


CSE 573: Artificial Intelligence

Agents and environments




slides adapted from
Stuart Russel, Dan Klein, Pieter Abbeel from ai.berkeley.edu
And Hanna Hajishirzi, Jared Moore, Dan Weld

Outline


- Agents and environments
- Rationality
- PEAS (Performance measure, Environment, Actuators, Sensors)
- Environment types
- Agent types

Agents and environments




- An agent *perceives* its environment through *sensors* and *acts* upon it through *actuators* (or *effectors*, depending on whom you ask)

Agents and environments




- Are humans agents?
- Yes!
 - Sensors = vision, audio, touch, smell, taste, proprioception
 - Actuators = muscles, secretions, changing brain state

Agents and environments



- Are pocket calculators agents?
- Yes!
 - Sensors = key state sensors
 - Actuators = digit display


Agents and environments




- AI is more interested in agents with large computational resources and environments that require nontrivial decision making

Agent functions

- The **agent function** maps from percept histories to actions:
 - $f: \mathcal{P}^* \rightarrow \mathcal{A}$
 - I.e., the agent's actual response to any sequence of percepts




Example: Vacuum world



- Percepts: *[location, status]*, e.g., *[A, Dirty]*
- Actions: *Left, Right, Suck, NoOp*

Vacuum cleaner agent



Agent function

Percept sequence	Action
[A,Clean]	Right
[A,Dirty]	Suck
[B,Clean]	Left
[B,Dirty]	Suck
[A,Clean],[B,Clean]	Left
[A,Clean],[B,Dirty]	Suck
etc	etc

Agent program

```
function Reflex-Vacuum-Agent([location,status])
    returns an action
    if status = Dirty then return Suck
    else if location = A then return Right
    else if location = B then return Left
```

What is the *right* agent function?

Can it be implemented by a small agent program?

Rationality

- A fixed **performance measure** evaluates the environment sequence
 - one point per square cleaned up?
 - Basically, but details matter: agent can dump dirt then clean, repeatedly
 - Add large penalty for dumping dirt? Add small penalty for moving?
- A **rational agent** chooses whichever action maximizes the **expected** value of the performance measure
 - given the percept sequence to date and prior knowledge of environment


Does Reflex-Vacuum-Agent implement a rational agent function?

Yes, if movement is free, or new dirt arrives frequently

Rationality, contd.


- Are rational agents *omniscient*?
 - No – they are limited by the available percepts
- Are rational agents *clairvoyant*?
 - No – they may lack knowledge of the environment dynamics
- Do rational agents *explore* and *learn*?
 - Yes – in unknown environments these are essential
- Do rational agents *make mistakes*?
 - No – but their actions may be unsuccessful / suboptimal
- Are rational agents *autonomous* (i.e., transcend initial program)?
 - Yes – as they learn, their behavior depends more on their own experience

A human agent in Pacman



The task environment - PEAS

- Performance measure
 - -1 per step; + 10 food; +500 win; -500 die; ghost
- Environment
 - Pacman dynamics (incl ghost behavior)
- Actuators
 - Left Right Up Down
- Sensors
 - Entire state is visible / observable (except power pellet duration)



Pacman agent contd.

- Can we (in principle) extend this reflex agent to behave well in all standard Pacman environments?
 - No – Pacman is not quite fully observable (power pellet duration)
 - Otherwise, yes – we can (*in principle*) make a lookup table.....

PEAS: Automated taxi

- Performance measure
 - Income, happy customer, vehicle costs, fines, insurance premiums
- Environment
 - US streets, other drivers, customers, weather, police...
- Actuators
 - Steering, brake, gas, display/speaker
- Sensors
 - Camera, radar, accelerometer, engine sensors, microphone, GPS



Image: <http://nypost.com/2014/06/21/how-google-might-put-taxi-drivers-out-of-business/>


PEAS: Backgammon

- Performance measure
 - Move all checkers home first
- Environment
 - Game board, other player?
- Actuators
 - Roll dice, decide how to move pieces
- Sensors
 - See the full board



PEAS: Medical diagnosis system

- Performance measure
 - Patient health, cost, reputation
- Environment
 - Patients, medical staff, insurers, courts
- Actuators
 - Screen display, email
- Sensors
 - Keyboard/mouse, test results



Environment types

	Pacman	Backgammon	Diagnosis	Taxi
Fully or partially observable	F*	F	P	P
Single-agent or multiagent	M	M	S	M
Deterministic or stochastic	D	S	D*	S
Static or dynamic	D	D	S	D
Discrete or continuous	D	D	C	C
Known physics?	Y	Y	N	Y
Known perf. measure?	Y	Y	N	Y*

Agent design

- The environment type largely determines the agent design
 - *Partially observable* => agent requires *memory* (internal state)
 - *Stochastic* => agent may have to prepare for *contingencies*
 - *Multi-agent* => agent may need to behave *randomly*
 - *Static* => agent has time to compute a rational decision
 - *Continuous time* => continuously operating *controller*
 - *Unknown physics* => need for *exploration*
 - *Unknown perf. measure* => observe/interact with *human principal*

Summary

- An *agent* interacts with an *environment* through *sensors* and *actuators*
- The *agent function* describes what the agent does in all circumstances
- Rational agents choose actions that maximize their expected utility
- PEAS descriptions define task environments; precise PEAS specifications are essential and strongly influence agent designs
- More difficult environments require more complex agent designs and more sophisticated representations