

Introduction to Artificial Intelligence

IFT3335 Lecture 2: Intelligent Agents

Bang Liu, Jian-Yun Nie

2 Plan

- Agents and environments
- Good behavior: the concept of rationality
- Task environment: PEAS (Performance measure, Environment, Actuators, Sensors)
- Different types of environments
- The structure of agents


Agents and Environments

4 Agents

- **Agent:** any entity that can **perceive** its **environment** via **sensors** and **act** on the environment via **actuators**
- **Human agent:**
 - eyes, ears, and other sensory organs;
 - hands, legs, mouth and other limbs as actuators
- **Robotic agent:**
 - cameras and infrared range finder as sensors;
 - different motors as actuators
- **Q:** describe some other agents: environment, sensor, actuator?

5 Agents and environments

- **Agent function:** maps percept sequence to an action
- **Agent program:** operates on a physical architecture to produce agent function f
- Agent = architecture + program



6 Example: the vacuum-cleaner world

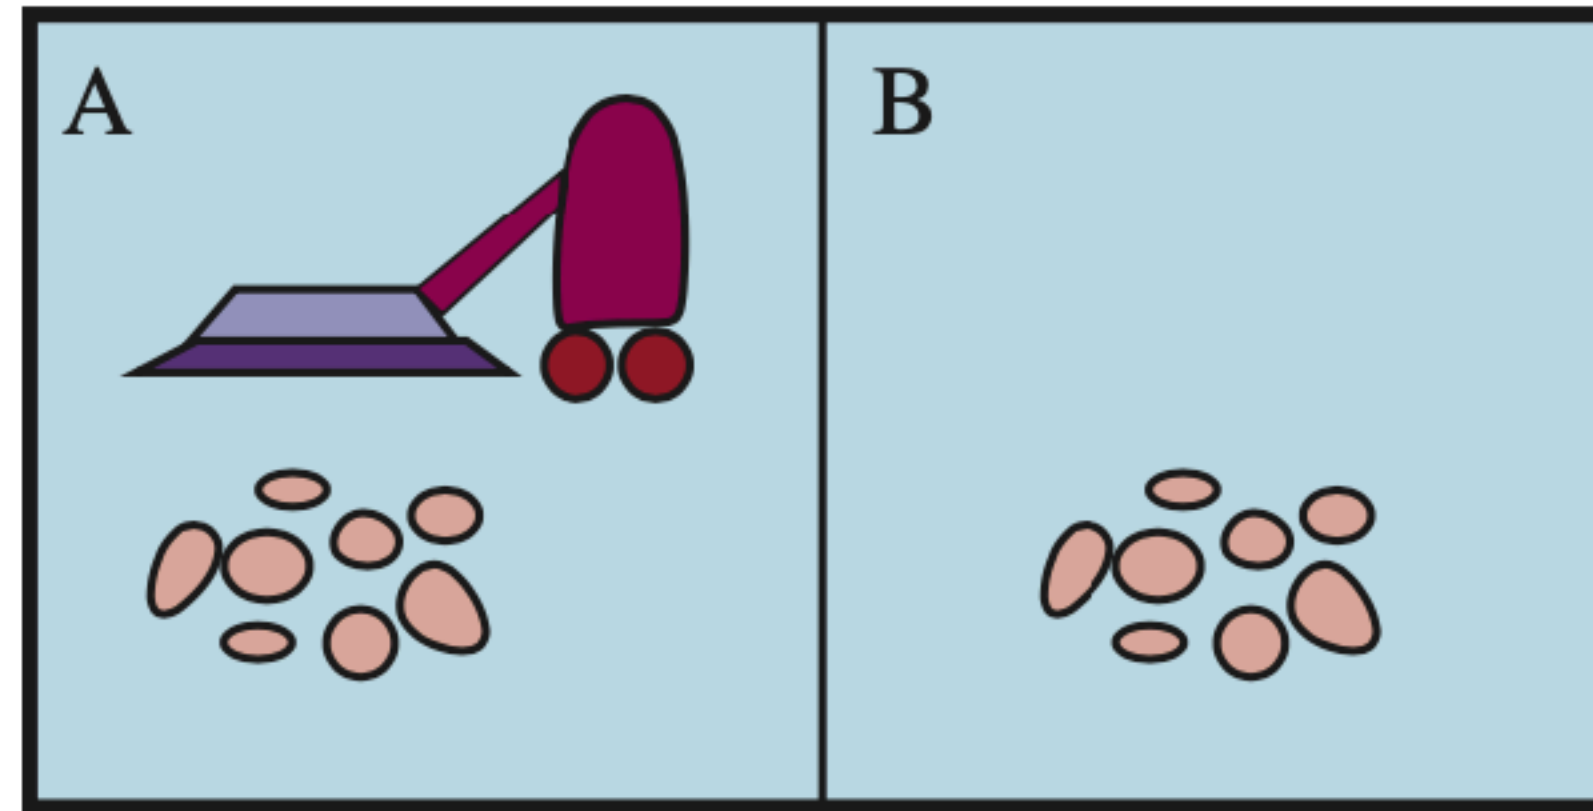


Figure 2.2 A vacuum-cleaner world with just two locations. Each location can be clean or dirty, and the agent can move left or right and can clean the square that it occupies. Different versions of the vacuum world allow for different rules about what the agent can perceive, whether its actions always succeed, and so on.

- Percepts: location and status, e.g., [A, dirty]
- Actions: Left, Right, Suck, NoOp

7 A vacuum-cleaner agent

Percept sequence	Action
<i>[A, Clean]</i>	<i>Right</i>
<i>[A, Dirty]</i>	<i>Suck</i>
<i>[B, Clean]</i>	<i>Left</i>
<i>[B, Dirty]</i>	<i>Suck</i>
<i>[A, Clean], [A, Clean]</i>	<i>Right</i>
<i>[A, Clean], [A, Dirty]</i>	<i>Suck</i>
<i>⋮</i>	<i>⋮</i>
<i>[A, Clean], [A, Clean], [A, Clean]</i>	<i>Right</i>
<i>[A, Clean], [A, Clean], [A, Dirty]</i>	<i>Suck</i>
<i>⋮</i>	<i>⋮</i>

Figure 2.3 Partial tabulation of a simple agent function for the vacuum-cleaner world shown in Figure 2.2. The agent cleans the current square if it is dirty, otherwise it moves to the other square. Note that the table is of unbounded size unless there is a restriction on the length of possible percept sequences.

- What is the **right** function?
- Can it be implemented in a small agent program?

Good Behaviour: The Concept of Rationality

9 Rational agent

- ⊙ A rational agent does the **right thing**.
 - What does it mean to do the right thing?
- ⊙ **Performance measurement:** An objective criterion of success for agent behavior
 - E.g., performance measure for a vacuum cleaner can be the volume of dirt picked up, the time taken, the electricity consumed, the noise produced, etc.
- ⊙ **Definition of a 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.

Task Environment: PEAS

11 Task environment: PEAS

- In designing an agent, first specify the task environment
- **PEAS: Performance, Environment, Actuators, Sensors**


● E.g.,

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver 	?	?	?	?

Task environment: PEAS

- In designing an agent, first specify the task environment
- **PEAS: Performance, Environment, Actuators, Sensors**

● E.g.,

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver 	Safe, fast, legal, comfortable trip, maximize profits, minimize impact on other road users	Roads, other traffic, police, pedestrians, customers, weather	Steering, accelerator, brake, signal, horn, display, speech	Cameras, radar, speedometer, GPS, engine sensors, accelerometer, microphones, touchscreen

Task environment: PEAS

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of questions, tests, diagnoses, treatments	Touchscreen/voice entry of symptoms and findings
Satellite image analysis system	Correct categorization of objects, terrain	Orbiting satellite, downlink, weather	Display of scene categorization	High-resolution digital camera
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arm and hand	Camera, tactile and joint angle sensors
Refinery controller	Purity, yield, safety	Refinery, raw materials, operators	Valves, pumps, heaters, stirrers, displays	Temperature, pressure, flow, chemical sensors
Interactive English tutor	Student's score on test	Set of students, testing agency	Display of exercises, feedback, speech	Keyboard entry, voice

Figure 2.5 Examples of agent types and their PEAS descriptions.

Different Types of Environments

15 Types of environment

◎ Fully observable v.s. partially observable


- Whether an agent's sensors give it access to the complete state of the environment (that are relevant to the choice of action) at each point in time



Types of environment

● Single agent v.s. multi agent

- Q: driving car, single agent or multi agent?
- Whether B's behaviour is best described as maximizing a performance measure whose value depends on A's behaviour. E.g., chess - competitive; driving - partially cooperative (drivers all avoid collisions).



17 Types of environment

⦿ **Deterministic v.s. nondeterministic**


- Whether the next state of the environment is completely determined by the current state and the action executed by the agent(s)



18 Types of environment

◎ Episodic v.s. sequential

- In an episodic task environment, the agent's experience is divided into atomic episodes. In each episode the agent receives a percept and then performs a single action. Crucially, the next episode does not depend on the actions taken in previous episodes.



19 Types of environment

◎ Static v.s. dynamic

- If the environment can change while an agent is deliberating, then we say the environment is dynamic for that agent; otherwise, it is static.



if play with a clock: **semidynamic**
(environment doesn't change, but
performance score changes with time)



20 Types of environment

◎ Discrete v.s. continuous

- A number of distinct and clearly defined states and actions.



Types of 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	?					
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
English tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

Types of 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
English tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

The Structure of Agents

Agent functions and programs

- An agent is specified completely by the agent function which maps the sequence of perceived to actions
- An agent function (or a small equivalence class) is rational
- Goal: find a way to implement the rational agent function in a concise way
- **Table-driven agent:** huge table, long time to build/search, hard to learn, ...

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*)

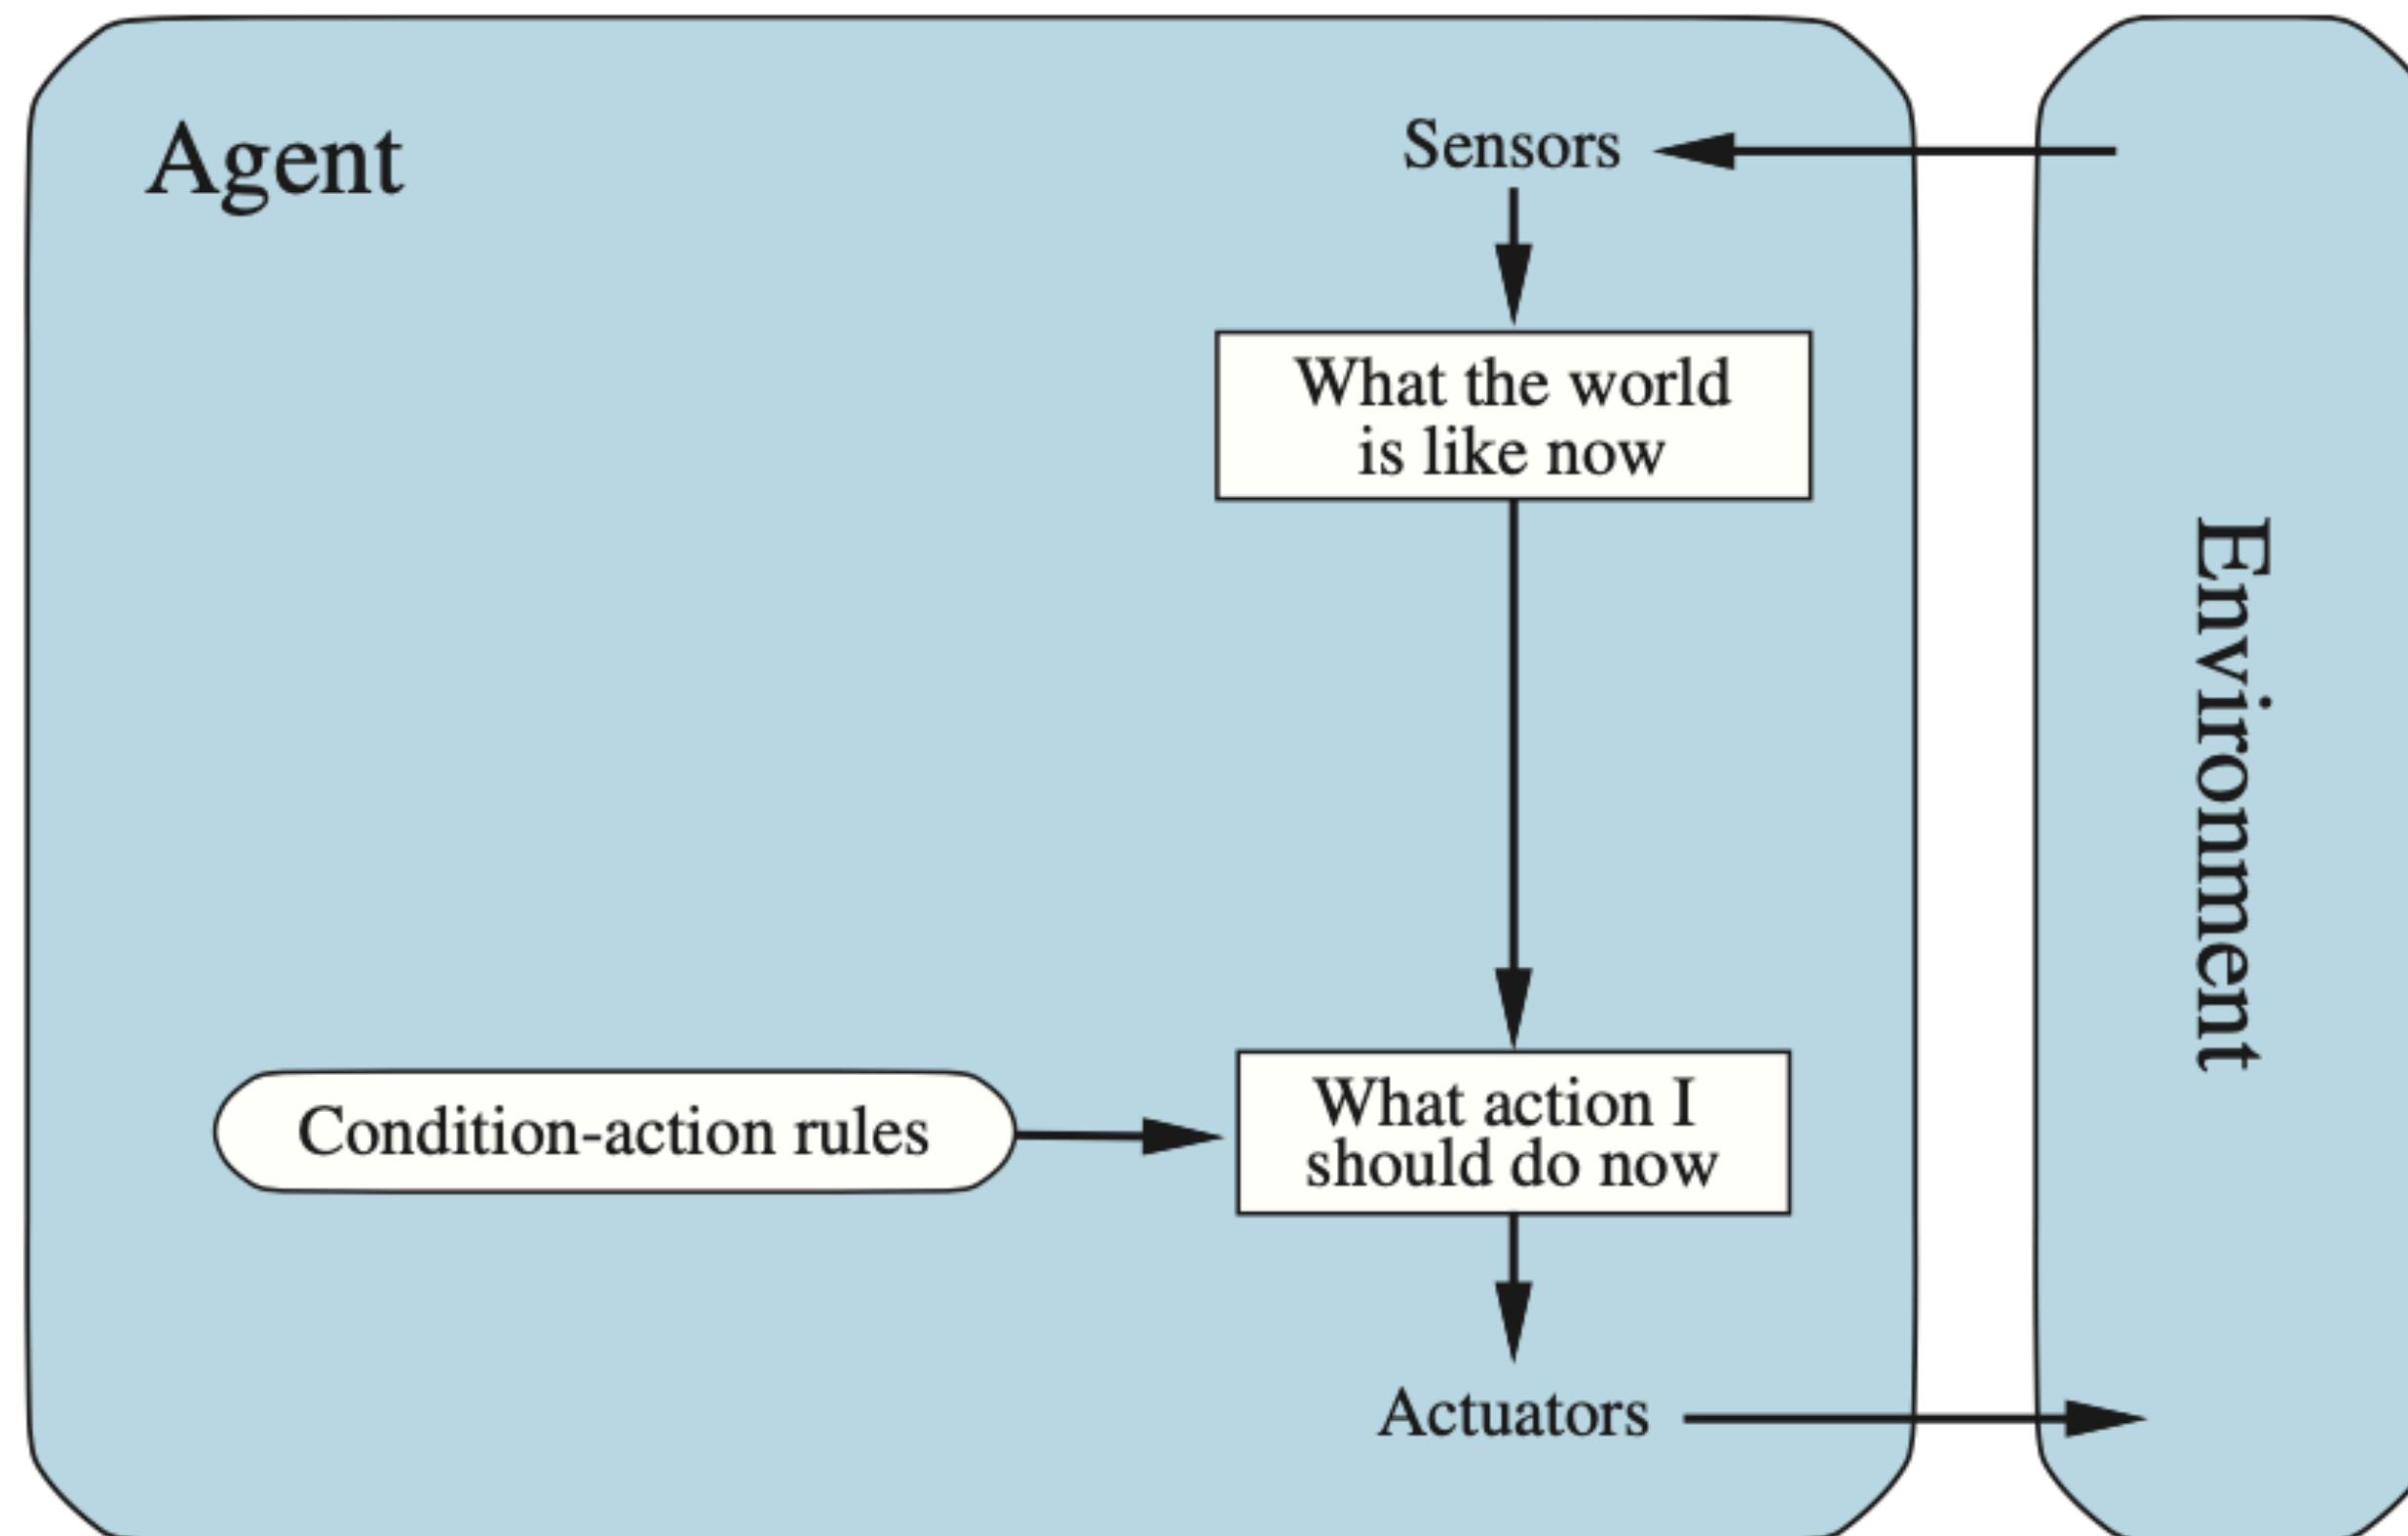
return *action*

Types of the agent

- Simple reflex agents
- Model-based reflex agents
- Goal-based agents
- Utility-based agents

Simple reflex agents

- Current percept
- Ignoring percept history
- Simple
- E.g., Vacuum environment



Work only if the env is fully observable

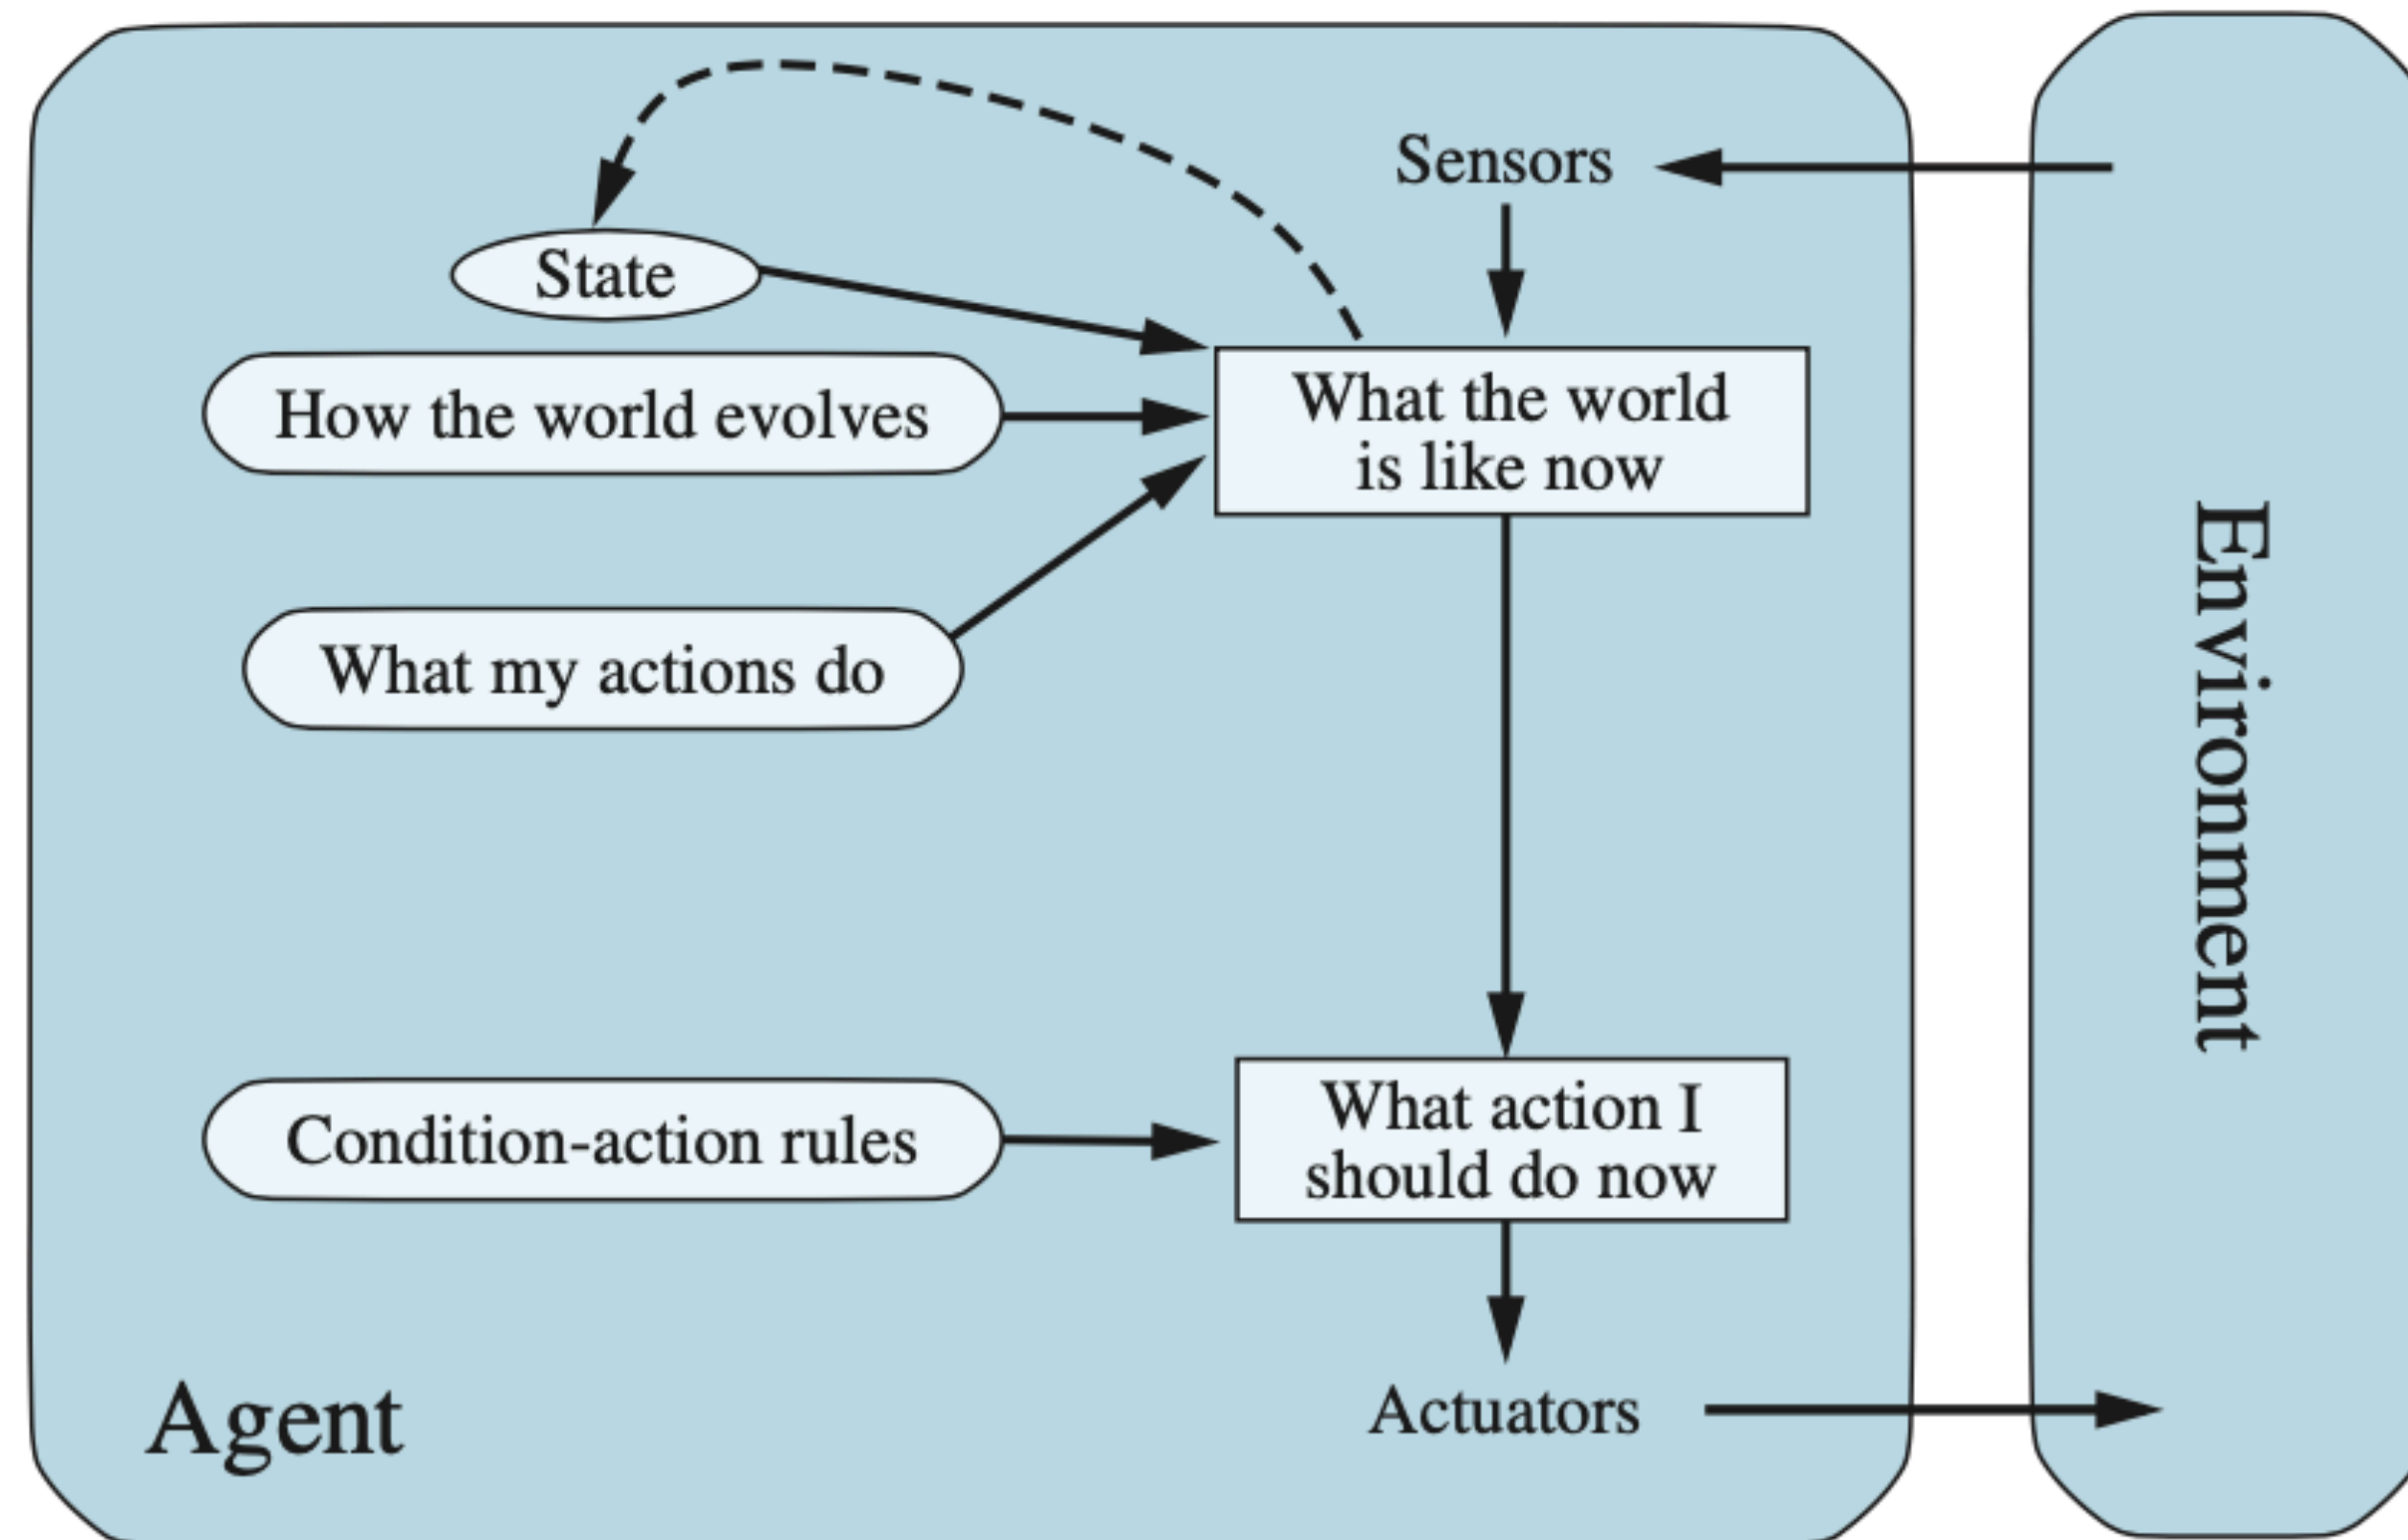
What if partially observable?

27 Model-based reflex agents

Internal state: keep track of the unseen world

Transition model: how the world works

Sensor model: how the world state is reflected in percepts

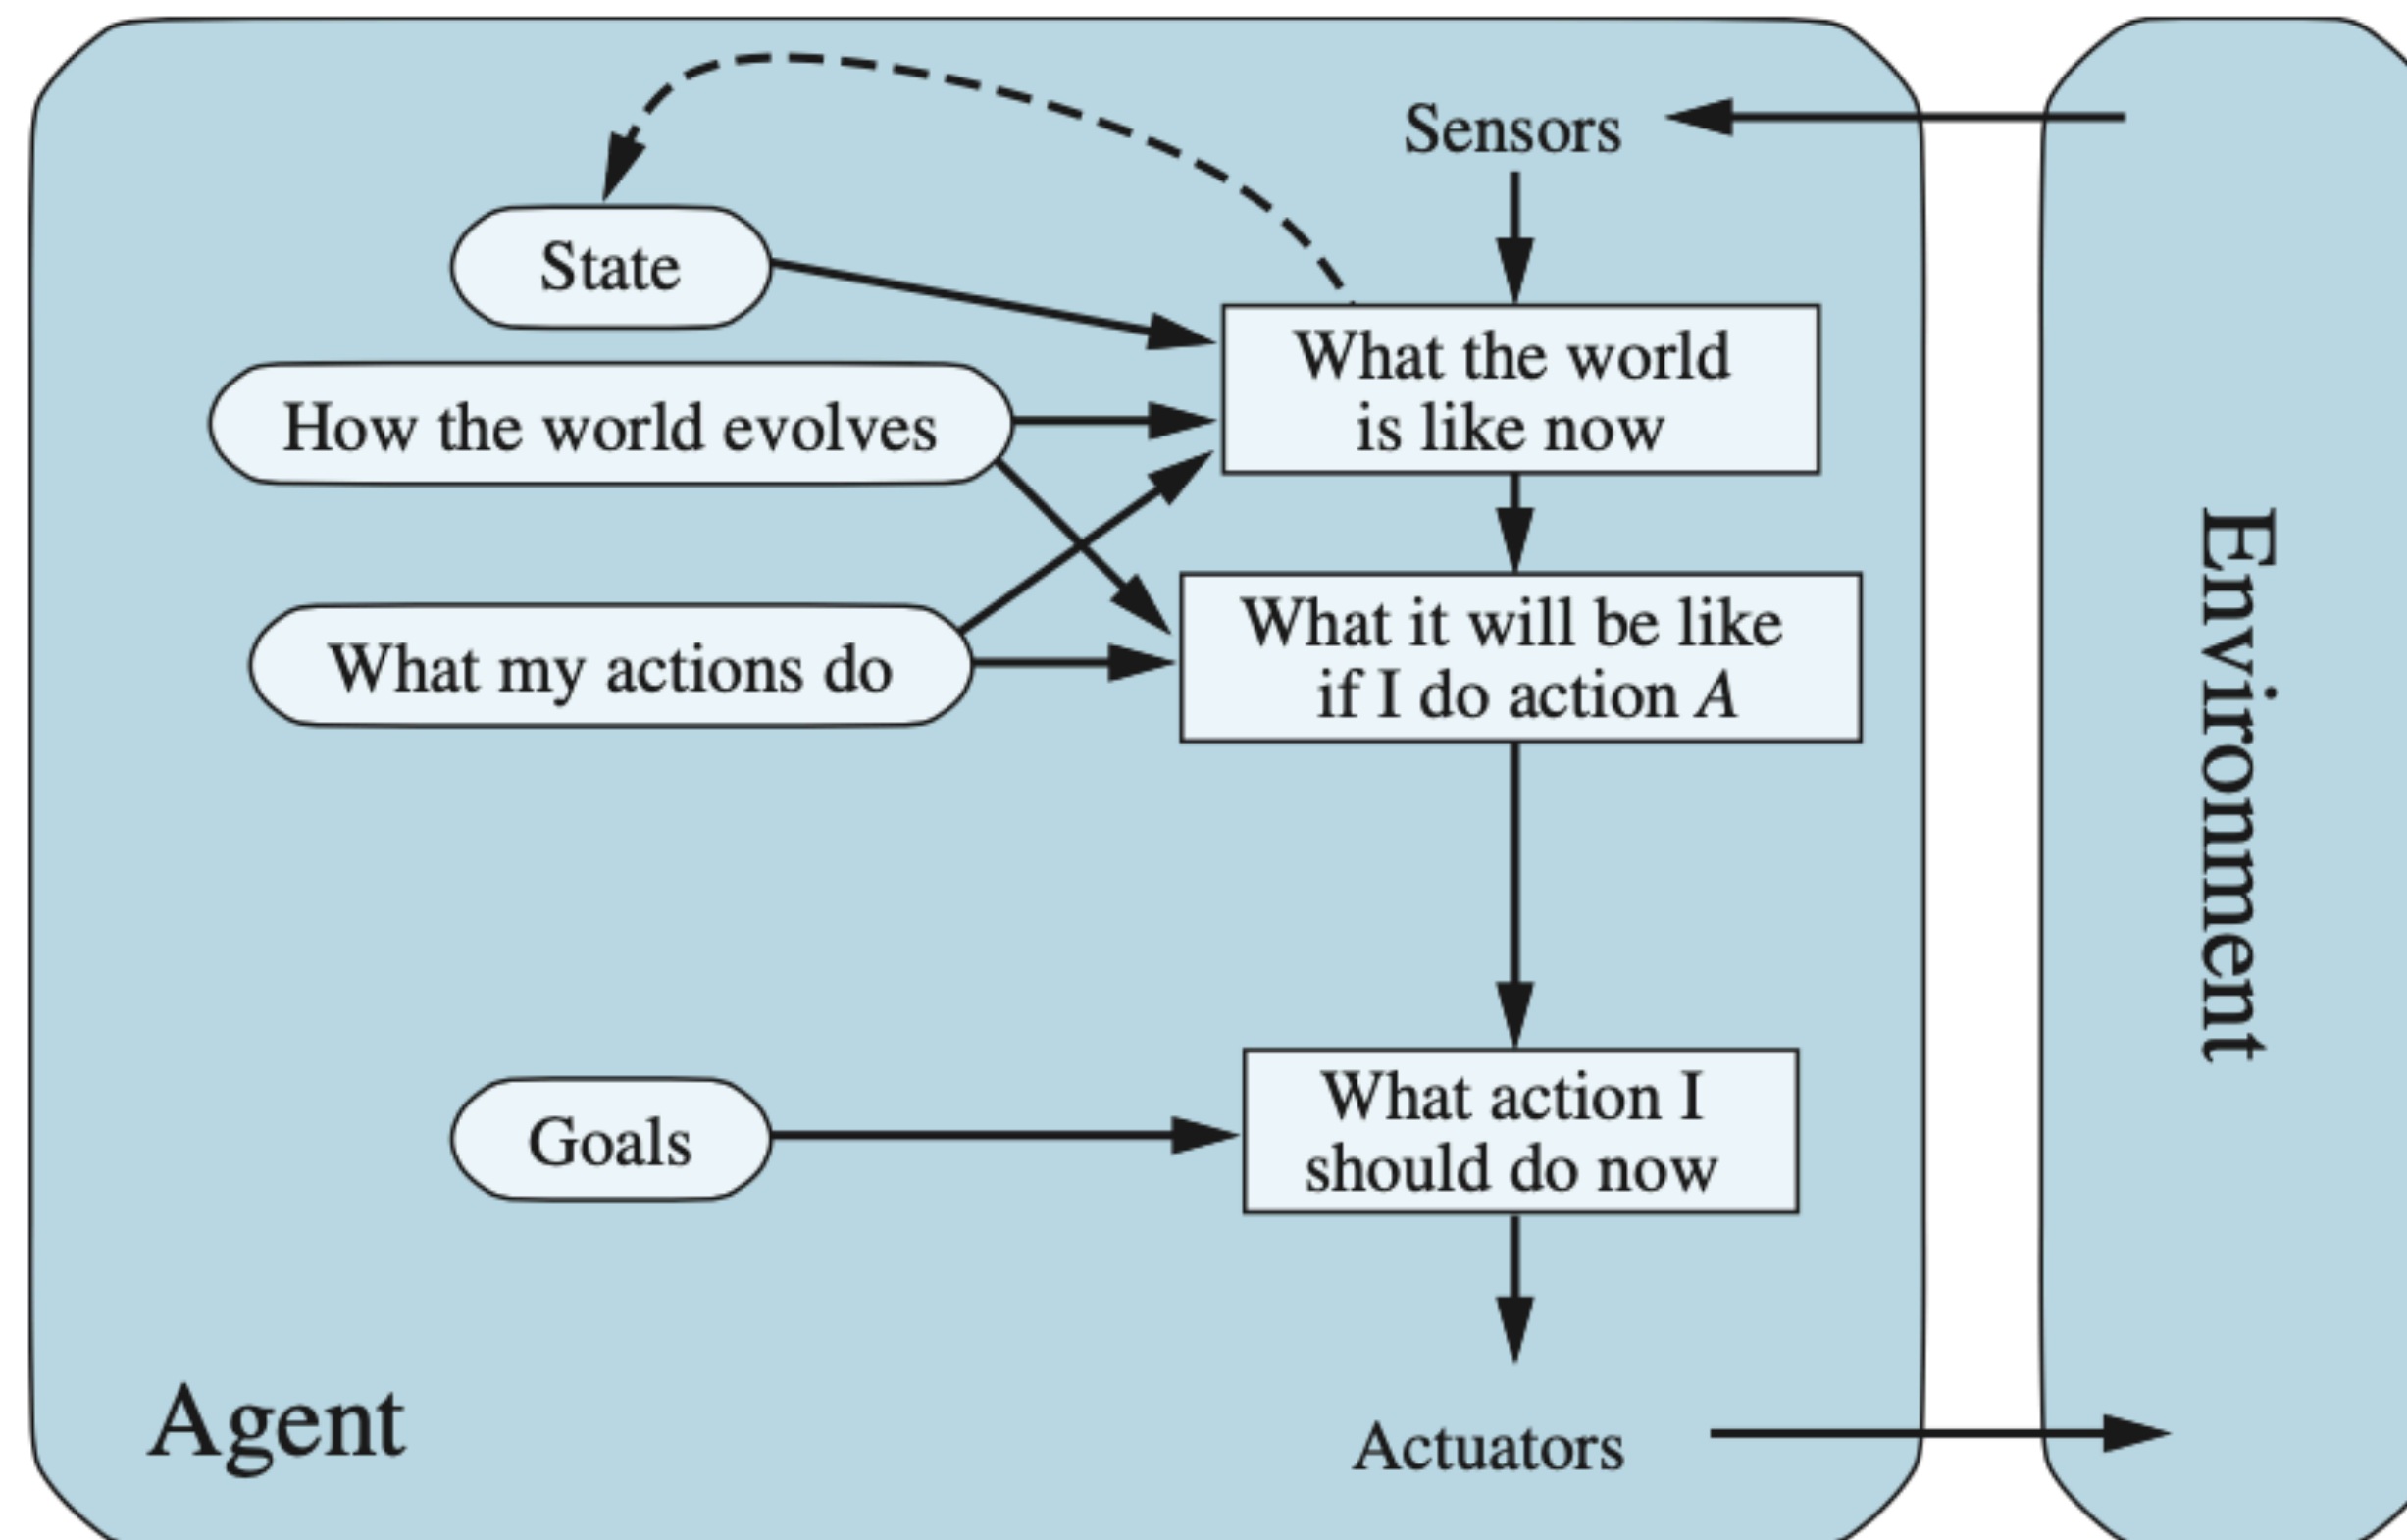


Model-based agent

Knowing something about the current environment is not enough. What is the goal?

Goal-based agents

Knowledge supports decision



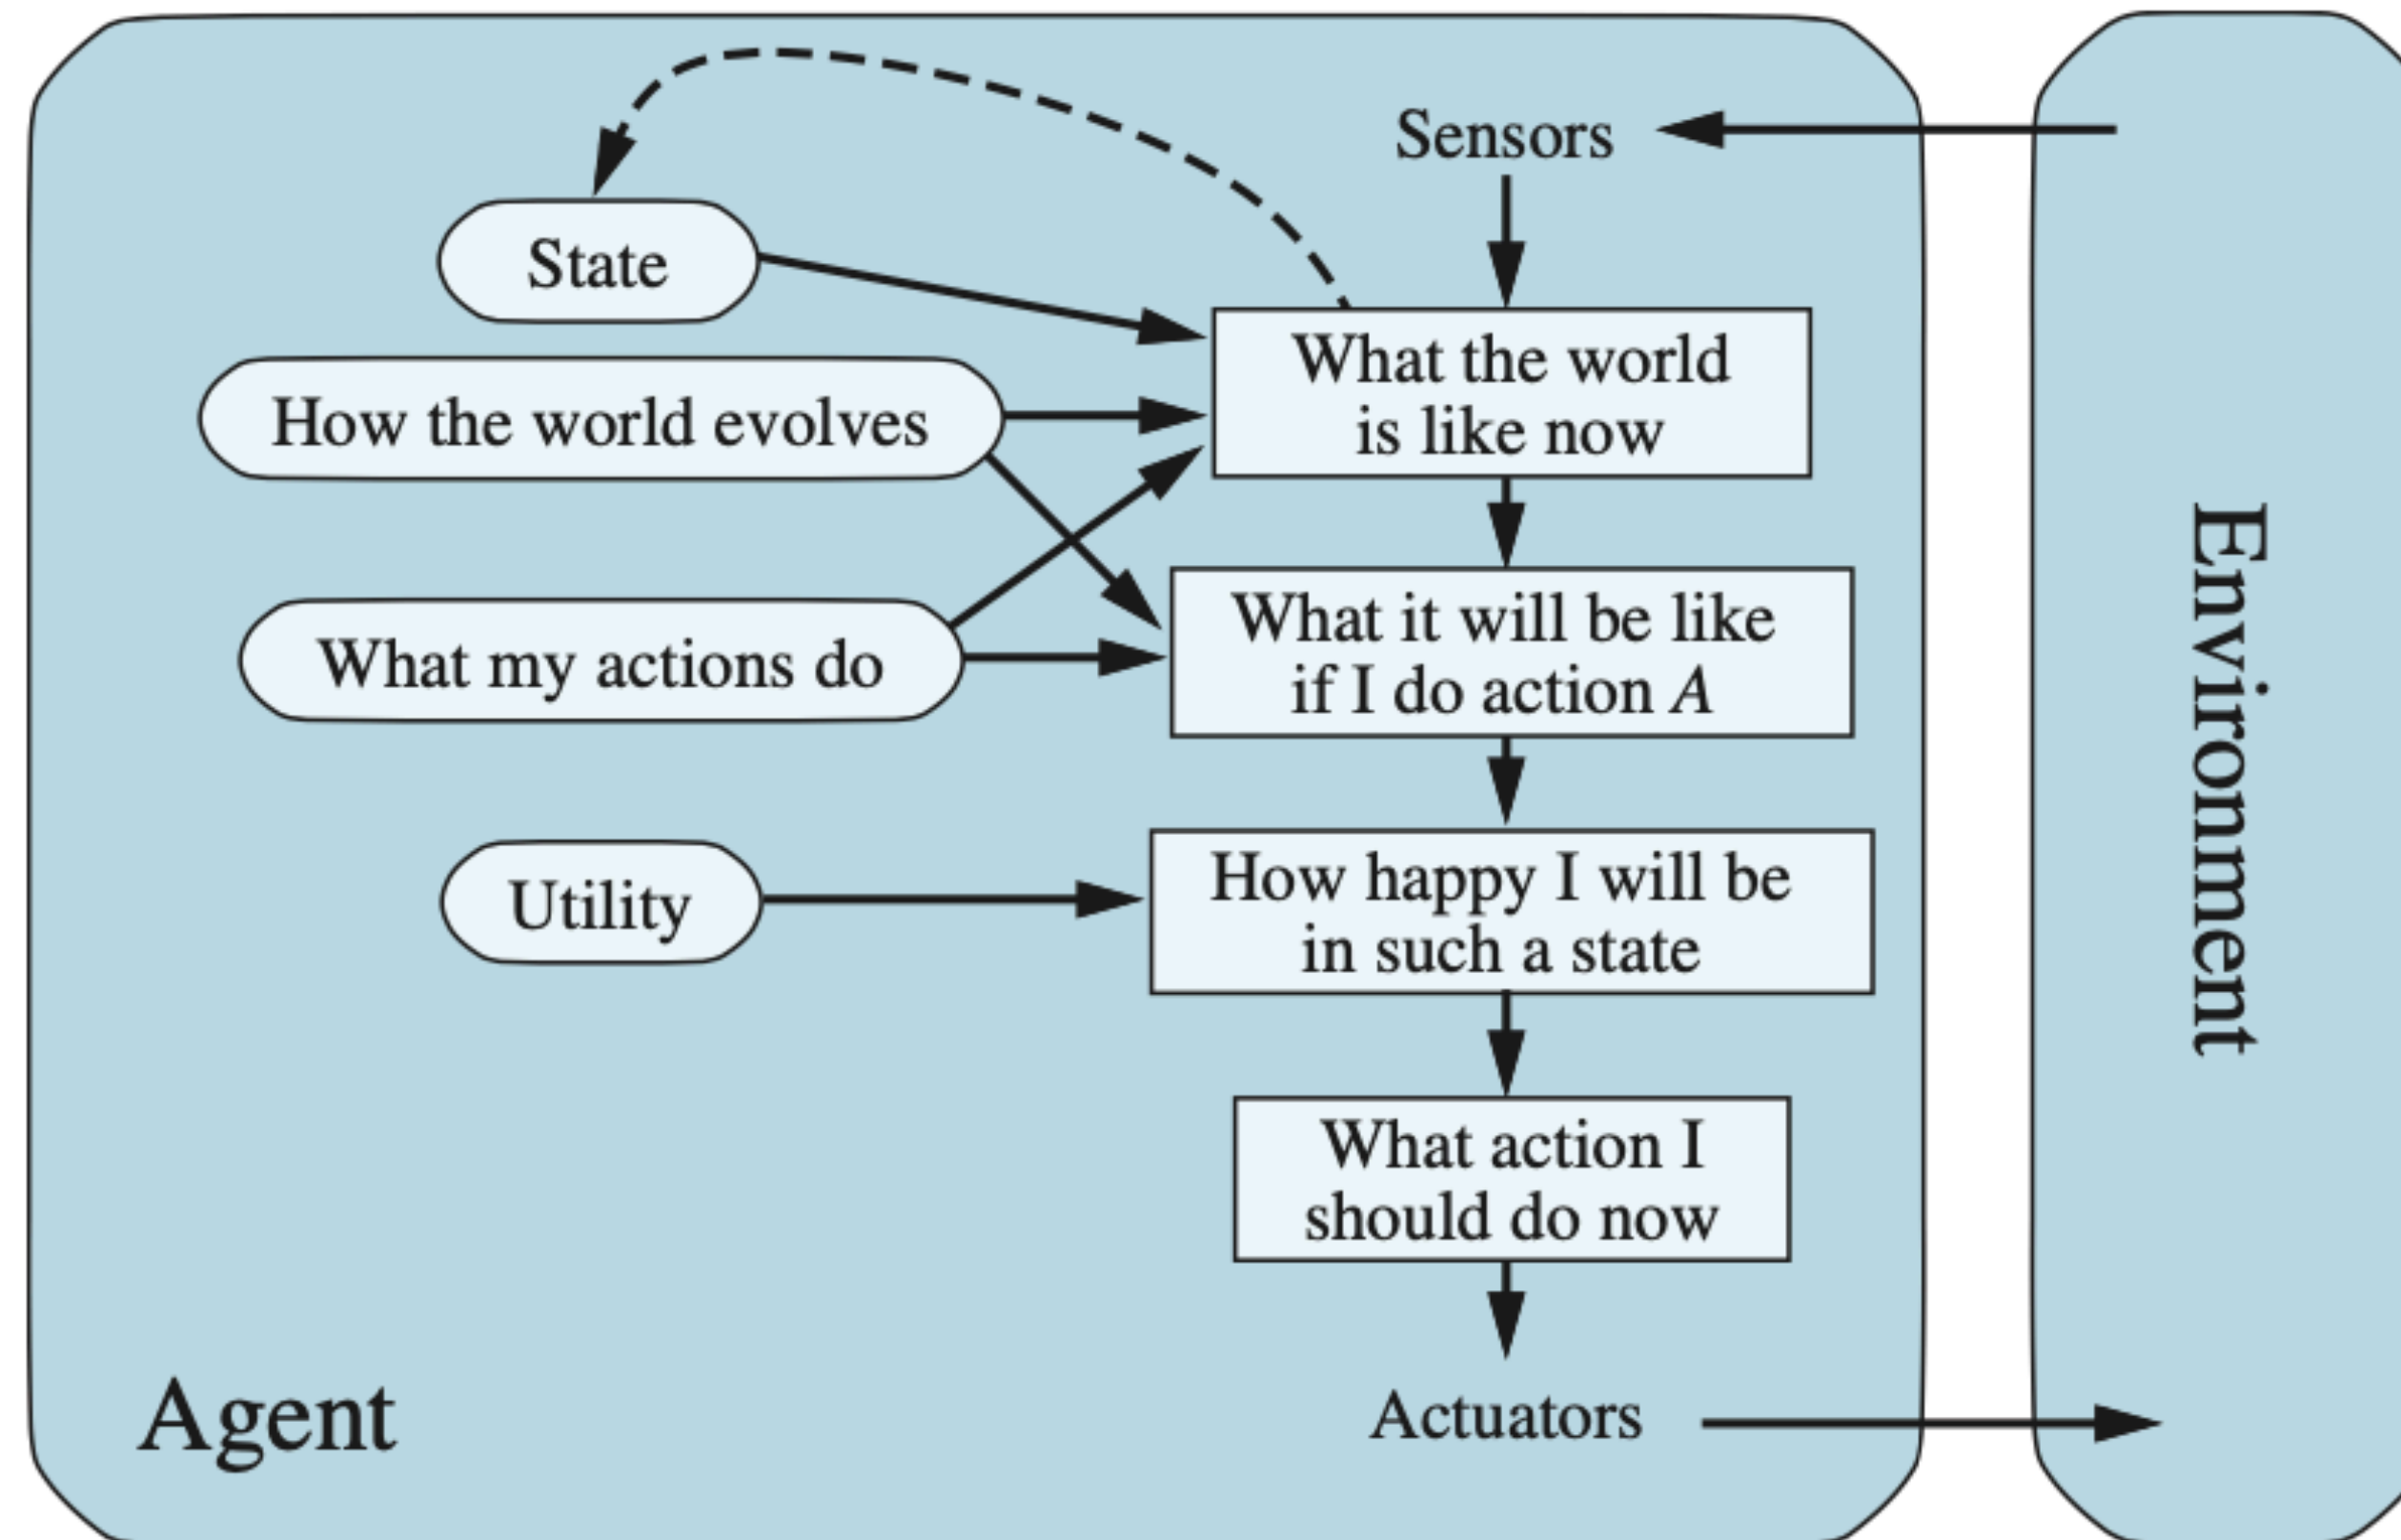
Can have different ways to achieve the same goal. E.g., traveling.

How to generate high-quality behaviour?

Utility-based agents

Utility function: essentially an internalization of the performance measure

Utility



We have described various ways for selecting actions.

But how the agent programs come into being?

Build learning machines.

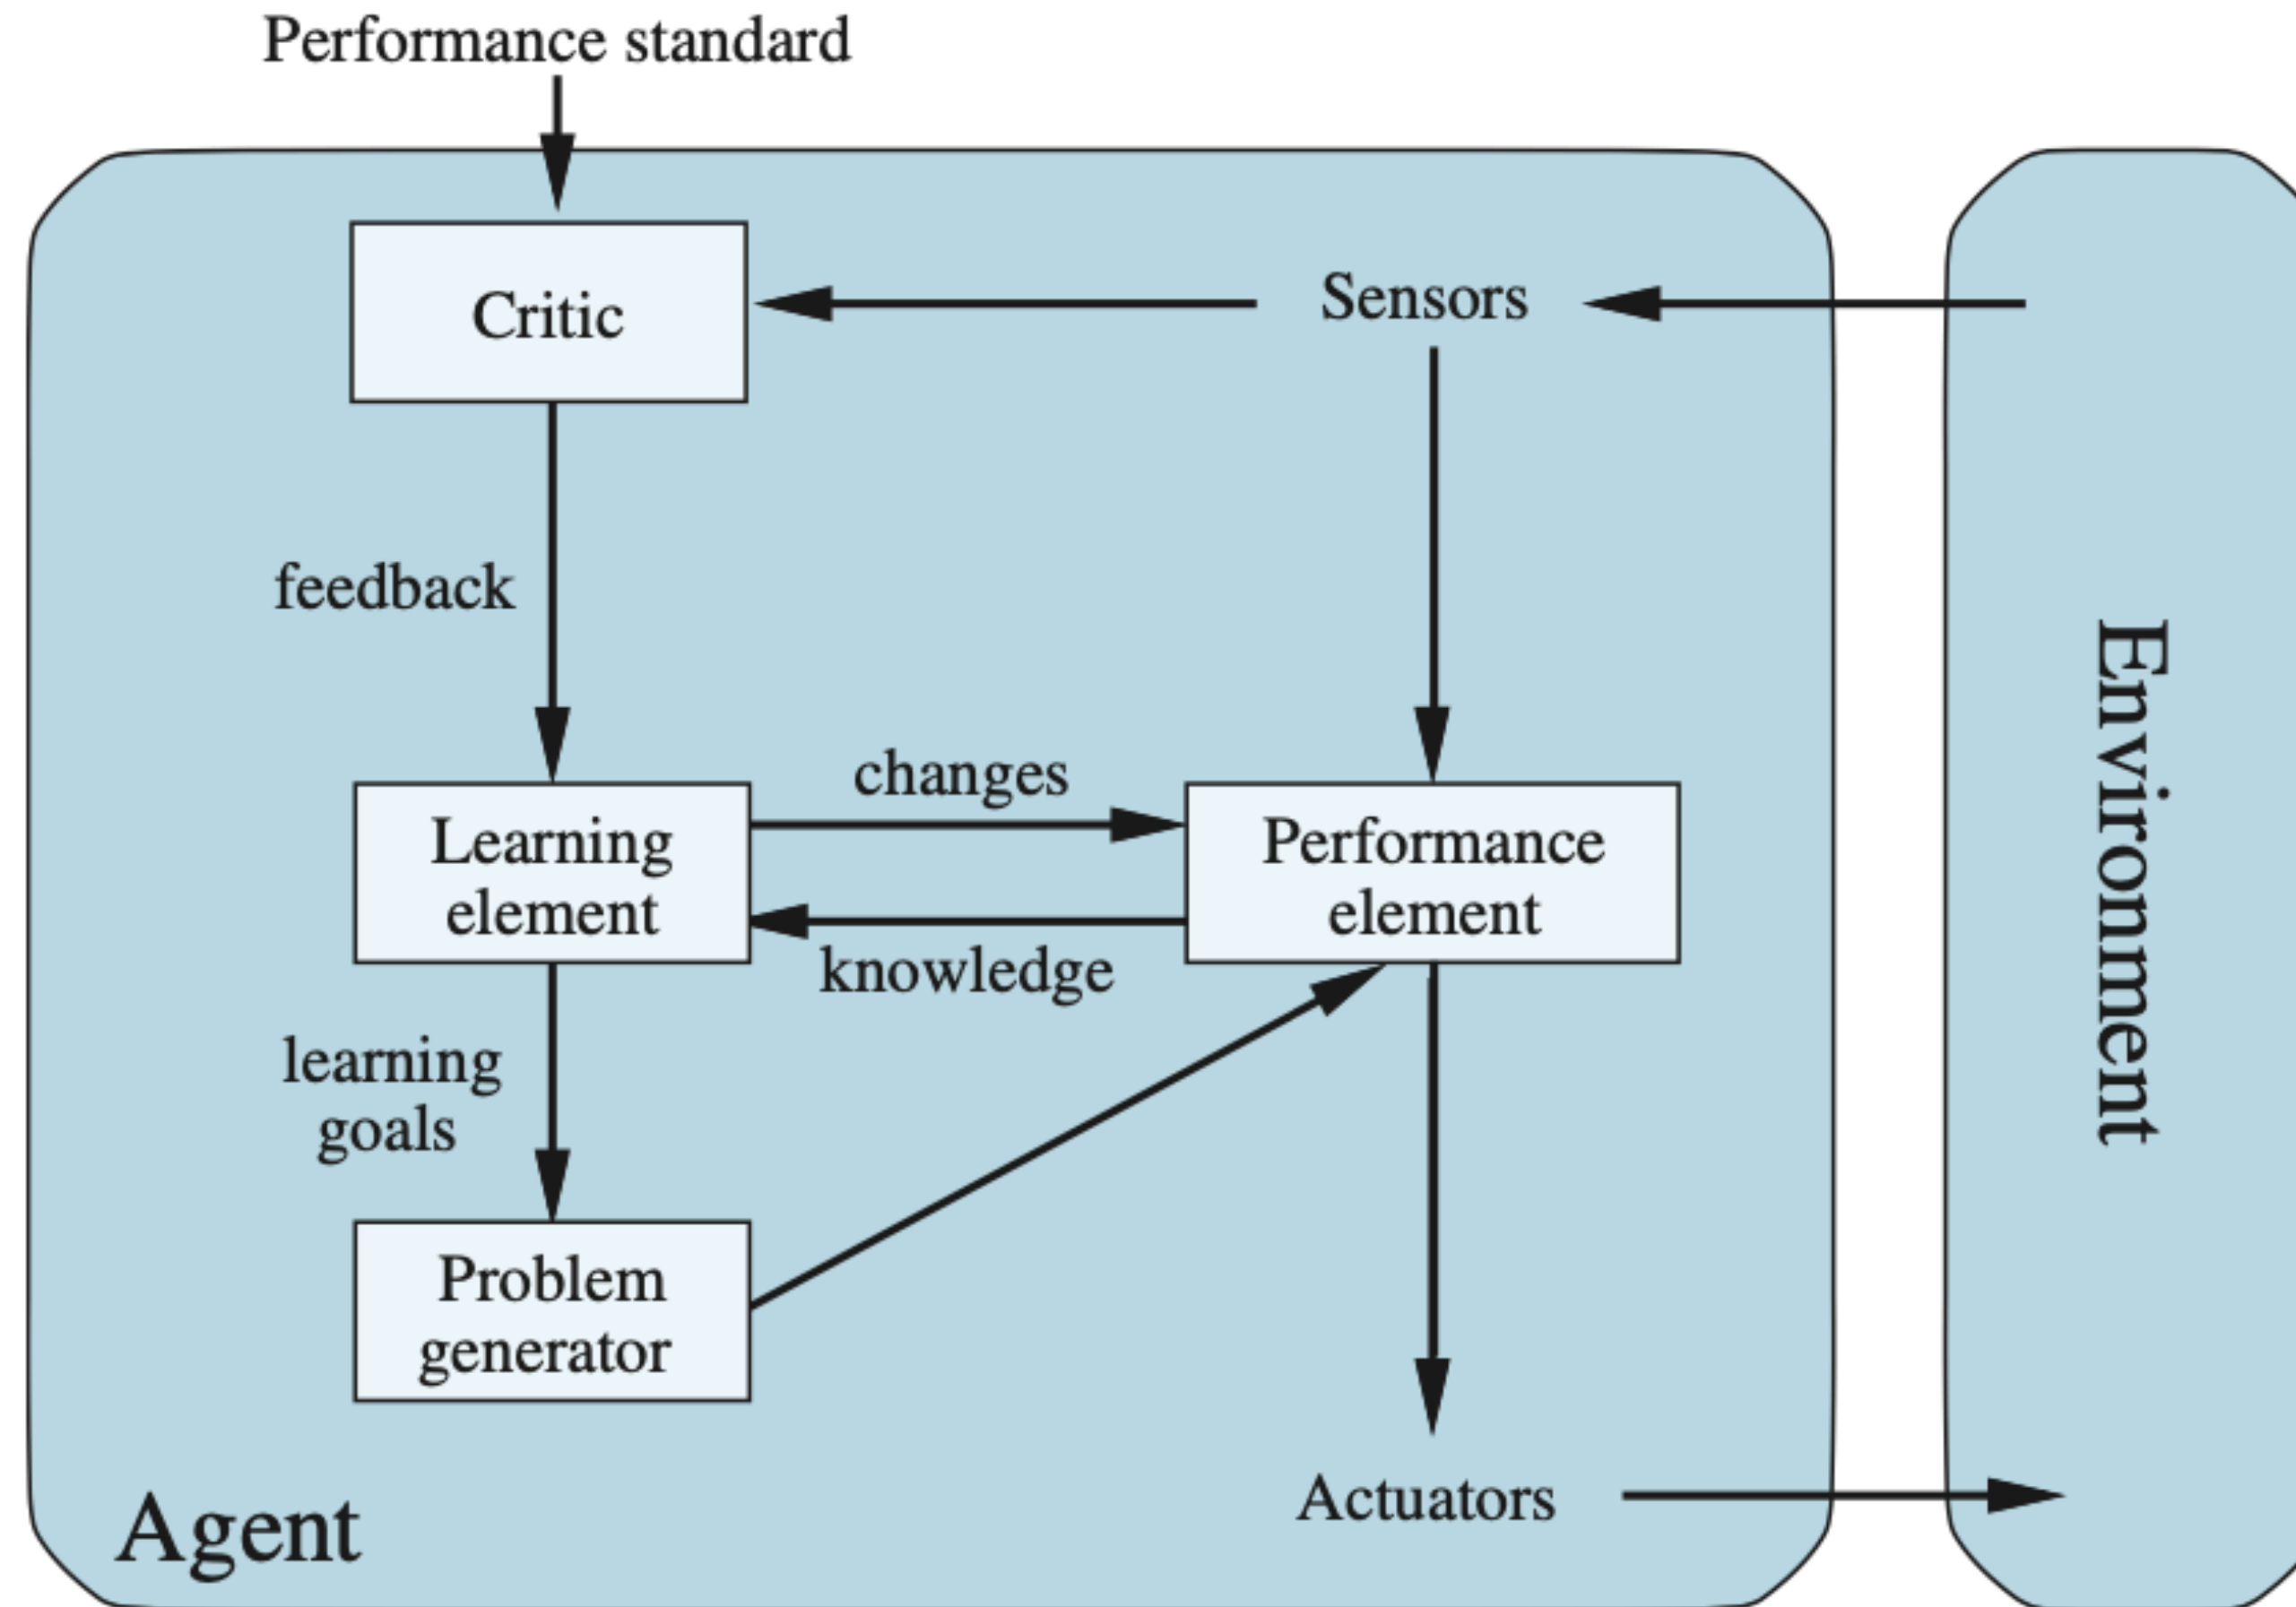
30 Learning agents

Critic: how well the agent is doing

Learning element: making improvements

Performance element: selecting actions

Problem generator: encouraging exploration



Reinforcement Learning!

Thanks! Q&A