Introduction to Non Monotonic Reasoning

Master Recherche SIS, Marseille

Nicola Olivetti

Professeur à la Faculté Econonomie Appliquée, Université Paul Cezanne

Laboratoire CNRS LSIS

2010-2011[a]

---

[a] I am indebted to Laura Giordano and Alberto Martelli for having provided me their course material.

Non-monotonic reasoning

- Often available knowledge is *incomplete*.
- However, to model commonsense reasoning, it is necessary to be able to jump to *plausible conclusions* from the given knowledge.
- To draw plausible conclusions it is necessary to make assumptions.
- The choice of assumptions is not *blind*: most of the knowledge on the world is given by means of general rules which specify typical properties of objects. For instance, "birds fly" means: birds typically fly, but there can be exceptions such as penguins, ostriches, ...

Non-monotonic reasoning

- **Nonmonotonic reasoning** deals with the problem of deriving plausible conclusions, but not infallible, from a knowledge base (a set of formulas).
- Since the conclusions are not certain, it must be possible to retract some of them if new information shows that they are *wrong*
- Classical logic is inadequate since it is monotonic: if a formula B is derivable from a set of formulas S, then B is also derivable from any superset of S:

  $S \vdash B$ implies $S \cup \{A\} \vdash B$, for any formula A.

Non-monotonic reasoning

- **Example:** let the KB contain:

  Typically birds fly.
  Penguins do not fly.
  Tweety is a bird.

- It is plausible to conclude that Tweety flies.
- However if the following information is added to KB

  Tweety is a penguin

  the previous conclusion must be retracted and, instead, the new conclusion that Tweety does not fly will hold.

Non-monotonic reasoning

- The statement "typically A" can be read as: "in the absence of information to the contrary, assume A".
- The problem is to define the precise meaning of "*in the absence of information to the contrary*".
- The meaning could be: "there is nothing in KB that is inconsistent with assumption A".
- Other interpretations are possible
- Different interpretations give rise to different non-monotonic logics

Inadequacy of Classical Logic

- We cannot represent a rule such as "typically birds fly" as

  $$\forall x(bird(x) \wedge \neg exception(x) \rightarrow fly(x))$$

  and then to add

  $$\forall x(exception(x) \leftrightarrow penguin(x) \vee ostrich(x) \vee canary(x) \vee \ldots)$$

- We do not know in advance all exceptions
- In order to conclude that "Tweety"' fly we should prove that "‘tweety is not an exception"', that is:

  $$\neg penguin(tweety), \neg ostrich(tweety), \ldots$$

Inadequacy of Classical Logic

- On the contrary we would like to prove that Tweety flies because we cannot conclude that it is an exception, not because we can prove that it is not an exception.

Closed World Assumption

- A basic understanding of database logic, is that only positive information is represented explicitly. Negative information is not represented explicitly.
- If a positive fact is not present in the database (DB), it is assumed that its negation holds.
- This is called Closed World Assumption: the only true facts are the provable ones.
- If $DB \not\vdash A$ then $DB \vdash_{CWA} \neg A$
- This inference *is not valid* in classical logic.

Closed World Assumption

- **Example**: suppose a DB contains facts of the form "practice(person, sport)"', for instance:

  $$practice(anne, tennis)$$
  $$practice(joe, tennis)$$
  $$practice(anne, sky)$$

- Then we have

  $$DB \vdash_{CWA} \neg practice(joe, sky)$$

- Trivially CWA is non-monotonic, since adding a fact may lead to withdraw the negative conclusion:

  $$DB \cup \{practice(joe, sky)\} \not\vdash_{CWA} \neg practice(joe, sky)$$

Frame Problem

- Problem of representing a dynamic world
- How to represent that objects are not affected by state change?
- Example: moving an object does not change its color
- In a representation based on a classical-logic , we must explicitly assert the *persistence* of object properties. We need a great number of frame axioms, such as:
  $\forall x \forall c \forall s \forall l (color(x, c, s) \rightarrow color(x, c, result(move, x, l, s)))$
  $\forall x \forall c \forall s (color(x, c, s) \rightarrow color(x, c, result(t\_light\_on, s)))$
  $\forall x \forall c \forall s (color(x, c, s) \rightarrow color(x, c, result(open\_door, s)))$
  $\vdots$

Frame Problem

- We would need a general meta-axiom of the form:

  $$\forall p \forall a \forall s(holds(p,s) \wedge \neg exception(p,a,s) \rightarrow holds(p, result(a,s)))$$

- But then we must be able to conclude that an action is not an *exception* to the preservation of a given property, unless we can show that it actually is.
- We need a non-monotonic reasoning mechanism.

NonMonotonic Logics

Non-Monotonic logics have been proposed at the beginning of the 80’s, here are historically the most important proposals:

- Non-monotonic logic, by McDermott and Doyle, ’80
- Default Logic, by Reiter, ’80
- Circumscription, by McCarthy, ’80
- Autoepistemic logic, Moore ’84

Default Logic

- Default logic extends classical logic by non-standard inference rules. These rules allows one to express default properties.
- Example:

$$\frac{bird(x) \ : fly(x)}{fly(x)}$$

that can be interpreted as: "‘if x is a bird and we can consistently assume that x flies then we can infer that x flies"’

Default Logic

- More generally we can have rules of the form:

$$\frac{\alpha(x) \ : \beta(x)}{\gamma(x)}$$

  that can be interpreted as: "“if $\alpha(x)$ holds and $\beta(x)$ can be consistently assumed then we can conclude $\gamma(x)$".
- terminology:
  - $\alpha(x)$: the prerequisite
  - $\beta(x)$: the justification
  - $\gamma(x)$: the consequent

Default Theory

- A default theory is a pair $< D, W >$, where D is a set of default rules and W is a set of first-order formulas.
- Example: let let $< D, W >$ be

$$D = \{\frac{bird(x) \ \ : fly(x)}{fly(x)}\}$$

$$\begin{aligned} W \quad = \quad & \{bird(tweety), \forall x(penguin(x) \rightarrow bird(x)), \\ & \forall x(penguin(x) \rightarrow \neg fly(x))\} \end{aligned}$$

Default Theory

Intuitively, in a default theory $< D, W >$:

- W represents the stable (but incomplete) knowledge of the world
- D rules for *extending* the knowledge W by plausible (but defeasible) conclusions.
- Notion of ***extension*** of a default theory: the theory (= deductively closed set of logical formulas) obtained by extending W by the rules in D.

Default Theory

- Example: let $< D, W >$ be as in the previous example
- Since $bird(tweety)$ is true, and it is consistent to assume $fly(tweety)$, then $fly(tweety)$ is true in the (unique) extension of $< D, W >$.
- Consider now the the default theory $< D, W' >$, where

$$W' = W \cup \{penguin(tweety)\}$$

then the assumption $fly(tweety)$ is no longer consistent, and the application of the default rule is blocked.

Default Theory

- Example2: let $< D, W >$ be as follows:

$$D = \{d_1 = \frac{Rep(x) \ \ : \neg Pac(x)}{\neg Pac(x)}, d_2 = \frac{Quack(x) \ \ : Pac(x)}{Pac(x)}, \}$$

$$W = \{Rep(Nixon), Quack(Nixon)\}$$

  For both default rules d_i, the prerequisite is derivable from W. What can be concluded from $< D, W >$?

- If we apply d_1, we conclude $\neg Pac(Nixon)$; therefore $Pac(Nixon)$ cannot be assumed consistently, so that d_2 is blocked.
- If we apply d_2, we conclude $Pac(Nixon)$; therefore $\neg Pac(Nixon)$ cannot be assumed consistently, so that d_1 is blocked.

Default Theory

- There are two extensions: one containing $\neg Pac(Nixon)$ and the other containing $Pac(Nixon)$.
- An extension (to be defined next) represents the set of plausible conclusions.
- As we shall see, a default-theory may have zero, one, or many extensions.

Extensions (propositional case)

Given a default theory $\Delta =< D, W >$, a set of formulas E is an extension of Δ, if:

- E is deductively closed: $E = Th(E)$
- all applicable defaults with respect to E have been applied, that is for all $\frac{\alpha : \beta}{\gamma} \in D$

  if $\alpha \in E$ and $\neg\beta \notin E$ then $\gamma \in E$
- Deductive closure operator: $Th(S) = \{C \in \mathcal{L} \mid S \vdash C\}$

Extensions: semi-inductive definition

Given a default theory $\Delta =< D, W >$, a set of formulas E is an extension of Δ, if it can be obtained as follows:

- $S_0 = W$
- $S_{i+1} = Th(S_i) \cup \{\gamma \mid \frac{\alpha : \beta}{\gamma} \in D, \alpha \in S_i, \neg\beta \notin E\}$
- $E = \bigcup_i S_i$

Extensions: semi-inductive definition

- The definition is not really inductive, since the definition of S_{i+1} makes reference to the whole E.
- The order in which defaults are considered in step S_{i+1} is significant: different orders give rise to different extensions.
- In the propositional case every extension can be "generated" in at most k stages where k is the number of defaults in the default theory.

Extensions

Example 1: let $\Delta = \{b, p \rightarrow \neg f\}, \{\frac{b : f}{f}\}$, then there is a unique extension $E = Th(\{b, p \rightarrow \neg f, f\})$

- $S_0 = \{b, p \rightarrow \neg f\}$
- $S_1 = S_0 \cup \{f\}$, since $S_0 \vdash b$ and $\neg f \notin E$

Extensions

Example 1': let $\Delta = \{b, p \rightarrow \neg f, p\}, \{\frac{b : f}{f}\}$, then there is a unique extension $E = Th(\{b, p \rightarrow \neg f, p\})$

- $S_0 = \{b, p \rightarrow \neg f, p\}$
- $S_1 = S_0$, since $S_0 \vdash b$ but $\neg f \in E$

Extensions

Example 2: let $\Delta = \{r, q\}, \{d_1 = \frac{r : \neg p}{\neg p}, d_2 = \frac{q : p}{p}\}$.

- Let $E_1 = Th(\{r, q, \neg p\})$
  - $S_0 = \{r, q\}$
  - $S_1 = S_0 \cup \{\neg p\}$, by applying d_1, since $S_0 \vdash r$ and $\neg\neg p \notin E_1$
  - $S_2 = S_1$, since d_2 cannot be applied $\neg p \in E_1$
  - for $i \geq 2$, $S_i = S_2$

Extensions

Example 2 (continued)

- Let $E_2 = Th(\{r, q, p\})$
  - $S_0 = \{r, q\}$
  - $S_1 = S_0 \cup \{p\}$, by applying d_2, since $S_0 \vdash q$ and $\neg p \notin E_2$
  - $S_2 = S_1$, since d_1 cannot be applied: $\neg\neg p \in E_2$
  - for $i \geq 2$, $S_i = S_2$

Extensions

Example 3 Let $\Delta =< W, D >$, where $W = \emptyset$ and $D = \{\frac{: a}{\neg a}\}$.
Suppose there is an extension E

- if $\neg a \notin E$, then it must be $\neg a \in E$ (we must apply the default)
- but if $\neg a \in E$, the default become inapplicable: thus it must be $\neg a \notin E$
- Δ has no extensions!

Extensions

Example 4 Let $\Delta =< W, D >$, where $W = \emptyset$ and
$D = \{d_1 = \frac{: \neg p}{q}, d2 = \frac{: \neg q}{p}\}$.

- Let $E_1 = Th(\{q\})$
  - $S_0 = \emptyset$
  - $S_1 = S \cup \{q\}$, , since $\neg\neg p \notin E$.
  - $S_2 = S_1$, since d_2 becomes inapplicable.
- Similarly, we get another extension $E_2 = Th(\{p\})$

Extensions

Example 4 Let $\Delta =< W, D >$, where $W = \emptyset$ and $D = \{\frac{a : b}{b}, \frac{b : a}{a}\}$. Then there is a unique extension $E = Th(\emptyset)$

- $S_0 = \emptyset$
- $S_1 = S_0$, since $S_0 \not\vdash a$, and $S_0 \not\vdash b$

Normal defaults

- A default d is ***normal*** if has the form $\frac{\alpha : \beta}{\beta}$
- A normal default theory $\Delta =< W, D >$ is a default theory where all defaults in D are normal
- **Theorem**: A normal default theory has always an extension.

Inference relation

Since a default theory $\Delta =< W, D >$ may have multiple extensions (including none), how to define a notion of inference? There are two natural notions:

- **(credulous inference)** $\Delta \vdash_c A$ if there exists an extension E of Δ such that $A \in E$.
- **(skeptical inference)** $\Delta \vdash_s A$ if for all extensions E of Δ, we have $A \in E$.
- Since a default theory may have no extensions $\Delta \vdash_s A$ does not imply $\Delta \vdash_c A$.

A simple algorithm

An algorithm to compute ***any*** extension of a theory $\Delta =< W, D >$

- (0) Let $< S_0, D_0 >=< W, \emptyset >$.

  (i+1) Let $< X, Y, Z >=< S_i, \emptyset, D - D_i >$

  for every $d \in Z$, $d = \frac{\alpha_d : \beta_d}{\gamma_d}$

  if $S_i \cup X \vdash \alpha_d$ and $S_i \cup X \not\vdash \neg\beta_d$

  then $< X, Y >=< X \cup \{\gamma_d\}, Y \cup \{d\} >$

  let $< S_{i+1}, D_{i+1} >=< S_i \cup X, D_i \cup Y >$
- stop with the least k such that $< S_k, D_k >=< S_{k+1}, D_{k+1} >$
- check whether for each $d = \frac{\alpha_d : \beta_d}{\gamma_d} \in D_k$, $S_k \not\vdash \neg\beta_d$.

  If "yes", return S_k.

Problems with default logic

- Unwanted transitivity: let $\Delta =< W, D >$, where $W = \{student\}$ and

$$D = \{d_1 = \frac{student : adult}{adult}, d_2 = \frac{adult : works}{works}\}$$

- it is easy to see that Δ has a unique extension including $\{student, works, adult\}$.
- it is rather unintuitive (as students usually do not work).
- if we add the default $\frac{student : \neg work}{\neg work}$, the theory has then two extensions:
  $E_1 = \{student, adult, works\}$
  $E_2 = \{student, adult, \neg works\}$
- But E_2 is more plausible than E_1

Problems with default logic

- Solution: replace d_2 by:

$$\frac{adult : works \wedge \neg student}{works}$$

- then the only extension is $E_2 = \{student, adult, \neg works\}$
- this default is not normal
- it is ***semi-normal***: the justification implies the consequent
- a semi-normal default theory (= a theory where all default are semi-normal) may have no extensions

Problems with default logic

- Handling specificity: let $\Delta =< W, D >$, where $W = \{user, blacklisted\}$ and

$$D = \{d_1 = \frac{user : login}{login}, d_2 = \frac{user \wedge blacklisted : \neg login}{\neg login}\}$$

- the theory has then two extensions:
  $E_1 = \{user, blacklisted, login\}$
  $E_2 = \{user, blacklisted, \neg login\}$
- But of course only E_2 is the intended one.

Problems with default logic

- The problem of *specificity* can be handled by assigning a priority to defaults on the base of their specificity. The priority order is taken into account for calculating extensions.
- Reiter's Default logic has also other problems (e.g. cumulativity)
- Many variants have been proposed, such as Brewka's one and Lukaszewicz's one.