

Intelligent Agents



Acting rationally: rational agent

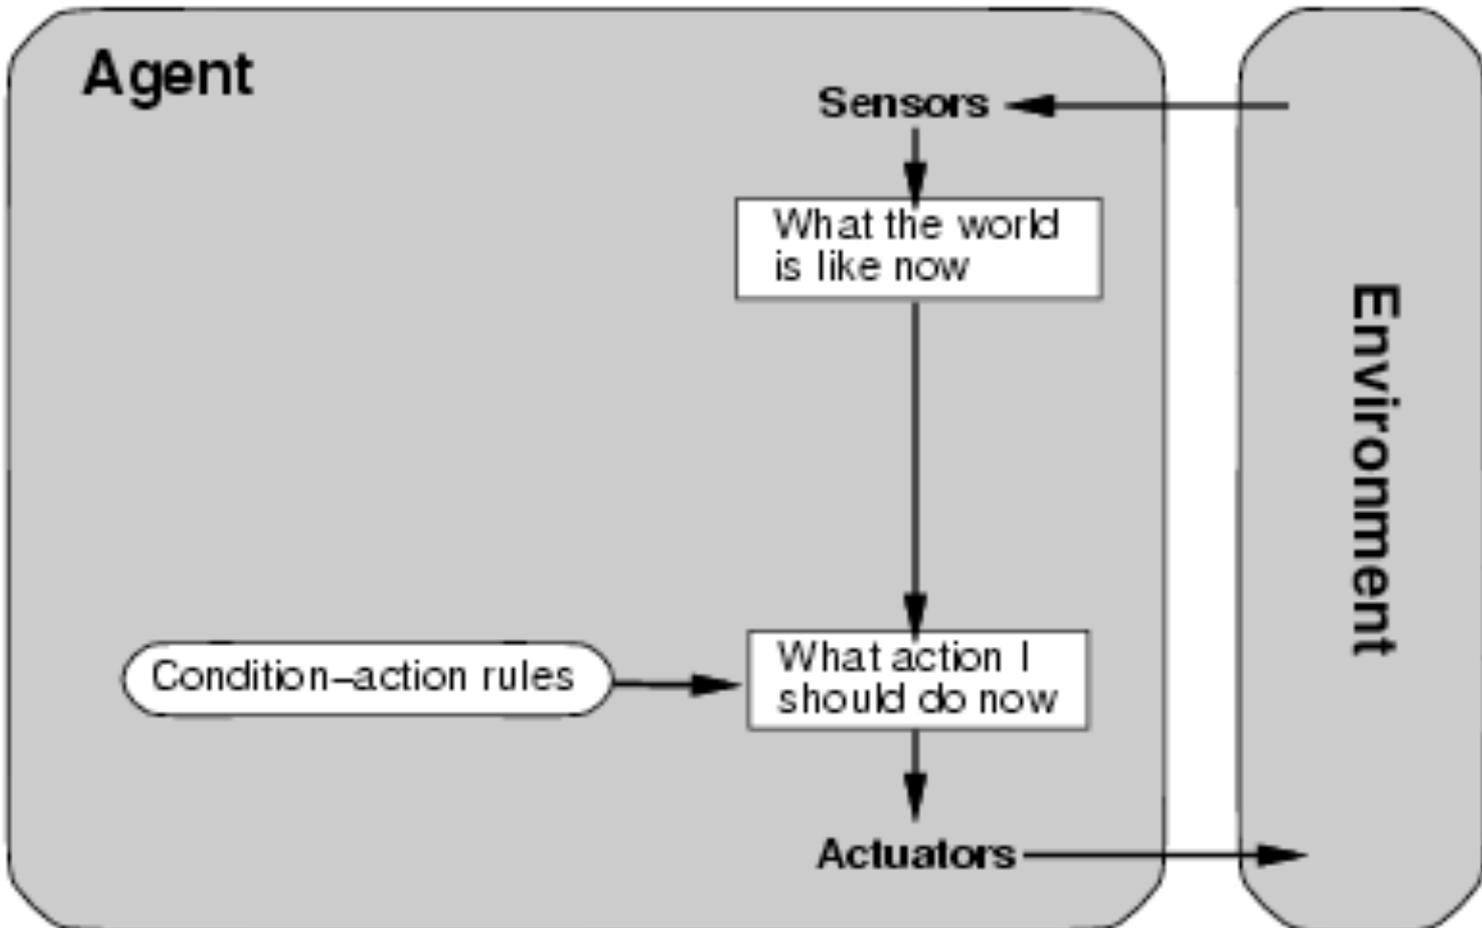
Agent: entity that perceives and acts

Rational behavior:

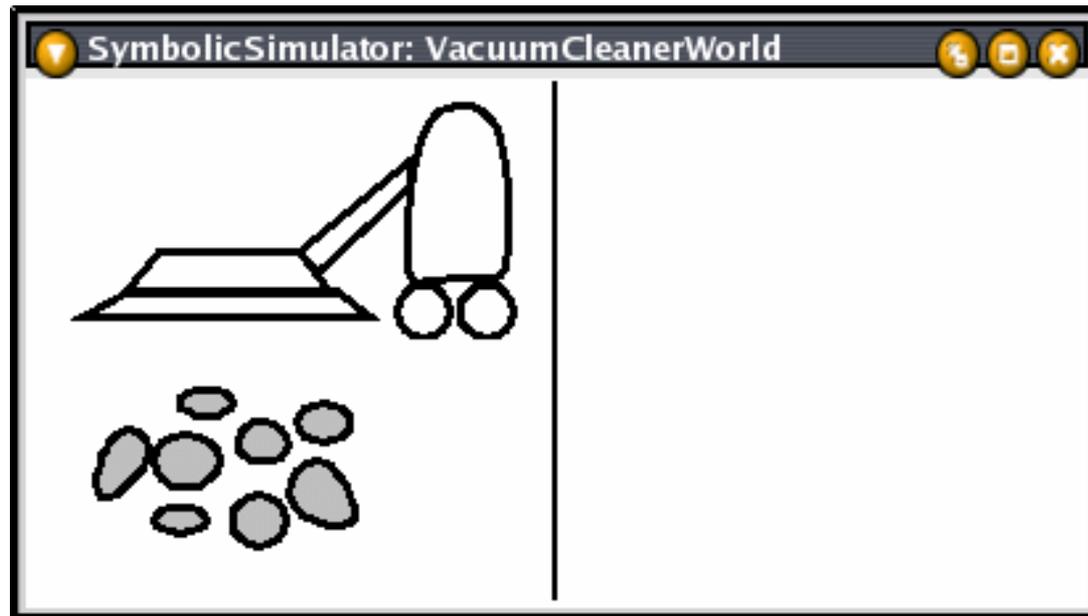
choose behavior that

maximize goal achievement, given the available information

Agent-Based AI

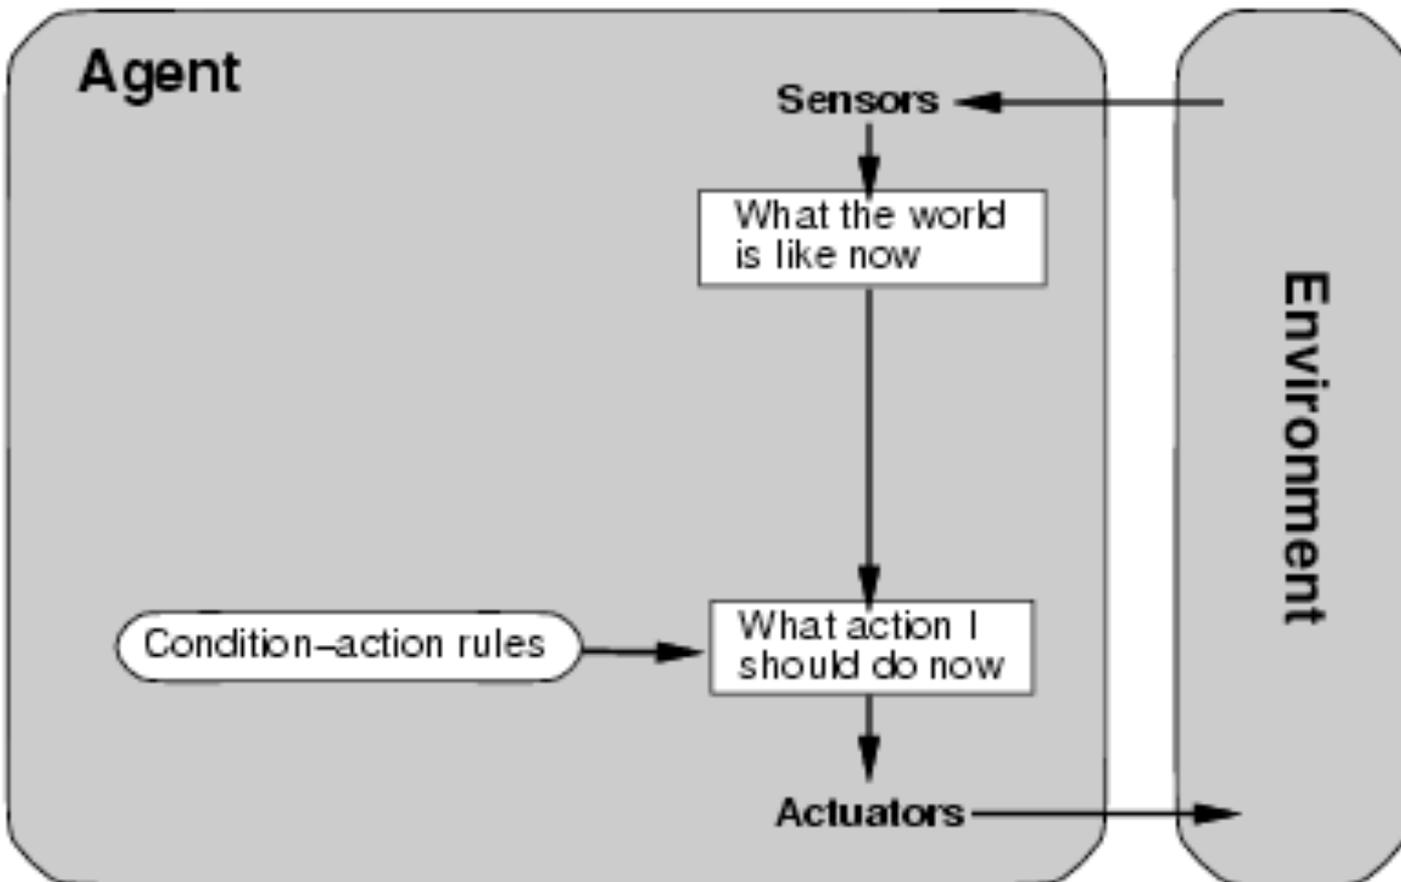


Example: Vacuum Cleaner World



What are the **actions**? What are the **percepts**?

Kinds of Agents: Simple Reflex Agent



Agent Design

- What can the agent do?
 - Range of **actions**
- What is the environment? (Input: **percepts**)
 - How is it **interpreted**?
- What does the agent know?
 - History of previous inputs and actions (how far back?)
 - Properties of environment: **world knowledge**
 - Knowledge of its own **goals** and preferences
 - **Strategies** for behavior
- How does the agent choose to act?
 - Mapping from percept sequence -> action called an **agent function**

Task Environment

Before we design an intelligent agent, we must specify its “task environment”:

PEAS:

Performance measure

Environment

Actuators

Sensors

Can you identify them for: Conversational Character, Medical diagnosis system, or NPC in FPS game?

➔ Example: Agent = Conversational Character

Performance measure:

Environment:

Actuators:

Sensors:

➔ Example: Agent = Medical diagnosis system

Performance measure:

Environment:

Actuators:

Sensors:

➔ Example: Agent = NPC in an FPS game

Performance measure:

Environment:

Actuators:

Sensors:

Environment types

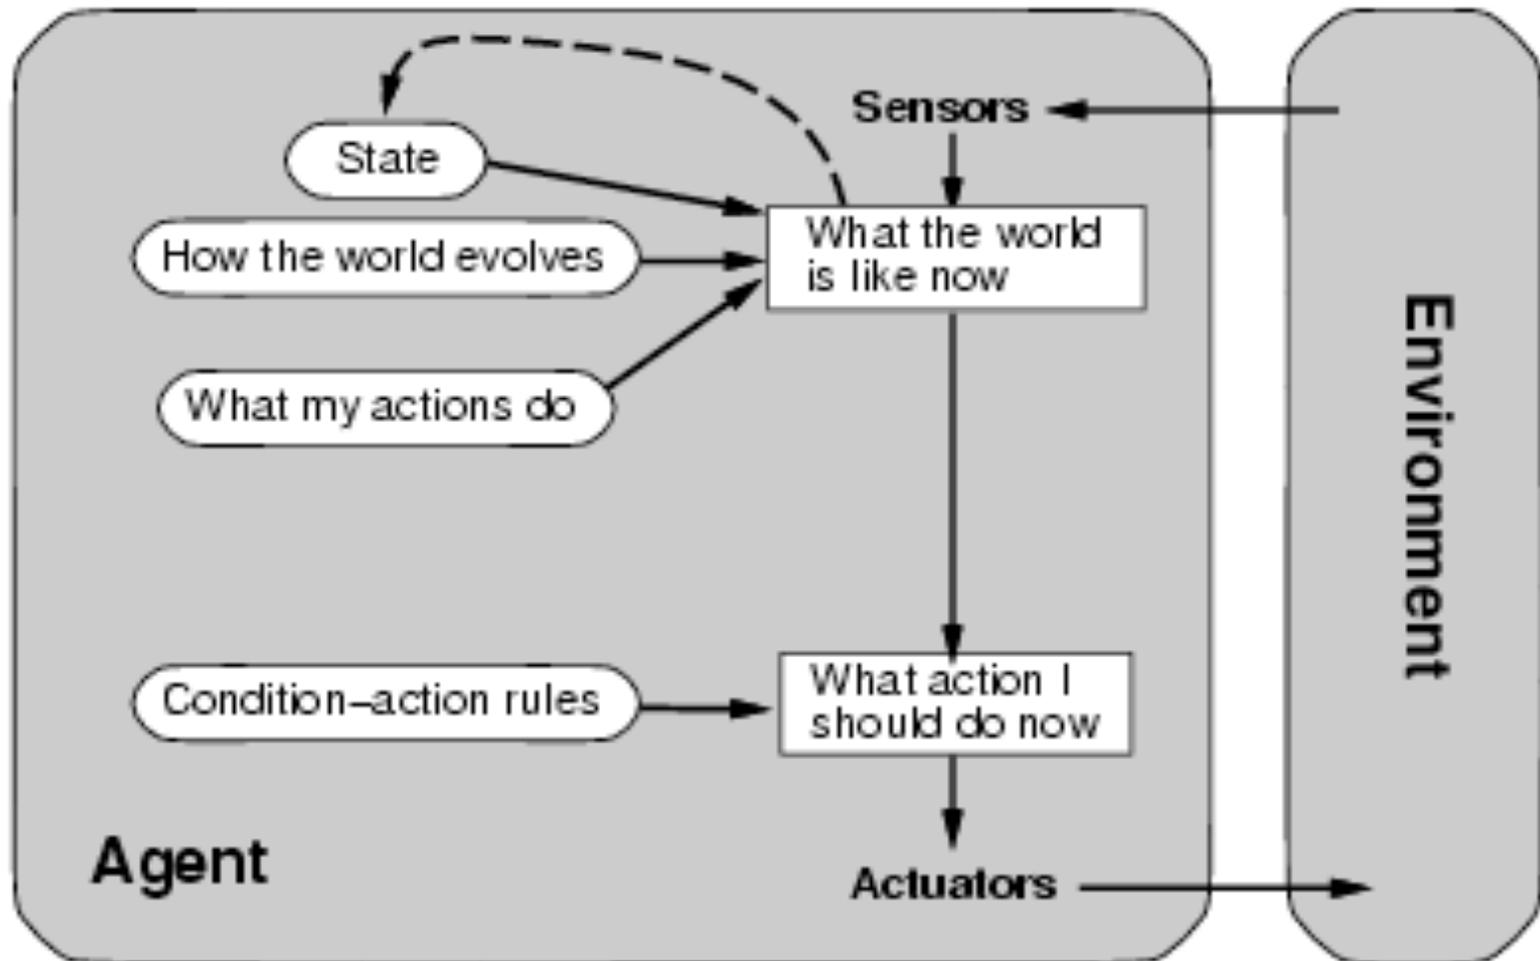
- **Fully observable** (vs. **partially observable**): An agent's sensors give it access to the complete state of the environment at each point in time.
- **Deterministic** (vs. **stochastic**): The next state of the environment is completely determined by the current state and the action executed by the agent. (If the environment is deterministic except for the actions of other agents, then the environment is **strategic**)
- **Episodic** (vs. **sequential**): An agent's action is divided into atomic episodes. Decisions do not depend on previous decisions/actions.

Environment types

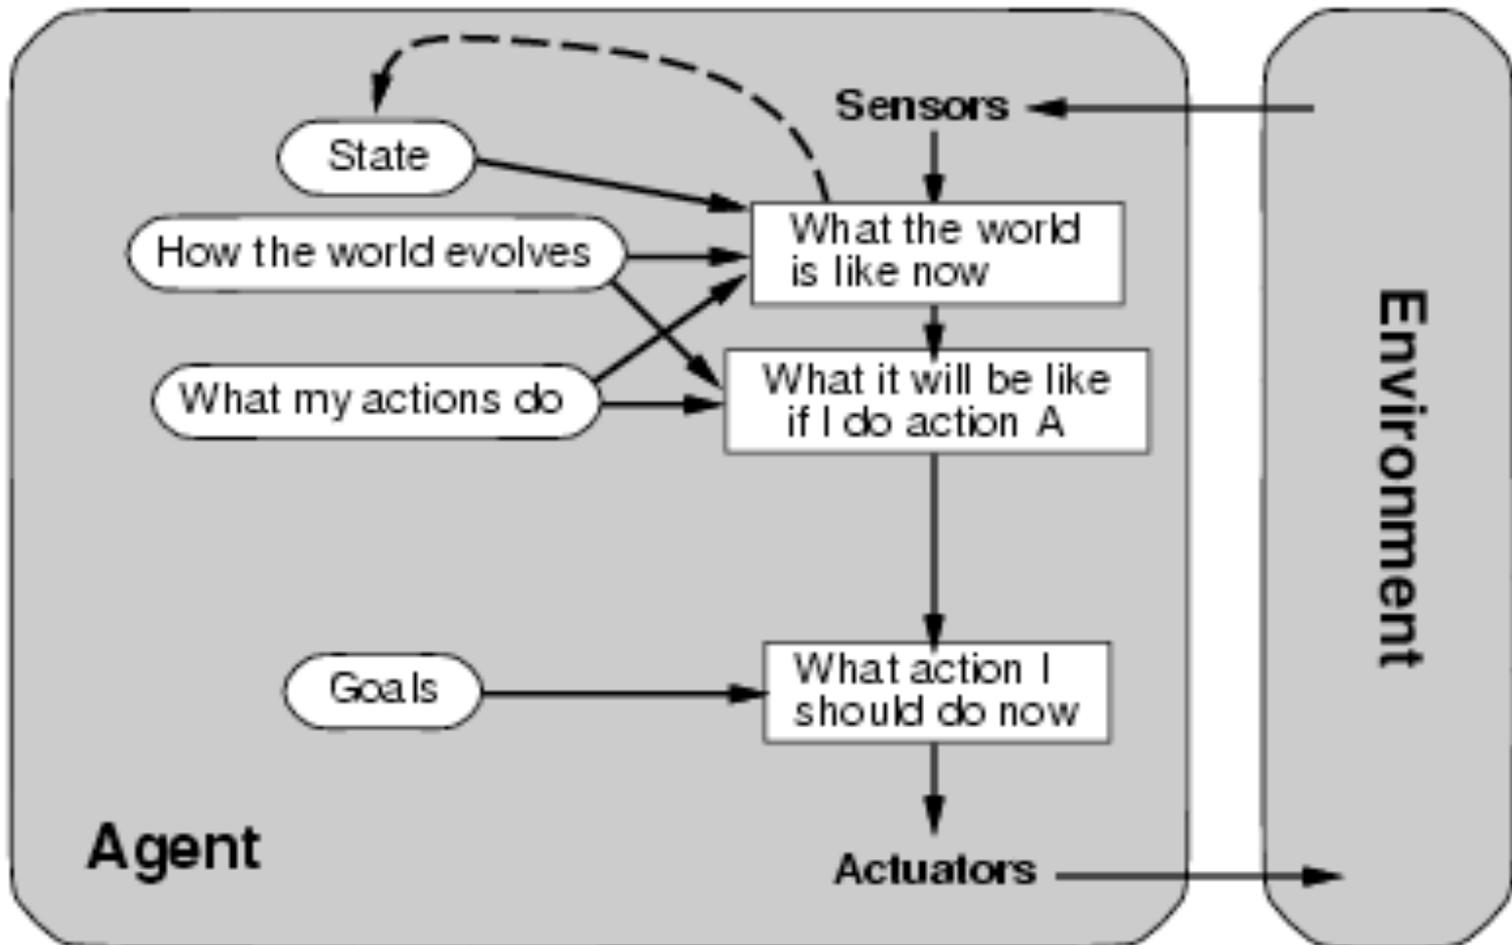
- **Static** (vs. **dynamic**): The environment is unchanged while an agent is deliberating. (The environment is **semidynamic** if the environment itself does not change with the passage of time but the agent's performance score does)
- **Discrete** (vs. **continuous**): A limited number of distinct, clearly defined percepts and actions.
- **Single agent** (vs. **multi-agent**): An agent operating by itself in an environment. Does the other agent interfere with my performance measure?

task environm.	observable	determ./ stochastic	episodic/ sequential	static/ dynamic	discrete/ continuous	agents
crossword puzzle	fully	determ.	sequential	static	discrete	single
chess with clock	fully	strategic	sequential	semi	discrete	multi
poker						
back gammon						
taxi driving	partial	stochastic	sequential	dynamic	continuous	multi
medical diagnosis	partial	stochastic	sequential	dynamic	continuous	single
image analysis	fully	determ.	episodic	semi	continuous	single
partpicking robot	partial	stochastic	episodic	dynamic	continuous	single
refinery controller	partial	stochastic	sequential	dynamic	continuous	single
interact. Eng. tutor	partial	stochastic	sequential	dynamic	discrete	multi

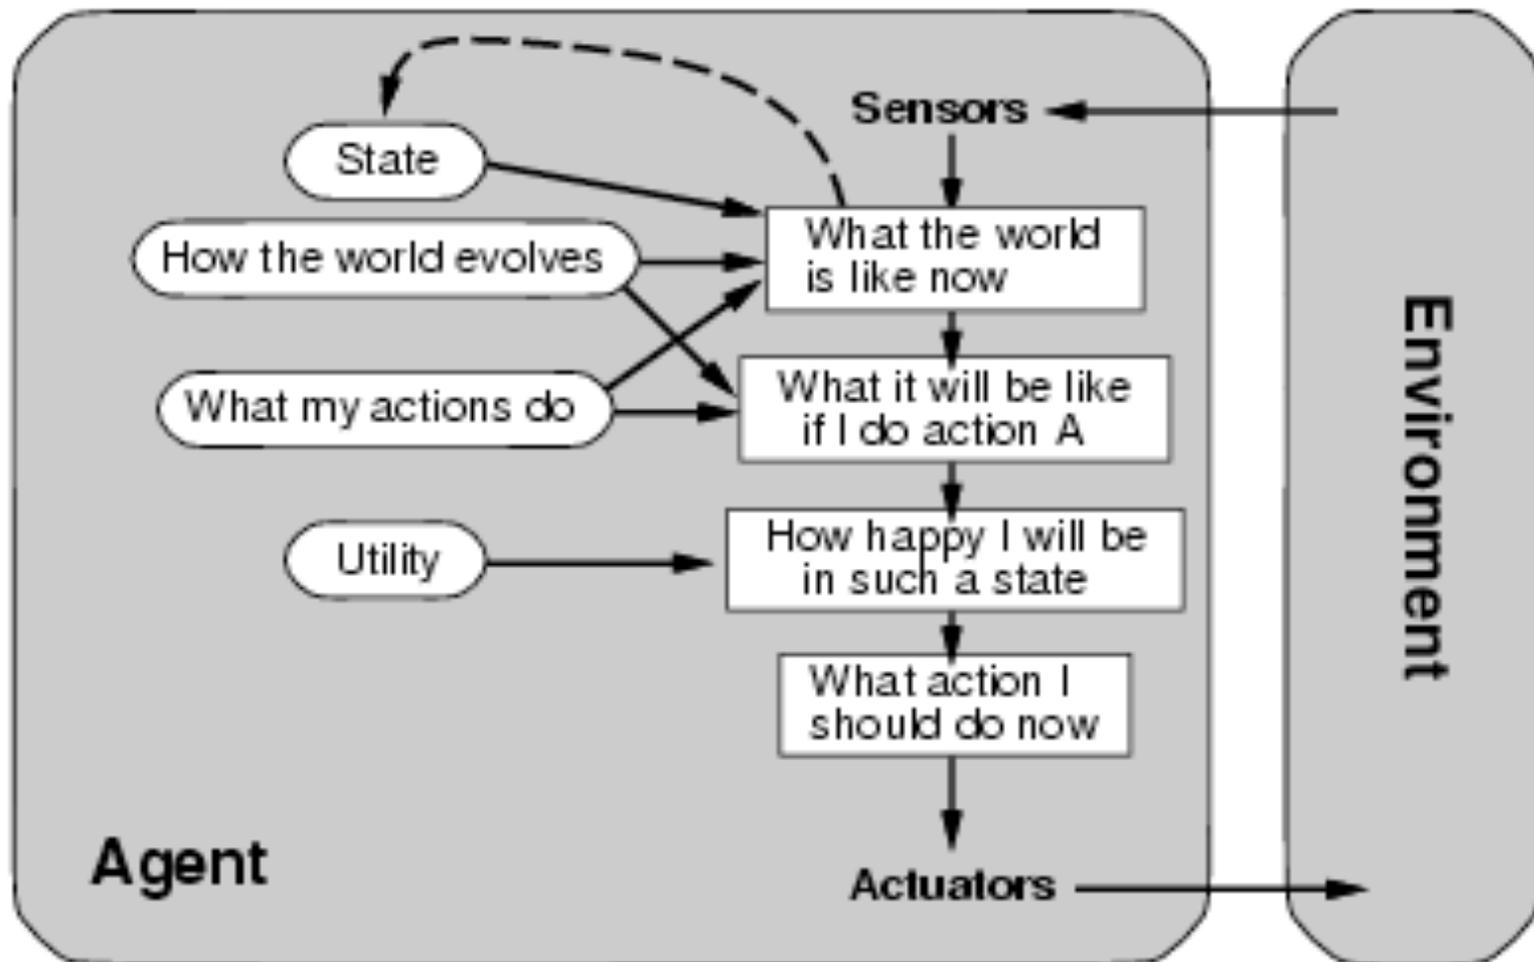
Kinds of Agents: Model-Based Agent



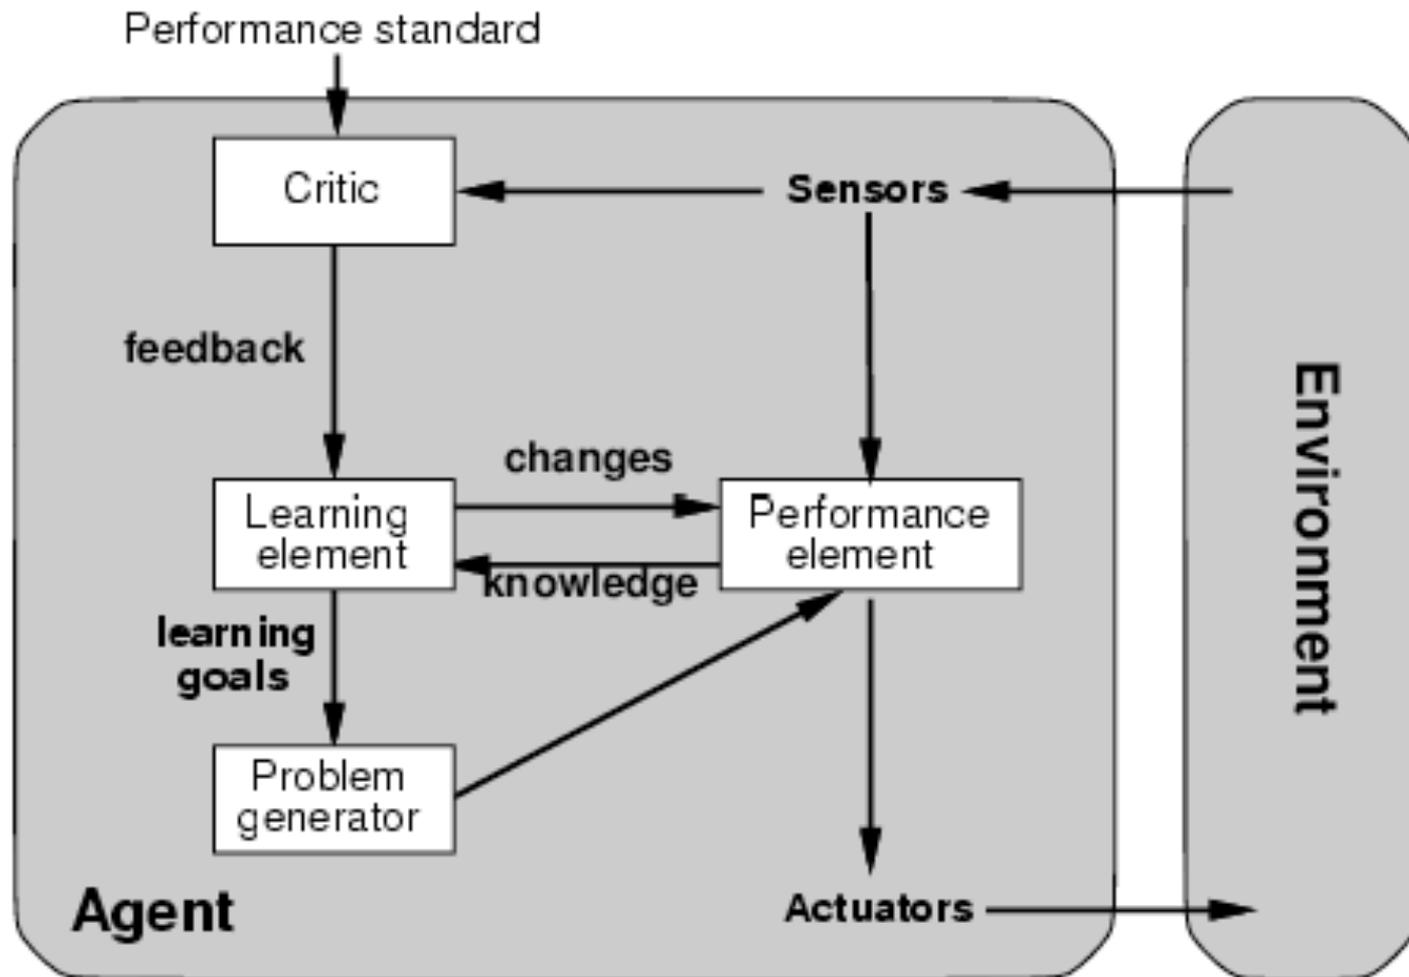
Kinds of Agents: Goal-Based Agent



Kinds of Agents: Utility-Based Agent



Kinds of Agents: Learning Agent



Learning agent

➤ Is this the best or most optimal agent given any problem ?

Python