Entailment defined

Semantic rules of interpretation tell us how to understand all wffs in terms of specification for non-logical symbols.

But some connections among sentences are independent of the non-logical symbols involved.

e.g. If α is true under $\mathfrak{I}$, then so is $\neg(\beta \wedge \neg\alpha)$,
no matter what $\mathfrak{I}$ is, why α is true, what β is, ...

$S \models \alpha$ iff for every $\mathfrak{I}$, if $\mathfrak{I} \models S$ then $\mathfrak{I} \models \alpha$.

Say that S entails α or α is a logical consequence of S:

In other words: for no $\mathfrak{I}$, $\mathfrak{I} \models S \cup \{\neg\alpha\}$. $S \cup \{\neg\alpha\}$ is unsatisfiable

Special case when S is empty: $\models \alpha$ iff for every $\mathfrak{I}$, $\mathfrak{I} \models \alpha$.

Say that α is valid.

Note: $\{\alpha_1, \alpha_2, ..., \alpha_n\} \models \alpha$ iff $\models (\alpha_1 \wedge \alpha_2 \wedge ... \wedge \alpha_n) \supset \alpha$

finite entailment reduces to validity

Why do we care?

We do not have access to user-intended interpretation of non-logical symbols

But, with entailment, we know that if S is true in the intended interpretation, then so is α.

If the user's view has the world satisfying S, then it must also satisfy α.

There may be other sentences true also; but α is logically guaranteed.

So what about ordinary reasoning?

Dog(fido) ➠ Mammal(fido) ??

Not entailment!

There are logical interpretations where $I[\text{Dog}] \not\subseteq I[\text{Mammal}]$

Key idea of KR:

include such connections explicitly in S
$\forall x[\text{Dog}(x) \supset \text{Mammal}(x)]$
Get: $S \cup \{\text{Dog(fido)}\} \models \text{Mammal(fido)}$

the rest is just details...

Knowledge bases

KB is set of sentences

explicit statement of sentences believed (including any assumed connections among non-logical symbols)

KB $\models \alpha$ α is a further consequence of what is believed

- explicit knowledge: KB
- implicit knowledge: $\{ \alpha \mid \text{KB} \models \alpha \}$

Often non trivial: explicit ➠ implicit

Example:

Three blocks stacked.

Top one is green.

Bottom one is not green.

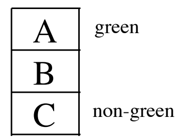

Is there a green block directly on top of a non-green block?

A formalization

S = $\{\text{On(a,b)}, \ \text{On(b,c)}, \ \text{Green(a)}, \ \neg\text{Green(c)}\}$

all that is required

$\alpha = \exists x \exists y [\text{Green}(x) \wedge \neg\text{Green}(y) \wedge \text{On}(x,y)]$

Claim: $S \models \alpha$

Proof:

Let $\mathfrak{I}$ be any interpretation such that $\mathfrak{I} \models S$.

Case 1: $\mathfrak{I} \models \text{Green(b)}$.

$\therefore \ \mathfrak{I} \models \text{Green(b)} \wedge \neg\text{Green(c)} \wedge \text{On(b,c)}$.

$\therefore \ \mathfrak{I} \models \alpha$

Case 2: $\mathfrak{I} \not\models \text{Green(b)}$.

$\therefore \ \mathfrak{I} \models \neg\text{Green(b)}$

$\therefore \ \mathfrak{I} \models \text{Green(a)} \wedge \neg\text{Green(b)} \wedge \text{On(a,b)}$.

$\therefore \ \mathfrak{I} \models \alpha$

Either way, for any $\mathfrak{I}$, if $\mathfrak{I} \models S$ then $\mathfrak{I} \models \alpha$.

So $S \models \alpha$.   QED

Knowledge-based system

Start with (large) KB representing what is explicitly known

e.g. what the system has been told or has learned

Want to influence behaviour based on what is implicit in the KB (or as close as possible)

Requires reasoning

deductive inference:

process of calculating entailments of KB

i.e given KB and any α, determine if KB $\models \alpha$

Process is sound if whenever it produces α, then KB $\models \alpha$

does not allow for plausible assumptions that may be true in the intended interpretation

Process is complete if whenever KB $\models \alpha$, it produces α

does not allow for process to miss some α or be unable to determine the status of α