Chapter 2
Agents & Environments

Outline

- Agents and environments
- Rationality
- PEAS specification
- Environment types
- Agent types

Agents

- An agent is any entity that can perceive its environment through sensors and act upon that environment through actuators
- Human agent:
  - Sensors: Eyes, ears, and other organs
  - Actuators: Hands, legs, mouth, etc.
- Robotic agent:
  - Sensors: Cameras, laser range finders, etc.
  - Actuators: Motorized limbs, wheels, etc.

3

Types of Agents

- Immobots (Immobile Robots)
  - Intelligent buildings
  - Intelligent forests
  - Autonomous spacecraft

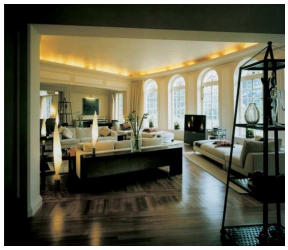

- Softbots
  - Jango (early softbot for shopping)
  - Microsoft Clippy
  - Askjeeves.com (now Ask.com)
  - Expert Systems
    - Cardiologist

Intelligent Agents

- Have sensors and actuators (effectors)
- Implement mapping from percept sequence to actions

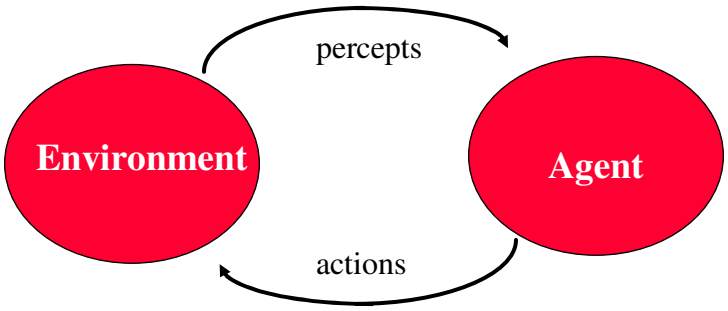

- Maximize a Performance Measure

Performance Measures

- Performance measure = An objective criterion for success of an agent's behavior
- E.g., vacuum cleaner agent
  performance measure: amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

6

Rational Agent

"For each possible percept sequence, does whatever action is expected to maximize its performance measure on the basis of evidence perceived so far and built-in knowledge."

- Rationality vs. omniscience
  - Rationality maximizes expected performance
  - Omniscience maximizes actual performance (but impossible to achieve in reality)
- Need to use information gathering actions and learning

Autonomy

A rational agent is autonomous if it can learn to compensate for partial or incorrect prior knowledge

Why is this important?

Task Environments

- The "task environment" for an agent is comprised of PEAS (Performance measure, Environment, Actuators, Sensors)
- E.g., Consider the task of designing an automated taxi driver:
  - Performance measure = ?
  - Environment = ?
  - Actuators = ?
  - Sensors = ?

9

PEAS

- PEAS for 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

10

PEAS

- PEAS for 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)

11

Properties of Environments

- Observability: full *vs.* partial
  Sensors detect all aspects of state of environment relevant to choice of action?
- Deterministic *vs.* stochastic
  Next state completely determined by current state and action?
- Episodic *vs.* sequential
  Current action independent of previous actions?
- Static *vs.* dynamic
  Can environment change over time?
- Discrete *vs.* continuous
  State of environment, time, percepts, and actions discrete or continuous-valued?
- Single *vs.* multiagent

Properties of Environments

- Observability: full *vs.* partial
- Deterministic *vs.* stochastic
- Episodic *vs.* sequential
- Static *vs.* dynamic
- Discrete *vs.* continuous
- Single *vs.* multiagent

  - **Chess**
  - **Poker**
  - **Coffee delivery mobile robot**

Agent Functions and Agent Programs

- An agent's behavior can be *described* by an agent function mapping percept sequences to actions taken by the agent
- An *implementation* of an agent function running on the agent architecture (e.g., a robot) is called an agent program
- Our goal: Develop concise agent programs for implementing rational agents

14

Example

Vacuum-cleaner world

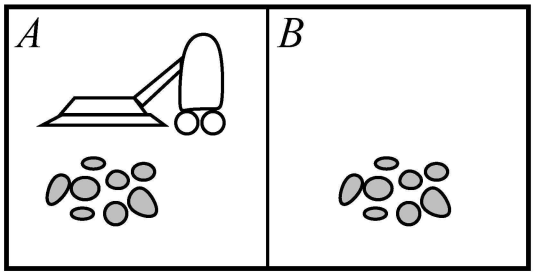

Percepts: location and contents, e.g., $[A, Dirty]$

Actions: $Left$, $Right$, $Suck$, $NoOp$

15

Example: vacuum world state space graph

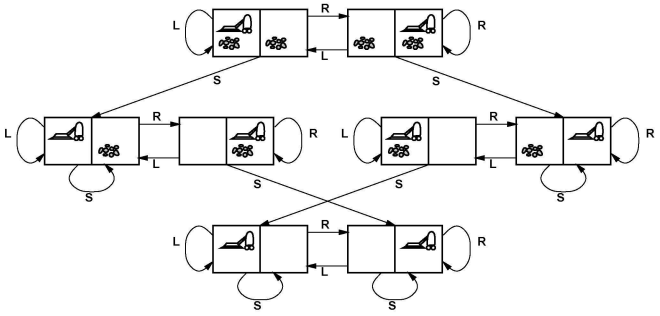

How should the agent be designed if…

- It has location and dirt sensors, but no internal state?
- It has no sensors, but knows the starting state?
- It has no sensors, and does not know the starting state?

16

Implementing Rational Agents

- Table lookup based on percept sequences
  Infeasible
- Agent programs:
  Simple reflex agents
  Agents with memory
  - Reflex agent with internal state
  - Goal-based agents
  - Utility-based agents

Simple Reflex Agents

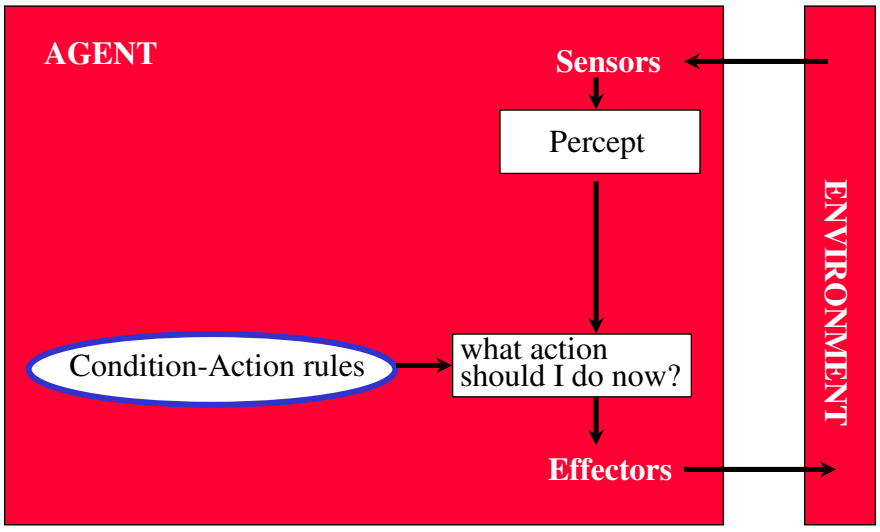

Simple Reflex Agents

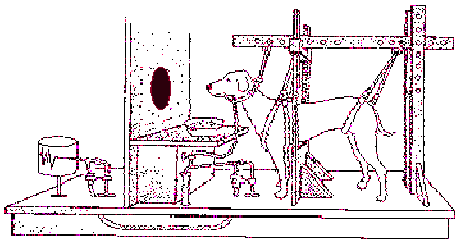

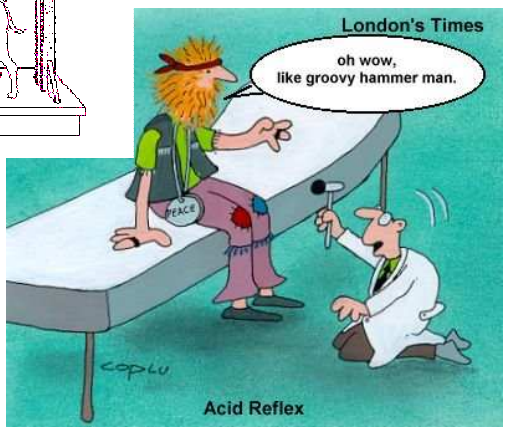

Reflex Agent with Internal State

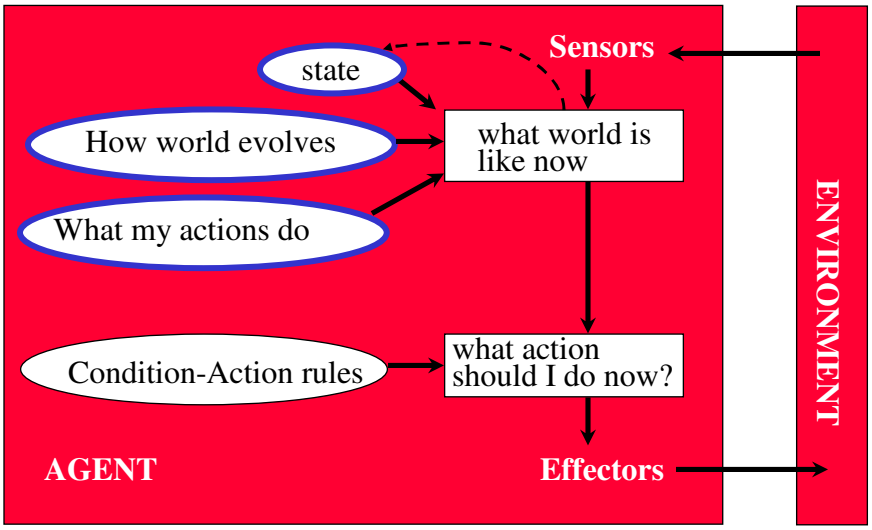

Goal-Based Agents

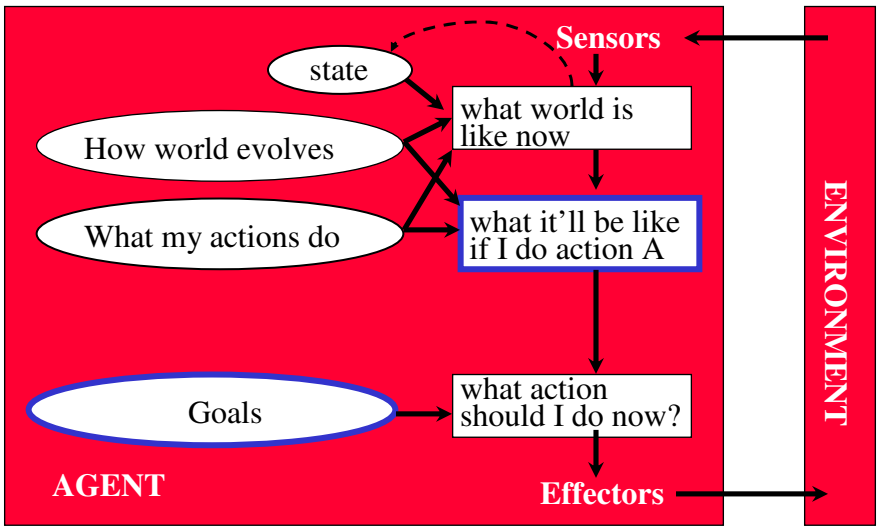

Utility-Based Agents

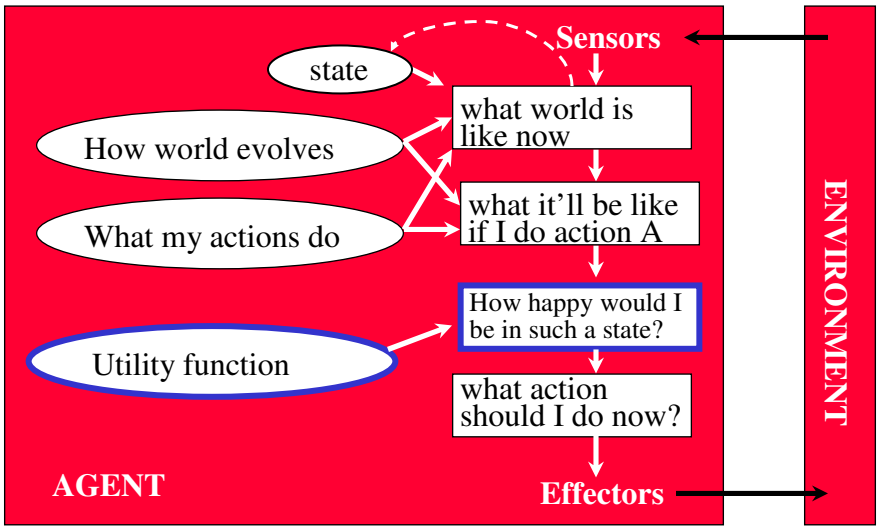

Learning Agents

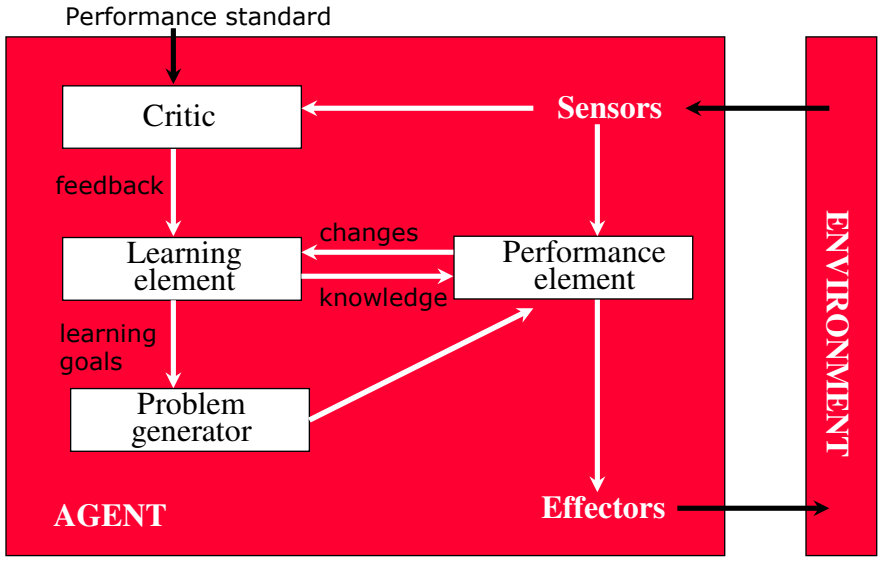

While driving, what's the best policy?

- Always stop at a stop sign
- Never stop at a stop sign
- Look around for other cars and stop only if you see one approaching
- Look around for a cop and stop only if you see one

- **What kind of agent are you?**
  - reflex, goal-based, utility-based?