2

DEFAULT REASONING AND DYNAMIC INTERPRETATION OF NATURAL LANGUAGE*

EWAN KLEIN & MARC MOENS
Centre for Cognitive Science
2 Buccleuch Place
University of Edinburgh
Edinburgh EH8 9LW, Scotland, UK

FRANK VELTMAN
Department of Philosophy
University of Amsterdam
15 Nieuwe Doelenstraat
Amsterdam, The Netherlands

Abstract

We present a proposal for treating default reasoning from the perspective of a dynamic approach to semantics, where meaning is a mapping between information states. Information states are identified with sets of possible worlds—the epistemic possibilities which those states admit. Generic rules, like *On weekdays, Giles normally gets up at 8.00* are then taken to induce a pre-order on possible worlds, where worlds complying with the rules are less exceptional than those which go against the rules. Thus, a particular weekday on which Giles gets up at 8.00 is less exceptional than one on which he stays in bed till noon. Unlike many other approaches to nonmonotonicity, we draw a distinction at the level of the object language between defeasible and indefeasible conclusions.

---

* This paper has been published before in *ESPRIT '90. Proceedings of the Annual ESPRIT Conference, Brussels, November 12-15, 1990*, Kluwer Academic Publishers, Dordrecht, pp 52-65. It is reprinted here with the kind permission of the Commission of the European Communities Directorate-General Telecommunications, Information Industries and Innovation, Scientific and Technical Communications Service, Luxembourg.

1 Introduction

The ESPRIT Basic Research Action DYANA —*Dynamic Interpretation of Natural Language*—is concerned with developing a formal theory of language interpretation and processing which models human cognitive abilities but is at the same time mathematically precise and admits computational interpretation. An important goal of DYANA is to go beyond particular, isolated problems occurring at individual levels of interpretation and to study the way these levels of interpretation interact in an integrated theory. The work programme is divided into three interdependent components:

- grammar development, speech and prosody
- meaning, discourse and reasoning
- logic and computation

For a more detailed discussion of the work in each of these areas, and for a general overview of the project, the reader is referred to Klein and Moens [KlM 89]. In this paper we will describe recent DYANA work on default reasoning. In § 2, we will sketch why this work is relevant to natural language understanding and show how it fits into the DYANA project. In § 3, we discuss some basic notions of updates semantics, while § 4 presents a key distinction between stable and non-stable sentences. § 5 sketches the mechanisms by which default rules induce a preference order on epistemic states. Finally, in § 6, we briefly discuss relations between our approach and that of semantic networks, and point to some directions for future research.

2 Partiality, Dynamics and Nonmonotonicity

DYANA focuses on two important themes, namely the dynamics of natural language interpretation, and theories of partial information. The two themes are connected: interpretation is dynamic since it involves the constant manipulation of information which is extracted, transduced and modified at all levels of representation—phonological, syntactic, semantic, and pragmatic. Meaning thus becomes a **dynamic** notion: at all levels of representation, it can be defined as a function from information states to information states.[1] Both the domain and the range of this function will be states of partial information, since complete information states hardly need an update. But partiality also arises as a result of the dynamics of the interpretation process itself: ambiguities and other indeterminacies will be encountered at each stage of the

[1]For a synoptic discussion of the dynamic perspective to logic, see van Benthem [1].

interpretation process. The result is that it is not always possible to pass complete and reliable information between levels of representation in a predefined way.

Defeasibility plays a pervasive role in natural language understanding. At the most global level, understanding a discourse involves the integration of new incoming information into an existing body of beliefs, assumptions and commitments. It is hardly surprising that information states evolve in a nonmonotonic fashion—assumptions which are plausible at one stage become rendered untenable later on, and even deeply-held commitments may have to be abandoned in the face of new, conflicting facts. However, defeasibility infuses the very texture of the human processing mechanisms which map linguistic input (whether speech or written text) into some kind of discourse representation. We briefly list below just a few examples where defeasible inferences are drawn on the basis of partial information:

Semantics:

- Generic sentences: *Tigers have four legs. Shere Khan is a tiger... Shere Khan has lost a leg.*
- Quantifier scope preferences: *Every student here speaks a foreign language... It is French.*
- Tense and aspect: *Lee was crossing the street... Unfortunately, she was hit by a truck before she reached the other side.*
- Lexical semantics: *This is a flower... In fact, it's a plastic flower.*

Morphology:

1. English verbs take a past tense by suffixing *-d* (e.g. *bake* ~ *baked*).
2. Verbs with roots of the form *Xing* (e.g. *sing*) take a past tense by changing to *Xang* (e.g. *sang*).
3. *bring* is a verb of the form *Xing* and takes a past tense *brought.*

Morphophonology:

1. The masculine singular form of the French lexeme BEAU is *bo*\\
2. The masculine singular form of the lexeme BEAU is *bɛl*\\ if the following word starts with a vowel.

Notice that our morphology example is analogous to the following set of statements:

1. Birds (normally) fly.
2. Penguins (which are birds) waddle (but don't fly).
3. Max is a penguin which hops (but doesn't waddle).

Formulated in terms of rule application within the framework of generative grammar, we would say that more specific rules (like rule (3) for *bring*) are deemed to take precedence over less specific ones (such as rule (2) for verbs of the form *Xing*, or rule (1) for verbs in general). The most general rule is said to be the 'elsewhere' case.

The problem of formalizing nonmonotonic inference is an active research topic in the area of Artificial Intelligence. Reasoning devices are supposed to derive conclusions that follow logically from the facts and rules stored in their databases. However, it has often been noted that reasoning devices are sometimes expected to draw conclusions that are not necessarily true but nevertheless seem reasonable given the circumstances. The Artificial Intelligence literature contains many examples of this type of default reasoning in domains other than natural language, and offers a plethora of techniques for the formalisation of the nonmonotonic behaviour of reasoning systems (see, e.g., [Rei 87]; [Sho 87]).

One branch of DYANA work on nonmonotonicity, carried out by Morreau [Mor 90], studies the dynamics of information states which support or contain, in the form of conditional sentences, meta-information about their own response to revision. In a second branch of research, [Vel 90] develops a modal semantics for default reasoning and, to the extent that these express default rules, for generic sentences. As such it moves on territory which is familiar from the Artificial Intelligence literature on the subject, while importing into it techniques which originated in philosophical logic. Given limitations of space, we will not try to discuss Morreau's work further here, but instead give a brief overview of Veltman's results. Before entering into more details, it is worth drawing attention right away to a central feature of our approach: the notion of default reasoning is captured by drawing a distinetion between defeasible and non-defeasible conclusions at the level of the **object language**. As a result, our task is provide an adequate semantics for special kinds of sentences, namely those which express default rules and defeasible conclusions.

3 Dynamic Interpretation

According to the standard explication of logical validity, an argument is valid if its premises cannot all be true without its conclusion being true as well. Crucial to this approach is the specification of truth conditions. The heart of the theory presented below consists instead in a specification of **update conditions**. That is, you know the meaning $[\phi]$ of a sentence ϕ if you know the change that it brings about in the information state of anyone who accepts the news conveyed by ϕ. Thus, as we suggested above, $[\phi]$ is an operation for updating the information states of an idealised agent.

Let σ be an information state and ϕ a sentence with meaning $[\phi]$. Then we write

$$\sigma[\phi]$$

for the information state that results when σ is updated with $[\phi]$. In most cases $\sigma[\phi]$ will be different from σ, but it is possible that the information conveyed by ϕ is already subsumed by σ; thus, updating σ by $[\phi]$ will simply result in σ. In such a case, i.e. when

$$\sigma[\phi] = \sigma,$$

we say that ϕ is **accepted in** σ, and write this as

$$\sigma \Vdash \phi.$$

It may be helpful to the reader if we give a preliminary example to show what updating rules look like. To begin with, let us be more specific about how to characterise an information state. For the sake of simplicity, an information state σ can be identified with a subset of the set W of possible worlds. Intuitively, σ represents everything that an agent takes to be true at a given time, and thus contains those worlds which may yet turn out to be the real one. If the agent happens to know nothing at all, then any world may be the actual one, and σ is just W. We shall use '**1**' to represent this minimal information state. As the agent's information increases, σ shrinks, until—in the limit—it just consists of a single world. Thus, the growth of information is understood as the elimination of possibilities. Moreover, we also admit an absurd information state '**0**', identified with the empty set. Thus, we will have $\sigma[\phi] = 0$ when ϕ is inconsistent with σ. We can now introduce some additional terminology. If $\sigma[\phi] \neq 0$, then we say that ϕ is **acceptable in** σ, whereas if $\sigma[\phi] = 0$, then ϕ is **not acceptable in** σ.

Notice finally that we do not assume information states to be 'veridical', in the sense that they must contain the actual world. We admit $\sigma[\phi] = 0$ when in fact ϕ is true, and equally we allow $\sigma[\phi] = \sigma$ when in fact ϕ is false. Suppose however that $\sigma[\phi] = 0$, for a true sentence ϕ. In this case, an agent cannot refuse to accept ϕ when confronted with the facts; rather, she should **revise** her information state in such a manner that ϕ becomes acceptable. However, we shall not attempt to say anything here about how such a revision is carried out.

Let us take as given a finite set $\mathcal{A}$ of atomic sentences, and let $\mathcal{L}(\mathcal{A})$ be a propositional language based on $\mathcal{A}$ whose sentences are built in the usual way. We can think of such sentences as expressing the kind of descriptive content which constitutes an information state. In addition, we add to the language a one-place sentential operator *might*. This can be prefixed to any sentence ϕ which does not already contain occurrences of *might*. *Might* ϕ sentences should not be thought of as expressing descriptive propositions. Rather, they have a meta-semantic character which tells us something about our current information state, namely whether ϕ is acceptable given what we already take to be true.

As pointed out by van Benthem [Ben 90], a dynamic approach to semantics makes it natural to postulate various **modes** of operating on an information state σ. For example:

Update: make a transition from σ to a new state σ' which extends the information in σ.

Downdate: revise an unsafisfactory information σ to produce a new state σ' which eliminates certain (mis)information in σ.

Test: check whether a given proposition is accepted in σ, and leave σ unchanged.

In terms of this taxonomy, descriptive sentences will perform updates, whereas sentences such as *might* ϕ carry out tests. In a fuller treatment, we would also need to allow an information state to be revised by a downdate statement—however, we will not consider problems of revision here.

Before turning to some explicit clauses to illustrate the dynamic approach, it is a useful technical detail to identify a possible world w with the set of atomic sentences from $\mathcal{A}$ which are true in w; hence W is the powerset of $\mathcal{A}$. With this clarification, the evaluation clauses for our language can now be stated as follows:

(1)
a $\sigma[p] = \sigma \cap \{w \in W \mid p \in w\}$, for any atom $p \in \mathcal{A}$
b $\sigma[\neg\phi] = \sigma \setminus \sigma[\phi]$
c $\sigma[\phi \wedge \psi] = \sigma[\phi] \cap \sigma[\psi]$
d $\sigma[\phi \vee \psi] = \sigma[\phi] \cup \sigma[\psi]$
e $\sigma[might\ \phi] = \sigma$ if $\sigma[\phi] \neq 0$
$\sigma[might\ \phi] = 0$ if $\sigma[\phi] = 0$

The analysis of *might* is motivated by the following considerations: an agent will accept *might* ϕ just in case ϕ is consistent with what she takes to be true. As pointed out above, clause (1e) tests σ rather than update it; if ϕ is acceptable in σ, then you have to accept *might* ϕ, while remaining in information state σ.

Although other notions of logical validity are possible in this context (cf. [Vel 90] for discussion), the one we shall employ here goes as follows: Let **1** be the minimal information state, where all epistemic possibilities are open. Then an argument is **valid** iff updating **1** with the premises $\psi_1 ... \psi_n$, in that order, yields an information state σ in which ϕ is accepted. Formally:

$\psi_1 ... \psi_n \vDash \phi$ iff $\mathbf{1}[\psi_1] ... [\psi_n] \Vdash \phi$

4 Stability

An important motivation of classical logic has been to abstract away from the context-dependence of ordinary discourse, with the goal of formalising arguments whose validity does not shift according to their position in a discourse. In particular, conclusions should be **stable** in the sense that if they are true at one stage, then they remain true regardless of what ensues subsequently. Equally, much attention has been paid within the framework of logic programming to find declarative formulations of data operations which abstract away from details of implementation, such as the order in which operations are carried out. Yet given that one of our goals is to formally model the context-dependence of natural language discourse, we wish to find ways of explicitly capturing the procedural aspects of informal argumentation. In pursuit of this, we will indeed allow into our formal system **non-stable** information—information which may become obsolete when more facts are acquired.

The distinction between stability and non-stability is one that we shall draw at the level of the object language. Thus, we say that some sentences ϕ are not stable, in the following sense:

Definition 1 (Stability) *A sentence* ϕ *is stable just in case for any* σ *and* $\psi_1 ... \psi_n$ *if* $\sigma \Vdash \phi$ *then* $\sigma[\psi_1] ... [\psi_n] \Vdash \phi$.[2]

Sentences involving *might* provide a simple example of non-stability. In the minimal information state **1**, *it might be raining* is accepted, since *it is raining* is certainly acceptable in **1**.

Suppose we now learn that it isn't raining, and update our information state accordingly. Then *It might be raining* becomes unacceptable. Reading p as *it is raining*, we have **1** $\Vdash$

[2]Notice that we are assuming that $\psi_1 ... \psi_n$ only express updates or tests on σ; if we were to admit downdates as well, then a different definition would be required.

might p by (1e), since $\mathbf{1}[p] \neq 0$, but we do not have $\mathbf{1}[\neg p] \Vdash$ *might p*, since $\mathbf{1}[\neg p][p] = 0$

There are other important epistemic operators which create non-stability. The one we wish to look at is *presumably*. Thus consider the following argument:

(2) Adults are normally employed
Wim is an adult
Presumably Wim is employed

According to the semantics developed in [Vel 90], this argument is valid, in the sense defined above. Notice that we do not conclude that Wim is employed—only that he presumably is. This qualification makes explicit that an unstable, and therefore defeasible, conclusion has been drawn.

The argument (2) can remain valid as one learns more about Wim, so long as there is no evidence that the new information is relevant to the conclusion:

(3) 1: Adults are normally employed
2: Wim is an adult and a student
Presumably Wim is employed

However, if we now adopt the default rule *Students normally aren't employed*, the argument is no longer valid. Thus, we can draw no conclusions about whether or not Wim is employed from the following premises:

(4) 1: Students normally aren't employed
2: Adults are normally employed
3: Wim is an adult and a student

Adding a fourth premise may make the balance tip. Thus, in (5), we draw a conclusion that is the opposite of what we previously inferred:

(5) 1: Students are normally adults
2: Students normally aren't employed
3: Adults are normally employed
4: Wim is an adult and a student
Presumably Wim isn't employed

In the presence of premise 1, the apparent incommensurability between Wim's being a student and his being an adult is lost, and premise 2 takes precedence over 3.

It should now be evident what we meant by our claim that defeasibility is made explicit at the level of the object language. In other theories, one may well infer from the premises in (2) that Wim is employed, only this time it is a different kind of inference (i.e. an nonmonotonic one). Default reasoning, we claim, is not a special kind of reasoning with ordinary sentences, but rather **ordinary reasoning with a special kind of sentence**.

5 Rules with Exceptions

Although there is no space to give an elaborated presentation of the intended semantics for sentences like those discussed in the preceding section, we shall attempt to sketch the basic mechanism by which default sentences are interpreted. The theory arose out of an attempt to give a dynamic twist to the theory developed by Delgrande [Del 88], who in turn took Lewis's [Lew 73] study of counterfactuals as his starting point.

When an agent adopts a sentence of the form *normally* ϕ, she adopts certain expectations: worlds where ϕ holds are less surprising than those where it doesn't. To capture this idea, we need to give more structure to an information state. It must not only contain the set of epistemic alternatives, as before, but also an **expectation pattern** which makes explicit what an agent would expect to happen in the absence of complete information. Of course, the dynamics of interpretation now includes two kinds of change on an information state σ:

- modifying σ's descriptive content
- modifying σ's expectation pattern

Operations available in the language can avail themselves of just one of these options, or both.

We formalise the notion of expectation pattern in terms of a pre-order $\leq$ (i.e. a reflexive, transitive relation of 'preference'), where $w \leq w'$ just in case w is at least as normal as w'. When an agent updates her information state σ with a default statement such as

(6) On weekdays, Giles normally gets up at 8.00,

her expectation pattern, encoded as the pre-order, will be modified in such a way that worlds in which the sentence holds are considered more normal than those in which it fails. Thus, given (6), a Monday on which Giles gets up at 8.00 is less exceptional than one on which he stays in bed till noon.

We increment our language with two new unary operators, *normally* and *presumably*. Again, we forbid iteration: sentences of the form *normally* ϕ and *presumably* ϕ are not allowed to contain further occurrences of any of the epistemic operators. *Presumably* ϕ perfoms a simple test on an information state σ to determine whether ϕ can be (defeasibly) concluded. *Normally* ϕ, which expresses a default rule, effects a subtle change on σ's expectation pattern.

An information state now involves a pair $(\leq,s)$, where s is again a subset of W and $\leq$ is a pre-order on W.We will call $\leq$ an **expectation pattern** on W.If $w \leq w'$ and $w' \leq w$, we write

$$w \equiv w'$$

Clearly, $\equiv$ is an equivalence relation.

Definition 2 *Let $\leq$ be an expectation pattern on W.Then*

1. *w is a* normal *world if $w \leq w'$ for every w' in W.*
2. *$\leq$ is* coherent *if there exist some normal worlds.*
3. *$NORM_{\leq}$ is the set of all normal worlds relative to $\leq$.*

That is, a pattern is considered to be coherent if there is at least one possible world in which every proposition that expresses how things should normally be does in fact hold. This does not mean, however, that the real world must satisfy all the default rules accepted by an agent—by definition, default rules allow for exceptions.

Definition 3 *Let W be as before. Then σ is an information state if $\sigma = (\leq,s)$ and one of the following conditions is satisfied:*

1. $\leq$ is a coherent pattern on W and s is a nonempty subset of W, or
2. $\leq\; = \{\langle w,w \rangle \mid w \in W\}$ *and* $s = \varnothing$.

We now have:
$1 = (W \times W, W)$ *is the* minimal *information state,*
$0 = (\{\langle w,w \rangle \mid w \in W\}, \varnothing)$ *is the* absurd *information state.*

In order to grasp the semantic treatment we are proposing, it is helpful to consider illustrations like those in Figure 1. We assume in figure 1 that $W = \{w_0, w_1, w_2, w_3\}$ where $w_0 = \varnothing$, $w_1 = \{p\}$, $w_2 = \{q\}$, and $w_3 = \{p,q\}$. If two worlds belong to the same $\equiv$ equivalence class, they are placed within the same oval, and if $w < w'$, then w is pictured to the left of w'. Finally, the worlds belonging to s are drawn within a dashed rectangle.

The information state $\sigma_1 = (\leq_1, s_1)$ on the left is the minimal state, **1**: the expectation pattern treats all worlds as equally normal. When σ_1 is updated to a new state σ_2 with the sentence *normally p*, the expectation pattern is **refined**, with the result that worlds where p does not hold are judged to be more exceptional than those where it does. In particular, the worlds w_2 and w_0 no longer stand in the 'as normal as' relation to w_3 and w_1, and hence the corresponding pairs are removed from the new expectation pattern $\leq_2$ which arises after the update. A similar refinement occurs when σ_2 is updated by *normally q*. Now the most normal world in σ_3 is w_3, where both p and q hold.

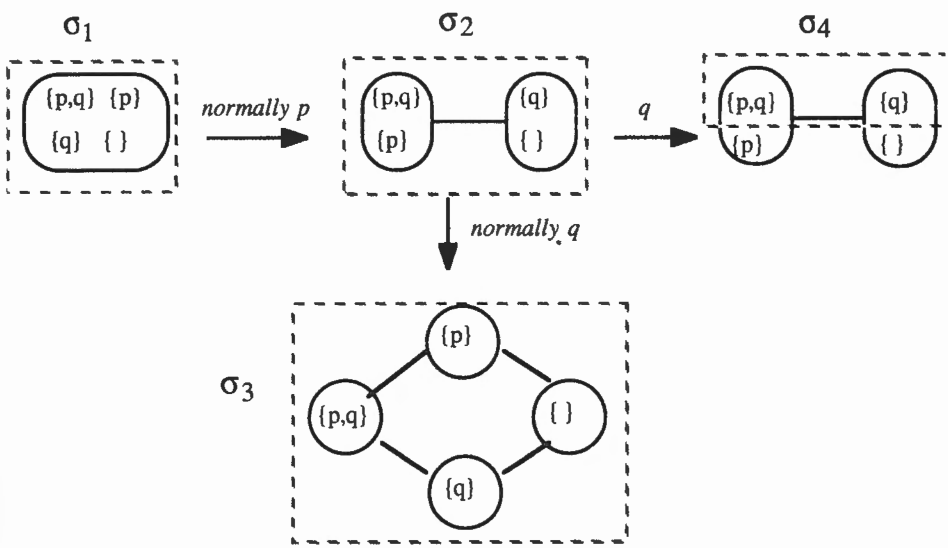

Figure 1: Expectation Patterns

As mentioned above, the dashed rectangles in Figure 1 demarcate the sets of worlds which are held to be true by the agent. It will be observed that default sentences only affect the expectation pattern in the information state; thus $s_3 = s_2 = s_1 = W$.However, when the information state σ_2 is updated with a descriptive sentence q, the s_2 is reduced accordingly, so that worlds w_1 and w_3 are excluded from the resulting set s_4.

We now define the notion of refinement.

Definition 4 (Refinement) *Let $\leq_0$, $\leq_1$ be expectation patterns on W and let $X \subset W$.*

1. $\leq_1$ is a refinement of $\leq_0$ if $\leq_1 \subset \leq_0$.
2. $\leq \circ X = \{\langle w,w' \rangle \in \leq \mid w \in X \text{ or } w' \notin X\}$.

As we already observed, refinement is brought into play when a new default rule is acquired. Suppose that we are currently in an information state σ where w is at least as normal as w'. That means that w satisfies at least as many default rules as w'. What happens when σ is updated with a new default rule *normally* ϕ ? Let us define

$$\|\phi\|$$

to be the set of ϕ-worlds; i.e. worlds in which ϕ holds. Then only worlds within $\|\phi\|$ can satisfy all the default rules. Assume that ϕ is compatible with all the preceding rules, and thus that $NORM_{\leq}[\phi] \neq \emptyset$. Then we have to refine $\leq$ to the new pattern $\leq \circ \|\phi\|$ by excluding from $\leq$ any pairs which render a non ϕ-world at least as normal as a ϕ-world; i.e. any pairs of the form $\langle w,w' \rangle$ such that $w \notin \|\phi\|$ but $w' \in \|\phi\|$. Consider, for example, $NORM_{\leq_2}$ of σ_2 in Figure 1. This will contain the pair $(\{p\},\{p,q\})$. But $\{p\} \notin \|q\| = \{\{p,q\},\{q\}\}$, though $\{p,q\}$ is. Consequently $\langle \{p\},\{p,q\} \rangle$ is removed from $\leq_2 \circ \|q\|$, as required by the update rule for *normally q*.

Although the expectation pattern $\leq_4$ of σ_4 is the same as that in σ_3, there is an obvious sense in which one of the normal worlds, namely w_1, is no longer relevant. The expectation patterns of an agent in information state σ_4 will be determined by the set $\{w_3\}$ of **optimal** worlds which result when $\leq_4$ is restricted to the set s_4 We define this notion as follows:

Definition 5 *Let $\leq$ be a pattern on W and let $s \subseteq W$.*

1. *w is* optimal in $(\leq,s)$ *if $w \in s$ and for every $w' \in s$, if $w' \leq w$ then $w' \equiv w$.*
2. *OPT is an* optimal *set in $(\leq,s)$ if there is some optimal w in $(\leq,s)$ such that $OPT = \{w' \in s \mid w \equiv w'\}$.*

Optimality, we see, is relative to two considerations: the default rules which are accepted in an information state, and the set s of worlds which constitute the current epistemic alternatives. Worlds which are less than optimal at one point become important when expectations have to be readjusted. As one's knowledge increases, and more and more alternatives are eliminated, the optimal worlds may disappear, and the best among the less than optimal worlds take over their role.

We bring this section to a close by stating update clauses for *normally* and *presumably*. Thus, if $\sigma = \langle\leq,s\rangle$ and ϕ is in $\mathcal{L}_2(\mathcal{A})$, $\sigma[\phi]$ is defined as follows:

(7) a. If ϕ is a sentence of $\mathcal{L}_0(\mathcal{A})$, then
   1. $\sigma[\phi] = 0$ if $s[\phi] = \varnothing$.
   2. Otherwise, $\sigma[\phi] = \langle\leq, s[\phi]\rangle$.
b. If ϕ is *normally* φ, then
   1. $\sigma[\phi] = 0$ if $NORM_{\leq}[\varphi] = \varnothing$.
   2. Otherwise, $\sigma[\phi] = \langle\leq \circ \|\varphi\|, s\rangle$.
c. If ϕ is *presumably* φ, then
   1. $\sigma[\phi] = \sigma$ if $OPT[\varphi] = OPT$ for every optimal set OPT in$\langle\langle\leq s,\rangle$.
   2. Otherwise, $\sigma[\phi] = 0$.

As we noted before, *presumably* ϕ resembles *might* ϕ in being an invitation to perform a test on σ rather than updating it. If the proposition expressed by ϕ holds in all optimal worlds in $\sigma = \langle\leq,s\rangle$, then *presumably* ϕ must be accepted, and σ is left unchanged. Unlike *normally* sentences, sentences of the form *presumably* ϕ are not in general stable. Even if it is a rule that *normally* ϕ, it may be wrong to expect that ϕ. Such non-stability can be illustrated schematically by the following example:

(8) $\mathbf{1}$[normally p] $\Vdash$ presumably p
   $\mathbf{1}$[normally p] [$\neg$p] $\nVdash$ p.

6. Conclusion

Important requirements for any approach to nonmonotonicity are the following:

- the defeasibility of conclusions drawn by default;
- scepticism in the face of conflicting defaults;
- the priority of more specific information before more general information.

The only other theory of nonmonotonic reasoning which seems to do justice to these requirements is the sceptical theory of nonmonotonic semantic networks due to

Thomason and his colleagues [Rei 87,CaT 89]. In a sense it is to that theory that the present one is most closely related: semantic nets may be regarded as sets of statements belonging to a proper sublanguage of the languages whose semantics is explicated in [Vel 90]. Broadly speaking, a link where P inherits form Q is formalised as a sentence of the form

$$P \sim Q$$

which we read as 'P normally implies Q'. Unfortunately there is not space here to present the semantics of this new binary operator; however, it essentially selects an expectation pattern that is appropriate for the property P. This semantics does not correspond exactly to the inference principles proposed for semantic networks, but the fit is quite close. In fact, it appears that the logic defined by [Vel 90] is preferable to that characterized by those inference principles. Cycles of defaults, for example, are excluded in the Thomason *et al.* theory for technical reasons, but pose no special problem for the current modal theory of default reasoning by update, like Veltman's. It is an interesting but still open question whether the low computational complexity that has been claimed for inferencing on semantic nets is lost when we move to the richer languages introduced in [Vel 90].

The DYANA work described above is intended initially as a contribution to the theories of nonmonotonic reasoning taking shape within philosophy and Artificial Intelligence. However, it does not address directly the more specific problems of nonmonotonic processing that arise in the context of computational linguistics. In itself this is not objectionable at the end of this first phase of the research into nonmonotonic reasoning, for much is still needed by way of general clarification of the various aspects of nonmonotonic reasoning and revision in the face of inconsistency. But ultimately it is one of the central purposes of DYANA to relate such general insights to the processing of language. It is linguistic applications such as the inheritance of morphological and phonological properties and the structure of the lexicon that we hope to address in future work.

References

[Ben 90] Benthem, J. van, General Dynamics, paper presented to Workshop on Semantics and Computation, CIS Universtät München, July 1990. To appear in *Theoretical Linguistics*.

[CaT 89] Carpenter, B., and Thomason, R.H., *Inheritance Theory and Path-Based Reasoning: an Introduction*, Ms., Intelligent Systems Program, Pittsburgh, PA: University of Pittsburgh, 1989.

[Del 88] Delgrande, J., An Approach to Default Reasoning Based on a First-Order Conditional Logic: Revised Report, *Artificial Intelligence*, 36, 63-90, 1988.

[HTT 87] Horty, J.F., Thomason, R.H., and Touretzky, D.S., A Skeptical Theory of Inheritance in Nonmonotonic Semantic Networks, *Proceedings of the 6th National Conference on Artificial Intelligence*, Seattle, Washington, 358-363, 1987.

[KlM 89] Klein, E., and Moens, M., The Dynamic Interpretation of Natural Language in *ESPRIT '89; Proceedings of the 6th Annual ESPRIT Conference,* Brussels, 27 November-1 December, 1100-1107, 1989.

[Lew 73] Lewis, D., *Counterfactuals*, Basil Blackwell, Oxford, 1973.

[Mor 90] Morreau, M., Epistemic Semantics for Counterfactuals, in Kamp, H. (ed.) *Conditionals, Defaults and Belief Revision*. Edinburgh; DYANA Deliverable R2.5.A, 1-27, 1990.

[Rei 87] Reiter, R., Nonmonotonic Reasoning, *Annual Review of Computer Science*, 2, 147-187, 1987.

[Sho 87] Shoham, Y., Nonmonotonic Logics: Meaning and Utility, *IJCAI87*, 388-393, 1987.

[Vel 90] Veltman, F., Defaults in Update Semantics, in Kamp, H. (ed.) *Conditionals, Defaults and Belief Revision*. Edinburgh: DYANA Deliverable R2.5.A, 28-64, 1990.