Artificial Intelligence

Intelligent Agents

Outline

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

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, mouth, and other body parts for actuators

- Robotic agent: cameras and infrared range finders for sensors; various motors for actuators

- This course is about designing rational agents.

Agents and environments

- Abstract description of an agent (external) — The agent function maps from percept histories to actions:

  [f: P* —> A]

- Concrete implementation (internal) — The agent program runs on the physical architecture to produce f

- agent = architecture + program

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

- Caveat: computational limitations make perfect rationality unachievable —> design best program for given machine resources

Example

- Nail Gun Agent (not very smart)

- Table for f

Percept Sequence	Action
[board, full]	punch
[board, empty]	load
[blank, full]	wait
[blank, empty	load
[board, full] [board, full]	punch
[board, full] [board, empty]	load

- Function f
  if empty return load
  else if board return punch
  else return wait

Rational agents (I)

- Goal: design rational agents

- 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.

- Rationality at any moment depends on:
    - the performance measure that defines the criterion of success
    - the agent's prior knowledge of the environment
    - the actions that the agent can perform
    - the agent's percept sequence to date

Rational agents (II)

- Rationality is distinct from omniscience (all-knowing with infinite knowledge)
- Exploration, learning and autonomy are all important part of rationality.
  - exploration, or information gathering, is to do actions in order to modify the future percepts so that the agent would not do any uninformed action.
  - learning---adapt to changing environment. A successful agent computes its agent function in 3 periods:
    - when it is designed by its designer;
    - when it decides the next action;
    - when it learns from experiences to decide how to modify its behaviour (improve the rules).
- autonomy---it should learn what it can to compensate for partial or incorrect prior knowledge, relies more on its percepts than on the prior knowledge of its designer.

Task Environment

- To design a rational agent, we must specify the task environment (the problem to which the rational agents are the “solutions”).

- It includes (PEAS):
  - Performance measure
  - Environment
  - Actuators
  - Sensors

- The environment type largely determines the agent design.

PEAS - 1

- Agent: Self-driving car
- Performance measure: Safe (for both passengers and pedestrians), fast, legal, comfortable trip, efficiency
- Environment: Roads, other traffic, pedestrians, passengers
- Actuators: Steering wheel, accelerator, brake, signal, horn
- Sensors: Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard

PEAS - 2

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

PEAS - 3

- Agent: Assembly line robot
- Performance measure: Completeness, precision, speed
- Environment: Conveyor belt, parts
- Actuators: Jointed arm and hand, tools
- Sensors: Camera, joint angle sensors

PeAS - 4

- Agent: Machine translator
- Performance measure: Minimize word error rate, easy to understand
- Environment: Documents, cultures, personal preference
- Actuators: Output device, such as screen display, speaker, and printer
- Sensors: Input device, such as scanner and microphone

Environment types (I)

- 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): The agent's experience is divided into atomic "episodes" (each episode consists of the agent perceiving and then performing a single action), and the choice of action in each episode depends only on the episode itself.

Environment types (II)

- Static (vs. dynamic): The environment is unchanged while an agent is deliberating. (The environment is semi-dynamic 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.

Environment Types — Examples

	Chess with a clock	Chess without a clock	Self-driving car
Fully observable	Yes	Yes	No
Deterministic	Strategic	Strategic	No
Episodic	No	No	No
Static	Semi	Yes	No
Discrete	Yes	Yes	No
Single agent	No	No	No

- The environment type largely determines the agent design
- The real world is (of course) partially observable, stochastic, sequential, dynamic, continuous, multi-agent

Agent functions and programs

- An agent is completely specified by the agent function mapping percept sequences to actions. Percept sequence summarizes the history of what the agent has perceived.

- One agent function (or a small equivalence class) is rational.

- An agent program is an implementation or execution of the agent function. The agent function is realized by an agent program executing on the physical architecture.

- Aim: find an optimal agent program for the rational agent function.

Agent Types

- Four basic agent types in order of increasing generality:
  - Simple reflex agents
  - Model-based reflex agents
  - Goal-based agents
  - Utility-based agents
- All these can be turned into learning agents

Simple reflex agents

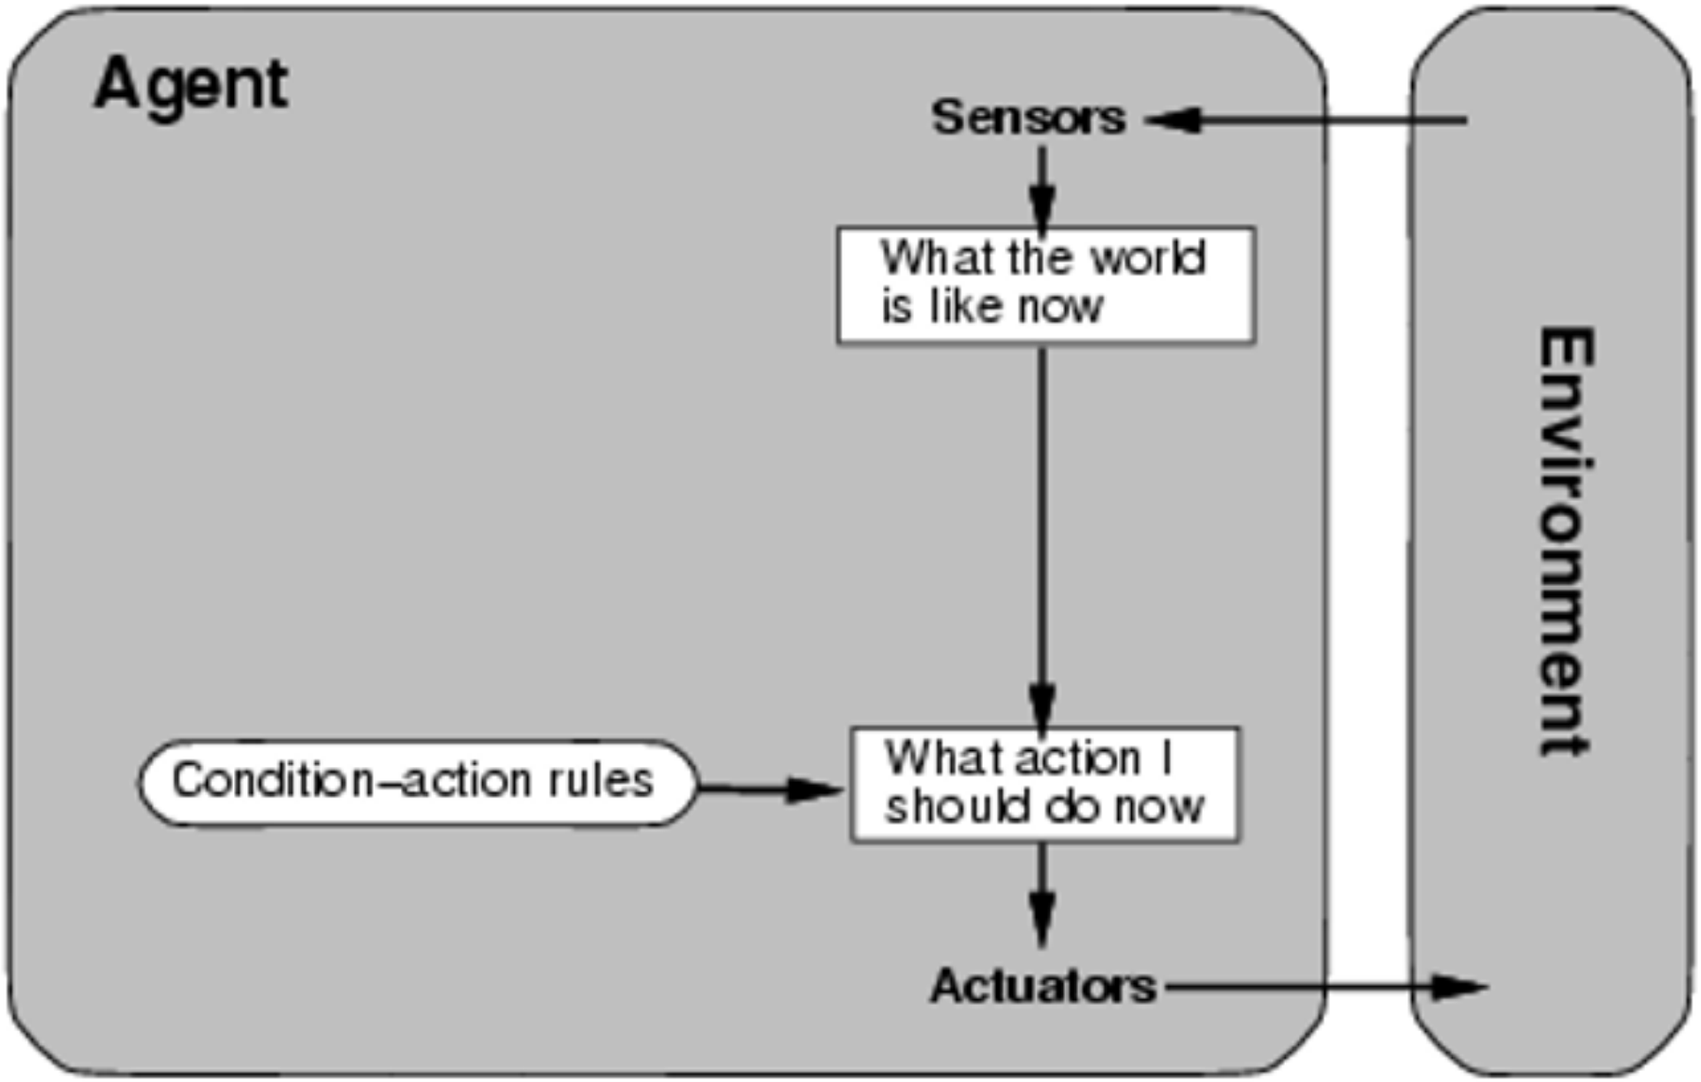

Model-based reflex agents

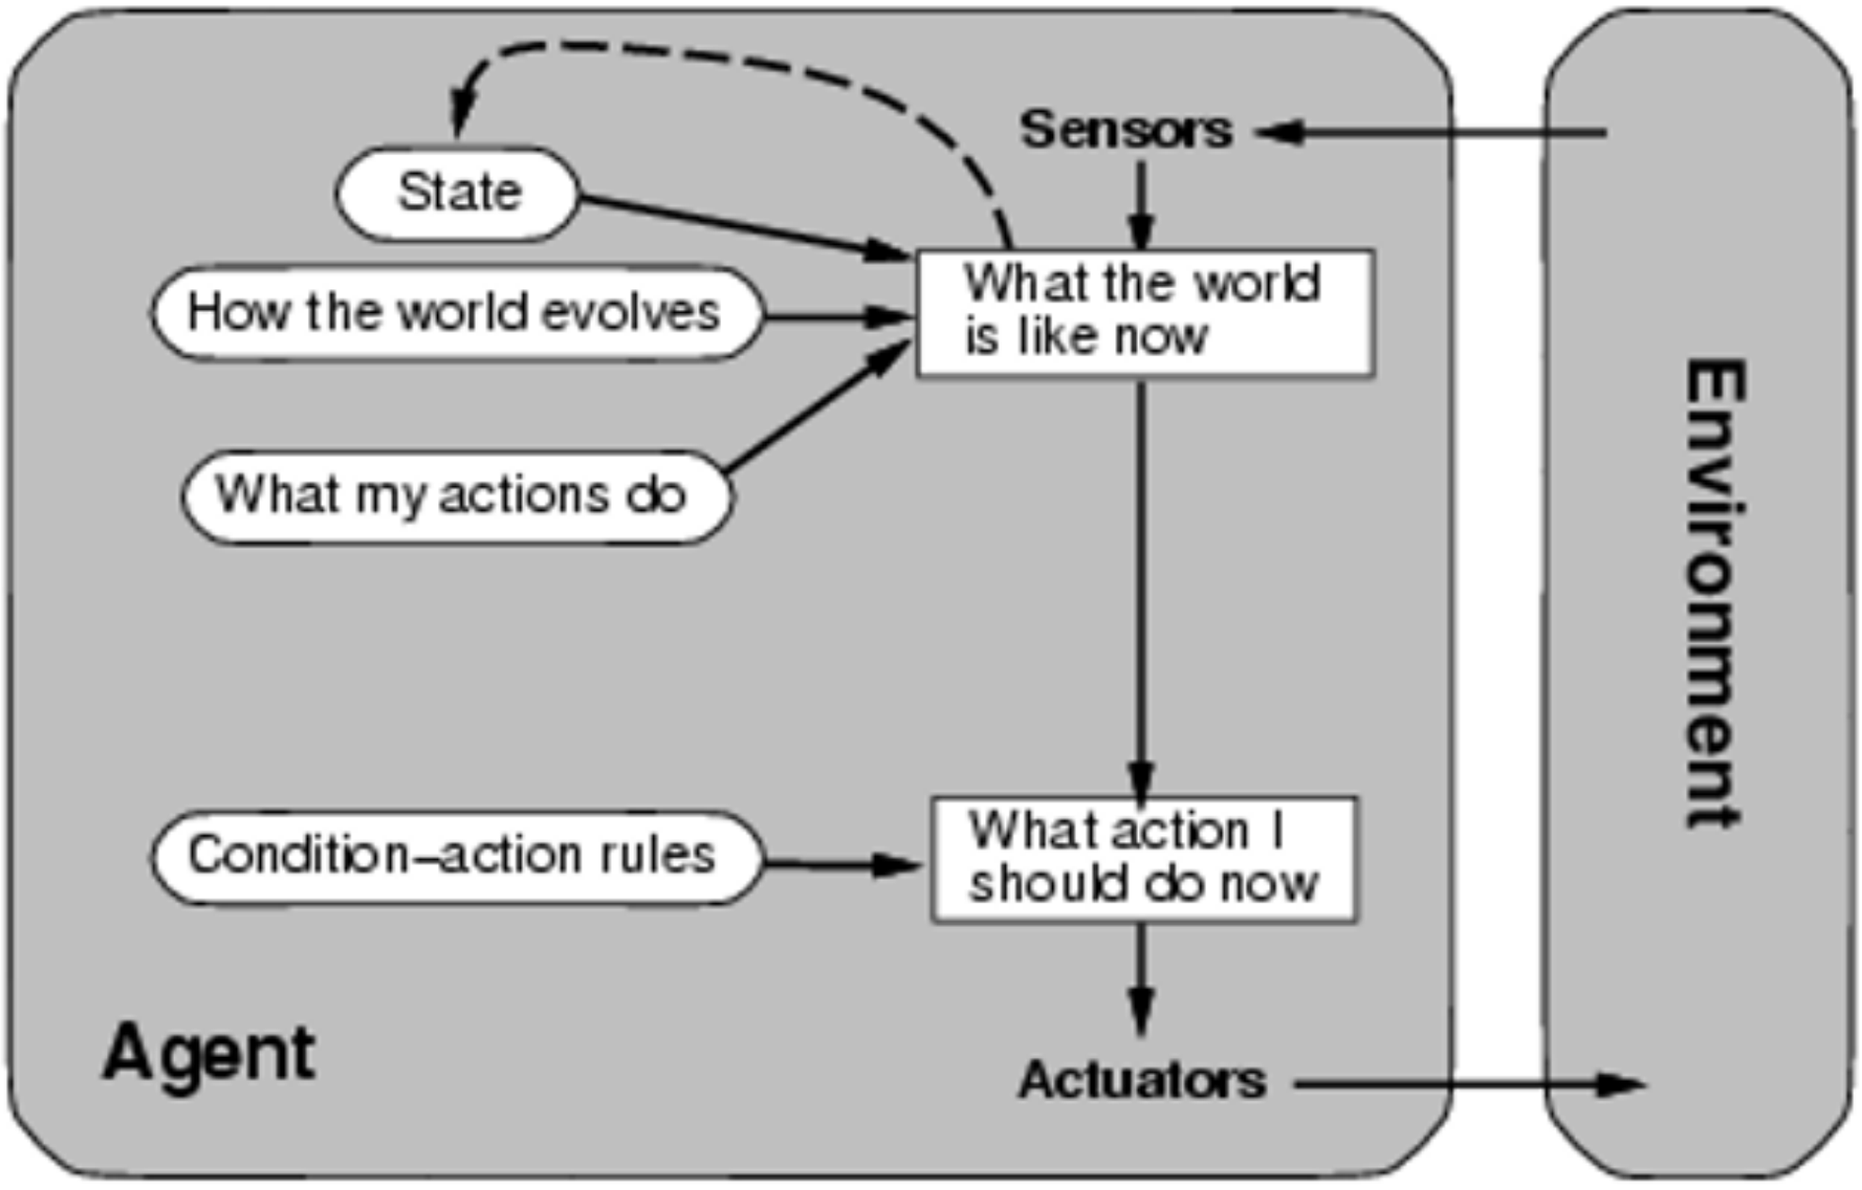

Goal-based agents

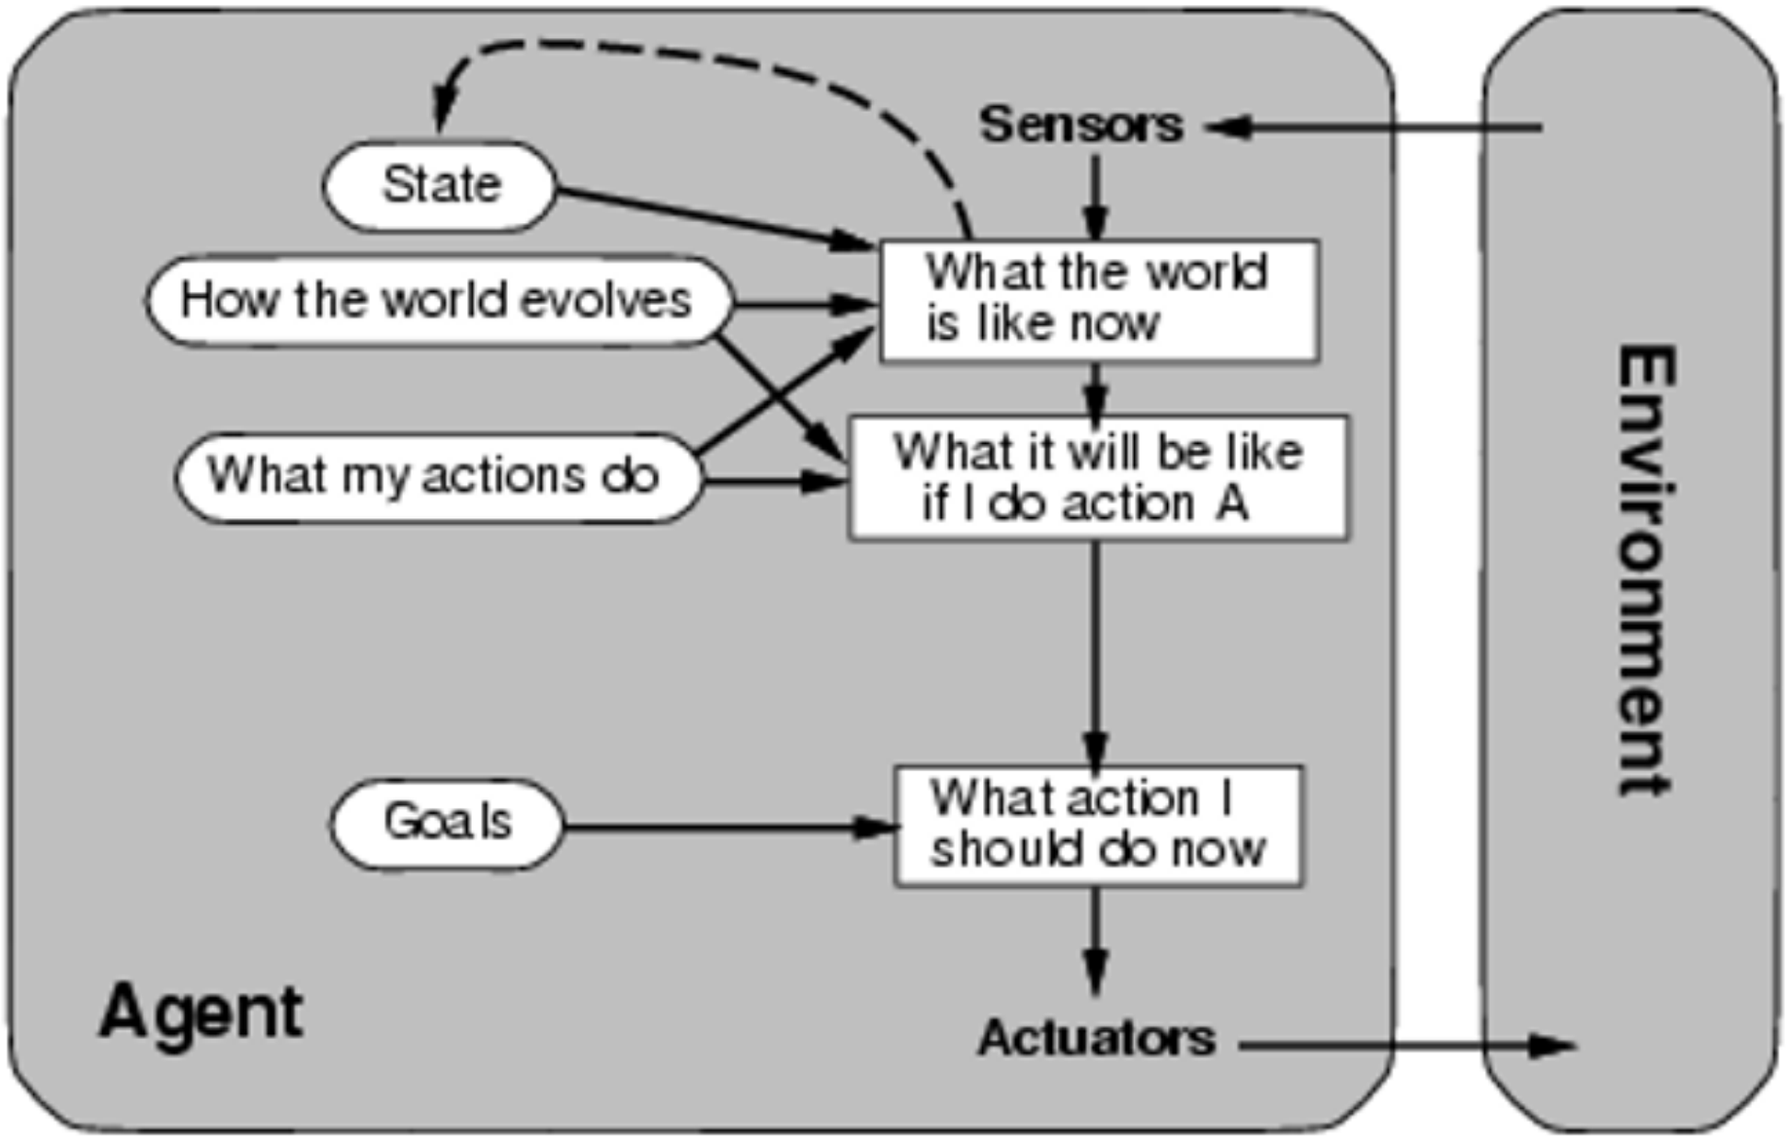

Utility-based agents

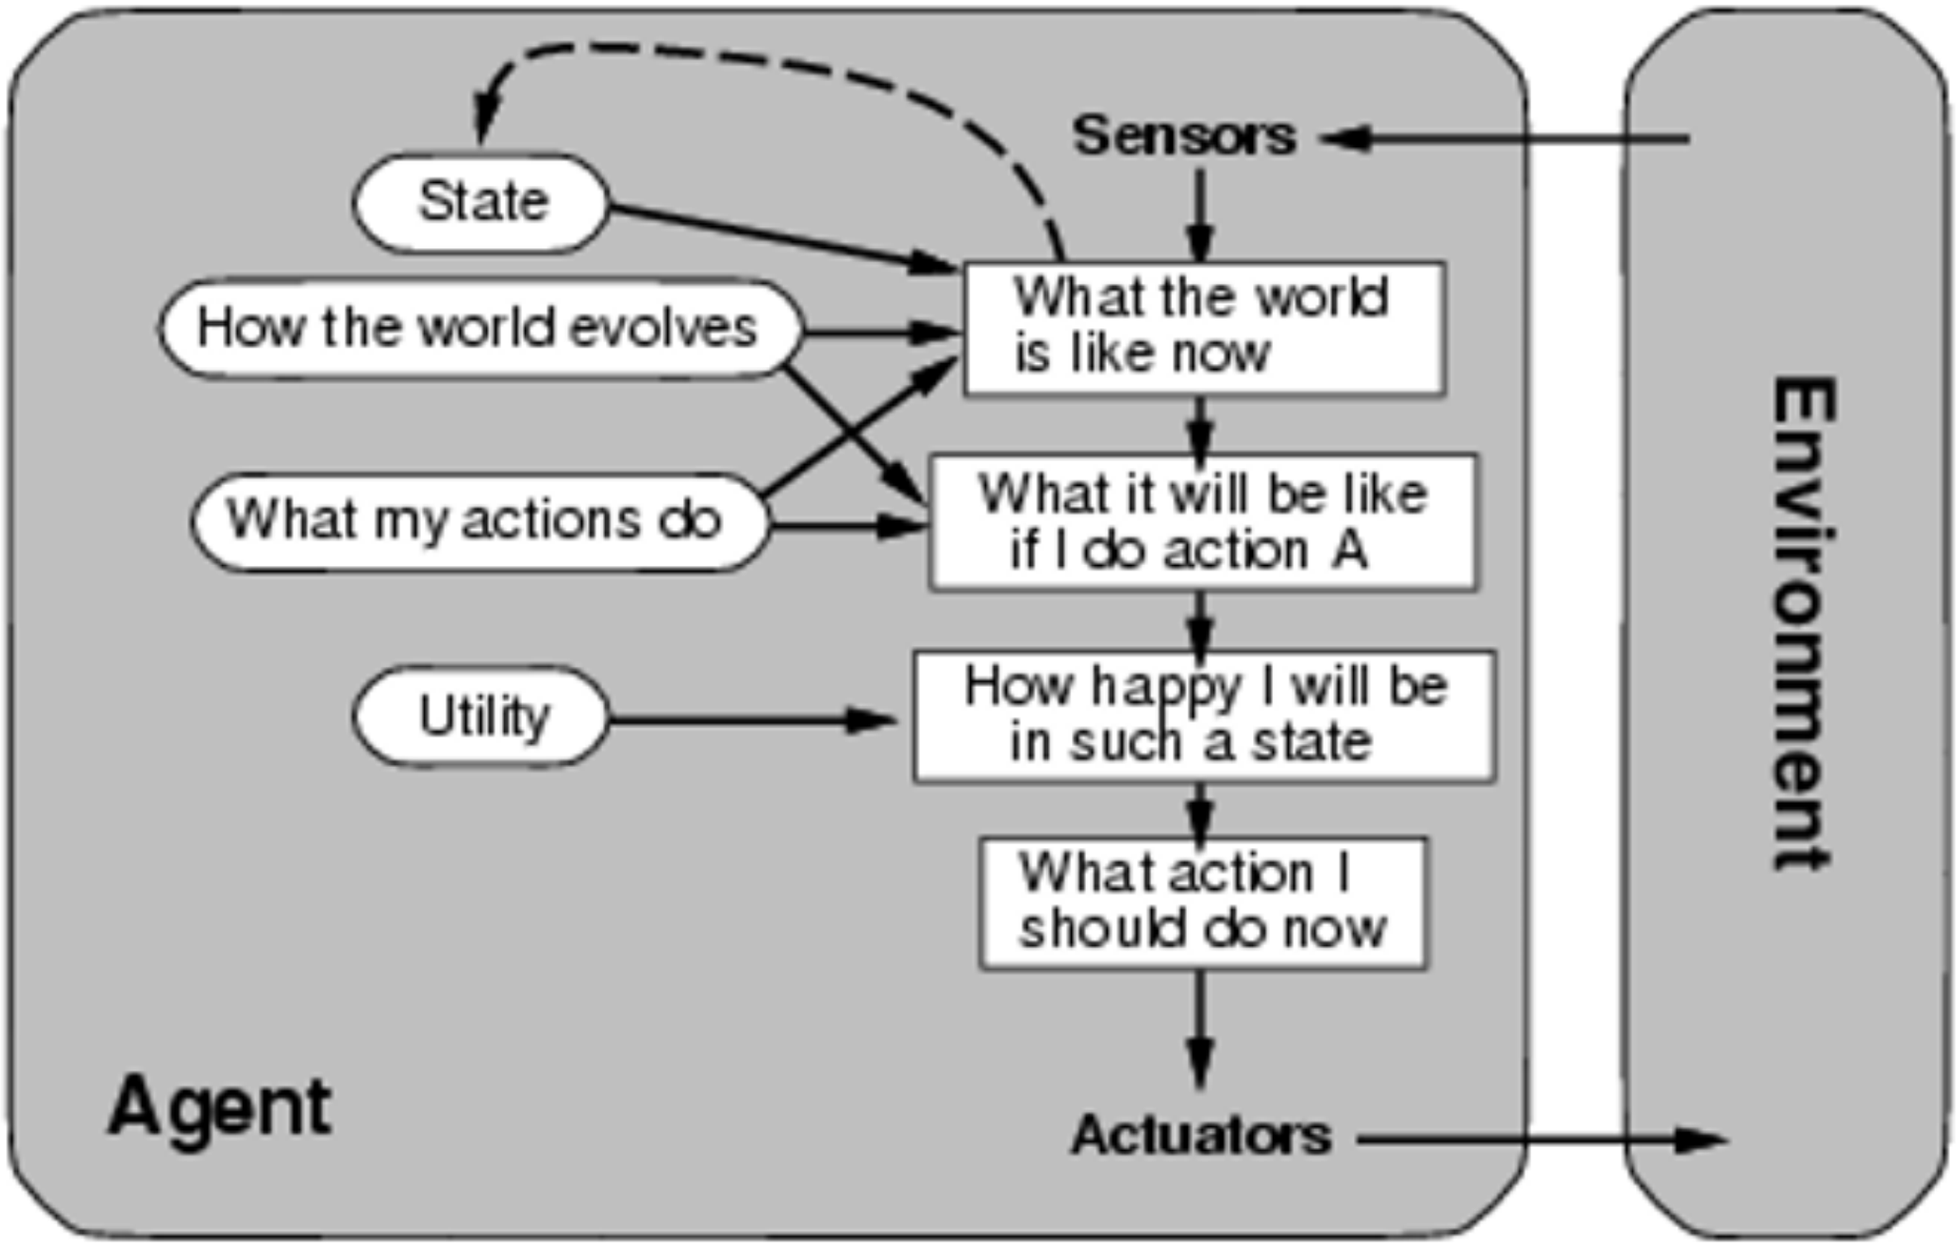

Learning agents