Introduction to Artificial Intelligence

Intelligent Agents

Bernhard Beckert

UNIVERSITÄT KOBLENZ-LANDAU

Wintersemester 2003/2004

Outline

- Agents and environments
- PEAS (Performance, Environment, Actuators, Sensors)
- Environment types
- Agent types
- Example: Vacuum world

Agents and environments

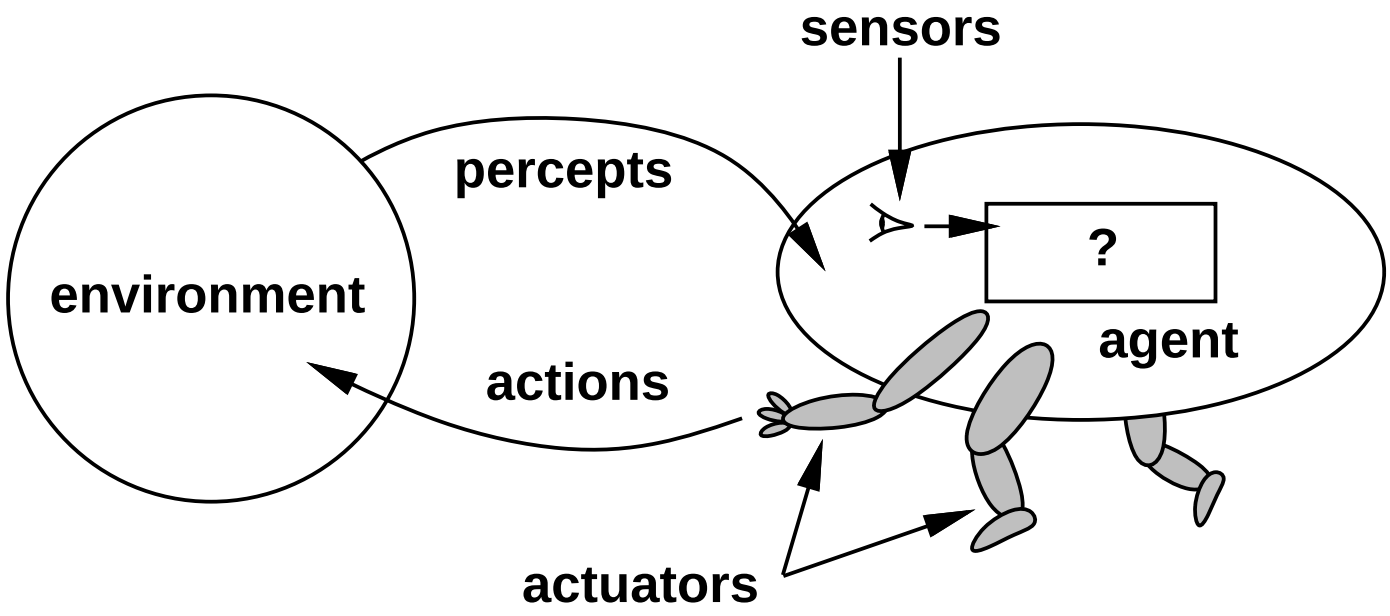

Agents include

- humans
- robots
- software robots (softbots)
- thermostats
- etc.

Agent functions and programs

Agent function

An agent is completely specified by the agent function

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

mapping percept sequences to actions

Agent functions and programs

Agent program

- **runs on the physical architecture to produce** f
- **takes a single percept as input**
- **keeps internal state**

```
function SKELETON-AGENT(percept) returns action
    static: memory     /* the agent's memory of the world */

    memory ← UPDATE-MEMORY(memory, percept)
    action ← CHOOSE-BEST-ACTION(memory)
    memory ← UPDATE-MEMORY(memory,action)
    return action
```

AIMA code

Available at

`http://aima.cs.berkeley.edu/code.html`

in different languages (Java, Lisp, . . .)

Code for each topic divided into four directories

`agents`:	code defining agent types and programs
`algorithms`:	code for the methods used by the agent programs
`environments`:	code defining environment types, simulations
`domains`:	problem types and instances for input to algorithms

For experiments

Often algorithms on domains rather than agents in environments

Rationality

Goal

Specified by performance measure,
defining a numerical value for any environment history

Rational action

Whichever action maximizes the expected value
of the performance measure given the percept sequence to date

Rationality

Goal

Specified by performance measure,
defining a numerical value for any environment history

Rational action

Whichever action maximizes the expected value
of the performance measure given the percept sequence to date

Note

$$\text{rational} \neq \text{omniscient}$$
$$\text{rational} \neq \text{clairvoyant}$$
$$\text{rational} \neq \text{successful}$$

Rationality

Goal

Specified by performance measure,
defining a numerical value for any environment history

Rational action

Whichever action maximizes the expected value
of the performance measure given the percept sequence to date

Note

$$\text{rational} \neq \text{omniscient}$$
$$\text{rational} \neq \text{clairvoyant}$$
$$\text{rational} \neq \text{successful}$$

Agents need to: gather information, explore, learn, ...

PEAS: The setting for intelligent agent design

Example: Designing an automated taxi

Performance

Environment

Actuators

Sensors

PEAS: The setting for intelligent agent design

Example: Designing an automated taxi

Performance	**safety, reach destination, maximize profits, obey laws, passenger comfort, . . .**
Environment	
Actuators	
Sensors	

PEAS: The setting for intelligent agent design

Example: Designing an automated taxi

Performance	**safety, reach destination, maximize profits, obey laws, passenger comfort, . . .**
Environment	**streets, traffic, pedestrians, weather, customers, . . .**
Actuators	
Sensors	

PEAS: The setting for intelligent agent design

Example: Designing an automated taxi

Performance	**safety, reach destination, maximize profits, obey laws, passenger comfort, . . .**
Environment	**streets, traffic, pedestrians, weather, customers, . . .**
Actuators	**steer, accelerate, brake, horn, speak/display, . . .**
Sensors	

PEAS: The setting for intelligent agent design

Example: Designing an automated taxi

Performance	**safety, reach destination, maximize profits, obey laws, passenger comfort, . . .**
Environment	**streets, traffic, pedestrians, weather, customers, . . .**
Actuators	**steer, accelerate, brake, horn, speak/display, . . .**
Sensors	**video, accelerometers, gauges, engine sensors, keyboard, GPS, . . .**

PEAS: The setting for intelligent agent design

Example: Medical diagnosis system

Performance

Environment

Actuators

Sensors

PEAS: The setting for intelligent agent design

Example: Medical diagnosis system

Performance Healthy patient, minimize costs, avoid lawsuits, . . .

Environment

Actuators

Sensors

PEAS: The setting for intelligent agent design

Example: Medical diagnosis system

Performance Healthy patient, minimize costs, avoid lawsuits, . . .

Environment patient, hospital, staff, . . .

Actuators

Sensors

PEAS: The setting for intelligent agent design

Example: Medical diagnosis system

Performance	**Healthy patient, minimize costs, avoid lawsuits, . . .**
Environment	**patient, hospital, staff, . . .**
Actuators	**questions, tests, diagnoses, treatments, referrals, . . .**
Sensors	

PEAS: The setting for intelligent agent design

Example: Medical diagnosis system

Performance	**Healthy patient, minimize costs, avoid lawsuits, . . .**
Environment	**patient, hospital, staff, . . .**
Actuators	**questions, tests, diagnoses, treatments, referrals, . . .**
Sensors	**keyboard (symptoms, test results, answers), . . .**

Environment types

Fully observable **(otherwise: partially observable)**

Agent's sensors give it access to the complete state of the environment at each point in time

Environment types

Fully observable (otherwise: **partially observable**)

Agent's sensors give it access to the complete state of the environment at each point in time

Deterministic (otherwise: **stochastic**)

The next state of the environment is completely determined by the current state and the action executed by the agent
(strategic: deterministic except for behavior of other agents)

Environment types

Fully observable (otherwise: **partially observable**)

Agent's sensors give it access to the complete state of the environment at each point in time

Deterministic (otherwise: **stochastic**)

The next state of the environment is completely determined by the current state and the action executed by the agent
(strategic: deterministic except for behavior of other agents)

Episodic (otherwise: **sequential**)

The agent's experience is divided atomic, independent episodes
(in each episode the agent perceives and then performs a single action)

Environment types

Static **(otherwise: dynamic)**

Environment can change while the agent is deliberating

(semidynamic: not the state but the performance measure can change)

Environment types

Static (otherwise: **dynamic**)

Environment can change while the agent is deliberating
(semidynamic: not the state but the performance measure can change)

Discrete (otherwise: **continuous**)

The environment's state, time, and the agent's percepts and actions have discrete values

Environment types

Static (otherwise: **dynamic**)

Environment can change while the agent is deliberating
(semidynamic: not the state but the performance measure can change)

Discrete (otherwise: **continuous**)

The environment's state, time, and the agent's percepts and actions have discrete values

Single agent (otherwise: **multi-agent**)

Only one agent acts in the environment

Environment types

	Crossword puzzle	Chess	Back-gammon	Internet shopping	Taxi	Part-picking robot
Fully observable						
Deterministic						
Episodic						
Static						
Discrete						
Single agent						

Environment types

	Crossword puzzle	Chess	Back-gammon	Internet shopping	Taxi	Part-picking robot
Fully observable	yes					
Deterministic	yes					
Episodic	no					
Static	yes					
Discrete	yes					
Single agent	yes					

Environment types

	Crossword puzzle	Chess	Back-gammon	Internet shopping	Taxi	Part-picking robot
Fully observable	yes	yes				
Deterministic	yes	strat.				
Episodic	no	no				
Static	yes	semi				
Discrete	yes	yes				
Single agent	yes	no				

Environment types

	Crossword puzzle	Chess	Back-gammon	Internet shopping	Taxi	Part-picking robot
Fully observable	yes	yes	yes			
Deterministic	yes	strat.	no			
Episodic	no	no	no			
Static	yes	semi	yes			
Discrete	yes	yes	yes			
Single agent	yes	no	no			

Environment types

	Crossword puzzle	Chess	Back-gammon	Internet shopping	Taxi	Part-picking robot
Fully observable	yes	yes	yes	no		
Deterministic	yes	strat.	no	(yes)		
Episodic	no	no	no	no		
Static	yes	semi	yes	semi		
Discrete	yes	yes	yes	yes		
Single agent	yes	no	no	no		

Environment types

	Crossword puzzle	Chess	Back-gammon	Internet shopping	Taxi	Part-picking robot
Fully observable	yes	yes	yes	no	no	
Deterministic	yes	strat.	no	(yes)	no	
Episodic	no	no	no	no	no	
Static	yes	semi	yes	semi	no	
Discrete	yes	yes	yes	yes	no	
Single agent	yes	no	no	no	no	

Environment types

	Crossword puzzle	Chess	Back-gammon	Internet shopping	Taxi	Part-picking robot
Fully observable	yes	yes	yes	no	no	no
Deterministic	yes	strat.	no	(yes)	no	no
Episodic	no	no	no	no	no	yes
Static	yes	semi	yes	semi	no	no
Discrete	yes	yes	yes	yes	no	no
Single agent	yes	no	no	no	no	yes

Environment types

The real world is

- **partially observable**
- **stochastic**
- **sequential**
- **dynamic**
- **continuous**
- **multi-agent**

Agent types

Four basic types

(in order of increasing generality)

- **simple reflex agents**
- **reflex agents with state**
- **goal-based agents**
- **utility-based agents**

All these can be turned into learning agents

Simple reflex agents

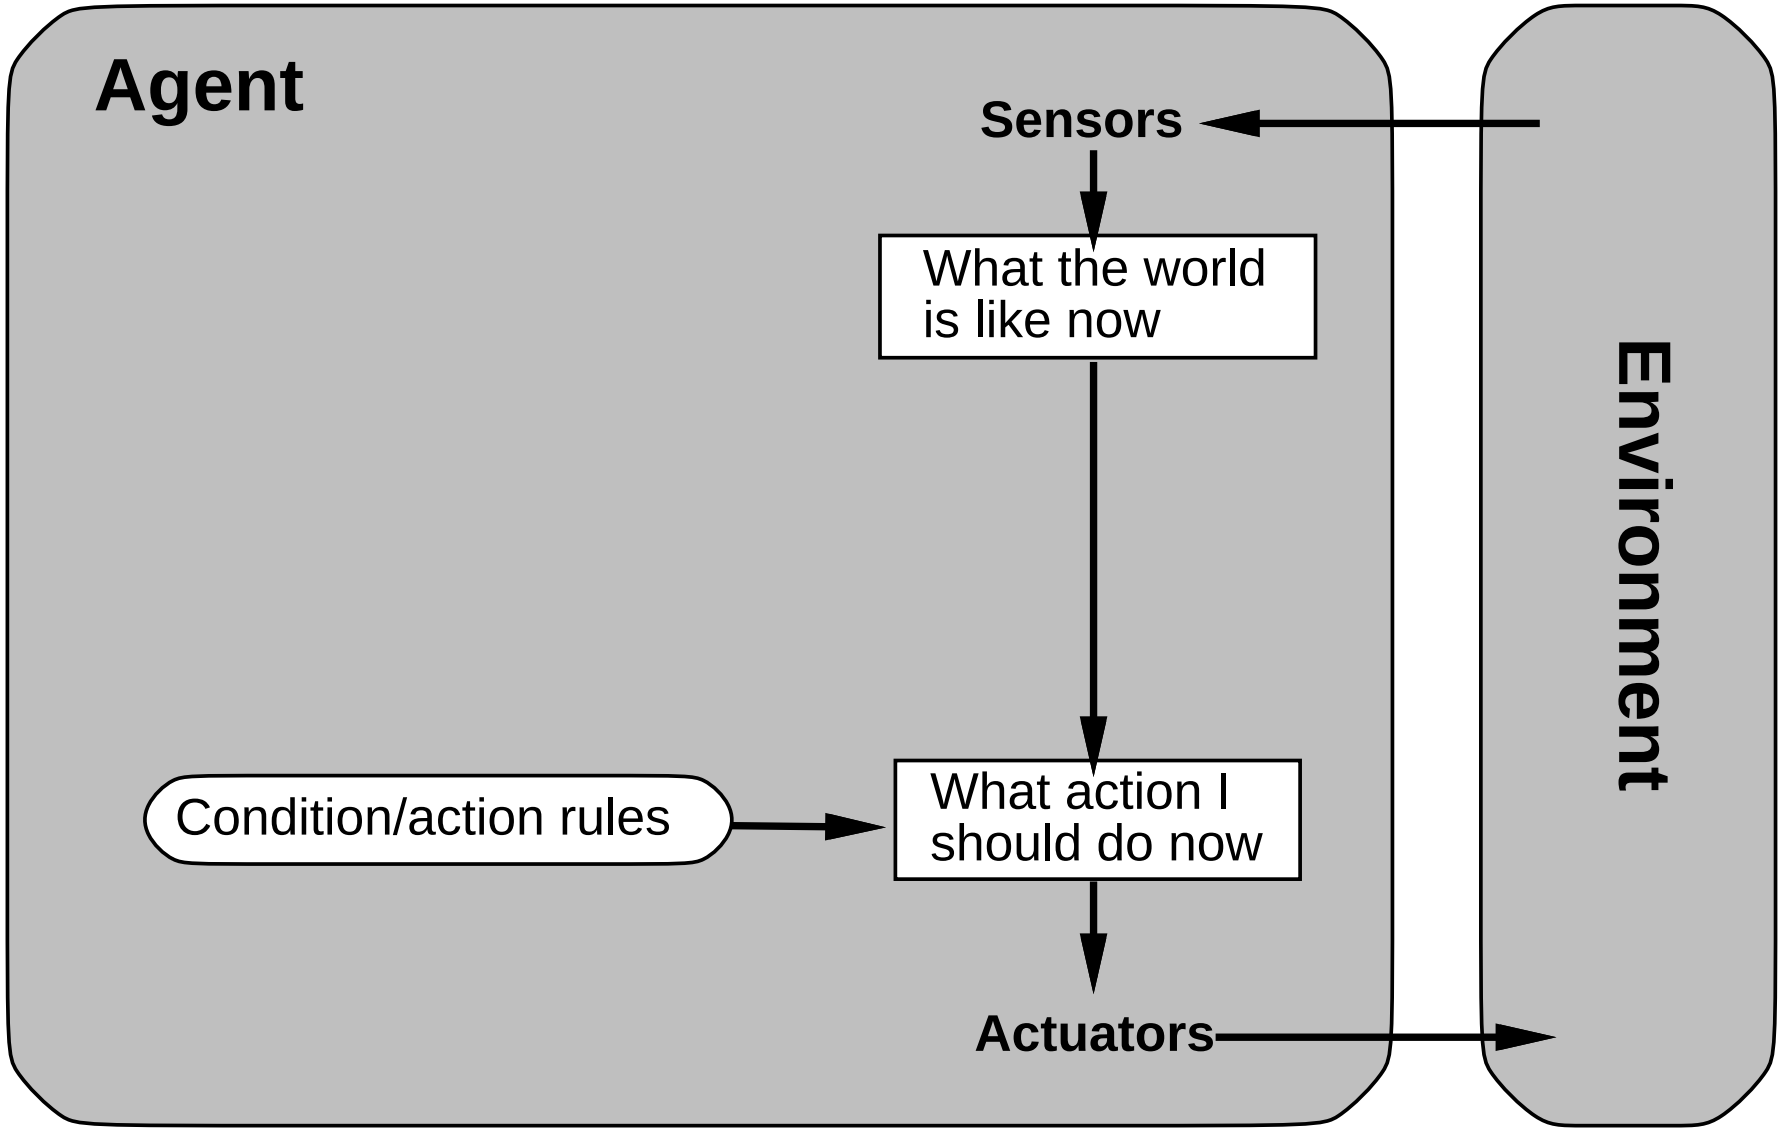

Model-based reflex agents

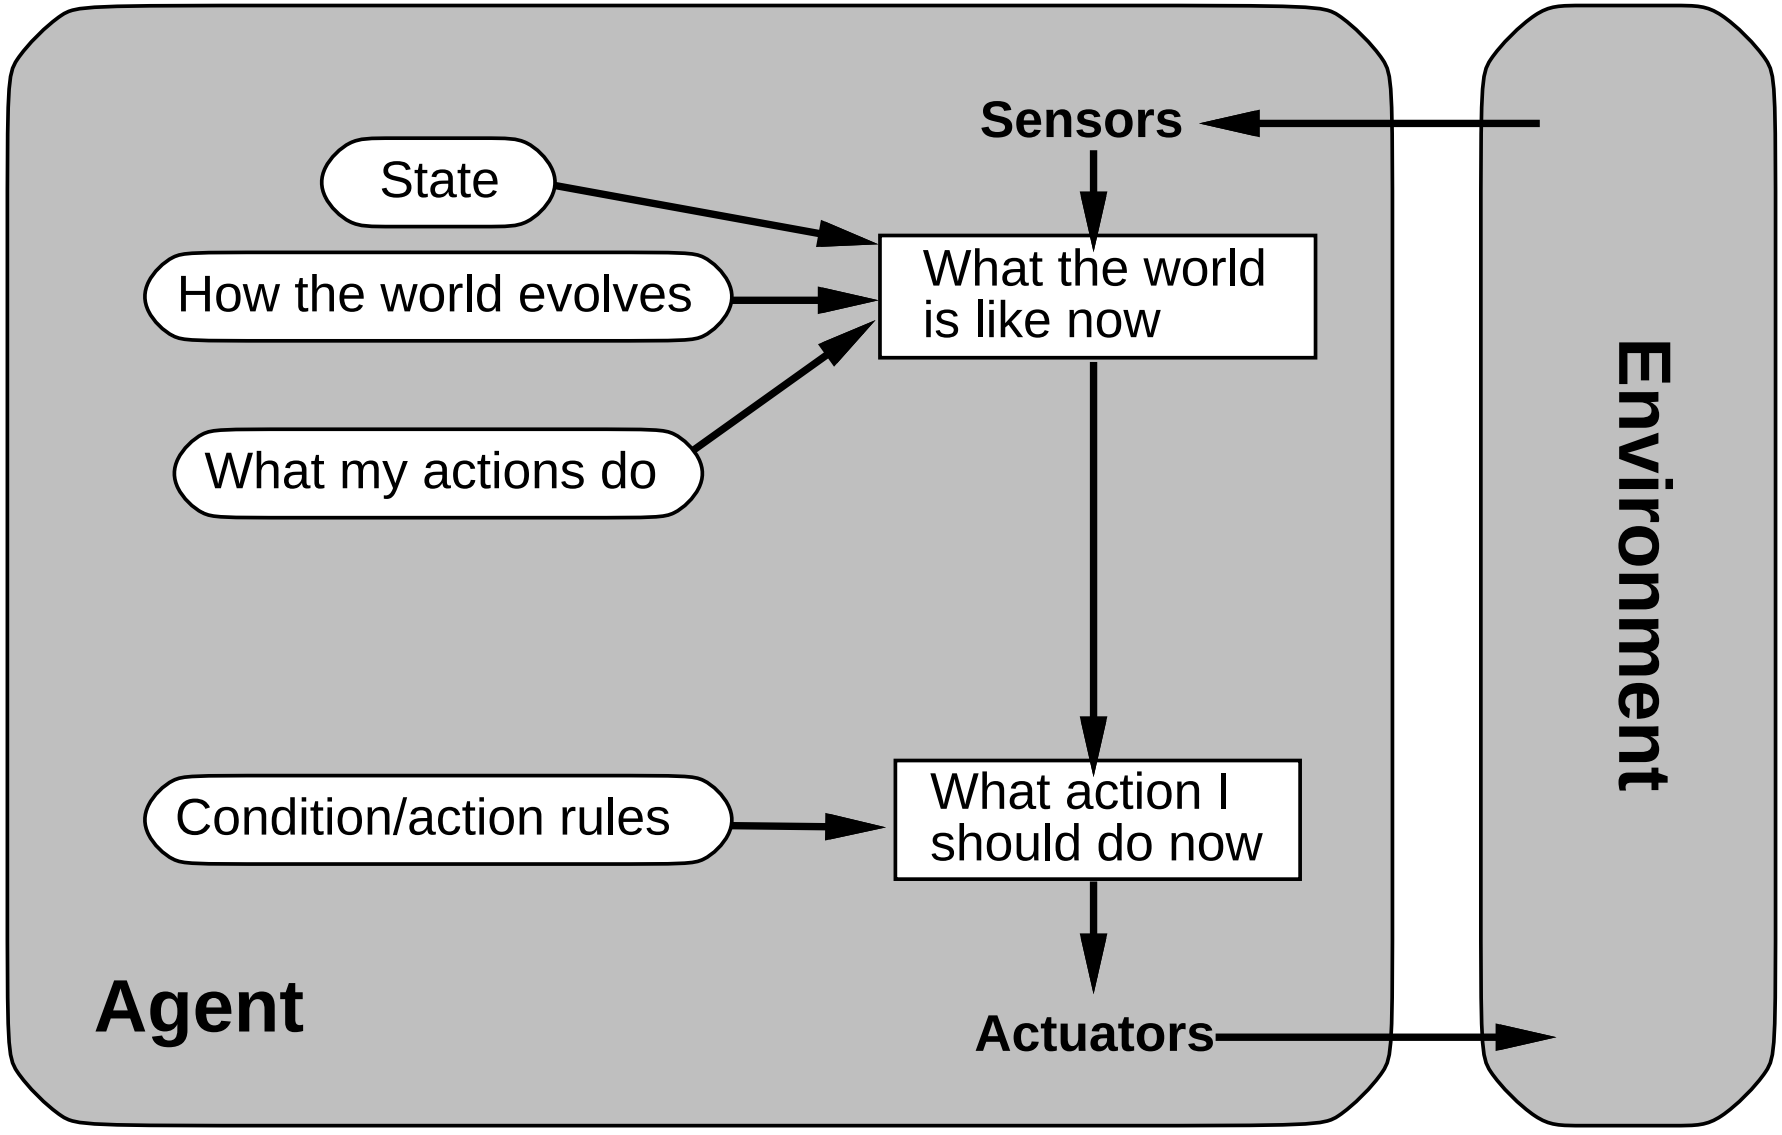

Goal-based agents

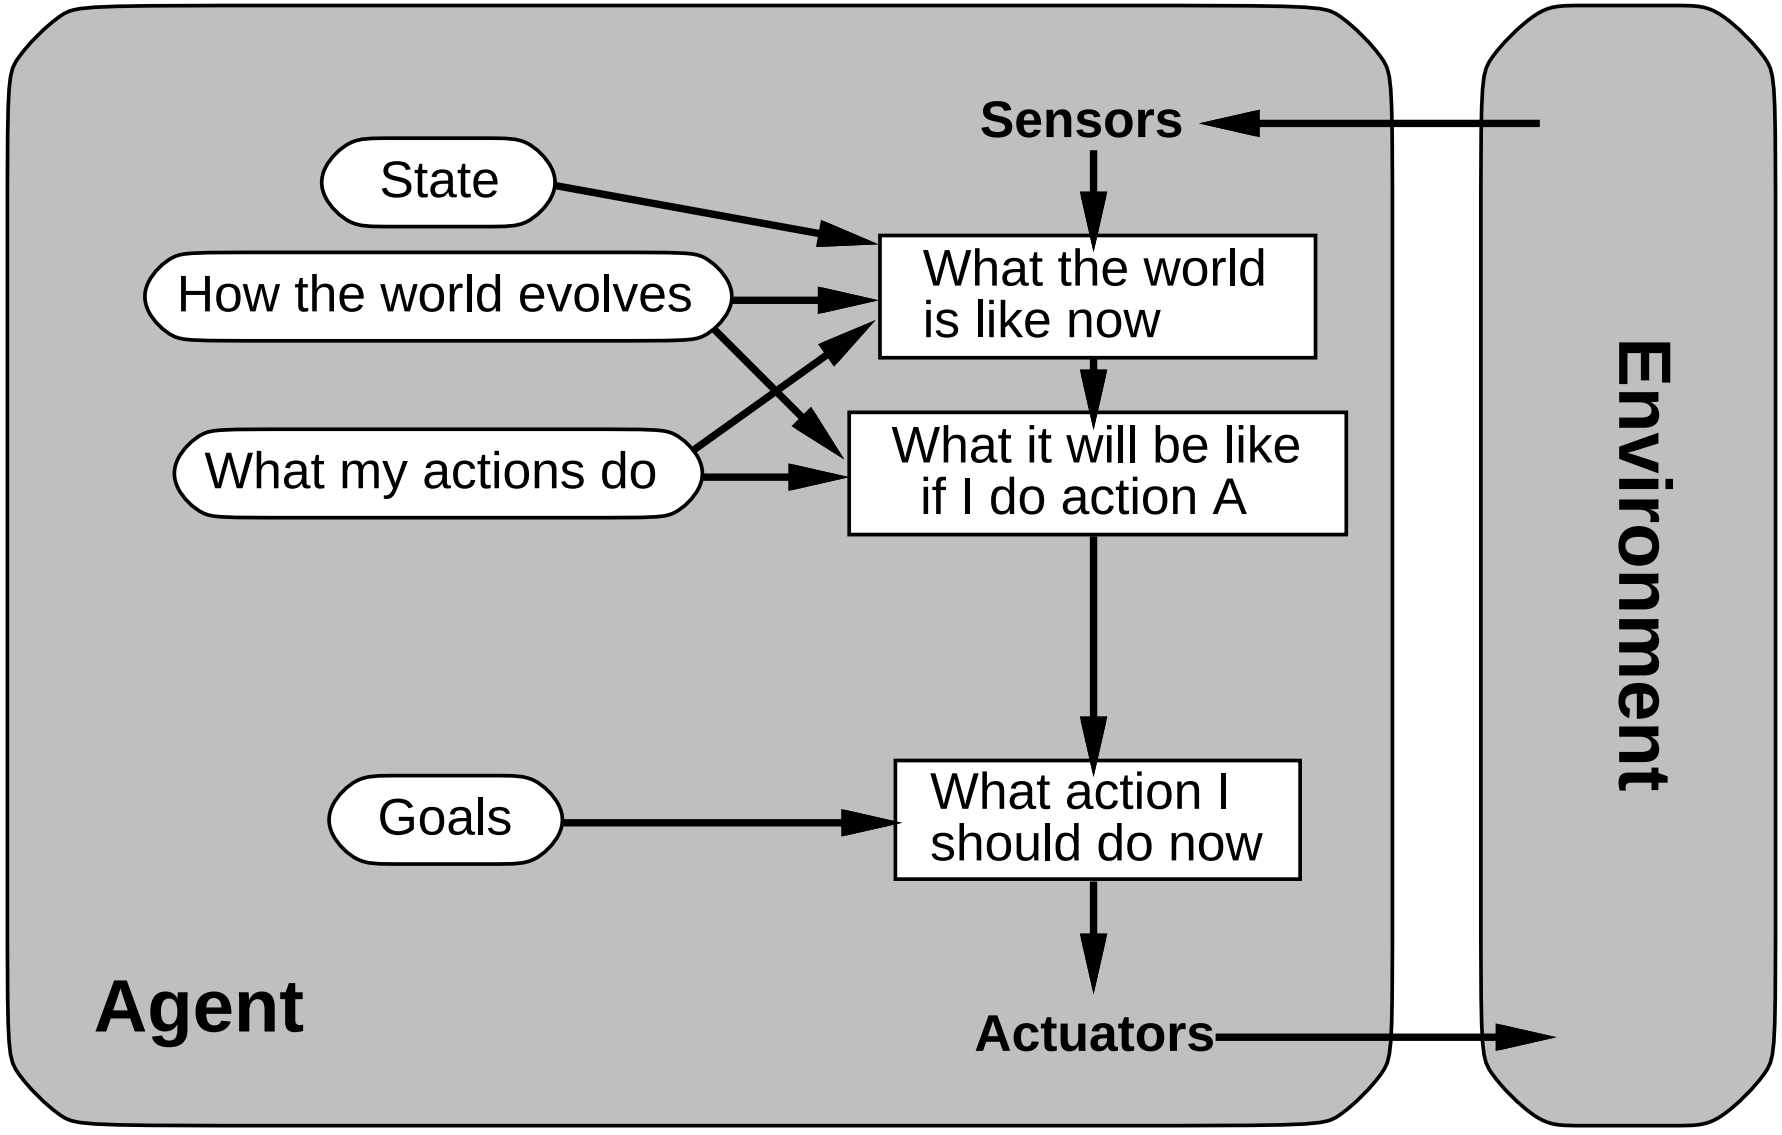

Utility-based agents

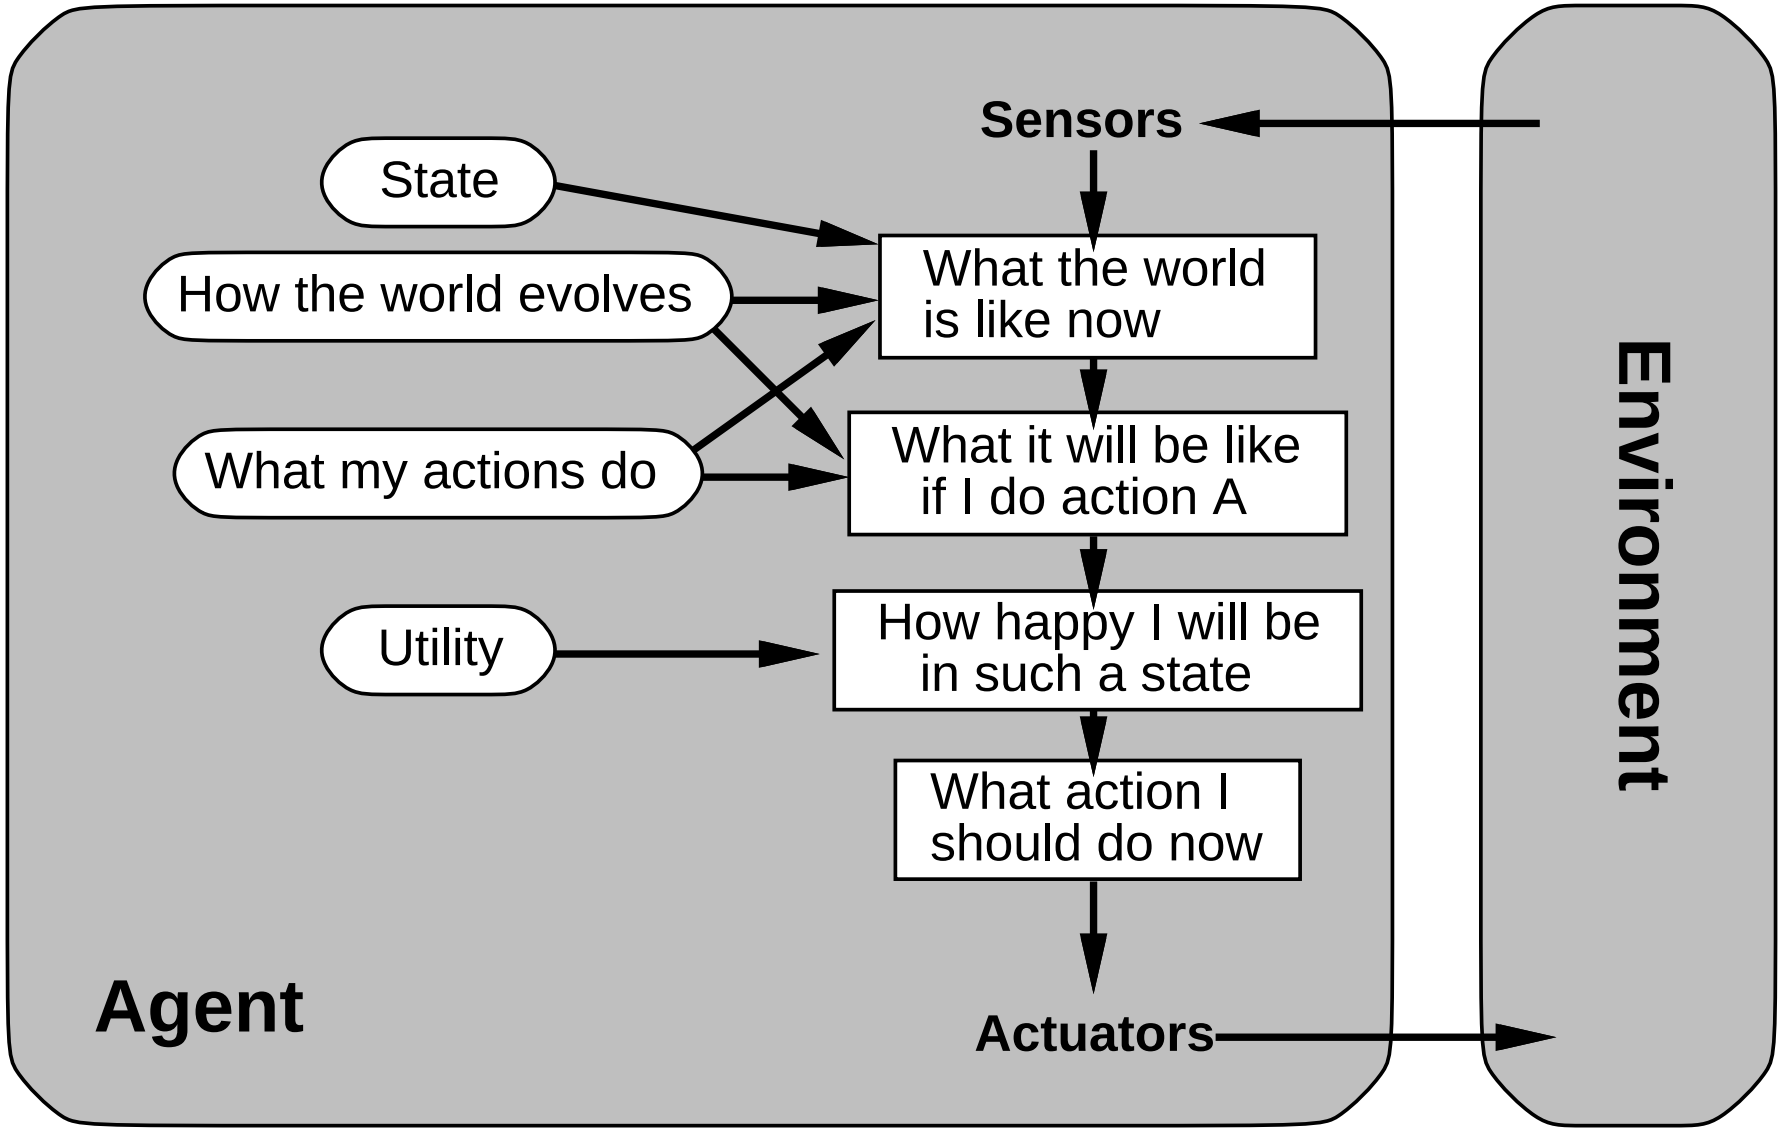

The vacuum-cleaner world

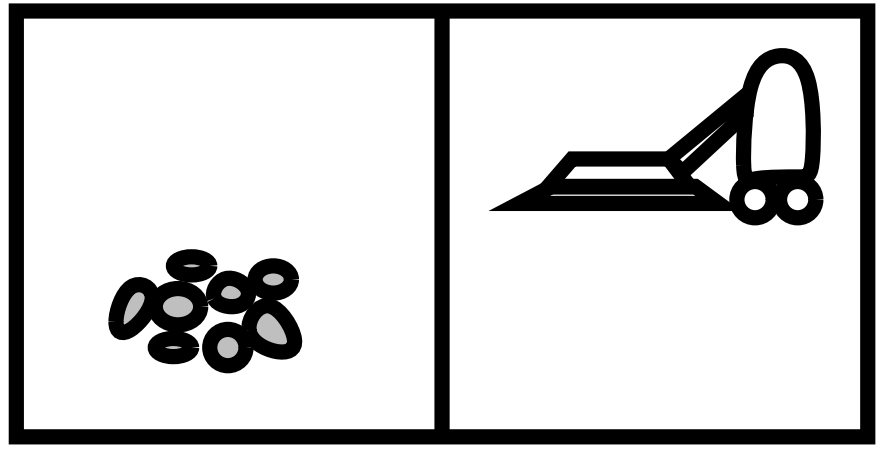

The vacuum-cleaner world

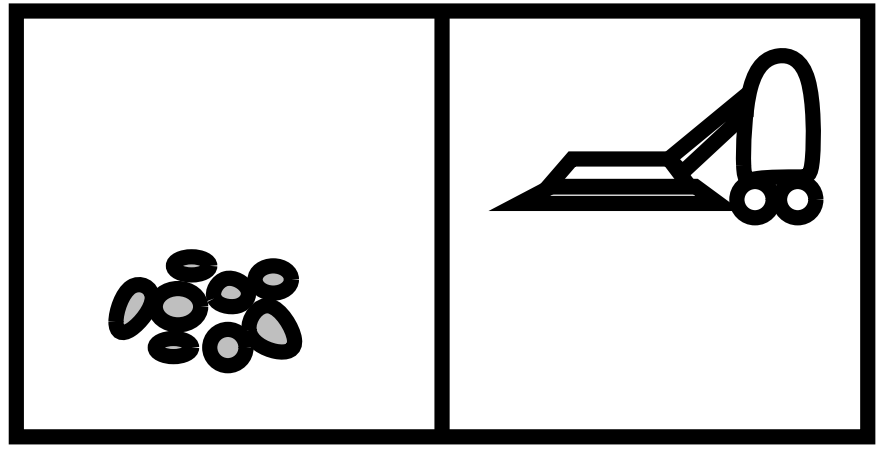

Percepts

- **location**
- ***dirty / not dirty***

The vacuum-cleaner world

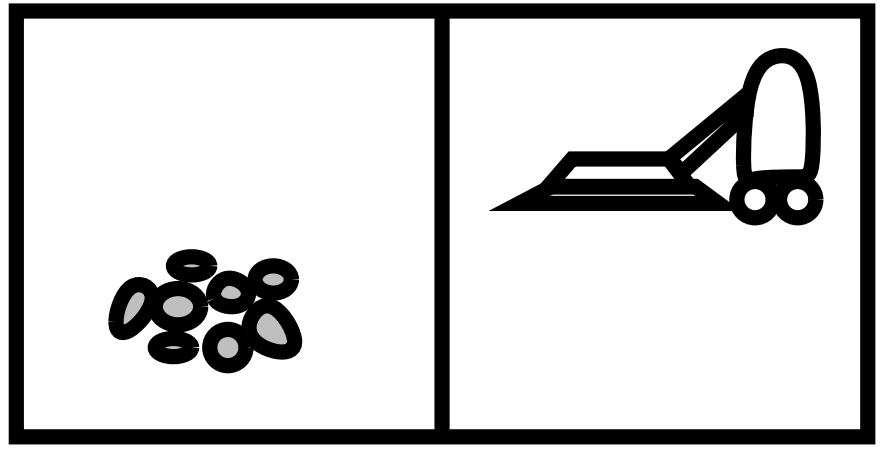

Percepts

- **location**
- ***dirty / not dirty***

Actions

- ***left***
- ***right***
- ***suck***
- ***noOp***

The vacuum-cleaner world

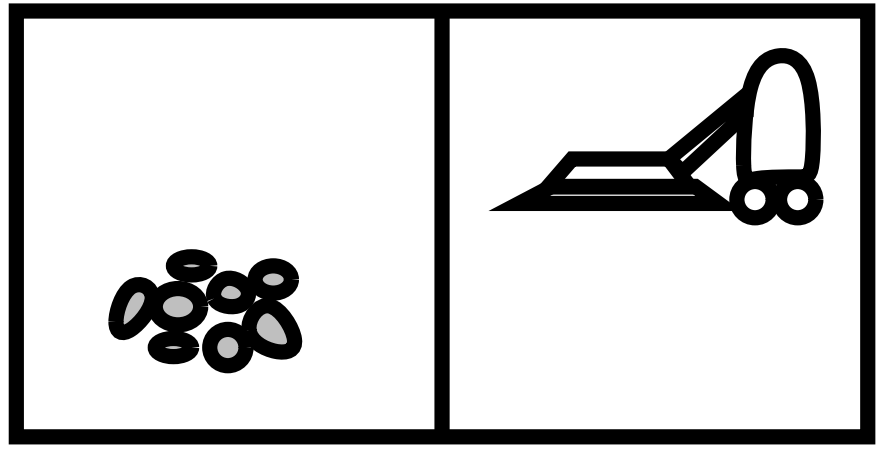

Performance measure

$+100$	**for each piece of dirt cleaned up**
-1	**for each action**
-1000	**for shutting off away from home**

Environment

- **grid**
- **dirt distribution and creation**

The vacuum-cleaner world

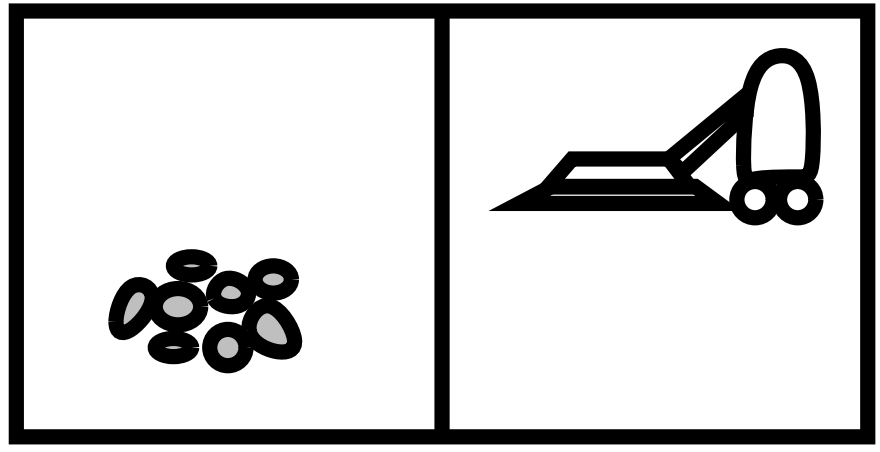

Observable? **Deterministic?** **Episodic?** **Static?** **Discrete?**