Non-Monotonic Logics

Marco Degano

Philosophical Logic 2024
3 December 2024

Readings

Required:

- Frank Veltman lecture notes on counterfactuals (sec. 5). `https://staff.fnwi.uva.nl/f.j.m.m.veltman/papers/Notes_Counterfactuals.pdf`
- Lecture notes: chapter 7

Plan

1. Defeasible Reasoning
2. Preferential Consequence Relation
3. Applications

Outline

Defeasible Reasoning

Birds fly

Our reasoning is often based on some notion of **normality/typicality**: we state facts which are rationally compelling but not deductively valid.

But this reasoning is **defeasible**: conclusions or beliefs can be overturned (or 'defeated') when new evidence or arguments arise.

Broad Cases of Defeasible Reasoning

- **Legal reasoning:** A witness provides an alibi for a suspect, but surveillance footage later invalidates the alibi.
- **Everyday decision-making:** Assuming a restaurant is closed because the lights are off, but revising the assumption upon noticing a 'We're Open' sign.
- **Scientific Hypotheses**: Newtonian mechanics were long accepted as accurate but were revised with the advent of Einstein's theory of relativity.
- **Medical Diagnosis:** Suspecting strep throat and prescribing antibiotics, but switching the diagnosis to mononucleosis after further testing.
- …

Non-monotonic reasoning

Defeasible Reasoning is not monotonic. From $\phi \rightarrow \psi$, we cannot infer $(\phi \wedge \chi) \rightarrow \psi$.

(1)
a. If it is a mammal, then it gives birth to live young.
b. If it is a mammal and it is a platypus, then it gives birth to live young.

(2)
a. If you press the brake pedal, then the car will stop.
b. If you press the brake pedal and the brake lines are cut, then the car will stop.

(3)
a. If the ground is icy, then it is cold outside.
b. If the ground is icy and you are inside an ice rink, then it is cold outside.

(4)
a. If the dog is barking, then someone is at the door.
b. If the dog is barking and a squirrel is in the yard, then someone is at the door.

From Aristotle to AI

Aristotle: generalizations are important for practical reasoning and distinct from the certainty of deductive logic.

In philosophy, defeasibility has been typically assumed when it comes to moral or political issues. Only in the second half of last century, precise logical system aimed at formalized defeasible reasoning emerged.

Ray Reiter's **Default Logic**: Reiter introduced a system where defaults allow provisional conclusions unless contradicted by known facts.

Circumscription: John McCarthy proposed a system where the 'least abnormal'/'minimal' models of a situation are selected.

Today's focus: Kraus et al. (1990) and **preferential consequence relation**.

Outline

Setting the stage

Notice: similarity analyses of counterfactuals invalidate monotonicity.

Certain worlds are deemed ‘more normal/typical/preferred’ models based on a preference ordering, capturing the idea that defaults apply unless overridden by abnormal conditions.

A conclusion ψ follows defeasibly from ϕ if ψ holds in all the most preferred models where ϕ holds.

Preferential Models

We can provide a framework for non-monotonic logics based on the similarity framework we introduced for counterfactuals, with the following additional assumptions:

- The limit assumption: $\prec$ is a well-founded partial order on W.
- **Absoluteness:** for every $u, w \in W : \prec_u = \prec_w$ [$\prec_w$ is independent of w]
- **Universality:** for every $w \in W$, $W_w = W$ [the ordering is on W]

Logical consequence

Recall the original clause for counterfactuals.

$M \models (\phi \rightsquigarrow \psi)$ iff for every world $w \in W$, $M, u \models \psi$ for every closest $[\![\phi]\!]$-world u to w.

With Universality and Absoluteness, we can simplify it as follows:

$M \models (\phi \rightsquigarrow \psi)$ iff $M, u \models \psi$ for every $\prec$-minimal $[\![\phi]\!]$-world u.

Logical consequence

Instead of $\phi \rightsquigarrow \psi$, we write $\phi \mid\!\sim \psi$ and more generally $(\phi_1 \wedge \ldots \wedge \phi_n) \rightsquigarrow \psi$ as $\phi_1, \ldots, \phi_n \mid\!\sim \psi$

$\phi_1, \ldots, \phi_n \mid\!\sim_M \psi$ iff $M, u \models \psi$ for every $\prec$-minimal world u in $[\![\phi_1]\!] \cap \ldots \cap [\![\phi_n]\!]$.

Importantly, while $\rightsquigarrow$ can be nested, the consequence relation $\mid\!\sim$ can not, as it belongs to the metalanguage.

Preferential Models

We can restate preferential models in the following way.

A triple $M = \langle W, \prec, I \rangle$, with W as set of worlds, $\prec$ a irreflexive and transitive preference ordering on W, $I : W \times P \to \{0, 1\}$ an interpretation function for the formulas in our language.

Smoothness conditions: For all formulas ϕ, $[\![\phi]\!] := \{w \in W : I(w, \phi) = 1\}$ is smooth.

Given $A \subseteq W$, $w \in W$ is $\prec$-minimal in A if $w \in A$ and $\forall a \in A : \neg(a \prec w)$.

A subset $A \subseteq W$ is smooth if for all $a \in A$, either a is $\prec$-minimal in A or there is $a' \in A$ with $a' \prec a$ and a' is $\prec$-minimal in A.

$\phi \mid\!\sim_M \psi$ iff for all $\prec$-minimal elements w of $[\![\phi]\!]$, we have $w \in [\![\psi]\!]$.

Entailment

A set K of plausible consequences preferentially entails a plausible consequence $\phi \mid\!\sim \psi$ if, for all preferential models M, if every plausible consequence of K is validated by M, $K \subseteq \mid\!\sim_M$, then $\phi \mid\!\sim_M \psi$.

Think of K as your Knowledge base.

Proof Theory

The following takes the name of system $\mathbf{P}$ and characterizes the consequence relation $\mid\!\sim$:

1. **Reflexivity:** $\phi \mid\!\sim \phi$
2. **Left logical equivalence:** if $\phi \models \psi$ and $\psi \models \phi$, and $\phi \mid\!\sim \chi$, then $\psi \mid\!\sim \chi$
3. **Right weakening:** if $\phi \models \psi$ and $\chi \mid\!\sim \phi$, then $\chi \mid\!\sim \psi$
4. **Cut:** if $\phi \wedge \psi \mid\!\sim \chi$, and $\phi \mid\!\sim \psi$, then $\phi \mid\!\sim \chi$
5. **Cautious monotonicity:** if $\phi \mid\!\sim \psi$ and $\phi \mid\!\sim \chi$, then $\phi \wedge \psi \mid\!\sim \chi$
6. **Or:** if $\phi \mid\!\sim \chi$ and $\psi \mid\!\sim \chi$, then $\phi \vee \psi \mid\!\sim \chi$

Reflexivity

Reflexivity: $\phi \mid\sim \phi$

This seems to satisfied by any reasoning based on some notion of consequence.

Left logical equivalence

Left logical equivalence: if $\phi \models \psi$ and $\psi \models \phi$, and $\phi \mid\sim \chi$, then $\psi \mid\sim \chi$

Logical equivalent formulas have the same consequences. In other words, the consequence of a formula depend on its meaning, not on its form.

Right weakening

Right weakening: if $\phi \models \psi$ and $\chi \mid\!\sim \phi$, then $\chi \mid\!\sim \psi$

Plausible consequence is closed under logical consequence.

Cut

Cut: if $\phi \wedge \psi \mathrel{|\!\sim} \chi$, and $\phi \mathrel{|\!\sim} \psi$, then $\phi \mathrel{|\!\sim} \chi$

You can rely on intermediate conclusions ψ as stepping stones to reach further conclusions χ. Once an intermediate step ψ is established as plausible from ϕ, it can be effectively 'cut out' of the reasoning chain.

Is **Cut** a rule we should have?

(5)
a. We expect it will be raining tonight.
b. If it rains tonight, normally birds should sing tomorrow.
c. Normally, birds should sing tomorrow.

Cautious monotonicity

Cautious monotonicity: if $\phi \mid\!\sim \psi$ and $\phi \mid\!\sim \chi$, then $\phi \wedge \psi \mid\!\sim \chi$

Learning a new fact, the truth of which could have been plausibly concluded, should not invalidate previous conclusions.

(6) a. We expect it will be raining tonight.
b. Normally, birds should sing tomorrow.
c. Even if it rains tonight, normally birds should sing tomorrow.

Or

Or: if $\phi \mid\sim \chi$ and $\psi \mid\sim \chi$, then $\phi \vee \psi \mid\sim \chi$

If we do not assume this rule, we get a weaker system, called $\mathbf{C}$

If a formula is a plausible consequence of two different formulas, then it should also be a plausible consequence of their disjunction.

Some Derived Rules (+ Exercise)

Equivalence: if $\phi \mathrel{|\!\sim} \psi$ and $\psi \mathrel{|\!\sim} \phi$ and $\phi \mathrel{|\!\sim} \chi$, then $\psi \mathrel{|\!\sim} \chi$

Union: if $\phi \mathrel{|\!\sim} \psi$ and $\chi \mathrel{|\!\sim} \gamma$, then $\phi \vee \chi \mathrel{|\!\sim} \psi \vee \gamma$

Exercise: Show also that

1. if $\phi \vee \psi \mathrel{|\!\sim} \phi$ and $\psi \vee \chi \mathrel{|\!\sim} \psi$, then $\phi \vee \chi \mathrel{|\!\sim} \phi$
2. if $\phi \vee \psi \mathrel{|\!\sim} \phi$ and $\psi \vee \chi \mathrel{|\!\sim} \psi$, then $\phi \mathrel{|\!\sim} \chi \rightarrow \psi$

Model and Proof Theory

For a set K of plausible consequences and a plausible consequence $\phi \mid\!\sim \psi$, the following are equivalent

1. K preferentially entails $\phi \mid\!\sim \psi$ (i.e., for all preferential models M, if $K \subseteq \mid\!\sim_M$, then $\phi \mid\!\sim_M \psi$)
2. Taking the elements of K as additional axioms, one can derive in $\mathbf{P}$, the plausible consequence $\phi \mid\!\sim \psi$.

Some Underivable Rules

Monotonicity: if $\phi \models \psi$ and $\psi \mid\sim \chi$, then $\phi \mid\sim \chi$

Rational Monotonicity: if $\phi \mid\sim \chi$ and $\phi \not\mid\sim \neg\psi$, then $\phi \wedge \psi \mid\sim \chi$

To show the failure of rational monotonicity, take a model $M = \langle W, \prec, I \rangle$ s.t. $W = \{w_1, w_2, w_3\}$ with $\prec = \{\langle w_2, w_3 \rangle\}$ and I specified s.t. we have p, q, r in w_1, we have $p, q, \neg r$ in w_2, and we have $p, r, \neg q$ in w_3. Now set $\phi = p$, $\chi = q$ and $\psi = r$.

However, suppose that we hold (7-a), but not (7-b), shouldn't we hold (7-c)?

(7) a. Normally, the party will be great.
b. Normally, Peter will not come to the party.
c. Even if Peter comes, normally the party will be great.

Some Underivable Rules (Exercise)

Check that the following rules cannot be derived in P:

1. if $\phi \mid\sim \psi \to \chi$, then $\phi \wedge \psi \mid\sim \chi$
2. if $\phi \vee \psi \mid\sim \chi$, then $\phi \mid\sim \chi$ or $\psi \mid\sim \chi$

The Penguin Triangle (+ Exercise)

Suppose K contains the assertions below:

1. $p \mid\!\sim b$
2. $p \mid\!\sim \neg f$
3. $b \mid\!\sim f$

Show that $p \not\mid\!\sim f$ in K

Show that all preferential models that satisfy K satisfy also $p \wedge b \mid\!\sim \neg f$

Exercise: Show also that

1. $f \mid\!\sim \neg p$
2. $b \mid\!\sim \neg p$
3. $b \vee p \mid\!\sim f$
4. $b \vee p \mid\!\sim \neg p$

Outline

1. Defeasible Reasoning
2. Preferential Consequence Relation
3. Applications

The Frame Problem

Imagine designing how a robot should reason: what are the adequate axioms that specify the way the robot should act?

(8) a. If the daylight sensor is low, turn on the light.
b. If the temperature is low, turn on the heating.

Daylight sensor low $\rightarrow$ turn on the light.

Temperature low $\rightarrow$ turn on the heating.

But should the robot keep the light on? *Normally*, yes.

Grice and Implicatures

Does John speak English?

Well, he knows the colors.

From this conversation, we *normally* infer that John does not speak English. But this inference is defeasible.

In linguistics, conversational implicatures follow from the assumption that speakers adhere to certain maxims of conversation.

Pronoun resolution

(9) John met Bill at the station. *He* greeted *him*.

Preferred interpretation follows the order of the sentence.

But this can be overruled if additional information is presented, as in *Then John greeted him as well.*

Temporal anaphora

The event of a sentence is in the simple past *normally* takes place before the event of a sentence that follows.

But this can be overruled as well:

(10) John fell. Mary pushed him.

Asher and Lascarides (2003): systematic account of temporal anaphora, lexical disambiguations, . . . , within a non-monotonic framework.