Knowledge Representation and Reasoning

George Konidaris
gdk@cs.duke.edu

DUKE COMPUTER SCIENCE

Spring 2016

Knowledge

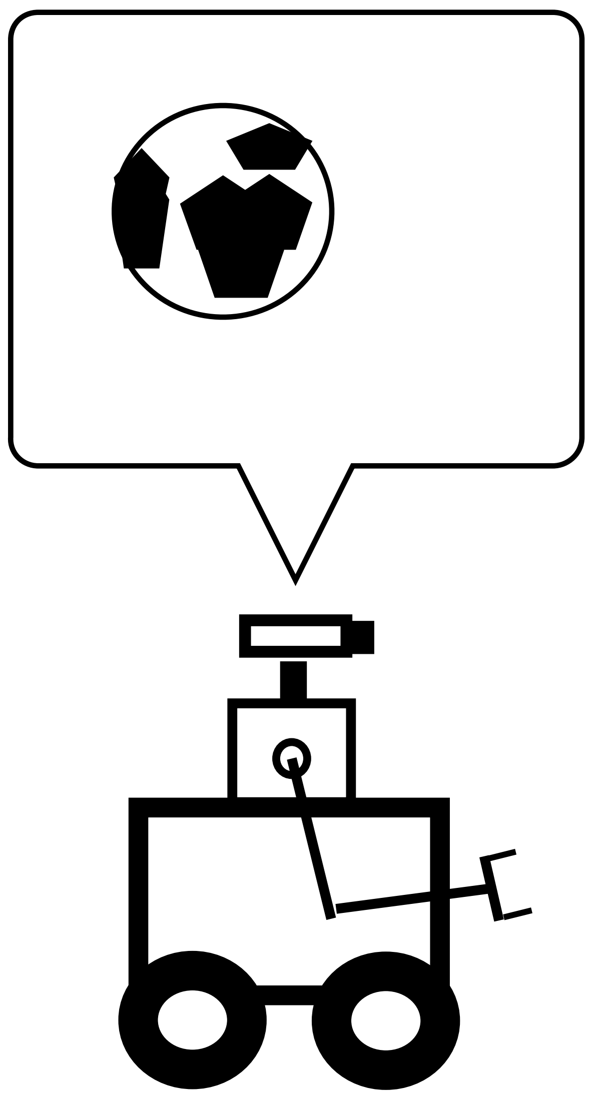

Representation and Reasoning

Represent knowledge about the world.

- Representation language.
- Knowledge base.
- Declarative - *facts* and *rules*.

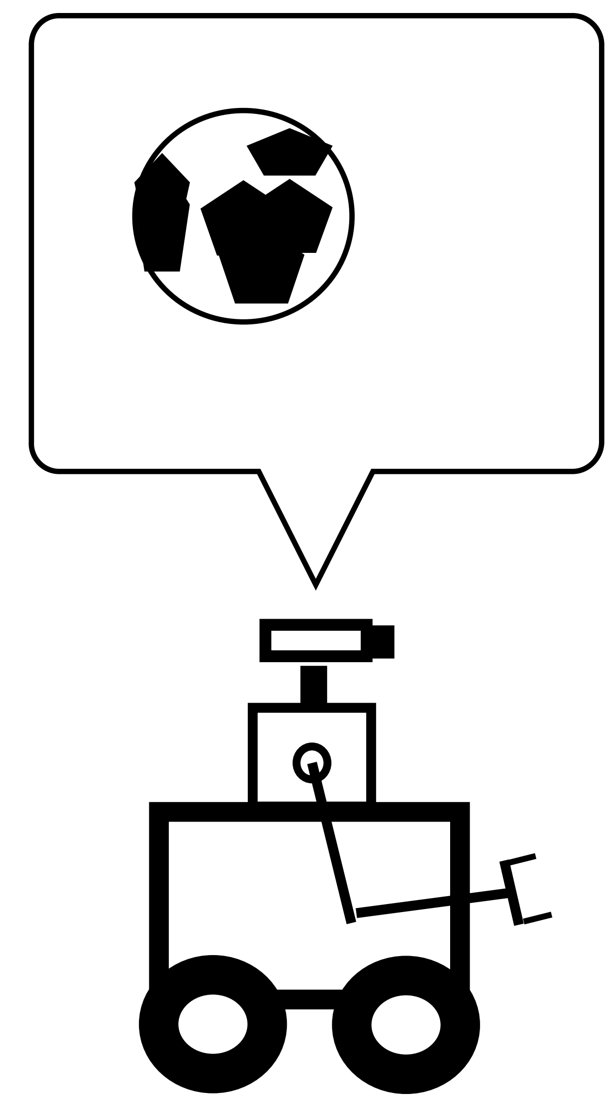

Reason using that represented knowledge.

- Often *asking questions*.
- Inference procedure.
- Heavily dependent on language.

Propositional Logic

Representation language and *set of inference rules* for reasoning about facts that are either true or false.

Model the world as a set of *propositions*:

- *Raining*
- *Cloudy*
- *ExamToday*

Each proposition is either *True* or *False* (though we may not know which).

Propositional Logic

Can combine propositions using **logical operators** to make **sentences** *(syntax vs. semantics)*:

$\neg A$	(not A - **A is** ***False***)
$A \vee B$	(A or B - **one (or both) of A or B is** ***True***)
$A \wedge B$	(A and B - **both A and B are** ***True***)
$A \implies B$	(A implies B - **if A is** ***True*****, so is B**)
$A \iff B$	(A iff B - **A and B both** ***True*** **or both** ***False***)

Two uses of sentence:

- Fact
- Question

Knowledge Base

A list of *sentences that apply to the world*.

For example:

$$Cold$$

$$\neg Raining$$

$$(Raining \vee Cloudy)$$

$$Cold \iff \neg Hot$$

A knowledge base describes *a set of worlds in which these facts and rules are true.*

Knowledge Base

A *model* is a formalization of a "world":

- Set the value of every variable in the KB to *True* or *False*.
- 2^n models possible for n propositions.

Proposition	**Value**
Cold	False
Raining	False
Cloudy	False
Hot	False

Proposition	**Value**
Cold	True
Raining	False
Cloudy	False
Hot	False

...

Proposition	**Value**
Cold	True
Raining	True
Cloudy	True
Hot	True

Models and Worlds

Each sentence has a *truth value* in each model.

Proposition	Value
Cold	True
Raining	False
Cloudy	True
Hot	True

If sentence *a* is true in model *m*, then *m* **satisfies** (or **is a model of**) *a*.

$Cold$	True
$\neg Raining$	True
$(Raining \vee Cloudy)$	True
$Cold \iff \neg Hot$	False

Models and Worlds

The KB specifies a subset of all possible models - those that satisfy all sentences in the KB.

Proposition	Value
Cold	False
Raining	False
Cloudy	False
Hot	False

Proposition	Value
Cold	True
Raining	False
Cloudy	False
Hot	False

...

Proposition	Value
Cold	True
Raining	True
Cloudy	True
Hot	True

Each new piece of knowledge narrows down the set of possible models.

Inference

So if we have a KB, then what?

We'd like to ask it *questions*.

Given:

$$Cold$$

$$\neg Raining$$

$$(Raining \vee Cloudy)$$

$$Cold \iff \neg Hot$$

… we can ask:

$$Hot?$$

Inference: process of deriving new facts from given facts.

Inference (Formally)

DUKE COMPUTER SCIENCE

KB *A* **entails** sentence *B*
if and only if:
every model which satisfies *A*, satisfies *B*.

$$A \models B$$

In other words: if *A* is true then *B* **must be true**.

That's nice, but how do we compute?
Could just enumerate worlds …

Logical Inference

Take a KB, and produce new sentences of knowledge.
Most frequently, determine whether $KB \models Q$

Inference algorithms: search process to find a proof of *Q* using a set of *inference rules*.

Desirable properties:

- Soundness (or truth-preserving)
- Complete

Inference Rules

Form	Description
$(A \land B) \equiv (B \land A)$	Commutivity of $\land$.
$(A \lor B) \equiv (B \lor A)$	Commutivity of $\lor$.
$((A \land B) \land C) \equiv (A \land (B \land C))$	Associativity of $\land$.
$((A \lor B) \lor C) \equiv (A \lor (B \lor C))$	Associativity of $\lor$.
$\neg(\neg A) \equiv A$	Double negative elimination.
$(A \Rightarrow B) \equiv (\neg B \Rightarrow \neg A)$	Contraposition.
$(A \Rightarrow B) \equiv (\neg A \lor B)$	Implication elimination.
$(A \Leftrightarrow B) \equiv (A \Rightarrow B) \land (B \Rightarrow A)$	Biconditional elimination.
$\neg(A \land B) \equiv (\neg A \lor \neg B)$	De Morgan.
$\neg(A \lor B) \equiv (\neg A \land \neg B)$	De Morgan.
$(A \land (B \lor C)) \equiv ((A \land B) \lor (A \land C))$	Distributivity of $\land$ over $\lor$.
$(A \lor (B \land C)) \equiv ((A \lor B) \land (B \lor C))$	Distributivity of $\lor$ over $\land$.

If $A \implies B$, and *A* is true, then *B* is true.

If $A \land B$, then *A* is true, and *B* is true.

Inference Rules

Often written in form:

Start with

Given this knowledge

$$\frac{A \vee B, \neg B}{A}$$

can infer this

Proofs

For example, given KB:

$Cold$

$\neg Raining$

$(Raining \vee Cloudy)$

$Cold \iff \neg Hot$

Inference:

$Cold = True$

$True \iff \neg Hot$

$\neg Hot = True$

$Hot = False$

We ask:

$Hot?$

Resolution

The following inference rule is **both sound and complete:**

$$\frac{l_1 \vee ... \vee l_k, m_1 \vee ... \vee m_n}{l_1 \vee ...l_{i-1} \vee l_{i+1} \vee ...l_k \vee m_1 \vee ...m_{j-1} \vee m_{j+1} \vee ...m_n}$$

where $l_i = \neg m_j$

This is called **resolution**. It is sound and complete when combined with a sound and complete search algorithm.

The World and the Model

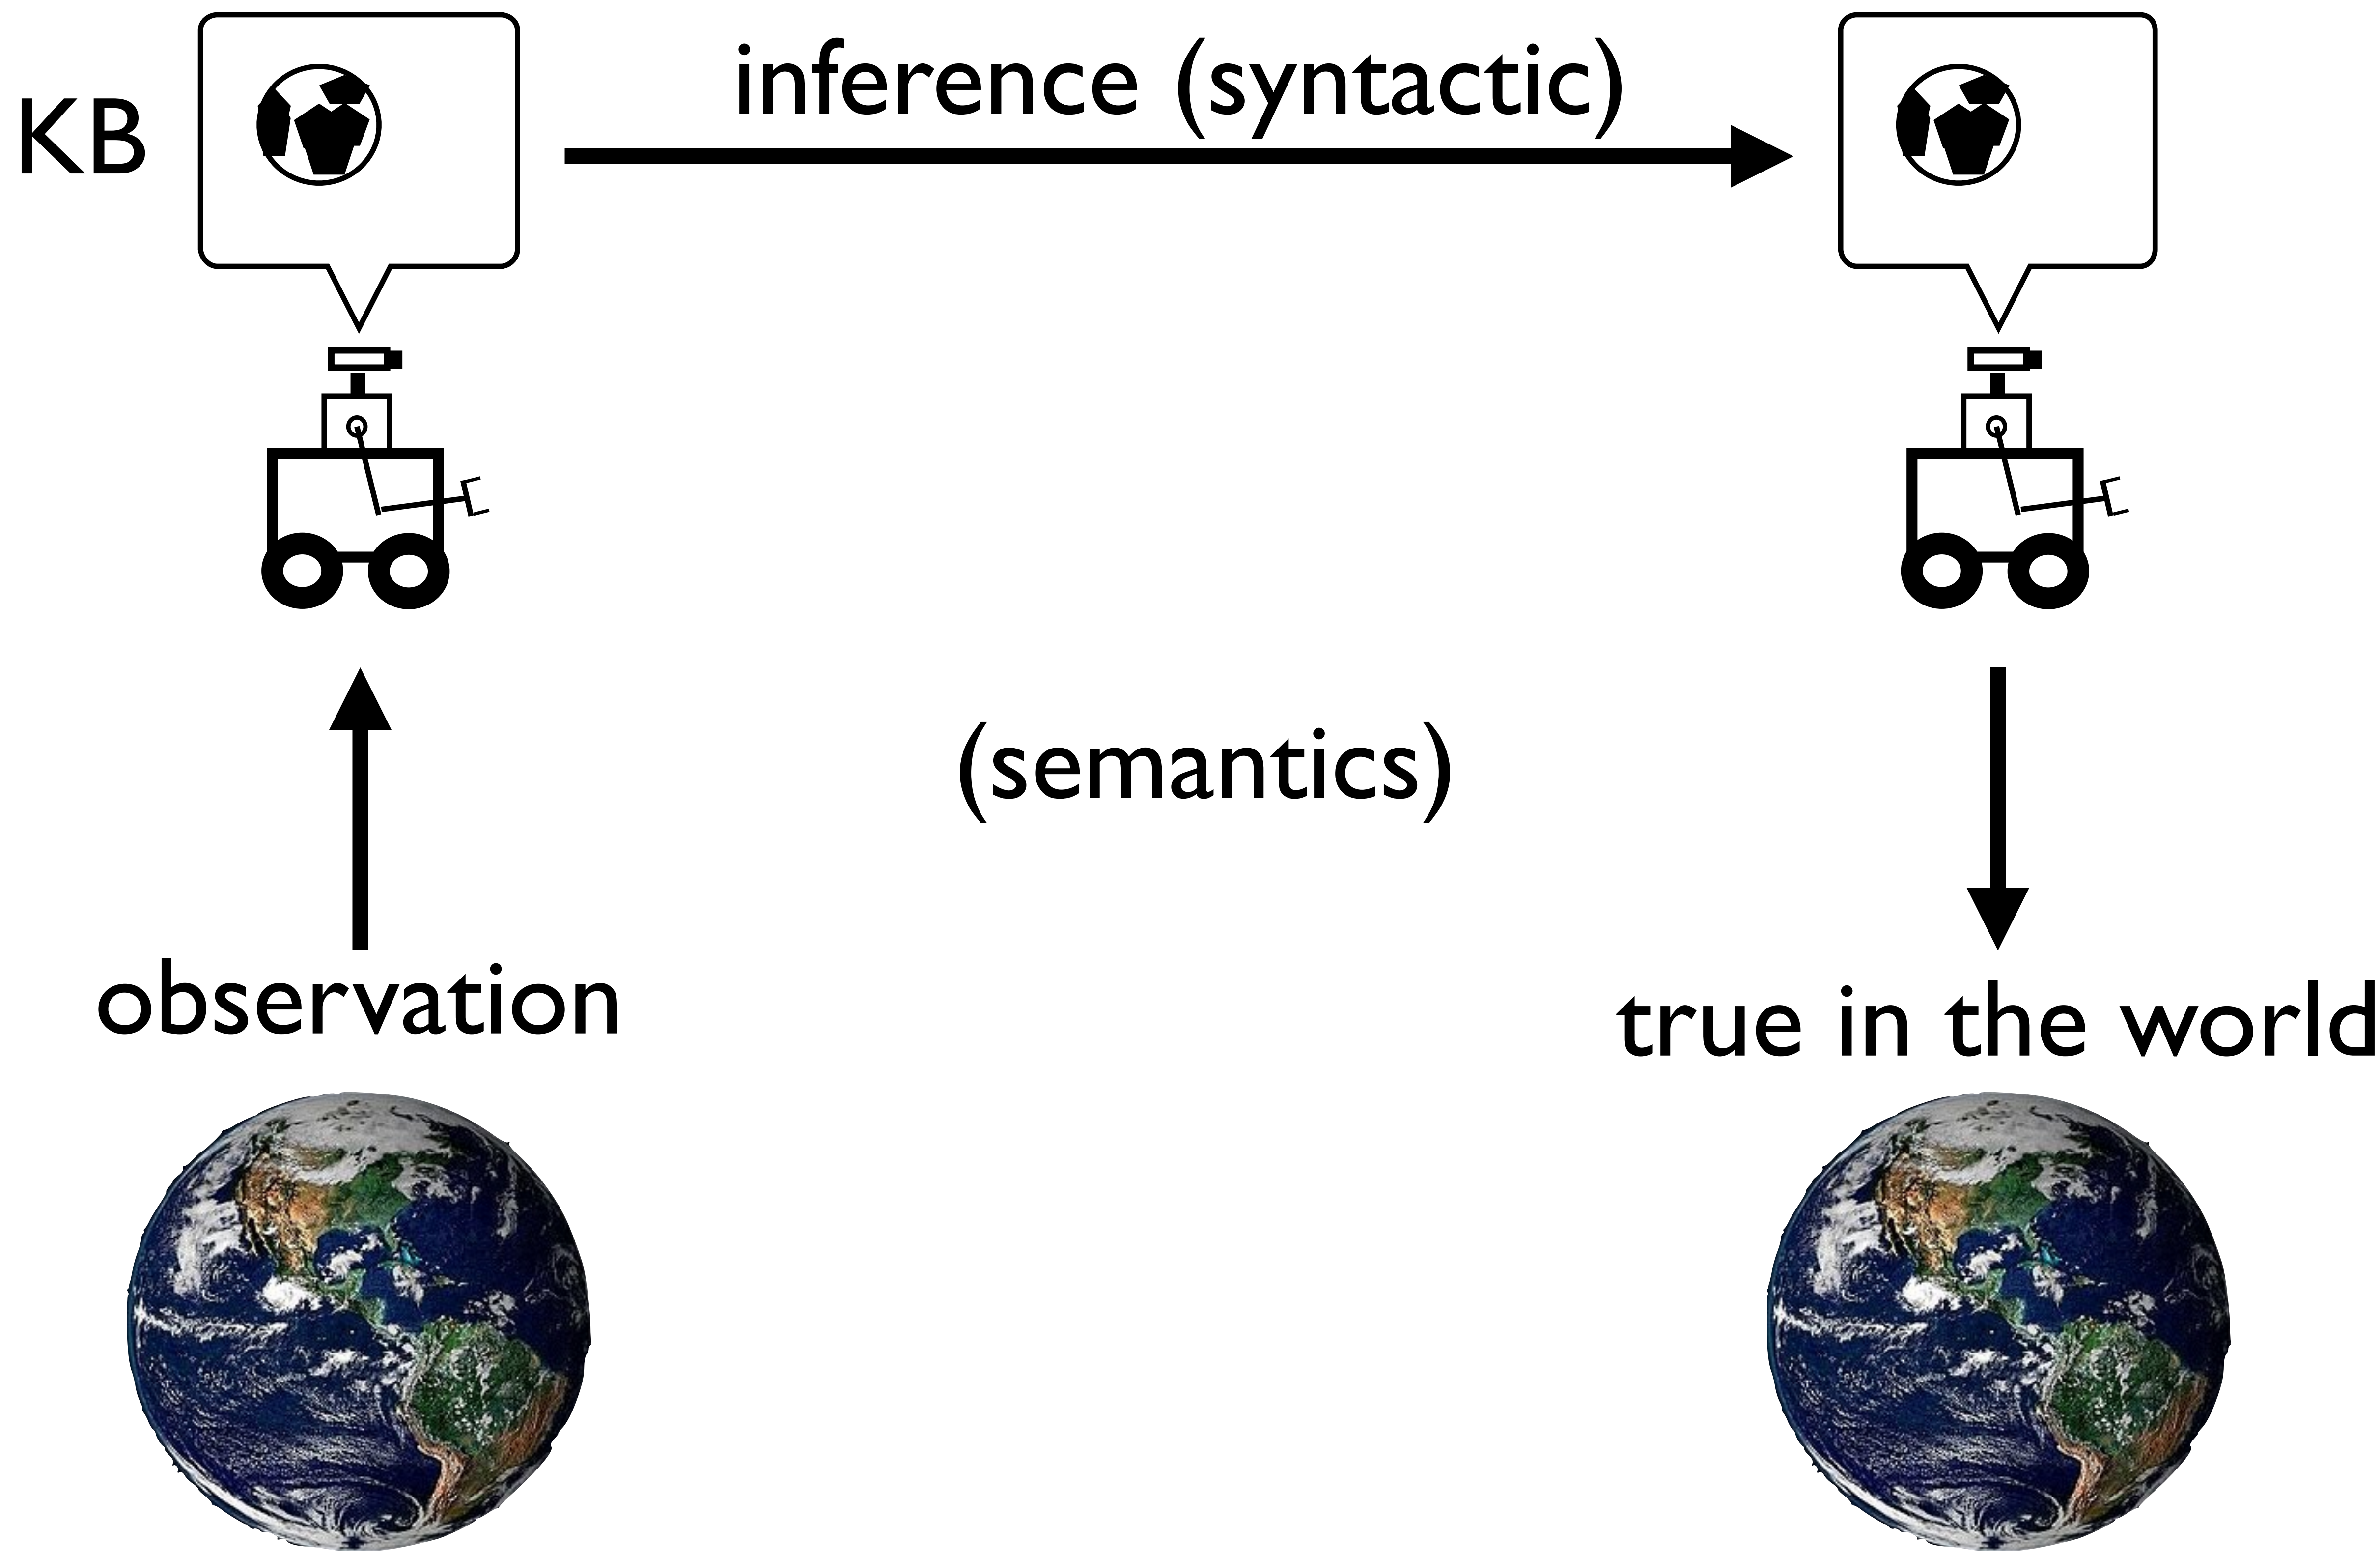

DENDRAL and MYCIN

"**Expert Systems**" - knowledge based.

DENDRAL: (Feigenbaum et al. ~1965)

- Identify unknown organic molecules
- Eliminate most "chemically implausible" hypotheses.

MYCIN: (Shortliffe et al., 1970s)

- Identify bacteria causing severe infections.
- "research indicated that it proposed an acceptable therapy in about 69% of cases, which was *better than the performance of infectious disease experts*."

Major issue: **the Knowledge Bottleneck.**