

Intelligent Agents

Burr H. Settles
CS-540, UW-Madison
www.cs.wisc.edu/~cs540-1
Summer 2003

1

Announcements

- Read Chapter 3 in *AI: A Modern Approach* for next time
- Feel free to interrupt with questions!
- Office hours are posted on the website, and lecture slides will be added each day after class (PDF format)

2

What is an Agent?

Well, maybe!



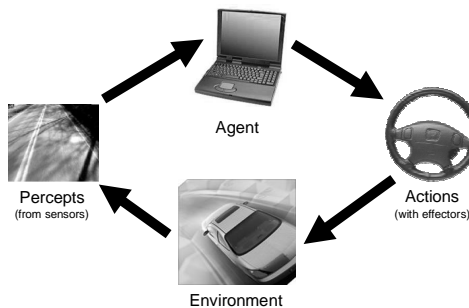
3

What is an Agent?

- An intelligent agent perceives its environment with *sensors* and acts rationally upon that environment with its *effectors*
- Agents perform the reasoning (or *thinking*) task in our AI framework from last time... it consists of:
 - Percepts (input from sensors)
 - Actions (output with effectors)
 - Goals
 - Environment

4

Agent Architecture



5

Examples of Intelligent Agents

Agent	Percepts	Actions	Goals	Environment
ALVINN	Images, Position	Steering, Control speed	Drive from A to B	Roads, Vehicles, Hazards
Google engine	Web pages, Search queries	Navigate web, Gather info.	Find best search matches	Internet
Deep Blue	Current board state	Choose next move	Win the game!	Opponent, Game board
Financial Forecaster	Stock market data	Pick stocks to buy & sell	Maximize investments	Stock market, company reports
Medical Diagnostic	Symptoms, Test results	Tests, Treatments	Healthy patient	Patient, Hospital

6

Agent Program



* *Agents map percepts to actions*

■ Basic cycle

- Perceive
- Reason, decide, or *think*
- Act
- Repeat

7

Another Definition



■ An agent is a computer software system whose main characteristics are:

- Situatedness
- Autonomy
- Adaptivity
- Sociability

8

Features of Agents



■ Situatedness

- Emphasizes the agent's direct connection to its environment through some set of sensors (percepts) and effectors (actions)

9

Features of Agents



■ Autonomy

- The agent acts without direct intervention by humans or other agents; it has control over its own actions and internal state
- Some aspect of the *current situation* must trigger a response (preprogramming to know when and how to respond doesn't count)

10

Features of Agents



■ Adaptivity

- React flexibly to changes in its environment
- Take goal-directed initiatives
- Learn from its own experience, its environment, and its interactions with others (be they people or other agents)

11

Features of Agents



■ Sociability

- The agent is capable of interacting in a peer-to-peer manner with humans or other agents

12

How should an Agent Act?



- * *A rational agent does the right thing!*
- A *performance measure* is needed say how successful a task has been achieved
- Rational agents try to *maximize success*

13

How should an Agent Act?



- An ideal rational agent should, for each possible sequence of perceptions, do whatever action(s) will maximize its performance measure based on:
 - The percepts
 - Its built knowledge
 - Acquired knowledge

14

Types of Agent Programs



- There are several kinds of agent programs which try to meet these conditions:
 - Lookup table
 - Simple reflex agent
 - Model-based reflex agent
 - Goal-based agent
 - Utility-based agent

15

Lookup Table



- Map percepts to actions *explicitly*
 - For interesting problems, the table would be too big to generate and store
 - Not adaptive to changes in the environment; *entire* table must be updated
 - Matches only the current percept... doesn't track its history

16

Simple Reflex Agent



- Use *if-then rules* to match certain percepts to an action
 - No need to consider all percepts
 - Can generalize percepts by mapping them to the same action
 - Can adapt to changes in the environment by adding new rules
 - Reacts only the current percept... *still* doesn't track its history

17

Model-Based Reflex Agent



- Internally encode state (maintain a model) of the world based on past percepts and prior knowledge
 - Actions can be based on a sequence of percepts; knowledge of non-perceptual world state
 - Must be able to represent change in the world; could represent just current state but this limits reasoning ability

18

Goal-Based Agent



- Actions chosen to reach a *desired goal*
 - Goals help decide which actions are good
 - Proactive rather than reactive
 - May have to consider long sequences of possible actions before deciding if goal is achieved: “what will happen if I do this...?”
 - Doesn’t allow for *degrees* of goodness

19

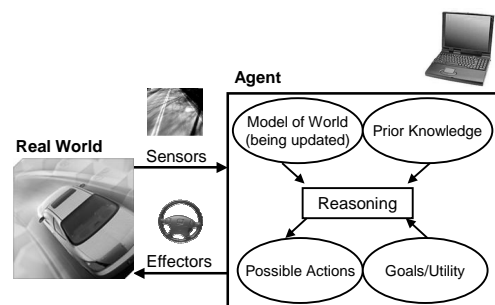
Utility-Based Agent



- Choose which alternative is *best*
 - Utility function gives a measure of success or goodness for a given state
 - Allows decisions comparing choice between conflicting goals, and choice between likelihood of success and importance of goal (if achievement is uncertain)

20

Agent Architecture Example



21

Features of Environments



- Is the world fully or partially observable?
- An environment is fully observable if the agent is provided with full knowledge about the state of the world

22

Features of Environments



- Is the world deterministic or stochastic?
- An environment is deterministic if the next state of the world is completely determined by the current state and the agent’s actions
- Randomness and chance cause stochastic environments

23

Features of Environments



- Is the world episodic or sequential?
- An environment is episodic if each percept-action episode does not depend on the actions in prior episodes
 - e.g. Picking defective parts off an assembly line
- Driving a car or playing a game are sequential: short-term actions have long-term consequences

24

Features of Environments



- Is the world static or dynamic?
- An environment is static if it doesn't change between the time of perceiving and acting
- Time is an important factor in dynamic environments, since percepts are time-sensitive and can become obsolete

25

Features of Environments



- Is the world discrete or continuous?
- An environment is discrete if there are a limited number of distinct, clearly-defined percepts and actions

26

Features of Environments



- Is the world single agent or multi-agent?
- Solving a crossword puzzle is clearly a single agent environment
- Playing chess, driving a car, working with other robots to assemble a product are all multi-agent
 - Some competitive and some cooperative multi-agent environments

27

Features of Environments



- Let's think about the features for these environments:
 1. Chess
 2. Poker
 3. Driving a taxi
 4. Classifying web pages
 5. Reading addresses on mail envelopes

Fully observable	Partially observable
Episodic	Sequential
Deterministic	Stochastic
Static	Dynamic
Discrete	Continuous
Single-agent	Multi-agent

28

Environments



- In practice there is often uncertainty due to partially observable, stochastic, and dynamic environments
- In AI, we commonly use *toy environments*, which approximate the world but with simpler states and actions

29

Summary

* *An agent perceives and acts in an environment to achieve specific goals.*

- Types of agents
 - Lookup table
 - Simple reflex agent
 - Model-based reflex agent
 - Goal-based
 - Utility-based

30

Summary

■ Features of Environments

- Fully vs. partially observable
- Deterministic vs. stochastic
- Episodic vs. sequential
- Static vs. dynamic
- Discrete vs. continuous
- Single vs. multi-agent

31