

Artificial Intelligence

CS482, CS682, MW 1 – 2:15, SEM 201, MS 227

Prerequisites: 302, 365

Instructor: Sushil Louis, sushil@cse.unr.edu, <http://www.cse.unr.edu/~sushil>

Question

- Are reflex actions rational? Are they intelligent?

Question

- How could introspection – reporting on one's inner thoughts – be inaccurate? Could I be wrong about what I am thinking?

Question

- To what extent are the following computer systems instances of artificial intelligence?
 - Supermarket bar code scanner
 - Web search engines
 - Voice-activated telephone menus
 - Internet routing algorithms that respond dynamically to the state of the network


Question

- Which tasks can currently be solved by computers?
 - Playing a decent game of table tennis
 - Driving in the center of Cairo, Egypt
 - Driving in Victorville, CA
 - Buying a week's worth of groceries at the market
 - Buying a week's worth of groceries on the web
 - Playing a decent game of bridge at a competitive level
 - Discovering and proving mathematical theorems
 - Writing an intentionally funny story
 - Giving competent legal advice in a specialized area of law
 - Translating spoken English into Swedish in real-time
 - Performing a complex surgical operation

Agents



- What is a rational agent?
 - Optimizes performance
- What are design principles for building rational agents (intelligent agents)?
- Agent performance will depend on their operating **Environment**. Some environments will be more difficult than others
- Types of environments
- Design considerations for agents in these different types of environments



What is an agent?

- Perceives an **environment** through **sensors** and acts on the environment through **actuators**

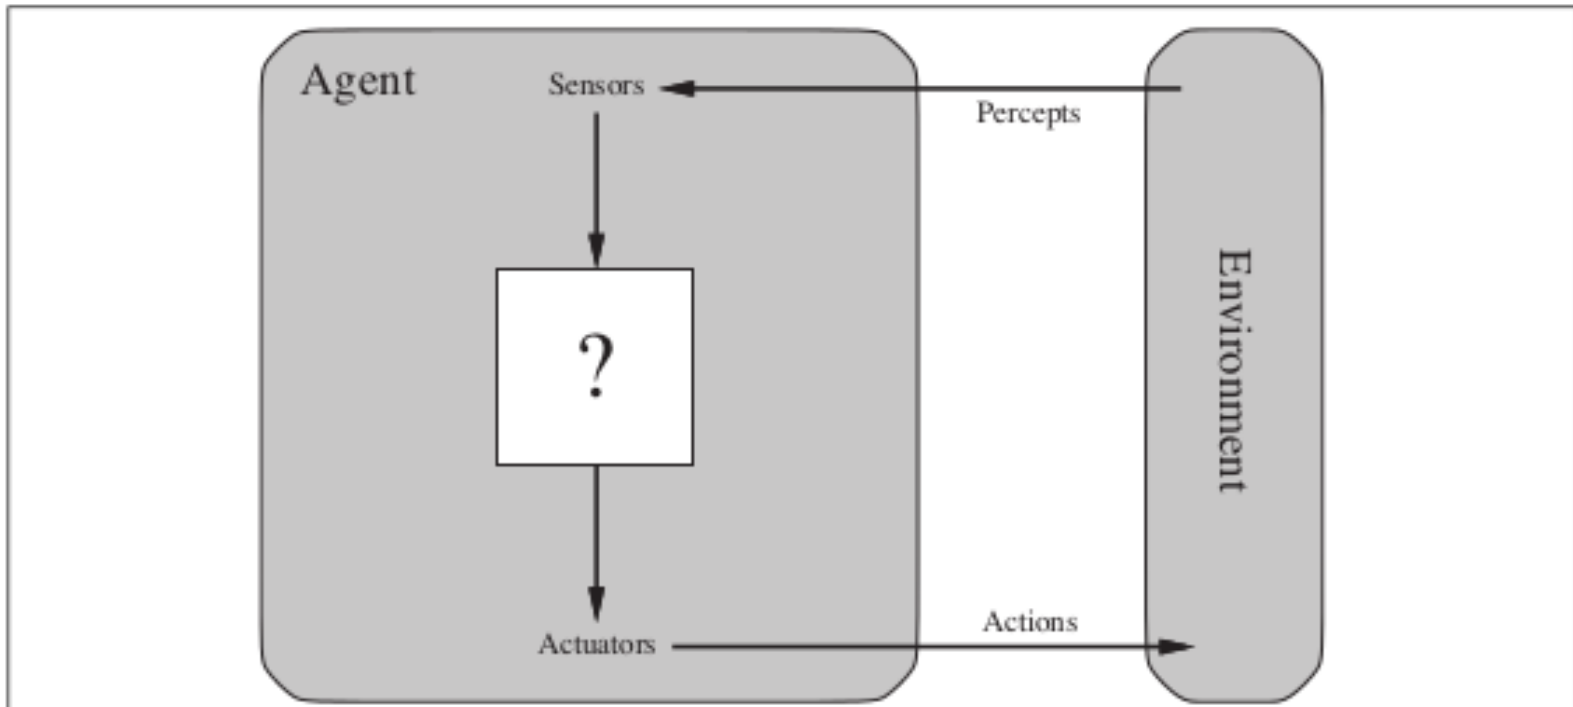


Figure 2.1 FILES: figures/agent-environment.eps (Tue Nov 3 16:22:19 2009). Agents interact with environments through sensors and actuators.

What is an Agent?

- **Percept:** Agent's perceptual inputs at any given **instant** of time
- Percept sequence is the complete history of everything the agent has ever perceived
- Agent's choice of action usually depends on percept sequence but not on anything it has not perceived
- Behavior is governed by an **agent function** that maps **percept sequence** to actions

Agent example




Figure 2.2 FILES: figures/vacuum2-environment.eps (Tue Nov 3 16:24:01 2009). A vacuum-cleaner world with just two locations.

Agent (Behavior) function

F (Percept sequence) \rightarrow Action

Percept Sequence	Action
[A, Clean]	Right
[A, Dirty]	Suck
[B, Clean]	Left
[B, Dirty]	Suck

Agent **programs** implement agent functions on some architecture


This is **just** a table, with percept sequences on the left and actions on the right

Behavior function

Percept Sequence	Action
[A, Clean]	Right
[A, Dirty]	Suck
[B, Clean]	Left
[B, Dirty]	Suck
[A, Clean], [A, Clean]	Right
[A, Clean], [A, Dirty]	Suck
...	
[A, Clean], [A, Clean], [A, Clean]	Right
[A, Clean], [A, Clean], [A, Dirty]	Suck
...	

This is **just** a table, with percept sequences on the left and actions on the right
But how big?

Vacuum Cleaner world



```
function REFLEX-VACUUM-AGENT([location,status]) returns an action
  if status = Dirty then return Suck
  else if location = A then return Right
  else if location = B then return Left
```

Add action "No-Op"

Rational Behavior versus Good behavior

- Agent's action affect environment → changes environment's state
- A sequence of agent actions → a sequence of environment states
- A performance measure on **environment states** differentiates good behavior from bad
 - Is our vacuum cleaner agent rational?
 - What is the performance metric?
 - What is the agent's prior knowledge?
 - What percept sequence has the agent seen?
 - What actions can agent perform?
 - Suppose the performance measure is just concerned with the first T time steps of the environment, show that a rational agent's action may depend not just on the state of the environment but also on time step


Rationality

- Performance metric
 - 1 point per square cleaned?
 - 1 point per square cleaned in time T ?
 - 1 point per square per time step minus one per move?
 - Penalty for $> k$ dirty squares?
 - ...
- A rational agent chooses whichever action maximizes expected value of performance measure given the precept sequence to date

Rationality

- Rational != omniscient
- Rational != clairvoyant
- Rational != successful

- Rational implies
 - Exploration
 - Learning
 - Autonomy



PEAS

- To design a rational agent, we need to specifying a **task environment**
- Consider automated taxi
 - Performance metric?
 - Environment?
 - Actuators?
 - Sensors?

PEAS

- To design a rational agent, we need to specifying a **task environment**
- Consider automated taxi
 - Performance metric?
 - Safety, destination, profits, legal, comfort, speed, cost
 - Environment?
 - US streets/freeways, traffic, pedestrians, weather, ...
 - Actuators?
 - Steering, accelerator, brake, horn, speaker/display, ...
 - Sensors?
 - Video, accelerometers, gauges, engine sensors, keyboard, GPS, ...

Internet shopping agent

- Performance metric?
- Environment?
- Actuators?
- Sensors?

Internet shopping agent

- Performance metric?
 - Price, quality, efficiency, appropriateness
- Environment?
 - Current and future websites, vendors, shippers, shoppers
- Actuators?
 - Display to use, follow URL, fill forms
- Sensors?
 - HTML pages (text, graphics, scripts)

Interactive English Tutor

- Performance Metric
- Environment
- Actuators
- Sensors

Interactive English Tutor

- Performance Metric
 - Score on test
- Environment
 - Students, testing agency
- Actuators
 - Display of exercises, suggestions, corrections, ...
- Sensors
 - Keyboard, mouse



Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword						
Chess with Clk						
Poker						
Taxi Driving						
Medical Diagnosis						
Image analysis						
Part-picking robot						
Refinery Controller						
Interactive English Tutor						

Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk						
Poker						
Taxi Driving						
Medical Diagnosis						
Image analysis						
Part-picking robot						
Refinery Controller						
Interactive English Tutor						

Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker						
Taxi Driving						
Medical Diagnosis						
Image analysis						
Part-picking robot						
Refinery Controller						
Interactive English Tutor						

Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Taxi Driving						
Medical Diagnosis						
Image analysis						
Part-picking robot						
Refinery Controller						
Interactive English Tutor						

Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Taxi Driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical Diagnosis						
Image analysis						
Part-picking robot						
Refinery Controller						
Interactive English Tutor						

Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Taxi Driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical Diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis						
Part-picking robot						
Refinery Controller						
Interactive English Tutor						

Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Taxi Driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical Diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot						
Refinery Controller						
Interactive English Tutor						

Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Taxi Driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical Diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery Controller						
Interactive English Tutor						

Types of task environments

Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Taxi Driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical Diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery Controller	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Interactive English Tutor						

Types of task environments

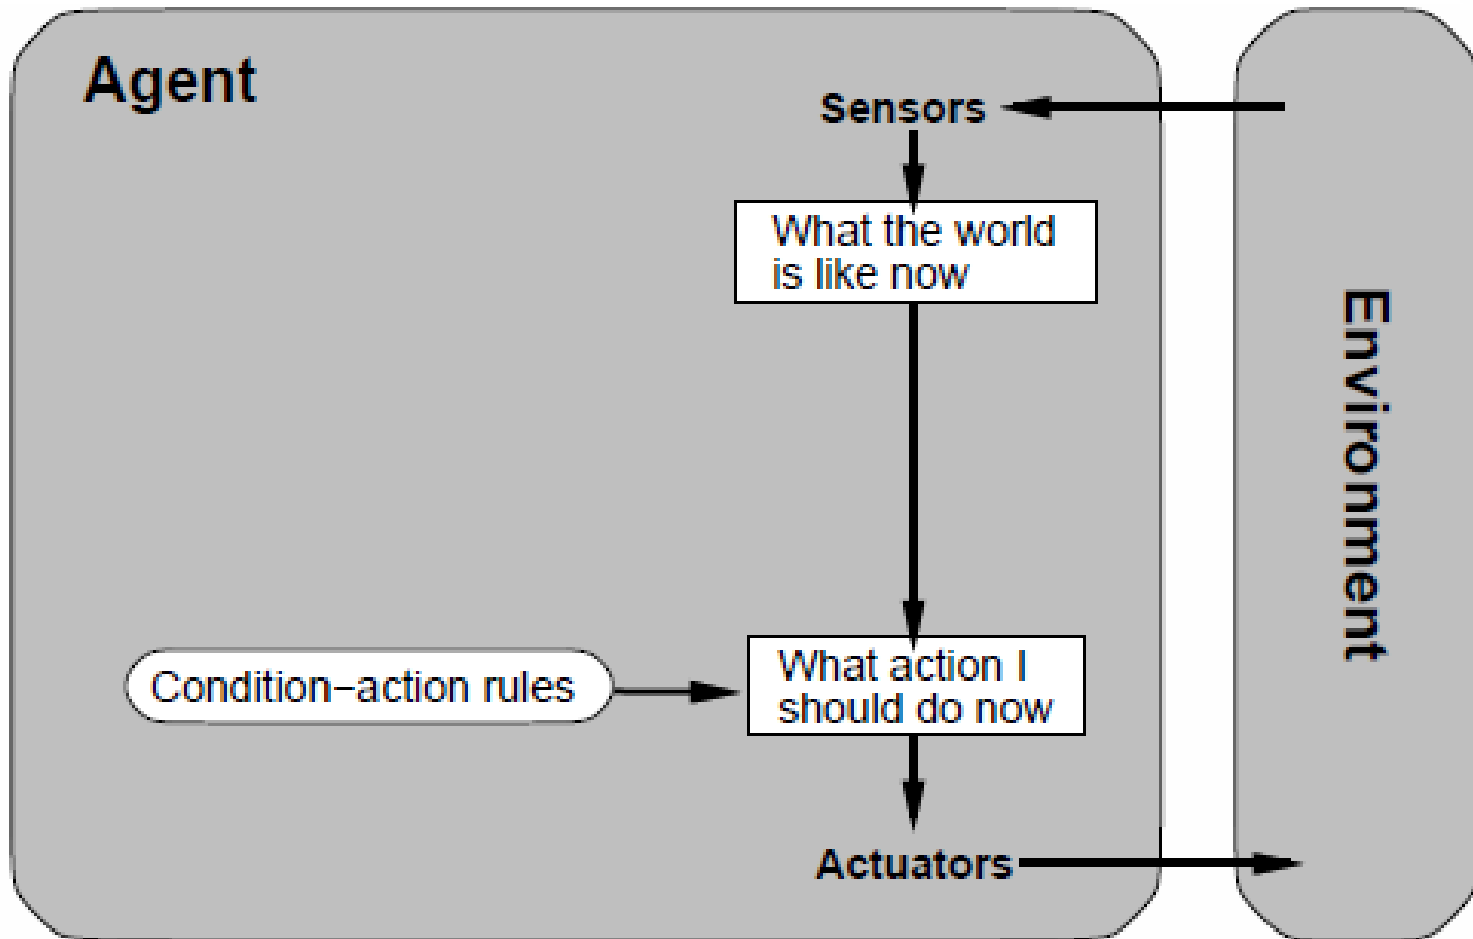
Task Env	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with Clk	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Taxi Driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical Diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery Controller	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Interactive English Tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

Types of agents

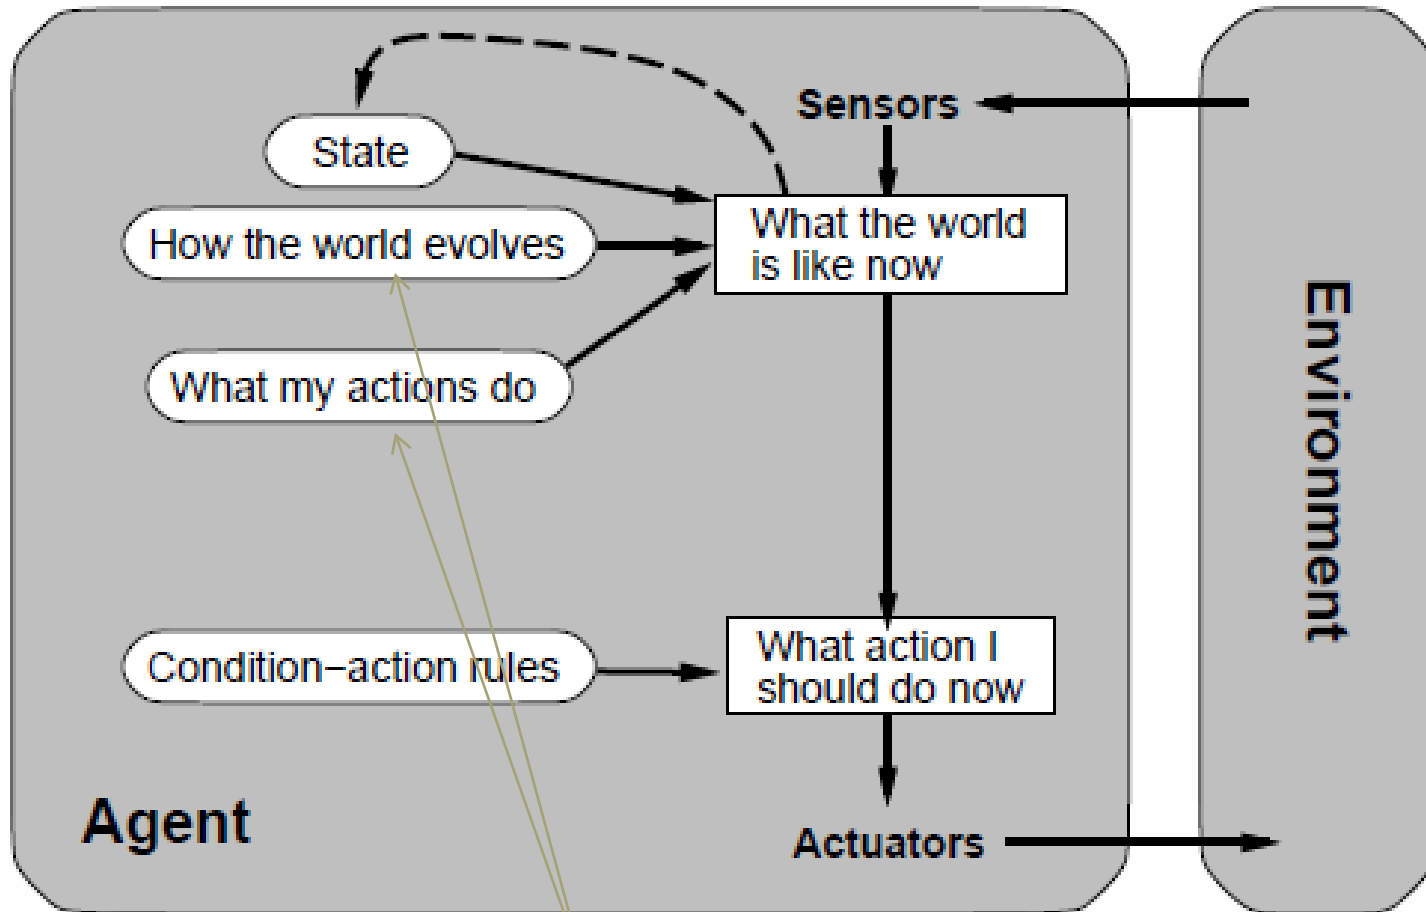
- Simple reflex agents
- Reflex agents with state
- Goal based agents
- Utility-based agents
- All can be turned into learning agents

```
function REFLEX-VACUUM-AGENT([location,status]) returns an action
  if status = Dirty then return Suck
  else if location = A then return Right
  else if location = B then return Left
```

Simple reflex agents

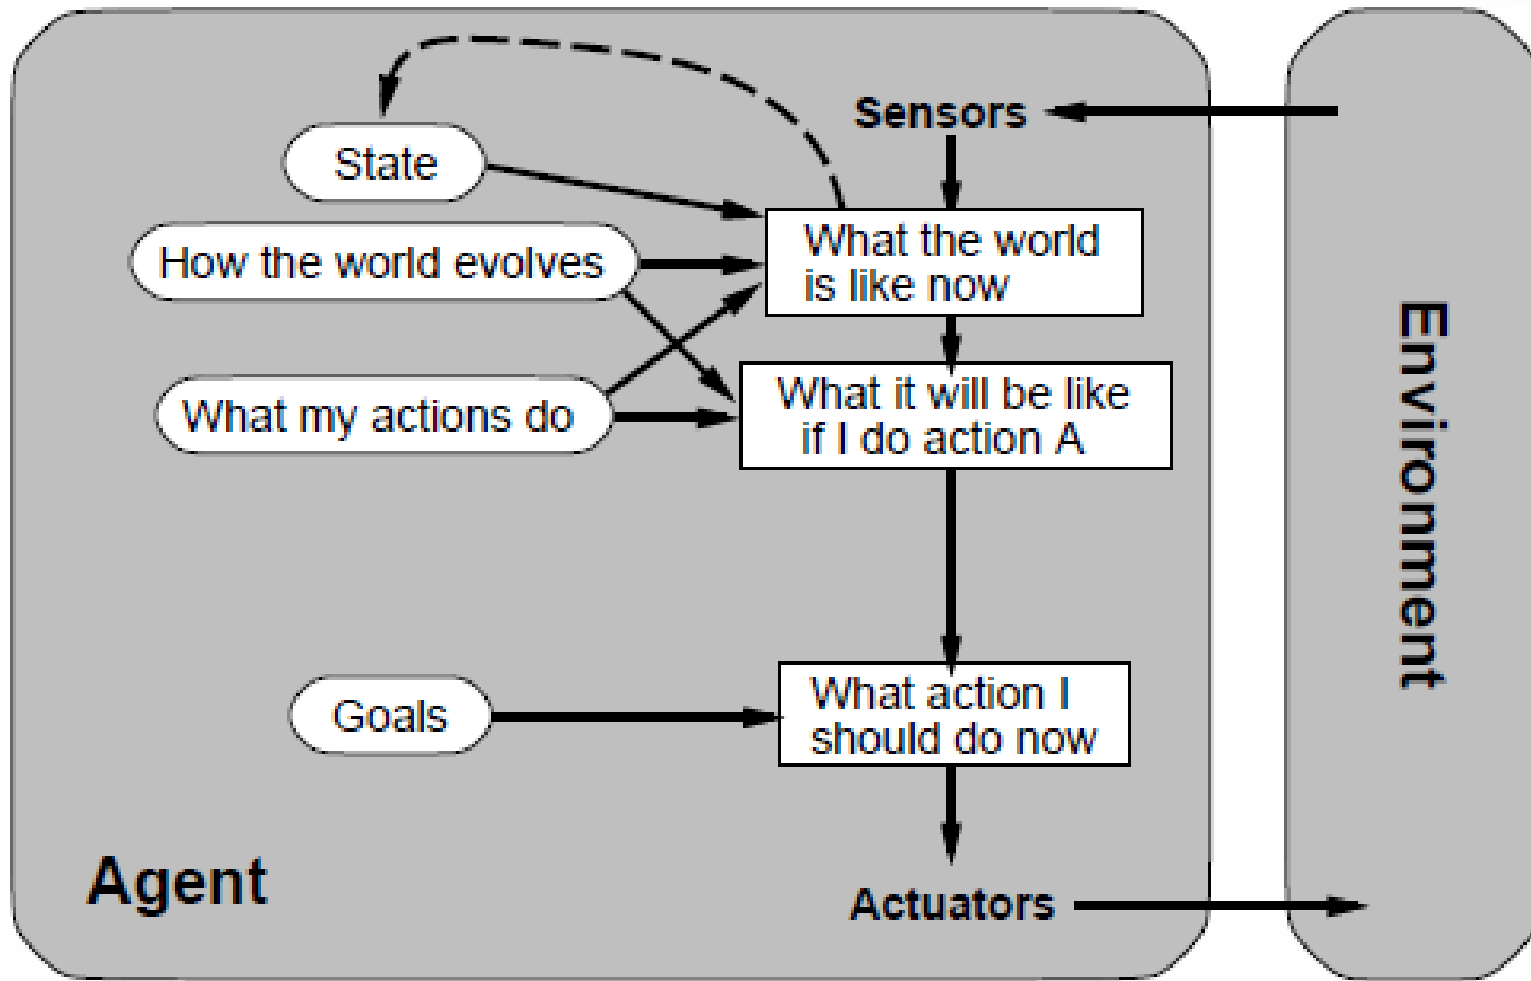


Reflex agent with state



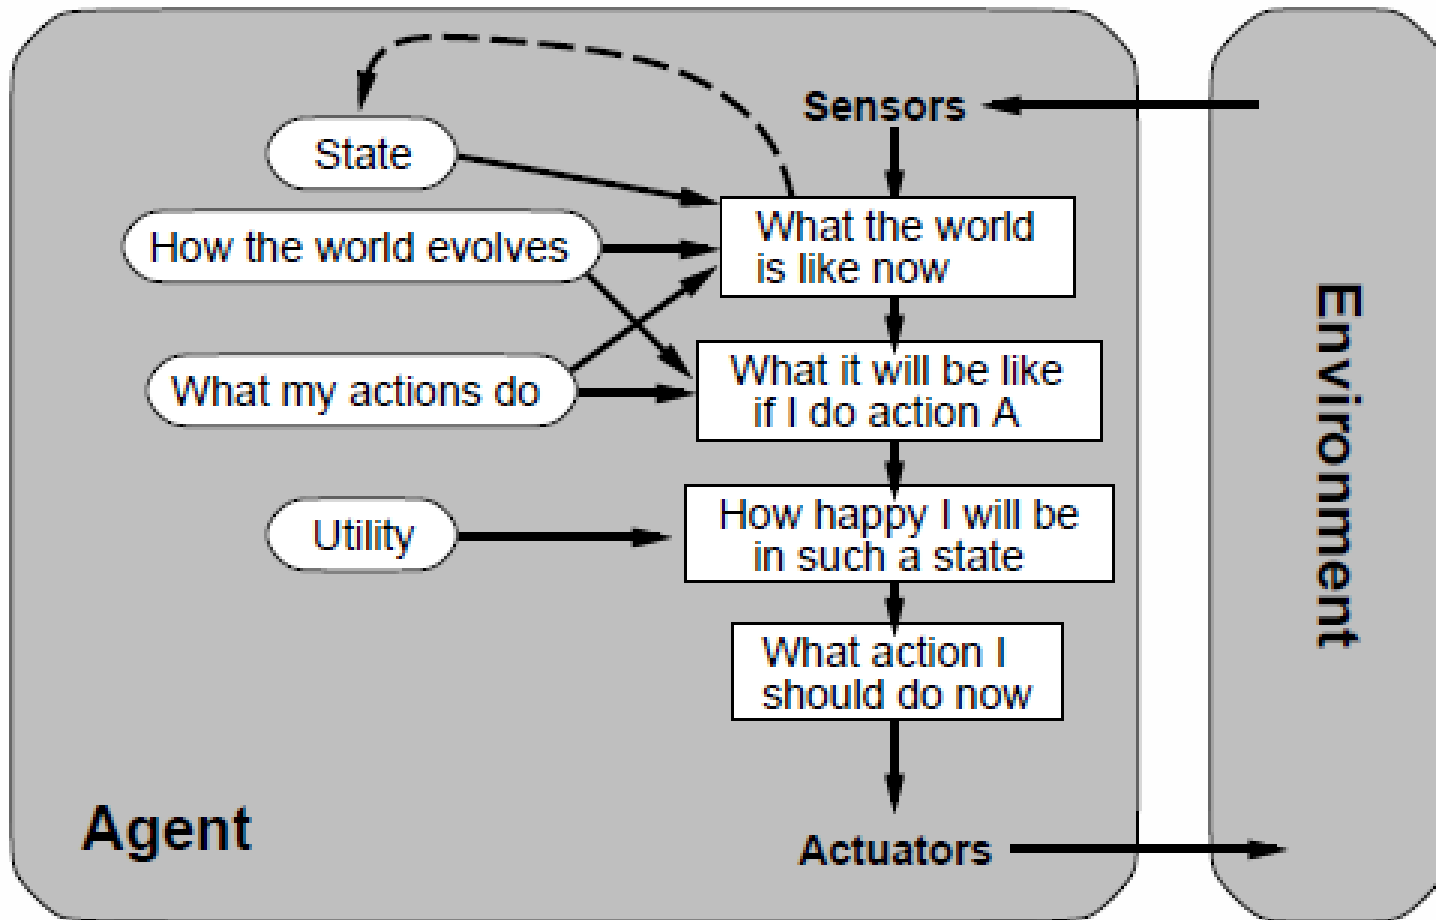
Model-based agent

Goal-based agent



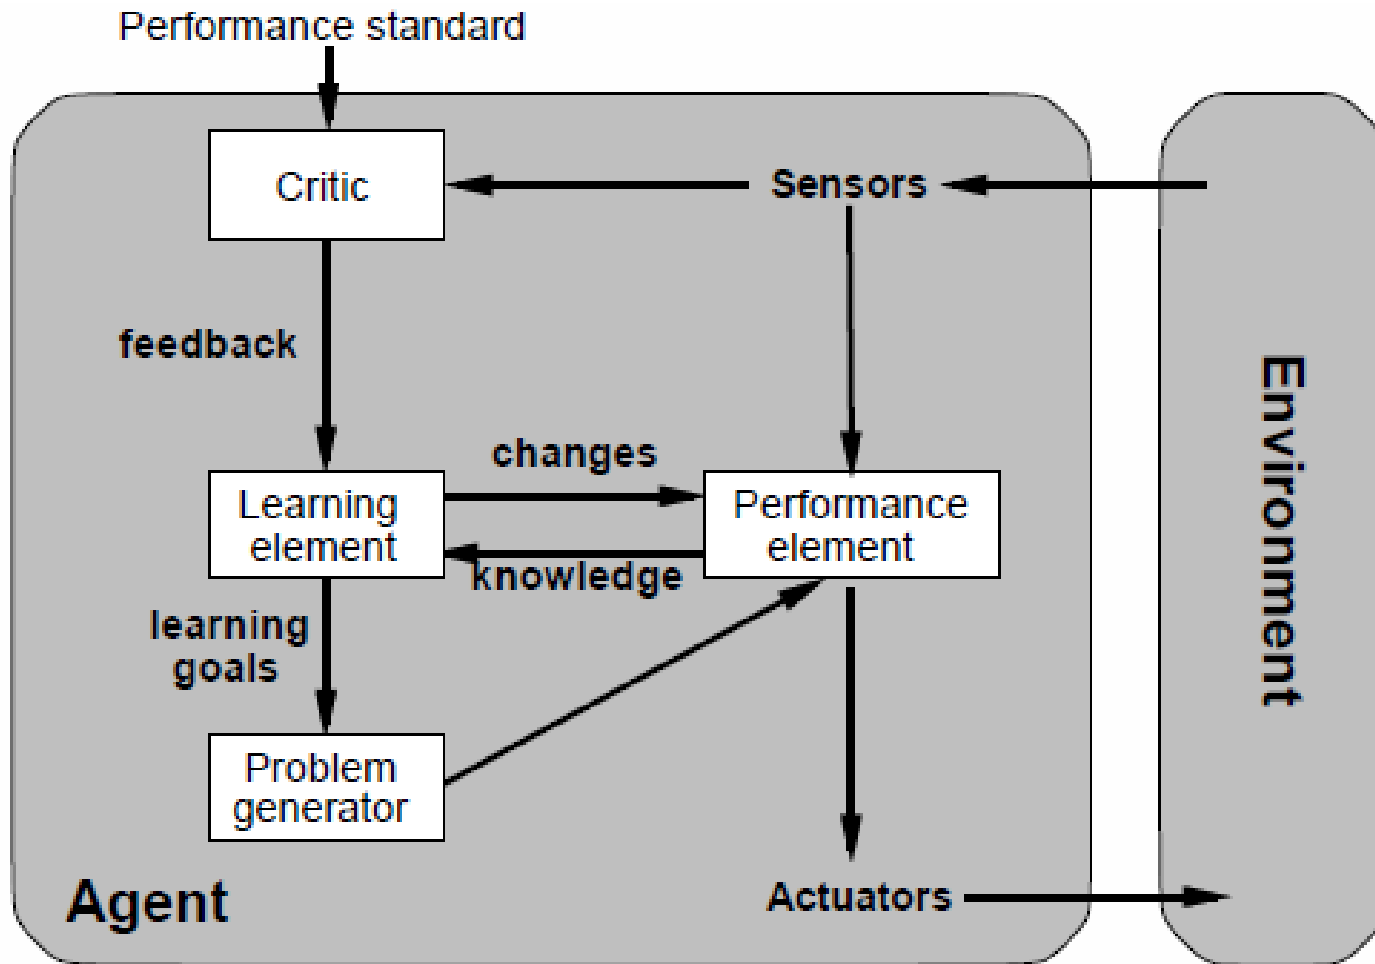
Search for ways to achieve goals. Make plans to achieve goals. Searching for plans

Utility-based agent




Maximizes **expected** utility


Learning agents




Representing environmental states



(a) Atomic



(b) Factored



(b) Structured

Less expressive



More Expressive

Summary

- Agents interact with environment with actuators and sensors
- Agent function describes agent behavior
- Performance measure evaluates the environment sequence produced by agent actions
- A perfectly rational agent maximizes expected performance
- Agent programs implement agent functions on some architecture
- PEAS descriptions define task environments
- Environments can be categorized along
 - Observable, deterministic, episodic, static, discrete, single-agent
- Several basic single-agent architectures exist
 - Reflex, reflex with state, goal-based, utility-based, learning