Logical agents

Chapter 6

Outline

- ♢ Knowledge bases
- ♢ Wumpus world
- ♢ Logic in general
- ♢ Propositional (Boolean) logic
- ♢ Normal forms
- ♢ Inference rules

Knowledge bases

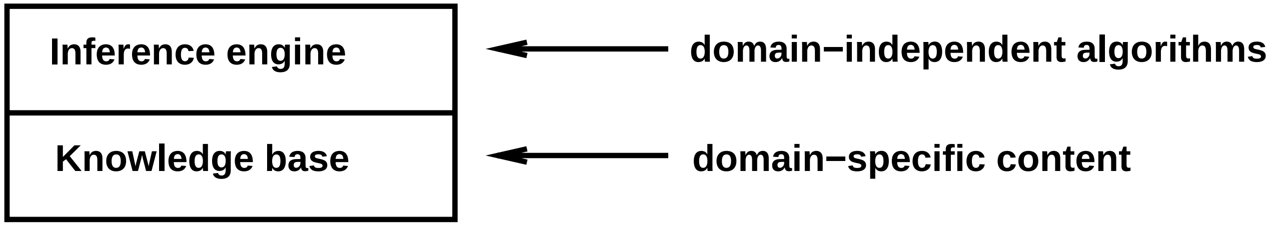

Knowledge base = set of <u>sentences</u> in a <u>formal</u> language

<u>Declarative</u> approach to building an agent (or other system):
    TELL it what it needs to know

Then it can ASK itself what to do—answers should follow from the KB

Agents can be viewed at the <u>knowledge level</u>
    i.e., what they know, regardless of how implemented

Or at the <u>implementation level</u>
    i.e., data structures in KB and algorithms that manipulate them

A simple knowledge-based agent

function KB-AGENT(*percept*) **returns** an *action*
  static: *KB*, a knowledge base
    t, a counter, initially 0, indicating time

  TELL(*KB*, MAKE-PERCEPT-SENTENCE(*percept*, *t*))
  action ← ASK(*KB*, MAKE-ACTION-QUERY(*t*))
  TELL(*KB*, MAKE-ACTION-SENTENCE(*action*, *t*))
  t ← *t* + 1
  return *action*

The agent must be able to:
- Represent states, actions, etc.
- Incorporate new percepts
- Update internal representations of the world
- Deduce hidden properties of the world
- Deduce appropriate actions

Wumpus World PAGE description

<u>Percepts</u> Breeze, Glitter, Smell

<u>Actions</u> Left turn, Right turn,
Forward, Grab, Release, Shoot

<u>Goals</u> Get gold back to start
without entering pit or wumpus square

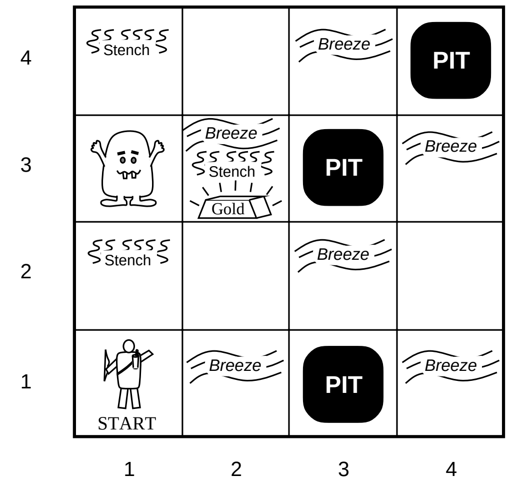

<u>Environment</u>

Squares adjacent to wumpus are smelly
Squares adjacent to pit are breezy
Glitter if and only if gold is in the same square
Shooting kills the wumpus if you are facing it
Shooting uses up the only arrow
Grabbing picks up the gold if in the same square
Releasing drops the gold in the same square

Wumpus world characterization

Is the world deterministic??

Is the world fully accessible??

Is the world static??

Is the world discrete??

Wumpus world characterization

Is the world deterministic?? Yes—outcomes exactly specified

Is the world fully accessible?? No—only local perception

Is the world static?? Yes—Wumpus and Pits do not move

Is the world discrete?? Yes

Exploring a wumpus world

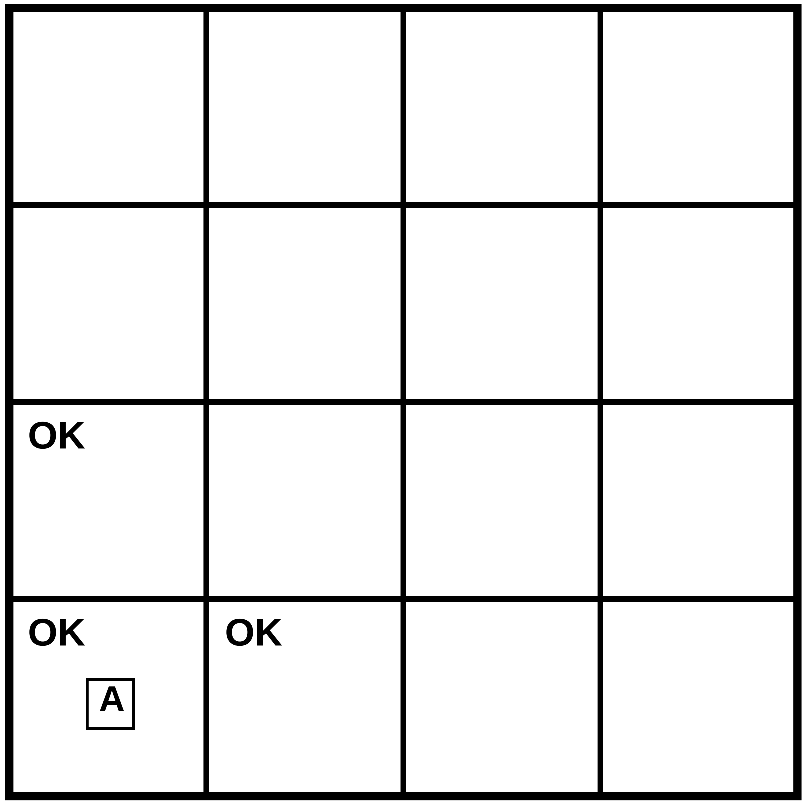

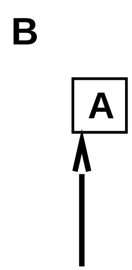
B
A

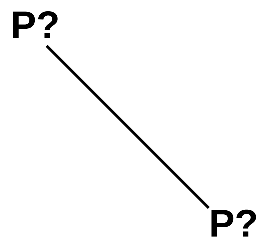
P?
P?

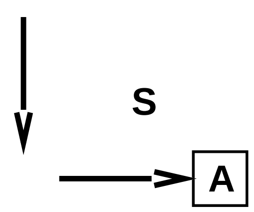
S
A

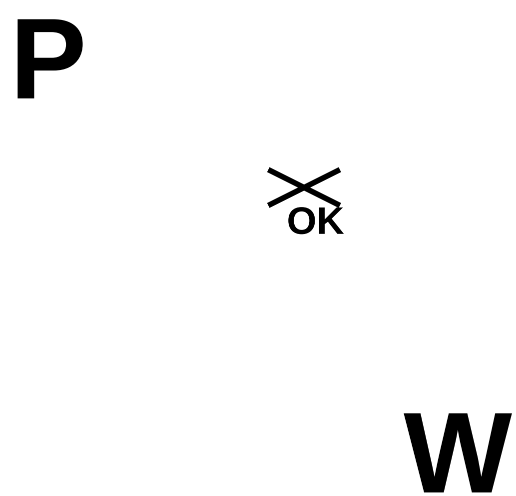
P
~~OK~~
W

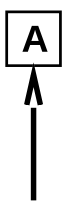
A

OK

OK

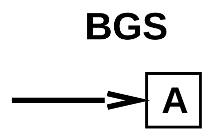
BGS
A

Other tight spots

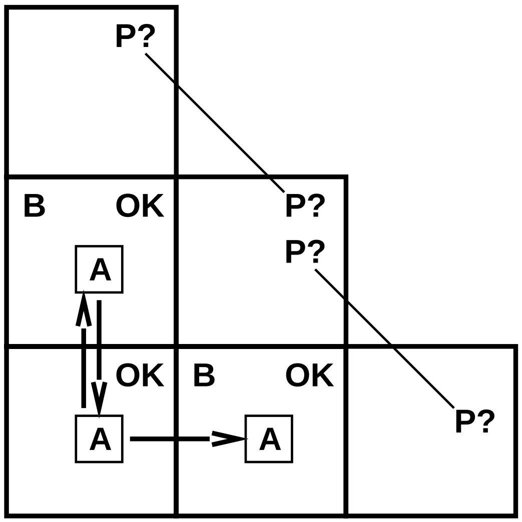

Breeze in (1,2) and (2,1)
$\Rightarrow$ no safe actions

Assuming pits uniformly distributed,
(2,2) is most likely to have a pit

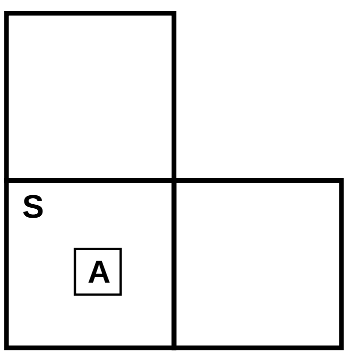

Smell in (1,1)
$\Rightarrow$ cannot move
Can use a strategy of <u>coercion</u>:
shoot straight ahead
wumpus was there $\Rightarrow$ dead $\Rightarrow$ safe
wumpus wasn't there $\Rightarrow$ safe

Logic in general

Logics are formal languages for representing information
such that conclusions can be drawn

Syntax defines the sentences in the language

Semantics define the "meaning" of sentences;
i.e., define truth of a sentence in a world

E.g., the language of arithmetic

$x+2 \geq y$ is a sentence; $x2+y>$ is not a sentence

$x+2 \geq y$ is true iff the number $x+2$ is no less than the number y

$x+2 \geq y$ is true in a world where $x=7,\ y=1$
$x+2 \geq y$ is false in a world where $x=0,\ y=6$

Types of logic

Logics are characterized by what they commit to as "primitives"

Ontological commitment: what exists—facts? objects? time? beliefs?

Epistemological commitment: what states of knowledge?

Language	Ontological Commitment	Epistemological Commitment
Propositional logic	facts	true/false/unknown
First-order logic	facts, objects, relations	true/false/unknown
Temporal logic	facts, objects, relations, times	true/false/unknown
Probability theory	facts	degree of belief $0 \ldots 1$
Fuzzy logic	degree of truth	degree of belief $0 \ldots 1$

Entailment

$$KB \models \alpha$$

Knowledge base KB <u>entails</u> sentence α
if and only if
α is true in all worlds where KB is true

E.g., the KB containing "the Giants won" and "the Reds won" entails "Either the Giants won or the Reds won"

Models

Logicians typically think in terms of models, which are formally structured worlds with respect to which truth can be evaluated

We say m is a model of a sentence α if α is true in m

$M(\alpha)$ is the set of all models of α

Then $KB \models \alpha$ if and only if $M(KB) \subseteq M(\alpha)$

E.g. $KB =$ Giants won and Reds won
$\alpha =$ Giants won

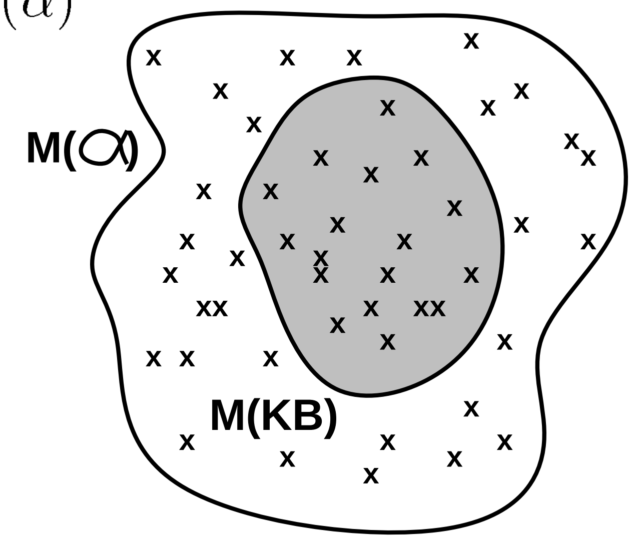

Inference

$KB \vdash_i \alpha$ = sentence α can be derived from KB by procedure i

Soundness: i is sound if
whenever $KB \vdash_i \alpha$, it is also true that $KB \models \alpha$

Completeness: i is complete if
whenever $KB \models \alpha$, it is also true that $KB \vdash_i \alpha$

Preview: we will define a logic (first-order logic) which is expressive enough to say almost anything of interest, and for which there exists a sound and complete inference procedure.

That is, the procedure will answer any question whose answer follows from what is known by the KB.

Propositional logic: Syntax

Propositional logic is the simplest logic—illustrates basic ideas

The proposition symbols P_1, P_2 etc are sentences

If S is a sentence, $\neg S$ is a sentence

If S_1 and S_2 is a sentence, $S_1 \wedge S_2$ is a sentence

If S_1 and S_2 is a sentence, $S_1 \vee S_2$ is a sentence

If S_1 and S_2 is a sentence, $S_1 \;\Rightarrow\; S_2$ is a sentence

If S_1 and S_2 is a sentence, $S_1 \;\Leftrightarrow\; S_2$ is a sentence

Propositional logic: Semantics

Each model specifies true/false for each proposition symbol

E.g.	A	B	C
	$True$	$True$	$False$

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
i.e.,	is false iff	S_1	is true <u>and</u>	S_2	is false
$S_1 \Leftrightarrow S_2$	is true iff	$S_1 \Rightarrow S_2$	is true <u>and</u>	$S_2 \Rightarrow S_1$	is true

Propositional inference: Enumeration method

Let $\alpha = A \vee B$ and $KB = (A \vee C) \wedge (B \vee \neg C)$

Is it the case that $KB \models \alpha$?
Check all possible models—α must be true wherever KB is true

A	B	C	$A \vee C$	$B \vee \neg C$	KB	α
$False$	$False$	$False$				
$False$	$False$	$True$				
$False$	$True$	$False$				
$False$	$True$	$True$				
$True$	$False$	$False$				
$True$	$False$	$True$				
$True$	$True$	$False$				
$True$	$True$	$True$				

Propositional inference: Solution

A	B	C	$A \vee C$	$B \vee \neg C$	KB	α
$False$	$False$	$False$	$False$	$True$	$False$	$False$
$False$	$False$	$True$	$True$	$False$	$False$	$False$
$False$	$True$	$False$	$False$	$True$	$False$	$True$
$False$	$True$	$True$	$True$	$True$	$True$	$True$
$True$	$False$	$False$	$True$	$True$	$True$	$True$
$True$	$False$	$True$	$True$	$False$	$False$	$True$
$True$	$True$	$False$	$True$	$True$	$True$	$True$
$True$	$True$	$True$	$True$	$True$	$True$	$True$

Normal forms

Other approaches to inference use syntactic operations on sentences, often expressed in standardized forms

<u>Conjunctive Normal Form</u> (CNF—universal)

$\quad$ $\mathit{conjunction}$ of $\underbrace{\mathit{disjunctions} \text{ of } \mathit{literals}}_{\mathit{clauses}}$

$\quad$ E.g., $(A \vee \neg B) \wedge (B \vee \neg C \vee \neg D)$

<u>Disjunctive Normal Form</u> (DNF—universal)

$\quad$ $\mathit{disjunction}$ of $\underbrace{\mathit{conjunctions} \text{ of } \mathit{literals}}_{\mathit{terms}}$

$\quad$ E.g., $(A \wedge B) \vee (A \wedge \neg C) \vee (A \wedge \neg D) \vee (\neg B \wedge \neg C) \vee (\neg B \wedge \neg D)$

<u>Horn Form</u> (restricted)

$\quad$ $\mathit{conjunction}$ of $\mathit{Horn\ clauses}$ (clauses with ≤ 1 positive literal)

$\quad$ E.g., $(A \vee \neg B) \wedge (B \vee \neg C \vee \neg D)$

$\quad$ Often written as set of implications:

$\quad$ $B \;\Rightarrow\; A$ and $(C \wedge D) \;\Rightarrow\; B$

Validity and Satisfiability

A sentence is valid if it is true in all models

e.g., $A \vee \neg A$, $\quad A \Rightarrow A$, $\quad (A \wedge (A \Rightarrow B)) \Rightarrow B$

Validity is connected to inference via the Deduction Theorem:

$KB \models \alpha$ if and only if $(KB \Rightarrow \alpha)$ is valid

A sentence is satisfiable if it is true in some model

e.g., $A \vee B$, $\quad C$

A sentence is unsatisfiable if it is true in no models

e.g., $A \wedge \neg A$

Satisfiability is connected to inference via the following:

$KB \models \alpha$ if and only if $(KB \wedge \neg \alpha)$ is unsatisfiable

i.e., prove α by *reductio ad absurdum*

Proof methods

Proof methods divide into (roughly) two kinds:

Model checking

truth table enumeration (sound and complete for propositional)

heuristic search in model space (sound but incomplete)

e.g., the GSAT algorithm (Ex. 6.15)

Application of inference rules

Legitimate (sound) generation of new sentences from old

Proof = a sequence of inference rule applications

Can use inference rules as operators in a standard search alg.

Inference rules for propositional logic

Resolution (for CNF): complete for propositional logic

$$\frac{\alpha \vee \beta, \qquad \neg\beta \vee \gamma}{\alpha \vee \gamma}$$

Modus Ponens (for Horn Form): complete for Horn KBs

$$\frac{\alpha_1, \ldots, \alpha_n, \qquad \alpha_1 \wedge \cdots \wedge \alpha_n \Rightarrow \beta}{\beta}$$

Can be used with forward chaining or backward chaining

Summary

Logical agents apply <u>inference</u> to a <u>knowledge base</u>
to derive new information and make decisions

Basic concepts of logic:

- <u>syntax</u>: formal structure of <u>sentences</u>
- <u>semantics</u>: <u>truth</u> of sentences wrt <u>models</u>
- <u>entailment</u>: necessary truth of one sentence given another
- <u>inference</u>: deriving sentences from other sentences
- <u>soundess</u>: derivations produce only entailed sentences
- <u>completeness</u>: derivations can produce all entailed sentences

Wumpus world requires the ability to represent partial and negated information, reason by cases, etc.

Propositional logic suffices for some of these tasks

Truth table method is sound and complete for propositional logic