Handout Lecture 8: Non-monotonic logics

Xavier Parent and Leon van der Torre
University of Luxembourg

April 27, 2016

Abstract. This handout is devoted to non-monotonic logics–a family of logics devised to model defeasible reasoning. The focus is on the semantics for non-monotonic reasoning based on preferential models.

1 Introduction

The term "non-monotonic logics" (in short, NMLs) covers a family of logics developed in AI in order to capture and represent common-sense reasoning. One of the distinctive features of common-sense reasoning is that it is defeasible, meaning that the acquisition of new knowledge can cause earlier conclusions to be withdrawn.

Such logics are called non-monotonic, because they reject the so-called principle of monotony. This is the principle, where Γ is a set of formulas:

$$\Gamma \vdash \phi \Rightarrow \Gamma \cup \{\psi\} \vdash \phi \qquad \text{(Mon)}$$

When Γ is finite, (Mon) is equivalent to the law of strengthening of the antecedent, which (as we have seen in lecture 7) is characteristic of material implication:

$$\vdash \xi \to \phi \quad \Rightarrow \quad \vdash \xi \wedge \psi \to \phi \qquad \text{(SA)}$$

Thus, the distinctive feature of non-monotonic logics is that the body of inferred knowledge does not grow monotonically with the body of received knowledge.

Why reject (Mon), or equivalently, (SA)? This is because in everyday life situations, available information is usually incomplete. One still needs to be able to reason about such situations, and make decisions.

Four examples of defeasible reasoning are given below. The symbol "$|\!\sim$" (*snake*) denotes the construct that is distinctive of NMLs. For "$\phi \mid\!\sim \psi$", read "If ϕ, normally ψ".

Reasoning based on generalizations allowing for exceptions

- bird $|\!\sim$ flies
- bird, penguin$|\!\not\sim$ flies

Reasoning based on normality Consequences are drawn based on what is most likely to be the case:

- cake $|\!\sim$ good
- cake, soap $|\!\not\sim$ good

Reasoning by default Consequences are derived only because of lack of evidence to the contrary:

- cf. presumption of innocence (innocent until proven guilty)

Diagnosis (also called "abductive reasoning") Consequences are deduced as most likely explanations of a given fact:

- gras wet, sunny $|\!\sim$ sprinkler on
- gras wet, sunny, street wet $|\!\not\sim$ sprinkler on

Classical logic is inadequate since it is monotonic.

2 Negation as failure

If you want a concrete example from computer science, think of a logic programming language like Prolog. It uses non-monotonic logic.

Remember that a program written in Prolog is a set of sentences in logical form, expressing facts and rules about some domain. This, we call it a "knowledge base" (in short, **KB**). We use a Prolog program, by running queries. That is, by asking questions about the information stored in the knowledge base. Prolog answers queries using the following assumption, known as the Closed World Assumption (CWA):

What cannot be shown to be true is false (CWA)

For example, if an airline database does not contain a seat assignment for a traveler, so that it cannot be proved that the traveler has checked in, it is assumed that the traveler has not checked in. (CWA) translates into a rule called "Negation as failure" (NAS), where the symbol $\vdash$ means "is provable":

$$\text{If } \nvdash \phi \text{ then } \vdash \neg\phi \qquad \text{(NAS)}$$

This inference is not warranted in classical logic. A logic that supports (NAS) is non-monotonic: a new fact may yield to withdraw a previously obtained conclusion.

(NAS) has two uses, a trivial one and a non-trivial one. First, a "yes or no" query always gets answered−this is its trivial use. Second, (NAS) can be used to express and reason about exceptions to the rules in the knowledge base−this is its non-trivial use. Let me explain how this is done.

Consider the following small program:

```
penguin(opus).    /* Opus is penguin */
eagle(hedwig).  /* Hedwig is an eagle */
bird(X) :- eagle(X). /* an eagle is a bird */
bird(X) :- penguin(X). /* a penguin is a bird */
```

Here are a few queries one can run along with their answers:

```
?-penguin(opus).   /* is Opus a penguin? */
true.
?-bird(opus).   /* is Opus a bird?*/
true.
?-penguin(hedwig).  /* is Hedwig a penguin?*/
false.
```

The last query illustrates the trivial use of (NAS). It cannot be proved that "penguin(hedwig)" is true, and hence it is considered as false.

Let us add the following rule–it is defeasible, because it leaves room for exceptions:

```
fly(X) :- bird(X),\+ penguin(X).
     /* a bird flies unless it is a penguin */
```

"/+" is the "not" operator, but its behavior is more subtle than negation as defined in logic. Indeed its behavior is regimented by (NAS). Compare:

```
?-fly(opus).      /*Does Opus fly? */
false.
?-fly(hedwig).    /* Does Hedwig fly? */
true.
```

Prolog's query-answering procedures is based on modus-ponens. In the logic programming notation, the rule of interest is:

$$f \leftarrow b, \neg p \tag{1}$$

The query "**?-fly(opus).**" fails because Opus is known to be an exception; modus-ponens cannot be applied:

$$\vdash p \quad \stackrel{\text{MP?}}{\longrightarrow} \quad \nvdash f \quad \stackrel{\text{NAS}}{\longrightarrow} \quad \vdash \neg f$$

The query "**?-fly(hedwig).**" succeeds because Hedwig is not known to be an exception; modus-ponens can be applied modulo (NAS);

$$\nvdash p \quad \stackrel{\text{NAS}}{\longrightarrow} \quad \vdash \neg p \quad \stackrel{\text{MP}}{\longrightarrow} \quad \vdash \neg f$$

3 Preferential semantics

This section describes one of the most well-known semantics for non-monotonic reasoning, the preferential semantics, due to Shoham [7], Makinson [5, 6], Kraus, Lehman & Magidor [2], and Lehmann & Magidor [4].

3.1 Preferential models

The language used in such an approach is based on conditional expressions $\phi \mid\sim \psi$, read as "Normally, if ϕ then ψ", where ϕ and ψ are formulas of propositional logic. From now onwards the construct $\phi \mid\sim \psi$ will be referred to as a default rule, or just a default.

Reasoning in such a framework means to be able, given a set **KB** of such defaults, to derive new defaults.

Definition 1. *A preferential model is a pair $(W, <)$ where W is a set of valuations v_1, v_2 … on the propositional letters, and $<$ is a binary relation over W. Intuitively, one reads $v_1 < v_2$ as saying that v_1 is (strictly) preferred over v_2, because it describes a more normal (more typical, more plausible, etc.) situation.*

Notes:

- W is not necessarily the entire set of all the valuations
- We use $\mathbf{1}$ and $\mathbf{0}$ (bold fonts) to denote the truth-values "true" and "false"
- $<$ is often called a preference relation
- One might allow for multiple copies of valuations, by indexing valuations v by elements s of an arbitrary index set S. We do not do it, and take the notion of valuation in its usual sense, as a function from the set of all elementary letters $\{p, q, ...\}$ into $\{\mathbf{1}, \mathbf{0}\}$

Definition 2. *Given a preferential model $(W, <)$ we say $\phi \mid\sim_{(W,<)} \psi$ iff $v(\psi) = \mathbf{1}$ for every valuation $v \in W$ that is minimal (w.r.t. $<$) among those in W that satisfy ϕ. The requirement of minimality means: there is no $v' \in W$ with $v'(\phi) = \mathbf{1}$ and $v' < v$.*

When $\phi \mid\sim_{(W,<)} \psi$ is the case, we say that the default $\phi \mid\sim \psi$ holds in the preferential model $(W, <)$, or just that $\phi \mid\sim \psi$ if it is clear what model is intended.

Intuitively, Definition 2 says: $\phi \mid\sim \psi$ holds if ψ holds in the preferred (most normal, most typical, etc) situations where ϕ holds.

It can be convenient to express the definition in a more compact manner. When W is a set of valuations and ϕ is a formulas, write $|\phi|_W$ for the set of all valuations in W that satisfy ϕ, i.e. $|\phi|_W = \{v \in W : v(\phi) = \mathbf{1}\}$. Write $\min_<(|\phi|_W)$ for the set of all minimal elements of $|\phi|_W$. In this notation, $\phi \mid\sim \psi$ whenever $\min_<(|\phi|_W) \subseteq |\psi|_W$.

Example 1 explains how to verify whether a set of defaults holds in a given preferential model.

Example 1. *Let $(W, <)$ be such that*

- *$W = \{v_1, v_2, v_3\}$, where*
  - *$v_1(p) = \mathbf{1}$ and $v_1(q) = \mathbf{0}$*
  - *$v_2(p) = \mathbf{0}$ and $v_2(q) = \mathbf{1}$*
  - *$v_3(p) = \mathbf{1}$ and $v_3(q) = \mathbf{1}$*
- *$v_1 < v_2$ and $v_3 < v_2$*

We have $\top \mid\sim p$, because

- *$\min_<(|\top|_W) = \min_<(\{v_1, v_2, v_3\}) = \{v_1, v_3\}$*
- *$|p|_W = \{v_1, v_3\}$.*

Hence:

$$\min_{<}(|\top|_W) \subseteq |p|_W = \{v_1\}$$

We also have $\neg p \mid\sim q$.

Note that in all these definitions, no constraints are placed on the relation $<$ over W. But to guide intuitions it is useful to keep at the back of one's mind the typical case that it is both irreflexive, transitive, smooth and modular (alias virtually connected). Smoothness says: if v_1 satisfies ϕ, then either $v_1 \in \min_<(|\phi|_W)$ or there is $v_2 \in \min_<(|\phi|_W)$ with $v_2 < v_1$. Modularity says that whenever $v_1 < v_2$ then either $v_1 < v_3$

or $v_3 < v_2$. A preferential model in which $<$ meets these four conditions is usually called a "ranked" model, because one may equivalently use a ranking function associating to each valuation a unique natural number, called its rank. Typically, the most preferred (normal, typical, etc) valuations get rank 0, the second-most preferred ones get rank 1, the third-most preferred ones rank 2, and so-forth.

Note that any irreflexive and transitive relation is also assymetric (never both $v < v'$ and $v' < v$) and more generally acyclic (never$v_1 < v_2 < ..., v_n < v_1$, for $n > 1$).

Note also that, if W is finite, then $<$ is necessarily smooth (given transitivity of $<$).

Informally, we can describe a ranked preferential model by a diagram with levels as we did for counterfactuals (cf. handout 7). For instance, here is a diagrammatic representation of the ranked model given in Example 1:

		ranking
v_2		1
v_1	v_3	0

The convention is that valuations lower down in the ordering are more preferred (more normal, more typical, etc) than those higher up, while those on the same level are incomparable under $<$.[1] Intuitively, all the valuations within the same level have the same rank.

Instead of listing the valuations, it is also sometimes more convenient to list the propositional letters that each valuation satisfies, and use the bar notation for those it makes false:

$\overline{p}q$	
$p\overline{q}$	pq

Exercise 1. *Show that antisymmetry of $<$ follows from irreflexivity and transitivity of $<$.*

Exercise 2. *Show that, if W is finite, then $<$ is necessarily smooth (given transitivity of $<$).*

Exercise 3. *Consider the following ranked preferential model:*

$pq\overline{r}\overline{s}$	$pqr\overline{s}$	
$p\overline{q}r\overline{s}$	$\overline{p}q\overline{r}s$	$p\overline{q}rs$

Show that $p \mid\sim r$, $p \wedge q \not\mid\sim r$, $p \wedge \neg q \mid\sim r$.

Exercise 4. *Show that $\phi \mid\sim \psi$ does not imply $\phi \wedge \phi' \mid\sim \psi$.*

The proof-theory is similar to that for conditional logic. It is known that the class of ranked preferential models corresponds to the following system of rules, called system **R** (for "Rational"):

$$\text{REF}\ \frac{}{\phi \mid\sim \phi} \qquad \text{LLE}\ \frac{\phi \mid\sim \psi \qquad \vdash \phi \leftrightarrow \phi'}{\phi' \mid\sim \psi}$$

$$\text{RW}\ \frac{\phi \mid\sim \psi \qquad \vdash \psi \rightarrow \psi'}{\phi \mid\sim \psi'} \qquad \text{AND}\ \frac{\phi \mid\sim \psi \qquad \phi \mid\sim \psi'}{\phi \mid\sim \psi \wedge \psi'}$$

$$\text{OR}\ \frac{\phi \mid\sim \psi \qquad \phi' \mid\sim \psi}{\phi \vee \phi' \mid\sim \psi} \qquad \text{CM}\ \frac{\phi \mid\sim \psi \qquad \phi \mid\sim \xi}{\phi \wedge \xi \mid\sim \psi}$$

[1] Two valuations v and v' are incomparable under $<$, if $v \not< v'$ and $v' \not< v$.

$$\text{RM}\ \frac{\phi \mid\sim \psi \qquad \phi \not\mid\sim \neg\xi}{\phi \wedge \xi \mid\sim \psi}$$

The abbreviations are read as follows: REF–reflexivity; LLE–left logical equivalence; RW–right weakening; CM–cautious monotony; RM–rational monotony. AND and OR are self-explanatory.

Note that CM calls for smoothness, and RM calls for modularity.

Exercise 5. *Show that the above rules are valid as long as $<$ meets the required properties as stated.*

3.2 A more fine-grained approach: lexicographic ordering

Preferential models as described in the previous section faces what has been called the "drowning problem". This one illustrates the difficulty to find the right balance between unrestricted strengthening of the antecedent and no strengthening at all. As Horty puts it, "what is needed [...] is a certain amount of strengthening, but not too much" [1, p.56]. One needs to find a middle way between these the aforementioned two extremes, and this task is not as as straightforward as one might think.

We first describe the drowning problem, and then present one solution to it, in terms of lexicographic ranking, due to Lehmann [3].

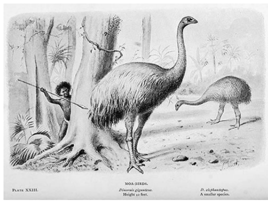

Figure 1. Moa: the only wingless bird that ever existed

Example 2 (Drowning problem)**.** *Suppose* **KB** *contains the following*

i) $b \mid\sim f$ *(most birds fly)*

ii) $b \wedge p \mid\sim \neg f$ *(penguins do not fly)*

iii) $b \mid\sim w$ *(birds have wings)*

Strengthening of the antecedent should not hold in its plain form— on pain of generating a contradiction. From i) one does not want to be able to derive $b \wedge p \mid\sim f$. But a certain of amount of strengthening of the antecedent is needed, because from iii) one would like to be able to infer

iv) $b \wedge p \mid\sim w$ *(penguins have wings)*

The intuition is that ii) overrides i), but not iii), so that one should be able to conditionalize iii) to $b \wedge p$.

Here is a ranked preferential model in which all the defaults in **KB** hold, and $b \wedge p \mid\sim w$ does not.

$bfpw$	
$b\overline{f}p\overline{w}$	$b\overline{f}pw$
$bf\overline{p}w$	

3.2.1 Procedure

Valuations are ranked based on two criteria:

(C_1) number of defaults that they violate: the less, the better
(C_2) seriousness of the violations: it is less serious to violate a less specific default than a more specific default

To make it work, we need two definitions.

Definition 3. *Given some* **KB***, the violation set V of a valuation $v \in W$ is the set of defaults in* **KB** *that are falsified by v:*

$$V(v) = \{\phi \mid\sim \psi \in \mathbf{KB} : v(\phi) = \mathbf{1}, v(\psi) = \mathbf{0}\}$$

Definition 4. *$\phi \mid\sim \psi$ is more specific than $\phi' \mid\sim \psi'$ if $\vdash \phi \rightarrow \phi'$ but not conversely.*

Semantically, given the ranked preferential models of **KB**, the construction leads us to consider only those models in which every valuation is considered as typical as possible, that is, it is "pushed downward" in the model as much as possible, modulo the satisfaction of **KB**.

Syntactically speaking, the conjunction of (C_1) and (C_2) gives us a defeasible version of the law of strengthening of the antecedent. If $\mathbf{KB} = \{\phi \mid\sim \psi\}$, then $\phi \wedge \xi \mid\sim \psi$ follows. If $\mathbf{KB} = \{\phi \mid\sim \psi, \phi \wedge \xi \mid\sim \neg\psi\}$, then $\phi \wedge \xi \mid\sim \psi$ does not follow. In other words, one may apply the strengthening of the antecedent principle, but only in the absence of evidence to the contrary. (Cf. Exercise 6.)

Another way to put it is to say that the construction specifies how (RM) should be satisfied: we have $\phi \mid\sim \psi$ in **KB**; in order to satisfy (RM) we have to add either $\phi \mid\sim \neg\xi$ or $\phi \wedge \xi \mid\sim \psi$. The procedure imposes that, whenever it is possible, we prefer the latter (that corresponds to a constrained application of monotony) over the former.

Now, the construction. There are three main steps.

Step 1 Divide **KB** into levels of specificity. That is, define a partition $\{\Delta_i\}_{0\leq i\leq m}$ of **KB**, where each Δ_i gathers the defaults with the same degree of specificity. Δ_0 enumerates the most specific defaults, Δ_1 the second most specific ones, and so on, up to the most general ones.[2]

Step 2 Assign to each valuation $v \in W$ a n-tuple $\langle n_0, n_1, ...n_m\rangle$, where $n_i = \mid V(v) \cap \Delta_i \mid$, for each $i \in \{0, ..., m\}$

Step 3 Rank the valuations using the lexicographic ordering on their associated n-tuples:

$$\langle n_0, n_1, ...n_m\rangle < \langle n'_0, n'_1, ...n'_m\rangle \text{ iff}$$
$$n_i < n'_i \text{ for the first } i \text{ s.t. } n_i \neq n'_i$$

3.2.2 Drowning problem revisited

We show how the construction resolves the drowning problem.

$$\mathbf{KB} = \{b \mid\sim f, b \wedge p \mid\sim \neg f, b \mid\sim w\}$$

[2] Lehmann uses a more complex definition than Definition 4. This one suffices for present purposes.

Step 1 We get

$$\Delta_0 = \{b \wedge p \mid\sim \neg f\} \text{ and } \Delta_1 = \{b \mid\sim f, b \mid\sim w\}$$

Step 2 We get

	b	f	p	w	V	n-tuple
v_1	**1**	**1**	**1**	**1**	$\{b \wedge p \mid\sim \neg f\}$	$\langle 1, 0\rangle$
v_2	**1**	**1**	**1**	**0**	$\{b \mid\sim w, b \wedge p \mid\sim \neg f\}$	$\langle 1, 1\rangle$
v_3	**1**	**1**	**0**	**1**	$\emptyset$	$\langle 0, 0\rangle$
v_4	**1**	**1**	**0**	**0**	$\{b \mid\sim w\}$	$\langle 0, 1\rangle$
v_5	**1**	**0**	**1**	**1**	$\{b \mid\sim f\}$	$\langle 0, 1\rangle$
v_6	**1**	**0**	**1**	**0**	$\{b \mid\sim f, b \mid\sim w\}$	$\langle 0, 2\rangle$
v_7	**1**	**0**	**0**	**1**	$\{b \mid\sim f\}$	$\langle 0, 1\rangle$
v_8	**1**	**0**	**0**	**0**	$\{b \mid\sim f, b \mid\sim w\}$	$\langle 0, 2\rangle$
v_9	**0**	**1**	**1**	**1**	$\emptyset$	$\langle 0, 0\rangle$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Step 3 We get the following ranking:

v_2	4	$\langle 1, 1\rangle$
v_1	3	$\langle 1, 0\rangle$
v_6 v_8	2	$\langle 0, 2\rangle$
v_4 v_5 v_7	1	$\langle 0, 1\rangle$
v_3 $v_9 - v_{16}$	0	$\langle 0, 0\rangle$

Note

$$\min_<(|b|_W) = \{v_3\}$$
$$\min_<(|b \wedge p|_W) = \{v_5\}$$

Hence the following conditionals hold:

$$b \mid\sim f \quad b \wedge p \mid\sim \neg f \quad b \mid\sim w$$

Also

$$b \wedge p \mid\sim w$$

We have resolved the drowing problem!

Exercise 6. *Suppose the language contains three propositional letters p, q and r. Compare what happens in the following two situations under the lexicographic account.*

$$\mathbf{KB} = \{p \mid\sim q\}$$

$$\mathbf{KB} = \{p \mid\sim q, p \wedge r \mid\sim \neg q\}$$

REFERENCES

[1] J. Horty, 'Moral dilemmas and nonmonotonic logic', *Journal of Philosophical Logic*, **23**(1), 35–65.
[2] S. Kraus, D. Lehmann, and M. Magidor, 'Nonmonotonic reasoning, preferential models and cumulative logics', *Artificial Intelligence*, **44**, 167–207, (1990).
[3] D. Lehmann, 'Another perspective on default reasoning', *Annals of Mathematics and Artificial Intelligence*, **15**, 61–82, (1995).
[4] D. Lehmann and M. Magidor, 'What does a conditional knowledge base entail?', *Artificial Intelligence*, **55**(1), 1–60, (1992).
[5] D. Makinson, 'Five faces of minimality', *Studia Logica*, **52**(3), 339–379, (1993).

[6] D. Makinson, ‘General patterns in nonmonotonic reasoning’, in *Handbook of logic in art. intell. and logic programming*, volume 3, 35–110, Oxford University Press, Inc., New York, NY, USA, (1994).

[7] Y. Shoham, *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*, MIT Press, Cambridge, 1988.