

First-Order Logic Does Not Always Work

We dealt before with *All X are Y*.

Now let's think about *Normally, X are Y*.

Main example:

Normally, birds fly.

Normally, penguins are birds.

Normally, penguins do not fly.

Normally, red penguins do not fly.

It would be a mistake to translate *Normally*, by *For all*.

All X are Y.

Normally, X are Y.

All Y are Z.

Normally, Y are Z.

All X are Z.

Normally, X are Z.

Good

Bad

Which Look Ok, Which Do Not?

Normally, X are Y .

Normally, Y are Z .

Normally, X are Z .

Normally, X are Y .

Normally, $X \vee Z$ are Y .

Normally, X are Y .

Normally, Z are Y .

Normally, $X \vee Z$ are Y .

Which Look Ok, Which Do Not?

Normally, X are Y .

Normally, Z are Y .

Normally, $X \wedge Z$ are Y .

Normally, X are Y .

Normally, X are Z .

Normally, X which are Z are Y .

Questions, and a Plan

The main question is whether we can make any sense of this phenomenon at all.

People have solid intuitions about *typicality* judgments, and they are far more common than complex deductive logical reasoning of the kind we have been studying (and doing!).

So what we want to do is to build two *mathematical models* of sentences like *Normally penguins do not fly*.

This gives two formal semantics.

Then we study the models to see whether the inferences that we have already considered come out or not.

Of course, we also hope that the mathematical models have *something* going for them!

A Semantics for “Normally, X are Y ”

A *preferential model* is a tree W whose nodes are called *worlds*.

Each world w comes with a list of which atomic sentences are true w (and hence which are false at w .)

The order on the tree is indicated by an arrow \rightarrow .

(There are no agents here, and the arrow has a different intuitive meaning than in modal logic.)

w is *maximal* if there is no v such that $w \rightarrow v$.

If X is a logic formula, we define $\|X\|$ to be the set of worlds where X is true.

$\|X\|_{max}$ is the set of worlds $w \in \|X\|$ is true and there is no x such that $w \rightarrow x$ and $x \in \|X\|$.

The Intuitions

Official terminology	Cog Sci Equivalent
preferential model	prototype model
world	prototype (??)
$x \rightarrow y$	y is more basic than x
maximal	most typical, most basic

The primary sources:

Yoav Shoham, "A semantical approach to nonmonotonic logics", 1987.

Judea Pearl, "System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning", 1990.

Ernest Adams, The Logic of Conditionals, 1975.

All of this work shows the connections between core problems in AI and traditional problems in philosophical logic.

Two fine points

We assume that \rightarrow is *transitive*: if $x \rightarrow y$ and $y \rightarrow z$, then $x \rightarrow z$.

As usual with trees, the transitive arrows are not indicated in pictures.

We also assume that $x \rightarrow x$ is *never* true.

We will only deal with *finite* models.

At any rate, the important point is that if $w \in \|\mathcal{X}\|$ and w is not maximal in $\|\mathcal{X}\|$, then there is some y such that $w \rightarrow y$ and y is maximal in $\|\mathcal{X}\|$.

A Semantics for “Normally, X are Y ”

Given a model \mathcal{W} , we say that “Normally, X are Y ” is true in our model if

$$\|X\|_{max} \subseteq \|Y\|.$$

Otherwise, it is false.

We also say that “All X are Y ” is true in the model if

$$\|X\| \subseteq \|Y\|.$$

Otherwise, it is false.

A Preferential Model

w_1

w_4

w_2

w_3

w_1 : Bird, Fly, \neg Penguin


w_2 : Bird, \neg Fly, Penguin

w_3 : Bird, Fly, Penguin

w_4 : Bird, \neg Fly, \neg Penguin

The idea is that going upward means a more “normal” or “prototypical” world.

A Preferential Model



$$\|\text{Bird}\| = \{w_1, w_2, w_3, w_4\}.$$

$$\|\text{Fly}\| = \{w_1, w_3\}. \quad \|\neg\text{Fly}\| = \{w_2, w_4\}.$$

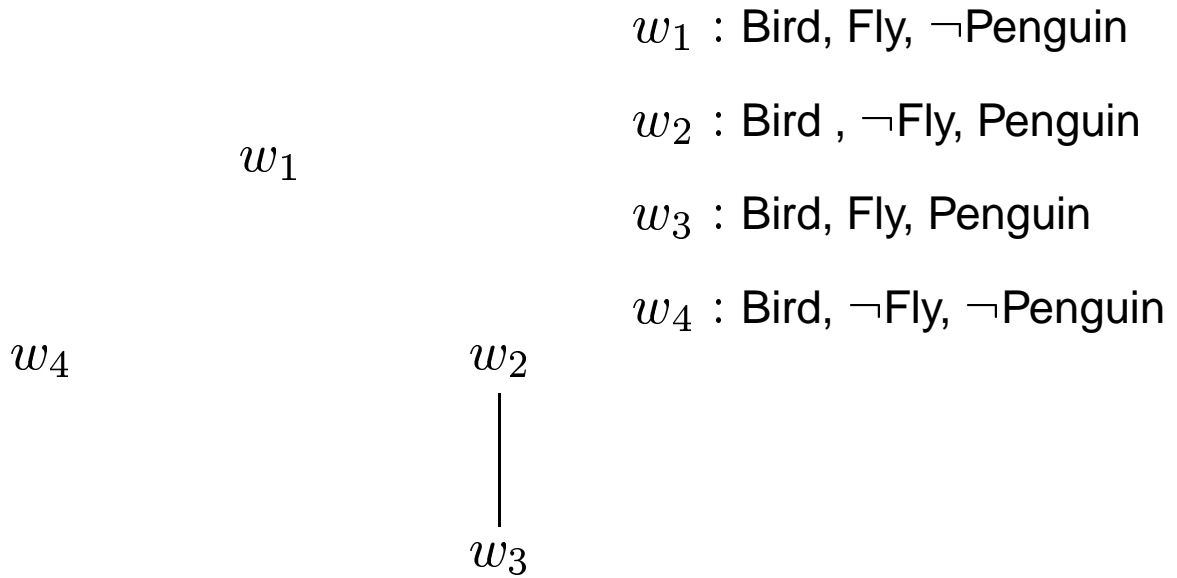
$$\|\text{Penguin}\| = \{w_2, w_3\}.$$

$$\|\text{Bird}\|_{max} = \{w_1\}.$$

$$\|\text{Fly}\|_{max} = \{w_1\}. \quad \|\neg\text{Fly}\|_{max} = \{w_2, w_4\}.$$

$$\|\text{Penguin}\|_{max} = \{w_2\}.$$

A Preferential Model



All penguins are birds is true in this model.

All birds fly is false in this model (because of w_2).

Normally, birds fly is true here because $\|\text{Bird}\|_{max} = \{w_1\}$
and $\|\text{Fly}\| = \{w_1, w_3\}$.

Normally, penguins (do) not fly is true here since
 $\|\text{Penguin}\|_{max} = \{w_2\}$ and $\|\neg\text{Fly}\| = \{w_2, w_4\}$.

Now That We Have a Semantics

We have a fragment with sentences

Normally, X are Y

where here X might be atomic, or a “boolean combination” of atomic sentences. For example

Normally, $X \cap Y$ are $\neg Z$.

We also have the notion of a *preferential model* W .

We say that $\Gamma \models S$ if every model of all sentences in Γ is also a model of S .

The idea here is that $\Gamma \models S$ should mean that “ S follows from Γ in virtue of general features of preferential models.”

What one wants to do now is to investigate whether the principles from earlier in this lecture hold for this semantics.

Then one would like to turn the semantic notion $\Gamma \models S$ into a proof-theoretic notion $\Gamma \vdash S$.

The Core Default Logic (System P)

$$\frac{\text{Normally, } X \text{ are } Y. \quad \text{Normally, } Z \text{ are } Y.}{\text{Normally, } X \cup Z \text{ are } Y.}$$
$$\frac{\text{Normally, } X \text{ are } Y. \quad \text{Normally, } X \text{ are } Z.}{\text{Normally, } X \text{ are } Y \cap Z.}$$
$$\frac{\text{Normally, } X \text{ are } Y. \quad \text{Normally, } X \text{ are } Z.}{\text{Normally, } X \cap Y \text{ are } Z.}$$
$$\frac{\text{Normally, } X \text{ are } Y. \quad \text{All } Y \text{ are } Z.}{\text{Normally, } X \text{ are } Z.}$$
$$\frac{\text{Normally, } X \text{ are } Y. \quad \text{All } X \text{ are } U. \quad \text{All } U \text{ are } X.}{\text{Normally, } U \text{ are } Y.}$$

Soundness

In every preferential model, if the assumptions above the line are true, then the conclusion below the line is also true.

Let's check this for one of the rules:

$$\frac{\text{Normally, } X \text{ are } Y. \quad \text{Normally, } X \text{ are } Z.}{\text{Normally, } X \cap Y \text{ are } Z.}$$

Let \mathcal{W} be a preferential model. We assume $\|X\|_{max} \subseteq \|Y\|$ and $\|X\|_{max} \subseteq \|Z\|$.

We have to check that $\|X \cap Y\|_{max} \subseteq \|Z\|$.

For this, let $w \in \|X \cap Y\|_{max}$. We prove $w \in \|Z\|$.

w certainly belongs to $\|X\|$.

If w is maximal in X , then it must be in $\|Z\|$ also.

Otherwise, let $u \in \|X\|_{max}$ be strictly above x .

But then $u \in \|Y\|$. And we see that $u \in \|X \cap Y\|$, contradicting the maximality of w .

Another Soundness Fact

Normally, X are Y . Normally, X are Z .

Normally, $X \cap Y$ are Z .

Let W satisfy the assumptions.

Let $w \in \|X \cap Y\|_{max}$.

If we only knew that $w \in \|X\|_{max}$, then we would have our goal: $w \in \|Z\|$.

So assume that $w \notin \|X\|_{max}$.

Let $w \rightarrow x$ and $x \in \|X\|_{max}$.

This x must belong to $\|Y\|$, since $\|X\|_{max} \subseteq \|Y\|$.

But then we have $w \rightarrow x$ and $x \in \|X \cap Y\|$.

This contradicts the assumption that $w \in \|X \cap Y\|_{max}$.

So indeed, $w \in \|X\|_{max}$. Hence $w \in \|Z\|$.

A Non-Soundness Fact

To show that a putative fact fails in this semantics, we can give a counterexample:

Normally, X are Y .

Normally, Y are Z .

Normally, X are Z .

This fails in the model from before.

Normally, X are Y .

Normally, Z are Y .

Normally, $X \wedge Z$ are Y .

Try to build a model where the assumptions are true and the conclusion is false.

A Logical System

We can also turn System P into a logical system, by making proof trees.

This is just what we did in the case of syllogisms.

(We could have also done it for the logic of announcements. The reason that I didn't do that is that the actual logical principles that we would need are complicated, more so than the homework problems.)

We write $\Gamma \vdash S$ if there is a tree whose root is S , whose leaves are labelled by sentences in Γ , and where at each point in the tree we must match one of our rules.

We then have already shown the soundness: If $\Gamma \vdash S$, then $\Gamma \models S$.

We could then go on to ask about the completeness.

But it is even more important to ask whether System P and the semantics are *right*.

Using the System: A Derived Rule

Normally, X are Y .

Normally, $X \wedge Z$ are $\neg Y$.

Normally, X are $\neg Z$.

$$(1) \quad \frac{\text{Norm } X \wedge Z \text{ are } \neg Y \quad \text{All } \neg Y \text{ are } \neg Y \vee \neg Z}{\text{Norm } X \wedge Z \text{ are } \neg Y \vee \neg Z}$$

$$\frac{\text{tree (1) above} \quad \frac{\text{All } X \wedge \neg Z \text{ are } \neg Y \vee \neg Z}{\text{Norm } X \wedge \neg Z \text{ are } \neg Y \vee \neg Z}}{\text{Norm } (X \wedge Z) \cup (X \wedge \neg Z) \text{ are } \neg Y \vee \neg Z}$$

$$\frac{\text{Norm } X \text{ are } Y \quad \text{Norm } X \text{ are } \neg Y \vee \neg Z}{\text{Norm } X \text{ are } Y \wedge (\neg Y \vee \neg Z)}$$

$$\frac{\text{Norm } X \text{ are } Y \wedge (\neg Y \vee \neg Z)}{\text{Norm } X \text{ are } Y \wedge \neg Z}$$

$$\frac{\text{Norm } X \text{ are } Y \wedge \neg Z}{\text{Norm } X \text{ are } \neg Z}$$

A Probabilistic Semantics for Defaults, I

Now we want to interpret our sentences on *probability spaces*.

First, we take a probability space \mathcal{S} . For each variable X , we associate an event $\|X\|$. (Recall that this means that $\|X\| \subseteq \mathcal{S}$.)

Second, we take we define a semantics using \mathcal{S} and the events:

$$\|\text{Normally, } X \text{ are } Y\| = Pr(\|Y\| \mid \|X\|)$$

This is conditional probability.

We also write

$$\|\text{All } X \text{ are } Y\| = \begin{cases} 1 & \text{if } \|X\| \subseteq \|Y\| \\ 0 & \text{if not} \end{cases}$$

A Probabilistic Semantics for Defaults, II

We say that $\Gamma \models S$ if

$$(\forall \epsilon > 0)(\exists \delta > 0)(\|T\| \geq 1 - \delta \text{ for all } T \in \Gamma) \rightarrow \|S\| > 1 - \epsilon).$$

The idea is that no matter how sure we want to be about the conclusion S , we can find a number δ so that

There is also an analogy to the definition of *continuity of a function f at a point x* in calculus:

$$(\forall \epsilon > 0)(\exists \delta > 0)(|y - x| < \delta \rightarrow |f(y) - f(x)| < \epsilon).$$

Example of a Soundness Proof

We show probabilistic soundness of

$$\frac{\text{Normally, } X \text{ are } Y. \quad \text{Normally, } X \text{ are } Z.}{\text{Normally, } X \cap Y \text{ are } Z.}$$

We are dealing here with $Pr(Z|X)$, $Pr(Y|X)$, and $Pr(Z|X, Y)$.

Recall that

$$Pr(Z|X) = Pr(Z|X, Y)Pr(Y|X) + Pr(Z|X, \neg Y)Pr(\neg Y|X).$$

Hence

$$Pr(Z|X, Y) = \frac{Pr(Z|X) - Pr(Z|X, \neg Y)Pr(\neg Y|X)}{Pr(Y|X)}$$

Assuming that $Pr(Z|X) \rightarrow 1$ and $Pr(Y|X) \rightarrow 1$,

we also have $Pr(\neg Y|X) \rightarrow 0$.

We see that the whole fraction $\rightarrow 1$.

Example of a Soundness Proof

We show that

$$\{\text{Norm } X \text{ are } Y, \text{ Norm } Z \text{ are } Y\} \models \text{Norm } X \cup Z \text{ are } Y.$$

We take some $\epsilon > 0$.

We need a $\delta > 0$, so that for all models and all events $\|X\|$, $\|Y\|$, and $\|Z\|$, if the assumptions have a value $> 1 - \delta$, then the conclusion will have a value $> 1 - \epsilon$.

Rather than $\|X\|$, let's just write X here.

Now the hypotheses have values $Pr(Y|X)$ and $Pr(Y|Z)$. The conclusion has probability $Pr(Y|X \cup Z)$.

Example of a Soundness Proof, Continued

So we assume that $Pr(X \cap Y)/Pr(X) > 1 - \delta$, and also $Pr(Z \cap Y)/Pr(Z) > 1 - \delta$.

Note that

$$Pr(Y|X \cup Z) = \frac{Pr(Y \cap (X \cup Z))}{Pr(X \cup Z)}.$$

Also, $Pr(Y \cap (X \cup Z)) = Pr((Y \cap X) \cup (Y \cap Z)) = Pr(Y \cap X) + Pr(Y \cap Z) - Pr(X \cap Y \cap Z)$.

The Interesting Fact

We started with a set of principles about *Normally* sentences that we believed.

We then had a semantics using preferential models, and a different one using probability.

The interesting fact is that the two different relations $\Sigma \models S$ that we get from these semantics *agree!*

By completeness, they also agree with the totally syntactic formulation $\Sigma \vdash S$.

Cog Sci conclusion:

Some Goals for Default Reasoning

1. Conclusions should be revokable given extra evidence.
2. New *but irrelevant* evidence should not force us to revoke anything.

System P does fine on the first task.

But the second task does not work out.

Normally, birds fly

Normally, red birds fly.

Normally, placophorans live underwater

Normally, red placophorans live underwater.

The Maximum Entropy Approach

Given a model W and background assumptions Γ , and some number ϵ , we want to say when

Normally X are Y

is true.

Plan: take all the probability distributions which make the assumptions true

(using ϵ as we have been doing),

then take the *one* with maximum entropy,

and evaluate the conclusion Normally X are Y with respect to that one probability distribution.

Maximum Entropy in Action

Let Γ be the following sentences:

Penguins are birds. Birds fly.

Penguins do not fly.

Penguins live in the arctic.

Birds have wings.

Animals that fly are mobile.

We can conclude from this in the maximum entropy approach:

White penguins have wings.

Penguins who do not live in the arctic do not fly.

penguins who do not live in the arctic have wings.