Artificial Intelligence

Lecture 5

Ph. D. Lect. Horia Popa Andreescu

2012-2013 3rd year, semester 5

- The slides for this lecture are based (partially) on chapter 7 of the Stuart Russel Lecture Notes [R, ch7], and on the same chapter from Russel & Norvig's book [RN, ch. 7]

Knowledge-Based agent

- A Knowledge Base (KB) is, informally a set of sentences,
  - each sentence is expressed in a “knowledge representation language”
  - and it represents some assertion about the world.
- Features of a KB:
  - It is possible to add new sentences and to query what is known (the tasks are called TELL and ASK)
  - Deriving new sentences from old ones is called inference
  - The new sentences are based on the knowledge already known to the KB, which was TELLed to the KB previously

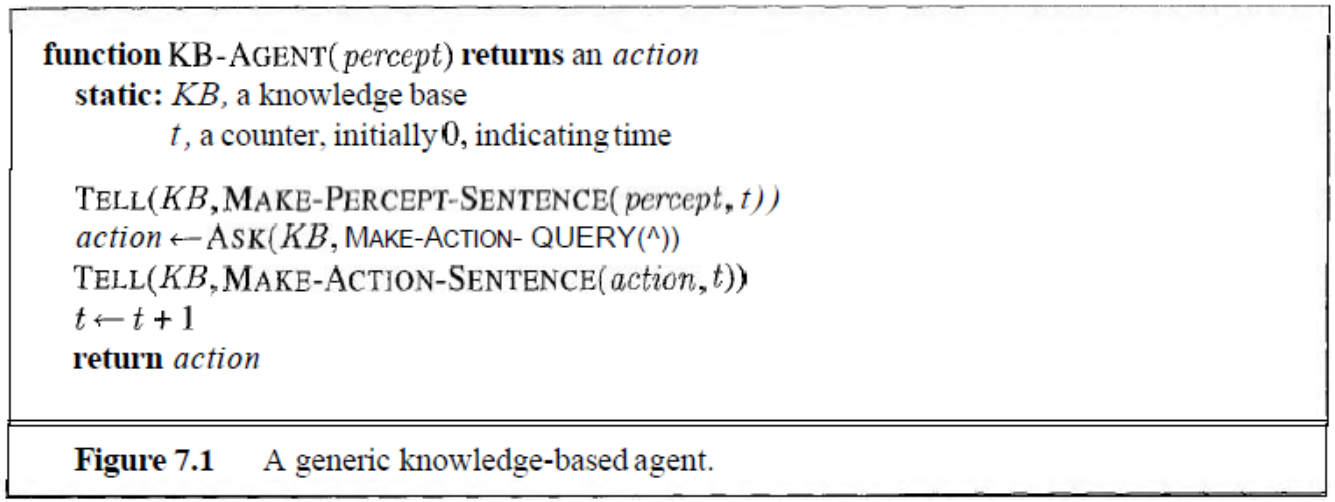

```
function KB-AGENT(percept) returns an action
  static: KB, a knowledge base
          t, a counter, initially 0, indicating time

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

Figure 7.1 A generic knowledge-based agent.

The agent from Figure 7.1 [RN, 196] takes as input a knowledge (a percept) and returns an action.
The agent maintains a knowledge base (KB), initially composed of background knowledge.
Each time the agent program is called, it does three things:
1) It TELLS the knowledge base what it perceives.
2) It ASKS the knowledge base what action it should perform.
3) The agent records its choice with TELL and executes the action.

Declarative Vs. Procedural knowledge

- Designing the representation language to make it easy to express this knowledge in the form of sentences simplifies the construction problem enormously**.**
- This is called **the declarative approach** to system building.[RN 197]
- In contrast, the **procedural approach** encodes desired behaviors directly as program code; minimizing the role of explicit representation and reasoning can result in a much more efficient system. [RN 197]
- The agents can be provided with a mechanism of **learning** new facts by themselves.
- The new knowledge incorporated in the KB can be used for decision making, making the agent f**ully autonomous.**

The wumpus world example [RN, 197]

- The **wumpus world** is a cave consisting of rooms connected by passageways.
- Lurking somewhere in the cave is the wumpus, a beast that eats anyone who enters its room.
- The wumpus can be shot by an agent, but the agent has only one ar,row.
- Some rooms contain bottomless pits that will trap anyone who wanders into these rooins (except for the wumpus, which is too big to fall in).
- The only mitigating feature of living in this environment is the possibility of finding a heap of gold.
- This game is an excellent test bed for intelligent agents.

The wumpus world example (II) [R, 7/5]

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

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

Sensors Breeze, Glitter, Smell

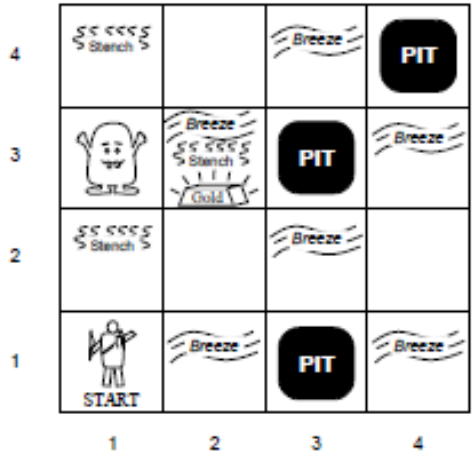

Wumpus world characterization [R, 7/6]

- Observable?? No – only local perception
- Deterministic?? Yes – outcomes exactly specied
- 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

Exploring a wumpus world [R, 7/16]

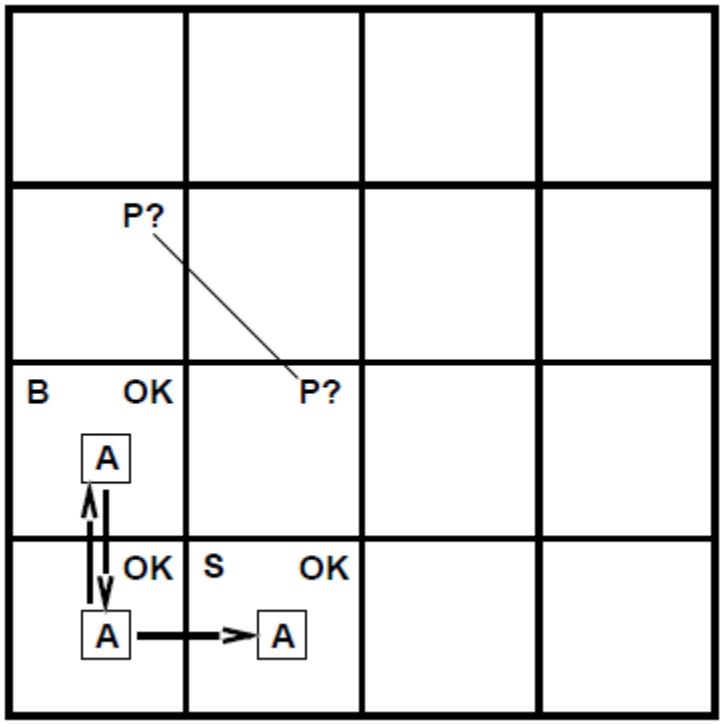

Exploring a wumpus world [R, 7/17]

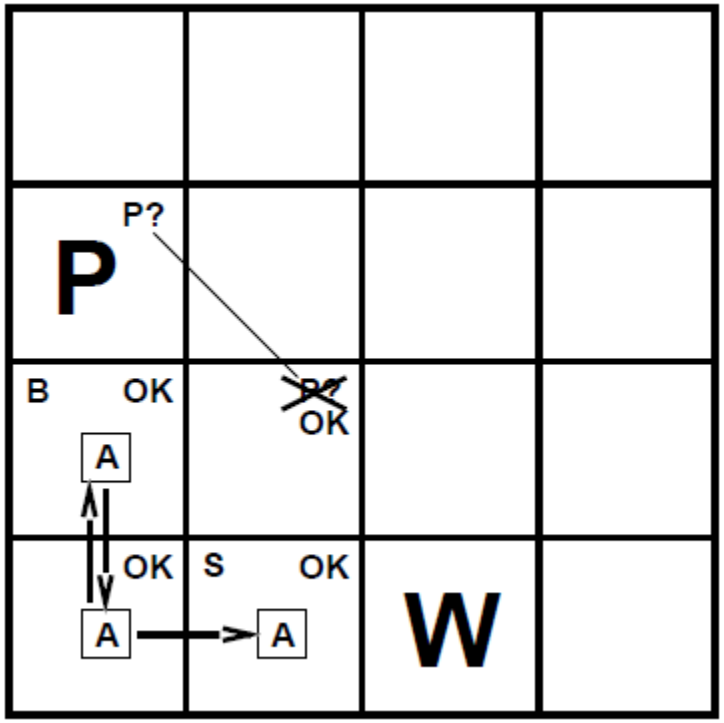

Exploring a wumpus world [R, 7/20]

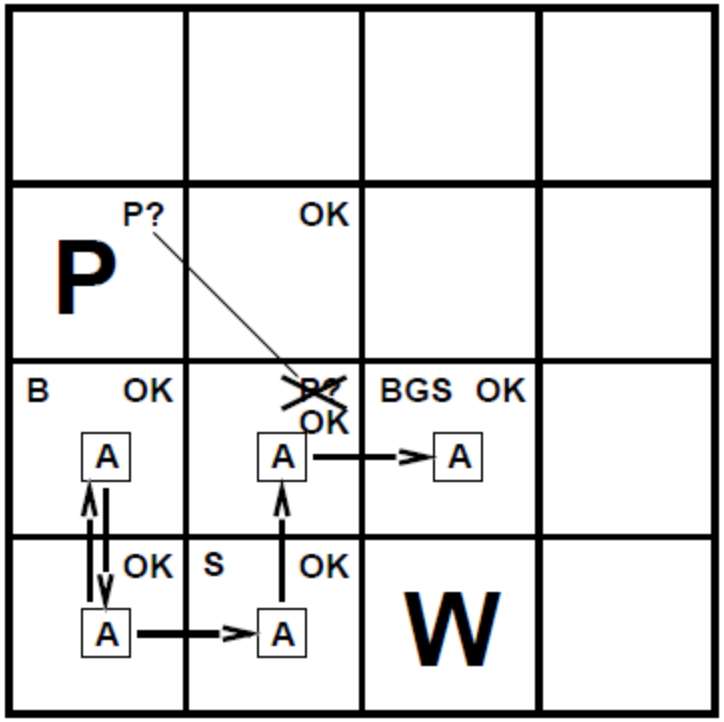

Reasoning in the Wumpus world [R, 7/21]

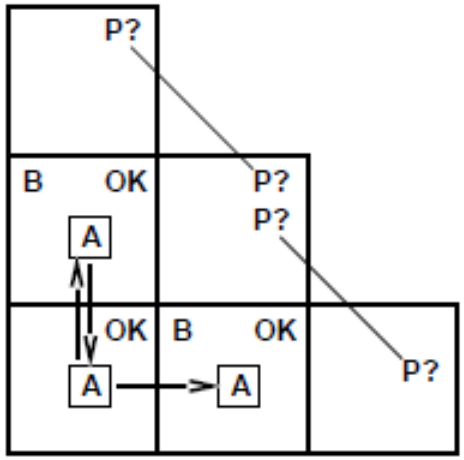

Breeze in (1,2) and (2,1)
⇒ no safe actions

Assuming pits uniformly distributed,
(2,2) has pit w/ prob 0.86, vs. 0.31

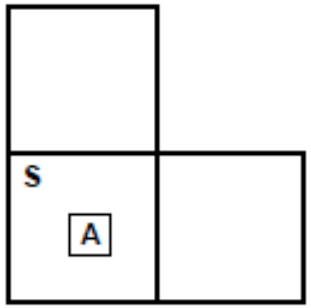

Smell in (1,1)
⇒ cannot move
Can use a strategy of coercion:
shoot straight ahead
wumpus was there ⇒ dead ⇒ safe
wumpus wasn't there ⇒ safe

- The lecture continues (until its final form is completed) with the slides from [R ch. 7] starting with slide 23

Bibliography

- [RN] Russel S., Norvig P. – Artificial Intelligence – A Modern Approach, 2nd ed. Prentice Hall, 2003 (1112 pages)
- [R] Stuart Russel – Course slides (visited oct. 2012 at http://aima.cs.berkeley.edu/instructors.html#homework)
- [W1] Mark Watson – Practical Artificial Intelligence Programming With Java AI 3rd ed., 2008
- [C] D. Cârstoiu – Sisteme Expert, Editura ALL, București, 1994