

Artificial Intelligence 1: Logic agents

Lecturer: Tom Lenaerts
SWITCH, Vlaams Interuniversitair Instituut voor Biotechnologie



"Thinking Rationally"

Computational models of human "thought" processes

Computational models of human behavior

Computational systems that "think" rationally

Computational systems that behave rationally

AI 1
24 maart
2008

Pag. 2

Logical Agents

Reflex agents find their way from Arad to Bucharest by dumb luck

Chess program calculates legal moves of its king, but doesn't know that no piece can be on 2 different squares at the same time

Logic (Knowledge-Based) agents combine general knowledge with current percepts to infer hidden aspects of current state prior to selecting actions

- Crucial in partially observable environments

AI 1
24 maart
2008

Pag. 3

Outline

Knowledge-based agents

Wumpus world

Logic in general

Propositional and first-order logic


- Inference, validity, equivalence and satisfiability
- Reasoning patterns
 - Resolution
 - Forward/backward chaining

AI 1
24 maart
2008

Pag. 4

Knowledge Base

Knowledge Base – set of sentences represented in a knowledge representation language and represents assertions about the world.



Inference rule: when one ASKs questions of the KB, the answer should *follow* from what has been TELLED to the KB previously.

AI 1
24 maart
2008

Pag. 5

Generic KB-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*

AI 1
24 maart
2008

Pag. 6

Abilities KB agent

Agent must be able to:

- Represent states and actions,
- Incorporate new percepts
- Update internal representation of the world
- Deduce hidden properties of the world
- Deduce appropriate actions

AI 1
24 maart
2008 Pag. 7

Description level


The KB agent is similar to agents with internal state

Agents can be described at different levels

- Knowledge level
 - What they know, regardless of the actual implementation. (Declarative description)
- Implementation level
 - Data structures in KB and algorithms that manipulate them e.g propositional logic and resolution.

AI 1
24 maart
2008 Pag. 8

A Typical Wumpus World



AI 1
24 maart
2008 Pag. 9

Wumpus World PEAS

Performance measure
gold +1000, death -1000
-1 per step, -10 for using the arrow

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

Sensors Breeze, Glitter, Smell

Actuators Left turn, Right turn,
Forward, Grab, Release, Shoot

AI 1
24 maart
2008 Pag. 10

Wumpus World Characterization

- Observable?
- Deterministic?
- Episodic?
- Static?
- Discrete?
- Single-agent?

AI 1
24 maart
2008 Pag. 11

Wumpus World Characterization

- Observable? No, only local perception
- Deterministic?
- Episodic?
- Static?
- Discrete?
- Single-agent?

AI 1
24 maart
2008 Pag. 12

Wumpus World Characterization

Observable? No, only local perception
Deterministic? Yes, outcome exactly specified
Episodic?
Static?
Discrete?
Single-agent?

AI 1
24 maart 2008 Pag. 13

Wumpus World Characterization

Observable? No, only local perception
Deterministic? Yes, outcome exactly specified
Episodic? No, sequential at the level of actions
Static?
Discrete?
Single-agent?

AI 1
24 maart 2008 Pag. 14

Wumpus World Characterization

Observable? No, only local perception
Deterministic? Yes, outcome exactly specified
Episodic? No, sequential at the level of actions
Static? Yes, Wumpus and pits do not move
Discrete?
Single-agent?

AI 1
24 maart 2008 Pag. 15

Wumpus World Characterization

Observable? No, only local perception
Deterministic? Yes, outcome exactly specified
Episodic? No, sequential at the level of actions
Static? Yes, Wumpus and pits do not move
Discrete? Yes
Single-agent?


AI 1
24 maart 2008 Pag. 16

Wumpus World Characterization

Observable? No, only local perception
Deterministic? Yes, outcome exactly specified
Episodic? No, sequential at the level of actions
Static? Yes, Wumpus and pits do not move
Discrete? Yes
Single-agent? Yes, Wumpus is essentially a natural feature.

AI 1
24 maart 2008 Pag. 17

Exploring the Wumpus World



[1,1] The KB initially contains the rules of the environment. The first percept is [none, none, none, none, none], move to safe cell e.g. 2,1
[2,1] breeze which indicates that there is a pit in [2,2] or [3,1], return to [1,1] to try next safe cell

AI 1
24 maart 2008 Pag. 18


Exploring the Wumpus World



[1,2] Stench in cell which means that wumpus is in [1,3] or [2,2]
 YET ... not in [1,1]
 YET ... not in [2,2] or stench would have been detected in [2,1]
 THUS ... wumpus is in [1,3]
 THUS [2,2] is safe because of lack of breeze in [1,2]
 THUS pit in [1,3]
 move to next safe cell [2,2]

AI 1
 24 maart 2008 Pag. 19

Exploring the Wumpus World



[2,2] move to [2,3]
 [2,3] detect glitter, smell, breeze
 THUS pick up gold
 THUS pit in [3,3] or [2,4]

AI 1
 24 maart 2008 Pag. 20

What is a logic?

A formal language

- Syntax - what expressions are legal (well-formed)
- Semantics - what legal expressions mean
 - in logic the truth of each sentence with respect to each possible world.

E.g the language of arithmetic

- $X+2 >= y$ is a sentence, $x+2+y$ is not a sentence
- $X+2 >= y$ is true in a world where $x=7$ and $y=1$
- $X+2 >= y$ is false in a world where $x=0$ and $y=6$

AI 1
 24 maart 2008 Pag. 21

Entailment

One thing follows from another

$$KB \models \alpha$$

KB entails sentence α if and only if α is true in worlds where KB is true.

E.g. $x+y=4$ entails $4=x+y$

Entailment is a relationship between sentences that is based on semantics.

AI 1
 24 maart 2008 Pag. 22

Models

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


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

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

AI 1
 24 maart 2008 Pag. 23

Wumpus world model

Situation after detecting nothing in [1,1], moving right, breeze in [2,1]




Consider possible models for ?s assuming only pits

3 Boolean choices \Rightarrow 8 possible models


AI 1
 24 maart 2008 Pag. 24

Wumpus world model



AI 1
24 maart 2008 Pag. 25


Wumpus world model



$KB = \text{wumpus-world rules} + \text{observations}$

AI 1
24 maart 2008 Pag. 26


Wumpus world model



$KB = \text{wumpus-world rules} + \text{observations}$
 $\alpha_1 = \text{"[1,2] is safe", } KB \models \alpha_1, \text{ proved by model checking}$

AI 1
24 maart 2008 Pag. 27


Wumpus world model



$KB = \text{wumpus-world rules} + \text{observations}$

AI 1
24 maart 2008 Pag. 28

Wumpus world model



$KB = \text{wumpus-world rules} + \text{observations}$
 $\alpha_2 = \text{"[2,2] is safe", } KB \not\models \alpha_2$

AI 1
24 maart 2008 Pag. 29

Logical inference

The notion of entailment can be used for logic inference.

- Model checking (see wumpus example): enumerate all possible models and check whether α is true.

If an algorithm only derives entailed sentences it is called *sound* or *truth preserving*.


- Otherwise it just makes things up.
 α is sound if whenever $KB \models \alpha$ it is also true that $KB \models \alpha$

Completeness: the algorithm can derive any sentence that is entailed.

α is complete if whenever $KB \models \alpha$ it is also true that $KB \models \alpha$

AI 1
24 maart 2008 Pag. 30

Schematic perspective



If KB is true in the real world, then any sentence α derived From KB by a sound inference procedure is also true in the real world.