

CS 561: Artificial Intelligence

Instructor: Sofus A. Macskassy, macskass@usc.edu

TAs: Nadeesha Ranashinghe (nadeeshr@usc.edu)

William Yeoh (wyeoh@usc.edu)

Harris Chiu (chiciu@usc.edu)

Lectures: MW 5:00-6:20pm, OHE 122 / DEN

Office hours: By appointment

Class page: <http://www-rcf.usc.edu/~macskass/CS561-Spring2010/>

This class will use <http://www.uscden.net/> and class webpage

- Up to date information
- Lecture notes
- Relevant dates, links, etc.

Course material:

[AIMA] Artificial Intelligence: A Modern Approach,
by Stuart Russell and Peter Norvig. (2nd ed)

CS 561: Artificial Intelligence

Course overview: foundations of symbolic intelligent systems. Agents, search, problem solving, logic, representation, reasoning, symbolic programming, and robotics.

Prerequisites: programming principles, discrete mathematics for computing, software design and software engineering concepts. Good knowledge of C++ and STL required for programming assignments.

Grading:

- 20% for homeworks (4 homeworks, 5% each)
- 20% for programming projects (2 projects, 10% each)
- 30% for midterms (2 midterms, 15% each) +
- 30% for final (cumulative)

1 day late = 25% reduction in score

2 days late = 50% reduction in score

NOTE: You have 1 week from getting a homework/project/midterm to get it reviewed if you feel it was wrongly graded

Practical issues


- **Class mailing list:**
will be setup on the blackboard system
- **Homeworks:** See class web page on blackboard
 - Jan 25 - HW1 out
 - Feb 10 - HW1 due, HW2 out
 - Feb 22 - HW2 due
 - Mar 8 - HW3 out
 - Mar 22 - HW3 due, HW4 out
 - Apr 5 - HW4 due
- **Projects:** See class web page on blackboard
 - Feb 1 - Project 1 out
 - Mar 8 - Project 1 due, Project 2 out
 - Apr 19 - Project 2 due
- **Exams:**
 - Mar 1 – midterm 1 (in class)
 - Apr 12 – midterm 2 (in class)
 - May 5 – final (room TBA)

Practical issues

- **Grading will be based on absolute scores**
 - A 90.0%
 - A- 87.5%
 - B+ 85.0%
 - B 80.0%
 - B- 77.5%
 - C+ 75.0%
 - C 70.0%
- **Exams will be open book and open notes**

Last Time: Acting Humanly: The Full Turing Test

- Alan Turing's 1950 article *Computing Machinery and Intelligence* discussed conditions for considering a machine to be intelligent
 - "Can machines think?" \leftrightarrow "Can machines behave intelligently?"
 - The Turing test (The Imitation Game): Operational definition of intelligence.



- Computer needs to possess: Natural language processing, Knowledge representation, Automated reasoning, and Machine learning
- Problem: 1) Turing test is not reproducible, constructive, and amenable to mathematic analysis. 2) What about physical interaction with interrogator and environment?
- Total Turing Test: Requires physical interaction and needs perception and actuation.

This time: Outline [AIMA Ch. 2]

- Intelligent Agents (IA)
- Environment types
- IA Behavior
- IA Structure
- IA Types

What is an (Intelligent) Agent?

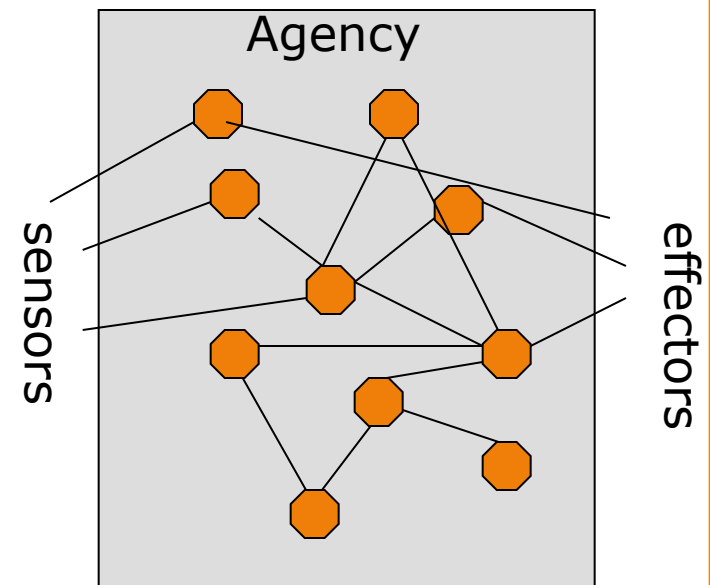
- An over-used, over-loaded, and misused term.
- Anything that can be *viewed* as **perceiving** its **environment** through **sensors** and **acting** upon that environment through its **actuators** to maximize progress towards its **goals**.

What is an (Intelligent) Agent?

- **PAGE** (Percepts, Actions, Goals, Environment)
- Task-specific & specialized: well-defined goals and environment
- The notion of an agent is meant to be a tool for analyzing systems,
 - It is not a different hardware or new programming languages

Intelligent Agents and Artificial Intelligence

- **Example:** Human mind as network of thousands or millions of agents working in parallel. To produce real artificial intelligence, this school holds, we should build computer systems that also contain many agents and systems for arbitrating among the agents' competing results.
- Distributed decision-making and control
- Challenges:
 - Action selection: What next action to choose
 - Conflict resolution

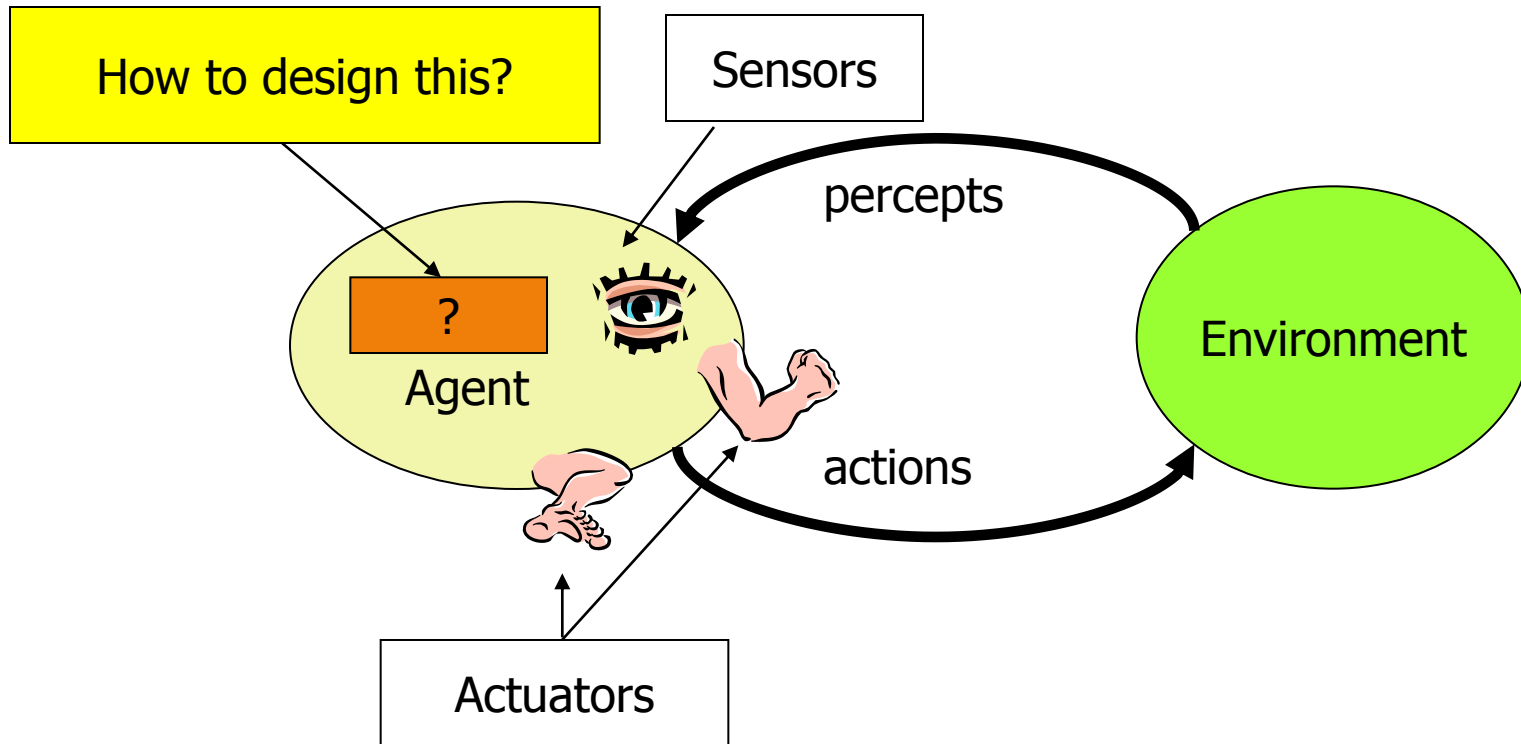


Agent Types

We can split agent research into two main strands:

- Distributed Artificial Intelligence (DAI) – Multi-Agent Systems (MAS) (1980 – 1990)
- Much broader notion of "agent" (1990's – present)
 - interface, reactive, mobile, information

Rational Agents



Agents include humans, robots, softbots, thermostats, etc.
The **agent function** maps from percept histories to actions:
 $f : P^* \rightarrow A$
The **agent program** runs on the physical **architecture** to produce f

A Windshield Wiper Agent (Cont'd)

How do we design an agent that can wipe the windshields when needed?

- Goals: ?
- Percepts: ?
- Sensors: ?
- Actuators: ?
- Actions: ?
- Environment: ?

A Windshield Wiper Agent (Cont'd)

How do we design an agent that can wipe the windshields when needed?

- Goals: Keep windshields clean & maintain visibility
- Percepts: Raining, Dirty
- Sensors: Camera (moist sensor)
- Actuators: Wipers (left, right, back)
- Actions: Off, Slow, Medium, Fast
- Environment: Inner city, freeways, highways, weather ...

Interacting Agents

Collision Avoidance Agent (CAA)

- Goals: Avoid running into obstacles
- Percepts: ?
- Sensors: ?
- Actuators: ?

- Actions: ?

- Environment: Freeway

Lane Keeping Agent (LKA)

- Goals: Stay in current lane
- Percepts: ?
- Sensors: ?
- Actuators: ?
- Actions: ?
- Environment: Freeway

Interacting Agents

Collision Avoidance Agent (CAA)

- Goals: Avoid running into obstacles
- Percepts: Obstacle distance, velocity, trajectory
- Sensors: Vision, proximity sensing
- Actuators: Steering Wheel, Accelerator, Brakes, Horn, Headlights
- Actions: Steer, speed up, brake, blow horn, signal (headlights)
- Environment: Freeway

Lane Keeping Agent (LKA)

- Goals: Stay in current lane
- Percepts: Lane center, lane boundaries
- Sensors: Vision
- Actuators: Steering Wheel, Accelerator, Brakes
- Actions: Steer, speed up, brake
- Environment: Freeway

Conflict Resolution by Action Selection Agents

- **Override:** CAA overrides LKA
- **Arbitrate:** if Obstacle is Close then CAA
else LKA
- **Compromise:** Choose action that satisfies both agents
- Any combination of the above
- **Challenges:** Doing the right thing

The Right Thing = The Rational Action

- **Rational Action:** The action that maximizes the expected value of the performance measure given the percept sequence to date
 - Rational = Best ?
 - Rational = Optimal ?
 - Rational = Omniscience ?
 - Rational = Clairvoyant ?
 - Rational = Successful ?

The Right Thing = The Rational Action


- **Rational Action:** The action that maximizes the expected value of the performance measure given the percept sequence to date
 - Rational = Best Yes, to the best of its knowledge
 - Rational = Optimal Yes, to the best of its abilities (incl. its constraints)
 - Rational \neq Omniscience
 - Rational \neq Clairvoyant
 - Rational \neq Successful

Behavior and performance of IAs

- **Perception** (sequence) to **Action Mapping**: $f : P^* \rightarrow A$
 - **Ideal mapping**: specifies which actions an agent ought to take at any point in time
 - **Description**: Look-Up-Table, Closed Form, etc.
- **Performance measure**: a *subjective* measure to characterize how successful an agent is (e.g., speed, power usage, accuracy, money, etc.)
- (degree of) **Autonomy**: to what extent is the agent able to make decisions and take actions on its own?

Look up table

Distance	Action
10	No action
5	Turn left 30 degrees
2	Stop



Closed form

- Output (degree of rotation) = $F(\text{distance})$
- E.g., $F(d) = 10/d$ (distance cannot be less than $1/10$)

How is an Agent different from other software?

- Agents are **autonomous**, that is, they act on behalf of the user
- Agents contain some level of **intelligence**, from fixed rules to learning engines that allow them to adapt to changes in the environment
- Agents don't only act **reactively**, but sometimes also **proactively**

How is an Agent different from other software?

- Agents have **social ability**, that is, they communicate with the user, the system, and other agents as required
- Agents may also **cooperate** with other agents to carry out more complex tasks than they themselves can handle
- Agents may **migrate** from one system to another to access remote resources or even to meet other agents

Environment Types

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

Environment Types

- Characteristics
 - Accessible vs. inaccessible
 - Sensors give access to **complete** state of the environment.
 - Deterministic vs. nondeterministic
 - The next state can be determined based on the current state and the action.
 - Episodic vs. nonepisodic (Sequential)
 - Episode: each perceive and action pairs
 - The quality of action does not depend on the previous episode.

Environment Types

- Characteristics
 - Hostile vs. friendly
 - Static vs. dynamic
 - Dynamic if the environment changes during deliberation
 - Discrete vs. continuous
 - Chess vs. driving

Environment types

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Operating System					
Virtual Reality					
Office Environment					
Mars					

The environment types largely determine the agent design.

Environment types

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Operating System	Yes	Yes	No	No	Yes
Virtual Reality					
Office Environment					
Mars					

The environment types largely determine the agent design.

Environment types

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Operating System	Yes	Yes	No	No	Yes
Virtual Reality	Yes	Yes	Yes/no	No	Yes/no
Office Environment					
Mars					

The environment types largely determine the agent design.

Environment types

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Operating System	Yes	Yes	No	No	Yes
Virtual Reality	Yes	Yes	Yes/no	No	Yes/no
Office Environment	No	No	No	No	No
Mars					

The environment types largely determine the agent design.

Environment types

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Operating System	Yes	Yes	No	No	Yes
Virtual Reality	Yes	Yes	Yes/no	No	Yes/no
Office Environment	No	No	No	No	No
Mars	No	Semi	No	Semi	No

The environment types largely determine the agent design.

Structure of Intelligent Agents

- Agent = architecture + program
- **Agent program:** the implementation of $f : P^* \rightarrow A$, the agent's perception-action mapping

```
function Skeleton-Agent(Percept) returns Action  
    memory ← UpdateMemory(memory, Percept)  
    Action ← ChooseBestAction(memory)  
    memory ← UpdateMemory(memory, Action)  
return Action
```


- **Architecture:** a device that can execute the agent program (e.g., general-purpose computer, specialized device, beobot, etc.)

Using a look-up-table to encode $f : P^* \rightarrow A$

- **Example:** Collision Avoidance

- Sensors: 3 proximity sensors
- Effectors: Steering Wheel, Brakes


- How to generate?
- How large?
- How to select action?



Using a look-up-table to encode $f : P^* \rightarrow A$

- **Example:** Collision Avoidance

- Sensors: 3 proximity sensors
- Effectors: Steering Wheel, Brakes



- **How to generate:** for each $p \in P_l \times P_m \times P_r$ generate an appropriate action, $a \in S \times B$

- **How large:** size of table = #possible percepts * #possible actions = $|P_l| |P_m| |P_r| |S| |B|$

E.g., $P = \{\text{close, medium, far}\}^3$

$A = \{\text{left, straight, right}\} \times \{\text{on, off}\}$

then size of table = $27 * 3 * 2 = 162$

- **How to select action?** Search.

Agent types

- Reflex agents
- Reflex agents with internal states
- Goal-based agents
- Utility-based agents

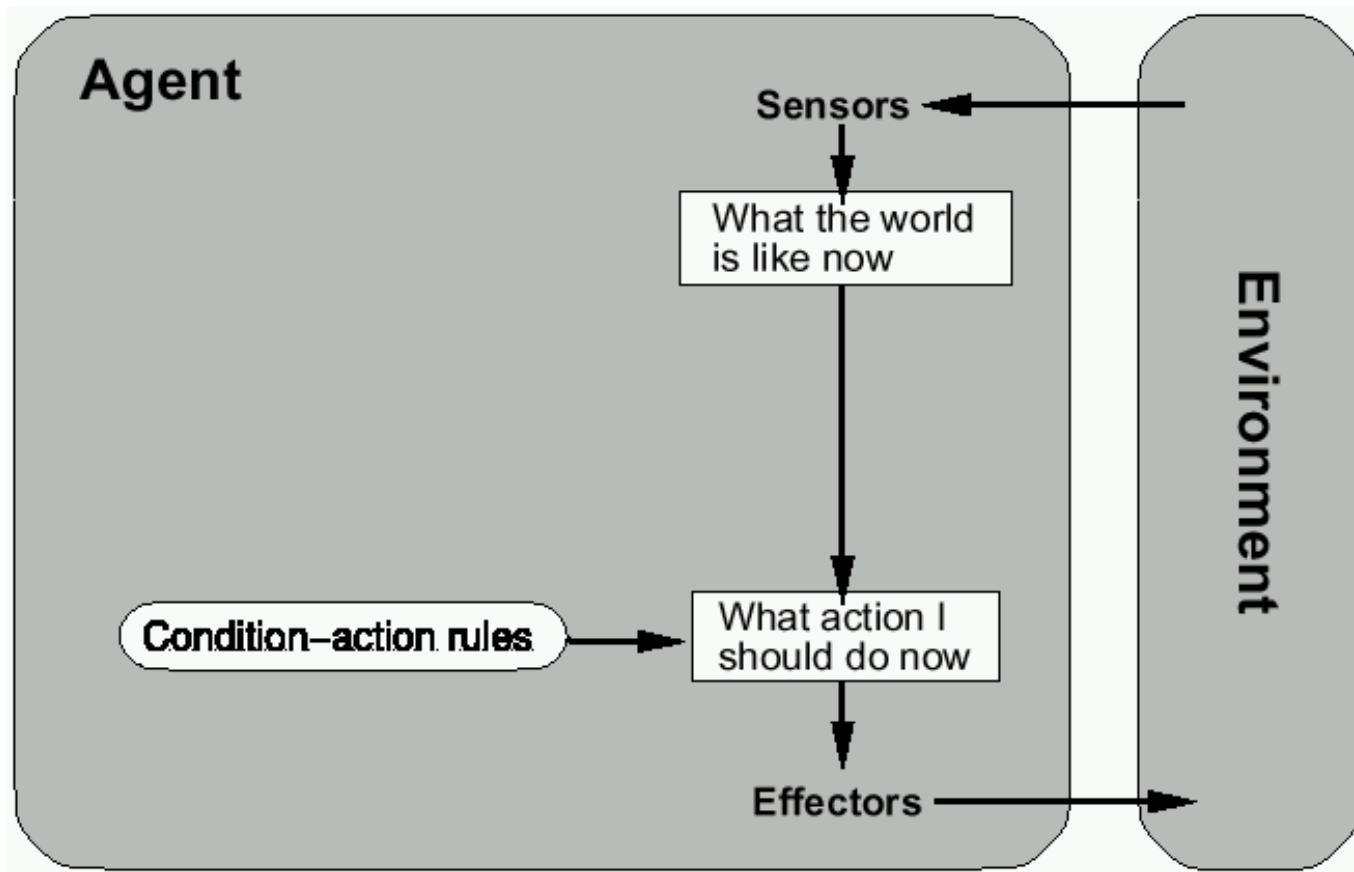
Agent types

- Reflex agents
 - Reactive: No memory
- Reflex agents with internal states
 - W/o previous state, may not be able to make decision
 - E.g. brake lights at night.
- Goal-based agents
 - Goal information needed to make decision

Agent types

- Utility-based agents
 - How well can the goal be achieved (degree of happiness)
 - What to do if there are conflicting goals?
 - Speed and safety
 - Which goal should be selected if several can be achieved?

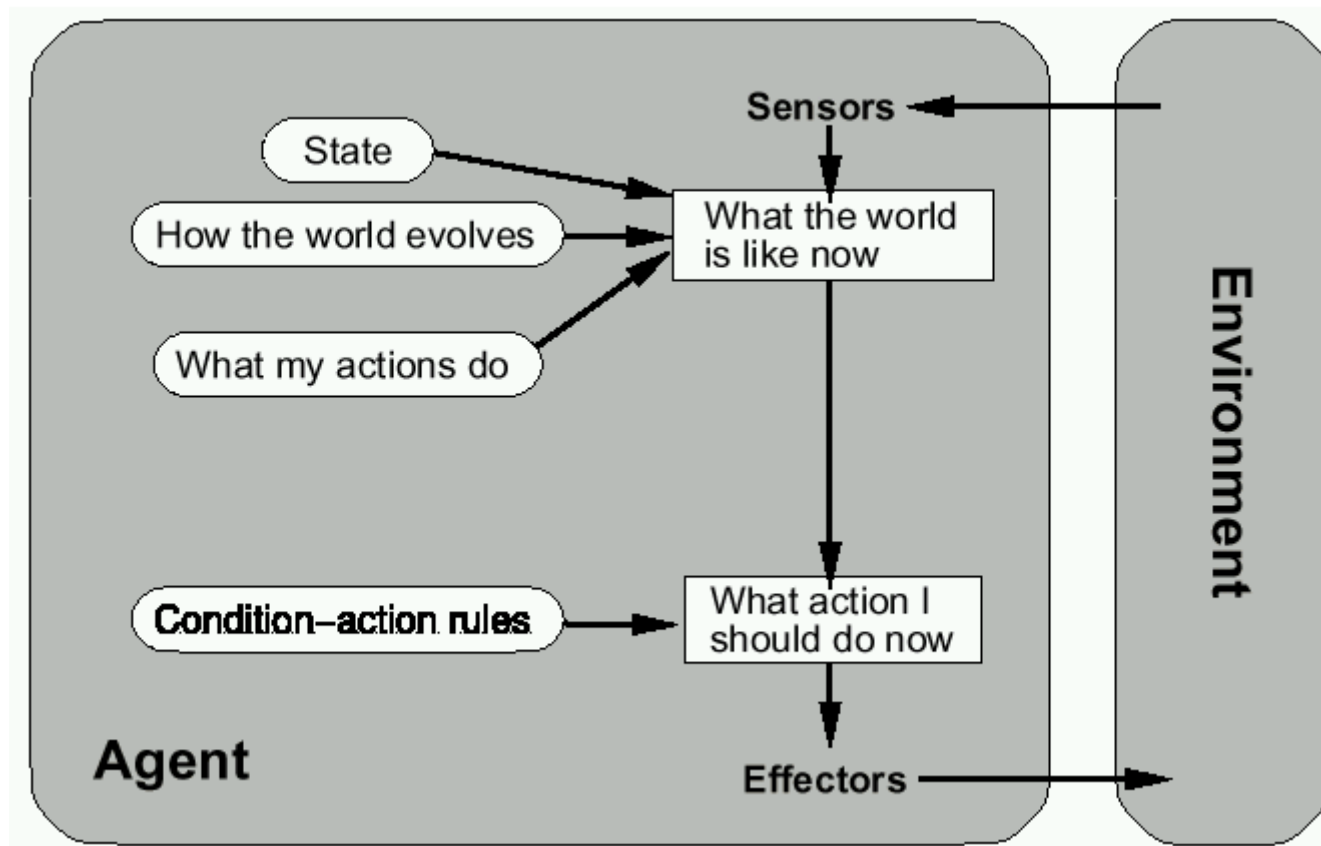
Reflex agents



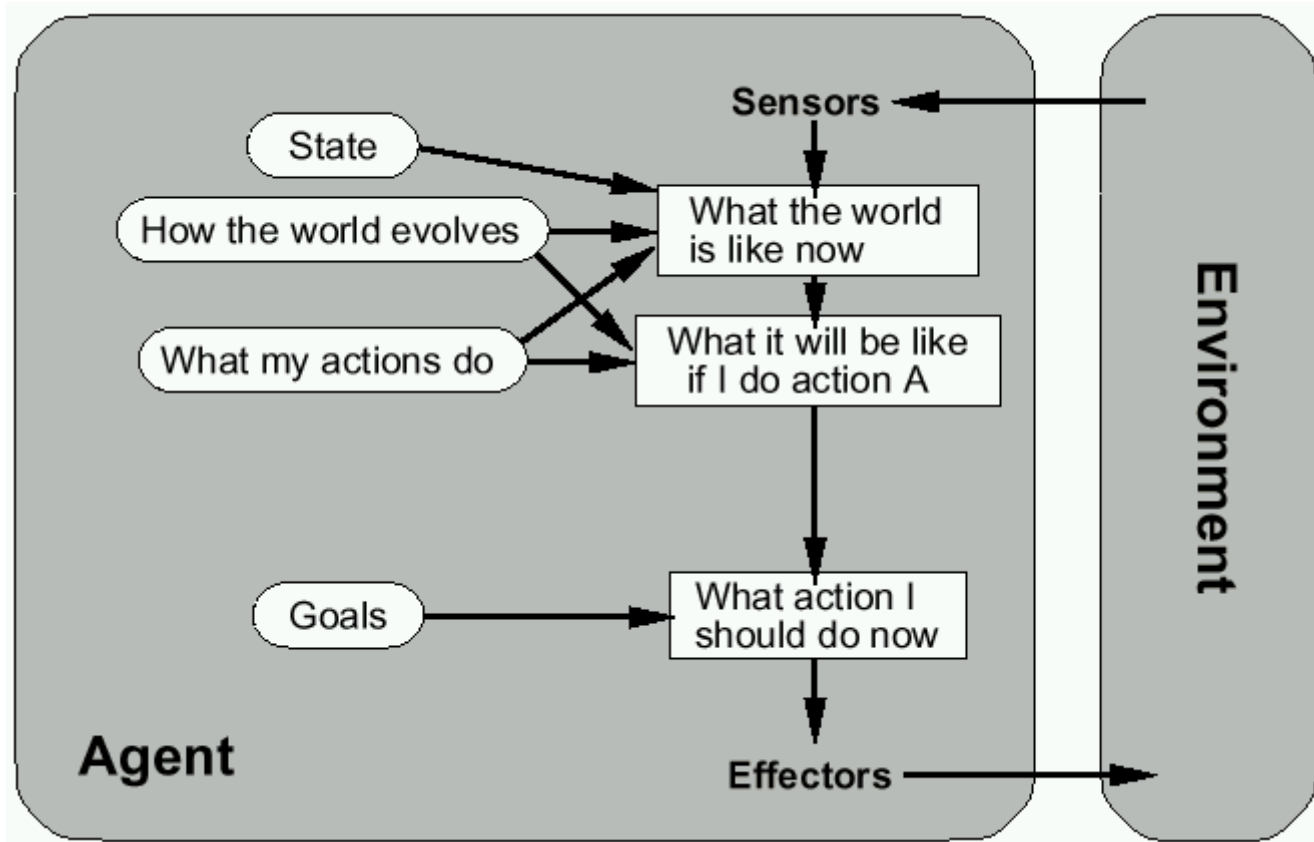
Reactive agents

- Reactive agents do not have internal symbolic models.
 - Act by stimulus-response to the current state of the environment.
 - Each reactive agent is simple and interacts with others in a basic way.
 - Complex patterns of behavior emerge from their interaction.
-
- **Benefits:** robustness, fast response time
 - **Challenges:** scalability, how intelligent?
and how do you debug them?

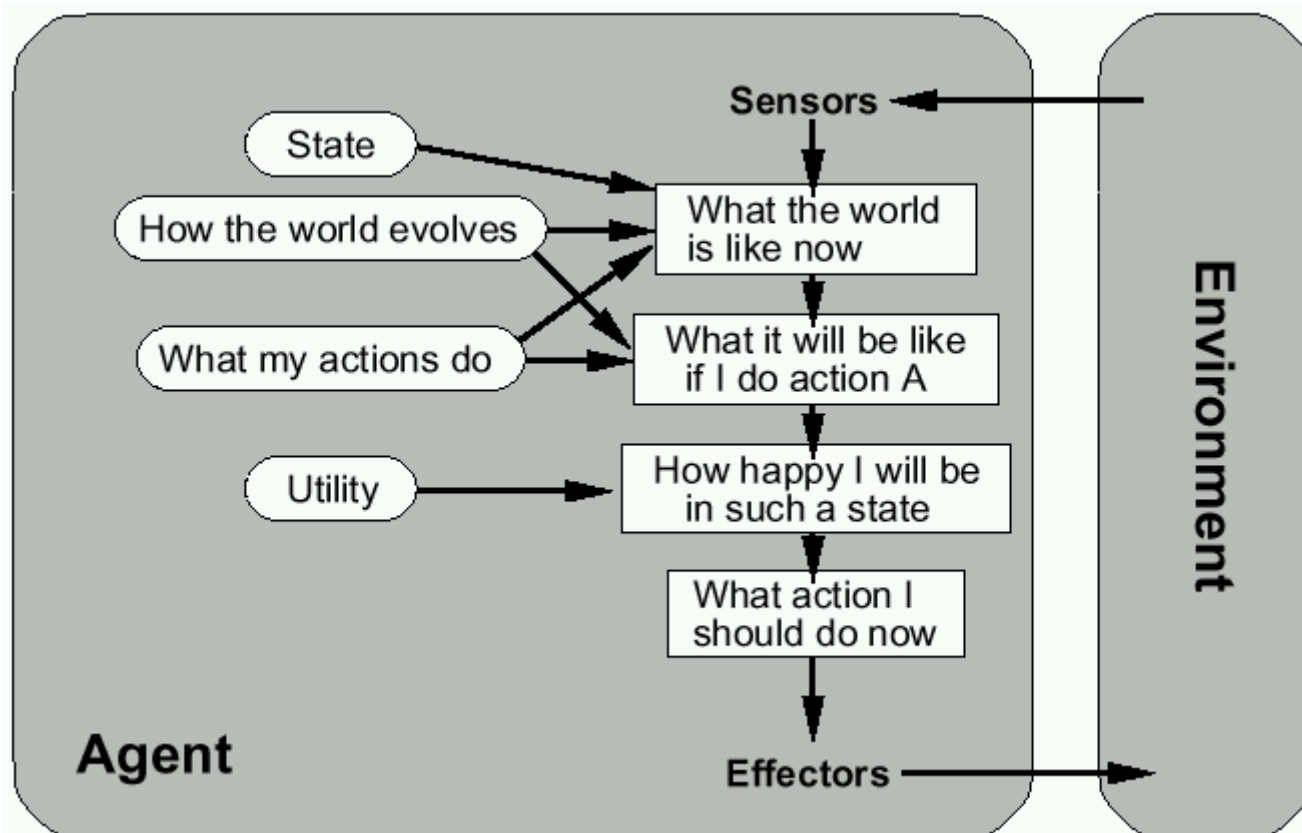
Reflex agents w/ state



Goal-based agents



Utility-based agents




Mobile agents

- Programs that can migrate from one machine to another.
- Execute in a platform-independent execution environment.
- Require agent execution environment (places).
- Mobility not necessary or sufficient condition for agenthood.
- Practical but non-functional advantages:
 - Reduced communication cost (eg, from PDA)
 - Asynchronous computing (when you are not connected)
- Two types:
 - One-hop mobile agents (migrate to one other place)
 - Multi-hop mobile agents (roam the network from place to place)
- Applications:
 - Distributed information retrieval.
 - Telecommunication network routing.

Mobile agents

- Programs that can migrate from one machine to another.
- Execute in a platform-independent execution environment.
- Require agent execution environment (places).
- Mobility not necessary or sufficient condition for agenthood.



A mail agent

Mobile agents

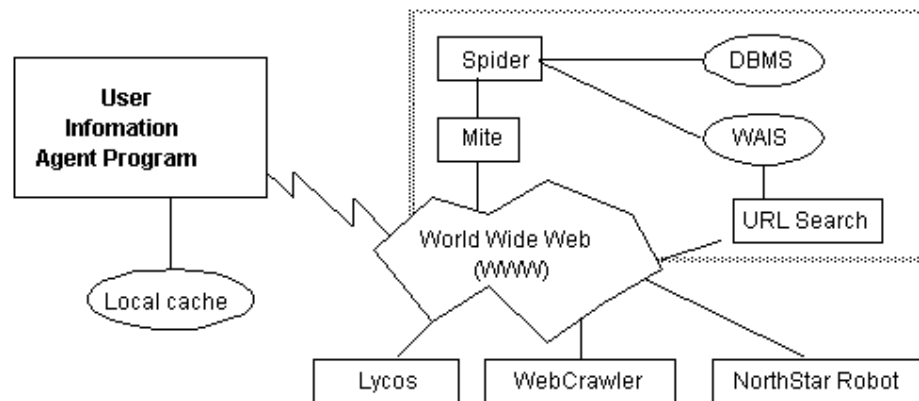
- Practical but non-functional advantages:
 - Reduced communication cost (e.g. from PDA)
 - Asynchronous computing (when you are not connected)
- Two types:
 - One-hop mobile agents (migrate to one other place)
 - Multi-hop mobile agents (roam the network from place to place)

Mobile agents

- Applications:
 - Distributed information retrieval.
 - Telecommunication network routing.

Information agents

- Manage the explosive growth of information.
- Manipulate or collate information from many distributed sources.
- Information agents can be mobile or static.
- Examples:
 - [BargainFinder](#) comparison shops among Internet stores for CDs
 - [FIDO](#) the Shopping Doggie (out of service)
 - [Internet Softbot](#) infers which internet facilities (finger, ftp, gopher) to use and when from high-level search requests.
- Challenge: ontologies for annotating Web pages



Summary

- **Intelligent Agents:**
 - Anything that can be *viewed* as **perceiving** its **environment** through **sensors** and **acting** upon that environment through its **effectors** to maximize progress towards its **goals**.
 - PAGE (Percepts, Actions, Goals, Environment)
 - Described as a Perception (sequence) to Action Mapping: $f : P^* \rightarrow A$
 - Using look-up-table, closed form, etc.
- **Agent Types:** Reflex, state-based, goal-based, utility-based
- **Rational Action:** The action that maximizes the expected value of the performance measure given the percept sequence to date