Artificial Intelligence

Intelligent Agents
Lecture 2

CS 444 – Spring 2019
Dr. Kevin Molloy
Department of Computer Science
James Madison University

Outline for Today

- Agents and Environments
- Rationality
- PEAS (Performance measure, Environment, Actuators, Sensors)
- Environment Types
- Agent Types

Agents and Environments

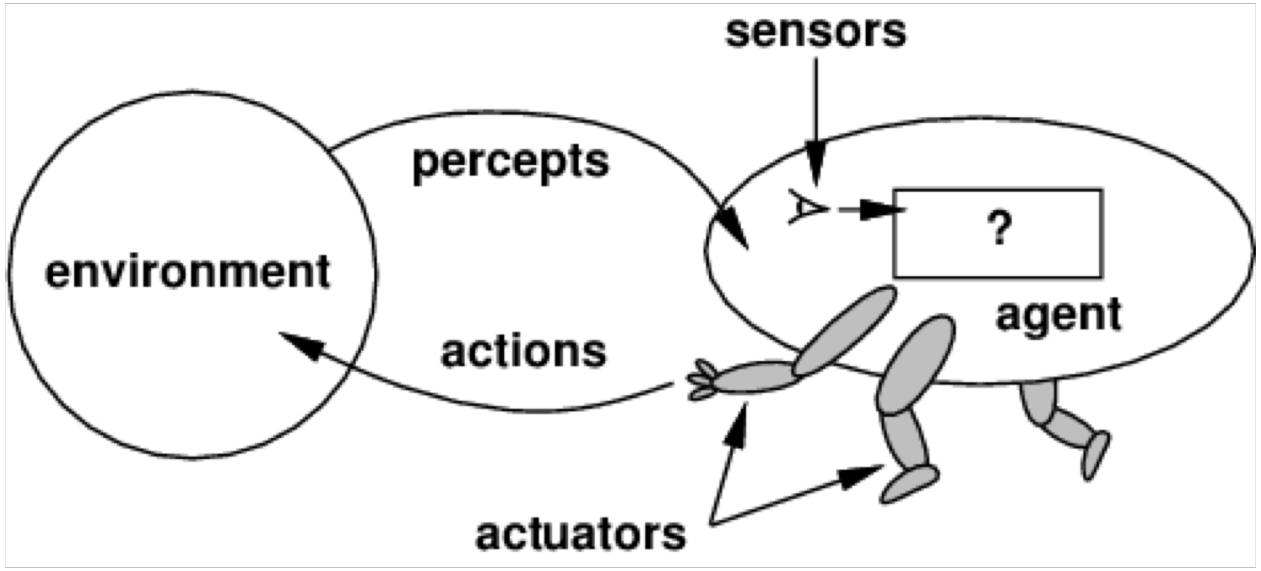

Agents include humans, robots, softbots, thermostats, etc.

The agent function maps from percept histories to actions: $f : \mathcal{P}^* \rightarrow \mathcal{A}$

The agent program runs on the physical architecture to produce f.

Vacuum-cleaner World

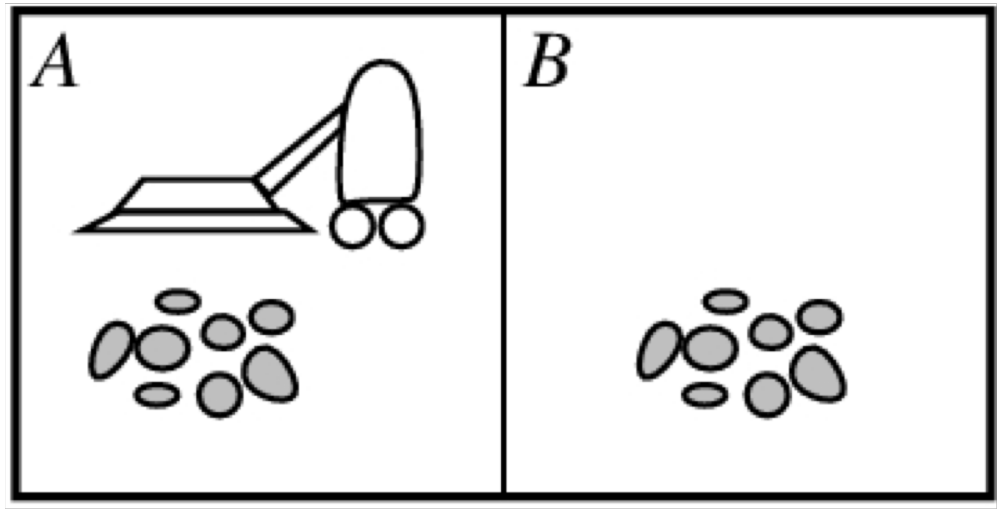

Percepts: location and contents, e.g., [A, Dirty]

Action: *Left, Right, Suck (remove the dirt) , NoOp*

A Vacuum-cleaner Agent

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
…..	…..

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 in a small agent program?

[note the difference between agent function and agent program]

Rationality

Fixed performance measure evaluates the sequence of environment states.

Possible performance measures:

- one point per square cleaned up in time *T* ?
- One point per clean square per time step, minus one per move?

A rational agent:

chooses which ever action maximizes the expected value of the performance measure given the percept sequence to date.

Rational ≠ omniscient

- Percepts may not supply all relevant information

Rational ≠ clairvoyant

- Action outcomes may not be as expected

Hence, rational does not always equal successful.

PEAS

To design a rational agent, we must first specific the task environment – PEAS.

Performance measure

Environment

Actuators

Sensors

PEAS – Automated Taxi

Performance measure:	safety, destination, profits, comfort,…
Environment:	US streets/freeways, traffic, pedestrians, weather, …
Actuators:	steering, accelerator, brake, horn, …
Sensors:	Video, acceleromters, gauges (gas, oil), GPS, keyboard, microphone

PEAS – Internet Shopping Agent

Performance measure:	price, quality, efficiency
Environment:	current and future web sites, vendors
Actuators:	display to user, follow URL, fill in form
Sensors:	parse HTML pages (text, graphic, scripts)

Environment Types

Do the agent's sensors give complete information (relevant to the choice of action) about the estate of the environment at each point in time?

- Fully vs. partially-observable

Does the agent operate in an environment with other agents?

- Single vs. multi-agent (competitive, cooperative)

Is the next state of the environment complete determined by the current state and agent action?

- Episodic vs. sequential

Can the environment change while the agent is deliberating?

- Static vs dynamic

What is the domain of values for variables racking environment state, agent state, and time?

- Discrete vs. continuous

Does the agent know outcomes of all its actions?

- Known vs unknown

Environment Types

	Solitaire	Poker	Backgammon	Internet Shopping	Automated Taxi
Observable	Yes	No	Yes	No	No
Deterministic	Yes	No	No	Partly	No
Episodic	No	No	No	No	No
Static	Yes	Yes	Yes	Semi	No
Discrete	Yes	Yes	Yes	Yes	No
Single-agent	Yes	No	No	Yes (except actions)	No

The environment type largely determines the agent design.

Agent Types

Four basic types of agents:

- Simple reflex agents
- Reflex agents w/state
- Goal-based agents
- Utility-based agents

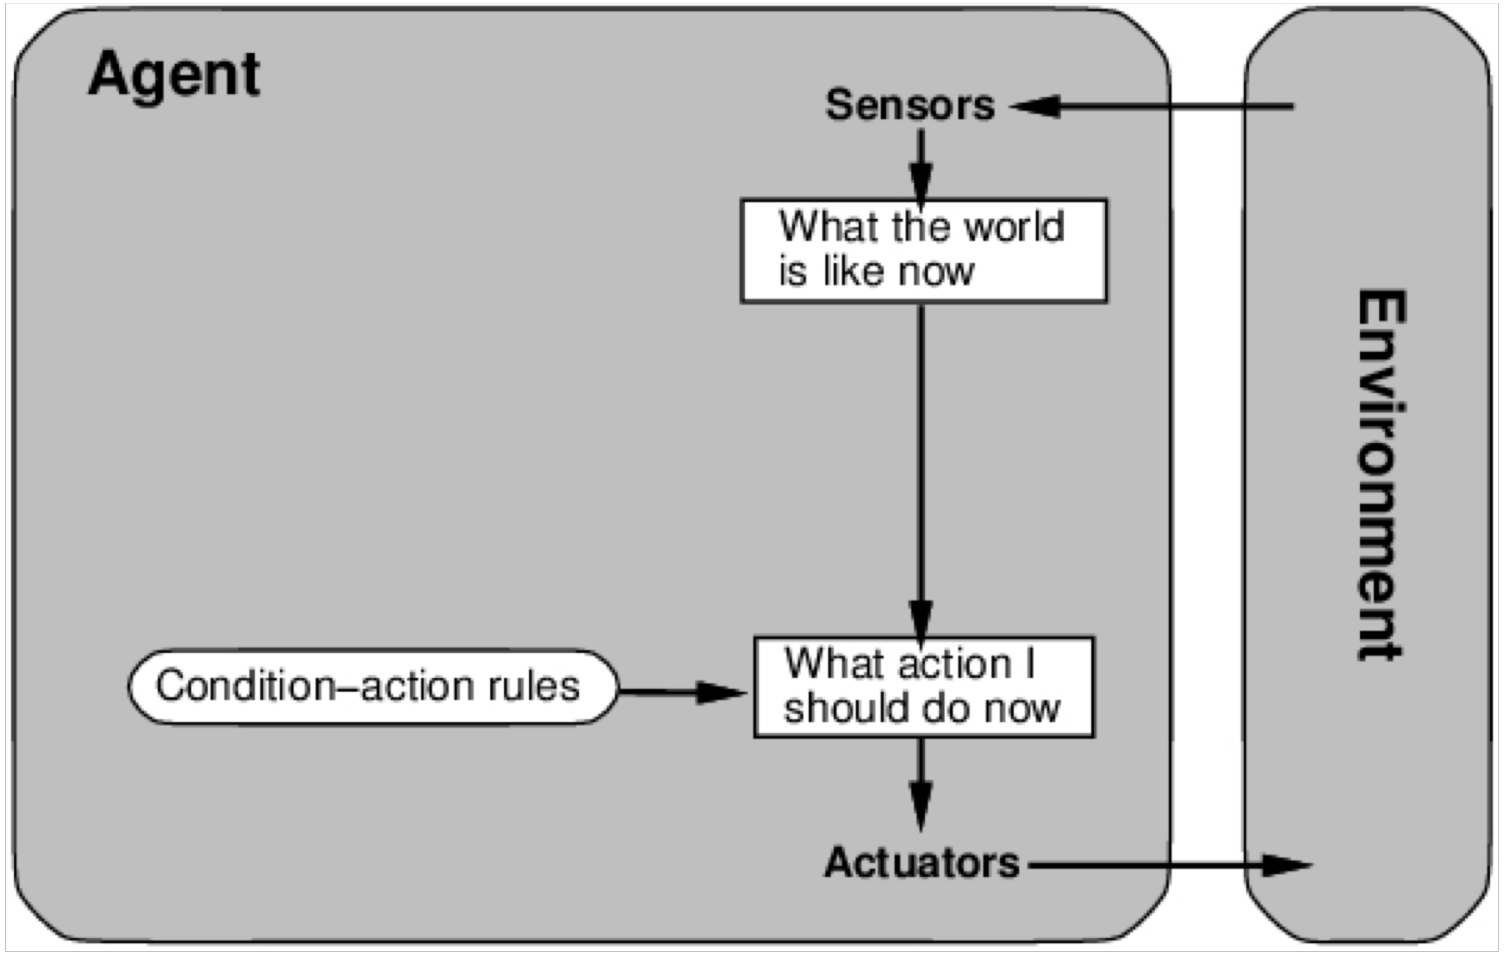

A simple reflex agent

Reflex agent example

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*

Can a reflex agent be rational?

A rational agent:

chooses which ever action maximizes the expected value of the performance measure given the percept sequence to date.

Depends, on the performance measure:

- 1 pt for each clean square in each time step
- Geography is known *a priori*
- Agent correctly perceives its location and dirt, and the cleaning mechanism works 100% of the time.

Reflex agent w/state (Model-based agents)

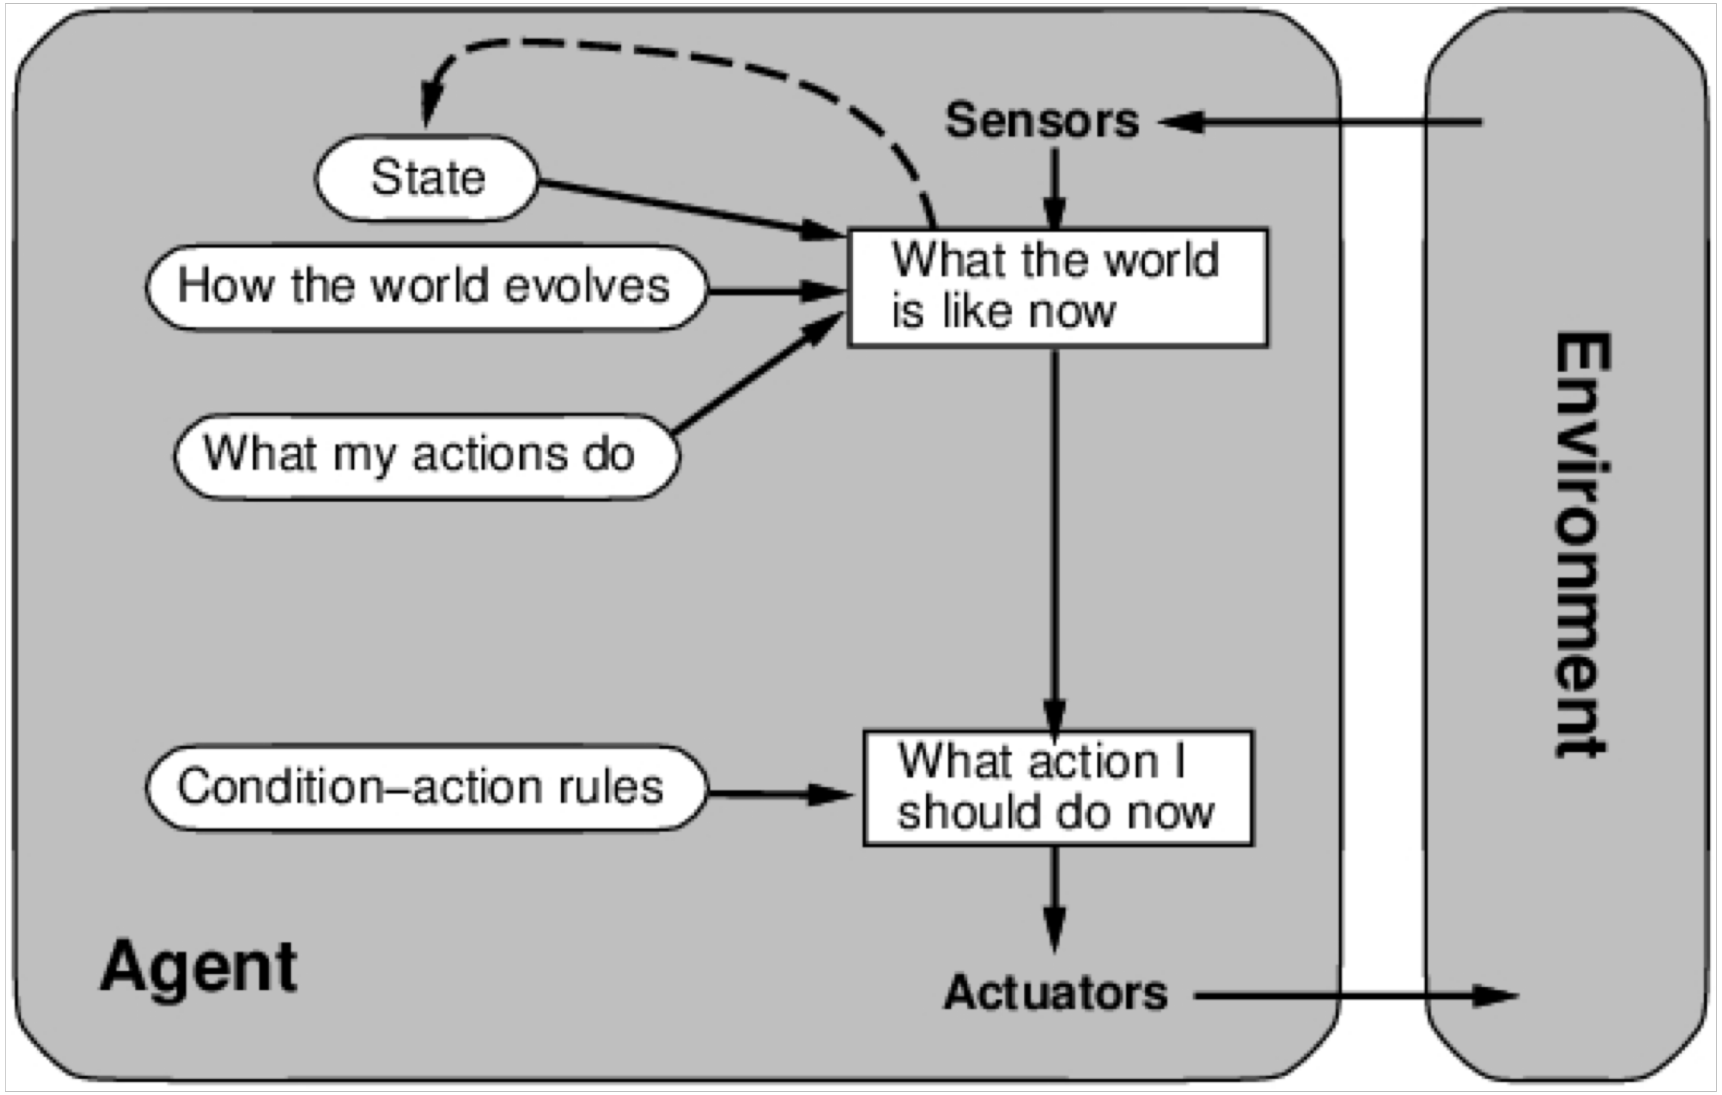

Goal based agents

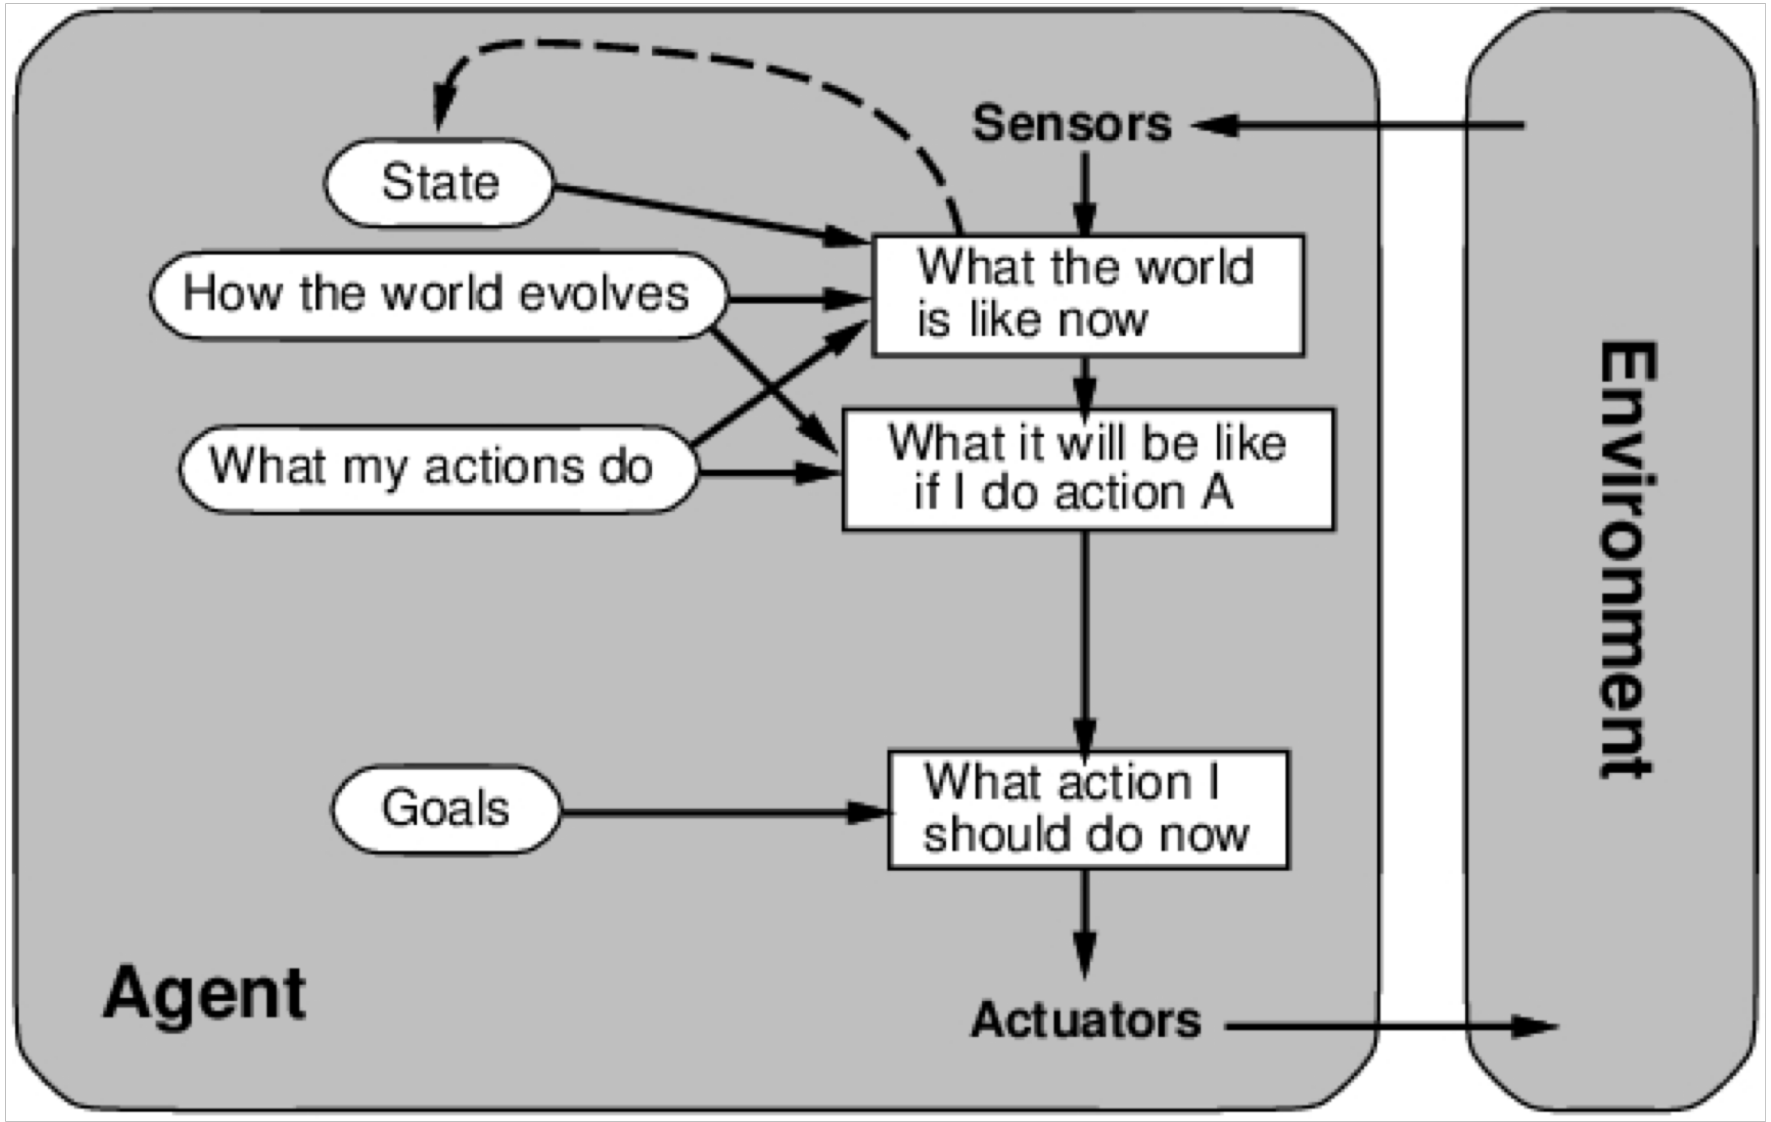

Utility-based Agents

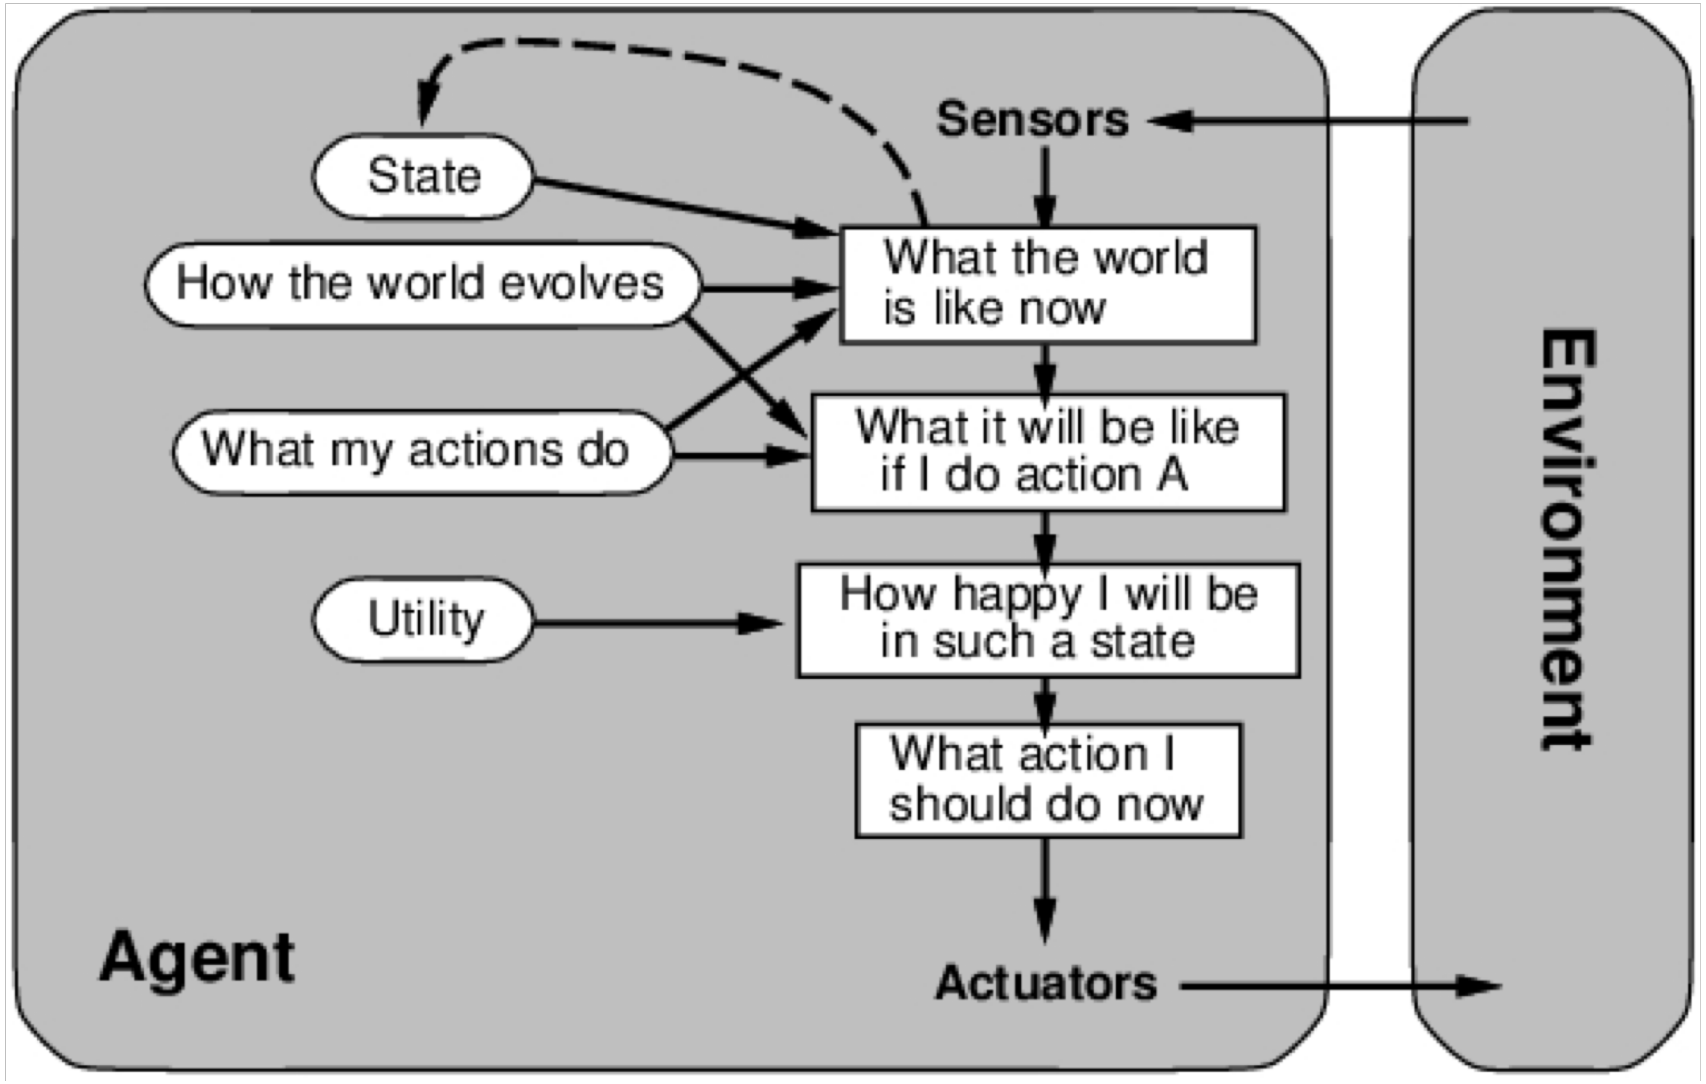

Learning-based agents

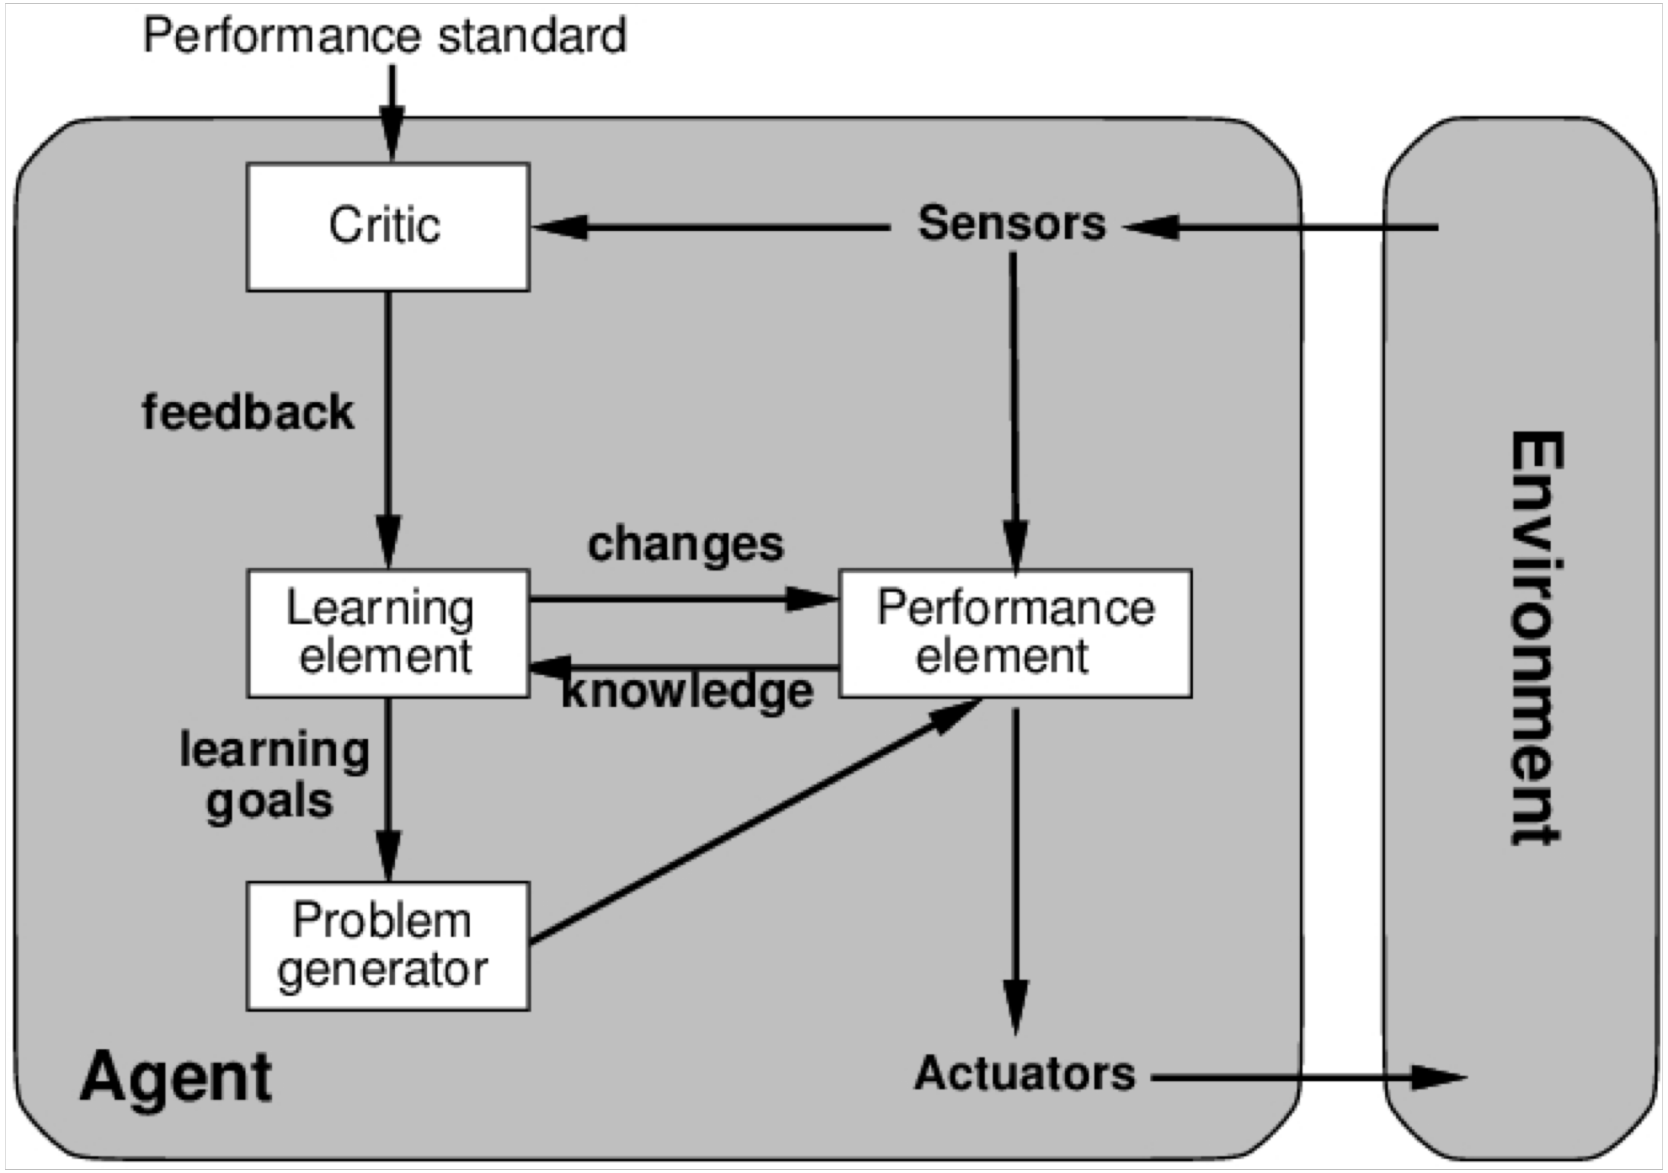

Summary

- Agents interact with environments through actuators and sensors
- The agent function describes that the agent does in all circumstances
- The agent program determines what to do next
- The performance measure evaluates the environment sequence
- A perfectly rational agent maximizes expected performance
- Agent programs implement (some) agent function
- PEAS descriptions define task environments. Environments are categorized as:
  - Fully- vs. partially observable
  - *Deterministic vs stochastic*
  - *Episodic vs sequential*
  - *Static vs dynamic*
  - *Discrete vs continuous*
  - *Single vs. multi-agent*

Problem 2.1

Suppose that the performance measure is concerned with just the first *T* time steps of the environment and ignores everything thereafter. Show that a rational agent's action may depend not just on the state of the environment but also on the time step it has reached (in other words, its nots necessarily episodic).

Problem 2.2

Show that the simple vacuum-cleaner agent is indeed rational,, given the following assumptions.

- The performance measure awards one pint for each clean square at each time step (1,000 time steps)
- The geography is known ahead of time.
- Only actions are left, right, and suck.
- Agent correctly perceived is location an weather that location contains dirt.

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*

Problem 2.3

For each of the following, state true or false and support your claim.

a) An agent that senses only partial information about the state cannot be perfectly rational.

d) The input to an agent program is the same as the input to the agent function.

i) A perfectly rational poker-playing agent never loses.

Problem 2.6

Differences between agent function and agent programs.

a) Can there be more than one agent program that implements a given agent function? Give an example or show why one is not possible.

b) Are there agent functions that can not be implemented by any agent program?

d) Given an architecture with *n* bits of storage, how many different possible agent programs are there?

e) Suppose we keep the agent program fixed but speed up the machine by a factor of two. Does that change the agent function?

Why?

All the problems in this class can be categorized using these terms.

Thus, we will be learning the tradeoffs between these approaches and what types of problems for which they are suited.

Next Class

- Read Chapter 3 pg 64 – 91
- Watch for instructions on how to configure python to use code from class.