Rational Agents (Chapter 2)

Agents

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators

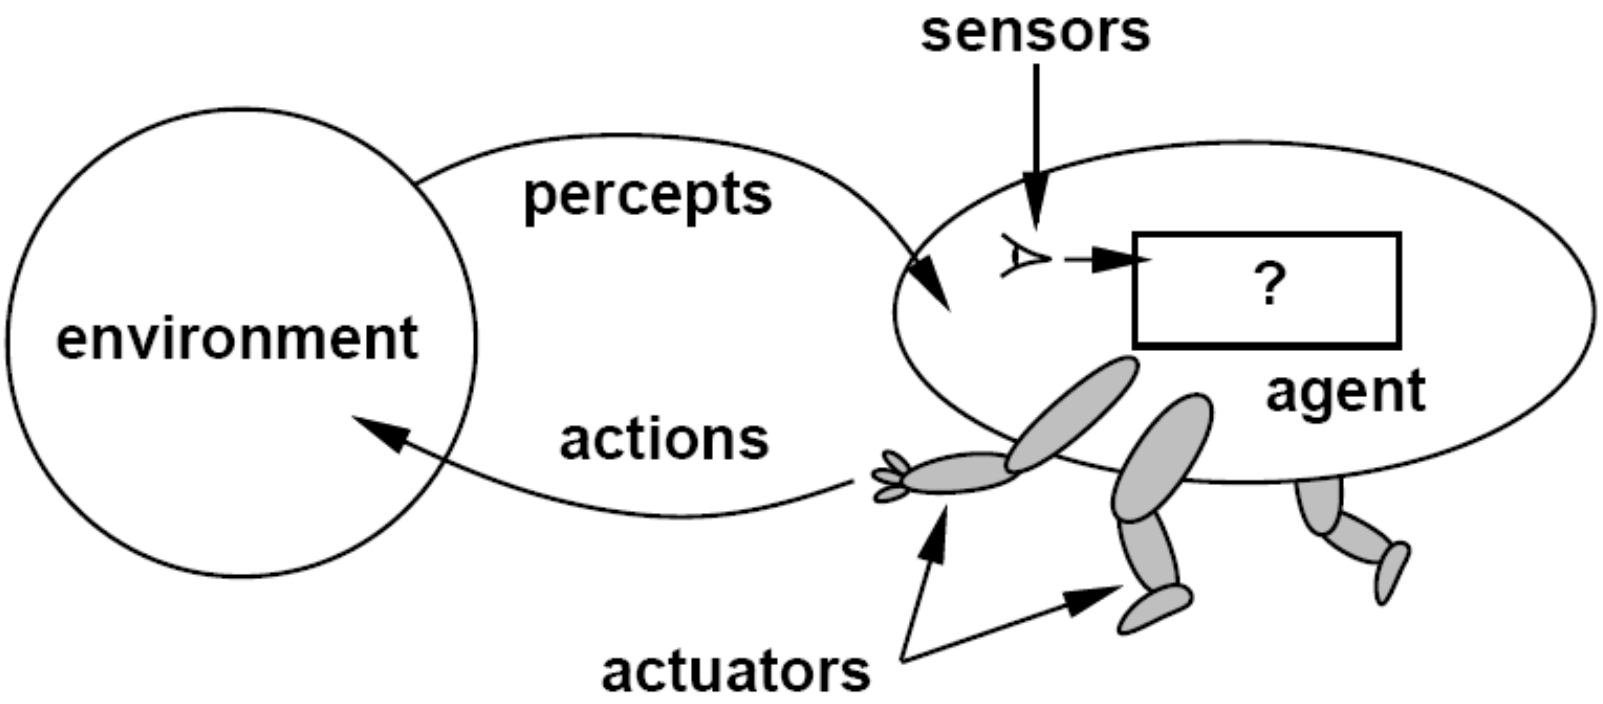

Example: Vacuum-Agent

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

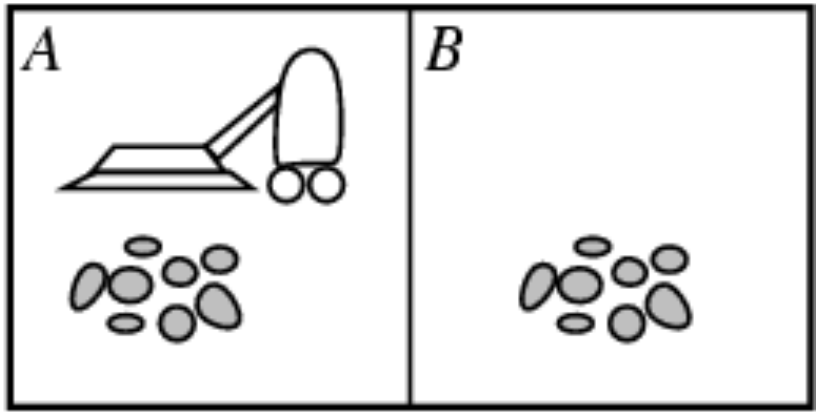

function Vacuum-Agent([location,status]) returns an action

- *if* status = Dirty *then* return Suck
- *else if* location = A *then* return Right
- *else if* location = B *then* return Left

Rational agents

- For each possible percept sequence, a **rational agent** should select an action that is expected to maximize its **performance measure**, *given the evidence provided by the percept sequence and the agent's built-in knowledge*
- **Performance measure (utility function):** An *objective* criterion for success of an agent's behavior
- **Expected utility:**

$$EU(action) = \sum_{outcomes} P(outcome \mid action) U(outcome)$$

- Can a rational agent make mistakes?

Back to Vacuum-Agent

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

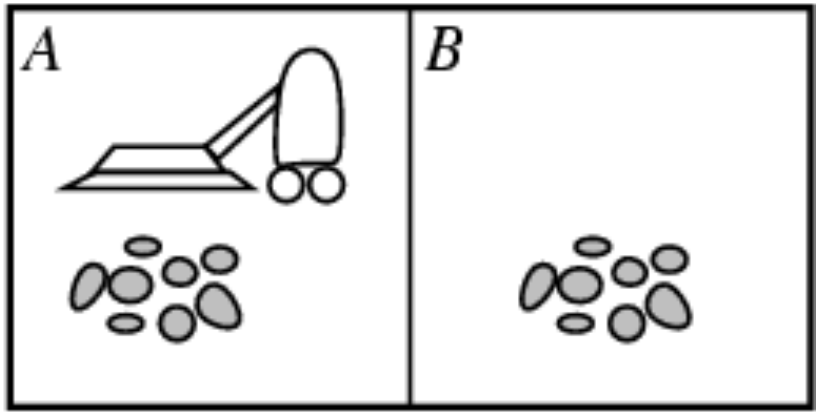

function Vacuum-Agent([location,status]) returns an action

- *if* status = Dirty *then* return Suck
- *else if* location = A *then* return Right
- *else if* location = B *then* return Left

- Is this agent rational?
  - Depends on performance measure, environment properties

Specifying the task environment

- **PEAS: Performance measure, Environment, Actuators, Sensors**
- **P:** a function the agent is maximizing (or minimizing)
  - Assumed given
  - In practice, needs to be computed somewhere
- **E:** a formal representation for *world states*
  - For concreteness, a tuple (var_1=val_1, var_2=val_2, … ,var_n=val_n)
- **A:** actions that change the state according to a *transition model*
  - Given a state and action, what is the successor state (or distribution over successor states)?
- **S:** observations that allow the agent to infer the world state
  - Often come in very different form than the state itself
  - E.g., in tracking, observations may be pixels and state variables 3D coordinates

PEAS Example: Autonomous taxi

- Performance measure
  - Safe, fast, legal, comfortable trip, maximize profits
- Environment
  - Roads, other traffic, pedestrians, customers
- Actuators
  - Steering wheel, accelerator, brake, signal, horn
- Sensors
  - Cameras, LIDAR, speedometer, GPS, odometer, engine sensors, keyboard

Another PEAS example: Spam filter

- Performance measure
  - Minimizing false positives, false negatives
- Environment
  - A user's email account, email server
- Actuators
  - Mark as spam, delete, etc.
- Sensors
  - Incoming messages, other information about user's account

Environment types

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

Fully observable vs. partially observable

- Do the agent's sensors give it access to the complete state of the environment?
  - For any given world state, are the values of all the variables known to the agent?

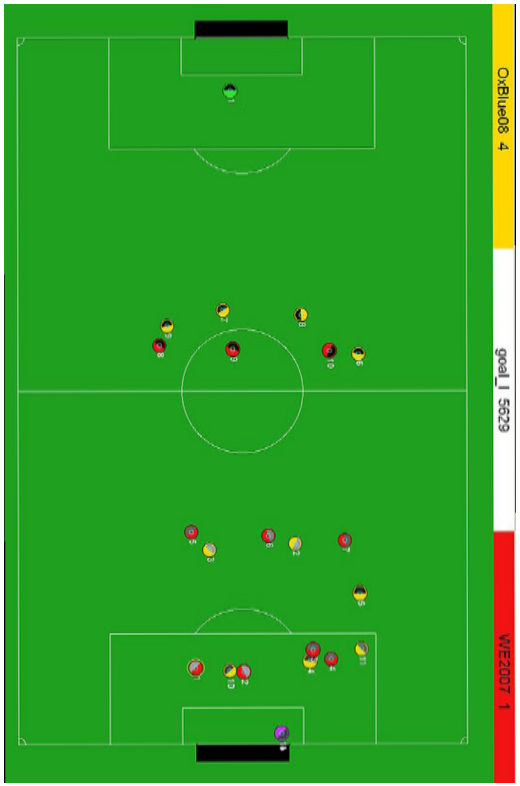

vs.

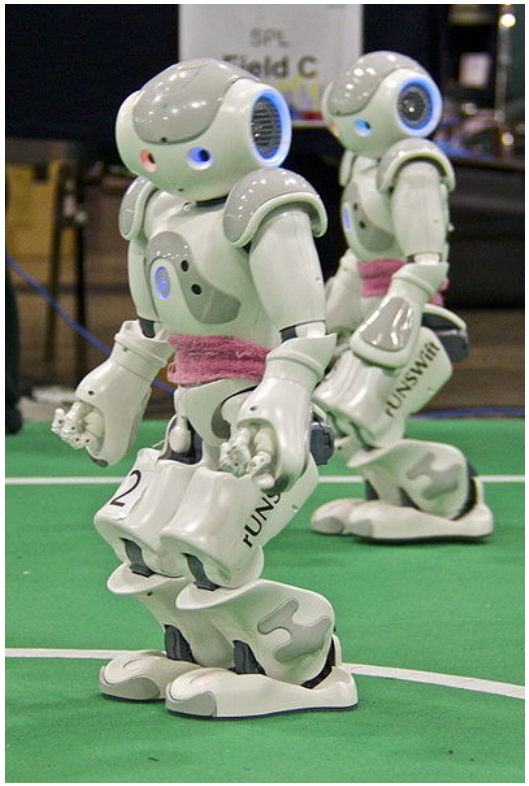

Source: L. Zettlemoyer

Deterministic vs. stochastic

- Is the next state of the environment completely determined by the current state and the agent's action?
  - Is the transition model deterministic (unique successor state given current state and action) or stochastic (distribution over successor states given current state and action)?
  - **Strategic:** the environment is deterministic except for the actions of other agents

vs.

Episodic vs. sequential

- Is the agent's experience divided into unconnected single decisions/actions, or is it a coherent sequence of observations and actions in which the world evolves according to the transition model?

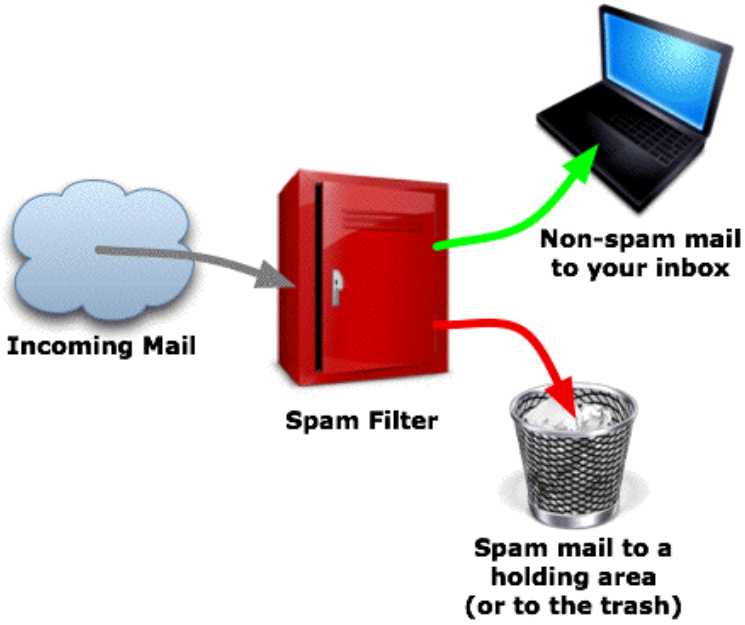

vs.

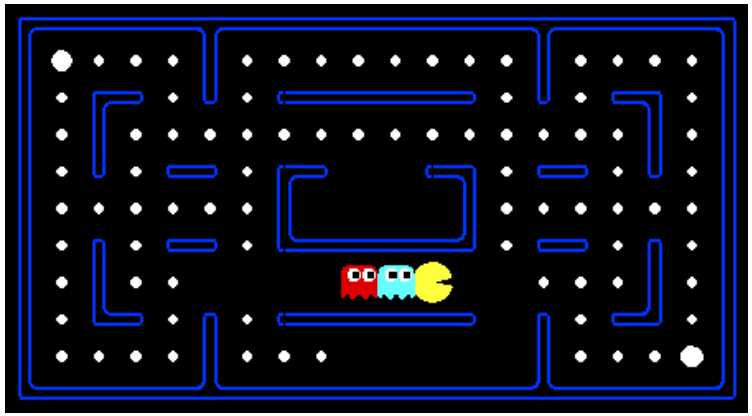

Static vs. dynamic

- Is the world changing while the agent is thinking?
  - **Semidynamic:** the environment does not change with the passage of time, but the agent's performance score does

vs.

Discrete vs. continuous

- Does the environment provide a fixed number of distinct percepts, actions, and environment states?
  - Are the values of the state variables discrete or continuous?
  - Time can also evolve in a discrete or continuous fashion

vs.

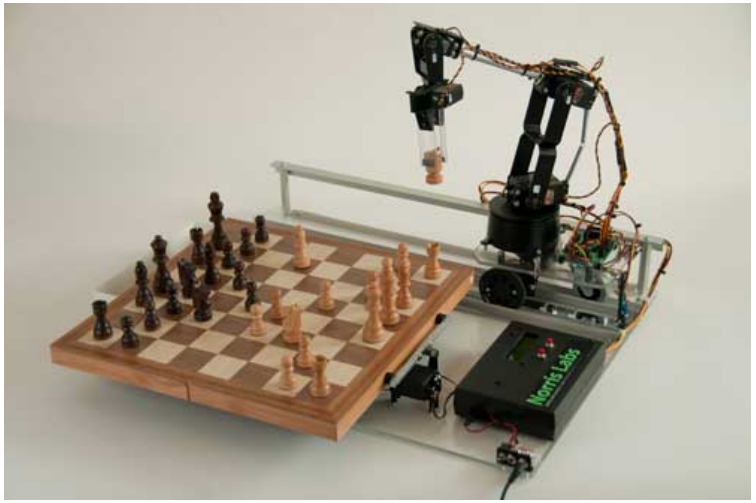

Single-agent vs. multiagent

- Is an agent operating by itself in the environment?

vs.

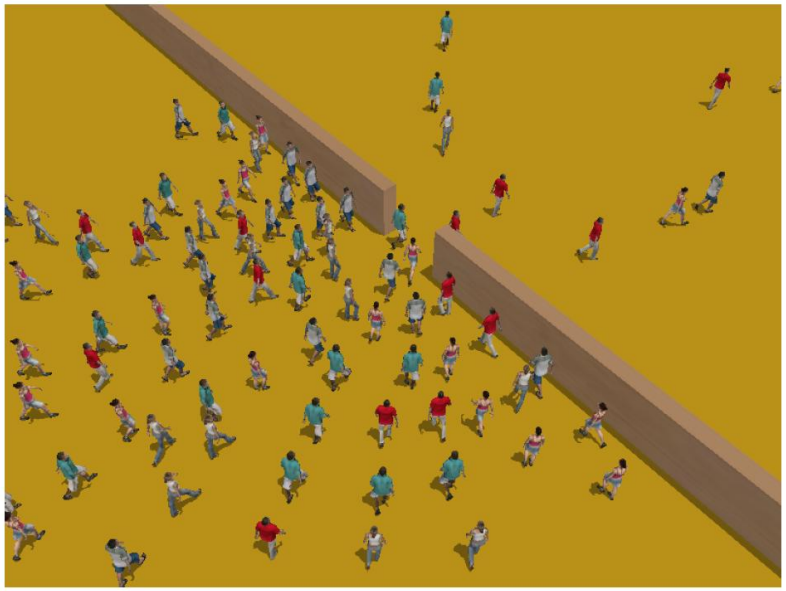

Known vs. unknown

- Are the rules of the environment (transition model and rewards associated with states) known to the agent?
  - Strictly speaking, not a property of the environment, but of the agent's state of knowledge

vs.

Examples of different environments

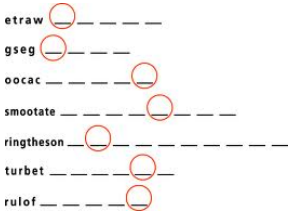

Word jumble solver

Chess with a clock

Scrabble

Autonomous driving

	Word jumble solver	Chess with a clock	Scrabble	Autonomous driving
Observable	Fully	Fully	Partially	Partially
Deterministic	Deterministic	Strategic	Stochastic	Stochastic
Episodic	Episodic	Sequential	Sequential	Sequential
Static	Static	Semidynamic	Static	Dynamic
Discrete	Discrete	Discrete	Discrete	Continuous
Single agent	Single	Multi	Multi	Multi

Preview of the course

- Deterministic environments: search, constraint satisfaction, classical planning
  - Can be sequential or episodic
- Multi-agent, strategic environments: minimax search, games
  - Can also be stochastic, partially observable
- Stochastic environments
  - Episodic: Bayesian networks, pattern classifiers
  - Sequential, known: Markov decision processes
  - Sequential, unknown: reinforcement learning

Review: PEAS

Review: PEAS

- **P: Performance measure**
  - Function the agent is maximizing (or minimizing)
- **E: Environment**
  - A formal representation for *world states*
  - For concreteness, a tuple (var_1=val_1, var_2=val_2, … ,var_n=val_n)
- **A: Actions**
  - *Transition model:* Given a state and action, what is the successor state (or distribution over successor states)?
- **S: Sensors**
  - Observations that allow the agent to infer the world state
  - Often come in very different form than the state itself

Review: Environment types

- Fully observable vs. partially observable
- Deterministic vs. stochastic (vs. strategic)
- Episodic vs. sequential
- Static vs. dynamic (vs. semidynamic)
- Discrete vs. continuous
- Single agent vs. multi-agent
- Known vs. unknown

Attribution

Slides developed by Svetlana Lazebnik based on content from Stuart Russell and Peter Norvig, Artificial Intelligence: A Modern Approach, 3rd edition