ECE 448 Lecture 3: Rational Agents

Slides by Svetlana Lazebnik, 9/2016

Modified by Mark Hasegawa-Johnson, 9/2017

Contents

- Agent = Performance, Environment, Actions, Sensors (PEAS)
- What makes an agent ***Rational***?
- What makes an agent ***Autonomous***?
- Types of Agents: Reflex, Internal-State, Goal-Directed, Utility-Directed (RIGU)
- Properties of Environments: Observable, Deterministic, Episodic, Static, Continuous (ODESC)

Agents

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

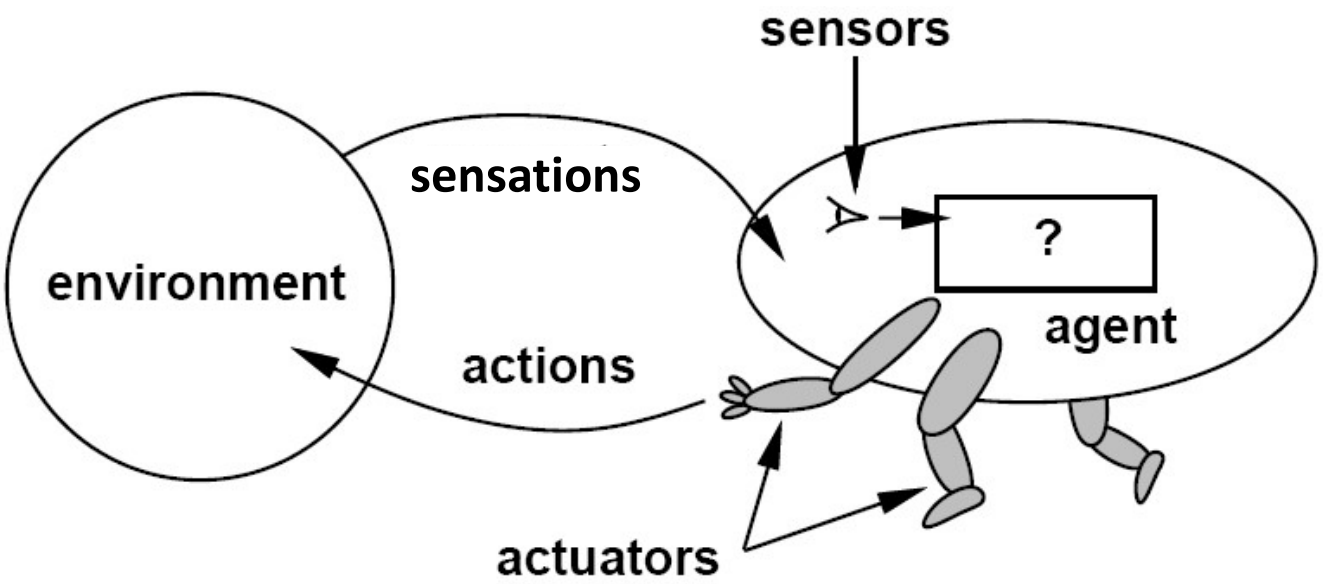

Example: Vacuum-Agent

- **Environment = tuple of variables:**
  - Location, status of both rooms,
    e.g., S = { Loc=A, Status=(Dirty, Dirty) }
- **Action = variable drawn from a set:**
  A ∈ { Left, Right, Suck, NoOp }
- **Sensors = tuple of variables:**
  - Location, and status of Current Room Only
    e.g., S = { Loc=A, Status = Dirty }

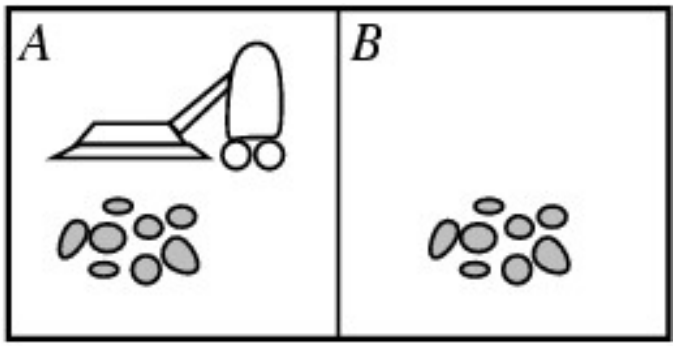

function Vacuum-Agent([location,status]) returns an action

- *if* Loc=A
  - *if* Status=Dirty *then* return Suck
  - *else if I have never visited* B *then return* Right
  - *else return* NoOp
- *else*
  - *if* Status=Dirty *then* return Suck
  - *else if I have never visited* A *then return* Left
  - *else return* NoOp

Specifying the task environment

- **PEAS: Performance, Environment, Actions, Sensors**
- **P:** a function the agent is maximizing (or minimizing)
  - Assumed given
- **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

Performance Measure

- An agent's performance is measured by some performance or **utility** measure
- Utility = function of the current environment E_t, and of the history of all actions from time 1 to time t, $A_{1:(t-1)}$:

$$U_t = f(E_t, A_{1:(t-1)})$$

- Example: U_t = # currently dirty rooms $-\frac{1}{2}$(# non-NoOp Actions)

What makes an agent *Rational*?

- 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:** the expected outcome of the action A_t:

$$EU(A_t) = \sum_{E_{t+1}} U_{t+1}(E_{t+1}, A_{1:t}) Pr\{E_{t+1} | A_{1:t}, S_{1:t}\}$$

- Can a rational agent make mistakes?

Back to the Vacuum-Agent

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

- *if* Loc=A
  - *if* Status=Dirty *then* return Suck
  - *else if I have never visited* B *then return* Right
  - *else return* NoOp
- *else*
  - *if* Status=Dirty *then* return Suck
  - *else if never visited* A *then* Left
  - *else return* NoOp

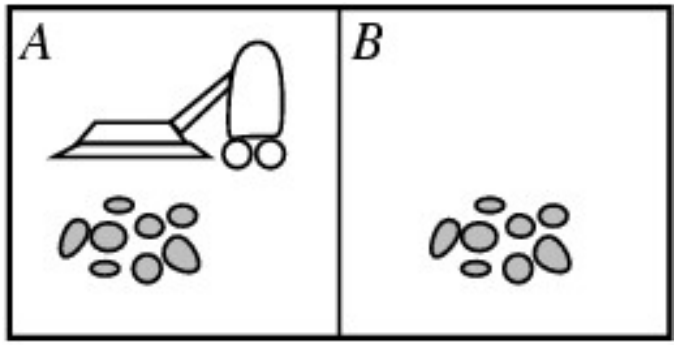

- Is this agent ***Rational***?

What makes an agent *Autonomous*?

- Russell & Norvig: “A system is autonomous to the extent that its behavior is determined by its own experience.”

- A Rational Agent might not be Autonomous, if its designer was capable of foreseeing the maximum-utility action for every environment.

- Example: Vacuum-Agent

Types of Agents

- Reflex agent: no concept of past, future, or value
  - Might still be Rational, if the environment is known to the designer with sufficient detail
- Internal-State agent: knows about the past
- Goal-Directed agent: knows about the past and future
- Utility-Directed agent: knows about past, future, and value

Reflex Agent

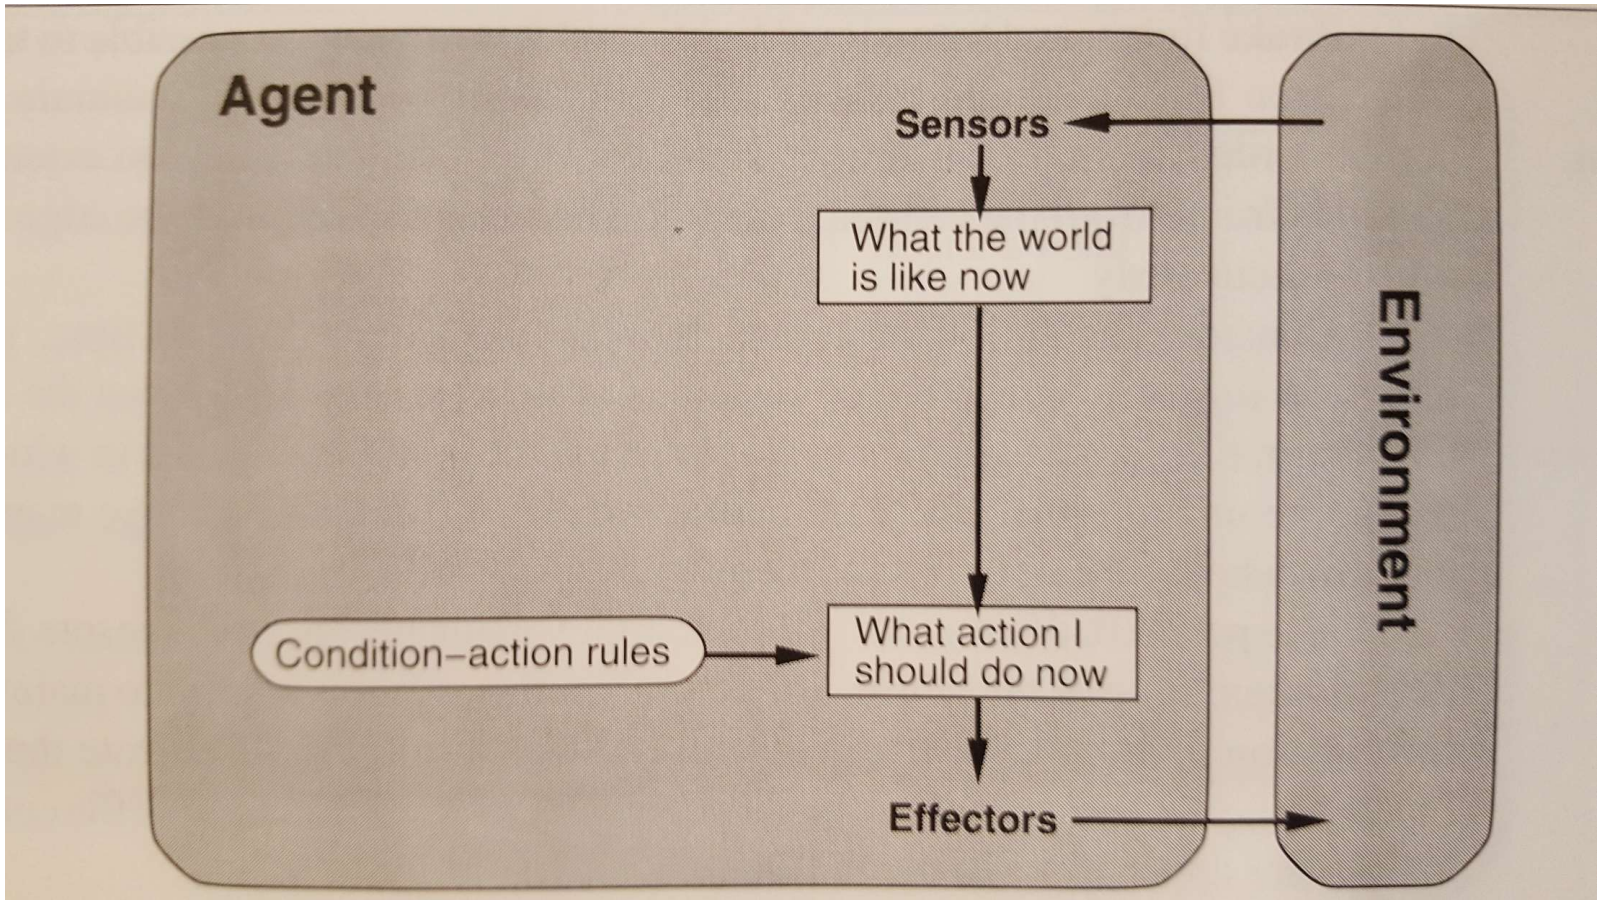

Internal-State Agent

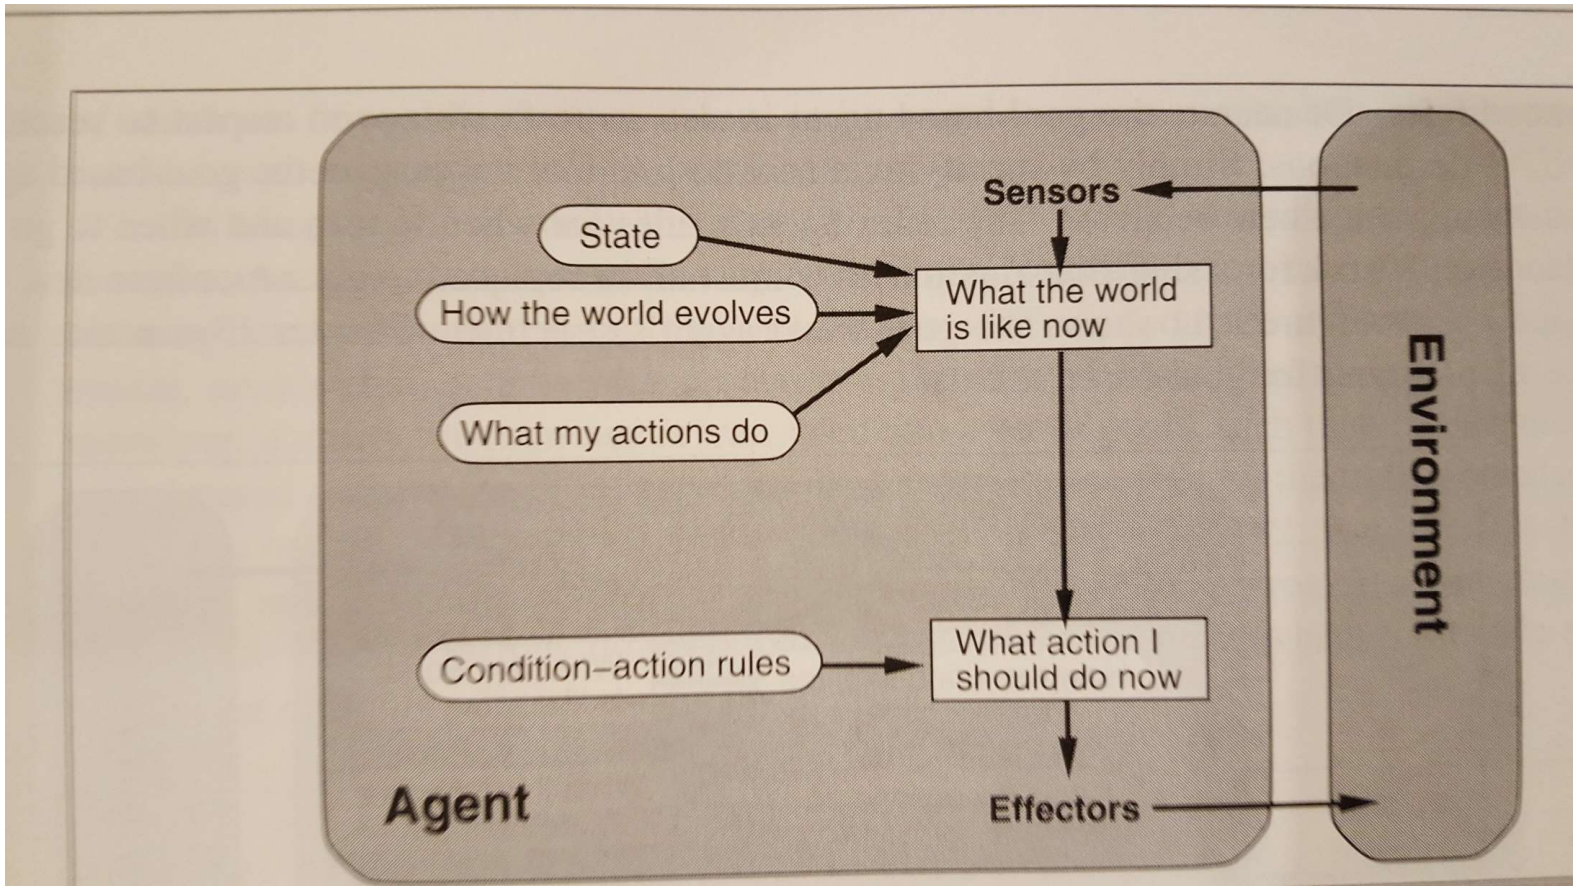

Goal-Directed Agent

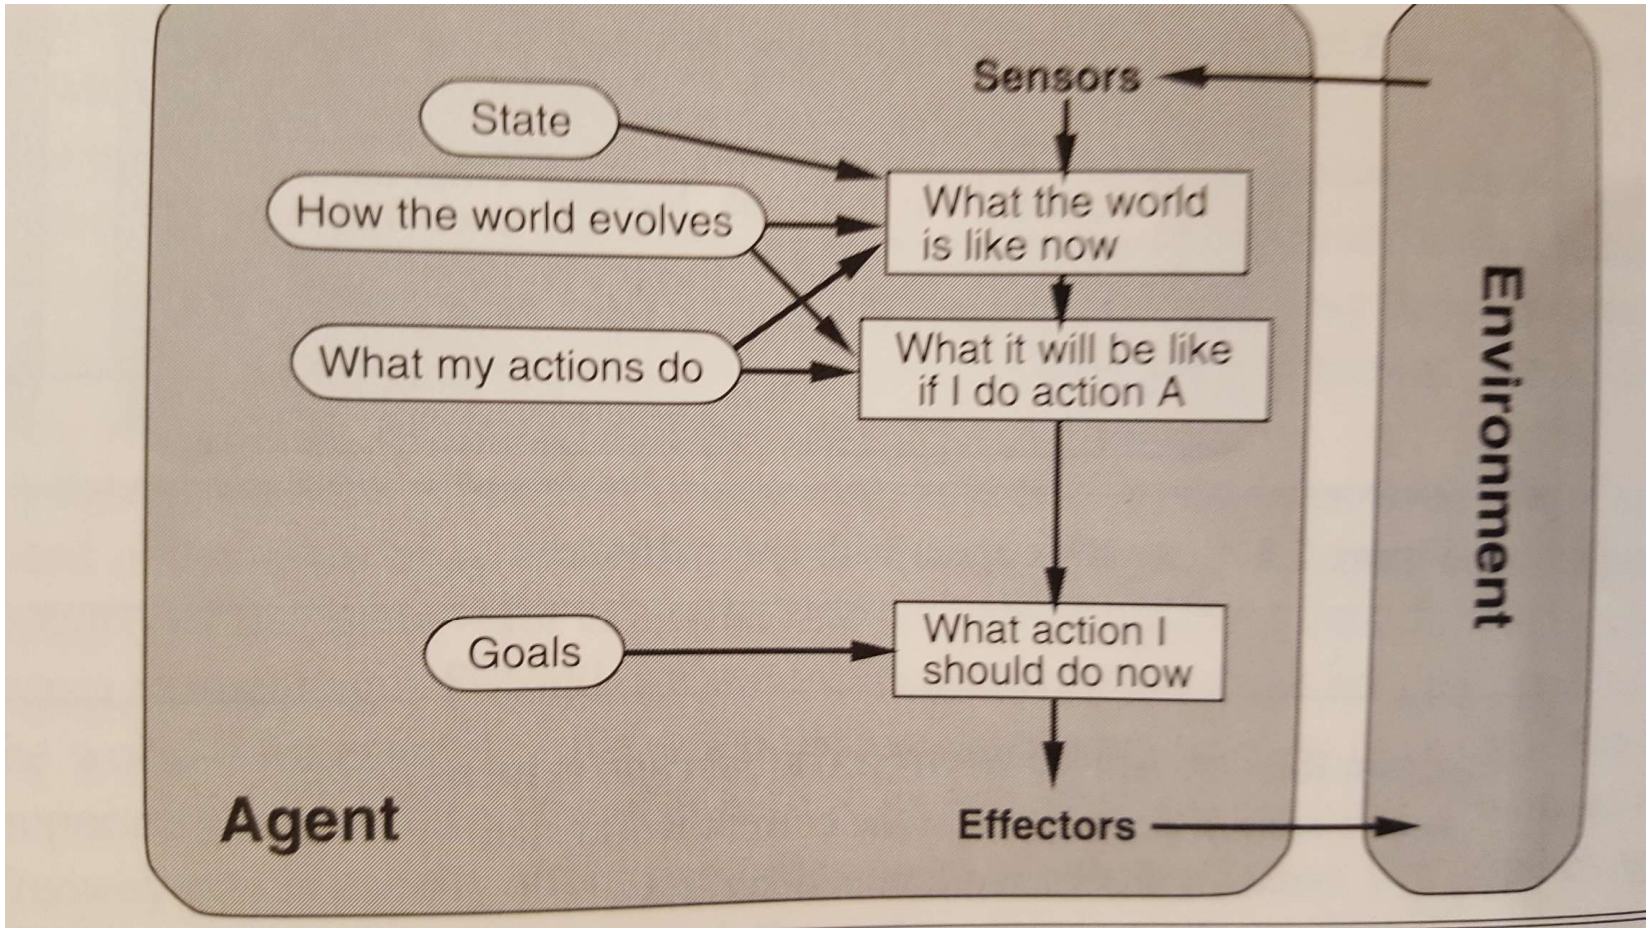

Utility-Directed Agent

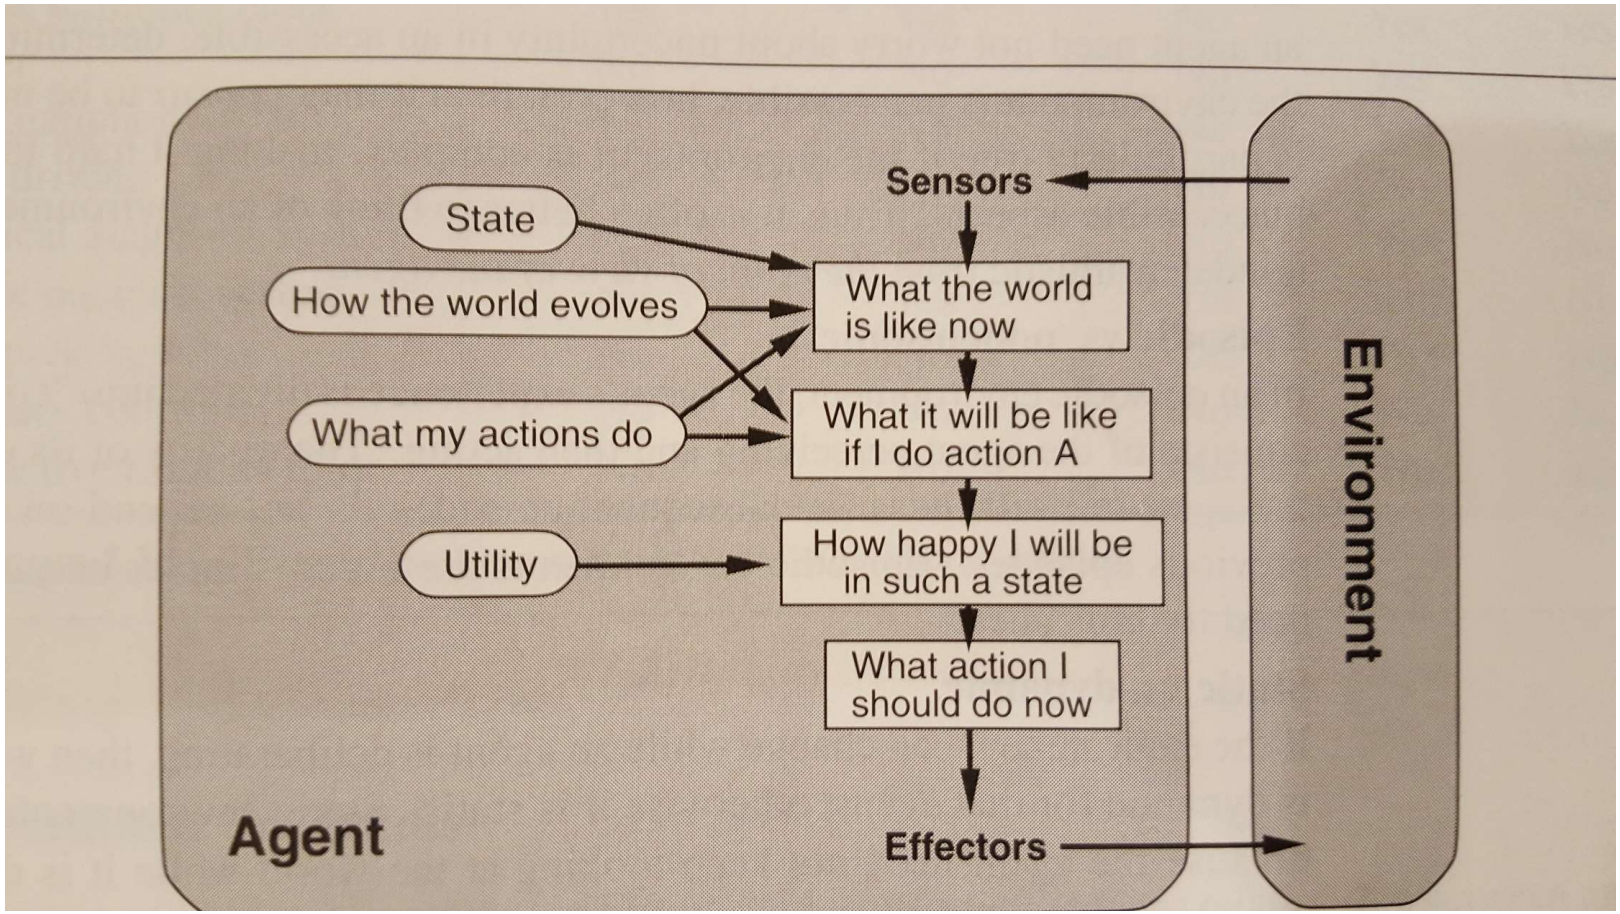

PEAS

- **P**erformance measure: Determined by the system designer, attempts to measure some intuitive description of behavior goodness.
- **A**ctions: Determined by the system designer, usually trades off cost versus utility
- **S**ensors: Determined by the system designer, usually trades off cost versus utility
- **E**nvironment: Completely out of the control of the system designer.

Properties of Environments

- 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?

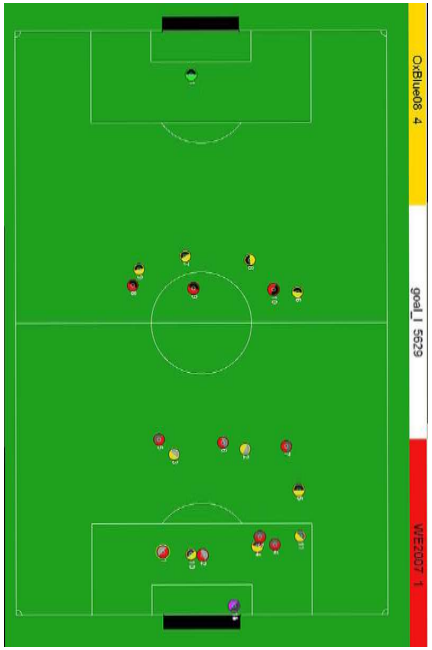

vs.

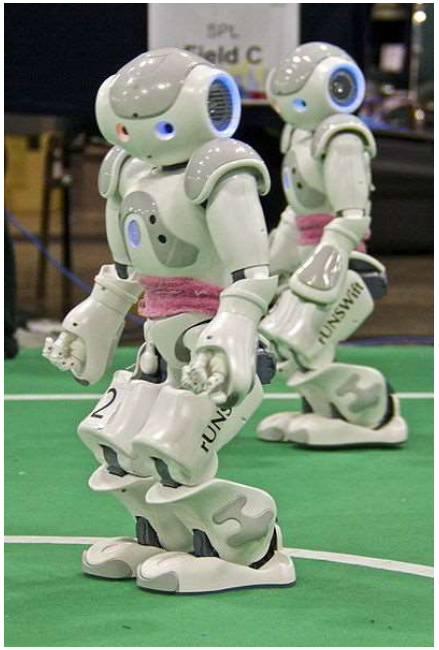

Source: L. Zettlemoyer

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

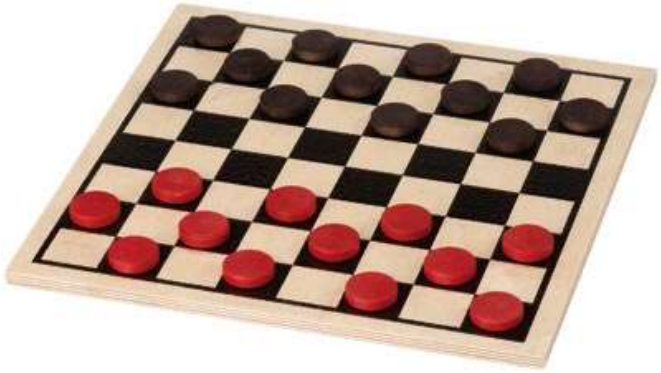

vs.

Episodic vs. sequential

- Is the agent's experience divided into unconnected episodes, or is it a coherent sequence of observations and actions?
  - Does each problem instance involve just one action or a series of actions that change the world state according to the transition model?

vs.

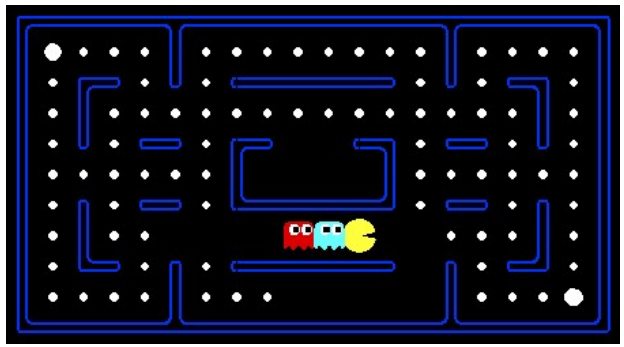

Static vs. dynamic

- 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

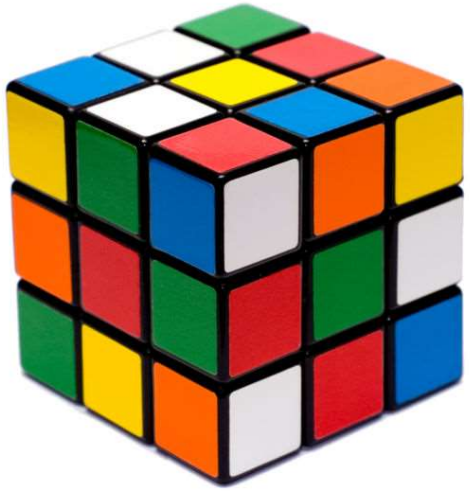

vs.

Discrete vs. continuous

- Does the environment provide a countable (discrete) or uncountably infinite (continuous) 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
  - "Distinct" = different values of utility

vs.

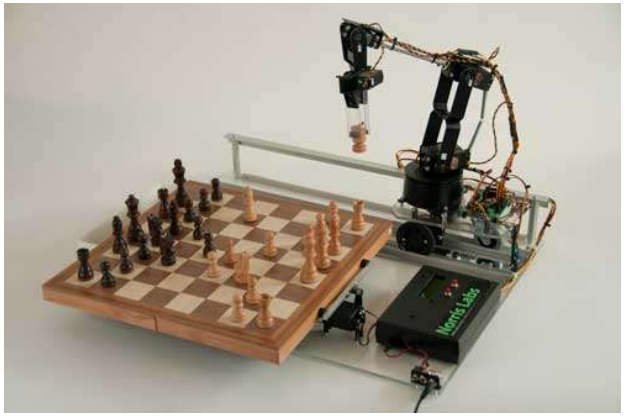

Single-agent vs. multiagent

- Is an agent operating by itself in the environment?

vs.

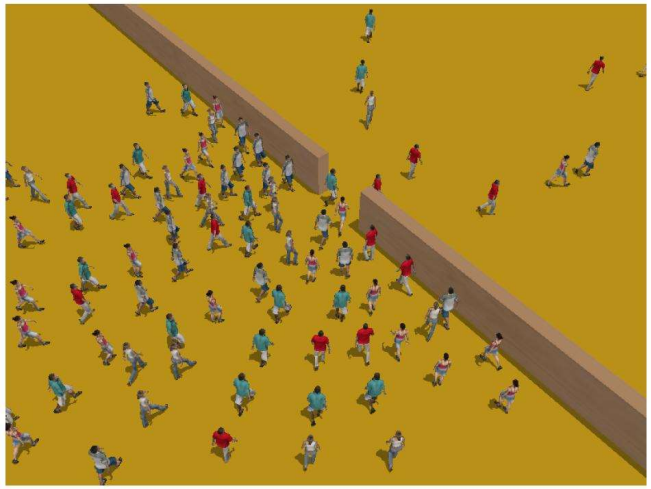

Known vs. unknown

- 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

	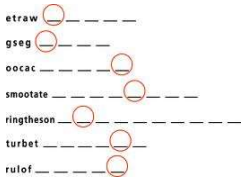 Word jumble solver	 Chess with a clock	 Scrabble	 Autonomous driving
Observable	Fully	Fully	Partially	Partially
Deterministic	Deterministic	Strategic	Stochastic	Stochastic
Episodic	Episodic	Sequential	Sequential	Sequential
Static	Static	Semidynamic	Static	Dynamic
Discrete	Discrete	Discrete	Discrete	Continuous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

Preview of the course

- Deterministic environments: search, constraint satisfaction, logic
  - 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