the ability to globally assert these preferences may be lost. We have three responses to this. First, many of these preferences are not well-defined, and break down when pressed (c.f. race conditions in [Fahlman et al 1981]). The inability to exhibit incorrect behaviour can hardly be called a liability. Second, given a well-defined preference scheme, it must preserve correctness: all inferences must lie in a single extension. If such a scheme exists which cannot be implemented within the confines outlined above, some other inference procedure will be required. Given the problems already observed with parallelism, we doubt that a parallel or quasi-parallel marker-passing algorithm can be found which takes global considerations into account (at least in unrestricted networks). Finally, if network structure is restricted, in the manner suggested by Touretzky [1981], so that resulting theories have unique extensions, the above algorithm produces the same results as any correct procedure.

1. Truth Maintenance

Another candidate for formalization using Default Logic is what has been variously called "Truth Maintenance" [Doyle, 1979; McAllester, 1978], "Tenability Maintenance" [McDermott & Doyle, 1980], "Belief Revision" [Doyle & London, 1980], and "Reason Maintenance" [Doyle, 1982a, 1982b]. A Truth Maintenance System (TMS) maintains a database of facts and justifications in such a way that its state agrees with the justifications. Each fact in the database is associated with a node in the TMS. The TMS tries to establish a maximal, contradiction free, set of nodes supported by the justifications. These nodes are made *in* and the corresponding fact is "believed"; the remaining nodes are *out* and not believed.

Like networks with exceptions and Default Logic, the TMS allows non-monotonicity. Justifications function similarly to defaults; they allow belief to be predicated on the lack of certain beliefs as well as on the acceptance of other beliefs. A justification for a node consists of two disjoint sets of nodes, the "inlist" and the "outlist". A node is brought in if and only if all the nodes on its inlist and none of those on its outlist are in. The truth maintenance process involves finding a well-founded set of nodes which satisfies the justifications. (The "well-founded" qualification rules out circularities, ensuring that no group of nodes is in, each for no other reason than because all the others are.) Whenever new facts or justifications are added to the knowledge base, the process is repeated.

Like many AI systems, TMS's have lacked rigorous semantic specifications, subsisting on intuitive descriptions and, sometimes, large program listings. The apparent correspondence of inlists and outlists to the prerequisites and justifications of defaults suggests that Default Logic can be used to provide a clear semantics for Truth Maintenance. Doyle himself [1982a, 1982b] has recently recognized the importance of "concise, exact specifications" for TMS's. He has addressed this problem using a formal system somewhat more abstract than Default Logic. While inconclusive, his results are encouraging. They have also highlighted facets of his TMS which could be improved.

A correspondence between truth maintenance and Default Logic should provide a number of benefits, including:

(1) A clearly defined, understandable, and correct TMS,

(2) A more efficient inference procedure for Default Logic, utilizing approporiate techniques from TMS's. For example, the extension construction procedure from Section 5 is superficially similar to the relaxation constraint propagation techniques used in some TMS's. Can better performance be obtained without sacrificing the procedure's generality?

(3) A better understanding of how default theories evolve when new information is added. The primary function of a TMS is to maintain a database as information is added over time. Little consideration has been given to the effects of adding information to a default theory.

9. Questions and Conclusions

We have suggested that Default Logic is an excellent tool for formalizing the reasoning processes involved in AI systems. Our intention in doing so has not been to diminish such systems in any way, but rather to provide a metric by which they can be measured and compared. A Default Logic specification of a system's semantics can provide both a more complete visualization of how the system performs and a guarantee that that performance is coherent.

To facilitate such applications, we have presented new results on Default Logic. We have characterized a large class of theories for which coherent reasoning is always possible (i.e. theories which always have at least one extension), and provided a totally correct inference algorithm for a subclass of these theories.

It might be — and has been — argued that a declarative formalism such as Default Logic is inadequate for the tasks of knowledge representation and reasoning. While we clearly disagree with this position, we expect Default Logic to be useful even to "proceduralists". Even if some system were fundamentally more than the sum of its declarative content, Default Logic could be used to formalize that declarative content. The non-declarative "control" information could then be treated as an inference algorithm for the resulting default theory. The correctness of the system would be determined by whether this inference algorithm was correct with respect to the proof theory of Default Logic.

By formalizing inheritance hierarchies with exceptions using Default Logic we have provided them with a precise semantics. This in turn allowed us to identify the notion of correct inference in such a hierarchy with that of derivability within a single extension of the corresponding default theory. We provided an inference algorithm for ordered inheritance hierarchies with exceptions which is provably correct with respect to this concept of derivability.

Our formalization suggests that *for unrestricted hierarchies,* it may not be possible to realize massively parallel marker-passing hardware of the kind envisaged by NETL. It appears that the best that can be achieved for such hierarchies is a restricted, quasi-parallel inference algorithm. We have sketched such an algorithm, but have shown that not every set of conclusions justified by the hierarchy is accessible to it. It remains to be seen whether the limitations imposed by the algorithm are acceptable. Fortunately, these pessimistic observations do not preclude parallel architectures for suitably restricted hierarchies.

There are a number of open problems:

(1) Is there a natural class of inheritance hierarchies with exceptions which admits a parallel inference algorithm yet does not preclude the representation of our common-sense knowledge about taxonomies?

(2) Define such a parallel algorithm and prove its correctness with respect to the derivability relation of Default Logic.

(3) In connection with (1), notice that it is natural to restrict attention to those hierarchies whose corresponding default theories have unique extensions. Characterize such hierarchies.

(4) In spite of its demonstrated shortcomings, the shortest path heuristic continues to be widely used. Is this because it approximates some common-sense inference mechanism for which no better characterization has been known? Touretzky [1983b] explores this hypothesis, and presents a refined heuristic which avoids some of the shortest path heuristic's failures. We have not attempted to analyze whether this new heuristic results in correct inferences, but we feel such analysis would be fruitful.

Defaults, in one form or another, are extremely common in AI. Reiter [1978, 1980] discusses a wide variety of common situations to which they can be applied, including several AI knowledge representation schemes. We expect that many of these may be amenable to analysis using the approach we have outlined. If some are not, two interesting possibilities arise: the features not so amenable may prove incorrect or inessential, or they may point out shortcomings of Default Logic. One hopes that either result will lead to progress in the field.

10. Acknowledgments

I would like to thank Raymond Reiter and Robert Mercer for their insightful listening, without which I probably would not have bothered, and for helpful comments on earlier drafts of this paper. Bob also went to the trouble of reading and *understanding* all the proofs, and then made me make them understandable. David Touretzky has done much to improve my understanding of networks and NETL. Thanks also to Akira Kanda for financing this report, and to UBC's Laboratory for Computational Vision for access to their coveted typesetting facilities.

11. References

AI (1980), Special issue on non-monotonic logic, *Artificial Intelligence 13 (1,2)*, April.

Brachman, R. (1982), "What 'IS-A' Is and Isn't", *Proc. Canadian Soc. for Computational Studies of Intelligence-82,* Saskatoon, Sask., May 17-19, pp 212-220.

Doyle, J. (1979), *A Truth Maintenance System,* AI Memo 521, MIT, Cambridge, Mass.

Doyle, J., and London, P. (1980), *A Selected Descriptor-Indexed Bibliography to the Literature on Belief Revision,* AI Memo 568, MIT, Cambridge, Mass.

Etherington, D.W. (1982), *Finite Default Theories,* M.Sc. Thesis, Dept. Computer Science, University of British Columbia.

Etherington, D.W., and Reiter, R. (1983), "On Inheritance Hierarchies With Exceptions", *Proc. American Assoc. for Artificial Intelligence-83,* Washington, D.C., August 24-26, pp 104-108.

Fahlman, S.E. (1979), *NETL: A System For Representing and Using Real-World Knowledge,* MIT Press, Cambridge, Mass.

Fahlman, S.E. (1982), "Three Flavors of Parallelism", *Proc. Canadian Soc. for Computational Studies of Intelligence-82,* Saskatoon, Sask., May 17-19, pp 230-235.

Fahlman, S.E., Touretzky, D.S., and van Roggen, W. (1981), "Cancellation in a Parallel Semantic Network", *Proc. IJCAI-81,* Vancouver, B.C., Aug. 24-28, pp 257-263.

Hayes, P.J. (1977), "In Defense of Logic", *Proc. IJCAI-77,* Cambridge, Mass., pp 559-565.

McAllester, D.A. (1978), *A Three-Valued Truth Maintenance System,* AI Memo 473, MIT, Cambridge, Mass.

McDermott, D., and Doyle, J. (1980), "Non-Monotonic Logic I", *Artificial Intelligence 13,* (April) pp 41-72.

Mendelson, E. (1964), *Introduction to Mathematical Logic,* Van Nostrand Reinhold, New York.

Reiter, R. (1978), "On Reasoning by Default", *Proc. Second Symposium on Theoretical Issues in Natural Language Processing,* Urbanna, Illinois.

Reiter, R. (1980), "A Logic for Default Reasoning", *Artificial Intelligence 13,* (April) pp 81-132.

Reiter, R., and Criscuolo, G. (1983), "Some Representational Issues in Default Reasoning", *Int. J. Computers and Mathematics,* (Special Issue on Computational Linguistics), to appear.

Schubert, L.K. (1976), "Extending the Expressive Power of Semantic Networks", *Artificial Intelligence 7(2),* pp 163-198.

Touretzky, D.S. (1981), (1983a), Personal Communication.

Touretzky, D.S. (1982), "Exceptions in an Inheritance Hierarchy", Unpublished Manuscript, Department of Computer Science, Carnegie-Mellon University.

Touretzky, D.S. (1983b), "Multiple Inheritance and Exceptions", Unpublished Manuscript, Department of Computer Science, Carnegie-Mellon University.

Winograd, T. (1980), "Extended Inference Modes in Reasoning", *Artificial Intelligence 13*, (April), pp 5-26.

Woods, W.A. (1975), "What's In A Link?", *Representation and Understanding,* Academic Press, pp 35-82.

Appendix I — Proofs of Theorems

Background Information

There are a few definitions and results due to Reiter [1980] on which we draw freely in the following proofs. We reproduce them here for the reader's convenience.

1) Theorem 0.1 [Reiter 1980, Theorem 2.1]

E is an extension for $\Delta = (D, W)$ if and only if $E = \bigcup_{i=0}^{\infty} E_i$, where

$E_0 = W$, and for $i > 0$

$$E_{i+1} = Th(E_i) \cup \{\omega \mid \frac{\alpha : \beta}{\omega} \in D, \alpha \in E_i \text{, and } \neg\beta \notin E\}^1$$

2) The *Generating Defaults* for E with respect to Δ are defined as:

$$GD(E,\Delta) = \{\frac{\alpha : \beta}{\omega} \in D \mid \alpha \in E, \neg\beta \notin E\}$$

3) If D is a set of defaults, then *CONSEQUENTS*(D) is defined, as one would expect, as:

$$CONSEQUENTS(D) = \{\omega \mid \frac{\alpha : \beta}{\omega} \in D\}$$

4) Theorem 0.2 [Reiter 1980, Theorem 2.5]

If E is an extension for $\Delta = (D, W)$, then

$$E = Th(W \cup CONSEQUENTS(GD(E,\Delta))).$$

5) Theorem 0.3 [Reiter 1980, Corollary 2.2]

If E is an extension for $\Delta = (D, W)$, then E is consistent if and only if W is.

In the proofs which follow, we will usually assume that formulae are in clausal form: i.e. expressed as a conjunction of disjunctions of literals. We define the functions *CLAUSES*(·) and *LITERALS*(·) as follows:

If $\beta = (\beta_{1,1} \vee \cdots \vee \beta_{1,m_1}) \wedge \cdots \wedge (\beta_{m,1} \vee \cdots \vee \beta_{m,m_m})$ then

$$CLAUSES(\beta) = \{(\beta_{i,1} \vee \cdots \vee \beta_{i,m_i}) \mid 1 \leq i \leq m\}$$

$$LITERALS(\beta) = \{\beta_{i,j} \mid 1 \leq i \leq m, 1 \leq j \leq m_i\}$$

Abusing the notation somewhat we sometimes use *CLAUSES*(Γ), where Γ is a set of

[1] Note the explicit reference to E in the definition of E_{i+1}.

formulae, to refer to $\bigcup_{\gamma \in \Gamma} CLAUSES(\Gamma)$.

We will define other notation as it is required.

Lemma 1.1

If E^i $(i \geq 0)$ is an extension for the default theory $\Delta_i = (D_i, E^{i-1})$ and $E^{-1} = W$, then the following are equivalent:

(1) $\alpha \in E^i$

(2) $E^i \vdash \alpha$

(3) $(W \cup \bigcup_{r=0}^{i} CONSEQUENTS(GD(E^r, \Delta_r))) \vdash \alpha$

Proof

(1) $\alpha \in E^i \leftrightarrow E^i \vdash \alpha$

This follows from the fact that E^i is an extension and thus logically closed.

(2) $E^i \vdash \alpha \leftrightarrow (W \cup \bigcup_{r=0}^{i} CONSEQUENTS(GD(E^r, \Delta_r))) \vdash \alpha$

If E is an extension for Δ, then by Theorem 0.2 we know that

$$E = Th(W \cup CONSEQUENTS(GD(E, \Delta))).$$

Hence

$$\begin{aligned} E^i &= Th(E^{i-1} \cup CONSEQUENTS(GD(E^i, \Delta_i))) \\ &= Th(Th(E^{i-2} \cup CONSEQUENTS(GD(E^{i-1}, \Delta_{i-1}))) \\ &\qquad \cup CONSEQUENTS(GD(E^i, \Delta_i))) \\ &= Th(Th \cdots (W \cup CONSEQUENTS(GD(E^0, \Delta_0))) \\ &\qquad \cup \cdots \cup CONSEQUENTS(GD(E^i, \Delta_i))) \end{aligned}$$

Since $Th(Th(A) \cup B) = Th(A \cup B)$,

$$E^i = Th(W \cup \bigcup_{r=0}^{i} CONSEQUENTS(GD(E^r, \Delta_r))).$$

From this, the result follows by the definition of Th.

QED Lemma 1.1

Definition 1.2

Let $\Delta = (D,W)$ be a closed, semi-normal default theory. Without loss of generality, assume all formulae are in clausal form. The partial relations, $\lesssim$ and $\ll$, on *Literals* $\times$ *Literals*, are defined as follows:

(1) If $\alpha \in W$ then $\alpha = (\alpha_1 \vee \cdots \vee \alpha_n)$, for some $n \geq 1$.
For all $\alpha_i, \alpha_j \in \{\alpha_1, \ldots, \alpha_n\}$, if $\alpha_i \neq \alpha_j$ let $\neg\alpha_i \lesssim \alpha_j$.
(Since: $(\alpha_1 \vee \cdots \vee \alpha_n) \equiv [(\neg\alpha_1 \wedge \cdots \wedge \neg\alpha_{j-1} \wedge \neg\alpha_{j+1} \wedge \cdots \wedge \neg\alpha_n) \supset \alpha_j])$

(2) If $\delta \in D$ then $\delta = \dfrac{\alpha : \beta \wedge \gamma}{\beta}$. Let $\alpha_1, \ldots \alpha_r, \beta_1, \ldots \beta_s$, and $\gamma_1, \ldots \gamma_t$ be the literals of the clausal forms of α, β, and γ, respectively. Then

(i) If $\alpha_i \in \{\alpha_1, \ldots, \alpha_r\}$ and $\beta_j \in \{\beta_1, \ldots, \beta_s\}$ let $\alpha_i \lesssim \beta_j$.
(ii) If $\gamma_i \in \{\gamma_1, \ldots, \gamma_t\}$, $\beta_j \in \{\beta_1, \ldots, \beta_s\}$ and $\gamma_i \notin \{\beta_1, \ldots, \beta_s\}$ let $\neg\gamma_i \ll \beta_j$.
(iii) Also, $\beta = \beta_1 \wedge \cdots \wedge \beta_m$, for some $m \geq 1$.
For each $i \leq m$, $\beta_i = (\beta_{i,1} \vee \cdots \vee \beta_{i,m_i})$, where $m_i \geq 1$.
Thus if $\beta_{i,j}, \beta_{i,k} \in \{\beta_{1,1}, \ldots, \beta_{m,m_m}\}$ and $\beta_{i,j} \neq \beta_{i,k}$ let $\neg\beta_{i,j} \lesssim \beta_{i,k}$.

(3) The expected transitivity relationships hold for $\ll$ and $\lesssim$. i.e.
(i) If $\alpha \lesssim \beta$ and $\beta \lesssim \gamma$ then $\alpha \lesssim \gamma$.
(ii) If $\alpha \ll \beta$ and $\beta \ll \gamma$ then $\alpha \ll \gamma$.
(iii) If $\alpha \ll \beta$ and $\beta \lesssim \gamma$ or $\alpha \lesssim \beta$ and $\beta \ll \gamma$ then $\alpha \ll \gamma$.

Definition 1.3

A semi-normal default theory is said to be ***ordered*** iff there is no literal, α, such that $\alpha \ll \alpha$.

Definition 1.4

For a closed, semi-normal default theory, $\Delta = (D, W)$, define the ***Universe of*** Δ, $U(\Delta)$, as follows:

$$U(\Delta) = \{\alpha \mid \alpha \in \textit{Literals} \text{ and } [\exists\, \xi.\ [(\alpha \vee \xi) \in \textit{CLAUSES}(W \cup \textit{CONSEQUENTS}(D))] \text{ or } [(\neg\alpha \vee \xi) \in \textit{CLAUSES}(W \cup \textit{CONSEQUENTS}(D))]]\}$$
$$\cup\ \{\alpha_i \mid \exists\, \alpha,\beta,\gamma.\ \frac{\alpha : \beta}{\gamma} \in D \text{ and } \alpha_i \in \textit{LITERALS}(\alpha)\}$$

$$\cup \{\neg\gamma_i \mid \exists\ \alpha,\beta,\gamma.\ \frac{\alpha : \beta \wedge \gamma}{\beta} \in D \text{ and } \gamma_i \in \mathit{LITERALS}(\gamma)\ \}$$

Observe that ξ may be the null clause.

Definition 1.5

For a closed, ordered, semi-normal default theory, $\Delta = (D, W)$, we define the function $l : U(\Delta) \mapsto \mathbf{N}$, as follows:

If $\alpha,\beta \in U(\Delta)$ and $\alpha \underline{\ll} \beta$ then $l(\alpha) \leq l(\beta)$. If $\alpha \ll \beta$ then $l(\beta) \geq l(\alpha)+1$.

If $\beta \in U(\Delta)$ and for no $\alpha \in U(\Delta)$ is $(\alpha \ll \beta)$ or $(\alpha \underline{\ll} \beta)$ then $l(\beta) = 0$.

If $n \in \mathbf{N}$, $\beta \in U(\Delta)$, and $l(\beta) > n$ then $\exists\ \alpha \in U(\Delta).\ (\alpha \ll \beta)$ and $l(\alpha) = n$.

Since Δ is ordered, l is well defined. Observe that l is a total function on $U(\Delta)$ which assigns a natural number to each literal in $U(\Delta)$. $l(\alpha)$ may be thought of as the length of the longest chain of semi-normal defaults which could figure in an inference of α.

Definition 1.6

If β is a closed formula, and the clausal form of β is

$$(\beta_{1,1} \vee \cdots \vee \beta_{1,m_1}) \wedge \cdots \wedge (\beta_{m,1} \vee \cdots \vee \beta_{m,m_m}),$$

then define $l_{MAX}(\beta) \equiv \text{MAX}(l(\beta_{i,j}))$

$$l_{MIN}(\beta) \equiv \text{MIN}(l(\beta_{i,j}))$$

Lemma 1.7

If $\Delta = (D, W)$ is an ordered, closed, semi-normal default theory, then there is a partition, $\{D_i\}$, for D induced by:

$$\forall\, \delta \in D.\ \delta = \frac{\alpha : \beta \wedge \gamma}{\beta} \text{ and } l_{MIN}(\beta) = i \leftrightarrow \delta \in D_i\ .$$

Proof

Clearly *LITERALS* (*CONSEQUENTS* ($\{\delta \in D\}$)) $\subseteq U(\Delta)$, and l is total on $U(\Delta)$.

Therefore: 1) $\forall\, \delta \in D.\ \forall\, i.\ \forall\, j.\ (\delta \in D_i \wedge \delta \in D_j) \rightarrow i = j$.

2) $\forall\, \delta \in D.\ \exists\, i.\ (\delta \in D_i)$.

QED Lemma 1.7

Corollary 1.8

If $\delta \in D_0$, then δ is a normal default.

Proof

If $\delta = \frac{\alpha : \beta \wedge \gamma}{\beta} \in D_0$ then $l_{MIN}(\beta) > l_{MAX}(\neg\gamma) \geq 0$.

QED Corollary 1.8

Corollary 1.9

If $i > 0$ and $D_i \neq \{\ \}$, there is at least one non-normal (i.e. semi-normal) default in D_i.

Proof

If D_i contains only normal defaults, then the minimality of l guarantees that l_{MIN}(*CONSEQUENTS* (D_i)) $< i$, which is a contradiction.

QED Corollary 1.9

Lemma 1.10

If Γ is consistent, if $l_{MAX}(\neg\beta) < j$, and if $l(\gamma)$ is defined for all $\gamma \in LITERALS(\Gamma)$, then there is a linear resolution refutation of β from Γ if and only if there is a linear resolution refutation of β from Ψ, where $\Psi \subseteq \Gamma$ and $\psi \in \Psi \leftrightarrow l_{MIN}(\psi) < j$.

Proof

($\rightarrow$)

The proof is by construction of such a refutation.

Since Γ is consistent, if there is a refutation of β from Γ, there is a refutation with top clause in CLAUSES (β). i.e.

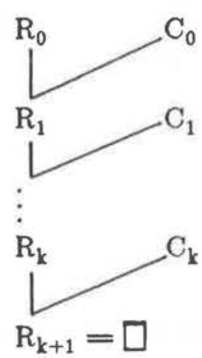

and $R_0 \in CLAUSES(\beta)$, $C_0 \in \Gamma$.

We proceed by induction on the steps in the refutation.

base

Assume β is in clausal form, i.e.

$$\beta = \beta_1 \wedge \cdots \wedge \beta_n \quad \text{and} \quad \beta_i = \beta_{i,1} \vee \cdots \vee \beta_{i,n_i} .$$

By hypothesis, $l(\neg\beta_{i,r}) < j$. Without loss of generality, assume that $R_0 = \beta_1 = \beta_{1,1} \vee \cdots \vee \beta_{1,n_1}$, that $C_0 = C_{0,1} \vee \cdots \vee C_{0,m_0}$, and that $C_{0,1}$ resolves on $\beta_{1,1}$ to produce R_1.

Thus $C_{0,1} = \neg\beta_{1,1}$ so $l(C_{0,1}) < j$ and $l_{MIN}(C_0) < j$. It follows that $C_0 \in \Psi$.

Since for i>1, $\neg C_{0,i} \lll C_{0,1}$, $l(\neg C_{0,i}) \leq l(C_{0,1}) < j$.

Thus, if $R_1 = R_{1,1} \vee \cdots \vee R_{1,t}$ then $\forall s.\ l(\neg R_{1,s}) < j$.

step

Assume that $R_i = R_{i,1} \vee \cdots \vee R_{i,m}$, that $\forall s.\ l(\neg R_{i,s}) < j$, and that $\forall r<i$. $C_r \in \Psi$ or $C_r \in \{R_0, \cdots, R_{i-1}\}$. Consider the resolution of R_i with C_i . $C_i = C_{i,1} \vee \cdots \vee C_{i,m_i}$. Without loss of generality, assume $C_{i,1} = \neg R_{i,1}$. Hence $l(C_{i,1}) = l(\neg R_{i,1}) < j$ and so $l_{MIN}(C_i) < j$. Thus $C_i \in \Psi$ or $C_i \in \{R_0, \cdots, R_i\}$. For $r>1$, $l(\neg C_{i,r}) \leq l(C_{i,1}) < j$. Thus $\forall s.\ l(\neg R_{i+1,s}) < j$.

By induction, for every clause, C_i, in the refutation of β, $C_i \in \Psi$ or C_i is a descendent of $\Psi \cup \{\beta\}$. Thus, there is a linear resolution refutation of β from Ψ.

($\leftarrow$)

Trivial: Since $\Psi \subseteq \Gamma$, the refutation from Ψ serves as a refutation from Γ.

QED lemma 1.10

Theorem 1

If $\Delta = (D, W)$ is an ordered, semi-normal default theory, then Δ has an extension.

Proof

If W is inconsistent, then Δ has the trivial extension, L . Hence assume W is consistent.

We proceed by constructing an extension, E for Δ. First, let $\{D_i\}$ be a partition of D induced by l, as described in Lemma 1.7. Recall that by Corollary 1.8, if $\delta \in D_0$ then δ is a normal default, and that by Corollary 1.9, for $i > 0$, D_i must contain at least one semi-normal default, say

$$\delta = \frac{\alpha : \beta \wedge \gamma}{\beta},$$

and $l_{MAX}(\neg\gamma) < l_{MIN}(\beta)$.

We now construct an extension for Δ.

Let $\Delta_0 = (D_0, W)$. Since Δ_0 is a normal default theory and W is consistent, Δ_0 has a consistent extension, say E^0.

For $i > 0$, construct Δ_i as follows:

$$D_i' = \{\frac{\alpha : \beta}{\beta} \mid \frac{\alpha : \beta}{\beta} \in D_i \vee \frac{\alpha : \beta \wedge \gamma}{\beta} \in D_i , \neg\gamma \notin E^{i-1}\}$$

$$\Delta_i = (D_i' , E^{i-1})$$

Where E^{i-1} is an extension for Δ_{i-1}. Since each Δ_i is a normal default theory, each Δ_i has at least one extension, E^i. Let $E = \bigcup_{i=0}^{\infty} E^i$. Since W is consistent, so is E^0, by Theorem 0.3. Since E^i is an extension for (D_i' , E^{i-1}), E^i is consistent if E^{i-1} is, and $E^{i-1} \subseteq E^i$. By induction E is consistent. We now show that E is an extension for Δ. By Theorem 0.1, it is sufficient to show that $E = \bigcup_{i=0}^{\infty} F_i$, where

$F_0 = W$, and for $i>0$

$F_{i+1} = Th(F_i) \cup \{\omega \mid \frac{\alpha : \beta}{\omega} \in D, \alpha \in F_i$, and $\neg\beta \notin E\}$.

(1) We first show that $\bigcup_{i=0}^{\infty} F_i \subseteq E$.

a) $F_0 = W \subseteq E^0 \subseteq E$.

b) Assume $F_i \subseteq E$. We show that $F_{i+1} \subseteq E$.

$F_{i+1} = Th(F_i) \cup \{\beta \mid \frac{\alpha : \beta \wedge \gamma}{\beta} \in D, \alpha \in F_i , (\neg\beta \vee \neg\gamma) \notin E\}$

i) Since $F_i \subseteq E$ and E is logically closed, $Th(F_i) \subseteq E$.

ii) Consider $\beta \in \{\beta \mid \frac{\alpha : \beta \wedge \gamma}{\beta} \in D, \alpha \in F_i , (\neg\beta \vee \neg\gamma) \notin E\}$.

Since $\alpha \in F_i$, $\alpha \in E$, and hence $\alpha \in E^j$ for some j .

Since $(\neg\beta \vee \neg\gamma) \notin E$, $\neg\gamma \notin E^{j-1}$, so $\frac{\alpha : \beta}{\beta} \in D_j'$.

But $\neg\beta \notin E$, so $\neg\beta \notin E^j$.

Therefore, since E^j is an extension for $\Delta_j = (D_j' , E^{j-1}))$ and $\alpha \in E^j$, $\beta \in E^j$.

Therefore $\beta \in E$.

By induction, $\bigcup_{i=0}^{\infty} F_i \subseteq E$.

(2) Finally, we show that $E \subseteq \bigcup_{i=0}^{\infty} F_i$.

A) Consider $\omega \in E^0$. E^0 is an extension for Δ_0, so by Theorem 0.1 $E^0 = \bigcup_{i=0}^{\infty} G_i$,

where

$G_0 = W$, and for i>0

$G_{i+1} = Th(G_i) \cup \{\omega \mid \frac{\alpha : \omega}{\omega} \in D_0 , \alpha \in G_i$, and $\neg\omega \notin E^0\}$.

It therefore suffices to show that $\bigcup_{i=0}^{\infty} G_i \subseteq \bigcup_{i=0}^{\infty} F_i$.

a) $G_0 = W = F_0 \subseteq \bigcup_{i=0}^{\infty} F_i$.

b) Assume $G_i \subseteq \bigcup_{i=0}^{\infty} F_i$, and consider $\omega \in G_{i+1}$.

$G_{i+1} = Th(G_i) \cup \{\omega \mid \frac{\alpha : \omega}{\omega} \in D_0 , \alpha \in G_i , \neg\omega \notin E^0\}$

i) If $\omega \in Th(G_i)$ then $\omega \in \bigcup_{i=0}^{\infty} F_i$ by hypothesis since $\bigcup_{i=0}^{\infty} F_i$ is logically closed.

ii) Otherwise $\omega \in \{\omega \mid \frac{\alpha : \omega}{\omega} \in D_0 , \alpha \in G_i , \neg\omega \notin E^0\}$.

But: 1) If $\omega \in G_{i+1}$ and $E^0 = \bigcup_{i=0}^{\infty} G_i$ then $\omega \in E^0 \subseteq E$.

Since E is consistent, $\neg\omega \notin E$.

2) If $\alpha \in G_i$ then $\alpha \in \bigcup_{i=0}^{\infty} F_i$ by hypothesis, so $\alpha \in F_k$ for some k.

3) $D_0 \subseteq D$

Thus $\omega \in F_{k+1} \subseteq \bigcup_{i=0}^{\infty} F_i$.

By induction, $\bigcup_{i=0}^{\infty} G_i \subseteq \bigcup_{i=0}^{\infty} F_i$.

B) Assume $E^{j-1} \subseteq \bigcup_{i=0}^{\infty} F_i$, and show $E^j \subseteq \bigcup_{i=0}^{\infty} F_i$.

Consider $\omega \in E^j$. E^j is an extension for $\Delta_j = (D_j{}', E^{j-1})$, so $E^j = \bigcup_{i=0}^{\infty} G_i$, where

$G_0 = E^{j-1}$, and for i>0

$G_{i+1} = Th(G_i) \cup \{\omega \mid \frac{\alpha : \omega}{\omega} \in D_j{}' , \alpha \in G_i$, and $\neg\omega \notin E^j\}$.

a) By hypothesis, $G_0 = E^{j-1} \subseteq \bigcup_{i=0}^{\infty} F_i$.

b) Assume $G_i \subseteq \bigcup_{i=0}^{\infty} F_i$ and consider $\omega \in G_{i+1}$.

i) If $\omega \in Th(G_i)$ then $\omega \in \bigcup_{i=0}^{\infty} F_i$ by hypothesis since $\bigcup_{i=0}^{\infty} F_i$ is logically closed.

ii) Otherwise $\omega \in \{\omega \mid \frac{\alpha : \omega}{\omega} \in D_j', \alpha \in G_i$, and $\neg\omega \notin E^j\}$.

Since $\alpha \in G_i$, we know that $\alpha \in E^j$ and $\alpha \in \bigcup_{i=0}^{\infty} F_i$. Also, if $\omega \in G_{i+1}$ then $\omega \in E^j$ so $\omega \in E$. Therefore $\neg\omega \notin E$, since E is consistent.

If $\delta = \frac{\alpha : \omega}{\omega} \in D_j'$ then either $\frac{\alpha : \omega}{\omega} \in D$ or $\exists\, \gamma.\ \frac{\alpha : \omega \wedge \gamma}{\omega} \in D$.

Thus there are two cases:

a) Either $\frac{\alpha : \omega}{\omega} \in D$, $\alpha \in \bigcup_{i=0}^{\infty} F_i$, and $\neg\omega \notin E$ and hence $\omega \in \bigcup_{i=0}^{\infty} F_i$,

b) Or $\frac{\alpha : \omega \wedge \gamma}{\omega} \in D$, $\alpha \in \bigcup_{i=0}^{\infty} F_i$, and $\neg\omega \notin E$.

Clearly, if $(\neg\gamma \vee \neg\omega) \notin E$ then $\omega \in \bigcup_{i=0}^{\infty} F_i$.

Since $\omega \in E$, it can be shown that $(\neg\gamma \vee \neg\omega) \in E \leftrightarrow \neg\gamma \in E$.

We show that $\neg\gamma \notin E$.

Clearly $l_{MAX}(\neg\gamma) < l_{MIN}(\omega) = j$.

Assume $\neg\gamma \in E$. Then $\exists\, r \geq j.\ (\neg\gamma \in E^r)$.

By Lemma 1.1, $(W \cup \bigcup_{i=0}^{r} \mathit{CONSEQUENTS}(GD(E^i, \Delta_i))) \vdash \neg\gamma$.

Thus there is a linear resolution refutation of γ from

$$\Gamma = (W \cup \bigcup_{i=0}^{r} \mathit{CONSEQUENTS}(GD(E^i, \Delta_i))).$$

Observe that if $\delta \in GD(E^i, \Delta_i)$ then $\delta \in D_i'$ and so $l_{MIN}(\mathit{CONSEQUENTS}(\delta)) = i$. By Lemma 1.10, the existence of a refutation of γ from Γ, given $l_{MAX}(\neg\gamma) < j$, implies that there is a refutation from $\Psi \subseteq \Gamma$ such that $\psi \in \Psi \leftrightarrow l_{MIN}(\lambda) < j$. Thus there is a refutation from

$$\Psi = (W \cup \bigcup_{i=0}^{j-1} \mathit{CONSEQUENTS}(GD(E^i, \Delta_i))).$$

Hence $\Psi \vdash \neg\gamma$ and, by Lemma 1.1, $\Psi \vdash \neg\gamma \leftrightarrow E^{j-1} \vdash \neg\gamma$. But if $\delta \in D_j'$ then $\neg\gamma \notin E^{j-1}$ and so $E^{j-1} \not\vdash \neg\gamma$ since E^{j-1} is logically closed. Hence we obtain a contradiction by assuming that $\neg\gamma \in E$, so $\neg\gamma \notin E$.

Thus $(\neg\gamma \vee \neg\omega) \notin E$ and so $\omega \in \bigcup_{i=0}^{\infty} F_i$.

We see that $G_{i+1} \subseteq \bigcup_{i=0}^{\infty} F_i$, and by induction $\bigcup_{i=0}^{\infty} G_i \subseteq \bigcup_{i=0}^{\infty} F_i$.

Therefore $E^j \subseteq \bigcup_{i=0}^{\infty} F_i$.

By induction, $E \subseteq \bigcup_{i=0}^{\infty} F_i$.

Together, (1) and (2) show that $E = \bigcup_{i=0}^{\infty} F_i$, so E is an extension for Δ.

QED Theorem 1

Before presenting the proof of Theorem 2, we repeat the definition of the procedure to generate extensions given earlier. Superscripts have been added which serve only as reference points in the proofs. They do not effect the computation.

$H_0 \leftarrow W$; $j \leftarrow 0$;
repeat
  $j \leftarrow j + 1$; $h_0^j \leftarrow W$; $GD_0^j \leftarrow \{ \}$; $i \leftarrow 0$;
  repeat
    $D_i^j \leftarrow \{ \frac{\alpha : \beta}{\gamma} \in D \mid (h_i^j \vdash \alpha), (h_i^j \nvdash \neg\beta), (H_{j-1} \nvdash \neg\beta) \}$;
    if $\neg$null($D_i^j - GD_i^j$) **then**
      choose δ **from** $(D_i^j - GD_i^j)$;
      $GD_{i+1}^j \leftarrow GD_i^j \cup \{\delta\}$;
      $h_{i+1}^j \leftarrow h_i^j \cup \{\text{CONSEQUENT}(\delta)\}$; **endif**;
    $i \leftarrow i + 1$;
  until null($D_{i-1}^j - GD_{i-1}^j$);
  $H_j = h_{i-1}^j$
until $H_j = H_{j-1}$

Lemma 2.1

If Δ is a finite default theory, then the algorithm can fail to converge only if one of the approximations is repeated. i.e. for some j and some $k > j+1$, $H_j = H_k$.

Proof

If Δ is finite, there are only a finite number of different combinations possible. Thus there are only a finite number of distinct H_i's which can be constructed. If $H_j = H_{j+1}$, the algorithm converges.

QED Lemma 2.1

Lemma 2.2

If Δ is a finite, semi-normal default theory, and W is consistent, then

$$H_i \vdash \beta \rightarrow H_i \nvdash \neg\beta.$$

Proof

Assume $H_i \vdash \beta, \neg\beta$. Let r, s be the smallest integers such that $h_r^i \vdash \beta$, $h_s^i \vdash \neg\beta$. Assume $r \leq s$, so $h_{s-1}^i \nvdash \neg\beta$. By hypothesis, $h_s^i \vdash \beta, \neg\beta$. Now $h_s^i = h_{s-1}^i \cup \{\omega\}$, where

$$\frac{\alpha : \omega \wedge \gamma}{\omega} \in D,\ \alpha \in h_{s-1}^i,\ H_{i-1} \nvdash (\neg\omega \vee \neg\gamma),\ \text{and}\ h_{s-1}^i \nvdash (\neg\omega \vee \neg\gamma)\ .$$

But if $h_s^i \vdash \beta, \neg\beta$, then $(h_{s-1}^i \cup \{\omega\}) \vdash \beta, \neg\beta$ so $h_{s-1}^i \vdash \neg\omega$ and hence $h_{s-1}^i \vdash (\neg\omega \vee \neg\gamma)$ which is a contradiction. The proof is similar if $s < r$.

QED Lemma 2.2

Definition 2.3

A default theory, $\Delta = (D, W)$, is ***hierarchical*** if it satisfies the following conditions:

(1) W contains only:

a) Literals (i.e. Atomic formulae or their negations), or

b) Disjuncts of the form $(\alpha \vee \beta)$ where α and β are literals.

(2) D contains only normal and semi-normal defaults of the form:

$$\frac{\alpha : \beta}{\beta} \quad \text{or} \quad \frac{\alpha : \beta \wedge \gamma_1 \wedge \cdots \wedge \gamma_n}{\beta}$$

where α, β, and γ_i are literals.

Lemma 2.4

If Δ is a finite, ordered, hierarchical default theory, if W is consistent, and if β is a literal, then $H_{i-1} \vdash \beta \rightarrow H_i \nvdash \neg\beta$.

Proof

Assume $H_{j-1} \vdash \beta$, and consider $H_j = \bigcup_{i=0}^{\infty} h_i^j$. Assume $H_j \vdash \neg\beta$. The proof proceeds by induction.

base

$h_0^j = W$. Since $H_{j-1} \nvdash \neg\beta$, clearly $W \nvdash \neg\beta$. Therefore $h_0^j \nvdash \neg\beta$.

step

Assume $h_i^j \nvdash \neg\beta$ and $h_{i+1}^j \vdash \neg\beta$. $h_{i+1}^j = h_i^j \cup \{\omega\}$, where

$$\frac{\alpha : \gamma \wedge \omega}{\omega} \in D,\ h_i^j \vdash \alpha,\ h_i^j \nvdash (\neg\gamma \vee \neg\omega), \text{ and } H_{j-1} \nvdash (\neg\gamma \vee \neg\omega).$$

Clearly, $\omega \neq \neg\beta$ or else $H_{j-1} \vdash \neg\omega$.

Note that:

i) H_j contains only disjunctions of two literals.

ii) $h_i^j = W \cup \mathit{CONSEQUENTS}(GD_i^j)$

iii) $GD_i^j \subseteq D$

iv) $\mathit{CONSEQUENTS}(GD_i^j) \subseteq$ ***Literals*** .

Consider a linear resolution refutation of β (i.e. a proof of $\neg\beta$) from h_{i+1}^{j}, with top clause β. We continue by induction on the structure of this refutation.

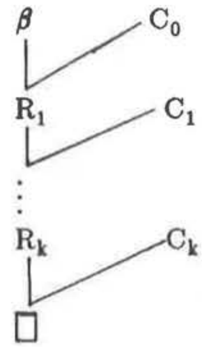

base

$\omega \in$ *Literals* and $\omega \neq \neg\beta$ so $C_0 \neq \omega$. Clearly, $C_0 \neq \beta$. Thus $C_0 \in h_i^j$. If $C_0 \in h_i^j - W$, then $C_0 \in$ *Literals*. But then $C_0 = \neg\beta$ which leads to the contradiction that $h_i^j \vdash \neg\beta$. Thus $C_0 \in W$. Clearly $C_0 \notin$ *Literals*, as above. Hence $C_0 = (\neg\beta \vee \xi)$, with $\xi \in$ *Literals*. Thus $R_1 = \xi \neq \square$.

step

Assume: i) $\omega \notin \{C_0, \ldots, C_{n-1}\}$

ii) $\{C_0, \ldots, C_{n-1}\} \subseteq W$

iii) $\{R_1, \ldots, R_n\} \subseteq$ *Literals*.

Let $R_n = \eta \in$ *Literals*. If $C_n = \omega$ then $\omega = \neg\eta$ so $W \cup \{\omega\} \vdash \neg\beta$ but $W \subseteq H_{j-1}$ and $H_{j-1} \vdash \beta$, so $H_{j-1} \vdash \beta, \neg\beta$, which contradicts Lemma 2.2. Clearly $\eta \neq \neg\beta$, so $C_n \neq \beta$, or else $W \vdash \neg\beta$ which is false. Thus $C_n \in W$. Clearly $C_n \notin$ *Literals*, as above, hence $C_n = (\neg\eta \vee \lambda)$ with $\lambda \in$ *Literals*. Therefore $R_{n+1} = \lambda \neq \square$.

So: i) $\omega \notin \{C_0, \ldots, C_n\}$

ii) $\{C_0, \ldots, C_n\} \subseteq W$

iii) $\{R_1, \ldots, R_{n+1}\} \subseteq$ *Literals*.

By induction, there is no such resolution refutation and the required result is proved.

QED Lemma 2.4

Lemma 2.5

If Δ is a finite, ordered, hierarchical default theory, and $\{\alpha_1, \cdots, \alpha_n\} \subseteq$ *Literals*, then $H_i \vdash (\alpha_1 \vee \cdots \vee \alpha_n)$ if and only if $W \vdash (\alpha_1 \vee \cdots \vee \alpha_n)$ or $H_i \vdash \alpha_j$, for some j.

Proof

($\leftarrow$) Trivial.

($\rightarrow$) Assume false, and consider a linear resolution proof of $(\alpha_1 \vee \cdots \vee \alpha_n)$ (i.e. a refutation of $(\neg\alpha_1 \wedge \cdots \wedge \neg\alpha_n)$) from H_i, with top clause $R_0 \in \{\neg\alpha_1, \cdots, \neg\alpha_n\}$.

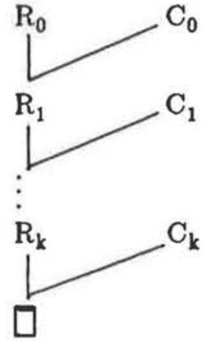

We know that $C_0 \in H_i \cup \{\neg\alpha_1, \cdots, \neg\alpha_n\}$, and that, for $i>0$, $C_i \in H_i$ or $C_i \in \{R_j \mid j \leq i\}$ or $C_i \in \{\neg\alpha_1, \cdots, \neg\alpha_n\}$. We proceed by induction.

base

Without loss of generality, assume $R_0 = \neg\alpha_1$. Clearly $\alpha_1 \notin \{\neg\alpha_1, \cdots, \neg\alpha_n\}$, or else $W \vdash (\alpha_1 \vee \cdots \vee \alpha_n)$, so $C_0 \notin \{\neg\alpha_1, \cdots, \neg\alpha_n\}$. Clearly $C_0 \neq \alpha_1$ or else $H_i \vdash \alpha_1$ which contradicts our assumption. Hence $C_0 = (\alpha_1 \vee \gamma) \in W$, for some $\gamma \in$ *Literals*, and so $R_1 = \gamma \neq \Box$.

step

Assume a) $\{R_0, \ldots, R_n\} \subseteq$ *Literals*

b) $\{C_0, \ldots, C_{n-1}\} \subseteq W$.

Let $R_n = \eta \in$ *Literals*. If $C_n = \neg\eta \in \{\neg\alpha_1, \cdots, \neg\alpha_n\}$ then $W \vdash (\alpha_1 \vee \cdots \vee \alpha_n)$ which contradicts our hypothesis. If $C_n = \neg\eta \in H_i \cup \{R_0, \cdots, R_n\}$ then $H_i \vdash \alpha_1$

which also contradicts the hypothesis. Hence $C_n = (\neg\eta \vee \xi) \in W$, with $\xi \in$ *Literals* and $R_{n+1} = \xi \neq \Box$.

By induction, there is no such resolution refutation, and the lemma is proved.

QED Lemma 2.5

Lemma 2.6

If Δ is a finite, ordered, hierarchical default theory, and $\alpha \in$ *Literals*, then $H_i \vdash \alpha$ if and only if $W \vdash \alpha$ or $\exists\ \beta \in$ Literals. $l(\beta) \leq l(\alpha)$, and $\beta \in H_i$, and $W \vdash (\beta \supset \alpha)$.

Proof

($\leftarrow$) Trivial.

($\rightarrow$) Assume false and consider a linear resolution proof of α (i.e. a refutation of $\neg\alpha$) from H_i, with top clause $\neg\alpha$. We proceed by induction.

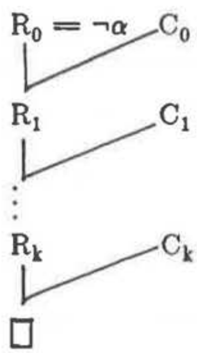

base

Clearly $C_0 \neq \alpha$ or else $\alpha \in H_i$ and $l(\alpha) \leq l(\alpha)$ and $W \vdash (\alpha \supset \alpha)$ which contradicts the hypothesis. Hence $C_0 = (\alpha \vee \gamma) \in W$, for $\gamma \in$ *Literals*. By definition, $l(\neg\gamma) \leq l(\alpha)$. $R_1 = \gamma \neq \Box$. Clearly $W \vdash (\neg\gamma \supset \alpha)$.

step

Assume: a) $\{C_0, \ldots, C_{n-1}\} \subseteq W$

b) $\{R_0, \ldots, R_n\} \subseteq$ *Literals*

c) $l(\neg R_n) \leq l(\alpha)$

d) $W \vdash (\neg R_n \supset \alpha)$

Let $R_n = \eta$. If $C_n = \neg\eta \in H_i$ then $H_i \vdash \alpha$, $\neg\eta \in H_i$, $W \vdash (\neg\eta \supset \alpha)$, and $l(\neg\eta) = l(\neg R_n) \leq l(\alpha)$ which contradicts our assumption. If $C_n = \neg\eta = \neg\alpha$ then $W \vdash \alpha$ which is also a contradiction. Hence $C_n = (\neg\eta \vee \xi) \in W$, with $\xi \in$ *Literals*, $R_{n+1} = \xi \neq \Box$, and $l(\neg R_{n+1}) = l(\neg\xi) \leq l(\neg\eta) = l(\neg R_n) \leq l(\alpha)$. By Modus Ponens, $W \vdash (\neg\xi \supset \alpha)$.

Thus there is no such refutation, and the result is proved.

QED Lemma 2.6

Lemma 2.7

If Δ is a finite, ordered, hierarchical default theory, and $\alpha \in$ Literals, $\alpha \notin H_i$, $\alpha \notin H_j$, and $\alpha \in H_k$ for $i<k<j$, then

$$\exists\, \beta \in \text{Literals}. \ (l(\beta) < l(\alpha)) \text{ and } \beta \in \bigcup_{i<r\leq j} H_i \Delta H_r .$$

Proof

Let j be the least $j>k$ such that $\alpha \notin H_j$.

Define $D_\alpha = \{\delta \in D \mid \delta = \dfrac{\gamma : \alpha \wedge \omega_1 \wedge \cdots \wedge \omega_n}{\alpha}\}$

$$GD_\alpha^i = \bigcup_{r=0}^{\infty} GD_r^i \cap D_\alpha$$

Clearly $GD_\alpha^{j-1} \neq \{\ \}$ and $GD_\alpha^j = \{\ \}$. Consider $\delta \in GD_\alpha^{j-1}$. Since $\delta \notin GD_\alpha^j$ three cases are possible:

1) $H_{j-1} \vdash (\neg\omega_1 \vee \cdots \vee \neg\omega_n)$. By Lemma 2.5, there is an ω_r, say ω, such that $H_{j-1} \vdash \neg\omega$. By Lemma 2.6, there is a $\beta \in H_{j-1}$ such that $l(\beta) \leq l(\omega)$ and $W \vdash (\beta \supset \omega)$. But then $l(\beta) < l(\alpha)$. Clearly $\beta \notin H_{j-2}$, so β is the required literal.

2) $H_j \vdash (\neg\omega_1 \vee \cdots \vee \neg\omega_n)$. The argument for case 1 applies.

3) $H_j \not\vdash \gamma$. By recursively applying the foregoing arguments to γ, we can construct a set of γ_r's which were in H_{j-1} and are not in H_j. The first of these to go into H_{j-1} must also go into H_j, unless $H_{j-1} \cup H_j$ contains a $\beta \ll \gamma_r \lesssim \alpha$ which was not in H_i .

QED Lemma 2.7

Lemma 2.8

If Δ is a finite, ordered, hierarchical default theory, and $\alpha \in$ Literals, $\alpha \in H_i$, $\alpha \in H_j$, and $\alpha \notin H_k$ for $i<k<j$, then either

1) $\exists\, \beta \in$ Literals. $(l(\beta) < l(\alpha))$ and $\beta \in \bigcup_{i<r\leq j} H_i \Delta H_r$, or

2) $\exists\, \beta \in$ Literals. $(l(\beta) \leq l(\alpha))$ and $\beta \in H_j$ and $\beta \notin H_i$.

Proof

Let k be the least $k>i$ such that $\alpha \notin H_k$. Let j be the least $j>k$ such that $\alpha \in H_j$.

Consider $\delta = \dfrac{\gamma : \alpha \wedge \beta}{\alpha} \in GD_\alpha^j$. Clearly $GD_\alpha^j \neq \{\ \}$, and $\delta \notin GD_\alpha^k$.

Cases: 1) $H_k \vdash \neg\beta$, $H_j \not\vdash \neg\beta$. This gives the first of the required conditions, by Lemmas 2.5 and 2.6.

2) $H_{k-1} \vdash \neg\beta$, $H_j \not\vdash \neg\beta$. The argument for case 1 applies.

3) $H_k \not\vdash \gamma$, $H_j \vdash \gamma$. By Lemma 2.6, $\exists\, \gamma_1 \lesssim \alpha$. $\gamma_1 \in H_j$, $\gamma_1 \notin H_k$.

Cases: a) $\gamma_1 \notin H_i$. This is the second of the required conditions.

b) $\gamma_1 \in H_i$. Repeating the above arguments for γ_1 yields a (possibly cyclic) chain of γ_r's such that $\gamma_r \in H_{k-1}$, $\gamma_r \notin H_k$. Consider the first γ_r to go into H_{k-1} . It must also go into H_k , which is a contradiction.

QED Lemma 2.8

Theorem 2

The procedure presented above always converges when applied to a finite, ordered, hierarchical default theory.

Proof

By Lemma 2.1, non-convergence implies there is a cycle. i.e. for some i and some j>i, $H_i = H_j$ and $H_i \neq H_{i+1}$.

Choose $\alpha \in \bigcup_{i<k\leq j} (H_i \Delta H_k)$ such that $\alpha \in$ *Literals* and for every $\beta \in \bigcup_{i<k\leq j} (H_i \Delta H_k)$, $\neg(l(\beta) < l(\alpha))$. Thus α is the "least" literal to change state between H_i and H_j. There are two cases:

(1) If $\alpha \notin H_i$ and $\alpha \in H_k$ then, by Lemma 2.7, $\exists\ \beta \in \bigcup_{i<k\leq j} (H_i \Delta H_k).\ l(\beta) < l(\alpha)$, so α is not the least such α, which is a contradiction.

(2) If $\alpha \in H_i$ and $\alpha \notin H_k$ then, by Lemma 2.8, either

a) $\exists\ \beta \in \bigcup_{i<k\leq j} (H_i \Delta H_k).\ l(\beta) < l(\alpha)$

so α is not the least such α, which is a contradiction, or

b) $\exists\ \beta.\ \beta \in H_j$ and $\beta \notin H_i$

which implies that $H_i \neq H_j$ which is also a contradiction.

Therefore, there is no cycle, and so the procedure converges.

QED Theorem 2

Theorem 3

The procedure given above always converges immediately when applied to a finite, normal default theory $\Delta = (D, W)$ — i.e. $Th(H_1)$ is an extension.

Proof

Etherington [1982] shows that $H_1 = H_2$ if and only if $Th(H_1)$ is an extension for Δ. If W is inconsistent, then $Th(H_1) = L$ which is an extension for Δ. Hence assume W is consistent. To show that $Th(H_1)$ is an extension for Δ, we invoke Theorem 0.1 and show that $Th(H_1) = \bigcup_{i=0}^{\infty} E_i$, where

$$E_0 = W$$

$$E_{i+1} = Th(E_i) \cup \{\omega \mid \frac{\alpha : \omega}{\omega} \in D, \alpha \in E_i, \neg\omega \notin Th(H_1)\} .$$

a) We first show that $\bigcup_{i=0}^{\infty} E_i \subseteq Th(H_1)$. Recall that $H_1 = \bigcup_{i=0}^{\infty} h_i^1$.

base

Clearly $E_0 = W = h_0^1 \subseteq Th(H_1)$.

step

Assume $E_i \subseteq Th(H_1)$ and consider $\omega \in E_{i+1}$.

i) If $\omega \in Th(E_i)$ then $\omega \in Th(H_1)$, by hypothesis and closure.

ii) Otherwise $\omega \in \{\omega \mid \frac{\alpha : \omega}{\omega} \in D, \alpha \in E_i, \neg\omega \notin Th(H_1)\}$. Therefore $H_1 \not\vdash \neg\omega$.

Hence $H_0 \not\vdash \neg\omega$ since $H_0 = W \subseteq H_1$. Also, $\alpha \in E_i$, so $\alpha \in Th(H_1)$, by hypothesis. It follows by [Etherington 1982, Lemma 3.3] that $H_1 \vdash \omega$.

Hence $E_{i+1} \subseteq Th(H_1)$.

b) Finally, we show that $Th(H_1) \subseteq \bigcup_{r=1}^{\infty} E_r$.

Since $\bigcup_{r=1}^{\infty} E_r$ is logically closed, it suffices to show that $H_1 \subseteq \bigcup_{r=1}^{\infty} E_r$.

base

Clearly $h_0^1 = W = E_0 \subseteq \bigcup_{r=1}^{\infty} E_r$.

step

Assume that $h_i^1 \subseteq \bigcup_{r=1}^{\infty} E_r$, and consider h_{i+1}^1 .

$h_{i+1}^1 = h_i^1 \cup \{\omega\}$, for some $\omega \in CONSEQUENTS(D_i^1)$.

Since $h_i^1 \subseteq \bigcup_{r=1}^{\infty} E_r$ by hypothesis, we need only show that $\omega \in \bigcup_{r=1}^{\infty} E_r$.

Since $\omega \in CONSEQUENTS(D_i^1)$, for some $\delta = \frac{\alpha : \omega}{\omega} \in D$, $\alpha \in h_i^1$,

$H_0 \not\vdash \neg\omega$, and $h_i^1 \not\vdash \neg\omega$.

By hypothesis, since $\alpha \in h_i^1$, $\alpha \in \bigcup_{r=1}^{\infty} E_r$, so $\alpha \in E_j$ for some j.

Since $\omega \in h_{i+1}^1 \subseteq H_1$, it follows by Lemma 2.2 that $H_1 \not\vdash \neg\omega$.

But then by definition of E_{j+1}, $\omega \in E_{j+1} \subseteq \bigcup_{r=1}^{\infty} E_r$.

Combining (a) and (b), we have the desired result.

QED Theorem 3

Appendix II — Dictionary of Symbols

Symbol	Definition
$\in$	Set membership
$\notin$	Set non-membership
$\bigcup$	Set union
$\bigcap$	Set intersection
$\{\}$	The empty set
$-$	Set difference: $\Psi - \Gamma = \{\alpha \mid \alpha \in \Psi \text{ and } \alpha \notin \Gamma\}$
Δ	Symmetric set difference: $\Psi \,\Delta\, \Gamma \equiv (\Psi - \Gamma) \bigcup (\Gamma - \Psi)$
$\vdash$	First order provability
$\nvdash$	First order non-provability
$\supset$	Logical implication
$\neg$	Logical negation
$\wedge$	Logical and
$\vee$	Logical or
$\exists$	Existential quantifier
$\forall$	Universal quantifier
.	Scope indicator: Preceding quantifier binds to end of formula.
$\Box$	The null clause
Th	Logical closure operator
$\rightarrow$	"It follows that" or "Implies"
$\leftrightarrow$	If and only if
iff	If and only if
$\therefore$	Therefore
$\ll$	Strong precedence relation on *Literals* $\times$ *Literals*
$\underline{\ll}$	Weak precedence relation on *Literals* $\times$ *Literals*
$\mapsto$	Function mapping
L	The first-order language (i.e. all well-formed formulae)
N	The set of all Natural numbers
Literals	The set of all atomic formulae and their negations
Ψ	A set of formulae
Γ	A set of formulae
Δ	A default theory