

Intelligent Agents


Chapter 2

Outline

- Agents and environments
- Rationality
- PEAS (Performance measure, Environment, Actuators, Sensors)
- Environment types
- Agent types

Intelligent Agents


- *“An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors.”* (Russell and Norvig)



Agents

- An **agent** is anything that can be viewed as **perceiving** its **environment** through **sensors** and **acting** upon that environment through **actuators**
- Human agent:
 - eyes, ears, and other organs for sensors;
 - hands, legs, fingers, etc for actuators
- Robotic agent:
 - cameras and infrared range finders for sensors;
 - various motors for actuators

Vacuum-cleaner world



- An example of a microworld
 - Small limited problem domain that appear to require intelligence
- Percepts: location and contents, e.g., [A,Dirty]
- Actions: *Left, Right, Suck, NoOp*

A few definitions

- Percept:
 - An agent's perceptual inputs at any given instant in time (i.e., what it perceives now)
- Percept sequence:
 - Complete history of everything the agent has ever perceived

Intelligent Agents

- Some properties of an intelligent agent:
 - Autonomous: operates without direct intervention; some form of control over actions and internal state
 - Reactive: reactions to external changes in environment
 - Pro-active: goal-directed
 - Social: interacts with others to achieve goals

Necessities to Achieve these Properties

- Knowledge Representation
 - How can we formally represent an intelligent agent's knowledge?
 - Percepts need to be represented
 - An agent's internal state needs to be represented
- Reasoning
 - Problem solving (AI term: Search)
 - Planning: what actions should the agent take in order to achieve some goal?

Rational agents

- An agent should strive to "do the right thing", based on what it can perceive and the actions it can perform. The right action is the one that will cause the agent to be most successful
- Performance measure: An objective criterion for success of an agent's behavior
- E.g., performance measure of a vacuum-cleaner agent could be amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

Rational agents

- **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.

Rational agents

- Rationality is distinct from omniscience (all-knowing with infinite knowledge)
- Agents can perform actions in order to modify future percepts so as to obtain useful information (information gathering, exploration)
- An agent is **autonomous** if its behavior is determined by its own experience (with ability to learn and adapt)

Specifying the Task Environment

- We specify the task environment with PEAS



- Just kidding

PEAS

- PEAS: Performance measure, Environment, Actuators, Sensors
- Must first specify the setting for intelligent agent design
- Consider, e.g., the task of designing an automated taxi driver:
 - Performance measure
 - Environment
 - Actuators
 - Sensors

PEAS

- Must first specify the setting for intelligent agent design
- Consider, e.g., the task of designing an automated taxi driver:
 - Performance measure: Safe, fast, legal, comfortable trip, maximize profits
 - Environment: Roads, other traffic, pedestrians, customers
 - Actuators: Steering wheel, accelerator, brake, signal, horn
 - Sensors: Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard

PEAS

- Agent: Medical diagnosis system
- Performance measure: Healthy patient, minimize costs, lawsuits
- Environment: Patient, hospital, staff
- Actuators: Screen display (questions, tests, diagnoses, treatments, referrals)
- Sensors: Keyboard (entry of symptoms, findings, patient's answers)

PEAS

- Agent: Part-picking robot
- Performance measure: Percentage of parts in correct bins
- Environment: Conveyor belt with parts, bins
- Actuators: Jointed arm and hand
- Sensors: Camera, joint angle sensors

PEAS

- Agent: Interactive English tutor
- Performance measure: Maximize student's score on test
- Environment: Set of students
- Actuators: Screen display (exercises, suggestions, corrections)
- Sensors: Keyboard

What about the environment?

- Where might our intelligent agent live and operate?
- What properties of its environment can play a role on its behavior?
- How does the environment impact the design of an intelligent agent?

Fully-observable vs. partially-observable

- A fully-observable environment is one in which the agent can obtain complete, accurate, up-to-date information about the environment's state
- Most moderately complex environments (including, for example, the everyday physical world and the Internet) are only partially-observable
- The more observable an environment is, the simpler it is to build agents to operate in it

Deterministic vs. stochastic

- A deterministic environment is one in which any action has a single guaranteed effect — there is no uncertainty about the state that will result from performing an action
- Stochastic: random or probabilistic (directed)
- The physical world can to all intents and purposes be regarded as stochastic
- Stochastic and Nondeterministic environments present greater problems for the agent designer
 - Nondeterministic is not a synonym for stochastic
 - Stochastic typically means we have probabilities attached to possible outcomes

Episodic vs. sequential

- In an episodic environment, the performance of an agent is dependent on a number of discrete episodes, with no link between the performance of an agent in different scenarios
- An episode consists of an agent perceiving and then executing a single action
- Choice of action in an episode only depends on the episode itself
- Episodic environments are simpler from the agent developer's perspective because the agent can decide what action to perform based only on the current episode — it need not reason about the interactions between this and future episodes

Static vs. dynamic

- A static environment is one that can be assumed to remain unchanged except by the performance of actions by the agent
- A dynamic environment is one that has other processes operating on it, and which hence changes in ways beyond the agent's control
- Other processes can interfere with the agent's actions (as in concurrent systems theory)
- The physical world is a highly dynamic environment

Discrete vs. continuous

- An environment is discrete if there are a fixed, finite number of actions and percepts in it
- Russell and Norvig give a chess game as an example of a discrete environment, and taxi driving as an example of a continuous one
- Continuous environments have a certain level of mismatch with computer systems
- Discrete environments could *in principle* be handled by a kind of “lookup table”

More properties of Environments

- Single agent vs. multiagent):
 - An agent operating by itself in an environment is a single agent system.
 - Agents interacting with other agents is a multi-agent environment
- Known vs unknown
 - In an unknown environment, the agent (or its designer) does not know the “laws of physics” of the environment

Some examples

	Crossword Puzzle	Chess (no clock)	Chess (with clock)
Observable			
Agents			
Deterministic			
Episodic			
Static			
Discrete			

Some examples

	Crossword Puzzle	Chess (no clock)	Chess (with clock)
Observable	Fully	Fully	Fully
Agents	Single	Multi	Multi
Deterministic	Yes	Yes	Yes
Episodic	Sequential	Sequential	Sequential
Static	Yes	Yes	Semidynamic
Discrete	Yes	Yes	Yes

Some examples

	Poker	Blackjack	Backgammon
Observable			
Agents			
Deterministic			
Episodic			
Static			
Discrete			

Some examples

	Poker	Blackjack	Backgammon
Observable	Partially	Partially	Fully
Agents	Multi	Multi	Multi
Deterministic	Stochastic	Stochastic	Stochastic
Episodic	Sequential	Sequential	Sequential
Static	Yes	Yes	Yes
Discrete	Yes	Yes	Yes

Some examples

	Taxi driving	Medical diagnosis	Part picking robot
Observable			
Agents			
Deterministic			
Episodic			
Static			
Discrete			

Some examples

	Taxi driving	Medical diagnosis	Part picking robot
Observable	Partially	Partially	Partially
Agents	Multi	Single	Single
Deterministic	Stochastic	Stochastic	Stochastic
Episodic	Sequential	Sequential	Yes
Static	Dynamic	Dynamic	Dynamic
Discrete	Continuous	Continuous	Continuous

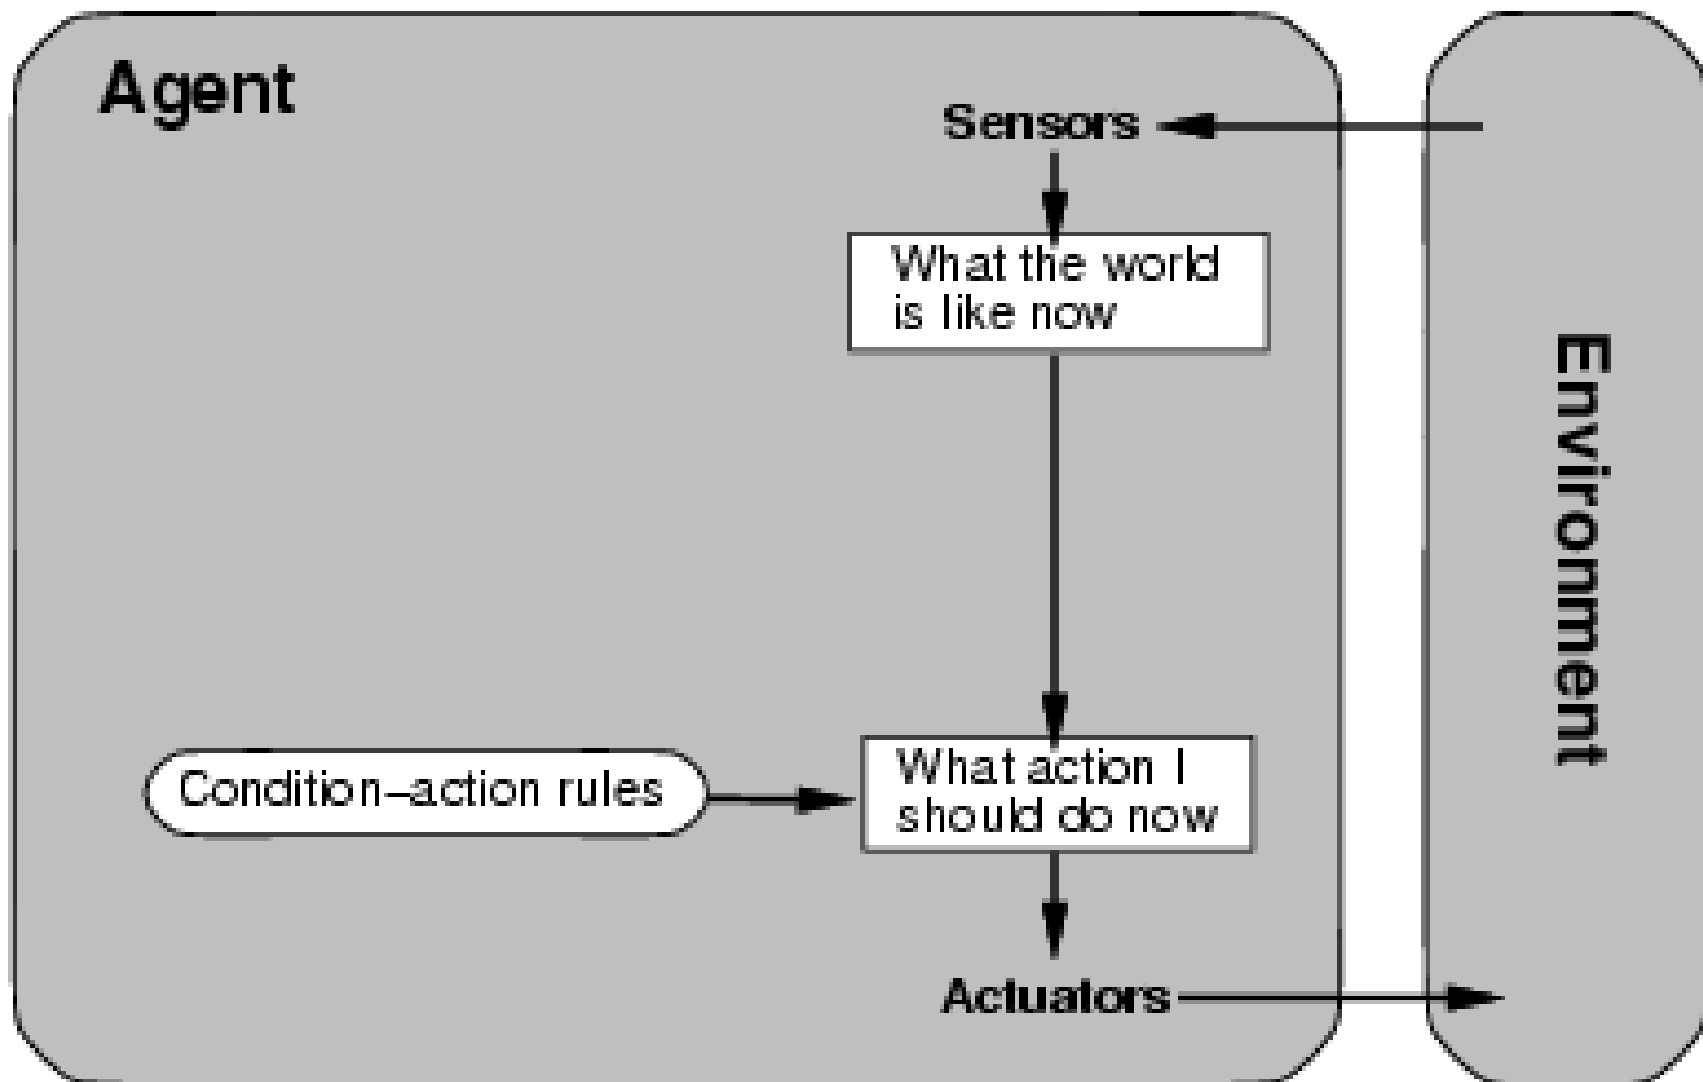
Properties of Task Environment

- The real world is (of course) partially observable, stochastic, sequential, dynamic, continuous, multi-agent

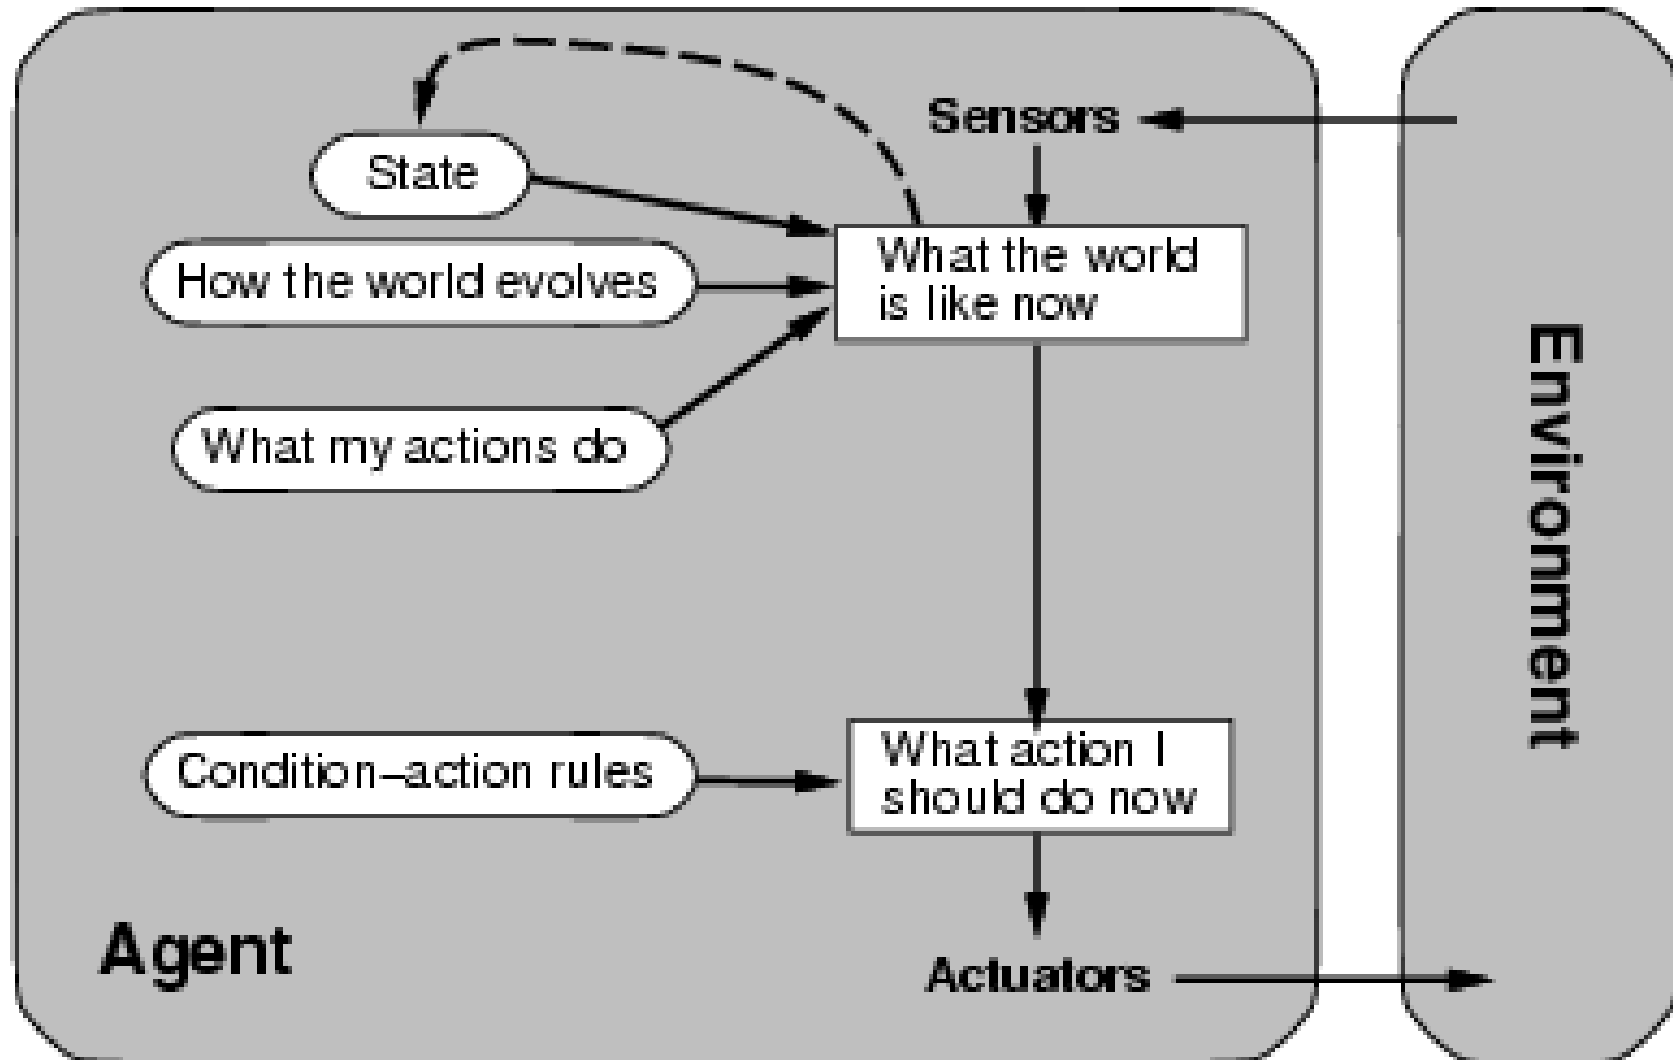
Agent types

- Four basic types in order of increasing generality:
 - Simple reflex agents
 - Model-based reflex agents
 - Goal-based agents
 - Utility-based agents

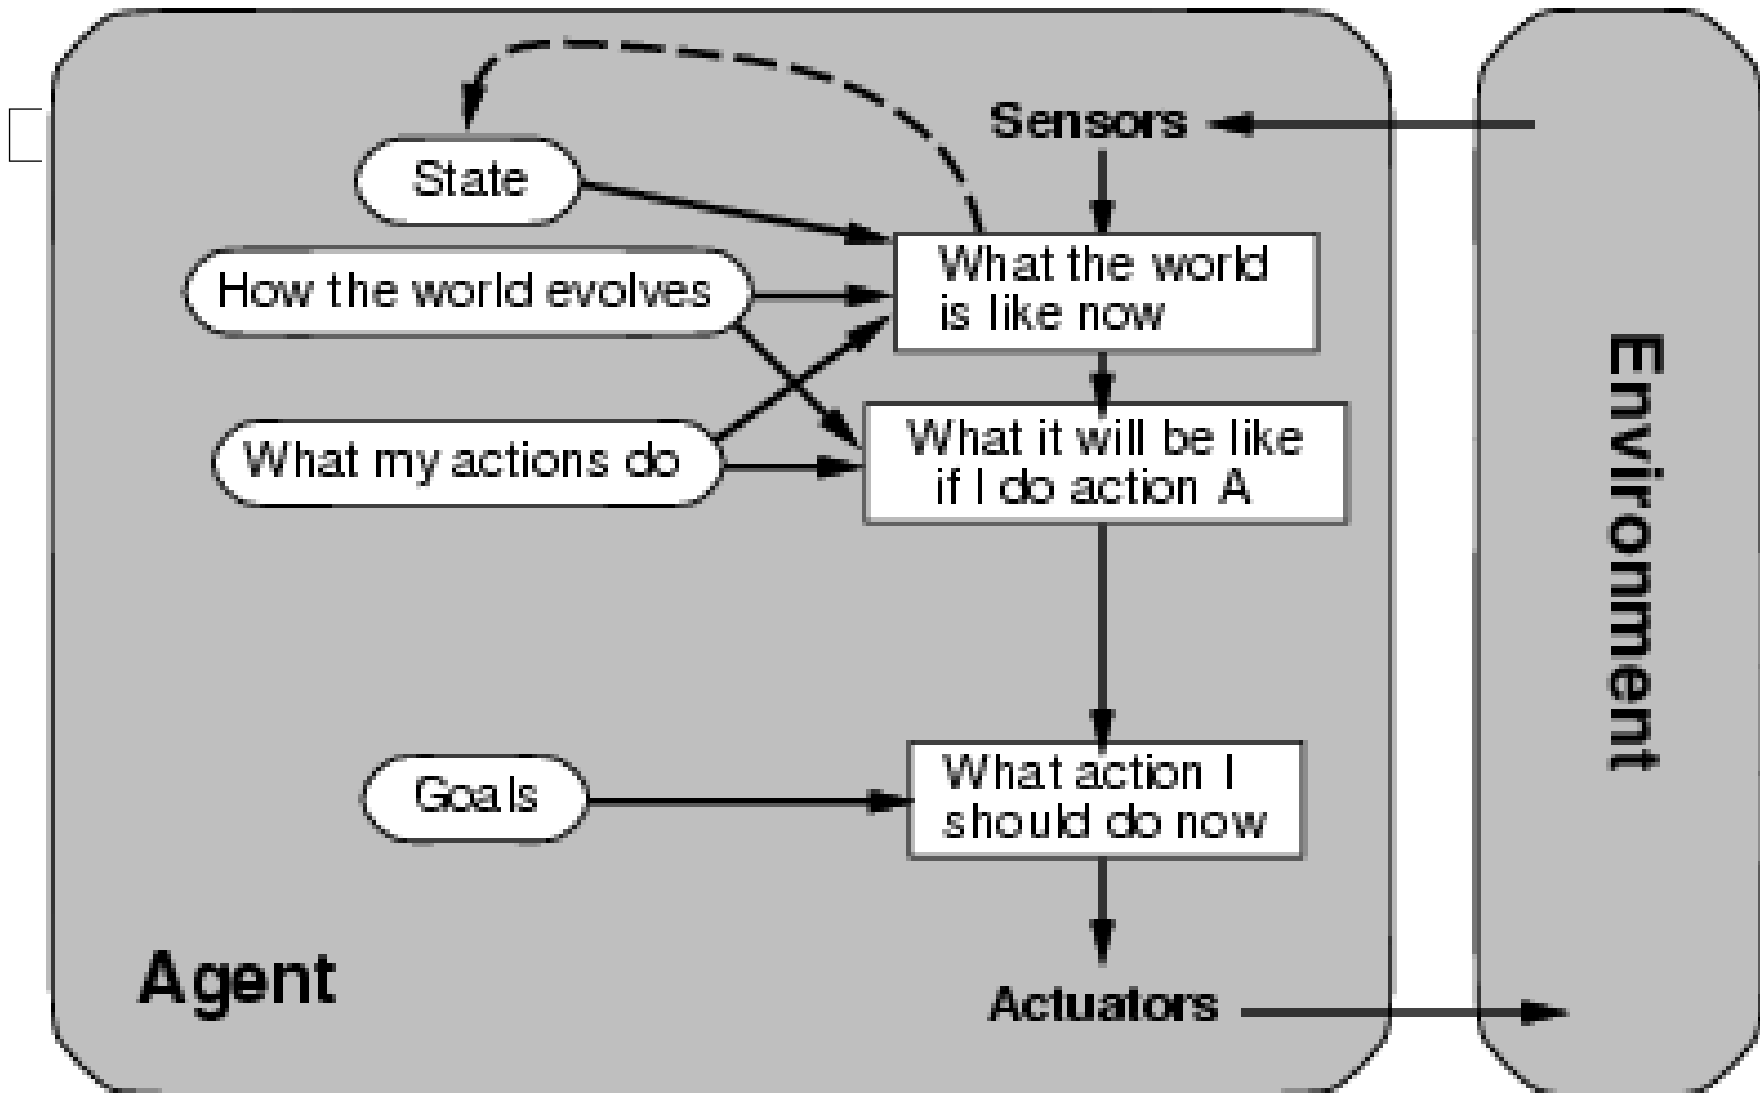
Simple reflex agents



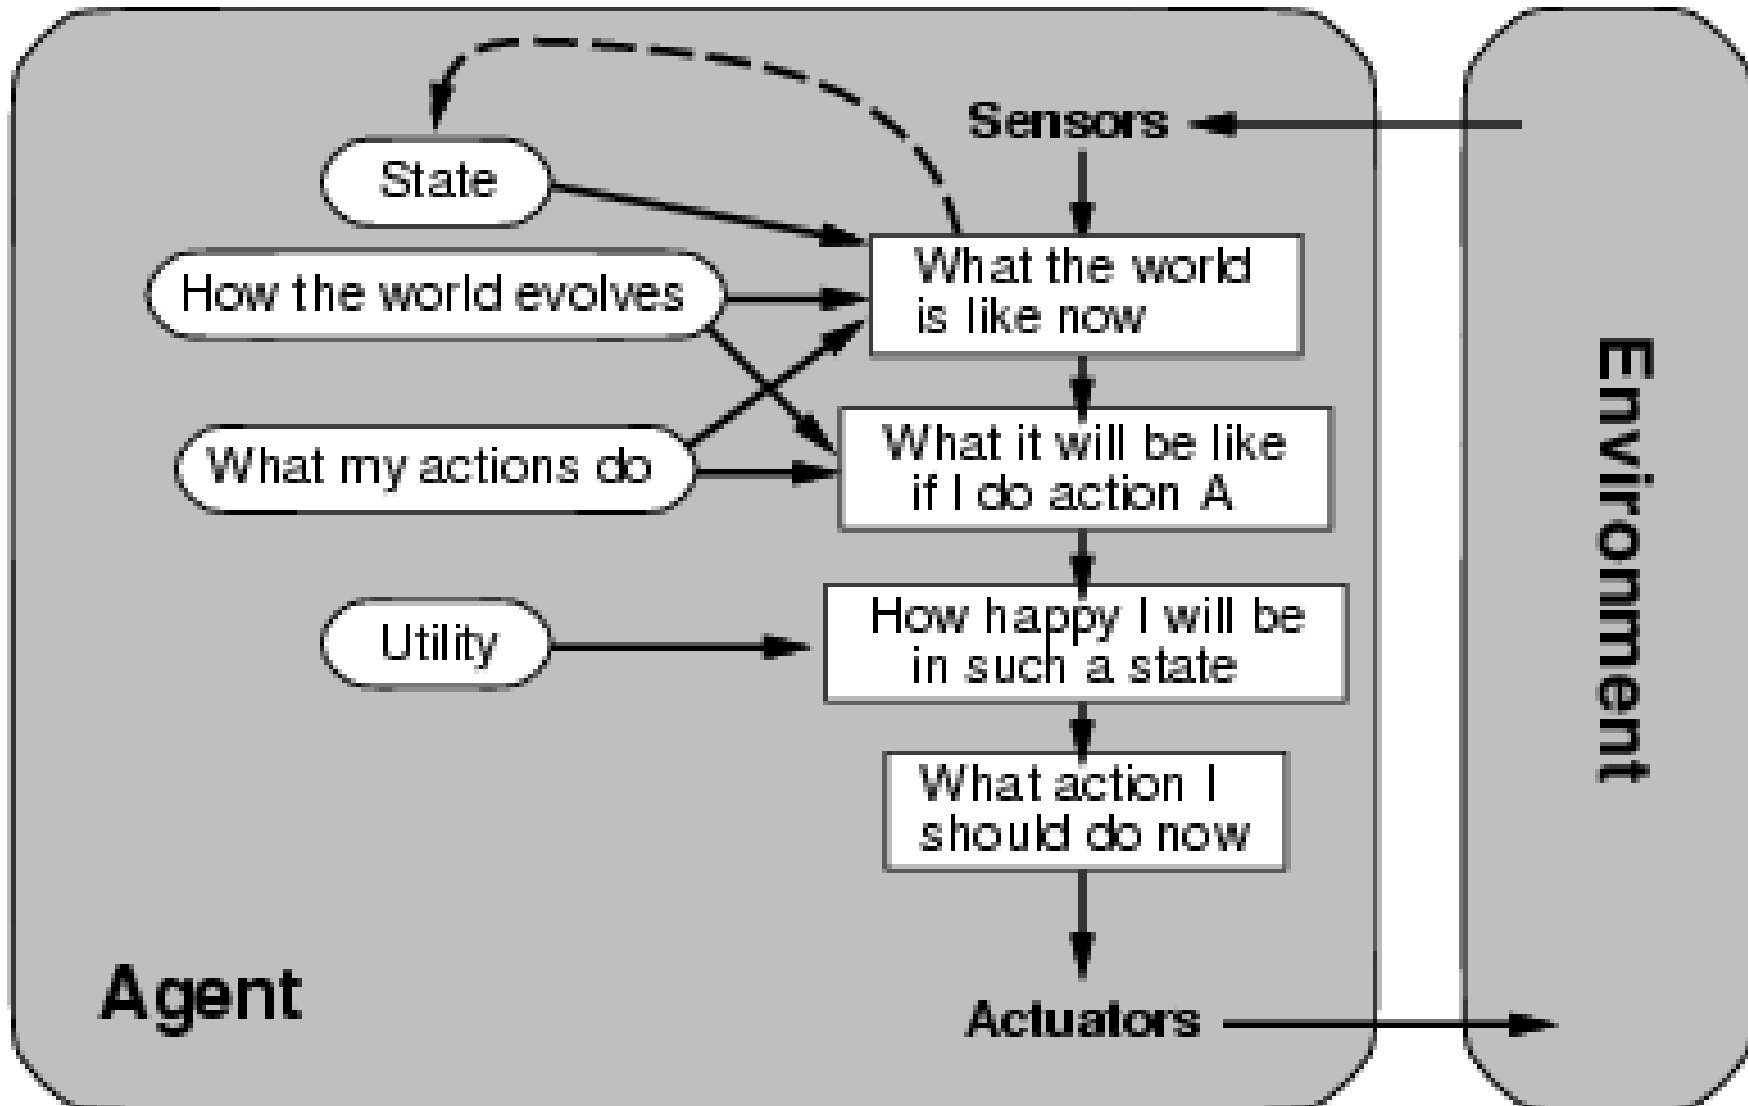
Model-based reflex agents



Goal-based agents



Utility-based agents



Learning agents

