Artificial Intelligence Review **8**: 3–16, 1994.

Operational Concepts of Nonmonotonic Logics Part 1: Default Logic

G. ANTONIOU and V. SPERSCHNEIDER

University of Osnabrück, Department of Mathematics and Computer Science, Albrechtstrasse 28, D-4500 Osnabrück, Germany

Abstract. We give an introduction to default logic, one of the most prominent nonmonotonic logics. Emphasis is given to providing an operational interpretation for the semantics of default logic that is usually defined by fixed-point concepts (extensions). We introduce a process model that allows to exactly calculate the extensions of a default theory in a quite easy way. We give a prototypical implementation of processes in Prolog able to handle the examples that can be found in literature. Finally, we develop some theoretical results about default logic and give new simple proofs using the process model as a theoretical tool.

Key Words: default logic, operational computation of extensions, prototypical implementation.

1. DEFAULT REASONING: MOTIVATIONAL EXAMPLES

Logic is a promising knowledge representation method (Bibel 1984; Moore 1982; Nilsson 1991). The best known logical system is predicate logic (Mendelson 1969), and its usefulness is out of question. It is a logic originally developed to capture mathematical reasoning, and provides thus a formal background for this discipline. One of the main properties of predicate logic is that derivation of new results does not destroy previous conclusions. This property is called *monotonicity* and can be formalized as follows:

If $M \subseteq M'$ and φ follows from M, then φ also follows from M'.

Now consider an agent acting in a heavily changing environment without having available complete knowledge about this environment. This situation is typical for acting in the real world. Reasons for the incompleteness of knowledge may differ from example to example: it may be lack of time to wait for further information, or economical reasons that prevent a collection of complete information.

1. EXAMPLE
Suppose I wake up in the morning and am asked how I shall get to work. My answer is by bus, because I usually drive to work by bus, and I do not have any information that this is not possible. This kind of reasoning is called *default reasoning*. So I leave home, get to the bus station and hear that there is surprisingly a bus strike called out in the night before. Now I have to revise

my previous conclusion of default reasoning which is invalidated now. My behaviour is thus *nonmonotonic*. Note that the rule of thumb (*default*) I used before does not have to be rejected. It is simply no more applicable, because now I have heard that there is a bus strike, so it is impossible to drive by bus. Of course, the situation could be modelled by a predicate logic formula like

$$(\neg\text{strike} \wedge \neg\text{snow} \wedge \text{oil-embargo} \wedge \ldots) \rightarrow \text{drive-by-bus}$$

instead. The problem with this axiom is twofold. First, all possible reasons why it is impossible to drive by bus would have to be listed; but there could always be an unforeseeable reason not captured by this axiom. Second, there remains the problem of proving(!) that all these reasons are not true in order to conclude that I can drive by bus. So, I would have to phone with news agencies, look out of my window etc. before making up my mind. If I have luck, I could draw my conclusion early in the afternoon (a little bit too late!), but often it is not even possible to gather the information needed.

2. EXAMPLE

Suppose you are told Saturday morning by telegram that you will receive a bird named *tweety* (the most prominent bird in the history of nonmonotonic logic!), and that tweety will arrive Sunday morning. Shops being closed on Sunday, you must quickly provide for some food and decide whether to buy a bird-cage. In the supermarket you get troubled whether to buy indeed the offered, quite expensive bird-cage, or to buy corn or meat as food. You definitely know:

Fact 1 Tweety is a red coloured bird.
Fact 2 By lack of enough money you may buy either the bird-cage, or else corn and meat, but not all of these things together.
Rule 1 If you get a bird that can fly, buy a bird-cage.
Rule 2 Eagles are carnivorous birds that can fly.
Rule 3 Penguins are carnivorous birds that cannot fly.
Rule 4 Ostriches are granivorous birds that can fly.

What do you decide to do? Classical logic would require to know what sort of bird tweety exactly is. By lack of such knowledge no decision would be possible, so tweety would either escape or starve. Default-reasoning might offer to you the following rules of thumb:

Default 1 Every bird may be assumed to fly unless the contrary is actually known.
Default 2 Buy corn unless it is known that tweety is carnivorous or you are out of money.
Default 3 Buy meat unless it is known that tweety is granivorous or you are out of money.

Since is it not known that tweety cannot fly (in particular, it is not known that tweety is a penguin or an ostrich) you may conclude from default 1 that tweety can fly. Thus, by rule 1, you buy the bird-cage. So far, it is not known that tweety is carnivorous. So you may apply default 2 and decide to buy corn. Now default 3 cannot be applied, since by current knowledge (you already bought a bird-cage and corn) and fact 2, you are out of money.

Applying the defaults in a different order, namely first 1, and then 3, you might decide to buy the bird-cage and meat. So there exits a concurrent set of beliefs leading to a different action.

Another effect occurs if you first apply default 2 and then 3. Then you buy corn and meat, and are thus out of money. Nevertheless, default 1 is applicable, since it is not known that tweety cannot fly. So, you might conclude that tweety can indeed fly and, by rule 1, buy a bird-cage. Now, having bought corn, meat and a bird-cage, you get trouble with fact 2. Logically spoken, a contradiction has occurred. So, application of default 1 after defaults 2 and 3 contradicts the prior justifications for the these applications. This disqualifies thereafter this considered application of default rule 1.

Summarizing, there are two concurrent *sets of beliefs* (they will later be called *extensions*) that an agent may hold. Note that the addition of further information (in form of new facts, rules, or defaults) may enforce a revision of any of these belief sets. If you get more information about tweety, for example, that tweety is a penguin, the application of some of the default rules in a certain order may be no longer justified.

3. EXAMPLE
A common situation in robot planning is that a particular action of a robot only slightly changes the environment. For example, taking away a box out of a room changes the presence of this particular box, but leaves unchanged temperature, colour of the floor, number of doors, and so on. In describing the altered environment after execution of such a robot action it is desirable to concentrate on the actual changes, and describe what has been left unchanged by an overall default rule instead of thousands of single rules.

> **Frame default rule:** for every situation S, ground atom β and action *act*, if holds(β, S) and if holds(β, act(S)) is consistent with current knowledge, then infer holds(β, act(S)).

Obviously, holds(β, act(S)) is inconsistent with current knowledge only if executing the action *act* in situation S yields a new situation act(S) satisfying $\neg\beta$.

In this paper we shall present the basics of a logic dealing with such defaults called *default logic*, with emphasis put on the applicability and operationality of its concepts. We assume familiarity with basic notions of predicate logic. We are using standard notation, representing predicate logic formulas by Greek letter ψ, φ etc. In case of any uncertainties, please refer to (Sperschneider and Antoniou 1991).

2. DEFAULT LOGIC: FORMAL DEFINITIONS

4. DEFINITION. A *default* δ is a string $\varphi{:}\psi_1, \ldots, \psi_n/\chi$ with closed first-order formulas φ, $\psi_1, \ldots, \psi_n$ and χ $(n > 0)$. We call φ the *prerequisite*, $\psi_1, \ldots, \psi_n$ the *justifications*, and χ the *consequent* of δ. A *default schema* is

a string of the form $\varphi:\psi_1, \ldots, \psi_n/\chi$ with arbitrary formulas. Such a schema defines a set of defaults, namely the set of all ground instances $\varphi\sigma:\psi_1\sigma, \ldots, \psi_n\sigma/\chi\sigma$ of $\varphi:\psi_1, \ldots, \psi_n/\chi$, where σ is an arbitrary ground substitution assigning values to all free variables of $\varphi, \psi_1, \ldots, \psi_n, \chi$.

5. DEFINITION. A *default theory* T is a pair (W, D) consisting of a set of closed formulas W (the set of truths) and a denumerable set of defaults D. The default set D may be defined using default schemata.

Remarks. 1. A default schema $p(X):q(X)/r(X)$ stands, for instance, for the set of all defaults $p(t):q(t)/r(t)$, for all ground terms t in the considered logical language (signature).

2. The informal reading of default $\varphi:\psi_1, \ldots, \psi_n/\chi$ is:

> If φ is currently known and it is consistent with the current knowledge to believe in each ψ_i, then conclude χ.

When we define the meaning (semantics) of defaults, we shall be primarily concerned with giving an appropriate interpretation of the notion 'current knowledge'.

The purpose of a default theory T is to lay down what an agent may believe in. The current belief of an agent forms a set E of closed formulas, called an *extension* for T. Usually, there will be several concurring (perhaps mutually excluding) extensions. Before giving the definition of extensions let us motivate it by postulating desirable properties.

- An extension E for $T = (W, D)$ should contain all truths of T: $W \subseteq E$.
- An extension E for $T = (W, D)$ should be closed with respect to logical conclusion (clearly, we do not prevent an agent from logical argument): $\mathrm{Th}(E) = E$.
- An extension E for $T = (W, D)$ should be closed with respect to application of defaults from D: if $\varphi:\psi_1, \ldots, \psi_n/\chi$ is a default in D, $\varphi \in E$ and $\neg\psi_1 \notin E, \ldots, \neg\psi_n \notin E$, then $\chi \in E$ should hold. Note that here we are already making use of E. This circularity of argument will be apparent in the following definitions that will be based on fixed-points. Our operational model in the next section is intended and will resolve this difficulty.
- An extension E for $T = (W, D)$ should be in some sense *grounded* on W w.r.t. D, i.e. that the formulas of E are obtained from the default theory with E as a proposed belief set. As an example, consider the default theory $(\emptyset, \{\text{true}:A/\neg A\})$ and let E be $\mathrm{Th}(\{\neg E\})$ (the set of all predicate logic formulas derivable from $\neg A$). E is not grounded in the sense above: If we take E as a belief set (relevant for the consistency conditions), it is not possible to obtain $\neg A$ from the default theory as the default is not applicable w.r.t. belief set E (its justification A is inconsistent with E).

6. DEFINITION. Let $\delta = \varphi:\psi_1, \ldots, \psi_n/\chi$ be a default, and E and F sets of formulas. We say that δ *is applicable to F with respect to belief set E* iff $\varphi \in F$, and $\neg\psi_1 \notin E, \ldots, \neg\psi_n \notin E$. For a set D of defaults, we say that F *is*

closed under D with respect to E iff, for every default $\varphi:\psi_1, \ldots, \psi_n/\chi$ in D that is applicable to F with respect to belief set E, its consequent χ is also contained in F.

7. DEFINITION. Given a default theory $T = (W, D)$ and a set of closed formulas E, let $\Lambda_T(E)$ be the least set of closed formulas that contains W, is closed under logical conclusion and closed under D with respect to E.

8. DEFINITION. Let T be a default theory. A set of closed formulas E is called an *extension* of T iff $\Lambda_T(E) = E$.

Note that condition $\Lambda_T(E) = E$, for a default theory $T = (W, D)$, realizes the above discussed groundedness property: given belief set E, we subsequently justify the presence in E of each of its formulas by deriving it from W via logical conclusion and application of defaults from D with E as assumed belief set governing the consistency check for the justifications in defaults. The other postulates for extensions are also fulfilled.

9. EXAMPLE
Consider the default theory $T = (\emptyset, \{\delta_0 = \text{true}:B/B,\ \delta_1 = \text{true}:\neg B/\neg B\})$. $E_1 = \text{Th}(\{B\})$ is an extension of T: It trivially includes the truths of T, $\text{Th}(E_1) = E_1$, and E_1 is closed under the default set w.r.t. E_1, as only δ_0 is applicable to E w.r.t. $\text{Th}(\{B\})$ (is it clear why?) and $B \in E_1$. Furthermore, E_1 is a minimal such set: the only deductively closed proper subset of E_1 is $\text{Th}(\emptyset)$, but $\text{Th}(\emptyset)$ is not closed under D w.r.t. E_1, as δ_0 is applicable to $\text{Th}(\emptyset)$ w.r.t. E_1, but $B \notin \text{Th}(\emptyset)$.

This shows that $\Lambda_T(E_1) = E_1$, so E_1 is an extension of T. Likewise, $E_2 = \text{Th}(\{\neg B\})$ is also an extension of T. Are there any more extension? We feel quite sure that this is not the case (as δ_0 'blocks' application of δ_1 and vice versa) but how could we prove our claim? And if we have such difficulties in this trivial example, what about really complicated theories? In the next section we shall present a method for obtaining an overview of all extensions of a default theory, thus giving an answer to these pessimistic questions.

3. AN OPERATIONAL MODEL OF EXTENSIONS BASED ON PROCESSES

10. DEFINITION. Let $T = (W, D)$ be a default theory and $\Pi = (\delta_0, \delta_1, \delta_2, \ldots)$ a finite or infinite sequence of defaults from D not containing any repetitions (modelling an application order of defaults from D). We denote by $\Pi[k]$ the initial segment of Π of length k, provided the length of Π is at least k. Then we define:

(a) $\text{In}(\Pi)$ is $\text{Th}(M)$, where M contains the formulas of W and all consequents of defaults occurring in Π.
(b) $\text{Out}(\Pi)$ collects the negations of justifications of defaults occurring in Π.
(c) Π is called a *process* of T iff δ_k is applicable to $\text{In}(\Pi[k])$ w.r.t. belief set $\text{In}(\Pi[k])$, for every k such that δ_k occurs in Π.

(d) Π is called a *successful process* of T iff $\mathrm{In}(\Pi) \cap \mathrm{Out}(\Pi) = \emptyset$, otherwise it is called a *failed process*.
(e) Π is a *closed process* of T iff every $\delta \in D$ which is applicable to $\mathrm{In}(\Pi)$ with respect to belief set $\mathrm{In}(\Pi)$ already occurs in Π.

For the default theory from example 9, $\Pi_1 = (\delta_0)$ is a process, but not $\Pi_2 = (\delta_0, \delta_1)$: $\mathrm{In}(\Pi_2[1]) = \mathrm{Th}(B)$, and δ_1 is not applicable to $\mathrm{Th}(B)$ w.r.t. $\mathrm{Th}(B)$.

$\mathrm{In}(\Pi)$ collects all formulas in which we believe after application of the defaults in Π, while $\mathrm{Out}(\Pi)$ consists of all those formulas which we should avoid to believe for the sake of consistency. The following theorem states the fundamental relationship between the extensions of a default theory T and the closed successful processes of T. It should not be surprising, since successful processes avoid contradictions to the justifications of an already applied default, while closed process Π guarantee that $\mathrm{In}(\Pi)$ is closed under application of defaults in D. Its proof can be found in Appendix A.

11. THEOREM. *Let $T = (W, D)$ be a default theory. If Π is a closed successful process of T, then* $\mathrm{In}(\Pi)$ *is an extension of T. Conversely, for every extension E of T there exists a closed, successful process Π of T with $E = \mathrm{In}(\Pi)$.*

The extensions of a default theory T are thus just the In-sets of closed, successful processes of T. So, to determine extensions, we may simply apply defaults in an arbitrary order hoping that never a failure occurs. Having arrived at a failure situation, we must give up the current knowledge base and trace back within our process. There is no reason to control the selection of defaults in order to avoid failed processes.

If D contains only a finite number of defaults, there is no problem with closedness at all: as long as defaults are applicable we apply them. These leads to a closed process. Closed processes must (obviously) be calculated differently when D contains an infinite number of defaults. The definition of closed processes seems to require a look ahead to the final extension $\mathrm{In}(\Pi)$ already while constructing Π. This problem is simply resolved by the following lemma.

12. LEMMA. *An infinite process Π of T is closed iff every $\delta \in D$ that is applicable to* $\mathrm{In}(\Pi[k])$ *with respect to belief set* $\mathrm{In}(\Pi[k])$, *for infinitely many numbers k, is already contained in* Π.

The simple proof is an application of the compactness theorem of predicate logic. So, to achieve closed processes we must eventually apply each default which is, from some stage on, constantly demanding for application. This is nothing more than *fairness*, a situation commonly known from various fields dealing with concurring processes. If a survey of all possible processes of a default theory T is aspired, we may arrange all possible processes in a canonical manner within the so-called process tree of T.

13. DEFINITION. Let $T = (W, D)$ be a default theory. The *process tree of T* is a finite or infinite tree with edges labelled with defaults from D and nodes labelled with a theory I (the In-set built up so far) and a formula set O (the Out-set obtained so far) as follows:

- The root node is labelled with Th(W) and $\varnothing$.
- Consider a node N labelled with theory I and formula set O.
  If $I \cap O = \varnothing$ then N possesses a successor node $N(\delta)$ for every default $\delta = \varphi{:}\psi_1, \ldots, \psi_n/\chi$ such that δ does not already occur along the path from the root to the considered node and δ is applicable to I with respect to belief set I. The edge from N to $N(\delta)$ is labelled with δ, and $N(\delta)$ is labelled with $\mathrm{Th}(I \cup \{\chi\})$ and $O \cup \{\neg\psi_1, \ldots, \neg\psi_n\}$.
  If $I \cap O \neq \varnothing$ then N is a leaf.
  A *path* is be any maximal sequence of nodes starting at the root of T.

Note that the process tree of T may contain four types of paths with respect to the sequence Π of defaults along it:
- failed (thus finite) paths
- successful paths of finite length (these are automatically closed)
- successful paths of infinite length which are closed
- successful paths of infinite length which are not closed.

14. EXAMPLE
Consider $T = (\varnothing, \{\mathrm{true}{:}A/\neg A\})$. The process tree of T is found in Figure 1. T has no extensions. This is an example of a default application that a postiori invalidates a previously successful consistency check.

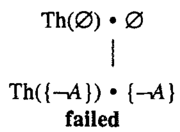

Fig. 1.

15. EXAMPLE
Consider $T = (\varnothing, D)$ with $D = \{\delta_0 = \mathrm{true}{:}P/\neg Q, \ \delta_1 = \mathrm{true}{:}Q/\neg P\}$. The process tree of T looks as shown in Figure 2. T has exactly two extension, $\mathrm{Th}(\{\neg P\})$ and $\mathrm{Th}(\{\neg Q\})$.

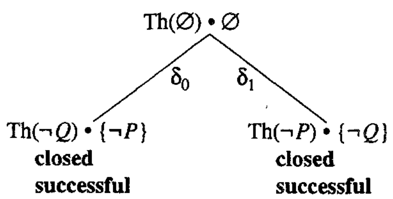

Fig. 2.

16. EXAMPLE
Consider the default theory $T = (\varnothing, \{\delta_0, \delta_1\})$ with $\delta_0 = \mathrm{true}{:}P/P$ and $\delta_1 = \mathrm{true}{:}Q/\neg P$. It has the only extension $\mathrm{Th}(\{\neg P\})$, as seen from its process tree in Figure 3.

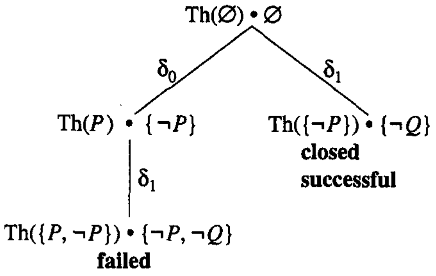

Fig. 3.

17. THEOREM. A default theory $T = (W, D)$ has an inconsistent extension iff W itself is inconsistent.

Proof. If W is inconsistent then $E = \text{Th}(W)$ coincides with $\Lambda_T(E)$, so E is an extension of T. Conversely, assume that T has the inconsistent extension E, that is, E is the set of all formulas. Since for every non-empty process Π of T with $E = \text{In}(\Pi)$ the intersection $\text{Out}(\Pi) \cap \text{In}(\Pi)$ is non-empty, it follows that the only successful process of T with $E = \text{In}(\Pi)$ is the empty process. Then, by definition of extensions, W is inconsistent. Note that in this case E is the only extension of T. ■

4. A STRAIGHTFORWARD PROLOG IMPLEMENTATION OF PROCESSES

Although we have now a method for operationally determining extensions, it may still be a little bit clumsy to write down the whole process tree of a bigger default theory. Fortunately, the operational model can be implemented in Prolog in a very simple way. For the sake of simplicity, we restrict ourselves to defaults with one justification. Defaults $A{:}B/C$ are represented as default(A, B, C), the negation symbol ¬ as ~. The Prolog code for determining extensions of a default theory looks as follows.

```
extension (W, D, E):– process(D, [ ], W, [ ], _, E, _).
process(D, Pcurrent, InCurrent, OutCurrent, P, In, Out):–
   element(default(A, B, C), D),
   not element(default(A, B, C), Pcurrent),
   sequent(InCurrent, [A]),
   not sequent(InCurrent, [~B]),
               % default A:B/C can be applied to In w.r.t. In
   process(D, [default(A, B, C)|Pcurrent], % extend current path P
               [C|InCurrent],
               [~B|OutCurrent], P, In, Out).
process(D, P, In, Out, P, In, Out):–   % P is a closed and successful process
   closed(D, P, In),
   success(In, Out).
```

```
closed (D, P, In):-                          % P is closed under D w.r.t. In
   not(element(default(A, B, C), D),
      not element(default(A, B, C), P),
      sequent(In, [A]),
      not sequent(In, [~B])).
success(In, Out):- not(element(B, Out), sequent(In,[B])).   %Th(In) ∩ Out = ∅
```

Missing are an input component for default theories, and a theorem prover (it is predicate sequent in the program text above; it could be Wang's algorithm for theories in pure propositional logic). Note that we do not build the deductive closure of In (it is clear why), so we must call sequent when testing success of a path P.

5. NORMAL DEFAULT THEORIES

Often, the consistency check in a default concerns just the formula to be derived, i.e. its consequent (see examples 1 and 3); such defaults with the same justification and consequent are called normal. Default theories consisting solely of normal defaults have some very desirable properties. Example 14 shows that general default theories need not have an extension. But theories consisting solely of normal defaults do always have at least one extension. In the proofs of the results below we shall see that processes are not only useful for the calculation of extensions, but also as a *theoretical tool* for proving properties of the logic.

18. DEFINITION. A default theory $T = (W, D)$ is called *normal* iff the defaults in D are of the form $\varphi:\psi/\psi$.

19. LEMMA. Let Π be a process for a normal default theory $T = (W, D)$ with consistent set W. Then Π is successful.

Proof. It is easy to see that for a process Π of a normal default theory (W, D), $\mathrm{In}(\Pi[i])$ is consistent, for all i such that $\Pi[i]$ is defined. If Π is non-empty then $\mathrm{In}(\Pi)$ is consistent. If Π is empty the assumption of consistency of W also implies that $\mathrm{In}(\Pi)$ is consistent. In any case, Π is successful. ■

COROLLARY. Let $T = (W, D)$ be a normal default theory. Then the process tree of T does not contain any failure path. In particular, every finite process for T can be extended to a closed, successful process of T. Thus, normal default theories always possess extensions.

20. THEOREM (Semi-monotonicity of normal theories). Let $T = (W, D)$ and $T' = (W, D')$ be normal default theories with $D \subseteq D'$. Then every extension E of T is contained in an extension E' of T'.

Proof. For inconsistent set W there is nothing to prove, so assume that W is consistent. Let E be an extension of T. Choose a closed process Π =

$(\delta_0, \delta_1, \delta_2, \ldots)$ of T such that $E = \mathrm{In}(\Pi)$. Obviously, Π is also a process of T'. (But note that Π is not necessarily a *closed process* of T'.)

Case 1. Π is a process of finite length. Then, by the corollary above Π can be extended to a closed, successful process Π' of T'.

Case 2. Π is a process of infinite length. Now we cannot extend Π as done in case 1, but must weave the application of the defaults in D' into Π *in such a way that the applicability of the old defaults in* Π *is not destroyed.* To do so we define a process $\Gamma = (\gamma_0, \gamma_1, \gamma_2, \ldots)$ of T' as follows:

(1) $\gamma_{2\cdot i} = \delta_k$, with k minimal such that $\delta_k \notin \{\gamma_0, \gamma_1, \ldots, \gamma_{2\cdot(i-1)}, \gamma_{2\cdot i-1}\}$
(2) $\gamma_{2\cdot i+1}$ is the first default δ in $D' - \{\gamma_0, \gamma_1, \ldots, \gamma_{2\cdot i}\}$ (referring to a fixed enumeration of D') such that δ is applicable to $\mathrm{In}(\gamma_0, \gamma_1, \ldots, \gamma_{2\cdot i})$ with respect to belief set $\mathrm{Th}(E \cup \mathrm{In}\,(\gamma_0, \gamma_1, \ldots \gamma_{2\cdot i}))$.

Since in case (2) we use belief set $\mathrm{Th}(E \cup \mathrm{In}(\gamma_0, \gamma_1, \ldots, \gamma_{2\cdot i}))$ instead of merely $\mathrm{In}(\gamma_0, \gamma_1, \ldots \gamma_{2\cdot i})$ the following properties are obvious from the definition of Γ:

- $\{\delta_0, \ldots, \delta_{i-1}\} \subseteq \{\gamma_0, \gamma_1, \ldots, \gamma_{2\cdot i-1}\}$
- the prerequisite of γ_i is contained in $\mathrm{In}(\gamma_0, \ldots, \gamma_{i-1})$
- $\mathrm{In}(\gamma_0, \gamma_1, \ldots, \gamma_{i-1}) \cup E$ is consistent.

So we obtain a process Γ of T' containing Π as a subsequence. We finally show that Γ is a closed process. Consider a default δ in D' that is applicable to $\mathrm{In}(\Gamma)$ with respect to belief set $\mathrm{In}(\Gamma)$. Since E is a subset of $\mathrm{In}(\Gamma)$ we may conclude that δ is applicable to $\mathrm{In}(\gamma_0, \ldots, \gamma_{i-1})$ with respect to belief set $\mathrm{In}(\gamma_0, \gamma_1, \ldots, \gamma_{i-1}) \cup E$, for infinitely many i. Having chosen in case (2) of the construction above the 'first default such that . . .' we are sure that δ will be selected at some stage of construction. ■

21. THEOREM. (Orthogonality of extensions). Let E and F be different extensions of a normal default theory $T = (W, D)$. Then $E \cup F$ is inconsistent.

Proof. Let E and F be different extensions of T. Assume that E is not contained in F and suppose that $E \cup F$ is consistent. Choose closed processes $\Pi = (\delta_0, \delta_1, \delta_2, \ldots)$ and $\Gamma = (\gamma_0, \gamma_1, \gamma_2, \ldots)$ of T such that $E = \mathrm{In}(\Pi)$ and $F = \mathrm{In}(\Gamma)$. Consider the least i such that $\delta_i \notin \Gamma$. Let δ_i be $\varphi{:}\psi/\psi$. By minimality of i, $\varphi \in \mathrm{In}(\delta_0, \ldots, \delta_{i-1}) \subseteq \mathrm{In}(\Gamma)$. Choose a number k such that $\varphi \in \mathrm{In}(\Gamma[k])$. Since $\psi \in E$, $\mathrm{In}(\Gamma[k]) \subseteq F$ and $E \cup F$ was assumed to be consistent it follows that δ_i is applicable to $\mathrm{In}(\gamma_0, \ldots, \gamma_k)$ with respect to $\mathrm{In}(\Gamma)$. But we know that $\delta_i \notin \Gamma$. This is only possible if $\neg\psi \in \mathrm{In}(\Gamma[1])$ for some $1 > k$. But then, $\psi \in E$ and $\neg\psi \in F$, a contradiction to our assumption that $E \cup F$ is consistent. ■

Orthogonality of extensions is not fulfilled in general for arbitrary default theories. As an example consider $T = (W, D)$ with $W = \emptyset$ and D $\{\mathrm{true}{:}\neg P/Q, \mathrm{true}{:}\neg Q/P\}$. Its extensions are $E = \mathrm{Th}(\{P\})$ and $F = \mathrm{Th}(\{Q\})$. $E \cup F$ is consistent.

6. A PROOF THEORY FOR NORMAL DEFAULT THEORIES

The process model in Section 3 applies defaults in a 'blind', bottom-up manner, meaning that no concrete goal exists. But what if a goal is given? The problem we are thinking about is the following: 'Given a formula φ, determine whether φ is included in at least one extension of default theory $T = (W, D)$' (in this case we say that φ is *weakly provable* from T). We would like to use a goal-oriented strategy, starting from the goal and working towards the truths using defaults.

22. EXAMPLE
Consider the default theory $T = (\{P\}, \{\delta_0 = P{:}Q/R,\ \delta_1 = R{:}Q/S,\ \delta_2 = \text{true}{:}\text{true}/\neg Q\})$. Given the goal ?- S, a goal-oriented approach would look as follows:

?-S
Use default δ_1; Q is consistent
?-R
Use default δ_0; Q is consistent
?-P: P is a fact, so we are through.

This could be a 'proof' we were looking for. Unfortunately, this approach cannot work, in general. The reason is that a process Π including the goal needs not be extendable to a closed, successful process of T. Indeed, even in our example (δ_0, δ_1) cannot be extended. Default δ_2 is applicable to $\text{In}(\Pi) = \text{Th}(P, R, S)$ w.r.t. $\text{In}(\Pi)$, leading to a failed process $\Pi' = (\delta_0, \delta_1, \delta_2)$, as $\neg Q \in \text{In}(\Pi') \cap \text{Out}(\Pi')$. In fact, S is not weakly provable from T, since it is not included in the only extension of T $\text{Th}(\{\neg Q\})$.

This kind of problem cannot occur when considering normal default theories, as we know from Section 5 that any process Π can be extended to a closed, successful process. So, if $\varphi \in \text{In}(\Pi)$, then $\varphi \in E$ for an extension E of the normal default theory. The following definition is a formalization of the approach outlined above.

23. DEFINITION. A *default proof* of φ in a normal default theory $T = (W, D)$ is a finite sequence $(D_0, \ldots, D_n)$ of subsets of D such that:

- φ follows from W and the consequents of D_0.
- For all $i < n$, the prerequisites of defaults in D_i follow from W and the consequents of D_{i+1}.
- D_n is the empty set.
- The set of all consequents of defaults in one D_i is consistent with W.

COROLLARY. φ has a default proof in a normal default theory T iff there is an extension E of T such that $\varphi \in T$.

24. EXAMPLE
Let $T = (\{P\}, (\delta_0{=}P{:}R/R, \delta_1{=}R{:}S/S\})$. $(\{\delta_1\}, \{\delta_0\}, \phi)$ is a default proof of S in T.

7. HISTORICAL AND BIBLIOGRAPHICAL REMARKS

Default logic was developed by Reiter (1980) and is established as one of the standard approaches to formalizing nonmonotonic reasoning. It is treated in all standard books on nonmonotonic logic, like (Besnard 1989, Brewka 1991a, Lucaszewicz 1990 or Marek and Truszczynski 1993). Through the basic concepts of the logic are treated in them, little attention has been given to *how to deal* with these concepts in an easy way. The idea of operational interpretation goes back to (Lucaszewicz 1990, Schwind 1990 or Levy 1991).

Several variants of default logic have been investigated in literature. Poole presented 1988 a simple but efficient system essentially equivalent with normal default theories without prerequisites, and demonstrated its applicability in many situations. Other variants of default logic include semi-normal defaults and ordered default theories (Etherington 1987), modified extensions (Lucaszewicz 1990), and disjunctive defaults Gelfond *et al.* (1991). A good overview is found in (Froidevaux and Menjin 1992).

Several implementational methods for default logic have been proposed, e.g. (Hopkins 1993, Antoniou and Langetepe 1994). They cover various paradigmata, like truth maintenance systems (Junker and Konolige 1990), resolution-based (Levy 1991), tableau-based (Schwind 1990) or graph-based (Dimopoulos and Magirou 1994). Of course, there are big differences in efficiency and representational power (most of them apply to some subclasses of default logic only).

Skeptical derivability of a formula φ from a default theory $T = (W, D)$ ($T \vdash \varphi$) means that φ is included in all extensions of T. One theoretical disadvantage of default logic is that skeptical derivability is *not cumulative*, as shown in Makinson (1989): It is possible that $(W, D) \vdash \varphi$, but not $(W, D) \vdash \psi \Leftrightarrow (W \cup \{\varphi\}, D) \vdash \varphi$ for some formula φ. This means that formulas skeptically derivable from a default theory may not be used as lemmata as in mathematics or predicate logic. The observation led to the development of cumulative versions of default logic (Brewka 1991b, Brewka 1992).

APPENDIX A: PROOF OF THEORM 11

Let Π be a closed, successful process of T and $E = \text{In}(\Pi)$. Then $\text{In}(\Pi) \cap \text{Out}(\Pi) = \emptyset$ and every $\delta \in D$ which is applicable to $\text{In}(\Pi) = E$ with respect to belief set E already occurs within Π. We show that $\Lambda_T(E) \subseteq E$. First, $W \subseteq E$ and $\text{Th}(E) = E$ by definition of the In-operator. Also, E is closed under D with respect to E, since Π was assumed to be closed. This show $\Lambda_T(E) \subseteq E$.

Next we show $E \subseteq \Lambda_T(E)$. By induction on k we show that $\text{In}(\Pi[k]) \subseteq \Lambda_T(E)$, for all k such that $\Pi[k]$ exists. For $k = 0$ we must show $\text{Th}(W) \subseteq \Lambda_T(E)$. This follows from the definition of $\Lambda_T(E)$ as a deductively closed set that contains W. Assume we have already shown $\text{In}(\Pi[k]) \subseteq \Lambda_T(E)$ and $\varphi:\psi_1, \ldots, \psi_n/\chi$ is the k-th element of Π. By definition of a process, $\varphi \in \text{In}(\Pi[k])$, and thus $\varphi \in E$. Also $\neg\psi_1, \ldots, \neg\psi_n \in \text{Out}(\Pi[k]) \subseteq \text{Out}(\Pi)$. Thus $\neg\psi_1, \ldots, \neg\psi_n \notin \text{In}(\Pi) = E$. This shows that $\varphi:\psi_1, \ldots, \psi_n/\chi$ is applicable to E with respect to belief set E. The definition of $\Lambda_T(E)$ implies that $\chi \in \Lambda_T(E)$. Thus $\mathrm{In}(\Pi[k+1]) = \mathrm{Th}(\mathrm{In}(\Pi[k]) \cup \{\chi\}) \subseteq \Lambda_T(E)$. Altogether, we have shown $E = \Lambda_T(E)$. Thus, E is an extension of T.

Conversely, consider an extension E of T. We choose an enumeration $\{\delta^0, \ldots, \delta^k, \ldots\}$ of the set of defaults D. Then we define a process Π of T such that $\mathrm{In}(\Pi[i]) \subseteq E$ and $\mathrm{Out}(\Pi[i]) \cap E = \emptyset$ for all i such that $\Pi[i]$ is defined. The definition is as follows (note that case 2 preserves the above property):

Let $\Pi[i]$ be already defined such that $\mathrm{In}(\Pi[i]) \subseteq E$ (*) and $\mathrm{Out}(\Pi[i]) \cap E = \emptyset$ (for $i = 0$, this is trivially true).

Case 1. Every $\delta \in D$ which is applicable to $\mathrm{In}(\Pi[i])$ with respect to belief set E already occurs in $\Pi[i]$. Then finish the construction of Π.

Case 2. There exists some $\delta \in D$ which is applicable to $\mathrm{In}(\Pi[i])$ with respect to E and does not occur in $\Pi[i]$. In the fixed enumeration of D choose the first such δ and append δ to $\Pi[i]$ to obtain $\Pi[i+1]$.

For case 1, we show that $E \subseteq \mathrm{In}(\Pi[i])$. Since $E = \Lambda_T(E)$ it suffices to show that $\mathrm{In}(\Pi[i])$ is a deductively closed set that contains W and is closed under D with respect to belief set E. The former properties are clear, the latter property by definition of case 1. Together with (*) we have $E = \mathrm{In}(\Pi[i])$. Thus, we may reformulate the property defining case 1 by saying that every $\delta \in D$ which is applicable to $\mathrm{In}(\Pi[i])$ with respect to belief set $\mathrm{In}(\Pi[i])$ already appears within $\Pi[i]$. Altogether we have shown that, in case 1, $\Pi[i]$ is a closed, successful process of T with $E = \mathrm{In}(\Pi[i])$.

Now we consider the case that the construction above yields a process Π of infinite length. We shown that Π is a closed process T. By construction, $\mathrm{In}(\Pi) \cap \mathrm{Out}(\Pi) = \emptyset$ since $\mathrm{In}(\Pi[i]) \subseteq E$ and $\mathrm{Out}(\Pi[i]) \cap E = \emptyset$ for all i. So Π is a successful process of T. Next we show that $\mathrm{In}(\Pi) = E$. One inclusion, namely $\mathrm{In}(\Pi) \subseteq E$, holds by construction. For the reversed inclusion, namely $E \subseteq \mathrm{In}(\Pi)$, we note that $\mathrm{In}(\Pi)$ is a deductively closed set that contains W and show that it is closed under D with respect to belief set E. For this, consider an arbitrary default $\delta = \varphi:\psi_1, \ldots, \psi_n/\chi$ in D with $\varphi \in \mathrm{In}(\Pi) \subseteq E$ and $\neg\psi_1, \ldots, \neg\psi \notin E$. Now we choose a number i such that $\varphi \in \mathrm{In}(\Pi[i])$ and every $\delta' \in \Pi$ that appears before δ in the fixed enumeration of D is already contained in $\Pi[i]$. By definition of case 2, δ is the default chosen in stage $i + 1$. Thus, its consequent χ is contained in $\mathrm{In}(\Pi)$. The argument above together with the just proved equation $E = \mathrm{In}(\Pi)$ also shows that Π is closed. ∎

REFERENCES

Antoniou, G. & Langetepe, E. (1994). Relating Some Classes of Default Logic to Normal Logical Programs with Standard Semantics. *Journal of Automated Reasoning* (submitted).

Besnard, P. (1989). *An Introduction to Default Logic.* Springer.

Bibel, W. (1984). Knowledge Representation from a Deductive Point of View. *Proc. IFAC Symposium Artificial Intelligence.* Pergamon Press: Oxford.

Brewka, G. (1991a). *Nonmonotonic Reasoning: Logical Foundations of Commonsense.* Cambridge University Press 1991.

Brewka, G. (1991b). Cumulative Default Logic – In Defence of Nonmonotonic Inference Relations. *Artificial Intelligence* **51**: 183–205.

Brewka, G. (1992). *A Framework for Cumulative Default Logics.* Technical Report TR-92-042, International Computer Science Institute, Berkeley.

Dimopoulos, Y. & Magirou, V. (1994). A Graph Theoretic Approach to Default Logic. *Information and Computation* (forthcoming).

Etherington, D. W. (1987). Formalizing Nonmonotonic Reasoning Systems. *Artificial Intelligence* **31**: 41–85.

Froidevaux, C. & Menjin, J. (1992). *A Framework for Default Logics.* Technical Report 755, Universite de Paris-Sud.

Gelfond, M., Przymusinska, H., Lifschitz, V. & Truszczynski, M. (1991). Disjunctive Defaults. *Proc. 2nd International Conference on Knowledge Representation and Reasoning.* Morgan Kaufmann.

Hopkins, M. S. (1993). Default Logic: Orderings and Extensions. In *Proc. European Conference on Symbolic and Quantitative Approaches to Uncertainty*, Springer LNCS 747.

Junker, U. & Konolige, K. (1990). Computing the Extensions of Autoepistemic and Default Logics with a TMS. In *Proc. AAAI-90*, 278–283.

Levy, F. (1991). Computing Extensions of Default Theories. In *Proc. European Conference on Symbolic and Quantitative Approaches to Uncertainty*, Springer LNCS 548.

Lucaszewicz, W. (1990). *Non-Monotonic Reasoning.* Ellis Horwood.

Makinson, D. (1989). General Theory of Cumulative Inference. In *Proc. 2nd International Workshop on Nonmonotonic Reasoning*, LNCS 346, Springer.

Marek, V. W. & Truszczynski, M. (1993). *Nonmonotonic Logic.* Springer 1993.

Mendelson, E. (1978). *Introduction to Mathematical Logic.* 2. ed., Van Nostrand: New York.

Moore, R. C. (1982). The Role of Logic in Knowledge Representation and Commonsense Reasoning. *Proc. AAAI-82.*

Nilsson, N. J. (1991). Logic and Artificial Intelligence. *Artificial* **47**: 31–56.

Poole, D. (1988). A Logical Framework for Default Reasoning. *Artificial Intelligence* **36**.

Reiter, R. (1980). A Logic for Default Reasoning. *Artificial Intelligence* **13**: 81–132.

Schwind, C. (1990). A Tableau-Based Theorem Prover for a Decidable Subset of Default Logic. In *Proc. 10th International Conference on Automated Deduction.* Springer LNC 449.

Sperschneider, V. & Antoniou, G. (1991). *Logic: A Foundation for Computer Science.* Addison-Wesley.