Module1

Source Book: Stuart J. Russell and Peter Norvig, Artificial Intelligence, 3rd Edition, Pearson,2015

Topics

- What is AI?
- The State of Art : **What can AI do today?**
- Agents, The structure of agents
- Concept of Rationality,
- The nature of environment.

What is AI?

- Artificial Intelligence (AI) is a field of computer science dedicated to develop systems capable of performing tasks that would typically require human intelligence.
- These tasks include **learning, reasoning, problem-solving, perception, understanding natural language**, and even interacting with the environment.
- AI aims to create machines or software that can mimic or simulate human cognitive functions.

According to Russell and Norvig, AI can be defined as follows:

" *AI (Artificial Intelligence) is the study of agents that perceive their environment, reason about it, and take actions to achieve goals. Each such agent implements a function that maps percept sequences to actions, and the study of these functions is the subject of AI.*"

Eight explanations of **AI shown in two groups**

Thinking Humanly	Thinking Rationally
"The exciting new effort to make computers think . . . *machines with minds*, in the full and literal sense." (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . ." (Bellman, 1978)	"The study of mental faculties through the use of computational models." (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Acting Humanly	**Acting Rationally**
"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)	"Computational Intelligence is the study of the design of intelligent agents." (Poole *et al.*, 1998) "AI . . . is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

Turing Test

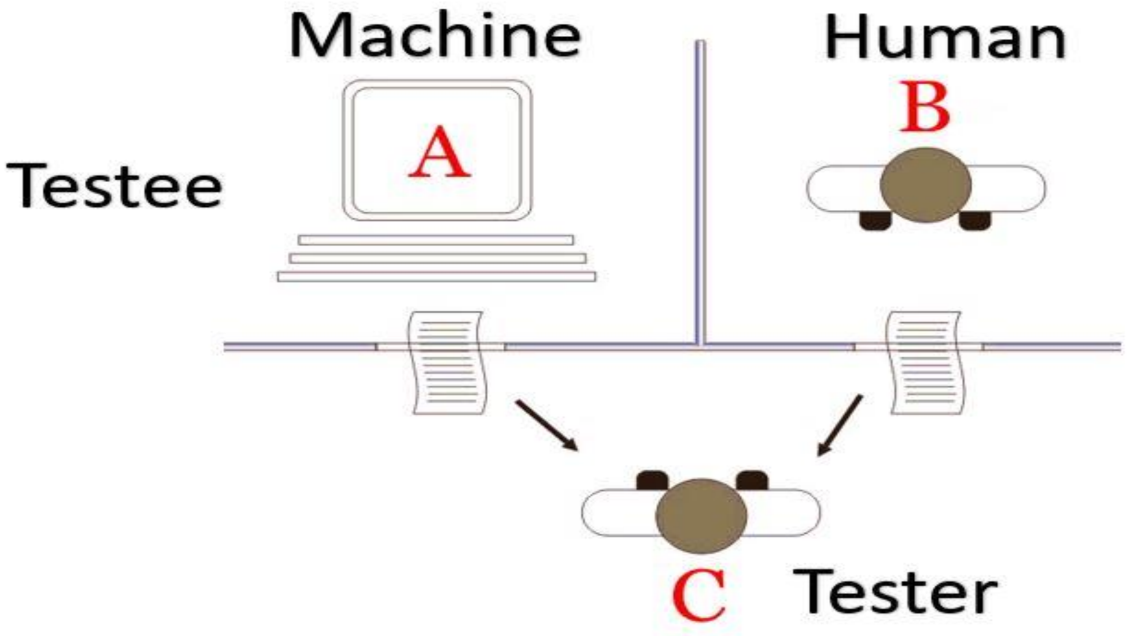

The State of Art: : What can AI do today?

- **Robotic Cars**
- **Speech Recognition**
- **Autonomous Planning and Scheduling**
- **Game Playing**
- **Spam Filtering**
- **Logistics Planning**

Agent

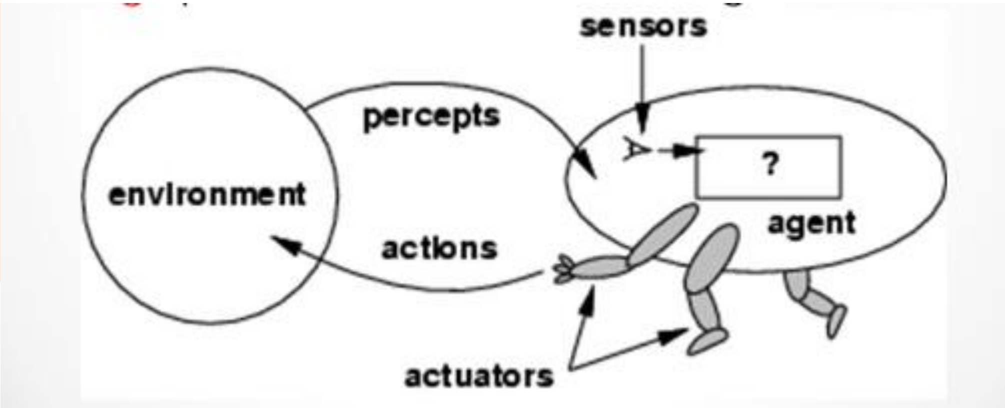

In AI, an agent is any entity that can perceive its environment and take actions to achieve its goals. These agents can be **physical robots, software programs, or any system capable of interacting** with the world. The agents sense the environment through sensors and act on their environment through actuators. An AI agent can have mental properties such as knowledge, belief, intention, etc.

Agent

An agent is defined as **anything capable** of perceiving its environment through sensors and acting upon that environment through actuators. This basic concept is depicted in Figure 2.1.

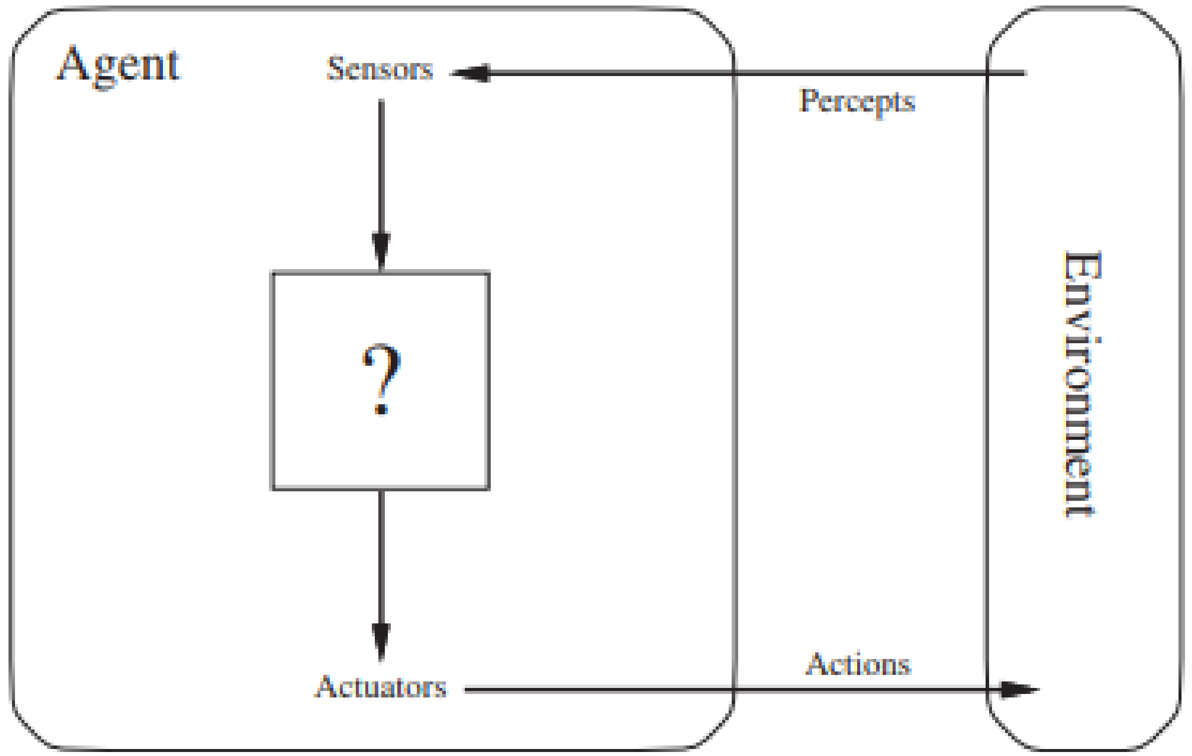

Figure 2.1 Agents interact with environments through sensors and actuators.

For instance, a **human agent** employs eyes, ears, and other sensory organs while utilizing hands, legs, vocal tract, etc., as actuators.

In contrast, a **robotic agent** might use cameras and infrared range finders as sensors and various motors as actuators.

A **software agent** takes in sensory inputs such as keystrokes, file contents, and network packets and responds by displaying on the screen, writing files, or sending network packets.

Percept: The term **"percept"** refers to an agent's perceptual inputs at any given moment, and a percept sequence encompasses the complete history of everything the agent has perceived

Table: A table serves as an external representation of an agent's behaviour, specifically presented in tabular form that outlines the agent's actions for every possible percept sequences.

Agent Program: An agent program represents the internal implementation of an agent's behaviour or actions.

Vacuum Cleaner World

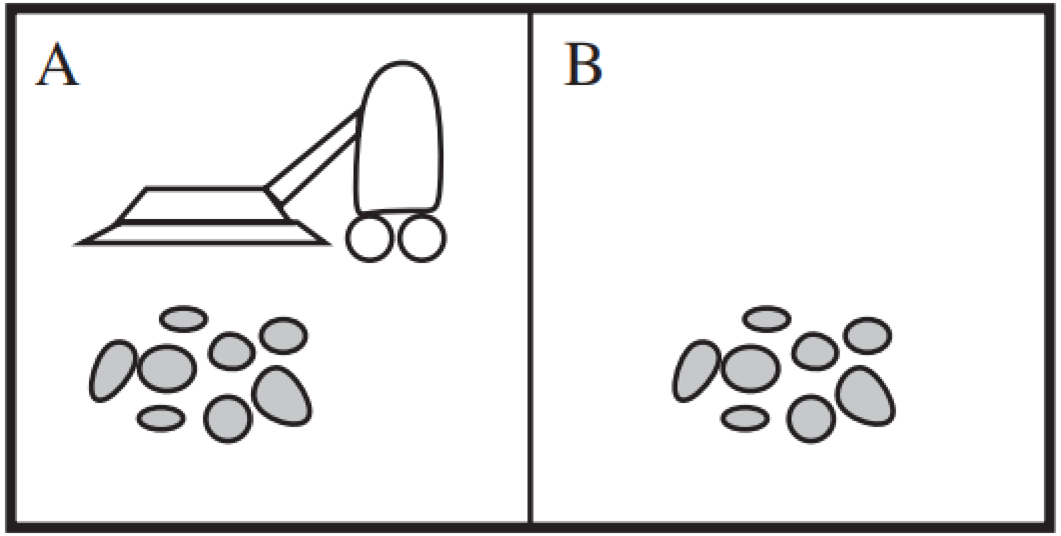

Table [Agent function]

Percept sequence	Action
$[A, Clean]$	$Right$
$[A, Dirty]$	$Suck$
$[B, Clean]$	$Left$
$[B, Dirty]$	$Suck$
$[A, Clean], [A, Clean]$	$Right$
$[A, Clean], [A, Dirty]$	$Suck$
$\vdots$	$\vdots$
$[A, Clean], [A, Clean], [A, Clean]$	$Right$
$[A, Clean], [A, Clean], [A, Dirty]$	$Suck$
$\vdots$	$\vdots$

Agent Program

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*

Characteristics of Intelligent agents

In artificial intelligence (AI), intelligent agents are entities that perceive their environment and take actions to achieve **specific goals**.

Characteristics

Perception: Agents have the ability to perceive or sense their environment. This involves gathering information from the surroundings through sensors or other means.

Action: Intelligent agents can take actions in response to their perceptions. These actions are chosen to influence the state of the environment or achieve specific goals.

Autonomy: Agents operate autonomously, making decisions and taking actions without direct human intervention. Autonomy allows them to adapt to changing conditions in their environment.

Goal-Directed Behaviour: Intelligent agents are typically designed to achieve specific goals. These goals can be predefined by a programmer or learned by the agent through experience.

Learning and Adaptation: Many intelligent agents have the capability to learn from experience and adapt their behaviour over time. This learning process may involve acquiring new knowledge or improving decision-making strategies.

Concept of Rationality

A rational agent is defined as one that consistently does right thing and makes decisions leading to favourable outcomes. This is reflected in the **correct completion of every entry in the agent function table**.

However, determining what constitutes the "**right thing**" involves the following:

- **Assessing Consequences**
- **Performance Measure**:
  - **Challenges and Considerations in Designing Performance Measures**
  - **Alignment with Environmental Goals**
- **Philosophical Considerations**

What is Rationality?

Rationality encompasses the state of being **reasonable, sensible, and possessing a sound judgment**. It pertains to anticipated actions and outcomes based on the agent's perceptions. Engaging in actions with the objective of acquiring valuable information is a crucial aspect of rationality.

Rationality at Any Given Time depends on the following four things:

1. **Performance Measure**: The criterion defining success.
2. **Agent’s Prior Knowledge**: Understanding of the environment.
3. **Available Actions**: The actions the agent can perform.
4. **Percept Sequence**: The agent’s historical sensory input.

Omniscience, Learning, and Autonomy in Rational Agents

Omniscience and Rationality: Distinguishing between rationality and omniscience is crucial. An omniscient agent knows the actual outcome of its actions and can act accordingly; but omniscience is impossible in reality. While rationality maximizes expected performance.

Information Gathering : Rational agents should perform actions to modify future percepts, known as information gathering. This concept is vital for maximizing expected performance.

Learning in Rational Agents: Rational agents are not only expected to gather information but also to learn from their perceptions. An agent's initial configuration may reflect prior knowledge, but as it gains experience, learning becomes imperative.

Dung Beetle

Think about the ***dung beetle****, a small insect that carefully constructs its nest, lays eggs, and uses a ball of dung to seal the entrance****. If the dung ball is taken away*** *while the beetle is carrying it, the beetle continues its job as if the dung ball is still there, not realizing its missing.*

Sphex Wasp

*The **sphex wasp** is a bit smarter. It digs a hole, stings a caterpillar, drags it into the hole, checks everything is fine, and then lays eggs. The caterpillar becomes food for the hatching eggs. However, if someone moves the caterpillar a little while the wasp is checking, it goes back to dragging the caterpillar, not realizing the change. Even after many attempts to move the caterpillar, the wasp doesn't learn that its plan isn't working and keeps doing it the same way.*

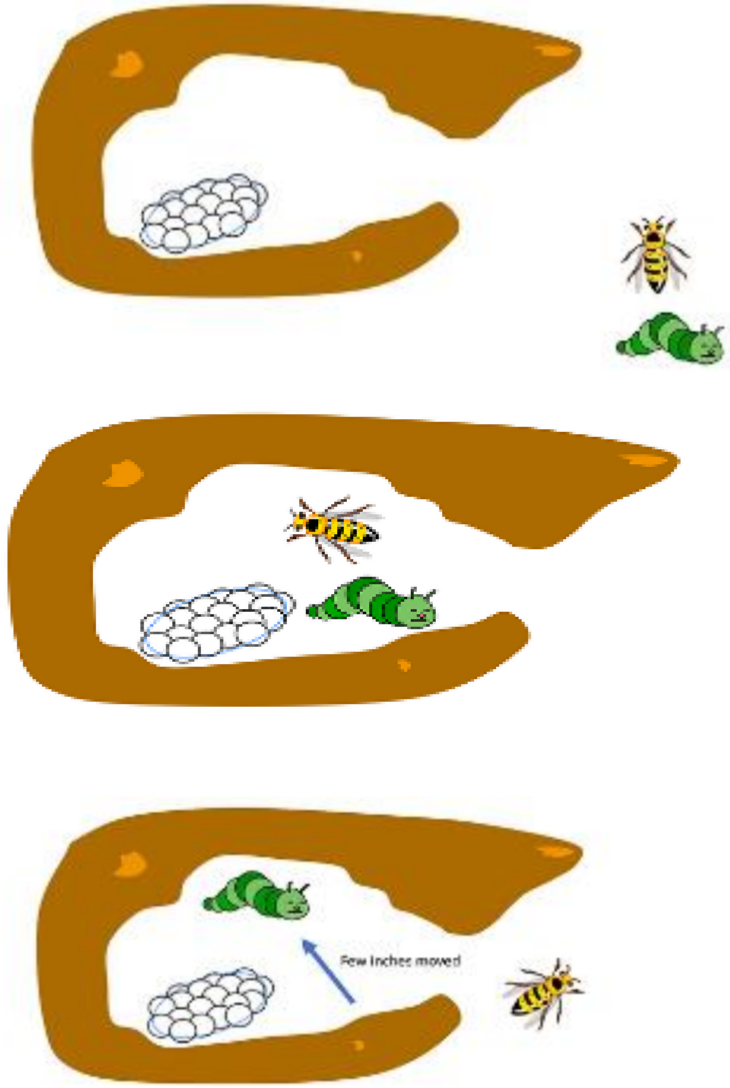

Autonomy in Rational Agents

Autonomy is emphasized as a crucial aspect of rationality. An autonomous agent learns to compensate for **partial or incorrect prior knowledge**, ensuring effective adaptation to the environment.

In essence, the combination of **rationality, information gathering, and learning enables the design of autonomous agents** capable of success across diverse environments.

Nature of Environments

- The task environment, refers to the **external system or surroundings** within which an **agent** operates and performs tasks.
- **Task environment is specified in terms of PEAS (Performance, Environment, Actuators, Sensors)**.

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of 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 of scene categorization	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 of exercises, suggestions, corrections	Keyboard entry

Properties of Task Environments

1.Fully Observable vs. Partially Observable:

Fully Observable: The agent can see the complete environment all the time.

Partially Observable: The agent's sensors may not capture all relevant aspects due to noise or missing information.

2. Single Agent vs. Multiagent:

Single Agent: An agent operates independently, like solving a crossword puzzle.

Multiagent: Agents interact with each other, introducing complexities in decision-making.

3. Deterministic vs. Stochastic:

Deterministic: The environment's next state is entirely determined by the current state and the agent's action.

Stochastic: Some uncertainty exists, making predictions challenging. For example, traffic behavior in taxi driving is stochastic.

4. **Episodic vs. Sequential:**

Episodic: Each decision is independent, not influenced by past decisions.

Sequential: Current decisions affect future ones, as seen in chess or taxi driving.

5. **Static vs. Dynamic**:

Static: The environment remains unchanged during decision-making.

Dynamic: The environment evolves continuously, demanding constant

6. **Discrete vs. Continuous**:

Discrete: Finite states, actions, or percepts.

Continuous: Values or actions vary smoothly.

7. **Known vs. Unknown:**

Known: The agent knows outcomes or probabilities.

Unknown: The agent needs to learn about the environment.

Task Environment	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword puzzle	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with a clock	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Backgammon	Fully	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

Structure of Intelligent Agents

Agent's structure can be viewed as –

- **Agent = Architecture + Agent Program**
- **Architecture = the machinery that an agent executes on.**
- **Agent Program = an implementation of an agent function.** **All agent programs share a common structure**: they take the current percept as input from the **sensors** and produce an action for the **actuators**.

Agent program

- The **agent program**, operates based on the current percept,
- The **agent programs** are presented in a simple pseudocode language

```
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
```

Agent Function

- The **agent function**, considers the entire percept history.
- Agent Function is represented using Table.

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*
⋮	⋮

Types of Agent Programs

There are, five basic types of agent programs as given below:

1. Simple reflex agents
2. Model-based reflex agents
3. Goal-based agents
4. Utility-based agents
5. Learning Agents

1.Simple reflex agents

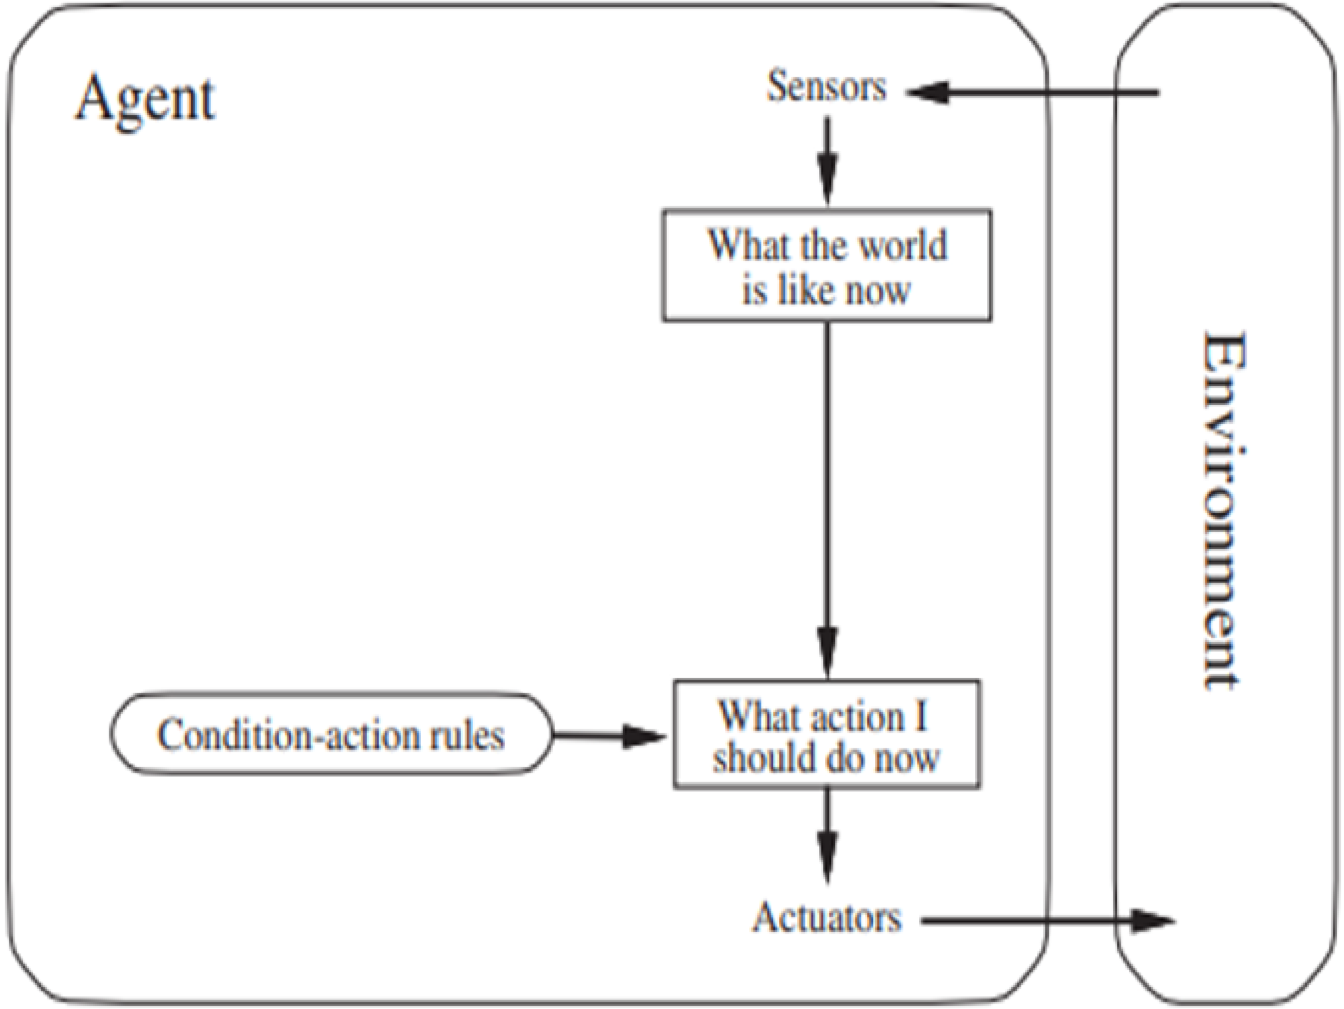

function REFLEX-VACUUM-AGENT([*location*,*status*]) **returns** an action

if *status* = *Dirty* **then return** *Suck*
else if *location* = *A* **then return** *Right*
else if *location* = *B* **then return** *Left*

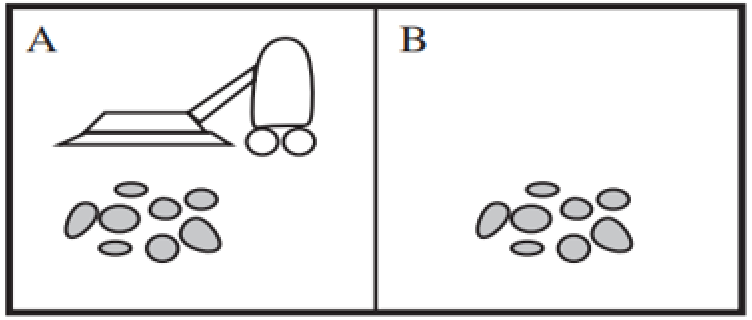

Agent Program Example 1

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*

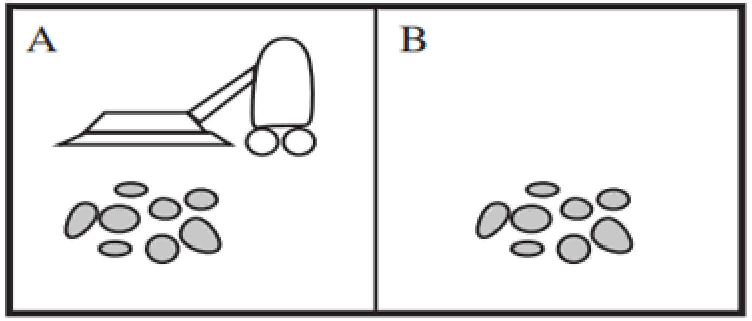

Agent Program Example 2

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*

if car-in-front-is-braking then initiate-braking

2. Model-based reflex agents

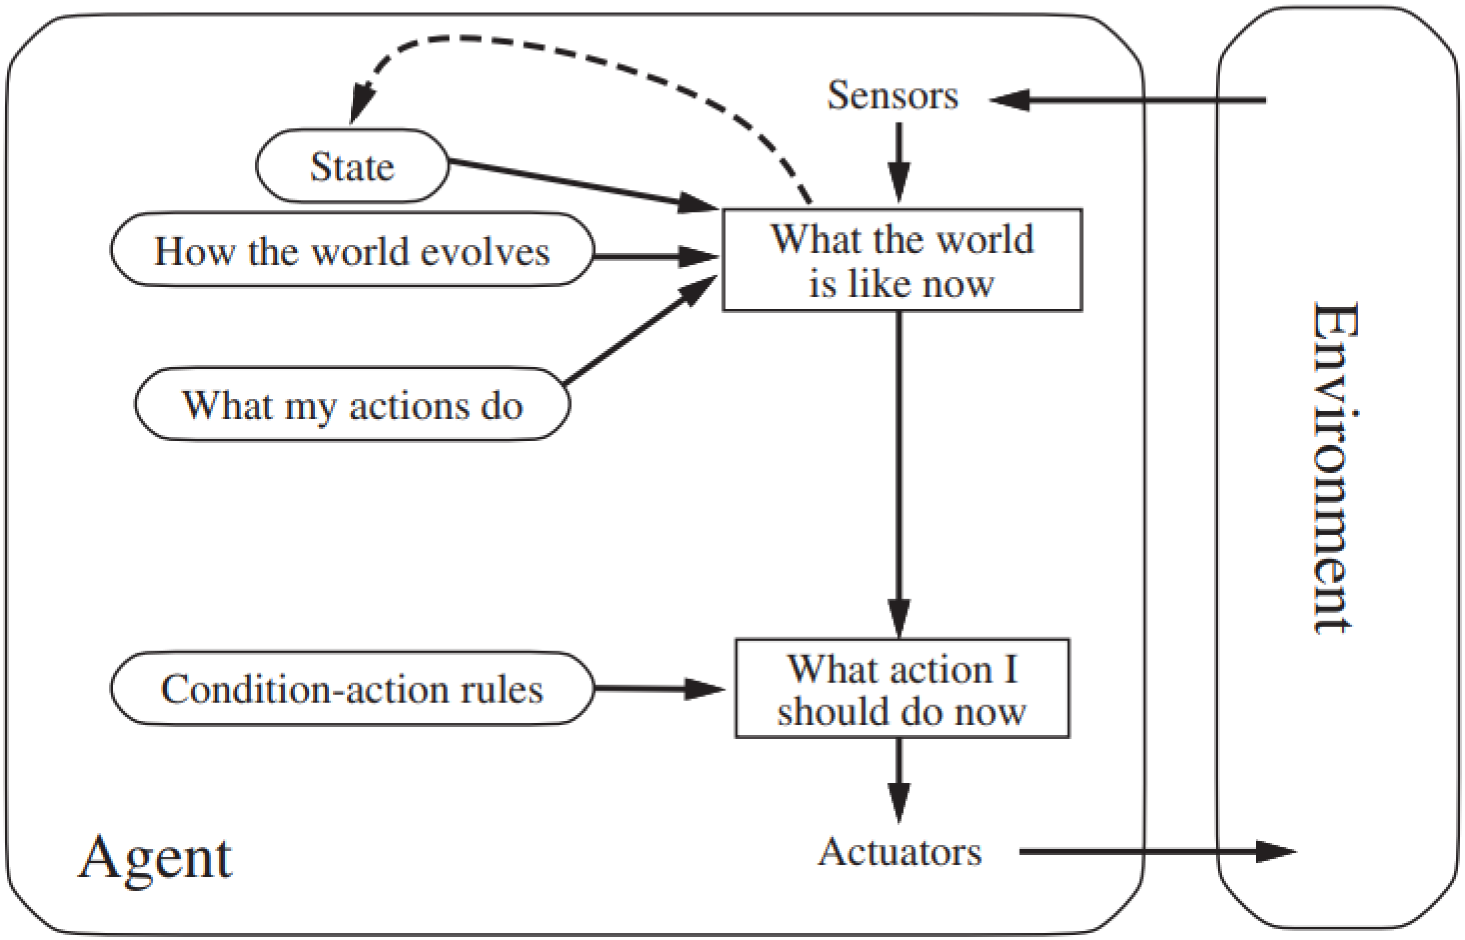

Agent Program

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*

3. Goal-based agents

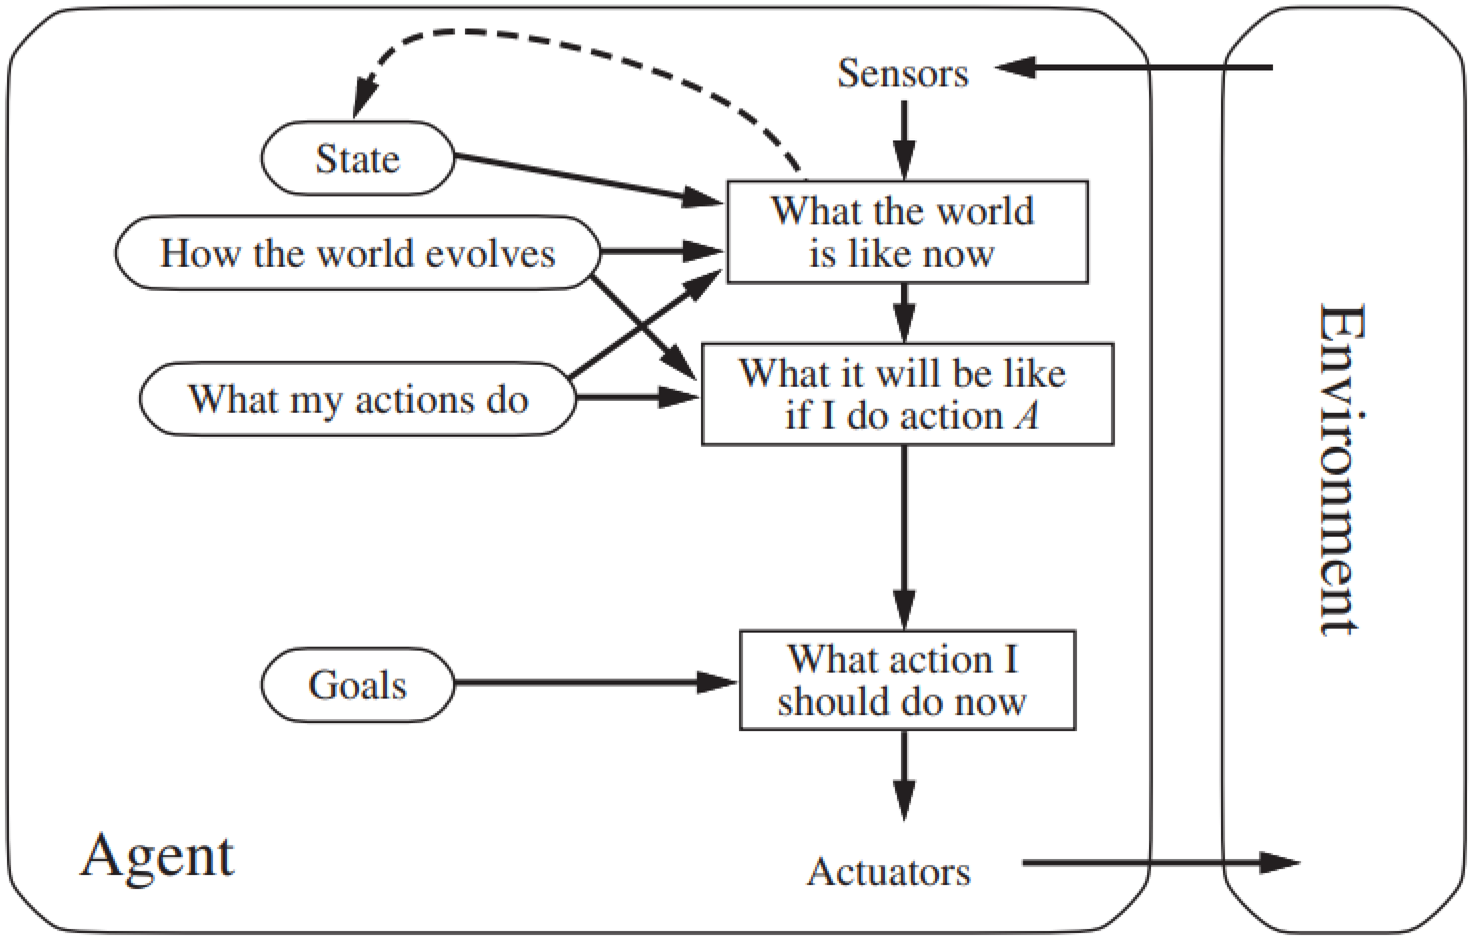

4. Utility based Agent

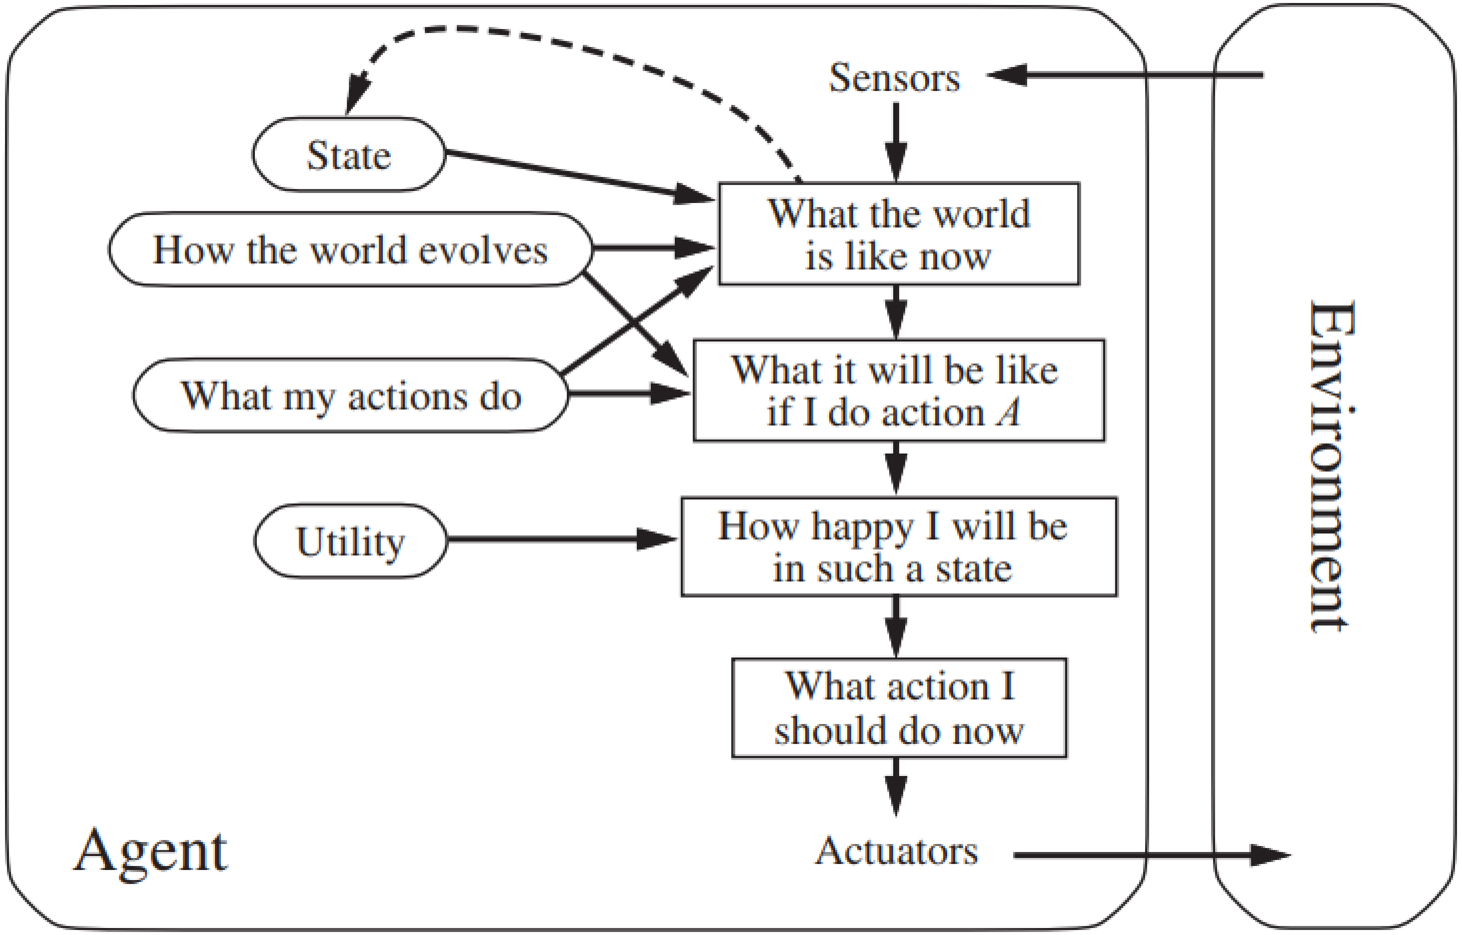

5. Learning agents

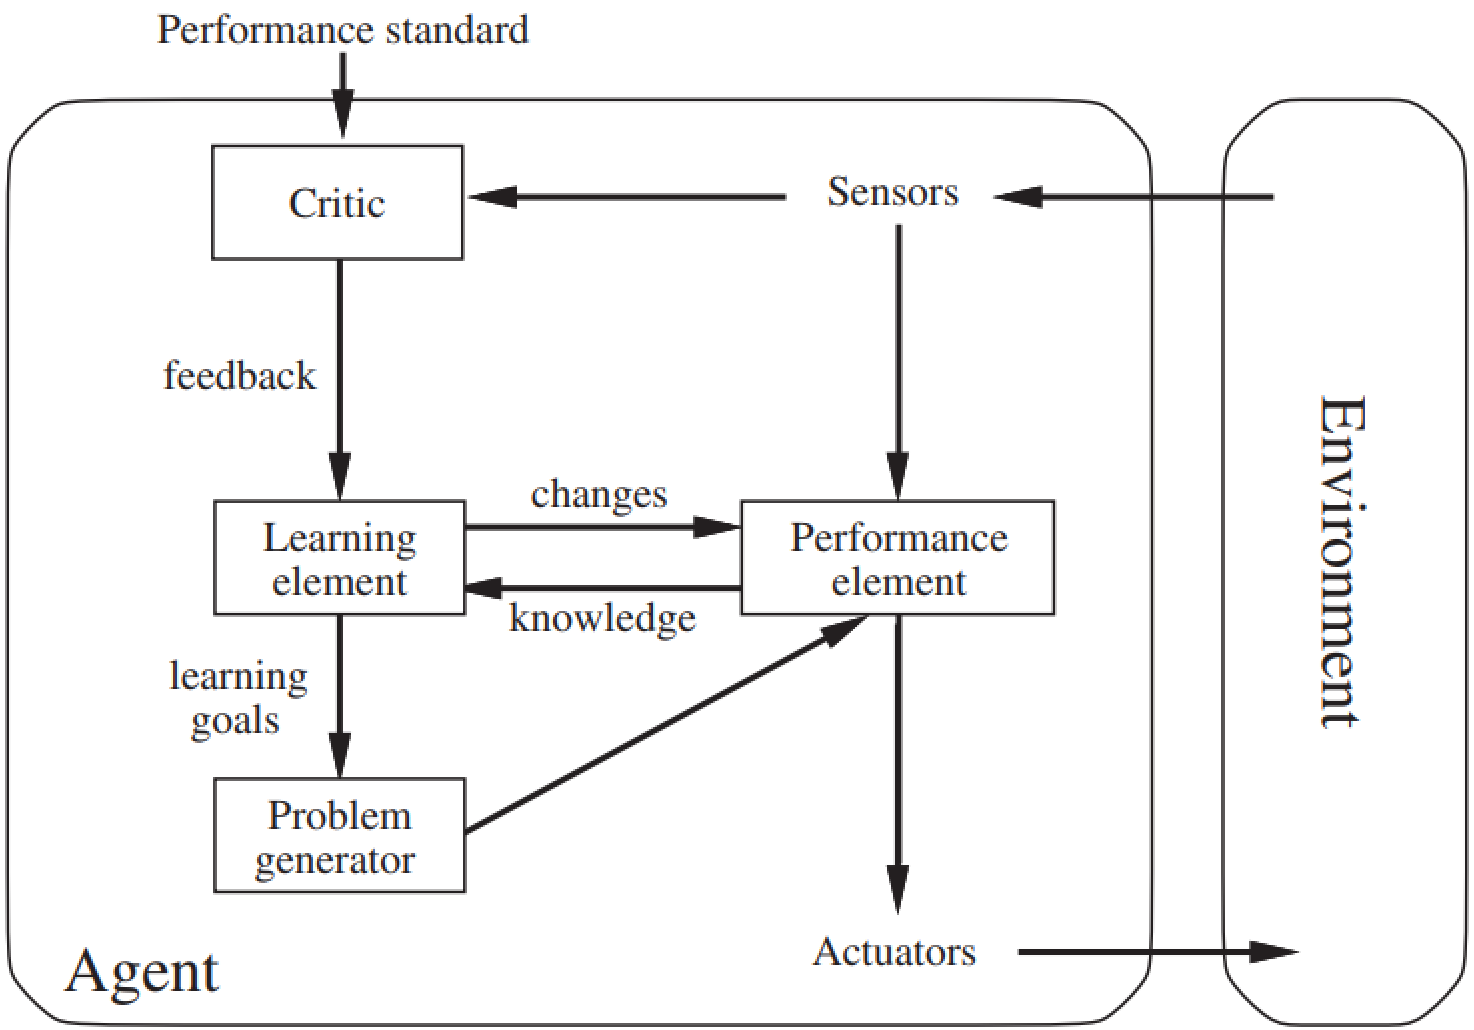

How the components of agent programs work ?

The representations of components of agents can be categorized along an axis of increasing complexity and expressive power: **atomic, factored, and structured**.

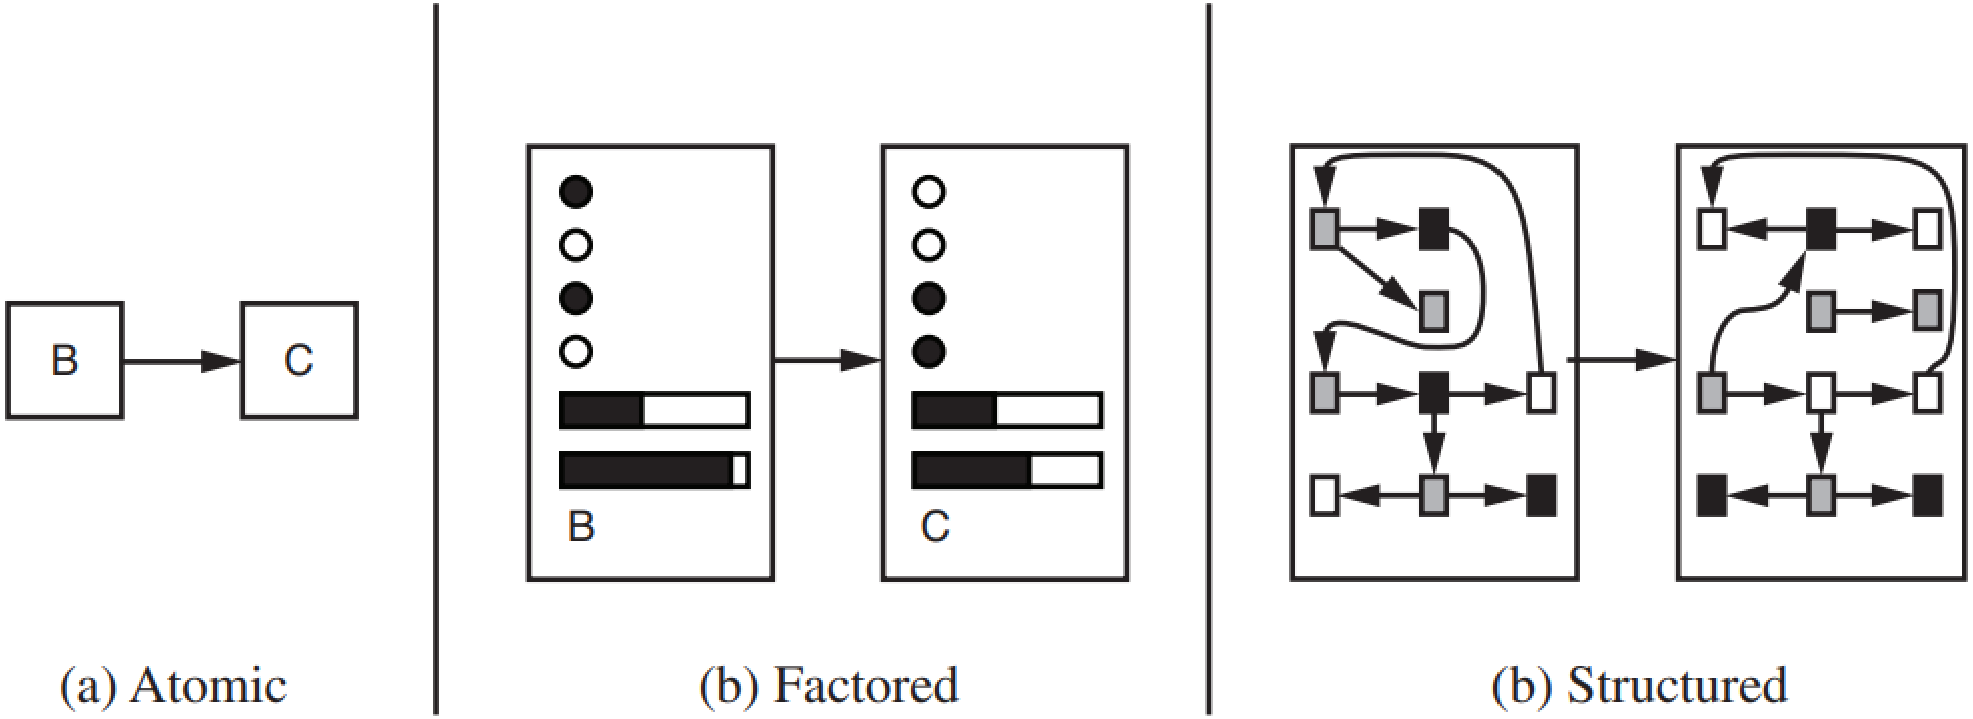

(a) Atomic (b) Factored (b) Structured

End of Module1