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 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 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.

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 section 2, 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 3, 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 4. Finally, we briefly discuss how the proposed observation logic relates to the multiple extension problem and the problem of sensor fusion. A more extensive discussion of these issues can be found in the extended version of this paper [11], which also contains the proofs of the propositions.

2 Normal Free Default Logic

We assume familiarity with default logic as introduced in [8], but we repeat some definitions to fix notation. 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:* $\alpha : \beta_1, \ldots, \beta_n / \omega$. *The formula* α *is called the* prerequisite, $\beta_1, \ldots, \beta_n$ *are called the* justifications, *and* ω *the* consequent *of the default rule. 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 6.)

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*.*

We assume, without loss of generality, that all formulas occuring in the default rules of a default theory are closed. In [8], an extension of a default theory $\langle D, \Gamma \rangle$ is proposed as s representation of a reasonable state of belief based on the default rules in D and facts in Γ, and an exact fixed point definition of the notion of an extension is given. For our purpose the following, rather informal characterization of default extensions, which is essentially a reformulation of the semantics for default logic given in [3], is sufficient. 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. However, there is a natural class of default theories with at least one extension, namely the class of *normal* default theories, which have only (normal) default rules with a single justification equivalent (in first-order logic) to the consequent of the rule. In addition to having at least one extension, a normal default theory has several other nice properties. See [8] or the full paper [11].

We are especially interested in a fragment of (normal) default logic in which the prerequisite of a default rule is assumed to be a tautology. Such a rule is called (prerequisite-)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, contrary to arbitrary (normal) default logic, allows reasoning by cases. See [1, 9, 10, 11]. 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 1 *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.*

Another interesting property of normal free default logic is that it induces a *cumulative* consequence operation. That is, if we define for any set D of default rules and for any set Γ of sentences of L, define $Cn_D(\Gamma) = \bigcap\{E : E \text{ is an extension of } \langle D, \Gamma \rangle\}$, then Cn_D is a cumulative consequence operation. ($Cn_D(\Gamma)$ is called the set of sceptical, or cautious, conclusions of $\langle D, \Gamma \rangle$.) 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.

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

3 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 $\langle \phi, e \rangle$ denote the observation that ϕ based on evidence e. We can interpret this, analogous to the rendition of perception reports in [5], as the conjunction of the following two statements.

1. ϕ is true *given the approximation of reality induced by e*. Notation: ϕ_e.

2. ϕ is expected to be true in more refined approximations, including reality itself. That is, it is expected that $\phi_{e'}$, for every e' that is more refined than e.

The expectation mentioned in the second statement is defeasible in the sense that it can be overridden by other observations. For example, suppose we have the observations $\langle \phi, e \rangle$ and $\langle \psi, e' \rangle$, where ϕ and ψ are incompatible (given the a priori knowledge) and e' is more refined than e. Then the expectation induced by $\langle \phi, e \rangle$ is defeated by $\langle \psi, e' \rangle$, since the latter observation implies $\psi_{e'}$.

The (defeasible) expectation accompanying the observation $\langle\phi, e\rangle$ intuitively (and informally) seems to imply the (normal) default rules $\langle\phi, e\rangle : \phi_{e'}/\phi_{e'}$, for each e' that is more refined than e. If one only adds these default rules to the theory when the observation $\langle\phi, e\rangle$ is made and the prerequisite is true, then one can use the normal free default rules $\top : \phi_{e'}/\phi_{e'}$ instead.

If we abstract in our formal language from the evidence on which the observation is based, then the defeasible expectation can simply be modelled as a normal free default rule $\top : \phi/\phi$. To make up for the loss of information caused by abstracting away the mention of the evidence, and 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 3 *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 4 *An* observation theory *is a prioritized normal free default theory.*

In this paper, the intended meaning of $d < d'$ is that the default rule d represents an observation based on evidence which is more refined, i.e., contains more information, than the evidence on which the observation represented by d' is based. This implies that d is stronger than d'. (This reading of $<$ is in accordance with the usually chosen direction of preference 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 restrictive 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 5 *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 $<$. Here $\prec$ is called* compatible *with $<$ iff for every $d, d' \in D$, $d < d'$ implies $d \prec d'$.*

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 more preferred defaults (with higher priorities) are applied first, thereby making less preferred incompatible defaults non-applicable.

Example 1 *Consider the observations* $\langle p, e\rangle$, $\langle q, e'\rangle$ *and* $\langle r, e''\rangle$, *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* $\top : q/q < \top : p/p$. *Both* $Th(\{\neg p, q, r\})$ *and* $Th(\{p, q, \neg r\})$ *are* $<$*-compatible extensions.*

In the above example, the same conclusion seems to be 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. One 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, since it does not seem appropriate for our interpretation of the preference relation between default rules as resulting from a refinement relation between the bodies of evidence underlying observations.

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 6 *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* $D - \cup_{0 \leq j \leq i-1} D_j$. *That is,* D_i *is the set of default rules* d *such that* i *is the maximal length of a sequence* $d > d_1 > \ldots > d_k$. E *is called a* layered extension *of* $\mathcal{D}$ *iff* $E = \cup_{i \geq 0} E_i$, *where* $E_0 = Th(\Gamma)$, *and for each* $i \geq 1$, E_i *is an (ordinary default) extension of* $\langle D_i, E_{i-1}\rangle$. *If one only takes the union of the first* n *layers, then one obtains what we call an* n*-extension of* $\mathcal{D}$. *In particular, a 1-extension of* $\mathcal{D}$ *is an (ordinary default) extension of* $\langle D_1, Th(\Gamma)\rangle$, *where* D_1 *is the set of* $<$*-minimal elements of* D.

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

Proposition 2 *Let* $\mathcal{D} = \langle D, \Gamma, <\rangle$ *be an observation theory. Every layered extension of* $\mathcal{D}$ *is a* $<$*-compatible extension of* $\mathcal{D}$.

Example 1 shows that the reverse of this proposition is not true, since the $<$-compatible extension $Th(\{p, q, \neg r\})$ is *not* a layered extension. 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 2 *Consider the observations $\langle p, e\rangle$, $\langle q, e'\rangle$ and $\langle r, e''\rangle$, 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 $\langle p, e\rangle$ 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 $\langle q, e'\rangle$ denotes the observation based on some sonar readings in addition to the videocamera image, that the door is closed, and that $\langle r, e''\rangle$ 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 $\langle p, e\rangle$ is defeated by $\langle q, e'\rangle$, and it is not clear whether the fact that $\langle q, e'\rangle$ is in turn defeated by $\langle r, e''\rangle$ justifies re-establishing the conclusions based on $\langle p, e\rangle$. 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 $\langle r, e''\rangle$, 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 5.

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 6) 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 6.

As is well known, normal default logic is semi-monotone in the sense that adding normal default rules to a normal default theory does not invalidate any conclusions of the old theory. (See [8].) At first sight, one might think that this semi-monotonicity results in a monotonic behaviour of the observation logic of 1-extensions. 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 3 *Consider the situation of example 2. 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 $\langle p, e\rangle$ (where p cannot be deduced from the background knowledge) by adding the observation $\langle \top, e'\rangle$ with e' more refined than e. (More generally, any $\langle \phi, e'\rangle$, where ϕ together with the background knowledge does not imply p and e' is 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.

4 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. (A partially ordered set X is called *directed* iff every finite subset of X has an upper bound in X.)

Definition 7 *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.

Example 4 *Consider the observation theory $\mathcal{D}$ of example 2, 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 $\top : r/r < \top : q/q < \top : p/p$. An example of an observation model $\mathbf{M}$ of $\mathcal{D}$ is the tuple $\langle M, \geq, obs \rangle$, where $M = \{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\}$, $\mathcal{M}_1 < \mathcal{M}_2 < \mathcal{M}_3$, $obs(\top : p/p) = \mathcal{M}_1$, $obs(\top : q/q) = \mathcal{M}_3$, $obs(\top : r/r) = \mathcal{M}_3$, p is true in $\mathcal{M}_1$ and false in the other first-order models, q is true in $\mathcal{M}_2$ and false in the other first-order models, and r is true in $\mathcal{M}_3$ and false in the other first-order models.*

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 7 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, in the sense that $obs(d) >' obs(d')$ whenever $d < d'$ will be called an *information image* of $<$.

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 *topless*. These topless 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.

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 8 *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. 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 9 *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* $obs(d') > obs(d)$.

Proposition 3 *The 1-extensions of an observation theory correspond with its minimal one level observation models. That is,* E *is a 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* $\mathcal{D}$.

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 10 *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 $\Gamma \cup CONSEQ(D' \cup \{d\})$ 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 11 *Let $\mathcal{D} = \langle D, \Gamma, < \rangle$ be an observation theory. The E is called an* unblocked extension *of $\mathcal{D}$ iff $E = Th(\Gamma \cup CONSEQ(D'))$, for some maximal unblocked $D' \subseteq D$ such that $\Gamma \cup CONSEQ(D')$ is consistent.*

Proposition 4 *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 $\mathcal{D}$.*

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 5 *Consider the situation of example 2, but without the second observation $\langle q, e' \rangle$. That is, we have the observations $\langle p, e \rangle$ and $\langle r, e'' \rangle$, 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 1-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, $\langle p, e \rangle$ 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 $\langle r, e'' \rangle$ 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 2, the observation $\langle q, e' \rangle$ blocks $\langle p, e \rangle$, although $\langle q, e' \rangle$ is defeated by the observation $\langle r, e'' \rangle$. 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 topless 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 obtain a semantics for the underlying default theory, without the priorities.

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 12 *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 $\prec$.

Proposition 5 *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* $\mathcal{D}$.

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.

5 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, which is frequently insufficiently available. Therefore, 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. In the full paper [11], we 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.

The interpretation of observations in a default logic implies a way of combining observations, where of course the choice for a particular notion of extension effects the combination behaviour. It is easy to see that the conclusions supported by 1-extensions are included in the set of conclusions supported by layered extensions, and 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 5, 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 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.

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 $\langle \phi, e \rangle$ and $\langle \psi, e' \rangle$ 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 6 *Consider two observations* $\langle \phi, e \rangle$ *and* $\langle \psi, e' \rangle$*, 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.

6 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.

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 $\langle \phi, e \rangle$ 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 $\langle \phi, e \rangle$ is defeated. If the other observations are incomparable to $\langle \phi, e \rangle$, then the supported conclusion is weaker than ϕ, but ϕ should still be considered possible. Thus, the default rule $\top : \phi / \phi$ representing $\langle \phi, e \rangle$ 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.

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. See the full paper [11] for an elaboration of this solution.

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).

[11] F. Voorbraak, A nonmonotonic observation logic, *research report* LP-1997-09 (ILLC, University of Amsterdam, 1997).