

AGENT ARCHITECTURES

BE4M36MAS - Multiagent systems

ORGANIZATION

Branislav Bošanský (branislav.bosansky@aic.fel.cvut.cz)

Karel Horák (karel.horak@aic.fel.cvut.cz)

Michal Šustr (michal.sustr@aic.fel.cvut.cz)

Petr Tomášek (petr.tomasek@aic.fel.cvut.cz)

Website:

<https://cw.fel.cvut.cz/wiki/courses/be4m36mas/start>

Attendance: **voluntary** (but tracked)

Assessment – 3 assignments:

1. Agent programming (max 11 pts)
2. Game theory (max 17 pts)
3. Distributed constraint satisfaction (12 pts)

AGENT ARCHITECTURES

Actions (A)

Ways for the agent to influence the environment

Percepts (P)

Observations about the state of the world

Decision making ($d : P^* \rightarrow A$)

Mapping perception history to actions

1. Reflex (reactive) Agent
2. Model-based Reflex Agent
3. Model-based Goal-based Agent
4. Model-based Utility-based Agent
5. Learning-based Agent

(Russell and Norvig)


Wumpus' World

- Grid world environment
- Agent has to find the gold brick and carry it to the bottom left square
- Problem: Entering a square occupied by Wumpus or containing a pit costs agent his life
(Wumpus does not move)




Wumpus' World — Percepts

- *Breeze* — whenever agent stands next to a pit
- *Stench* — whenever agent stands next to Wumpus
- *Gold* — when agent carries a gold brick



Wumpus' World — Actions

- Going to any neighboring square (only vertically and horizontally)




Agent conditions his decision solely on his **current** percepts.
(e.g. on the facts he can currently sense)

Task: Implement a reflex agent for Wumpus world. Beware, do not use any kind of memory or smarter reasoning ;-)

MODEL-BASED REFLEX AGENT

Agent uses percepts to gradually build a **model** of the environment.

Decisions are based on the expected state of the world according to his model.




Question: Does this approach allow us to overcome this issue?

MODEL-BASED REFLEX AGENT

Agent uses percepts to gradually build a **model** of the environment.

Decisions are based on the expected state of the world according to his model.



Question: Does this approach allow us to overcome this issue?

Task: Implement a model-based agent and reach the gold!

Question: Is the behaviour of the agent rational?

Question: Is the behaviour of the agent rational?

Definitely not!

Agent just exploits the model to stay alive. He does not intentionally pursue his goal.

Actions are chosen in order to reach a **declaratively** specified goal.

Techniques:

1. Planning

Planning in AI

2. Belief-Desire-Intention Architecture

this course

Question: What does it mean for an agent in Wumpus' world?

Not all ways to reach the goal are equally plausible. Some ways to reach the goal **should be preferred** against others. (e.g. cheaper or less risky ones)

Utility driven sequential decision making:

- Non-adversarial: MDPs, POMDPs
- Adversarial: Sequential games

Planning in AI
this course

Agent **does not fully know** the task he is facing.
(what his action does, what is his goal etc.)

He **learns** the task on the go — strategy reflecting these finds cannot be fixed in advance.

Learning **both** model and strategy.

NEXT TUTORIAL

- Belief-Desire-Intention architecture