Entailment

- Spse we assert γ_1, γ_2, ...

  `on(a,b), on(a,b)⇒above(a,b), ...`

Then... you should believe γ_1, γ_2, ...

What else should we believe?

(Should we believe a given β – eg, "`above(a,b)`"?)

- **Entailment** addresses this:
  You should believe all (and only) entailments.

Later: Present ALGORITHM for *computing* these entailments...

Semantics. . . based on Models

"Model" $\equiv$ "completely specified possible world"

Every claim is either true or false

$m \models \varphi$ means
formula φ holds in model m
Otherwise $m \not\models \varphi$

Propositional case: Complete assignment

Eg,

	A	B	C	D
m_1	+	0	0	+

$m_1 \models A$
"A is true in m_1"
"m_1 is a model of A"

Also. . .
$m_1 \models \neg B$ $\quad$ $m_1 \models \neg C$ $\quad$ $m_1 \models D$

- What about
  $A \vee B$? . . . $A \& \neg C \& D$?

Propositional logic: Semantics

- Each model specifies
  { true, false }   for each proposition symbol

Eg,

	A	B	C	D
m_1	+	0	0	+

- Rules for evaluating truth
  with respect to a model m:

$\neg S$	is true iff	S	is false		
$S_1 \wedge S_2$	is true iff	S_1	is true <u>and</u>	S_2	is true
$S_1 \vee S_2$	is true iff	S_1	is true <u>or</u>	S_2	is true
$S_1 \Rightarrow S_2$	is true iff	S_1	is false <u>or</u>	S_2	is true
"	is false iff	S_1	is true <u>and</u>	S_2	is false

$$m \models \neg S \quad \equiv \quad m \not\models S$$
$$m \models S_1 \wedge S_2 \quad \equiv \quad m \models S_1 \quad \text{and} \quad m \models S_2$$
$$m \models S_1 \vee S_2 \quad \equiv \quad m \models S_1 \quad \text{or} \quad m \models S_2$$
$$m \models S_1 \Rightarrow S_2 \quad \equiv \quad m \not\models S_1 \quad \text{or} \quad m \models S_2$$

Semantics of Connectives

P	Q	$\neg P$	$P\&Q$	$P \vee Q$	$P \Rightarrow Q$	$P \Leftrightarrow Q$
0	0	+	0	0	+	+
0	1	+	0	+	+	0
1	0	0	0	+	0	0
1	1	0	+	+	+	+

Just need $\wedge$, $\neg$:

- $P \vee Q$ means $\neg(\neg P \wedge \neg Q)$
- $P \Rightarrow Q$ means $\neg P \vee Q$

  ...counterintuitive:

  What's truth value of "5 is even $\Rightarrow$ Sam is smart"?

- $P \Leftrightarrow Q$ means $(P \Rightarrow Q) \wedge (Q \Rightarrow P)$

- "$\wedge$" – relatively easy
  complete knowledge...

  "$\vee$", "$\neg$" – more difficiult
  partial information

Interpretation

- Vocabulary: `(literals)`

$\texttt{on}_{ab}$:	if `a` is immediately on `b`
$\texttt{on}_{bc}$:	if `b` is immediately on `c`
	...
$\texttt{above}_{ab}$:	if `a` is (somewhere) above `b`
$\texttt{above}_{ac}$:	if `a` is (somewhere) above `c`

- An "interpretation" (aka "world")
  assigns `True` xor `False` to each literal

n literals $\Rightarrow$ 2^n possible worlds, $\mathcal{M}$

	$\texttt{on}_{ab}$	$\texttt{on}_{bc}$	$\cdots$	$\texttt{above}_{ab}$	$\texttt{above}_{ac}$	$\cdots$
m_1	+	+		+	+	
m_2	+	+		+	0	
m_3	+	+		0	+	
$\vdots$			$\vdots$			
m_m	+	0		+	0	
$\vdots$			$\vdots$			
m_{2^n}	0	0		0	0	

Models of a Formula

- Initially all 2^n models are possible

	$\texttt{on}_{ab}$	$\texttt{on}_{bc}$	$\texttt{above}_{ab}$	$\texttt{above}_{bc}$
m_1	0	0	0	0
m_2	0	0	0	+
m_3	0	0	+	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
m_{2^n-1}	+	+	+	0
m_{2^n}	+	+	+	+

- Assertion α ELIMINATES possible worlds

  Eg, $\neg\texttt{on}_{ab}$ eliminates models m where $\models_m \texttt{on}_{ab}$

- $\boxed{\mathcal{M}(\,\alpha\,) \quad = \quad \{\ m \mid\ \models_m \alpha\ \}}$

  is set of all models of α

$\mathcal{M}(\,\neg\texttt{on}_{ab}\,) \quad =$

	$\texttt{on}_{ab}$	$\texttt{on}_{bc}$	$\texttt{above}_{ab}$	$\texttt{above}_{bc}$
m_1	0	0	0	0
m_2	0	0	0	+
m_3	0	0	+	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$m_{2^{n-1}}$	0	+	+	+

Dealing with Boolean Combinations

	$\texttt{on}_{ab}$	$\texttt{on}_{bc}$	$\texttt{above}_{ab}$	$\texttt{above}_{ac}$	$\cdots$
m_1	+	+	+	+	...
m_3	+	+	0	+	...
m_5	+	0	+	+	...
m_7	+	0	0	+	...
m_9	0	+	+	+	...
m_{11}	0	+	0	+	...
m_{13}	0	0	+	+	...
m_{15}	0	0	0	+	...
$\vdots$				$\vdots$	

$$\mathcal{M}(\chi) = \{ m \mid \models_m \chi \}$$

- $\mathcal{M}(\texttt{on}_{ab}) = \{m_1, \ldots, m_8\}$
- $\mathcal{M}(\neg\texttt{on}_{ab}) = \{m_9, \ldots, m_{16}\}$
- $\mathcal{M}(\texttt{on}_{ab} \wedge \texttt{on}_{ac})$
  $= \{m_1, \ldots, m_8\} \cap \{m_1, \ldots, m_4, m_9, \ldots, m_{12}\}$
  $= \{m_1, \ldots, m_4\}$
- $\mathcal{M}(\neg\texttt{on}_{ab} \vee \texttt{above}_{ab})$
  $= \{m_9, \ldots, m_{16}\} \cup \{m_1, m_2, m_5, m_6, m_9, m_{10}, m_{13}, m_{14}\}$
  $= \{m_1, m_2, m_5, m_6, m_9, \ldots, m_{16}\}$

Note:
$$\begin{aligned} \mathcal{M}(\neg\chi) &= \mathcal{M} - \mathcal{M}(\chi) \\ \mathcal{M}(\chi \vee \xi) &= \mathcal{M}(\chi) \cup \mathcal{M}(\xi) \\ \mathcal{M}(\chi \wedge \xi) &= \mathcal{M}(\chi) \cap \mathcal{M}(\xi) \end{aligned}$$

What else to believe?

	$\texttt{on}_{ab}$	$\texttt{on}_{bc}$	$\texttt{above}_{ab}$	$\texttt{above}_{ac}$	$\cdots$
m_1	+	+	+	+	...
m_2	+	+	+	0	...
m_3	+	+	0	+	...
m_4	+	+	0	0	...
m_5	+	0	+	+	...
m_6	+	0	+	0	...
m_7	+	0	0	+	...
m_8	+	0	0	0	...
m_9	0	+	+	+	...
m_{10}	0	+	+	0	...
m_{11}	0	+	0	+	...
m_{12}	0	+	0	0	...
m_{13}	0	0	+	+	...
m_{14}	0	0	+	0	...
m_{15}	0	0	0	+	...
m_{16}	0	0	0	0	...
$\vdots$			$\vdots$		

- Assert $\neg\texttt{on}_{ab} \vee \texttt{above}_{ab}$
  RealWorld $\in \mathcal{M}(\ \neg\texttt{on}_{ab} \vee \texttt{above}_{ab}\ )$
  $= \{m_1, m_2, m_5, m_6, m_9, \ldots m_{16}\}$
- Also... assert $\texttt{on}_{ab}$
  Hence, RealWorld in $\mathcal{M}(\ \neg\texttt{on}_{ab} \vee \texttt{above}_{ab}\ ) \cap \mathcal{M}(\ \texttt{on}_{ab}\ )$
  $= \{m_1, m_2, m_5, m_6\}$

Note: In each such world, $\texttt{above}_{ab}$ also holds!

$$\mathcal{M}(\ [\neg\texttt{on}_{ab} \vee \texttt{above}_{ab}] \wedge \texttt{on}_{ab}\ ) \subset \mathcal{M}(\ \texttt{above}_{ab}\ )$$

$\Rightarrow$ We should believe $\texttt{above}_{ab}$!

Meaning of Entailment

- α is entailment of KB $\boxed{KB \models \alpha}$

  ⋆ α is true in all worlds where KB is true
  ⋆ All models of KB are models of α
  ⋆ $\mathcal{M}(KB) \subseteq \mathcal{M}(\alpha)$

- (As we believe everything in KB, . . .)

  Real world is a model of KB

  $$m_{RW} \in \mathcal{M}(KB)$$

- As $\mathcal{M}(KB) \subseteq \mathcal{M}(\alpha)$

  $\Rightarrow$ $m_{RW} \in \mathcal{M}(\alpha)$
  real-world is model of α

  ie, α must hold!

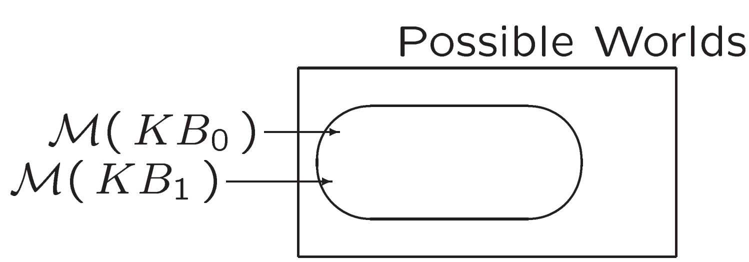

What to believe?

- **Suppose you believe KB and $KB \models \alpha$**

  Then you should believe α !

Why?

1. "Believe KB

   $\Rightarrow$ Real world m_{RW} in $\mathcal{M}(\,KB\,)$

2. $KB \models \alpha$ means $\mathcal{M}(\,KB\,) \subseteq \mathcal{M}(\,\alpha\,)$

$\Rightarrow m_{RW} \in \mathcal{M}(\,\alpha\,)$

$\ldots m_{RW} \models \alpha$

Ie, α holds in the Real World, so you should believe it!

Entailment vs Derivation

- Entailment $KB \models \alpha$
  Semantic Relation:
  α MUST hold whenever KB holds.

- Derivation $KB \vdash_i \alpha$
  Computational (Syntactic) Process:
  Maps $\langle KB, \alpha \rangle$ to $\{$Yes, No$\}$

- $\vdash_i$ can be arbitrary but...
  want $\vdash$ to correspond to $\models$!

GOAL: $\vdash_{SC}$ which returns only entailments:
For any KB, α,

$$KB \vdash_{SC} \alpha \quad \text{if-and-only-if} \quad KB \models \alpha$$

Answering Queries
1. Semantical Approach

- Entailment specifies what we should believe.

- Possible procedure for deciding whether
  $\{\gamma_i\} \models \beta$

MODEL CHECKING
Given n literals, write $n \times 2^n$ table.
Let W be set of all 2^n rows
For each assertion γ_i,
  let $W := W \cap \mathcal{M}(\gamma_i)$
  (ie, eliminate every row that does not satisfy γ_i)
Check β column.
If "+" in each row (of W)
  then Answer "YES: $\{\gamma_i\} \models \beta$"
  else Answer "NO: $\{\gamma_i\} \not\models \beta$"

- Problem: HUGE table!
  (∞ in predicate calculus)

- There is another approach:

Answering Queries
2. Syntactic Approach

- Proof Process (aka "derivation", "deduction") is mechanic process

  . . . for deciding whether
  conclusion (query) follows from
  premises (initial knowledge base)

First consider "forward chaining"

$$KB \quad \leadsto \quad KB'$$

Implemented by . . .
Applying sequence of
individual (sound) *Inference Rules*
to initial set of propositions,
to find new ones

$$\boxed{\text{Sound} \quad \Leftrightarrow \quad \text{preserves truth}}$$

If believe "antecedent", must believe conclusion

- Inference Rules $\neq$ HornClause

Sequence of Inference Rules

1.	$\alpha \ \& \ \beta$
2.	$\alpha \Rightarrow \ \gamma$
3.	$\beta \ \& \ \gamma \ \Rightarrow \ \delta$

$\xRightarrow{\&E \ 1}$

1.	$\alpha \ \& \ \beta$
2.	$\alpha \Rightarrow \ \gamma$
3.	$\beta \ \& \ \gamma \ \Rightarrow \ \delta$
4.	α
5.	β

$\xRightarrow{MP \ 4,2}$

1.	$\alpha \ \& \ \beta$
2.	$\alpha \Rightarrow \ \gamma$
3.	$\beta \ \& \ \gamma \ \Rightarrow \ \delta$
4.	α
5.	β
6.	γ

$\xRightarrow{\&I \ 5,6}$

1.	$\alpha \ \& \ \beta$
2.	$\alpha \Rightarrow \ \gamma$
3.	$\beta \ \& \ \gamma \ \Rightarrow \ \delta$
4.	α
5.	β
6.	γ
7.	$\beta \& \gamma$

$\xRightarrow{MP \ 7,3}$

1.	$\alpha \ \& \ \beta$
2.	$\alpha \Rightarrow \ \gamma$
3.	$\beta \ \& \ \gamma \ \Rightarrow \ \delta$
4.	α
5.	β
6.	γ
7.	$\beta \& \gamma$
8.	δ

New Facts from Old: A Sound Inference Rule

$$
\begin{array}{lllll}
\text{If} & & \text{``}P \Rightarrow Q\text{''} & & \in KB \\
 & & \text{``}P\text{''} & & \in KB \\
\hline
\text{Then} & \text{can add} & \text{``}Q\text{''} & \text{to } KB &
\end{array}
$$

Called "Modus Ponens"

Written:

$$
[MP] \quad \frac{\begin{array}{l} P \Rightarrow Q \\ P \end{array}}{Q}
$$

Verify Soundness

- **Modus Ponens:** $$\begin{array}{c} \alpha \ \Rightarrow \ \beta \\ \alpha \\ \hline \beta \end{array}$$

Truth table:

	α	β	$\alpha \Rightarrow \beta$
*	T	T	T
	T	F	F
	F	T	T
	F	F	T

Consider all worlds where $\left\{ \begin{array}{c} \alpha \\ \alpha \Rightarrow \beta \end{array} \right\}$ both hold.

Observe: β holds here as well!

- **And-Introduction [&I]** $$\begin{array}{c} \alpha \\ \beta \\ \hline \alpha \ \& \ \beta \end{array}$$

Truth table:

	α	β	$\alpha \& \beta$
*	T	T	T
	T	F	F
	F	T	F
	F	F	F

Why Prove Soundness?

- Why use truth-table?
  Isn't "soundness" obvious?

- Do you believe...

  Given
  `Some` **X**`s are` **Y**`s` &
  `Some` **Y**`s are` **Z**`s`
  conclude
  `Some` **X**`s are` **Z**`s` .

Eg,

Given
`Some polygons are rectangles` &
`Some rectangles are squares`
conclude
`Some polygons are squares` .

- If so...

  Given
  `Some men are doctors` &
  `Some doctors are women`
  conclude
  `Some men are women` ?

Sound Inference Rules

$[MP]$ $$\frac{P \Rightarrow Q \qquad P}{Q}$$

$[\forall E]$ $$\frac{\forall x.\, \varphi(x)}{\varphi(\mathtt{A})} \quad \text{for any } \mathtt{A}$$

$[\&I]$ $$\frac{P \qquad Q}{P \wedge Q}$$

$[RC]$ $$\frac{P \Rightarrow Q \qquad Q \Rightarrow R}{P \Rightarrow R}$$

$[MG]$ $$\frac{P \Rightarrow Q \qquad \neg P \Rightarrow Q}{Q}$$

$[\vee D]$ $$\frac{P \vee Q \qquad \neg P}{Q}$$

$[MT]$ $$\frac{P \Rightarrow Q \qquad \neg Q}{\neg P}$$

$[\&E]$ $$\frac{P \wedge Q}{P}$$

$[\vee I]$ $$\frac{P}{P \vee Q}$$

$[\bowtie]$ $$\frac{P \qquad \neg P}{\mathit{Contradiction}}$$

...

Sound Rules of Inference

$$\boxed{KB_0} \xrightarrow{ir_j} \boxed{KB_1}$$

Inference rule ir_j maps KB_0 into KB_1

If ir_j is sound [aka *Truth-Preserving*], then

$$KB_0 \models KB_1$$

Ie, $\mathcal{M}(KB_0) \subset \mathcal{M}(KB_1)$

Hence: If believe KB_0, must believe KB_1.

Possible Worlds

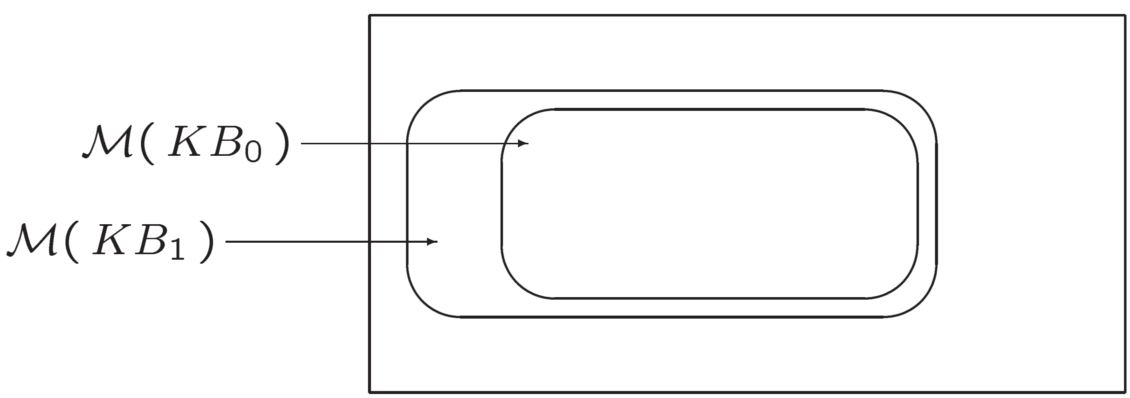

Sound Rules of Inference – con't

- In general, $KB_1 \;=\; KB_0 + \Delta$, so
  $KB_1 \;\models KB_0$

So KB_1, KB_0 hold in EXACTLY same worlds:

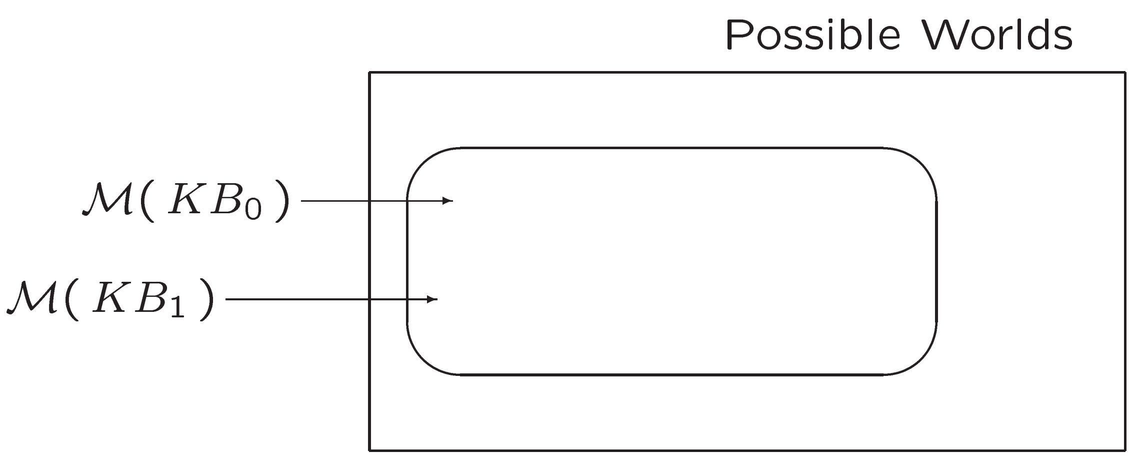

- If each ri_j is sound, then
  sequence $\langle ri_1 \cdots ri_n \rangle$ is sound.

$\boxed{\text{Answering Queries}}$

- Adding Truths (Forward Chaining)
  Given KB_0
  Produce KB_N s.t.

  $$\boxed{KB_0} \overset{ri_1}{\rightsquigarrow} \boxed{KB_1} \cdots \overset{ri_N}{\rightsquigarrow} \boxed{KB_N}$$

  (If $\{ri_j\}_j$ sound, then $KB_0 \models KB_N$)

- Answering Questions (Backward Chaining)
  Given KB_0, σ
  Determine if $KB \models \sigma$

  Requires sound $\langle ri_j \rangle_j$ s.t.

  $$\boxed{KB_0} \overset{ri_1}{\rightsquigarrow} \boxed{KB_1} \cdots \overset{ri_N}{\rightsquigarrow} \boxed{KB_N}$$

  and $\sigma \in KB_N$

$\boxed{\text{Properties of Derivation Process}}$

- $\vdash_\alpha$ is <u>SOUND</u>

  $\Sigma \vdash_\alpha \Psi \quad \Rightarrow \quad \Sigma \models \Psi$
  Produces only "true" results

- $\vdash_\alpha$ is <u>COMPLETE</u>

  $\Sigma \vdash_\alpha \Psi \quad \Leftarrow \quad \Sigma \models \Psi$
  Produces all "true" results

- $\vdash_\alpha$ is <u>DECIDABLE</u>

  $\Sigma \overset{?}{\vdash_\alpha} \Psi$ returns `Y` or `N` in finite time

$\boxed{\textbf{Degenerate } \vdash_\alpha}$

- For any Σ, $\Psi \subset$ WFFs:
  $\Sigma \not\vdash_N \Psi$
  $\Sigma \vdash_P \Psi$

- Notice:
  - $\vdash_N$ is SOUND
    (everything it returns is logically entailed)
  - $\vdash_P$ is COMPLETE
    (it returns everything logically entailed)
  - $\vdash_N$, $\vdash_P$ are DECIDABLE
    (answer every question)

$\boxed{\text{Fundamental Limitation}}$

- For any sufficiently complicated domain

  as complex as arithmetic

- **NO** $\vdash_\alpha$ can be
  SOUND, COMPLETE, DECIDABLE!!

- Reduction to Halting Problem.

$$\left(\begin{array}{l}\text{Not Predicate Calculus's fault:}\\ \quad\text{Reasoning is inherently undecidable,}\\ \quad\text{no manner what formalism used.}\end{array}\right)$$

Responses

- Deals only with WORST-Case!

  "Typical" case can be better.

TradeOffs (to increase efficiency):

? Sacrifice *SOUND*ness?

No — too severe.

? Sacrifice *COMPLETE*ness?

Reasonable... Specific proposals:

- Use only (incomplete set of) Inference Rules
- Use complete set of Inference Rules,

  but limit depth (...stop expanding nodes...)

? Sacrifice *EXPRESSIVE*ness?

[EXPRESSIVEness $\approx$ what can be distinguished.]

Common approach!

(After all, Logic's distinctions caused problems!)

Disallow "$\vee$" "$\neg$" "$\exists$" ...

Implemented Systems

- DataBase Systems

  ≈Sound, Complete, Limited Expressiveness

- **Prolog**

  ≈Sound, Complete, Limited Expressiveness

- General Theorem Provers

  Sound, Complete, Complete Expressiveness

- Production System (Emycin, OPS)

  ≈Sound, ≈Complete, Limited Expressiveness

- Frame Systems

  ?Sound?, ?Complete?, Limited Expressive

- Description Languages

  Sound, Complete, Limited Expressiveness

- Truth Maintenance

“Complexity Cliffs”

- Complexity Cliffs:

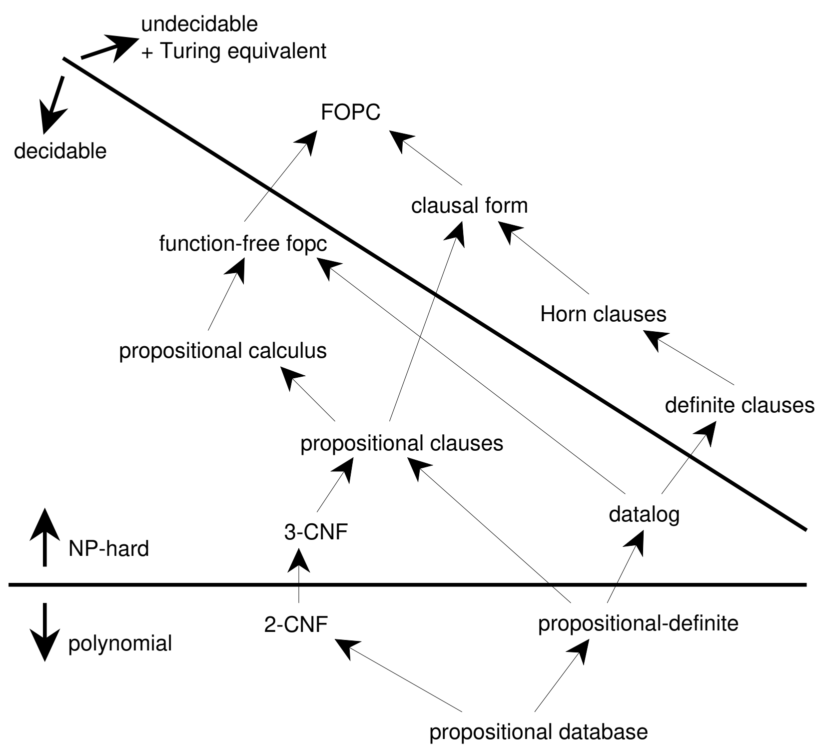

- Be as expressive as possible
  within “tractable” side

- Frames/SemanticNets tractable, but
  not very expressive

  full Predicate Calculus is very expressive
  by not tractable

Description Logics

- Undecidable if $\forall$, $\exists$, $\neg, \vee, \wedge$
- Just use "tractable" subset

  eg, avoid $\neg$, only some types of $\vee$, ...
- Define concepts

`HappyMother` $\equiv$
Woman whose children are all RICH, and all married to pediatricians.

`SuccessfulParent` $\equiv$ PERSON whose DAUGHTERS are married to doctors.

"Subsumption" questions like:

Is every `HappyMother` a `SuccessfulParent`?

What if...

at least 3 daughters?

either MDs or Profs?

Uses: CLASSIC (AT&T)

Financial Management, Database Interfaces, Software Information Systems

How to Reason?

Q: How to reason?

Given KB, q, determine if $KB \models q$?

A: Select Inference Rule IR

Select Fact(s) $\{F_i\}$ from KB

Apply rule IR to Facts $\{F_i\}$

to get new Fact ϕ

. . . Add ϕ to KB

Repeat until find $\phi = q$

Issues: 1. Lots of Inference Rules

Which one to use, when?

2. Is overall system "complete"?

If $\exists$ answer, guaranteed to find it?

Notation

- $\models_m \phi$
  m **is model of formula** ϕ
  Model (interpretation, possible world) m
  assigns formula ϕ to $+$

- Formula is **satisfiable** if it has ≥ 1 model

- Formula is **tautology**
  ("valid" in predicate logic)
  if true in every interpretation
  $a \vee \neg a$; $a \Rightarrow a$

- A formula is a **contradiction**
  (aka "inconsistent")
  if false in every interpretation,
  $a \& \neg a$; $\{a,\ a \Rightarrow b,\ \neg b\}$

- $\Sigma \models \phi$ means
  Whenever Σ is interpreted true, so is ϕ

  . . . ϕ is true in every model of Σ.