

Artificial Intelligence

2. Intelligent Agents

On the Structure of Intelligent Agents and their Environments

Jörg Hoffmann Wolfgang Wahlster



Summer Term 2012

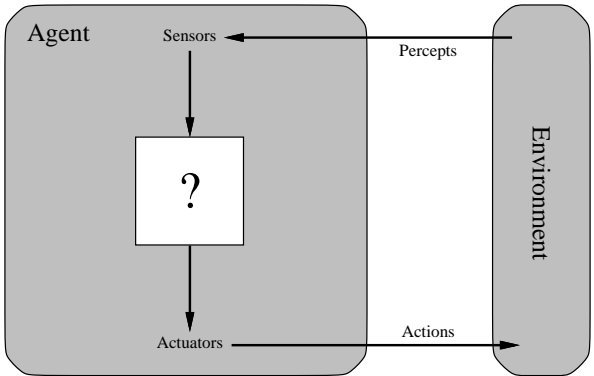
Thanks to Profs. Burgard, Nebel, Riedmiller (Univ. Freiburg) for slide sources

Agenda

- 1 What is an Agent?
- 2 What is a Rational Agent?
- 3 Different Classes of Agents
- 4 Different Types of Environments

Agents

- Perceive the environment through sensors (→ Percepts)
- Act upon the environment through actuators (→ Actions)



Examples: Humans and animals, robots and software agents (softbots), temperature control, ABS, ...

Rational Agents

... do the "right thing"!

In order to evaluate their performance, we have to define a **performance measure**. Autonomous vacuum cleaner example:

- m² per hour
- Level of cleanliness
- Energy usage
- Noise level
- Safety (behavior towards hamsters/small children)

Optimal behavior is often unattainable!

- Not all relevant information is perceivable
- Complexity of the problem is too high

Rationality vs. Omniscience

- An **omniscient agent** knows the **actual effects** of its **actions**.
- In comparison, a **rational agent** behaves according to its **percepts** and **knowledge** and attempts to **maximize the expected performance**.
- Example?

The Ideal Rational Agent

Rational behavior depends on:

- Performance measures (goals);
- Percept sequences;
- Knowledge of the environment;
- Possible actions.

Ideal rational agent

For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

$$\text{Percept Sequence} \times \text{World Knowledge} \rightarrow \text{Action}$$

Active perception is necessary to avoid trivialization!

Examples of Agents: PEAS Descriptions

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	healthy patient, costs, lawsuits	patient, hospital, staff	display questions, tests, diagnoses, treatments, referrals	keyboard entry of symptoms, findings, patient's answers
Satellite image analysis system	correct image categorization	downlink from orbiting satellite	display categorization of scene	color pixel arrays
Part-picking robot	percentage of parts in correct bins	conveyor belt with parts, bins	jointed arm and hand	camera, joint angle sensors
Refinery controller	purity, yield, safety	refinery, operators	valves pumps, heaters displays	temperature, pressure, chemical sensors
Interactive English tutor	student's score on test	set of students, testing agency	display exercises, suggestions, corrections	keyboard entry

Structure of Rational Agents

Realization of the ideal mapping through an

- **Agent program**, executed on an
- **Architecture** which also provides an interface to the environment (percepts, actions).

$$\rightarrow \text{Agent} = \text{Architecture} + \text{Program}$$

Question!

Which are Agents?

- (A): James Bond.
- (B): Vacuum cleaner.
- (C): Your dog.
- (D): Thermometer.

Question!

Who is rational?

- (A): James Bond, crossing the street after looking.
- (B): Vacuum cleaner, deciding to clean under your bed.
- (C): Your dog, crossing the street after not looking.
- (D): Thermometer, not doing anything.


```

function TABLE-DRIVEN-AGENT(percept) returns an action
  persistent: percepts, a sequence, initially empty
               table, a table of actions, indexed by percept sequences, initially fully specified

  append percept to the end of percepts
  action ← LOOKUP(percepts, table)
  return action
  
```

Problems:

- The table can become very large.
- It usually takes a very long time for the designer to specify it.
- ... **practically impossible!**



Direct use of perceptions is often not possible due to the large space required to store them (e.g., video images).

Input therefore is often interpreted before decisions are made.

Interpretative Reflex Agents

Since storage space required for actual perceptions is too large, **interpret** perceptions:

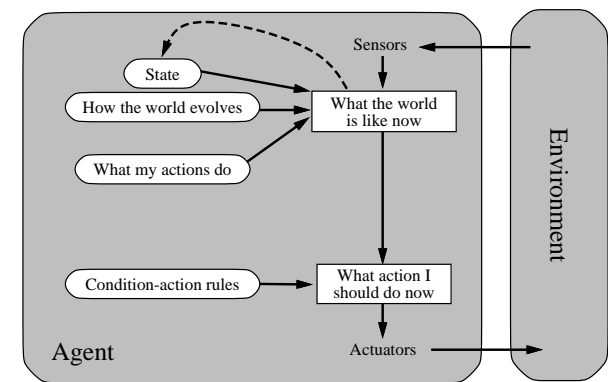
```

function SIMPLE-REFLEX-AGENT(percept) returns an action
  persistent: rules, a set of condition–action rules

  state ← INTERPRET-INPUT(percept)
  rule ← RULE-MATCH(state, rules)
  action ← rule.ACTION
  return action
    
```

Structure of Model-based Reflex Agents

In case the agent's history in addition to the actual percept is required to decide on the next action, it must be represented in a suitable form:



A Model-based Reflex Agent

```

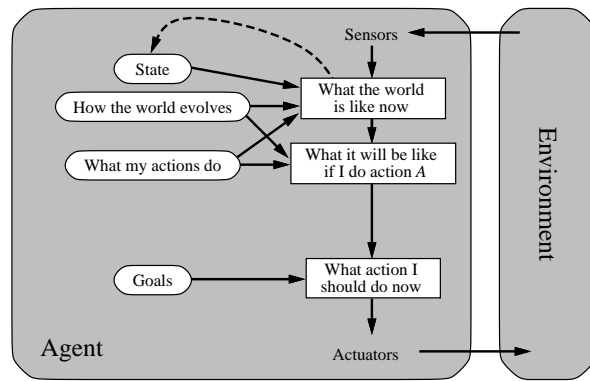
function MODEL-BASED-REFLEX-AGENT(percept) returns an action
  persistent: state, the agent's current conception of the world state
                 model, a description of how the next state depends on current state and action
                 rules, a set of condition–action rules
                 action, the most recent action, initially none

  state ← UPDATE-STATE(state, action, percept, model)
  rule ← RULE-MATCH(state, rules)
  action ← rule.ACTION
  return action
    
```

Model-based, Goal-based Agents

- Often, **percepts alone are insufficient** to decide what to do.
- This is because the correct action depends on the given **explicit goals** (e.g., go towards X).
- The **model-based, goal-based agents** use an explicit representation of goals and consider them for the choice of actions.

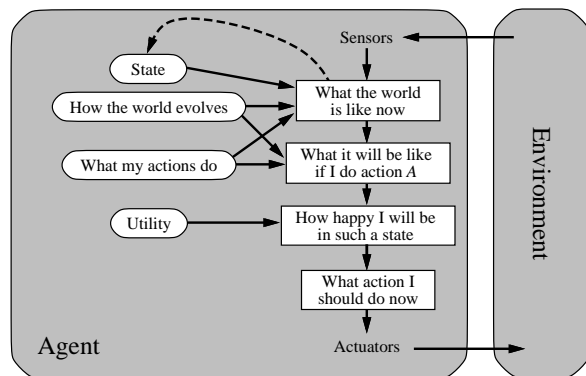
Model-based, Goal-based Agents



Model-based, Utility-based Agents

- Usually, there are **several possible actions** that can be taken in a given situation.
- In such cases, the **utility of the next achieved state** can come into consideration to arrive at a decision.
- A **utility function** maps a state (or a sequence of states) onto a real number.
- The agent can also use these numbers to **weigh the importance of competing goals**.

Model-based, Utility-based Agents



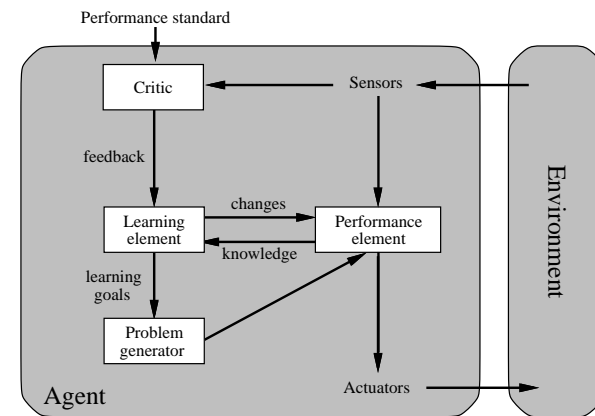
Learning Agents

- Learning agents can become **more competent over time**.
- They can start with an **initially empty knowledge base**.
- They can operate in initially **unknown environments**.

Components of Learning Agents

- Learning element (responsible for making improvements).
- Performance element (has to select external actions).
- Critic (determines the performance of the agent).
- Problem generator (suggests actions that will lead to informative experiences).

Learning Agents



Questionnaire

Question!
What kind of agent are you?

(A): Table-Driven.	(B): Reflex Agent.
(C): Utility-Based.	(D): Learning.

Questionnaire, ctd.

Question!
When are reflexes not good enough?

(A): Your goals are not always the same.	(B): Your goals are rather complicated.
(C): The world takes time to interpret.	(D): Your capabilities should grow over time.

The Environment of Rational Agents

- **Accessible vs. inaccessible** (fully observable vs. partially observable)
Are the relevant aspects of the environment accessible to the sensors?
- **Deterministic vs. stochastic**
Is the next state of the environment completely determined by the current state and the selected action?
 If the only non-determinism are actions of other agents, the environment is called **strategic**.
- **Episodic vs. sequential**
Can the quality of an action be evaluated within an episode (perception + action), or are future developments decisive?

The Environment of Rational Agents, ctd.

- **Static vs. dynamic**
Can the environment change while the agent is deliberating?
 If the environment does not change, but the agent's performance score changes, the environment is called **semi-dynamic**.
- **Discrete vs. continuous**
Is the environment discrete (e.g., chess) or continuous (e.g., a robot moving in a room)?
- **Single agent vs. multi-agent**
Is there just one agent, or several of them?
 There are **competitive** and **cooperative** multi-agent scenarios.

Examples of Environments

Task	Observable	Deterministic	Episodic	Static	Discrete	Agents
Chess with a clock	fully	strategic	sequential	semi	discrete	multi
Poker	partially	stochastic	sequential	static	discrete	multi
Backgammon	fully	stochastic	sequential	static	discrete	multi
Taxi driving	partially	stochastic	sequential	dynamic	continuous	multi
Medical diagnosis	partially	stochastic	sequential	dynamic	continuous	single
Image analysis	fully	deterministic	episodic	semi	continuous	single
Part-picking robot	partially	stochastic	episodic	dynamic	continuous	single
Refinery controller	partially	stochastic	sequential	dynamic	continuous	single
Interactive English tutor	partially	stochastic	sequential	dynamic	discrete	multi

Questionnaire

Question!

James Bond's environment is?

(A): Fully Observable. (B): Strategic.
 (C): Episodic. (D): Single-Agent.

Questionnaire, ctd.

Question!

Your own environment is?

- (A): Fully Observable.
- (B): Strategic.
- (C): Episodic.
- (D): Single-Agent.

Summary

- An **agent** is something that perceives and acts. It consists of an architecture and an agent program.
- An **ideal rational agent** always takes the action that maximizes its performance given the percept sequence and its environment knowledge.
- There are a variety of agent designs:
 - **Reflex agents** respond immediately to percepts.
 - **Goal-based agents** work towards goals.
 - **Utility-based agents** try to maximize their reward.
 - **Learning agents** improve their behavior over time.
- Some **environments** are more demanding than others ...
- ... your own, and that of James Bond, are the most difficult.