

Introduction to AI & Intelligent Agents

**This Lecture
Chapters 1 and 2**

**Next Lecture
Chapter 3.1 to 3.4**

(Please read lecture topic material before and after each lecture on that topic)

What is Artificial Intelligence?

- **Thought processes vs. behavior**
- **Human-like vs. rational-like**
- **How to simulate humans intellect and behavior by a machine.**
 - Mathematical problems (puzzles, games, theorems)
 - Common-sense reasoning
 - Expert knowledge: lawyers, medicine, diagnosis
 - Social behavior
 - Web and online intelligence
 - Planning for assembly and logistics operations
- **Things we call “intelligent” if done by a human.**

What is AI?


Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

The textbook advocates "acting rationally"

What is Artificial Intelligence

([John McCarthy](#), Basic Questions)



- **What is artificial intelligence?**
- It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.
- **Yes, but what is intelligence?**
- Intelligence is the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.
- **Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence?**
- Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others.
- More in: <http://www-formal.stanford.edu/jmc/whatisai/node1.html>

What is Artificial Intelligence

- **Thought processes**

- “The exciting new effort to make computers **think** .. Machines with minds, in the full and literal sense” (Haugeland, 1985)

- **Behavior**

- “The study of how to make computers **do things** at which, at the moment, people are better.” (Rich, and Knight, 1991)

- **Activities**

- The automation of activities that we associate with human thinking, activities such as decision-making, problem solving, learning... (Bellman)




AI as “Raisin Bread”

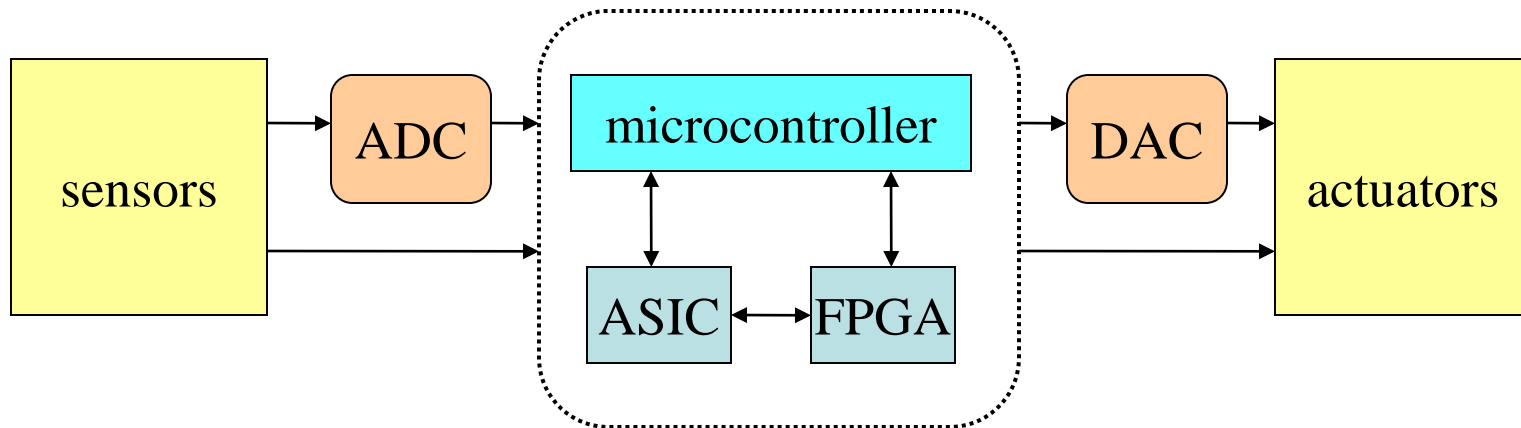


- **Esther Dyson [predicted] AI would [be] embedded in main-stream, strategically important systems, like raisins in a loaf of raisin bread.**
- **Time has proven Dyson's prediction correct.**
- **Emphasis shifts away from replacing expensive human experts with stand-alone expert systems toward main-stream computing systems that create strategic advantage.**
- **Many of today's AI systems are connected to large data bases, they deal with legacy data, they talk to networks, they handle noise and data corruption with style and grace, they are implemented in popular languages, and they run on standard operating systems.**
- **Humans usually are important contributors to the total solution.**
- **Adapted from Patrick Winston, Former Director, MIT AI Laboratory**

Agents and environments



Compare: Standard Embedded System Structure



The Turing Test

(Can Machine think? A. M. Turing, 1950)




Figure 1.1 The Turing test.

- **Requires:**

- Natural language
- Knowledge representation
- Automated reasoning
- Machine learning
- (vision, robotics) for full test

Acting/Thinking Humanly/Rationally

- **Turing test (1950)**
- **Requires:**
 - Natural language
 - Knowledge representation
 - automated reasoning
 - machine learning
 - (vision, robotics.) for full test
- **Methods for Thinking Humanly:**
 - Introspection, the general problem solver (Newell and Simon 1961)
 - Cognitive sciences
- **Thinking rationally:**
 - Logic
 - Problems: how to represent and reason in a domain
- **Acting rationally:**
 - Agents: Perceive and act

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

Agents and environments




- The **agent function** maps from percept histories to actions:

$$[f: \mathcal{P}^* \rightarrow \mathcal{A}]$$

- The **agent program** runs on the physical **architecture** to produce f
- **agent = architecture + program**

Vacuum-cleaner world



- **Percepts:** location and state of the environment, e.g., [A,Dirty], [B,Clean]
- **Actions:** *Left, Right, Suck, NoOp*

Rational agents

- **Rational Agent:** For each possible percept sequence, a rational agent should select an action that is *expected* to maximize its **performance measure**, based on the evidence provided by the percept sequence and whatever built-in knowledge the agent has.
- **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

- **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 percepts & experience (with ability to **learn and adapt**) without depending solely on build-in knowledge

Discussion Items

- **An realistic agent has finite amount of computation and memory available. Assume an agent is killed because it did not have enough computation resources to calculate some rare event that eventually that ended up killing it. Can this agent still be rational?**
- **The Turing test was contested by Searle by using the “Chinese Room” argument. The Chinese Room agent needs an exponential large memory to work. Can we “save” the Turing test from the Chinese Room argument?**

Task Environment

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

PEAS:

Performance measure

Environment

Actuators

Sensors

PEAS

- **Example: Agent = 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

- **Example: 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

- **Example: 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

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. How do we **represent** or **abstract** or **model** the world?
- **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

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	partial	stochastic	sequential	static	discrete	multi
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

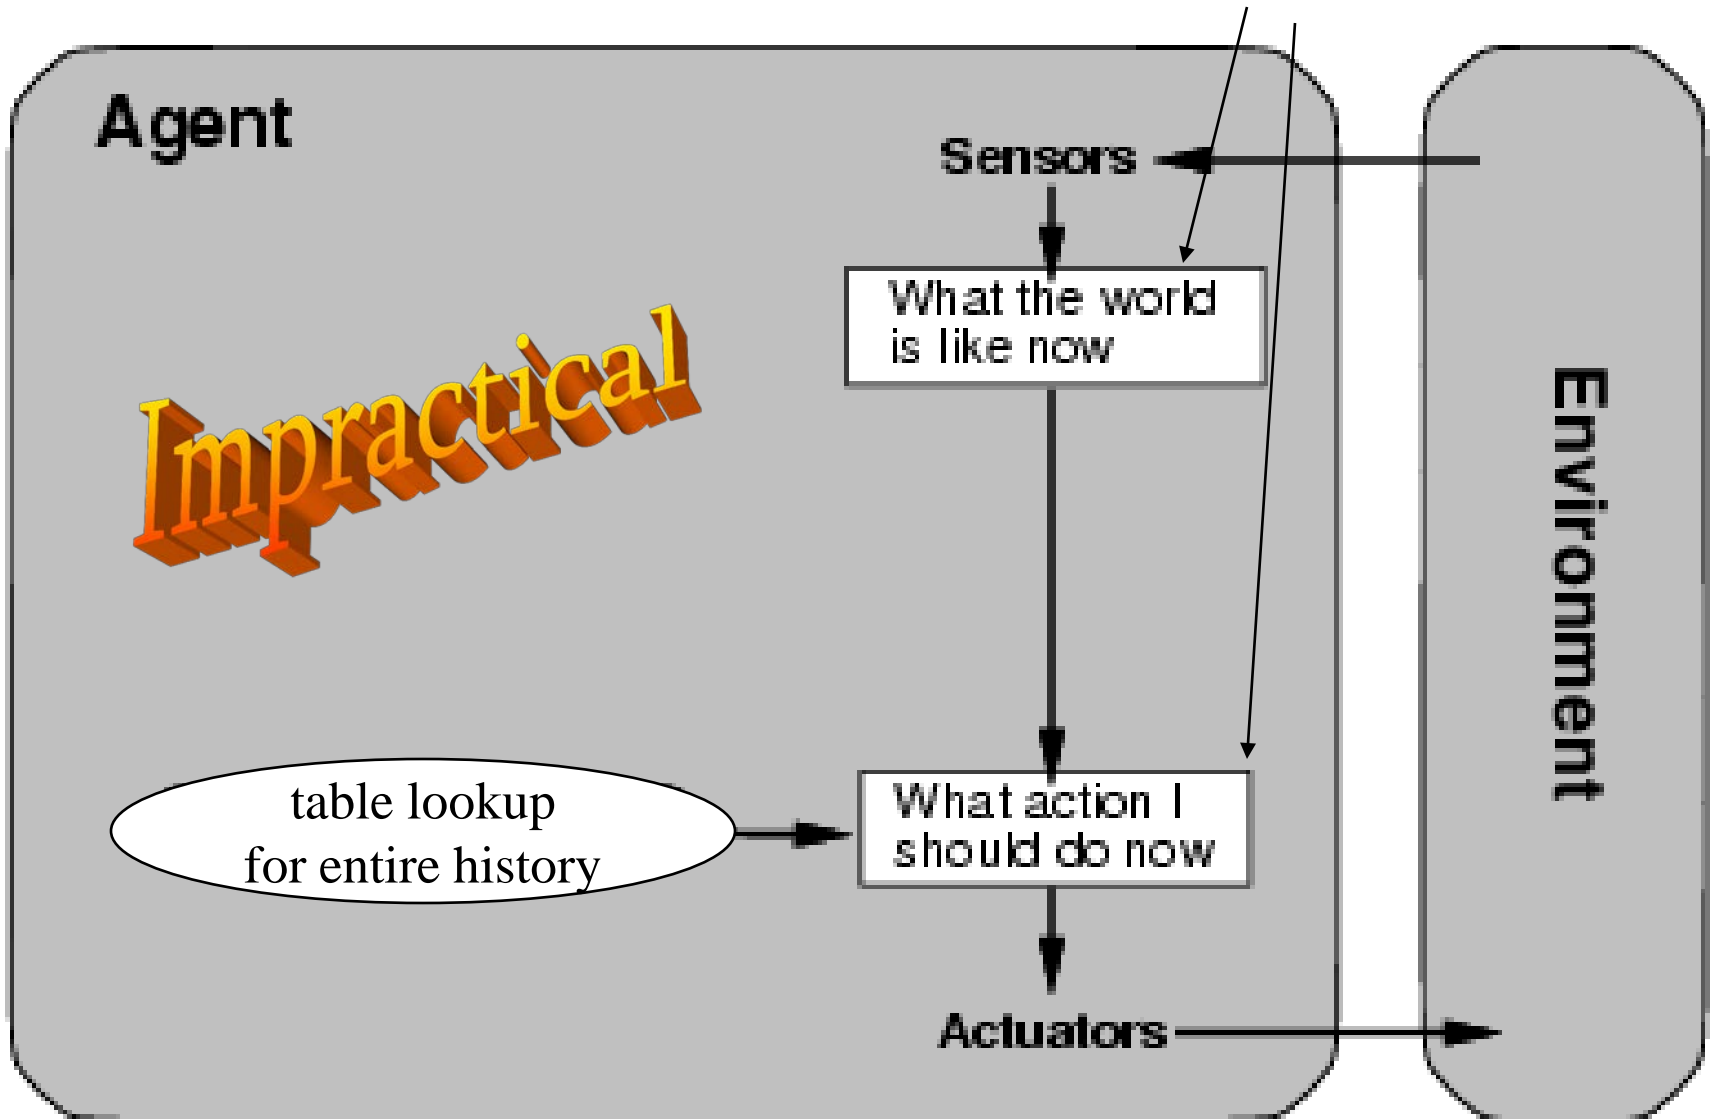
task environm.	observable	determ./ stochastic	episodic/ sequential	static/ dynamic	discrete/ continuous	agents
crossword puzzle	fully	determ.	sequential	static	discrete	single
chess with clock	fully	strategic	sequential	semi	discrete	multi
poker	partial	stochastic	sequential	static	discrete	multi
back gammon	fully	stochastic	sequential	static	discrete	multi
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

Agent types

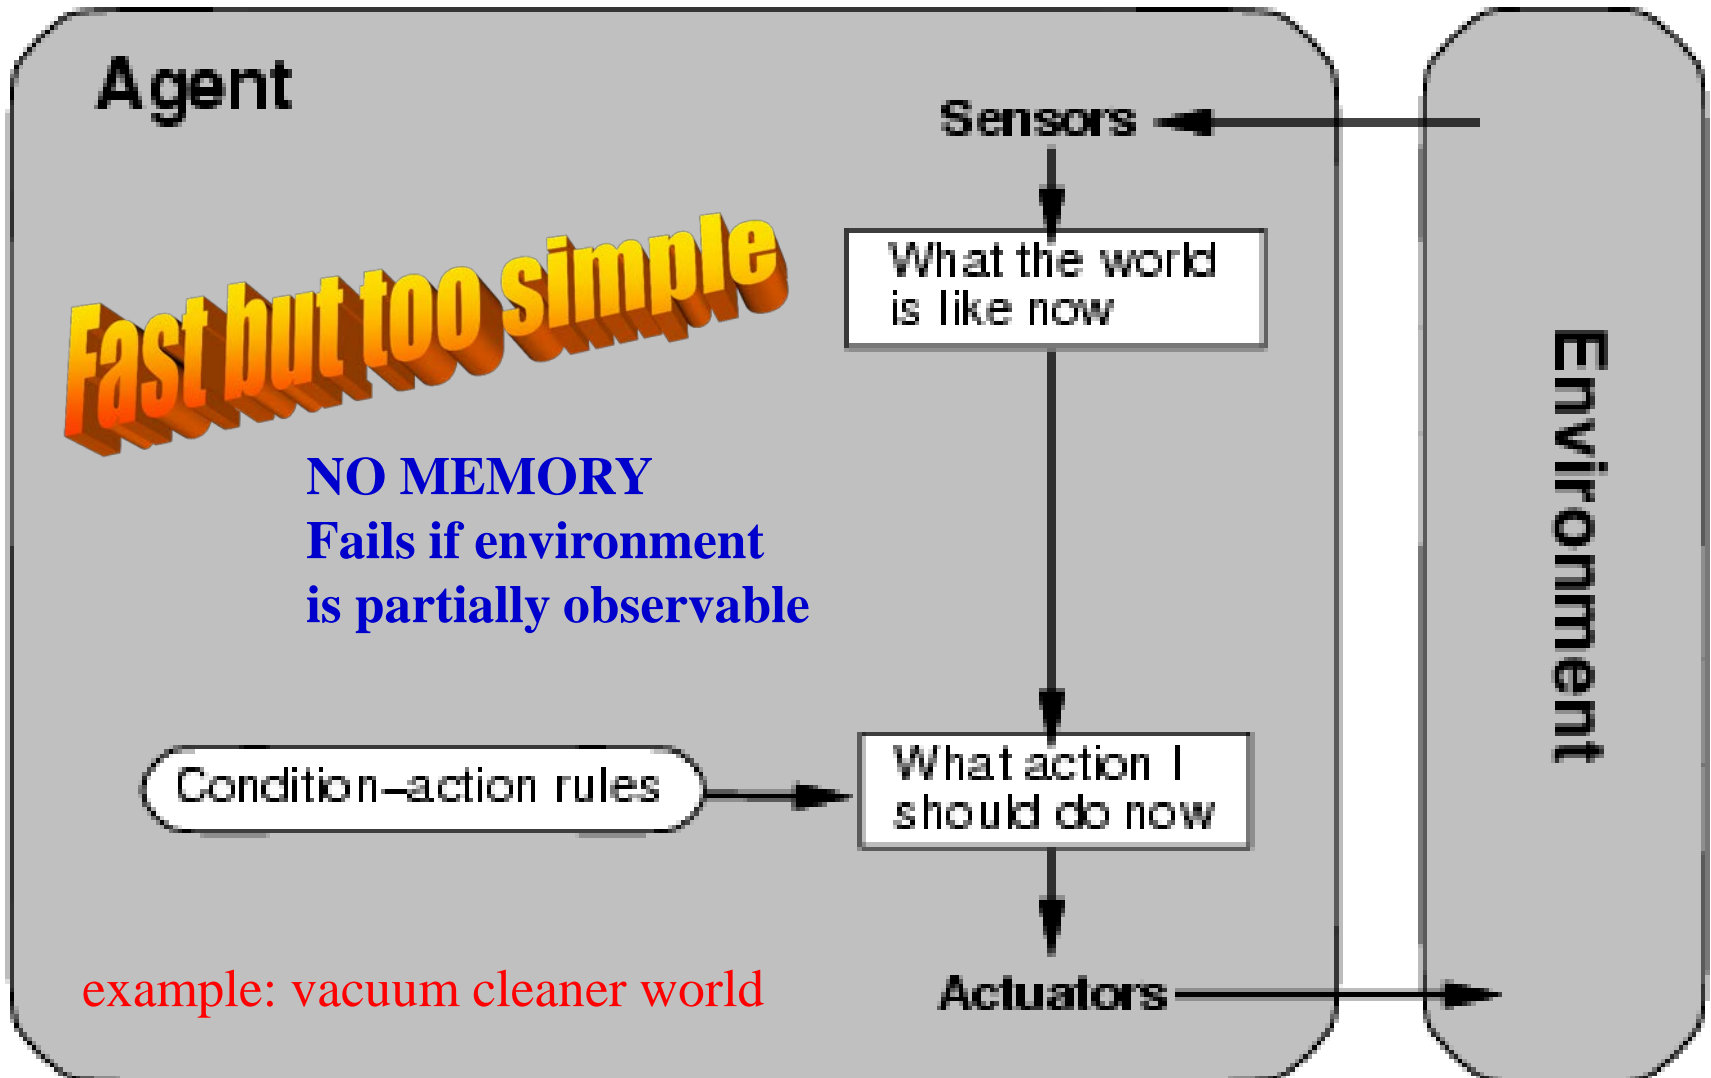
- Five basic types in order of increasing generality:
- **Table Driven** agents
- **Simple reflex** agents
- **Model-based** reflex agents
- **Goal-based** agents
- **Utility-based** agents

Table Driven Agent.

current state of decision process



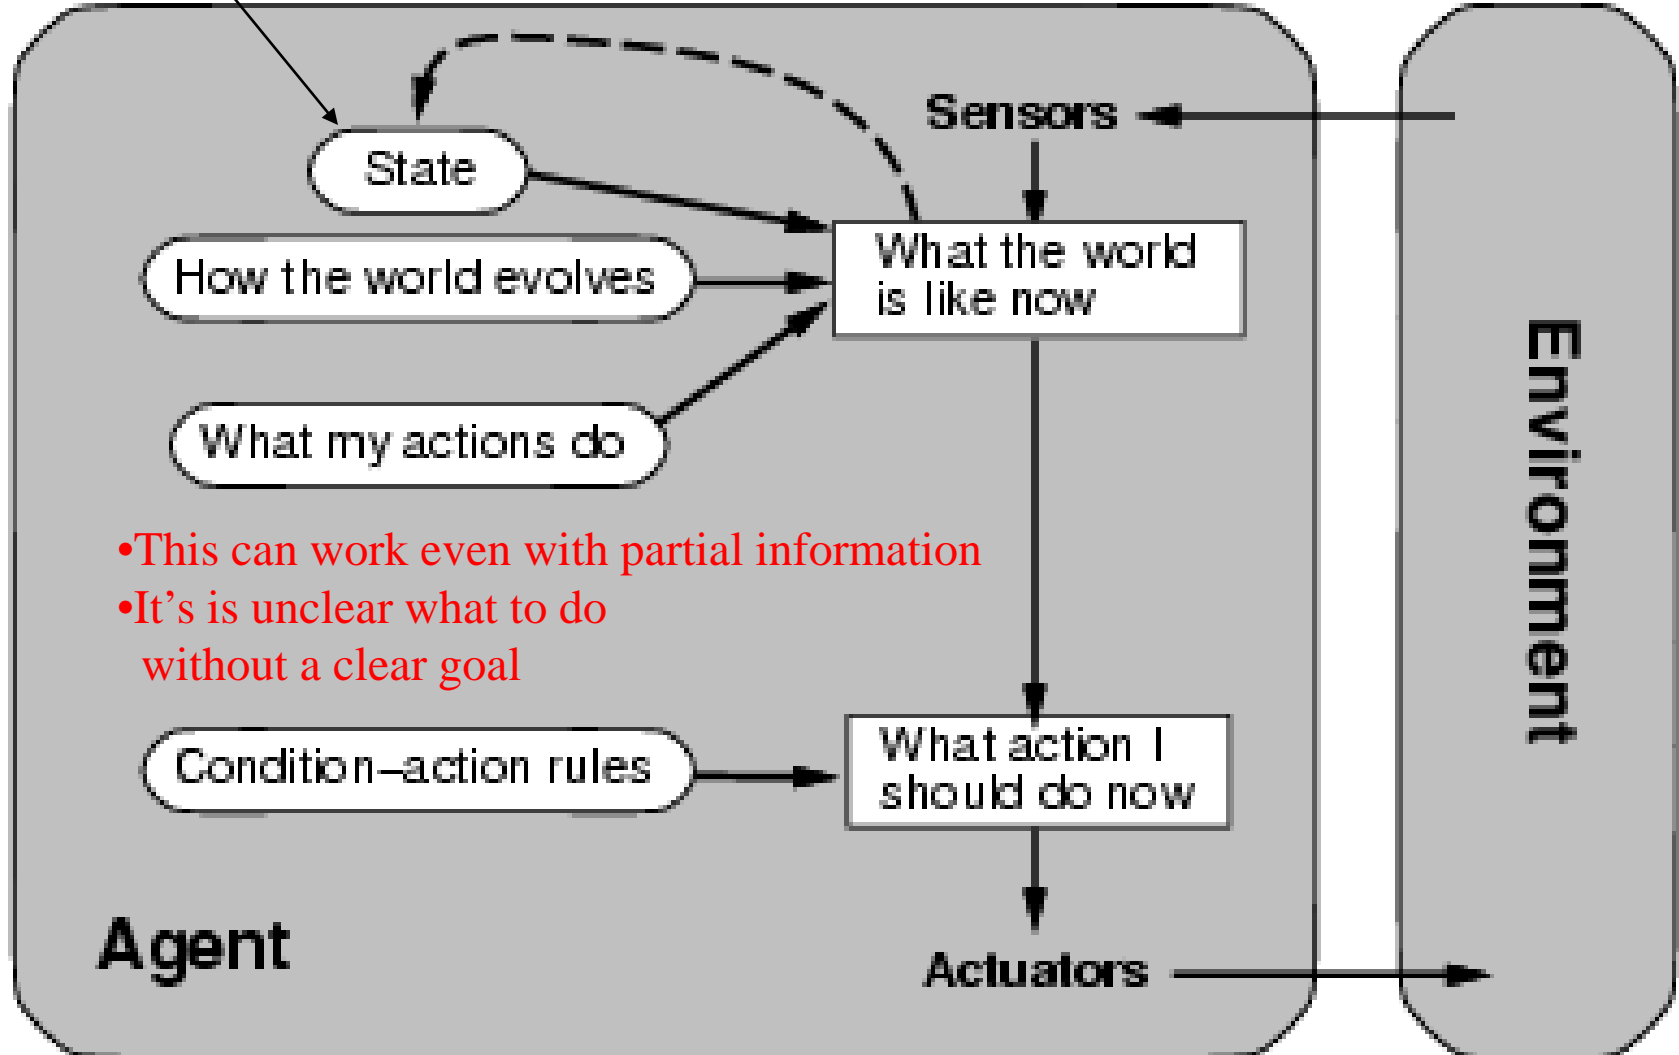
Simple reflex agents



Model-based reflex agents

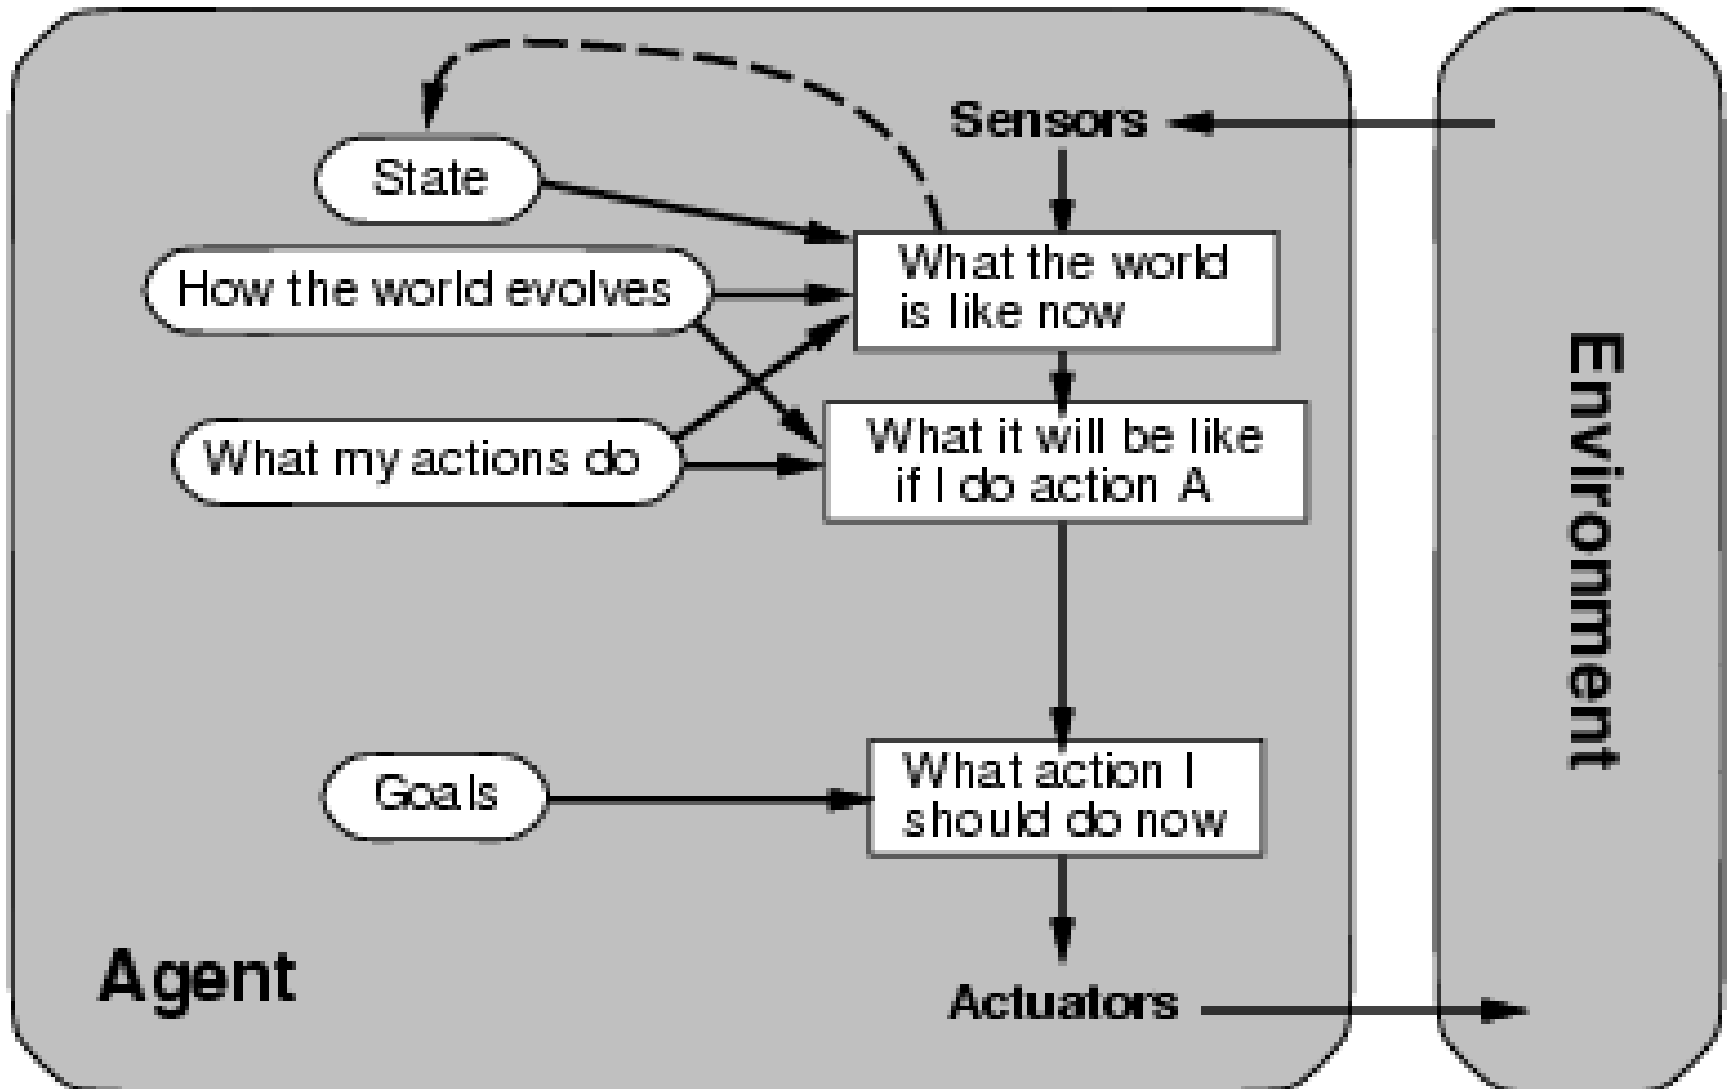
description of
current world state

Model the state of the world by:
modeling how the world changes
how it's actions change the world



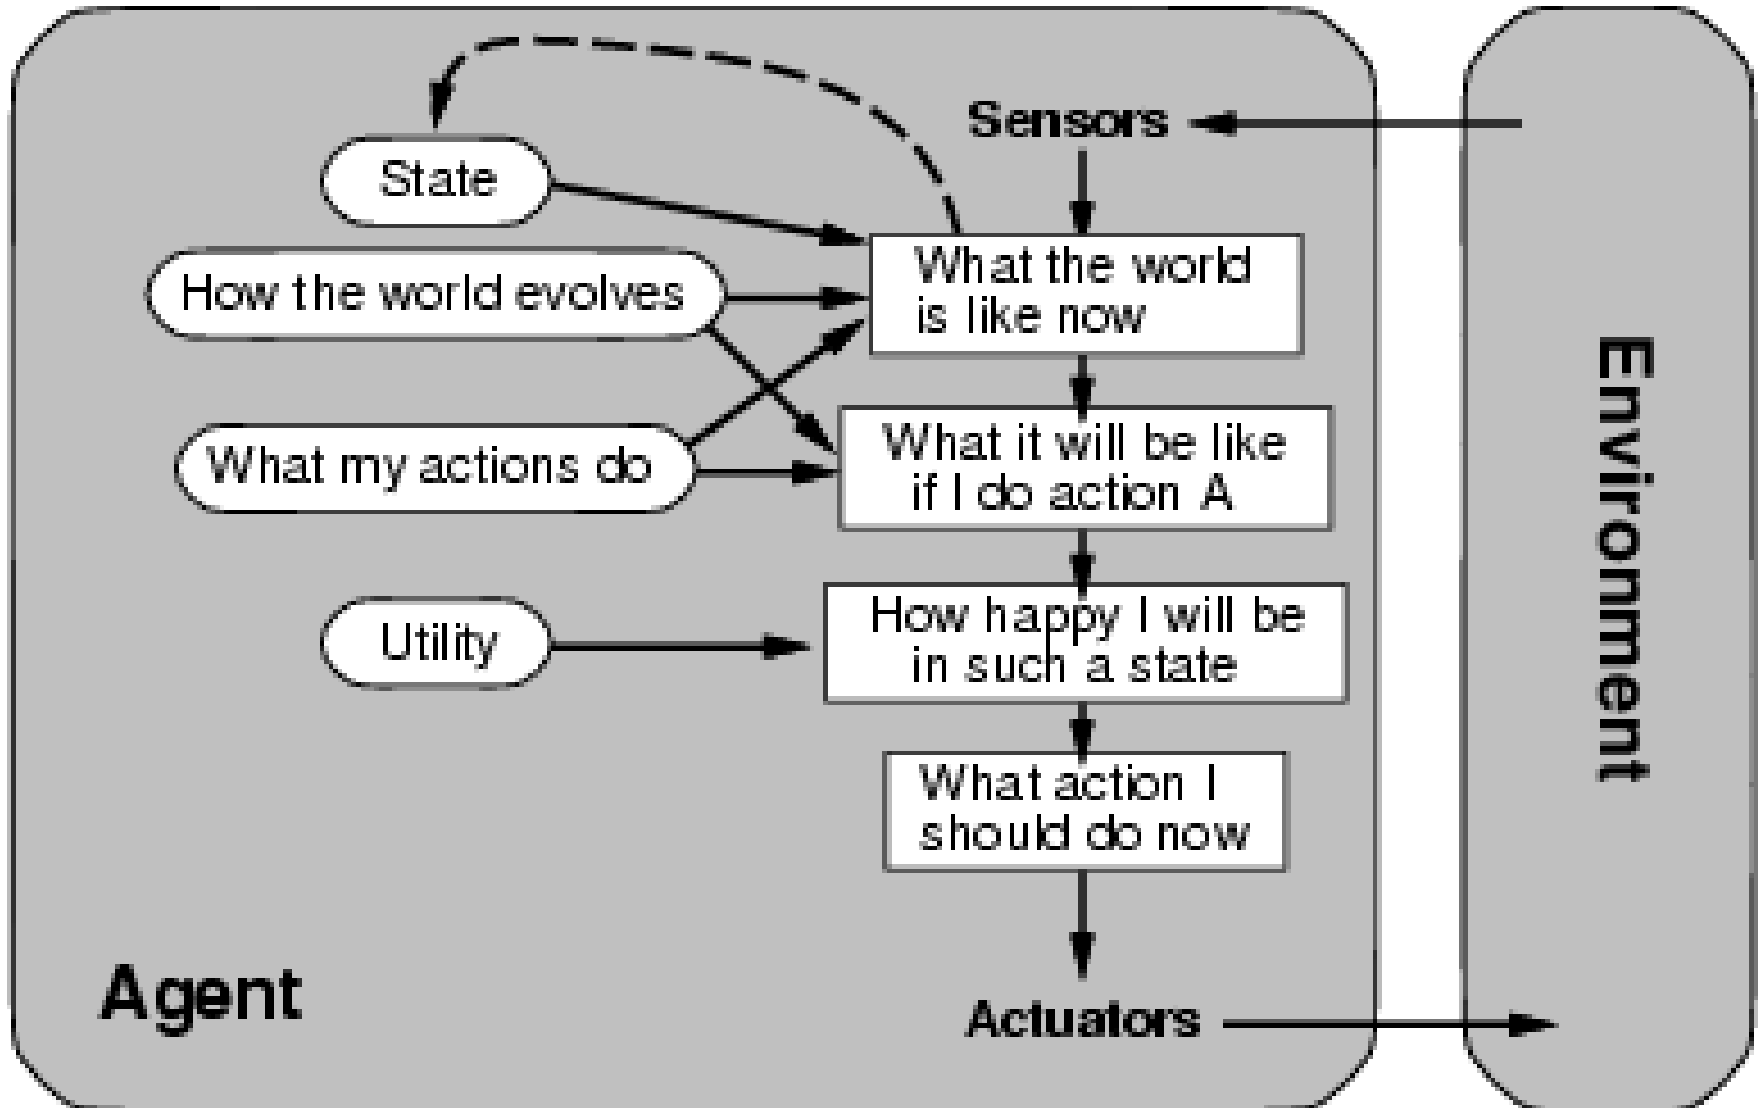
Goal-based agents

Goals provide reason to prefer one action over the other.
We need to predict the future: we need to plan & search



Utility-based agents

Some solutions to goal states are better than others.
Which one is best is given by a utility function.
Which combination of goals is preferred?



Learning agents

How does an agent improve over time?

By monitoring it's performance and suggesting

Performance standard **better modeling, new action rules, etc.**

