Lecture 2

Agents & Environments (Chap. 2)

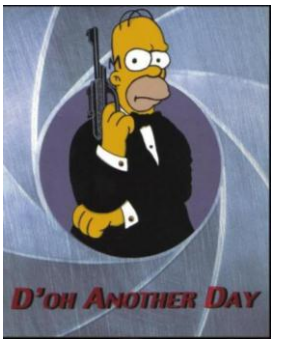

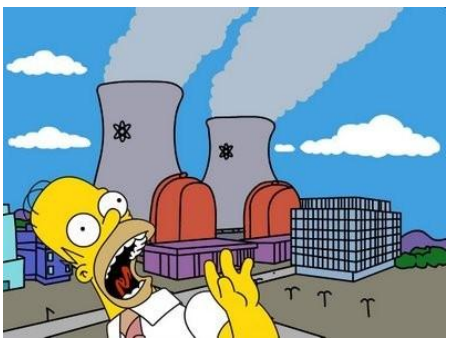

Based on slides by UW CSE AI faculty, Dan Klein, Stuart Russell, Andrew Moore

Outline

- Agents and environments
- Rationality
- PEAS specification
- Environment types
- Agent types
- Pac-Man projects

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

Other Types of Agents

- Immobots (Immobile Robots)
  Intelligent buildings
  Intelligent forests
- Softbots
  Askjeeves.com (now Ask.com)
  Expert Systems
  Microsoft Clippy

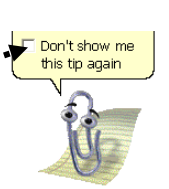

4

Intelligent Agents

- Have sensors and actuators (effectors)
- Implement mapping from percept sequence to actions
- Maximize a Performance Measure

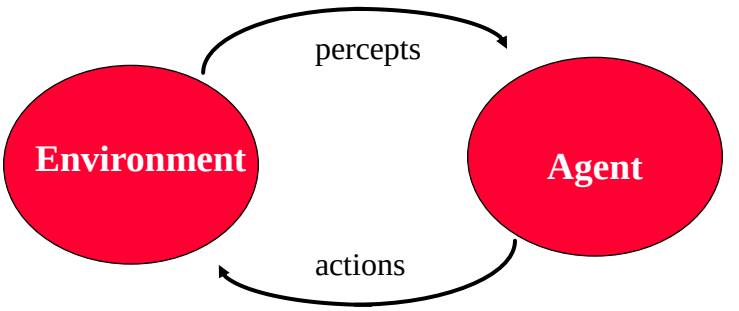

5

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 maximizes expected performance* on the basis of evidence perceived so far and built-in prior knowledge."

7

Autonomy

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

Why is this important?

8

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

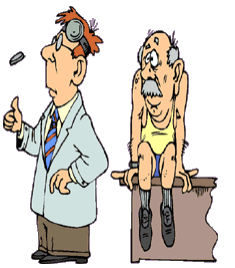

- PEAS for Medical diagnosis system

- 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

12

Fully observable vs. Partially observable

Can the agent observe the complete state of the environment?

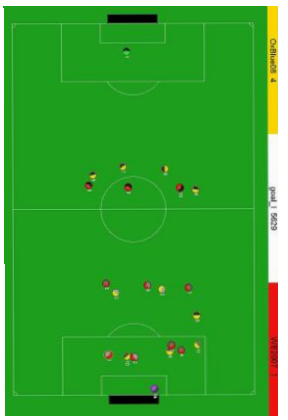

vs.

13

Single agent vs. Multiagent

Is the agent the only thing acting in the world?

vs.

14

Deterministic vs. Stochastic

Is there uncertainty in how the world works?

vs.

15

Episodic vs. Sequential

Does the agent take more than one action?

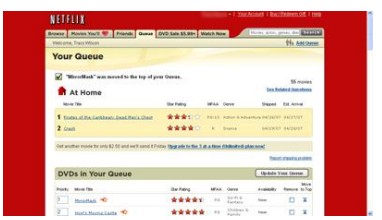

vs.

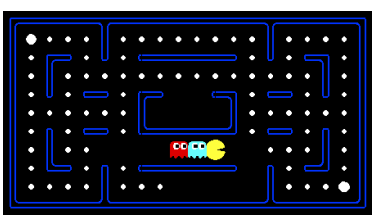

16

Discrete vs. Continuous

Are the states, actions etc. discrete or continuous?

vs.

17

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

18

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

19

Simple Reflex Agents

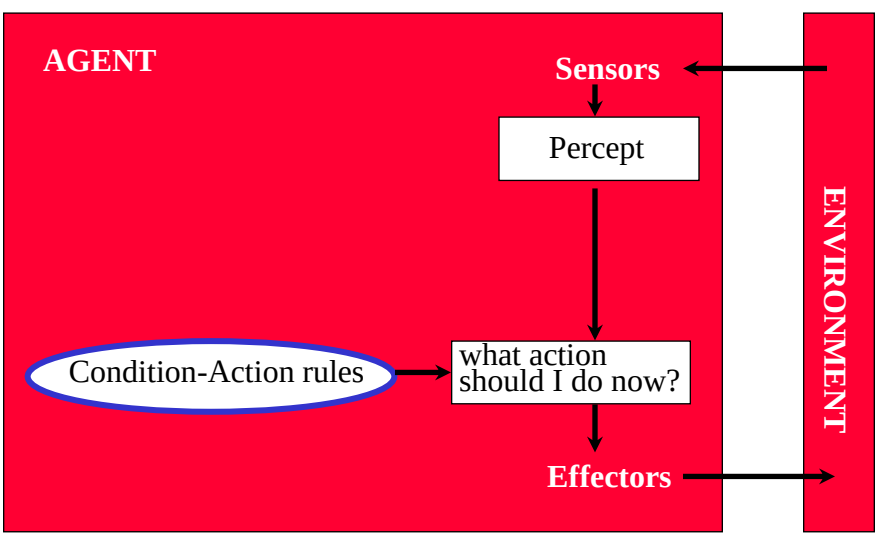

20

Simple Reflex Agents

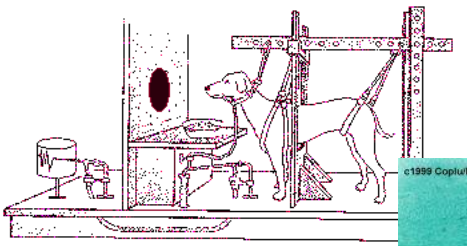

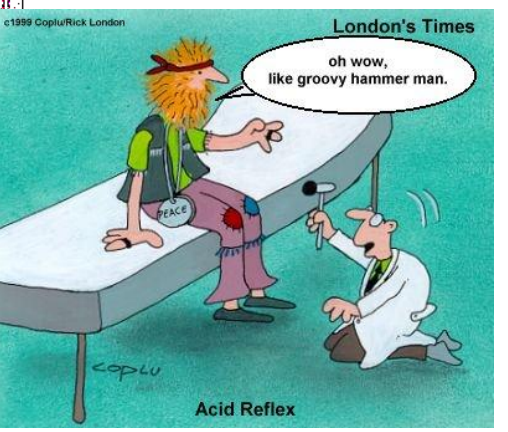

21

Famous Reflex Agents

22

Reflex Agent with Internal State

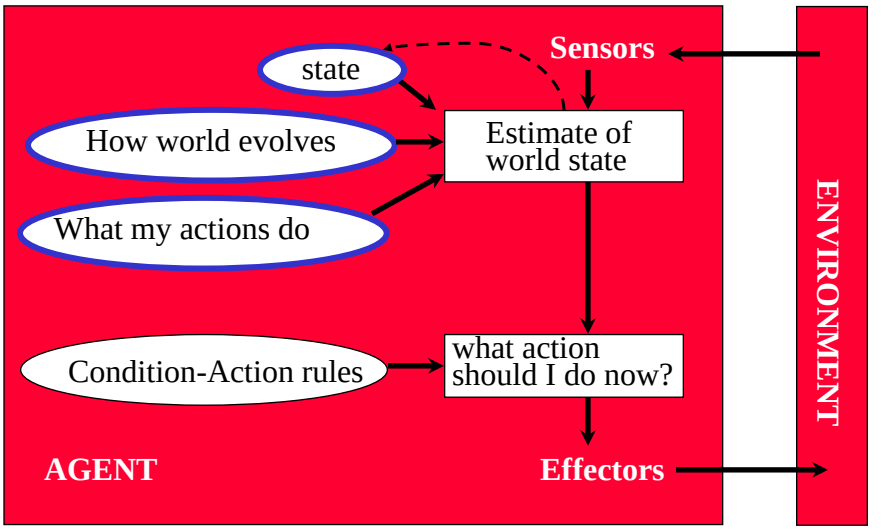

23

Goal-Based Agents

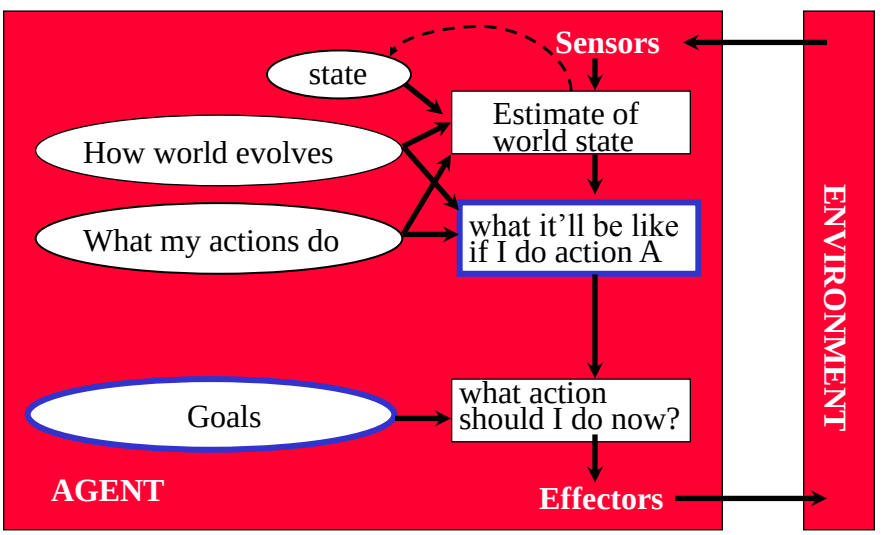

24

Utility-Based Agents

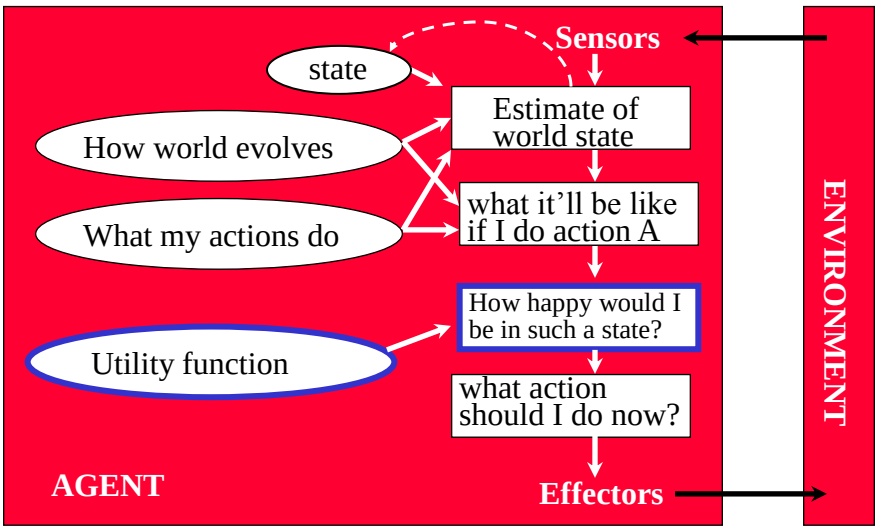

25

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?

26

Pac-Man as an Agent

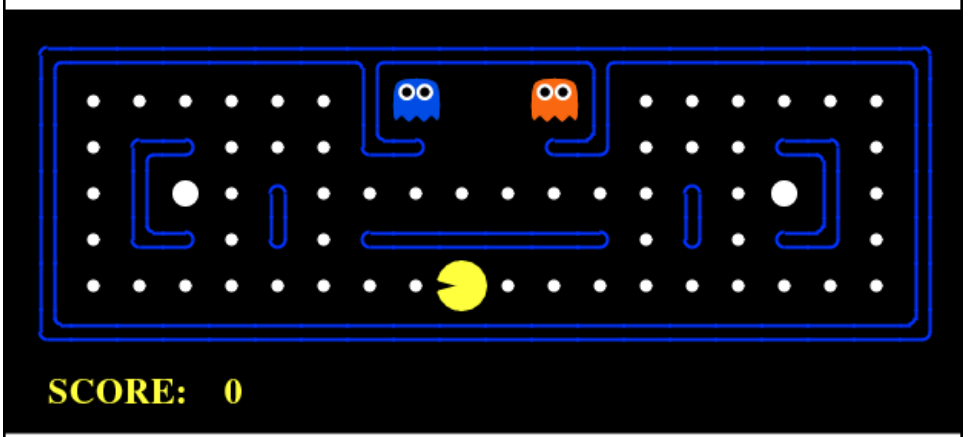

27

The CSE 473 Pac-Man Projects

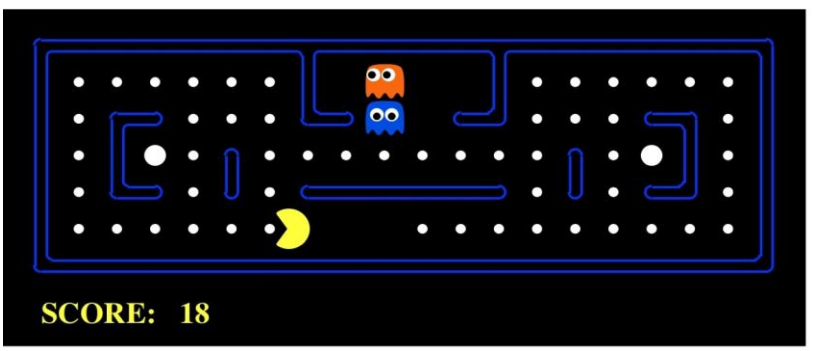

Originally developed at UC Berkeley:
http://www-inst.eecs.berkeley.edu/~cs188/pacman/pacman.html

28

Project 1: Search

Goal:

- Help Pac-man find its way through the maze

Techniques:

- Search: breadth-first, depth-first, etc.
- Heuristic Search: Best-first, A*, etc.

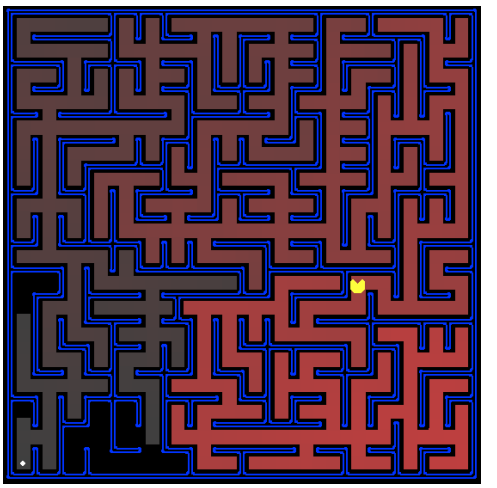

29

Project 2: Game Playing

Goal:
Build a rational Pac-Man agent!

Techniques:
Adversarial Search: minimax, alpha-beta, expectimax, etc.

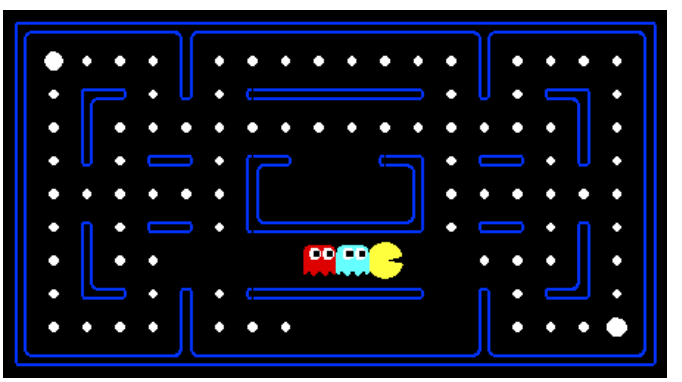

30

Project 3: Planning and Learning

Goal:
Help Pac-Man learn about its world

Techniques:
- Planning: MDPs, Value Iteration
- Learning: Reinforcement Learning

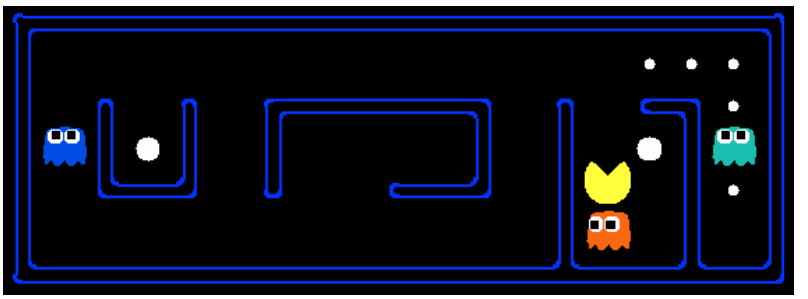

31

Project 4: Ghostbusters

Goal:
Help Pac-man hunt down the ghosts

Techniques:
- Probabilistic models: HMMs, Bayes Nets
- Inference: State estimation and particle filtering

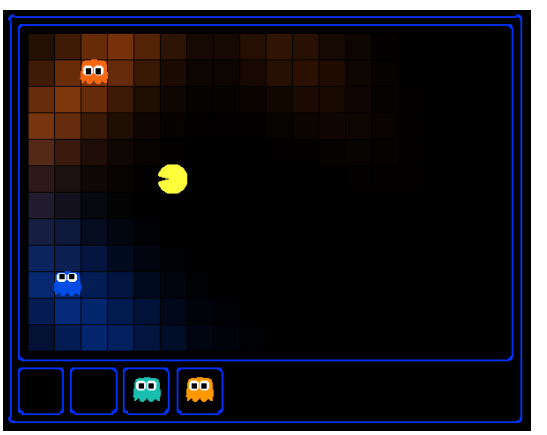

32

To Do

- Project 0: Python tutorial
- Finish chapters 1 and 2; start chapter 3