

Agent architectures

BE4M36MAS - Multiagent systems

Organization

José Ananías Hilario Reyes (jose.reyes@aic.fel.cvut.cz)

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

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

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

Website:

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

Agent programming

24 Sep	Introduction to multi-agent systems, Belief-Desire-Intention architecture	Pěchouček
01 Oct	Belief-Desire-Intention architecture	Jakob

Non-cooperative game theory

08 Oct	Introduction to Game Theory	Bošanský
15 Oct	Solving Normal-form Games	Bošanský
22 Oct	Games in Extensive Form	Bošanský
29 Oct	Solving Extensive-Form Games	Bošanský
05 Nov	Other Game Representations	Bošanský

Distributed constraint reasoning

12 Nov	Distributed constraint reasoning 1 (DCSP)	Božanský
19 Nov	Distributed constraint reasoning 2 (DCOP)	Božanský

Cooperative game theory

26 Nov	Cooperative Game Theory	Kroupa
03 Dec	Cooperative Game Theory 2	Kroupa

Other MAS topics

10 Dec	Social Choice, Voting	Jakob
17 Dec	Resource allocation, Auctions	Jakob
07 Jan	Multiagent Simulations	Jakob

Attendance: **voluntary** (but tracked)

Assessment – 3 assignments:

1. Agent programming (max 11 pts)
2. Competitive game theory (max 17 pts)
3. Cooperative game theory (max 12 pts)

Plagiarism is strictly forbidden

(Strong punishments would be applied)

Agent architectures

Components of agent architectures

Components of 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

Architecture types


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

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

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

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?

Model-based reflex agent

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.

Model-based Goal-based agent

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?

Model-based Utility-based agent

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

Learning-based agent

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