ARTIFICIAL INTELLIGENCE

- The term "Artificial Intelligence (AI)" was first coined in 1956.
- Difference in objective between AI and cognitive science (including philosophy and psychology) towards the study of intelligence:
  - AI:
    To understand what intelligence is *and* to build intelligent artifacts.
  - Cognitive science:
    To understand what intelligence is *only*.

AI and Computer Science

- Computer science:
  - To understand the physical and mathematical limits of computation.
  - To improve the ways computers are applied.
- Artificial intelligence:
  - To understand the mathematical and computational limits of intelligent behavior.
  - To improve the ability of computers to behave intelligently.

Some Definitions of AI

"The exciting new effort to make computers think ... *machines with minds*, in the full and literal sense" (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning ..." (Bellman, 1978)	"The study of mental faculties through the use of computational models" (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act" (Winston, 1992)
"The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991)	"A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkoff, 1990) "The branch of computer science that is concerned with the automation of intelligent behavior" (Luger and Stubblefield, 1993)

Some Definitions of AI (cont'd)

- Two dimensions:
  - Thinking (reasoning) vs. acting (behavior)
  - Human performance vs. rationality
- Four approaches:
  - The cognitive modeling approach (*thinking* in a *human* way)
  - The Turing Test approach (*acting* in a *human* way)
  - The laws of thought approach (*thinking* in a *rational* way)
  - The rational agent approach (*acting* in a *rational* way)

The Rational Agent Approach

- AI is the study and construction of rational agents.
- Advantage over the laws of thought approach:
  - *Less restrictive*, as making correct (logical) inferences is not necessarily the only means for achieving rationality (and hence intelligence).
- Advantage over human-centered approaches:
  - *More objective*, as the standard of rationality can be defined more generally.

Intelligent Agents

- Interactions between an agent and its environment:
  - *Perception* through sensors
  - *Action* through effectors

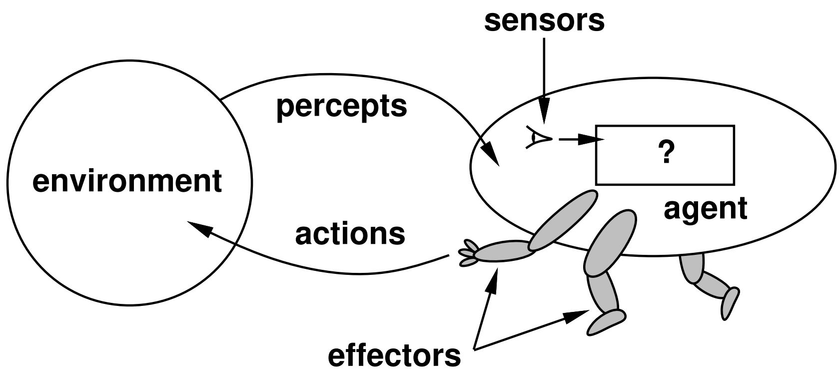

- Rational agents:
  - Choose to take the action that maximizes some "objective" performance measure, given the percept sequence seen so far.

Rationality

- Rationality depends on:
  - Performance measure
  - Percept sequence
  - Built-in knowledge about environment
  - Actions that can be performed
- Limited rationality:
  - Trying to do the best (i.e. most rational) under limited resources.
- Performing actions which lead to the perception of useful information (which otherwise would not be available) is an important aspect of rationality.
- Autonomous agents:
  - Do not rely entirely on built-in knowledge (i.e. not entirely pre-programmed).
  - Adapt to the environments through experience (i.e. learning).

Examples of Agents

Agent Type	Percepts	Actions	Goals	Environment
Medical diagnosis system	Symptoms, findings, patient's answers	Questions, tests, treatments	Healthy patient, minimize costs	Patient, hospital
Satellite image analysis system	Pixels of varying intensity, color	Print a categorization of scene	Correct categorization	Images from orbiting satellite
Part-picking robot	Pixels of varying intensity	Pick up parts and sort into bins	Place parts in correct bins	Conveyor belt with parts
Refinery controller	Temperature, pressure readings	Open, close valves; adjust temperature	Maximize purity, yield, safety	Refinery
Interactive English tutor	Typed words	Print exercises, suggestions, corrections	Maximize student's score on test	Set of students

An Example: A Texi Driver Agent

Agent Type	Percepts	Actions	Goals	Environment
Taxi driver	Cameras, speedometer, GPS, sonar, microphone	Steer, accelerate, brake, talk to passenger	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers

Agent Programs

- An *agent program* is a function that implements the agent mapping from percepts to actions.
- Different types of agent program are required to deal with different environment types effectively.

Agent Program Skeleton

function SKELETON-AGENT(*percept*) **returns** action
  static: *memory*, the agent's memory of the world

  memory ← UPDATE-MEMORY(*memory, percept*)
  action ← CHOOSE-BEST-ACTION(*memory*)
  memory ← UPDATE-MEMORY(*memory, action*)
  return *action*

A Simple Reflex Agent

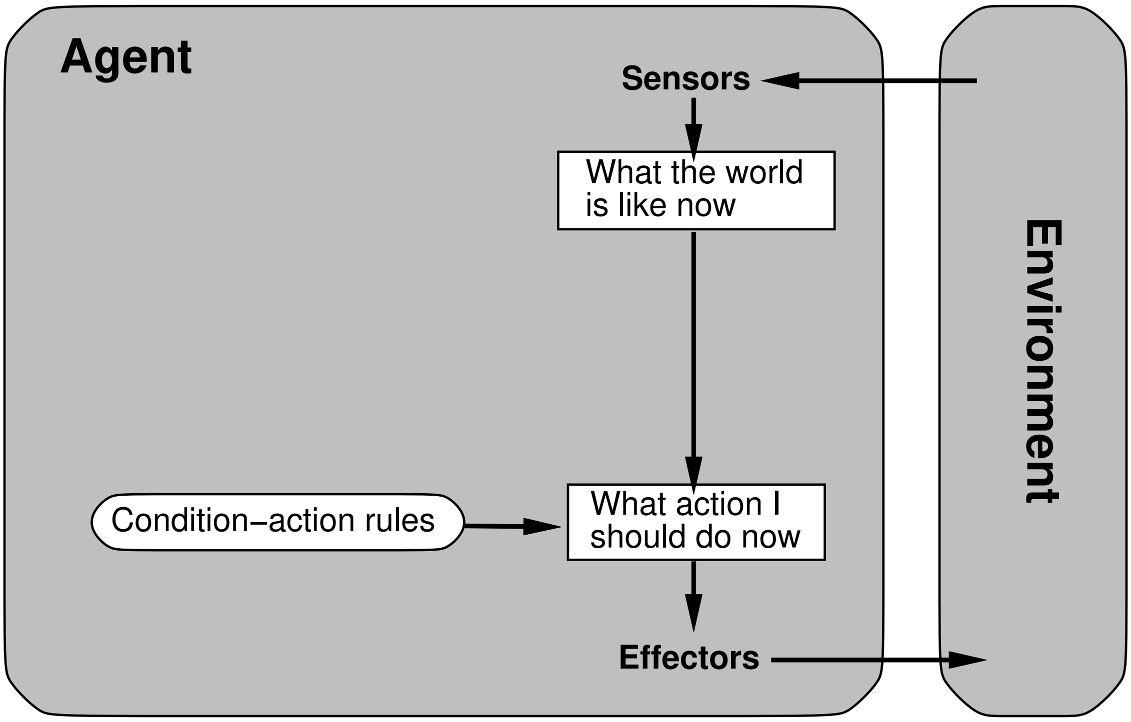

A Reflex Agent with Internal State

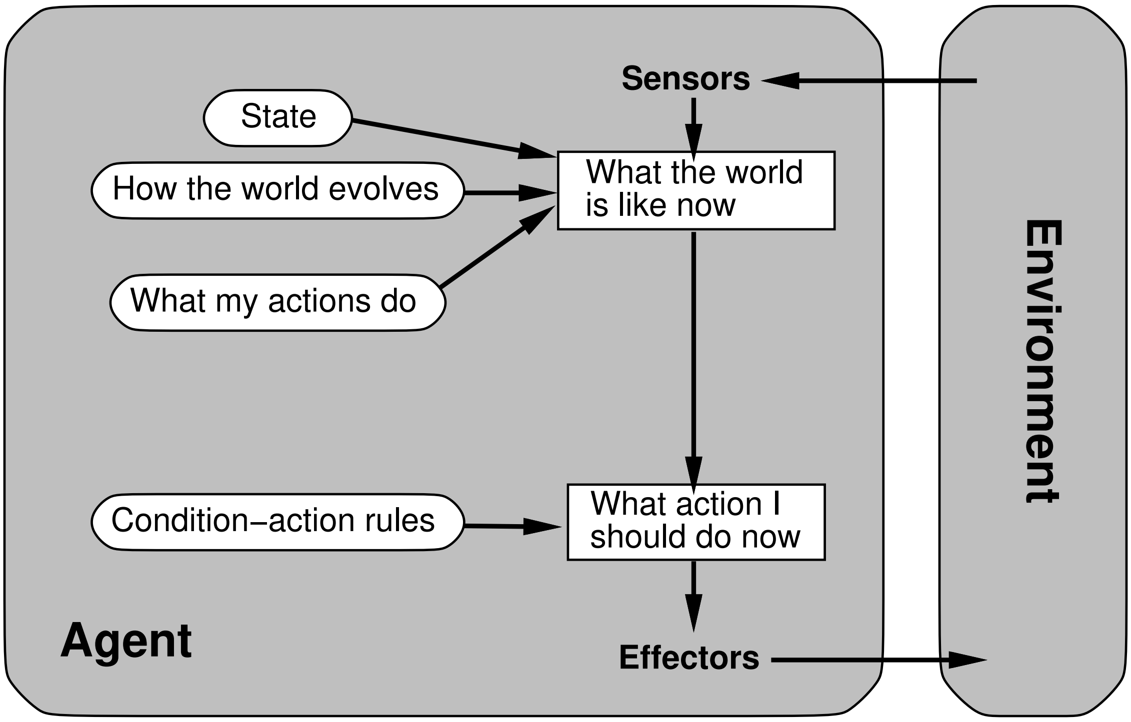

An Agent with Explicit Goals

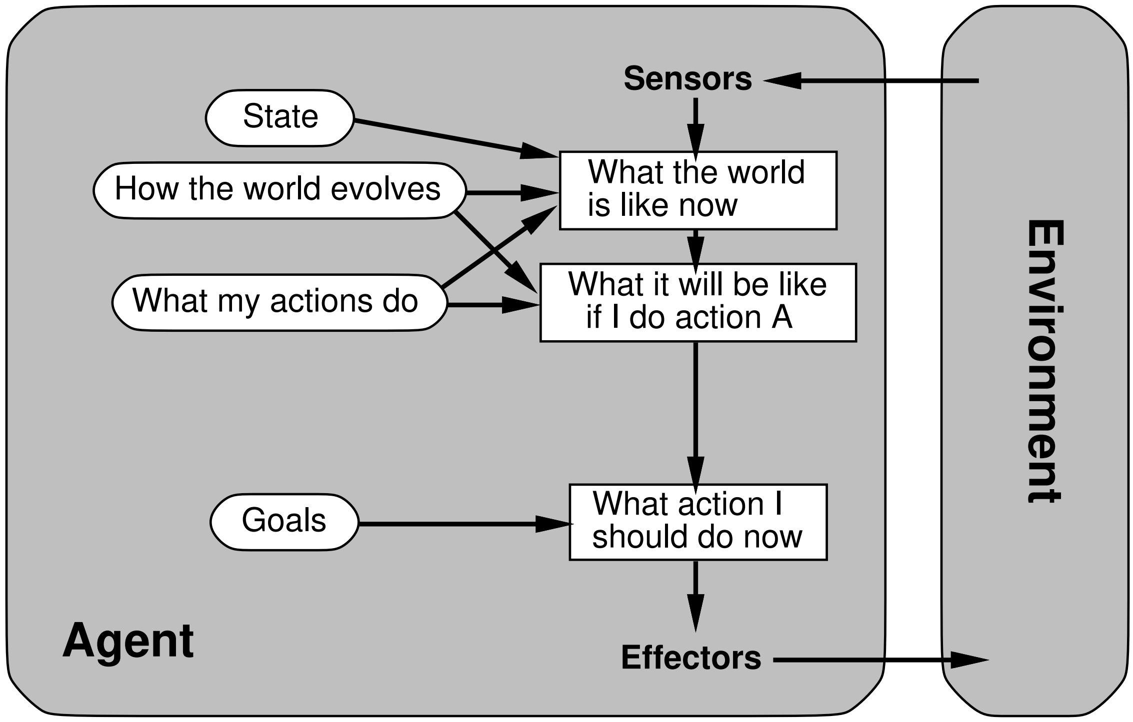

A Complete Utility-based Agents

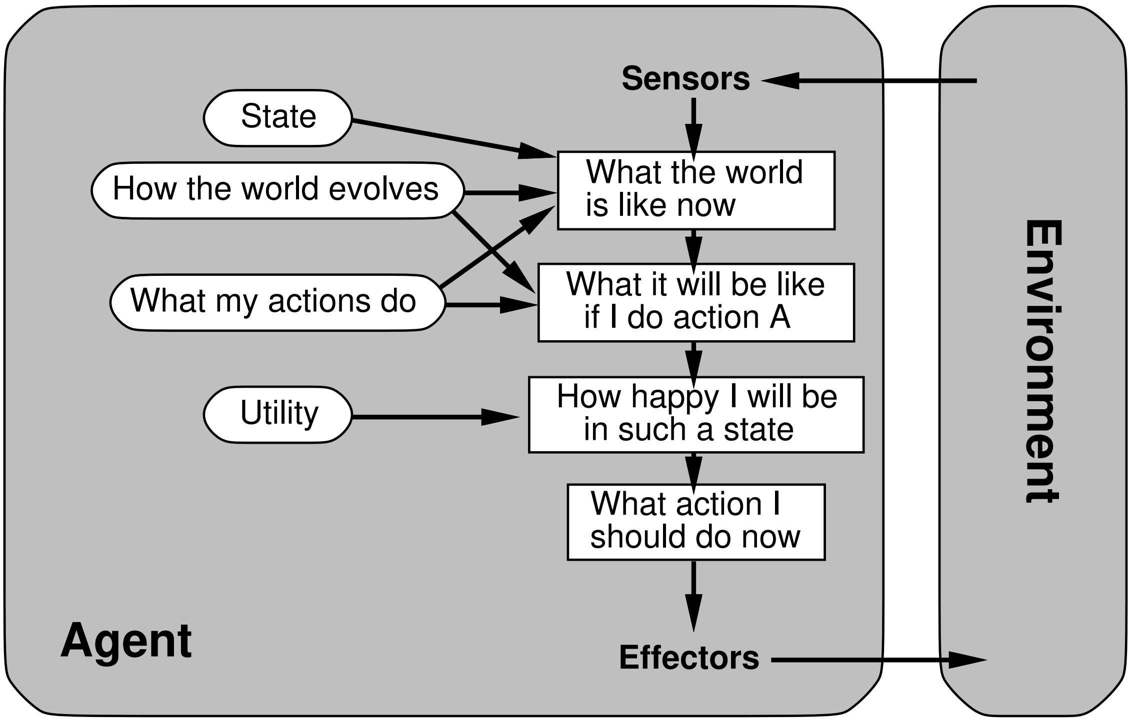

Properties of Environment

- Accessible vs. inaccessible
- Deterministic vs. nondeterministic
- Episodic vs. nonepisodic
- Static vs. dynamic
- Discrete vs. continuous

Examples of Environments

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Chess with a clock	Yes	Yes	No	Semi	Yes
Chess without a clock	Yes	Yes	No	Yes	Yes
Poker	No	No	No	Yes	Yes
Backgammon	Yes	No	No	Yes	Yes
Taxi driving	No	No	No	No	No
Medical diagnosis system	No	No	No	No	No
Image-analysis system	Yes	Yes	Yes	Semi	No
Part-picking robot	No	No	Yes	No	No
Refinery controller	No	No	No	No	No
Interactive English tutor	No	No	No	No	Yes