Intelligent agents

Artificial intelligence

Kristóf Karacs
PPKE-ITK

Recap

- What is intelligence?
- What can we use it for?
- How does it work? How to create it?
- How to control / repair / improve it?
- What are the consequences?
- Do we need to be afraid of it?
- What can we do?

Do we need to be afraid of it?

- You may, but it is better to take action

  learn – know – act

Reminders

- Project: February 27, next Wednesday
- Quick presentations
- Worked out problems

Program

- Problem solving by search
- Adversarial search
- Logic and inference
- Search in logic representation, planning
- Inference in case of constraints
- Bayesian networks
- Fuzzy logic
- Machine learning

Outline

- Agents and environments
- Rationality
- PEAS: performance measure, environment, actuators, sensors
- Models of agents
- Aspects of environments

Intelligent agents

- An **agent** is anything that can be viewed as
  - *perceiving* its environment through *sensors* and
  - *acting* upon the environment through *actuators*.

How do agents work?

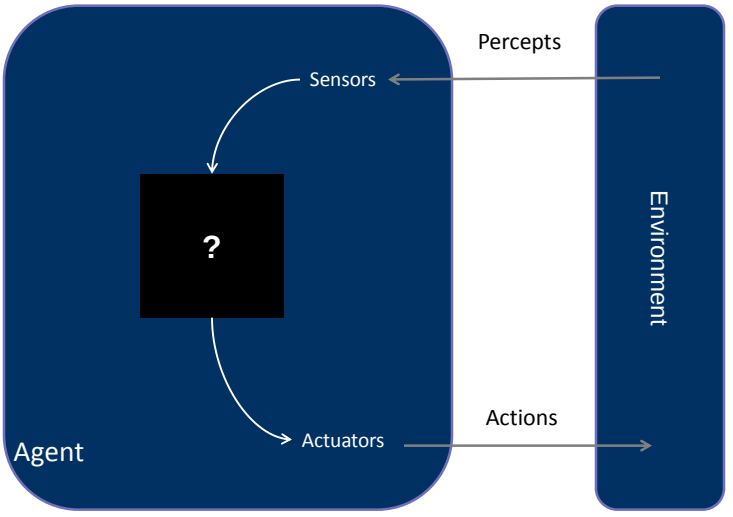

Type of agents

- Human

- Robot

- Software

Rational agent

- A **rational agent** is one that does the right thing
- Assessing the agent's performance
  - Performance measure: objectively tells how successful the agent is

Evaluation of rationality

- Goals and a performance measure
- Prior knowledge about the environment
- Abilities: possible primitive actions
- History
  - Percept sequence
  - Past experiences (data to learn from)

Internal Structure

- Agent = Architecture + Program
  HW, bg. SW + actual algorithm
- Knowledge of Environment
  - Source
    - Given a-priori
    - Learned from sensory input
  - May include
    - Present / past states of environment
    - Influence of actions on the environment

Complexity levels

- Reflex agents
  - Lookup table: if-then rules
  - Problems: size, time, flexibility
- Model-based reflex agents
  - Internal state
- Goal-based agents
  - Search and planning
- Utility-based agents
  - Non-binary measure

Reflexes

- Action depends only on sensory input
- Background knowledge not used
- Humans – flinching, blinking
- Chess – openings, endings
  - Lookup table (not a good idea in general)
  - 35^{100} entries required for the entire game

Reflex Agents

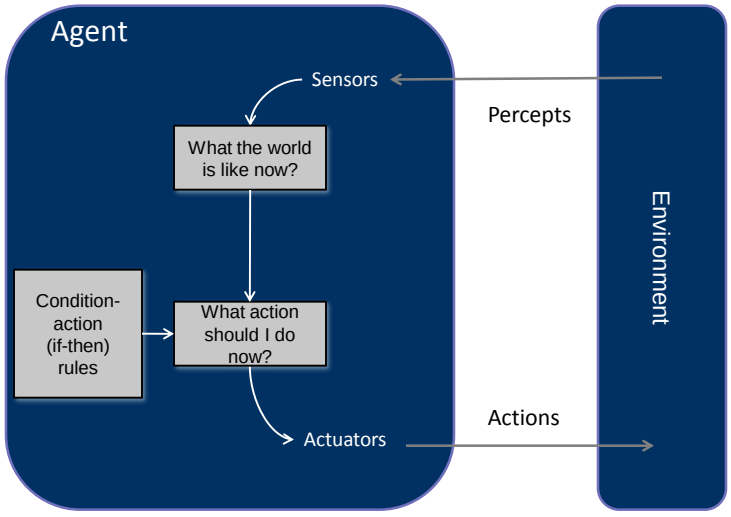

Model-based Reflex Agents

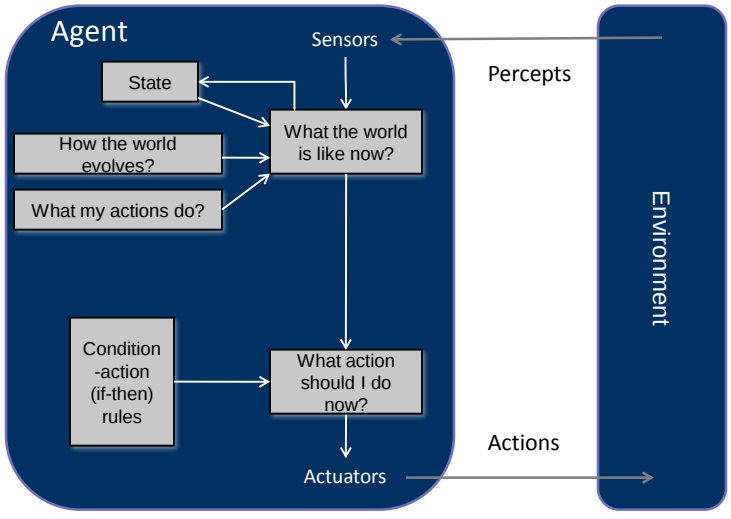

Goal of an agent

- Environment in itself is often not enough to decide what to do
- Goal is described by some properties
- A goal based agent
  - uses knowledge about a goal to guide its actions (search and planning)
  - compares the results of possible actions
- Principle: The action taken should modify the environment towards the goal

Goal-based Agents

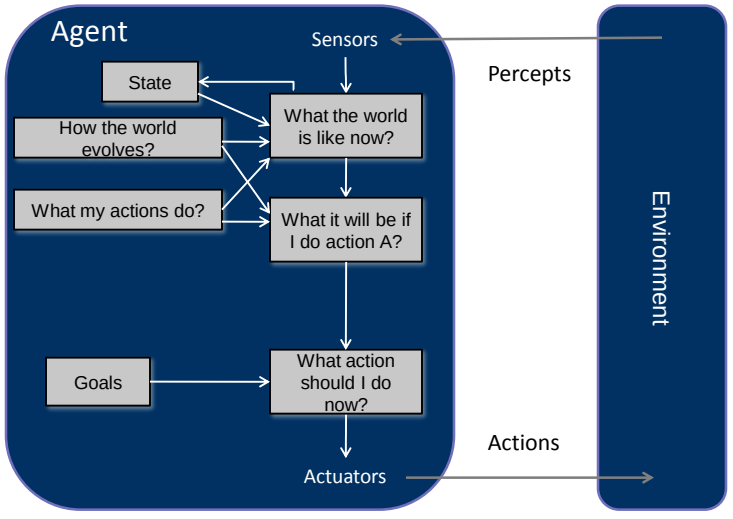

Search & Planning

Utility Functions

- Knowledge of a goal may be difficult to pin down (e.g. checkmate in chess)
- Agent may have multiple, controversial goals
- Comparing utility of states
  - Utility functions measure value of world states
  - Localized measures
- Choose action which best improves utility (Best First Search)

Utility-based Agents

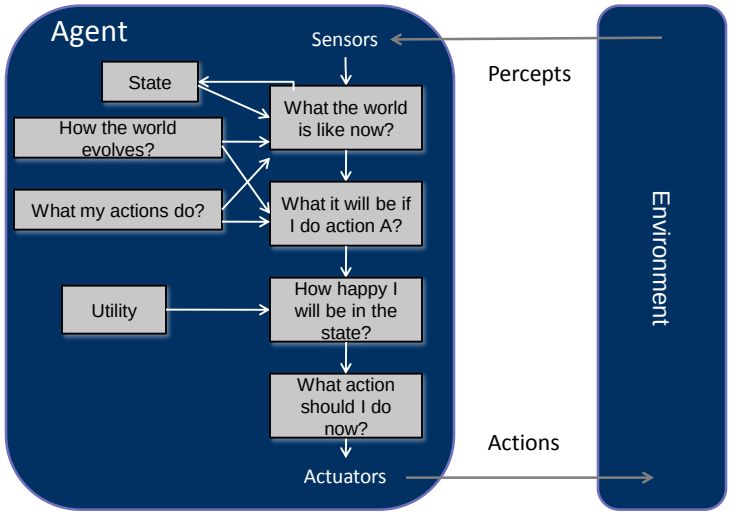

utility function: X (state space) $\rightarrow \Re$

Other aspects

- Hybrid agents
  - Hierarchical architecture
  - Trade-off between efficiency and flexibility
- Capability of learning
- Multi-agent systems
  - Competitive vs. cooperative relationship

Hierarchical control

- Delivery robot

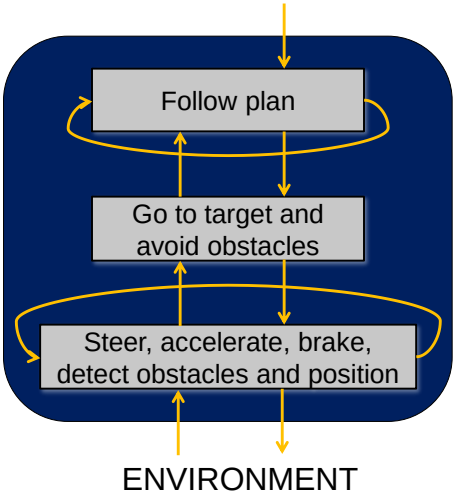

Learning Agents

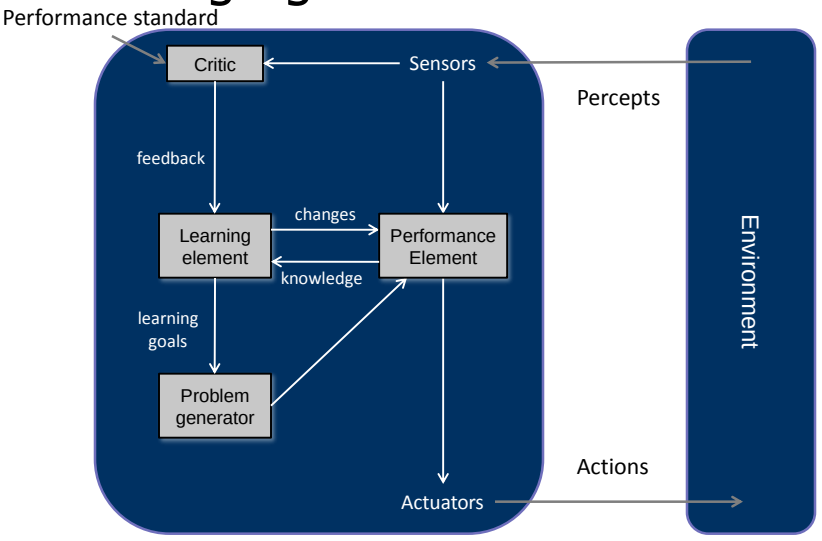

Autonomy of Agents

- Autonomy = extent to which the agent's behaviour is determined by its own experience
- Extreme cases
  - No autonomy – ignores input (environment)
  - Complete autonomy – acts randomly/no program
- Ideal agent: some autonomy
  - Gradually increasing over time
  - Example: baby learning to crawl & navigate

Details of the Environment

- Properties of the world are different (real-world robot vs. software agent)

Uncertainty

Fully observable	vs.	Partially observable
Deterministic	vs.	Stochastic
Episodic	vs.	Sequential
Static	vs.	Dynamic
Discrete	vs.	Continuous
Single agent	vs.	Multiple agents

Observability (sensing uncertainty)

- An environment is *fully observable*, if the agent can access every information in its environment it takes into account when choosing an action
- *Partially observable* if parts of the environment are not observable
- Unobservable information must be guessed → the agent needs a model
- Example: chess (fully) vs. poker (partially)

Determinism (effect uncertainty)

- An environment is *deterministic* if a change in the world state depends only on
  - current state of the world
  - agent's action
- *Non-deterministic* environments
  - have aspects beyond the control of the agent
  - non-observable can seem to be non-det.
  - can be treated as stochastic or probabilistic
- Example: chess (det.) vs. poker (non-det.)

Episodicity

- An environment is *episodic* if the choice of current action does not depend on previous actions
- In *sequential* environments
  - Agent has to plan ahead
  - Current choice affects future actions
- Example: mail sorting system (episodic) vs. poker, chess (sequential)

Time variance

- *Static* environments don't change over time
- *Dynamic* environments: changes have to be taken into account by either
  - sensing the change
  - predicting the change
  - neglecting the change (in the short run)
- Example: poker, chess (static) vs. taxi driving (dynamic)

Continuity

- Type of sensor data and choices of action
- *Discrete*: distinct, clearly defined set
- *Continuous*: non-sectionable
- Example: chess (discrete) vs. taxi driving (continuous)

Number of agents

- *Single agent*: the environment is not changed by other actors
- *Multi-agent*: the agent is aware of other agents, who also modifying the environment
  - Modelling question: multi-agent vs. stochastic single agent
- Examples: solitaire (single) vs. poker (multi-)

Summary

- Agent: defined in connection with the environment
  - Perceives and acts
- Rationality
- Basic mode: reflex, model, goal, and utility based
  - Hierarchical control
- Learning
- Environments: observable?, deterministic?, static?, episodic?, continuous?, multi-agent?