Chapter 2: Agents

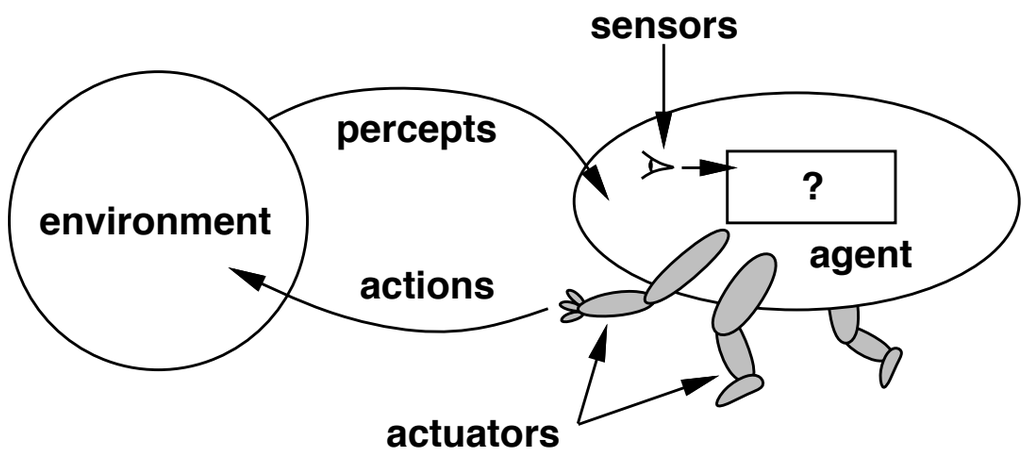

Agents: Something that does stuff. Includes humans, robots, thermostats, Google Maps, etc.

Vacuum cleaner world

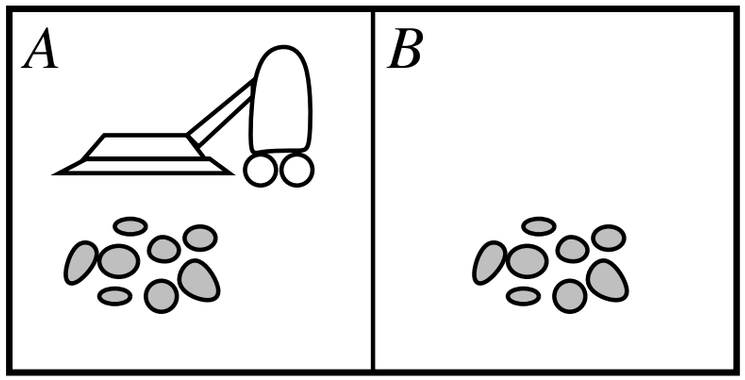

Actions: Left, Right, Suck, NoOp

Task environment: PEAS model

- Performance measure
- Environment
- Actuators
- Sensors

- To design an agent, we need to specify the task environment.
- Four parts components: PEAS.
- What is the task environment for the vacuum cleaner?
    - Environment: squares, dirt.
    - Sensors: Sees the whole state.
    - Actuators: R, L, Suck, NoOp
    - Performance: Points for cleaning? Negative points for dirty squares? Points as long as there are fewer than k dirty squares?

What is a PEAS model for a self-driving car?

- Performance measure
- Environment
- Actuators
- Sensors

- Environment: streets, traffic, pedestrians, weather, ...
- Sensors: video, accelerometers, LIDAR, engine sensors, driver input, GPS, ...
- Actuators: steering, accelerator, brake, horn, display, ...
- Performance: safety, legality, speed, comfort, ...

What is a PEAS model for an internet shopping agent?

- Performance measure
- Environment
- Actuators
- Sensors

- Environment: web sites
- Sensors: HTML pages, user input
- Actuators: follow URL, fill in form, display to user
- Performance: price, quality, relevance to query, shipping time

Agent function

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

f(percept sequence) -> action

Agent function vs agent program: Function is mathematical object. Program is an implementation (i.e. in Python).

AI is about creating agent functions.

What is the right/best function? Can we implement it?

Rationality / optimality

- "Rational" (/ "optimal") agent: One that chooses whatever action maximizes the performance.
- In random or unobserved environments: the expected value of performance.
- We're aiming to produce rational agents. Or as close to rational as we can.
- Rational != omniscient. Therefore, rational != successful.
- Rational can include learning, autonomy, etc.

Environment types

- Observable (vs partially-observable): We can see everything we need to know about the world.
- Deterministic (vs stochastic): Actions always have the same result.
- Single-agent (vs multi-agent): We are the only actor.
- Discrete (vs continuous): The environment has distinct states, rather than smoothly varying states.
- Episodic (vs sequential): The agent's experience is divided into independent episodes.

- What type of environment we are in will determine what algorithms we can apply.
- The type of environment is not well-defined. E.g:
  - We can think of randomness as being unobserved.
  - We can discretize continuous states.
- Episodic: Applies when we get to supervised machine learning.

Environment types

	Solitaire	Backgammon	Internet shopper	Self-driving taxi	Spam filter
Observable					
Deterministic					
Single-agent					
Discrete					
Episodic					