A Nonmonotonic Observation Logic

Frans Voorbraak*
Department of Computer Science
University of Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam
The Netherlands
fransv@wins.uva.nl

Abstract

A variant of Reiter's default logic is proposed as a logic for reasoning with (defeasible) observations. Traditionally, default rules are assumed to represent generic information and the facts are assumed to represent specific information about the situation, but in this paper, the specific information derives from defeasible observations represented by (normal free) default rules, and the facts represent (hard) background knowledge. Whenever the evidence underlying some observation is more refined than the evidence underlying another observation, this is modelled by means of a priority between the default rules representing the observations. We thus arrive at an interpretation of prioritized normal free default logic as an observation logic, and we propose a semantics for this observation logic. Finally, we discuss how the proposed observation logic relates to the multiple extension problem and the problem of sensor fusion.

1 Introduction

In this paper we propose a variant of Reiter's default logic [8] as a logic for reasoning with (defeasible) observations. A default theory consists of a set of facts and a set of default rules. Traditionally, default rules are assumed to represent general, or generic, information, such as 'typically, birds fly' and the facts are assumed to represent specific information about the situation, such as 'Tweety is a bird'. However, in this paper we consider the case where the specific information derives from defeasible observations, and the general information denotes (hard) background knowledge. These

*The investigations were carried out as part of the PIONIER-project Reasoning with Uncertainty, subsidized by the Netherlands Organization of Scientific Research (NWO), under grant pgs-22-262.

characteristics apply, for example, to the situation of an autonomous robot navigating through an environment using an a priori given map and its not completely reliable sensors.

The nonmonotonic logic proposed in this paper is motivated by the logic of vision proposed in [5], where perception reports are interpreted in inverse systems of first-order models approximating reality. A perception of ϕ is modelled as the truth of ϕ *given some approximation* plus a defeasible expectation that ϕ will also be true in more refined approximations. Perceptions based on more refined evidence can defeat this expectation.

We abstract from the particular language of [5] and propose prioritized normal free default logic for reasoning with defeasible observations. Each observation is modelled by a normal free default rule, i.e., a default rule which has a single justification equivalent to its consequent and a tautology as prerequisite. To take account of the fact that an observation can be defeated by another observation based on more refined evidence, we add a preference order on the default rules.

The remainder of this paper is organized as follows. In sections 2 and 3, we briefly review some basic properties of default logic, and in particular of the special case where default rules are normal and free. In section 4, we describe how normal free default logic with priorities between default rules can be interpreted as a logic for defeasible observations. A semantics for this observation logic is then presented in section 5. Finally, we discuss how the proposed observation logic relates to the multiple extension problem and the problem of sensor fusion.

2 Default Logic

We repeat some definitions and results (without proof) from [8]. Let L be some ordinary first-order language, and let Th denote an ordinary first-order consequence operation for that language. (The dependence on L and Th of the definitions and results below will mostly remain implicit.) Formulas of L are denoted by $\alpha, \beta, \ldots$. In examples, propositional letters $p, q, \ldots$ are used, which are supposed to be included in the language.

Definition 1 *A* default rule *is an expression of the form:*

$$\frac{\alpha : \beta_1, \ldots, \beta_n}{\omega} (n \geq 1).$$

The formula α is called the prerequisite*, $\beta_1, \ldots, \beta_n$ are called the* justifications*, and ω the* consequent *of the default rule. A default rule is called* closed *iff all the formulas appearing in the rule are closed. If D is a set of default rules, then $CONSEQ(D)$ denotes the set of consequents of the default rules of D.*

The default rule $\alpha : \beta_1, \ldots, \beta_n/\omega$ has the following intended meaning: If α can be derived and each of the $\beta_1, \ldots, \beta_n$ is consistent with what is derivable, then derive ω. (This is the standard interpretation. Our interpretation of the default rules representing observations will be slightly different. See section 7.)

Definition 2 *A* default theory $\mathcal{D}$ *is a pair* $\langle D, \Gamma\rangle$*, where* D *is a set of default rules and* Γ *is a set of sentences of* L*. A default theory is called* closed *iff all its default rules are closed.*

Definition 3 *An* extension *of the closed default theory* $\langle D, \Gamma\rangle$ *is a fixed point of the function* $f : 2^L \to 2^L$ *given by:* $f(\Sigma)$ *is the smallest set such that*

D1 $\Gamma \in f(\Sigma)$

D2 $Th(f(\Sigma)) = f(\Sigma)$

D3 *If* $\alpha : \beta_1, \ldots, \beta_n/\omega \in D$*,* $\alpha \in f(\Sigma)$ *and for all* $i \in \{1, \ldots, n\}$*,* $\neg\beta_i \notin \Sigma$*, then* $\omega \in f(\Sigma)$*.*

An extension of a default theory $\langle D, \Gamma\rangle$ is intended to represent a reasonable state of belief based on the default rules in D and facts in Γ. Below, we give a more informal characterization of default extensions, which is essentially a reformulation of the semantics for default logic given in [3]

Let Σ be a logically closed set of sentences of L. The default rule $\alpha : \beta_1, \ldots, \beta_n/\omega$ is called *applicable* in Σ iff $\alpha \in \beta_1$ and for every β_i, $\neg\beta_i \notin \Sigma$. If $d = \alpha : \beta_1, \ldots, \beta_n/\omega$ is applicable in Σ, then the result of applying d to Σ is the set $Th(\Sigma \cup \{\omega\})$. An extension E of a closed default theory $\langle D, \Gamma\rangle$ is the result of sequentially applying to $Th(\Gamma)$ the default rules of a subset D' of D such that every default rule of D' is applicable in E, and E contains the consequent of each default rule applicable in E. D' is called the set of *generating default rules* of E.

A default theory may have more than one extension, and there exist default theories without extensions. For example, $\langle\{\top : p/p, \top : \neg p/\neg p\}, \emptyset\rangle$ has two extensions, namely $Th(\{p\})$ and $Th(\{\neg p\})$, and $\langle\{\top : \neg p/p\}, \emptyset\rangle$ does not have an extension.

So far, we have only defined the notion of extension for closed default theories. In general, the extensions of a default theory $\langle D, \Gamma\rangle$ are defined to be the extensions of the closed default theory $\langle D', \Gamma\rangle$ where, roughly speaking, D' is obtained from D by taking all closed instances of D. (See [8, 10] for details.) From now on, we will assume, without loss of generality, that default theories are closed.

We have seen that not every default theory has an extension. However, there is a natural class of default theories with at least one extension:

Definition 4 *A default rule is called* normal *iff it has a single justification and this justification is equivalent (in first-order logic) to its consequent. A default theory is called* normal *iff all its default rules are normal.*

Proposition 1 (existence of extensions) *Every normal default theory has an extension.*

In addition to having at least one extension, a normal default theory has several other nice properties. Some of these properties are listed below. None of the listed properties holds for arbitrary default theories.

Proposition 2 (orthogonality of extensions) *If E and F are distinct extensions of a normal default theory, then $E \cup F$ is inconsistent.*

Proposition 3 (uniqueness of extensions) *If $\langle D, \Gamma \rangle$ is a normal default theory and $\Gamma \cup \mathit{CONSEQ}(D)$ is consistent, then $\langle D, \Gamma \rangle$ has a unique extension.*

Proposition 4 (semi-monotonicity) *Let D and D' be sets of normal defaults such that $D \subseteq D'$, and let E be an extension of $\langle D, \Gamma \rangle$. Then $\langle D', \Gamma \rangle$ has an extension $E' \supseteq E$.*

3 Free Default Logic

In this section, we consider a fragment of default logic in which the prerequisite of a default rule is assumed to be a tautology. Such a rule is called (prerequisite-)free.

Definition 5 *A default rule is called* free *iff its prerequisite is a tautology (in first-order logic). A default theory is called* free *iff all its default rules are free.*

Free default theories have been studied in [1, 9, 10]. (In [1] the free default rules are also assumed to be normal, which is the case we are especially interested in.) In [9, 10] it is argued that rational agents adhering to Savage's 'sure-thing principle' should be willing to replace arbitrary default rules by free default rules. In general, the resulting default theory will not be equivalent to the original one. However, it is not hard to show that any default theory is equivalent to a free default theory.

An advantage of free default logic is that it allows reasoning by cases.

Example 1 *Let $\mathcal{D} = \langle \{p : q/q, \neg p : q/q\}, \emptyset \rangle$. The default theory $\mathcal{D}$ does not allow the obvious conclusion q since its unique extension is $Th(\emptyset)$. In [9, 10] it is proposed to replace $\mathcal{D}$ by $\langle \{\top : p \wedge q/p \rightarrow q, \top : \neg p \wedge q/\neg p \rightarrow q\}, \emptyset \rangle$, which has the unique extension $Th(\{q\})$.*

In this paper, we are interested in *normal* free default logic, which is a particularly well-behaved fragment of default logic. For example, the extensions of normal free default theories have a relatively simple characterization.

Proposition 5 *Let $\mathcal{D} = \langle D, \Gamma \rangle$ be a normal free default theory. Then E is an extension of $\mathcal{D}$ iff $E = Th(\Gamma \cup CONSEQ(D'))$, for some maximal $D' \subseteq D$ such that $\Gamma \cup CONSEQ(D')$ is consistent.*

Proof. The result immediately follows from the informal characterization of extensions given in section 2 and from the fact that a normal free default rule $\top : \omega / \omega$ is applicable in Σ iff $\Sigma \cup \{\omega\}$ is consistent. ∎

Corollary 6 *If $\mathcal{D} = \langle D, \Gamma \rangle$ is a normal free default theory and E is an extension of $\mathcal{D}$, then E is complete with respect to $CONSEQ(D)$. That is, if $\omega \in CONSEQ(D)$, then either $\omega \in E$ or $\neg\omega \in E$.*

Another interesting property of normal free default logic is that it induces a *cumulative* consequence operation. The fact that the consequence operation induced by unrestricted default logic is not cumulative has been considered a drawback of default logic, and several authors have looked at ways to modify default logic to make its consequence operation cumulative.

Definition 6 *An operation $Cn : 2^L \to 2^L$ is called a* cumulative consequence operation *for L iff it satisfies the following four properties, for arbitrary $\Gamma, \Delta \subseteq L$.*

inclusion $\Gamma \subseteq Cn(\Gamma)$

Th-invariance $Cn(\Gamma) = Cn(Th(\Gamma))$

cumulative transitivity $\Gamma \subseteq \Delta \subseteq Cn(\Gamma) \Rightarrow Cn(\Delta) \subseteq Cn(\Gamma)$

cumulative monotonicity $\Gamma \subseteq \Delta \subseteq Cn(\Gamma) \Rightarrow Cn(\Gamma) \subseteq Cn(\Delta)$

Definition 7 *For any set D of default rules and for any set Γ of sentences of L, define $Cn_D(\Gamma) = \bigcap\{E : E$ is an extension of $\langle D, \Gamma \rangle\}$.*

$Cn_D(\Gamma)$ is the set of sceptical, or cautious, conclusions of $\langle D, \Gamma \rangle$. It is shown in [6] that for arbitrary normal default rules Cn_D does not satisfy cumulative monotonicity. However, this property does hold if one additionally assumes the default rules to be free.

Proposition 7 *Let D be a set of normal free default rules. Then Cn_D is a cumulative consequence operation for L.*

Proof. That Cn_D satisfies inclusion and Th-invariance is immediate. In [6] it is shown that Cn_D satisfies cumulative transitivity (even without assuming the defaults of D to be normal and free). To prove cumulative monotonicity, assume that $\Gamma \subseteq \Delta \subseteq Cn_D(\Gamma)$, and that E is an extension of $\langle D, \Delta\rangle$. It is sufficient to show that E is also an extension of $\langle D, \Gamma\rangle$.

By proposition 5, $E = Th(\Delta \cup CONSEQ(D'))$, for some maximal $D' \subseteq D$ such that $\Delta \cup CONSEQ(D')$ is consistent. Since $\Gamma \subseteq \Delta$, $\Gamma \cup CONSEQ(D')$ is also consistent. Let $D'' \supseteq D'$ be a maximal subset of D such that $\Gamma \cup CONSEQ(D'')$ is consistent. Then $E' = Th(\Gamma \cup CONSEQ(D''))$ is an extension of $\langle D, \Gamma\rangle$. We have $E \subseteq E'$, since $\Delta \subseteq Cn_D(\Gamma) \subseteq E'$ and $D' \subseteq D''$. We also have $E' \subseteq E$, since otherwise there would be a default rule in $D'' - D'$ such that its consequent is consistent with E, but not included in E. It follows that E is also an extension of $\langle D, \Gamma\rangle$. ∎

In the following section, we argue that an prioritized version of normal free default logic can be interpreted as a logic for defeasible observations.

4 Observation Logic

Consider a robot making an observation. In order to (symbolically) reason with the observation, the subsymbolic information acquired by the sensors has to be translated into a symbolic language. We assume that the observation systems of the robot translate the raw sensory data into propositions expressed in a first-order language L. Each observation is based on evidence which is limited in range and resolution, and which is therefore only directly related to (part of) an *approximation* of reality. An observation induces an expectation that what is observed will also hold for more refined approximations, but this expectation can be defeated by observations based on more refined evidence.

Let ϕ_e denote the observation that ϕ based on evidence e. We can interpret this, analogous to the rendition of perception reports in [5], as ϕ being true *given the approximation of reality induced by* e plus a defeasible expectation that ϕ will also be true in more refined approximations, including reality itself. For example, assume we have the observations ϕ_e and $\psi_{e'}$, where ϕ and ψ are incompatible (given the a priori knowledge) and e' is more refined than e. Then the expectation induced by ϕ_e is defeated by $\psi_{e'}$.

The (defeasible) expectation accompanying the observation ϕ_e intuitively seems to imply the (normal) default rules $\phi_e : \phi_{e'}/\phi_{e'}$, for each e' that is more refined than e. Since these default rules are only added to the theory when also the prerequisite

ϕ_e is added, one can use the normal free default rules $\top : \phi_{e'}/\phi_{e'}$ instead. In the case of the observations ϕ_e and $\psi_{e'}$, where ϕ and ψ are incompatible and e' is more refined than e, it then is clear that the expectation of ϕ_e is blocked by the truth of ψ at the approximation of reality induced by e'.

If we abstract in our formal language from the evidence on which the observation is based, then an observation can simply be modelled as a normal free default rule $\top : \phi/\phi$. To properly capture the interaction of different observations, one can mirror the relation between the underlying bodies of evidence in an ordering on the corresponding default rules.

Definition 8 *A* prioritized default theory *is a triple $\langle D, \Gamma, < \rangle$, where $\langle D, \Gamma \rangle$ is a default theory and $<$ is a well-founded strict partial order on D. (That is, $<$ is an irreflexive, transitive relation with no infinite descending chain $d_1 > d_2 > \ldots$.)*

Definition 9 *An* observation theory *is a prioritized normal free default theory.*

If d and d' are default rules, then the intended meaning of $d < d'$ is that d is stronger than d' because d represents an observation based on evidence which is more refined than the evidence on which the observation represented by d' is based. (This reading of $<$ is in accordance with the usually chosen direction of prefence relations in preferential semantics, where the preferred models are traditionally called minimal models. Unfortunately, if one has the context of information orderings in mind, then it is natural to expect the reverse reading.)

The assumption of $<$ being well-founded is necessary for the definitions of extension given below. The assumption (essentially the same as *stopperedness* of [6] and *smoothness* of [4]) seems not too restricitive for our purpose, since it is automatically satisfied if D is finite, and at any time a robot can have made only a finite number of observations. We first repeat from [2] the general definition of an extension of a prioritized default logic using a well-ordering $\prec$ compatible with $<$.

Definition 10 *A* well-ordering *$\prec$ on D is a well-founded (complete) order on D. A well-ordering $\prec$ on D is called* compatible *with $<$ on D iff for every $d, d' \in D$, $d < d'$ implies $d \prec d'$.*

Definition 11 *Let $\langle D, \Gamma, < \rangle$ be a prioritized default theory. The set E is called an $<$-compatible extension of $\langle D, \Gamma, < \rangle$ iff E is an extension of $\langle D, \Gamma \rangle$ generated by a well-ordering $\prec$ compatible with $<$.*

Intuitively, an extension of $\langle D, \Gamma \rangle$ generated by a well-ordering $\prec$ is obtained by sequentially applying, starting from $Th(\Gamma)$, the $\prec$-minimal default rule among the non-applied applicable default rules of D. (See [2, 7] for details.) By choosing a well-ordering compatible with $<$, the stronger defaults (with higher priorities) are applied first, thereby making weaker incompatible defaults non-applicable.

Example 2 *Consider the observations p_e and $q_{e'}$, and assume that it is a priori known that p and q are incompatible, i.e., $\neg(p \wedge q)$. Further assume that e' is more refined evidence than e. This situation is modelled by the observation theory $\mathcal{D}$ consisting of the default theory $\langle\{\top : p/p, \top : q/q\}, \{\neg(p \wedge q)\}\rangle$ with priority $\top : q/q < \top : p/p$. The unique extension of $\mathcal{D}$ is $Th(\{\neg p, q\})$. Without the priority, both $Th(\{\neg p, q\})$ and $Th(\{\neg q, p\})$ would have been extensions.*

To make the example more concrete, suppose that p_e denotes a robot's observation, based on some sonar readings, that a particular door is closed, and that $q_{e'}$ denotes the observation, based on a videocamera image in addition to the previously mentioned sonar readings, that the door is open. Then the robot should of course conclude that the door is open.

In the above example, the same conclusion is warranted if the evidence e' is just *stronger* or *more reliable* than e, rather than more refined. However, as we will illustrate later on, these different interpretations of the orderings can result in different conclusions in more complex examples.

On can imagine that having different interpretations of the ordering between default rules results in having different definitions of extensions of a prioritized default theory. In fact, the above standard definition of $<$-compatible extension will *not* be used as our notion of extension for observation theories.

Before we discuss which notion of extension we regard as appropriate for our purpose, let us consider an alternative way of obtaining extensions of prioritized default theories, namely by first computing an (ordinary default) extension using only the (unordered) default rules with maximal priority, and subsequently computing an extension using the default rules of the next priority level, et cetera. For observation theories this can be made more precise as follows.

Definition 12 *Let $\mathcal{D} = \langle D, \Gamma, <\rangle$ be an observation theory. Define $D_0 = \emptyset$, and let, for $i \geq 1$, D_i be the set of $<$-minimal elements of $D - \cup_{0 \leq j \leq i-1} D_j$. The set $Th(\Gamma)$ is called an 0-extension of $\mathcal{D}$. For $i \geq 1$, E is called an i-extension of $\mathcal{D}$ iff $E = Th(E' \cup CONSEQ(D'))$, for some $(i-1)$-extension E' of $\mathcal{D}$, and some maximal $D' \subseteq D_i$ such that $E' \cup CONSEQ(D')$ is consistent. Then E is called a layered extension of $\mathcal{D}$ iff $E = \cup_{i \geq 0} E_i$, where, for each $i \geq 0$, E_i is an i-extension of $\mathcal{D}$ and $E_i \subseteq E_{i+1}$.*

The set of layered extensions of an observation theory $\langle D, \Gamma, <\rangle$ is a subset of its set of $<$-compatible extensions.

Proposition 8 *Let $\mathcal{D} = \langle D, \Gamma, < \rangle$ be an observation theory. Every layered extension of $\mathcal{D}$ is a $<$-compatible extension of $\mathcal{D}$.*

Proof. Let E be a layered extension of $\mathcal{D} = \langle D, \Gamma, < \rangle$. Then E can be obtained from Γ by sequentially applying the default rules from a set $D' \subseteq D$ such that for every $d, d' \in D'$, $d < d'$ implies that d is applied before d'. Moreover, for any $d \in D - D'$, its justification is inconsistent with $\Gamma \cup CONSEQ(D''))$, where $D'' = \{d' \in D' : d' \leq d\}$. Let $\prec$ be a well-ordering obtained by refining $<$ in such a way that. whenever d and d' are incomparable according to $<$ and $d \in D'$ and $d' \notin D'$, it should hold that $d \prec d'$. The order between other pairs of $<$-incomparable default rules can be chosen arbitrarily. Then E is an extension of $\langle D, \Gamma \rangle$ generated by $\prec$, and since $\prec$ is compatible with $<$, E is a $<$-compatible extension of $\mathcal{D}$. ∎

The following example shows that the reverse of this proposition is not true.

Example 3 *Consider the observations p_e, $q_{e'}$ and $r_{e''}$, and assume that it is a priori known that p, q and r are jointly incompatible, i.e., $\neg(p \wedge q \wedge r)$. Further assume that e' is more refined evidence than e (and e'' is incomparable to the other bodies of evidence). This situation is modelled by the observation theory $\mathcal{D}$ consisting of the default theory $\langle \{\top : p/p, \top : q/q, \top : r/r\}, \{\neg(p \wedge q \wedge r)\} \rangle$ with priority $\top : q/q < \top : p/p$. The unique layered extension of $\mathcal{D}$ is $Th(\{\neg p, q, r\})$. However, both $Th(\{\neg p, q, r\})$ and $Th(\{p, q, \neg r\})$ are $<$-compatible extensions.*

The following example indicates a problem with both layered extension and $<$-compatible extensions in case $d < d'$ is interpreted to mean that default rule d is based on more refined evidence than d'.

Example 4 *Consider the observations p_e, $q_{e'}$ and $r_{e''}$, and assume that it is a priori known that q is incompatible both with p and with r, i.e., $\neg(p \wedge q)$ and $\neg(r \wedge q)$. Further assume that e'' is more refined evidence than e' and e' is more refined evidence than e. This situation is modelled by the observation theory $\mathcal{D}$ consisting of the default theory $\langle \{\top : p/p, \top : q/q, \top : r/r\}, \{(\neg(p \wedge q)) \wedge (\neg(r \wedge q))\} \rangle$ with priority $\top : r/r < \top : q/q < \top : p/p$. This theory has one $<$-compatible extension $Th(\{p, \neg q, r\})$ and this extension is also layered.*

To make the example more concrete, consider a robot in the process of finding out whether a particular door is open. Let us say that the door has to be open at least 70 degrees to be called open, since otherwise the robot cannot safely pass the door. Suppose that p_e denotes the robot's observation, based on some videocamera image taken from considerable distance from the door, that the door is 84 degrees open, that $q_{e'}$ denotes the observation based on some sonar readings in addition to the videocamera image, that the door is closed, and that $r_{e''}$ denotes the observation, *based in addition to the previously mentioned sensor information on a videocamera image taken at close proximity to the door, that the door is open (without being able to specify the exact angle of the door). Then both $<$-compatible and layered extensions allow the conclusion that the door is 84 degrees open.*

The problem with this conclusion is that the observation p_e is defeated by $q_{e'}$, and it is not clear whether the fact that $q_{e'}$ is in turn defeated by $r_{e''}$ justifies re-establishing the conclusions based on p_e. In fact, if e'' is more refined than e, then the information supporting p in the first observation is also present in e''. Therefore, if one assumes that r is the most specific proposition (relative to the background knowledge) supported by the observation $r_{e''}$, then p should not be derivable.

It is important to note that in the above example we introduced the assumption that the proposition mentioned in the observation is the *most specific* proposition (relative to the bacground knowledge) that is justified given the evidence on which the observation is based. This assumption and its relation to the treatment of the problem of sensor fusion will be discussed in more detail in section 6.

Given the above assumption, it is natural to consider only the default rules in the top layer of the ordering, and propose 1-extensions (as defined in definition 12) as the appropriate extensions of an observation theory. Of course, any 1-extension of an observation theory $\langle D, \Gamma, <\rangle$ is equivalent to an (ordinary Reiter) extension of the normal free default theory $\langle D_1, \Gamma\rangle$, where D_1 is the set of $<$-minimal default rules of D, as defined in definition 12.

At first sight, one might think that, by the semi-monotonicity of normal default logic (see proposition 4), the resulting logic will have a monotonic behaviour. However, this is prevented by the ordering on defaults, since adding an observation does not necessarily result in a monotonic increase in the $<$-minimal default rules.

Example 5 *Consider the situation of example 4. The observation theory $\mathcal{D}$ consisting of the default theory $\langle\{\top : p/p, \top : q/q, \top : r/r\}, \{(\neg(p \wedge q)) \wedge (\neg(r \wedge q))\}\rangle$ with priority $\top : r/r < \top : q/q < \top : p/p$ has a unique 1-extension, namely $Th(\{r, \neg q\})$. If we consider the situation without the observation modelled by $\top : r/r$, then the resulting observation theory has the unique 1-extension $Th(\{\neg p, q, \neg r\})$. That is, without the evidence of the videocamera at close proximity to the door, the conclusion that the door is closed would be justified.*

In fact, the logic obtained by considering 1-extensions of observation theories is an interesting nonmonotonic formalism, in which it is for example possible to 'withdraw' an observation p_e (where p cannot be deduced from the background knowledge) by adding the observation $\top_{e'}$ with e' more refined than e. (More generally, any $\phi_{e'}$, where ϕ together with the background knowledge does not imply p and e' more refined than e has a similar effect.)

However, since the assumption underlying the choice for 1-extensions may not be appropriate in all circumstances, we propose in the next section a rather general semantics for observation theories that allows us to consider different variants of observation logic next to the logic induced by 1-extensions.

5 Semantics

An observation theory $\langle D, \Gamma, < \rangle$ is called consistent iff Γ is consistent. Unless stated otherwise, we will from now on assume that observation theories are consistent, and moreover, we will assume that for any observation theory $\langle D, \Gamma, < \rangle$ the facts Γ are consistent with every consequent of a default rule $d \in D$. These assumptions are in accordance with our treatment of Γ as 'hard' background knowledge.

As models of observation theories we essentially use directed sets of first-order models, where the ordering between the first-order models representing the observations reflects the (refinement) ordering of (the evidence underlying) the observations. As a reminder we include the definition of a directed set.

Definition 13 *Let $\geq$ be a partial order on the set X. Then $\langle X, \geq \rangle$ is called a* directed set *iff every finite subset of X has an upper bound in X, i.e., for every finite $Y \subseteq X$, there exists an element $x \in X$ such that for every $y \in Y$, $x \geq y$.*

Definition 14 *Let $\mathcal{D} = \langle D, \Gamma, < \rangle$ be an observation theory. An* observation model $\mathbf{M}$ *of $\mathcal{D}$ is a tuple $\langle M, \geq, obs \rangle$, where $\langle M, \geq \rangle$ is a directed set of first-order models satisfying Γ, and obs is an injection $D \to M$ such that the following conditions are satisfied.*

- *$obs(d)$ satisfies the consequent of d*
- *if $d < d'$, then $obs(d) > obs(d')$.*

For any first-order sentence ϕ, we say that $\mathbf{M} \models \phi$ iff there exists a model $\mathcal{M} \in M$ such that every model $\mathcal{M}' \geq \mathcal{M}$ satisfies ϕ. For any set S of observation models, we write $Th(S)$ for the set of first-order formulas that are valid in every model of S.

There are many different ways to extend the scope of the satisfaction relation to include the default rules. Perhaps the most straightforward way is to define, for any default rule $d \in D$, $\mathbf{M} \models d$ iff every model $\mathcal{M} \geq obs(d)$ satisfies the consequent of d. Observation models where the satisfaction relation is thus extended are called straight. *More generally, observation models where the satisfaction relation is (in some way) extended to include default rules are called* extended.

An observation model is *not* a classical first-order model. In fact, an observation model plays a role which is roughly similar to a possible worlds, or Kripke, model in modal epistemic logic, namely as a model of an agent's belief set such that the agent's belief set corresponds to the set of formulas valid in all its models. More precisely, an observation model of an observation theory models the beliefs of an agent who believes all the facts and a subset of the conclusions of the default rules present in the observation theory. Since we consider agents who, by default, believe the observations, we will be interested in observation models which maximize the set of valid default conclusions, or rather, the set of applied or valid default *rules*.

Notice that the second condition of definition 14 involves a change in direction of the orderings. We feel this is appropriate, since the ordering on the models should be viewed as an information ordering. The intuitive reading of $\mathcal{M}' > \mathcal{M}$ is that $\mathcal{M}'$ is more refined than $\mathcal{M}$. Notice however, that we have not implemented a formal notion of refinement as in the inverse refining systems of models of [5]. An (information) ordering between models mirroring the preference relation $<$ between default rules will be called an information image of $<$.

Definition 15 *If $\mathcal{D} = \langle D, \Gamma, < \rangle$ is an observation theory, then a partial order $>'$ on the set M of first-order models of an observation model $\mathbf{M}$ of $\mathcal{D}$ is called an* information image *of $<$ iff it holds that $obs(d) >' obs(d')$ whenever $d < d'$.*

It is clear that there are two kinds of observation models. An observation model $\langle M, \geq, obs \rangle$ either has a *top*, i.e., a model $\mathcal{M}_\top \in M$ which is more refined than any other element of M, or M contains for any of its elements an infinite sequent of refinements. In the latter case, we call the observation model *top-less*. These top-less observation models are more like the refining inverse systems of models of [5] than the observation models with a top are. In [5], each first-order model in the refining system is viewed as an approximation to reality which itself is assumed to be infinitely precise and to refine each of the models in the refining system. However, in an observation model the top (or any other element) does not represent the reality (or an approximation of it) but (a partial description of) a world considered possible by the observing agent.

It is easy to see that in definition 14 we could have restricted ourselves to top-less observation models, since any observation model is equivalent to a top-less one. (Just add an infinite sequence of copies of the top-element on top of the top-element.) However, below we will consider observation models where the satisfaction of default rules is defined differently than for the straight observation models of definition 14, and for some of these variants it does make a difference whether observation models have a top or not.

The directedness condition ensures that default rules interact in the sense that a set of incompatible default rules cannot be satisfied in an observation model. As mentioned before, the most interesting observation models are those in which the set of valid default rules are maximized. In accordance with tradition, we will call these observation models *minimal.*

Definition 16 *Let* $\mathbf{M}$ *be a straight observation model of* $\mathcal{D} = \langle D, \Gamma, < \rangle$ *and let* $\mathbf{M}(\mathcal{D}) = \{d \in D : \mathbf{M} \models d\}$*.* $\mathbf{M}$ *is called* minimal *iff there exists no straight observation model* $\mathbf{M}'$ *of* $\mathcal{D}$ *such that* $\mathbf{M}(\mathcal{D}) \subset \mathbf{M}'(\mathcal{D})$*. We write* $[\mathbf{M}]_{\mathcal{D}}$ *for the equivalence class of straight observation models* $\mathbf{M}'$ *of* $\mathcal{D}$ *such that* $\mathbf{M}'(\mathcal{D}) = \mathbf{M}(\mathcal{D})$*.*

The above defined notions can be defined analogously for types of extended observation models other than the straight observation models. The following property is suggested as a minimal requirement on ways to extend the satisfaction relation of observation models to default rules.

Definition 17 *A class of extended observation models* $\mathbf{M}$ *of* $\mathcal{D} = \langle D, \Gamma, < \rangle$ *is called* regular *iff for every element* $\mathbf{M}$ *of the class,* $Th([\mathbf{M}]_{\mathcal{D}}) = Th(\Gamma \cup CONSEQ(\mathbf{M}(\mathcal{D})))$*.*

Lemma 9 *The class of straight observation models of an observation theory is regular.*

Proof. Assume $\mathbf{M}$ is a straight observation model of $\mathcal{D} = \langle D, \Gamma, < \rangle$. We will first show that $Th(\Gamma \cup CONSEQ(\mathbf{M}(\mathcal{D}))) \subseteq Th([\mathbf{M}]_{\mathcal{D}})$. Suppose $\phi \in Th(\Gamma \cup CONSEQ(\mathbf{M}(\mathcal{D})))$. Then $\phi \in Th(\Sigma)$, for some finite subset Σ of $\Gamma \cup CONSEQ(\mathbf{M}(\mathcal{D}))$. By the directedness condition, any $\mathbf{M}' \in [\mathbf{M}]_{\mathcal{D}}$ contains a model $\mathcal{M}$ such that for all $\mathcal{M}' \geq \mathcal{M}$, $\mathcal{M}' \models \Sigma$, and therefore $\mathcal{M}' \models \phi$. Thus $\phi \in Th([\mathbf{M}]_{\mathcal{D}})$.

To prove the inclusion in the other direction, it is sufficient to show that for any first-order model $\mathcal{M}$ of $\Gamma \cup CONSEQ(\mathbf{M}(\mathcal{D}))$, there exists a first-order equivalent observation model in $\mathbf{M}(\mathcal{D})$. This is easy, since adding $\mathcal{M}$ as top to any element of $\mathbf{M}(\mathcal{D})$ will give the desired observation model equivalent to $\mathcal{M}$. ∎

We now introduce an alternative kind of observation models, called the one level observation models, in order to model the 1-extensions of observations theories.

Definition 18 *Let* $\mathcal{D} = \langle D, \Gamma, < \rangle$ *be an observation theory. An* one level observation model $\mathbf{M}$ *of* $\mathcal{D}$ *is an observation model* $\langle M, \geq, obs \rangle$ *where we say that* $\mathbf{M} \models d$ *iff every model* $\mathcal{M} \geq obs(d)$ *satisfies the consequent of* d*, and there exists no* $d' \in D$ *such that* $obs(d') > obs(d)$*.*

Analogous to the proof of lemma 9 it can be shown that the class of one level observation models of an observation theory is regular. More importantly, there is an exact correspondence between minimal one level observation models of an observation theory and its 1-extensions.

Proposition 10 *The 1-extensions of an observation theory correspond with its minimal one level observation models. That is, E is an 1-extension of $\mathcal{D} = \langle D, \Gamma, < \rangle$ iff $E = Th([\mathbf{M}]_{\mathcal{D}})$ for some minimal one level observation model $\mathbf{M}$ of $\mathcal{D}$.*

Proof. Let E be an 1-extension of $\mathcal{D} = \langle D, \Gamma, < \rangle$. Then $E = Th(\Gamma \cup \mathit{CONSEQ}(D'))$, for some maximal $D' \subseteq D_1 = \{d \in D : d \text{ is } <\text{-minimal } \}$ such that $\Gamma \cup \mathit{CONSEQ}(D')$ is consistent. By the consistency of $\Gamma \cup \mathit{CONSEQ}(D')$ there exists an one level observation model $\mathbf{M}$ of $\mathcal{D}$ such that $\mathbf{M}(\mathcal{D}) = D'$. By the maximality of D', $\mathbf{M}$ is minimal, and by the regularity of one level observation models, $Th([\mathbf{M}]_{\mathcal{D}}) = Th(\Gamma \cup \mathit{CONSEQ}(D'))$ $= E$.

To prove the other direction, suppose $E = Th([\mathbf{M}]_{\mathcal{D}})$ for some minimal one level observation model $\mathbf{M}$ of $\mathcal{D}$. Then, by the regularity of one level observation models, $E = Th(\Gamma \cup \mathit{CONSEQ}(\mathbf{M}(\mathcal{D})))$. By the directedness condition, $\Gamma \cup \mathit{CONSEQ}(\mathbf{M}(\mathcal{D}))$ is consistent, and by the minimality of $\mathbf{M}$, there does not exist a set D'' such that $\mathbf{M}(\mathcal{D}) \subset D' \subseteq D_1$ and $\Gamma \cup \mathit{CONSEQ}(D'')$ is consistent. It follows that E is an 1-extension of $\mathcal{D}$. ∎

Since the notion of satisfying a default rule is more straightforward in the case of straight observation models than in the case of one level observation models, it is interesting to determine the kind of extensions corresponding with minimal straight observation models.

Definition 19 *Let $\mathcal{D} = \langle D, \Gamma, < \rangle$ be an observation theory. A set $D' \subset D$ is called* blocked *iff there exists a $d \in D$ such that $\Gamma \cup \mathit{CONSEQ}(D' \cup \{d\})$ is inconsistent and for every $d' \in D'$, $d < d'$. A set $D' \subseteq D$ is called* unblocked *iff D' contains no blocked subset.*

Definition 20 *Let $\mathcal{D} = \langle D, \Gamma, < \rangle$ be an observation theory. The E is called an* unblocked extension *of $\mathcal{D}$ iff $E = Th(\Gamma \cup \mathit{CONSEQ}(D'))$, for some maximal unblocked $D' \subseteq D$ such that $\Gamma \cup \mathit{CONSEQ}(D')$ is consistent.*

Proposition 11 *The unblocked extensions of an observation theory correspond with its minimal straight observation models. That is, E is an unblocked extension of $\mathcal{D} = \langle D, \Gamma, < \rangle$ iff $E = Th([\mathbf{M}]_{\mathcal{D}})$ for some minimal straight observation model $\mathbf{M}$ of $\mathcal{D}$.*

Proof. The proof is completely analogous to the proof of proposition 10. The only difference is that one now should use the fact that D' is a set of default rules satisfied by some straight observation model iff D' is unblocked and $\mathit{CONSEQ}(D')$ is consistent, and let the unblocked sets play the role that subsets of the $<$-minimal default rules play in the case of proposition 10. ∎

The difference between 1-extensions and unblocked extensions is that in the first case, an observation is assumed to defeat all observations based on less refined evidence, whereas in the latter case, observations that are not based on maximally refined evidence can still play a role (as long as they are not blocked by some incompatible observation based on more refined evidence).

Example 6 *Consider the situation of example 4, but without the second observation $q_{e'}$. That is, we have the observations p_e and $r_{e''}$, where e'' is more refined evidence than e. This situation is modelled by the observation theory $\mathcal{D}$ consisting of the default theory $\langle\{\top : p/p, \top : r/r\}, \emptyset\rangle$ with priority $\top : r/r < \top : p/p$. This theory has one one level extension $Th(\{r\})$ and one unblocked extension $Th(\{p, r\})$.*

In the concrete interpretation of a robot in the process of finding out whether a particular door is open, p_e denotes the robot's observation, based on some videocamera image taken from considerable distance from the door, that the door is 84 degrees open, and $r_{e''}$ denotes the observation, additionally based on a videocamera image taken at close proximity to the door, that the door is open (without being able to specify the exact angle of the door). In that case, the one level extension only allows the conclusion that the door is open (without being able to specify the exact angle), whereas the unblocked extension allows the more specific conclusion that the door is 84 degrees open.

The above example provides another illustration of our previously mentioned opinion that one level extensions are only appropriate under the assumption that observations mention the most specific proposition. If one uses unblocked extensions, one does not need this assumption, but it is not exactly clear how to justify the fact that an observation that itself is defeated can still defeat some observations. (In the original situation of example 4, the observation $q_{e'}$ blocks p_e, although $q_{e'}$ is defeated by the observation $r_{e''}$. Thus the unique unblocked extension coincides with the one level extension $Th(\{r\})$.)

If one believes that observations should not be blocked by defeated observations, then one should not use unblocked extensions. In that case, $<$-compatible extensions seem to be an obvious choice. In order to obtain a semantics for $<$-compatible extensions in terms of observation models, we should no longer require $\mathbf{M} \models d$ to imply that every model $\mathcal{M} \geq obs(d)$ satisfies the consequent of d. (Otherwise, a default rule d is automatically blocked by an incompatible default rule d' such that $obs(d') > obs(d)$.)

A first idea is to allow a finite number of exceptions (i.e., models $> obs(d)$ not satisfying the consequent of d) in the satisfaction clause for default rules. This does not work, since a finite extended observation model of $\langle D, \Gamma, <\rangle$ would then satisfy *all* default rules of D. Thus the minimal extended observation models would satisfy all the default rules, whether the rules are compatible or not. It is immediate that, in general, the class of the thus obtained extended observation models of an observation theory would not be regular.

Even if we restrict ourselves to top-less observation models, we do not obtain a semantics for $<$-compatible extensions, since the effect of the ordering is completely lost as soon as a finite number of exceptions is allowed in the satisfaction clause for default rules. In fact, in that case, we get a semantics for the underlying default theory, without the priorities.

Definition 21 *Let $\mathcal{D} = \langle D, \Gamma, < \rangle$ be an observation theory. A* cofinal observation model $\mathbf{M}$ *of $\mathcal{D}$ is an extended top-less observation model $\langle M, \geq, obs \rangle$ of $\mathcal{D}$, where we say that $\mathbf{M} \models d$ iff all but a finite number of the models $\mathcal{M} \geq obs(d)$ satisfy the consequent of d.*

It can be shown that the class of cofinal observation models of an observation theory is regular. (The proof of lemma 9 has to be slightly amended, since after adding a top to a cofinal model it is no longer top-less, and therefore no longer cofinal. However, instead of adding a single first-order model as a top, one can simply add an infinite sequence of copies of the first-order model.)

Proposition 12 *Let $\mathcal{D} = \langle D, \Gamma, < \rangle$ be an observation theory. The minimal cofinal observation models of $\mathcal{D}$ correspond with the (ordinary Reiter) extensions of $\langle D, \Gamma \rangle$. That is, E is an extension of $\langle D, \Gamma \rangle$ iff $E = Th([\mathbf{M}]_{\mathcal{D}})$ for some minimal cofinal observation model $\mathbf{M}$ of $\mathcal{D}$.*

Proof. The proof is completely analogous to the proof of proposition 10. The only difference is that one now should use the fact that D' is a set of default rules satisfied by some cofinal observation model of $\mathcal{D}$ iff $\Gamma \cup CONSEQ(D')$ is consistent, and let these compatible sets of default rules play the role that subsets of the $<$-minimal default rules play in the case of proposition 10. ∎

If one tries to model the situation where observations are not blocked by defeated observations, then one should not allow arbitrary exceptions in the satisfaction clause for default rules, but only allow exceptions among the first-order models representing defeated observations. This suggests a definition of the satisfaction of a default rule along the lines of $\mathbf{M} \models d$ iff for all $\mathcal{M} > obs(d)$ either $\mathcal{M}$ satisfies the consequent of d or there exist a $d' \in D$ such that $\mathcal{M} = obs(d')$ and $\mathbf{M} \models \neg d'$, where $\mathbf{M} \models \neg d'$ is a notation for 'd' is defeated'.

The problem with this approach is that the satisfaction and defeat of default rules interact in intricate ways, since the more default rules are satisfied, the more default rules are defeated, and vice versa. Therefore, it is not immediately clear what the correct definition of defeat would have to be. A simple solution is to introduce in the extended observation models an additional parameter for a particular well-ordering $\prec$ compatible with the partial order $<$ on default rules.

Definition 22 *Let $\mathcal{D} = \langle D, \Gamma, < \rangle$ be an observation theory. A* <-compatible *observation model $\mathbf{M}_\succ$ of $\mathcal{D}$ is an extended observation model $\langle M, \geq, obs \rangle$ of $\mathcal{D}$, where $\succ$ is an information image of a $<$-compatible well-ordering $\prec$, and the satisfaction and defeat of default rules are defined as follows.*

- *$\mathbf{M}_\succ \models d$ iff for all $\mathcal{M} \succ obs(d)$ either $\mathcal{M}$ satisfies the consequent of d or there exists a $d' \in D$ such that $\mathcal{M} = obs(d')$ and $\mathbf{M}_\succ \models \neg d'$.*
- *$\mathbf{M}_\succ \models \neg d'$ iff $\Gamma \cup CONSEQ(\{d'\} \cup \{d'' \in \mathbf{M}_\succ(\mathcal{D}) : d'' \prec d'\})$ is inconsistent.*

Here (as before) $\mathbf{M}_\succ(\mathcal{D}) = \{d \in D : \mathbf{M}_\succ \models d\}$, and we define $[\mathbf{M}_\succ]_\mathcal{D}$ to be the the equivalence class of $<$-compatible observation models $\mathbf{M'}_{\succ'}$ of $\mathcal{D}$ such that $\mathbf{M'}_{\succ'}(\mathcal{D}) = \mathbf{M}_\succ(\mathcal{D})$.

Notice that in the above definition the notions of satisfaction and defeat of default rules are well defined since their respective clauses only refer to other default rules that are more preferred according to $\prec$. Analogous to the proof of lemma 9 it can be shown that the class of $<$-compatible observation models of an observation theory is regular.

Proposition 13 *The minimal $<$-compatible observation models of an observation theory $\mathcal{D} = \langle D, \Gamma, < \rangle$ correspond with the $<$-compatible extensions of $\mathcal{D}$. That is, E is a $<$-compatible extension of $\mathcal{D}$ iff $E = Th([\mathbf{M}_\succ]_\mathcal{D})$ for some minimal $<$-compatible observation model $\mathbf{M}_\succ$ of $\mathcal{D}$.*

Proof. Suppose E is a $<$-compatible extension of $\mathcal{D} = \langle D, \Gamma, < \rangle$. Then E is an extension of $\langle D, \Gamma \rangle$ generated by some $<$-compatible well-ordering $\prec$. It follows that $E = Th(\Gamma \cup CONSEQ(D'))$, for some $D' \subseteq D$ such that $d \notin D'$ implies that $\Gamma \cup CONSEQ(\{d\} \cup \{d' \in D' : d' \prec d\})$ is inconsistent. Consider the class of $<$-compatible observation models $\mathbf{M}_\succ$ of $\mathcal{D}$ where $\succ$ is an information image of $\prec$. For every $\mathbf{M}_\succ$ in this class, the defeated default rules are exactly the elements of $D - D'$. Let $\mathbf{M}_\succ$ be a model of this class with a maximal set of satisfied defaults. Then $\mathbf{M}_\succ$ is minimal, $\mathbf{M}_\succ(\mathcal{D}) = D'$, and by regularity, $Th([\mathbf{M}_\succ]_\mathcal{D}) = E$.

To prove the other direction, let $E = Th([\mathbf{M}_\succ]_\mathcal{D})$, for some minimal $<$-compatible observation model $\mathbf{M}_\succ$ of $\mathcal{D}$. Suppose $\succ$ is the information image of the $<$-compatible well-ordering $\prec$ of D. By the minimality of $\mathbf{M}_\succ$, $\mathbf{M}_\succ(\mathcal{D})$ is the complement of the set of defeated default rules. Therefore, $d \notin \mathbf{M}_\succ(\mathcal{D})$ implies that the set $\Gamma \cup CONSEQ(\{d\} \cup \{d' \in \mathbf{M}_\succ(\mathcal{D}) : d' \prec d\})$ is inconsistent. Thus $Th(\Gamma \cup CONSEQ(\mathbf{M}_\succ(\mathcal{D})))$ is a $<$-compatible extension of $\mathcal{D}$. By the regularity of $<$-compatible observation models, it then follows that E is a $<$-compatible extension of $\mathcal{D}$. ∎

We conclude that the proposed observation models can provide a flexible semantics for different observation logics. In the following section we discuss in more detail the matter under what circumstances the different observation logics are most appropriate, and we relate this matter to the much discussed problem of sensor fusion in robotics.

6 The Sensor Fusion Problem

Evidence from one particular, isolated sensor reading is relatively well understood. One usually has at least sufficient partial or approximate knowledge of the behaviour, and in particular of the reliability, of a sensor to confidently relate possible sensor readings to conclusions about the state of the world. Unfortunately, since sensors are never completely reliable and often quite unreliable, these conclusions are typically fairly weak. Therefore, multiple sensor readings need to be combined in order to justify strong conclusions.

The problem with combining, or fusing, the evidence obtained from multiple sensor readings is that it is difficult to assess the interaction between multiple pieces of evidence acquired from different sensors, from different readings of the same sensor obtained at different locations, or even from readings of the same sensor obtained at the same location.

Typically, one uses some numeric uncertainty formalism, such as probability theory, which allows the reinforcement of conclusions whenever sensor readings agree and (partial) cancelling out of the individual effects of disagreeing sensor readings. This probabilistic, or in general numeric, reasoning is very powerful, but it requires a lot of data or strong assumptions. Since often the required data is insufficiently available and the chosen assumptions cannot adequately be justified, the conclusions obtained by numeric reasoning cannot always be trusted.

Since we use non-numeric default rules to represent observations, we cannot expect the same subtle reasoning. But default reasoning still may give reasonable and useful results, and, in particular in situations where the numeric data required for numeric reasoning cannot easily be obtained, the results of default reasoning are not necessarily inferior to those of the numeric approaches.

Let us briefly comment on the relation between expressing an observation in terms of (normal free) default rules and the representation of an observation in numeric, say probabilistic, terms. At first sight, it might seem reasonable to assume that a defeasible observation that ϕ, represented by $\top : \phi/\phi$, has a roughly equivalent probabilistic interpretation of the form that, given the evidence, ϕ is the most likely, or most probable, possibility. The problem with this assumption is t hat either the most likely possibility is trivially $\top$, or one has to compare only possibilities corresponding to elementary events or, more generally, possibilities corresponding to some partition of the sample space, and it is difficult to fix in advance an appropriate partition.

Example 7 *In the situation of example 4, the observation* $r_{e''}$*, based on a videocamera* *image taken at close proximity to the door, supports the possibility that the door is open (without supporting any exact angle of the door). This could be interpreted to mean that 'open' is the most likely element of* { *'open', 'closed'* }*. However, the observation* p_e*, based on some videocamera image taken from considerable distance from the door, supports the possibility that the door is 84 degrees open. If this possibility is supposed to be the most likely one given the evidence e, then possibilities are compared that are not elements of the set* { *'open', 'closed'* } *used in the case of* $r_{e''}$*.*

An alternative probabilistic interpretation of the default rule $\top : \phi/\phi$ representing a defeasible observation is to say that it is roughly equivalent to saying that, given the evidence, the probability of ϕ is sufficiently high. It is well known that, unless sufficiently high is interpreted as arbitrarily close to 1, both interpretations support different sets of inference rules. For example, whenever ϕ and ψ are compatible, $\top : \phi/\phi$ and $\top : \psi/\psi$ together support $\phi \wedge \psi$, whereas in the probabilistic interpretation this inference is not valid, since the probability of the conjunction $\phi \wedge \psi$ may be significantly lower than the probabilities of the individual conjuncts.

In spite of the fact that no exact agreement exists between the conclusion supported by default rules and those supported by the above mentioned probabilistic interpretation of default rules, we still think the interpretation is at least useful as a heuristic. For example, the fact that evidence that does not support a rather specific proposition (such as 'the door is 84 degrees open') can still support a less specific proposition (such as 'the door is open') is quite clear under the probabilistic interpretation.

For an adequate evaluation of the usefulness of the interpretation of observations by means of default rules, it is important to know how multiple observations are combined. To get an overview of how the choice for a particular notion of extension effect the combination behaviour, we summarize in table 1 the conclusions supported by the different defeasible observation logics in several examples.

It is easy to see that the conclusions supported by 1-extensions are included in the set of conclusions supported by layered extensions, since for every layered extension E of an observation theory $\mathcal{D}$ there exists an 1-extension E' of $\mathcal{D}$ such that $E' \subseteq E$. It follows from the examples summarized in table 1 that, in general, the set of conclusions supported by 1-extensions is not necessarily weaker (or stronger) than the set of conclusions supported by $<$-compatible, unblocked, or unprioritized extensions.

In example 6, all the discussed types of extensions, except the 1-extensions, support the conclusion that the door is 84 degrees open. As mentioned before, this strong conclusion is no longer justified if the observation that the door is open is assumed to be based on evidence which is really more refined than (and therefore includes) the evidence underlying the observation that the door is 84 degrees open, and if one additionally assumes that 'the door is open' is the most specific conclusion supported

extensions	example 4	example 6	new example
1-extensions	open	open	79 degr. open
$<$-compatible	84 degr. open	84 degr. open	79 or 84 degr. open
unblocked	open	84 degr. open	79 or 84 degr. open
layered	84 degr. open	84 degr. open	79 degr. open
without priorities	closed or 84 degr. open	84 degr. open	79 or 84 degr. open

Table 1: A comparison of the conclusions supported by different notions of extensions. In example 4 we have a refining sequence of observations that the door is 84 degrees open (videocamera, great distance), that the door is closed (sonar), and that the door is open (videocamera, small distance). Example 6 is the same, but without the sonar-based observation that the door is closed. In the new example, we add to example 6 a new observation that the door is 79 degrees open, and we assume that this observation is incomparable to both other observations.

by the more refined evidence.

Assuming that the evidence e underlying the observation that the door is open is more refined than the evidence e' underlying the observation that the door is 84 degrees open implies that e does not consist only of the videocamera image taken near the door, but also incorporates the videocamera image taken from considerable distance (e'). In other words, more refined evidence *combines* the less refined evidence (possibly) with new sensory information.

Although usually the evidence from one particular observation or sensor reading is better understood than the combined effect of several observations, it is also sometimes the case that in order to express a particular observation in an abstract, symbolic language, one has to take into account several other observations. For example, when interpreting a videocamera image and concluding from it that the door is open, one typically uses information obtained by other sensor readings about matters like the approximate distance between the camera and the door. Therefore, even if at one level of description sensor readings are initially interpreted in isolation, it may be the case that at some more abstract level, observations often refine other (previous) observations.

The pieces of evidence underlying incomparable (maximally preferred) observations can be viewed as pieces of evidence that (as yet) have not been combined, and the defeasible observation logic using 1-extensions can be viewed as a logic providing (defeasible) statements of what the results of the combination of these pieces of evidence are expected to be. In fact the process of combining two observations ϕ_e and $\psi_{e'}$ can be represented by adding an observation based on evidence refining both e and e'. Since it is not necessary for this latter observation to incorporate some new sensory information, one might consider the sensor fusion process itself as a kind of 'observation'.

It should be noted that, although we have used the sceptical, or cautious, interpretation of default consequence, it is *not* the case that the expectations concerning the results of combining evidence as expressed by the observation logic using 1-extensions are the weakest possible, or most cautious.

Example 8 *Consider two observations ϕ_e and $\psi_{e'}$, where ϕ denotes that the door is 84 degrees open and ψ denotes that the door is 79 degrees open, and where e and e' are incomparable. The observation logic using 1-extensions supports the conclusion that the door is 79 or 84 degrees open. However, after proper combination of e and e' it might be the case that both pieces of evidence are considered unreliable, and it might even become unclear whether the pieces of evidence provide any information concerning the state of the door at all.*

We conclude that reasoning with the observation logic using 1-extensions should *not* replace a careful sensor fusion process. However, as long as the defeasible conclusions of the observation logic are recognized as defeasible expectations, they form a potentially useful supplement to the conclusions based on the (at any particular time, probably incomplete) process of careful combination of obtained evidence.

Using $<$-compatible, unblocked, and layered extensions does not seem compatible with a reading of $<$ as a refining relation. In that case, $<$ is more naturally interpeted as a relation expressing relative strength or reliability of pieces of evidence. The observation logics using these different notions of extensions implement different intuitions concerning the issue of how to combine observations, but they all, in a sense, *replace* a careful sensor fusion process, rather than supplement it.

However, since we believe that, in general, a careful combination of the observations cannot adequately be represented, or even approximated, by a systematic interaction between the default rules representing the observations, we prefer 1-extensions above the other mentioned notions of extension.

7 The Multiple Extension Problem

The standard explanation of the meaning of a normal, free default rule $\top : \omega / \omega$ allows the derivation of ω, provided ω is consistent with what is derivable. Assuming that an extension is intended to represent a reasonable set of conclusions of the default theory, the existence of multiple extensions constitutes a problem, since any choice for a particular extension is arbitrary in the sense that it does not follow from the (default) theory.

The problem does not disappear in case the intersection of all extensions is proposed to be the set of conclusions of a default theory. The diffculty with this cautious or sceptical approach is that in general the intersection is itself not an extension and not closed under the default rules.

Example 9 *Consider the default theory consisting of the two default rules* $\top : p/p$ *and* $\top : q/q$ *and background knowledge* $\neg(p \wedge q)$*. This default theory has two extensions:* $Th(\{p, \neg q\})$ *and* $Th(\{\neg p, q\})$*. The intersection of the extensions is* $Th(\{p \leftrightarrow q\})$*. If this intersection is supposed to be the set of all derivable conclusions, then, under the standard interpretation of default rules,* $\top : p/p$ *is applicable, and* p *is derivable. In the same way,* $\top : q/q$ *allows the derivation of* q*. But* $Th(\{p \leftrightarrow q\})$ *contains neither* p *nor* q*.*

A possible solution of the problem is to come up with an alternative interpretation of default rules, where the justification does not refer to the set of derivable conclusions. When default rules are used to represent defeasible observations, such an alternative interpretation is actually rather natural.

An observation ϕ_e is a reason to conclude that ϕ, and this reason is sufficient, unless other (incomparable or more refined) observations provide reasons against ϕ. If these other observations are based on more refined evidence than e, then the observation ϕ_e is defeated. If the other observations are incomparable to ϕ_e, then the supported conclusion is weaker than ϕ, but ϕ should still be considered possible. Thus, the default rule $\top : \phi/\phi$ representing ϕ_e supports the conclusion ϕ only in case all evidence against ϕ is less refined than e.

Given the interpretation of a (normal, free) default rule as a (not necessarily sufficient) reason to conclude the consequent of the rule, the sceptical approach to default logic seems quite sensible. In the sceptical approach, extensions should not be interpreted as reasonable sets of conclusions of a default theory, but as maximally consistent sets of conclusions, mirroring the interaction between the default rules.

For example, the intersection $Th(\{p \leftrightarrow q\})$ of the extensions $Th(\{p, \neg q\})$ and $Th(\{\neg p, q\})$ seems to be a reasonable set of conclusion in case of example 9. The fact that there is a reason for believing p and a reason for believing q is compatible with this set of conclusions, since the reasons are only considered sufficient for actually drawing the conclusions in case all conflicting evidence is less refined.

If default rules are used to represent defeasible observation, then there is a second solution to the multiple extension problem. Since multiple extensions only occur in the presence of conflicting default rules, one can simply avoid the formalization of conflicting observations. For example, one can apply a process of sensor fusion to harmonize conflicting observations, before translating the observation into default rules. The resulting observation logic will have a unique extension.

In the context of 1-extensions, it is always possible to 'weaken' several conflicting observations to their disjunction. For example, p_e and $q_{e'}$, with p and q incompatible and e and e' incomparable, can be replaced by $(p \vee q)_{e''}$, with e'' more refined than e and e'. This can be viewed as a particular method of combining (possibly conflicting) observations, but the conclusions resulting from harmonizing conflicting observations with this particular method of sensor fusion are in general weaker than the conclusions supported by taking the intersection of all 1-extensions of the original observation theory.

Example 10 *Consider the default theory consisting of the three default rules $\top : p/p$, $\top : q/q$, $\top : r/r$, and background knowledge $\neg(p \wedge q \wedge r)$. This default theory has three extensions: $Th(\{p, q, \neg r\})$, $Th(\{p, \neg q, r\})$, and $Th(\{\neg p, q, r\})$. Each of these extensions contains the formula $\neg p \rightarrow (q \wedge r)$, but this formula is not contained in $Th(\{\neg(p \wedge q \wedge r), p \vee q \vee r\})$, which is the unique 1-extension of the observation theory obtained from the original default theory by adding a default rule $\top : p \vee q \vee r/p \vee q \vee r$ refining the original three default rules.*

One can of course try to represent other, perhaps more complex, sensor fusion methods in observation logic, but we will not pursue this matter here.

8 Conclusions

We have shown how defeasible observations can be formalized using prioritized normal free default logic, i.e., observation logic. Several notions of extension for this special case of default logic have been discussed, and it has been argued that the notion of 1-extension is the most appropriate, provided the priority relation is assumed to be a refining relation, and the formula mentioned in an observation statement is maximally specific, or informative.

The proposed semantics for observation logic is loosely based on the refining systems for the logic of vision presented in [5], and is sufficiently general to cover several variants of observation logic induced by various notions of extension. The difference between 1-extensions and the other notions is caused by a difference in the way the satisfaction relation of an observation model is extended to include default rules.

We argued that the observation logic induced by 1-extensions cannot replace a careful sensor fusion process, but that it can supplement sensor fusion by giving some statements of the expected combined effect of observations before the underlying pieces of evidence are actually combined in a careful (but difficult) sensor fusion process.

Finally, we claim that the existence of multiple extensions of an observation theory does not consitute a problem for observation logic, since the default rules representing defeasible observations have a natural interpretation where the justification of a default rule does not refer to the set of conclusions of the observation theory.

References

[1] P. Besnard, *An Introduction to Default Logic* (Springer, Berlin, 1989).

[2] G. Brewka, Preferred subtheories: An extended logical framework for default reasoning, *IJCAI-89* (Morgan Kaufmann, San Mateo CA, 1989) 1043–1048.

[3] D.W. Etherington, *Reasoning with Incomplete Information* (Pitman, London, 1988).

[4] S. Kraus, D. Lehmann and M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, *Artificial Intelligence* 44 (1990) 167–207.

[5] M. van Lambalgen and J. van der Does, A logic of vision, *research report* LP-97-07 (ILLC, University of Amsterdam, 1997). (Revised version of LP-96-14.)

[6] D. Makinson, General theory of cumulative inference. *Non-Monotonic Reasoning*, M. Reinfrank, J. de Kleer, M.L. Ginsberg, E. Sandewall (eds.) (Springer, Berlin, 1989) 1–18.

[7] V.W. Marek and M. Truszczyński, *Nonmonotonic Logic. Context-Dependent Reasoning* (Springer, Berlin, 1991).

[8] R. Reiter, A logic for default reasoning, *Artificial Intelligence* 13 (1980) 81–132.

[9] F. Voorbraak, Preference-based semantics for nonmonotonic logics, *IJCAI-93* (Morgan Kaufmann, San Mateo CA, 1993) 584–589.

[10] F. Voorbraak, *As Far as I Know. Epistemic Logic and Uncertainty*, (Dissertation, Utrecht University, Utrecht, 1993).