


Last update: August 30, 2012

# INTELLIGENT AGENTS

## CMSC 421: CHAPTER 2

# Agents and environments



- ◇ Russell & Norvig's book is based on the notion of *intelligent agents*
  - 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  $f$

# Outline

Chapter 2 describes some basic concepts relating to agents

- ◇ Agents and environments
- ◇ Rationality
- ◇ The PEAS model of an environment
- ◇ Environment types
- ◇ Agent types

## Vacuum-cleaner world



Percepts: location and contents, e.g.,  $[A, \textit{Dirty}]$

Actions: *Left*, *Right*, *Suck*, *NoOp*

## A vacuum-cleaner agent

Percept sequence	Action
$[A, Clean]$	<i>Right</i>
$[A, Dirty]$	<i>Suck</i>
$[B, Clean]$	<i>Left</i>
$[B, Dirty]$	<i>Suck</i>
$[A, Clean], [A, Clean]$	<i>Right</i>
$[A, Clean], [A, Dirty]$	<i>Suck</i>
$\vdots$	$\vdots$

```
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** function?

Can it be implemented in a small agent program?

# Rationality

- ◇ Fixed *performance measure* evaluates the *environment sequence*
  - one point per square cleaned up in time  $T$ ?
  - one point per clean square per time step, minus one per move?
  - penalize for  $> k$  dirty squares?
- ◇ A *rational agent* chooses whichever action maximizes the *expected value* of the performance measure, *given the percept sequence to date*
  - Rational  $\neq$  omniscient
    - percepts may not supply all relevant information
  - Rational  $\neq$  clairvoyant
    - action outcomes may not be as expected
  - Hence, rational  $\neq$  successful
  - Rational  $\implies$  exploration, learning, autonomy

# PEAS

To design a rational agent, we must specify the **task environment**

Consider, e.g., the task of designing an automated taxi:

Performance measure?

Environment?

Actuators?

Sensors?

# PEAS

To design a rational agent, we must specify the **task environment**

Consider, e.g., the task of designing an automated taxi:

*Performance measure?* safety, destination, profits, legality, comfort, ...

*Environment?* streets/freeways, traffic, pedestrians, weather, ...

*Actuators?* steering, accelerator, brake, horn, speaker/display, ...

*Sensors?* video, accelerometers, gauges, engine sensors, keyboard, GPS, ...



# Internet shopping agent

Performance measure?

Environment?

Actuators?

Sensors?

# Internet shopping agent

Performance measure? price, quality, appropriateness, efficiency

Environment? current and future WWW sites, vendors, shippers

Actuators? display to user, follow URL, fill in form

Sensors? HTML pages (text, graphics, scripts)

# Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
<u>Fully observable?</u>				
<u>Deterministic?</u>				
<u>Episodic?</u>				
<u>Static?</u>				
<u>Discrete?</u>				
<u>Single-agent?</u>				

## Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
<u>Fully observable?</u>	No	Yes	No	No
<u>Deterministic?</u>				
<u>Episodic?</u>				
<u>Static?</u>				
<u>Discrete?</u>				
<u>Single-agent?</u>				

## Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
<u>Fully observable?</u>	No	Yes	No	No
<u>Deterministic?</u>	Yes*	No	Partly	No
<u>Episodic?</u>				
<u>Static?</u>				
<u>Discrete?</u>				
<u>Single-agent?</u>				

\*After the cards have been dealt

*Episodic*: task divided into episodes, each to be considered only by itself

*Sequential*: Current decision may affect all future decisions

## Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
<u>Fully observable?</u>	No	Yes	No	No
<u>Deterministic?</u>	Yes*	No	Partly	No
<u>Episodic?</u>	No	No	No	No
<u>Static?</u>				
<u>Discrete?</u>				
<u>Single-agent?</u>				

\*After the cards have been dealt

*Static*: the world does not change while the agent is thinking

## Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
<u>Fully observable?</u>	No	Yes	No	No
<u>Deterministic?</u>	Yes*	No	Partly	No
<u>Episodic?</u>	No	No	No	No
<u>Static?</u>	Yes	Partly	Partly	No
<u>Discrete?</u>				
<u>Single-agent?</u>				

\*After the cards have been dealt

## Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
<u>Fully observable?</u>	No	Yes	No	No
<u>Deterministic?</u>	Yes*	No	Partly	No
<u>Episodic?</u>	No	No	No	No
<u>Static?</u>	Yes	Partly	Partly	No
<u>Discrete?</u>	Yes	Yes	Yes	No
<u>Single-agent?</u>				

\*After the cards have been dealt



## Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
<u>Fully observable?</u>	No	Yes	No	No
<u>Deterministic?</u>	Yes*	No	Partly	No
<u>Episodic?</u>	No	No	No	No
<u>Static?</u>	Yes	Partly	Partly	No
<u>Discrete?</u>	Yes	Yes	Yes	No
<u>Single-agent?</u>	Yes	No	No	No

\*After the cards have been dealt

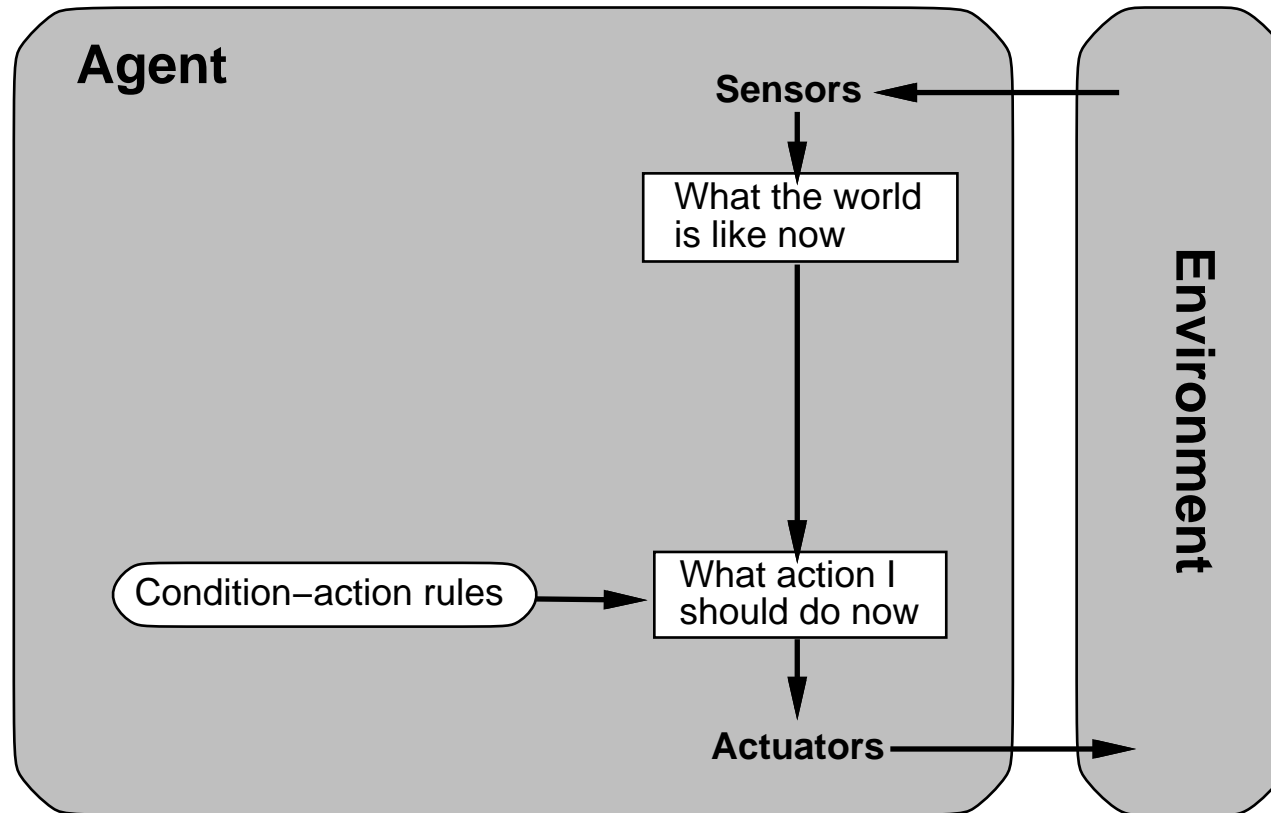
**The environment type largely determines the agent design**

Real world: partially observable, stochastic, sequential, dynamic, continuous, multi-agent

# Agent types

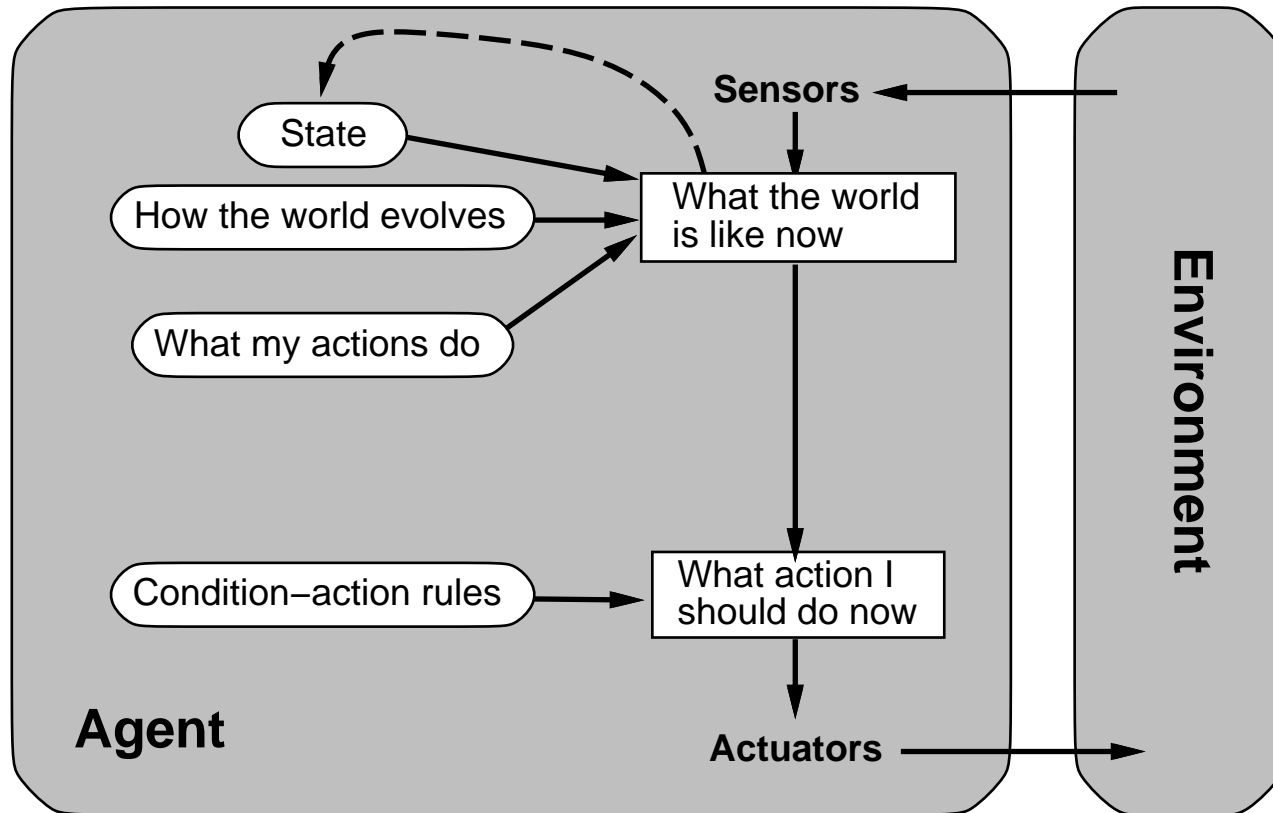
- ◇ Four basic types in order of increasing generality:
  - simple reflex agents
  - reflex agents with state
  - goal-based agents
  - utility-based agents
- ◇ All of these can be turned into learning agents

# Simple reflex agent



```
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
```

## Reflex agent with state



- ◇ E.g., a vacuum-cleaning agent that can't sense what room it's in, but can remember its location
  - next page ...

# Example

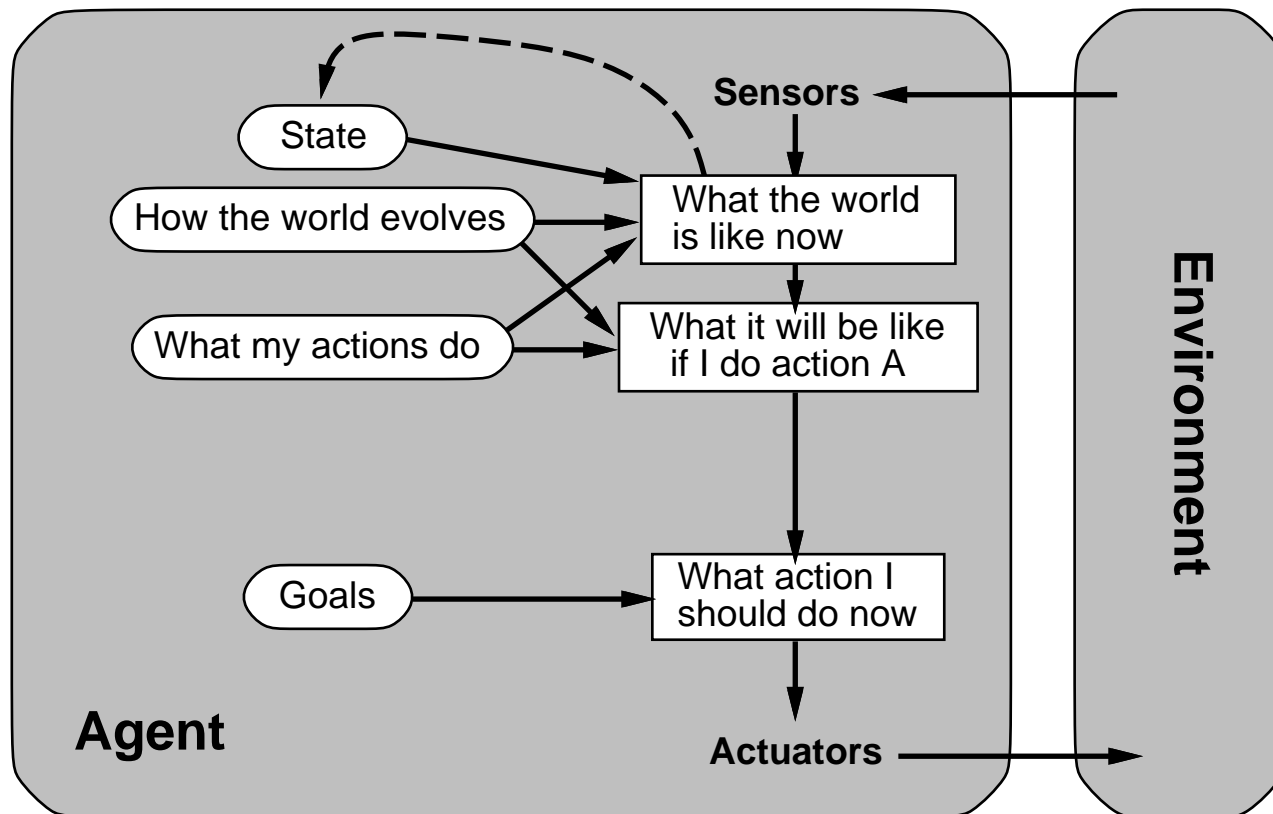
Vacuum-cleaning agent that remembers its location:

```
function VACUUM-AGENT-WITH-STATE( [location] ) returns an action
static: location, initially A

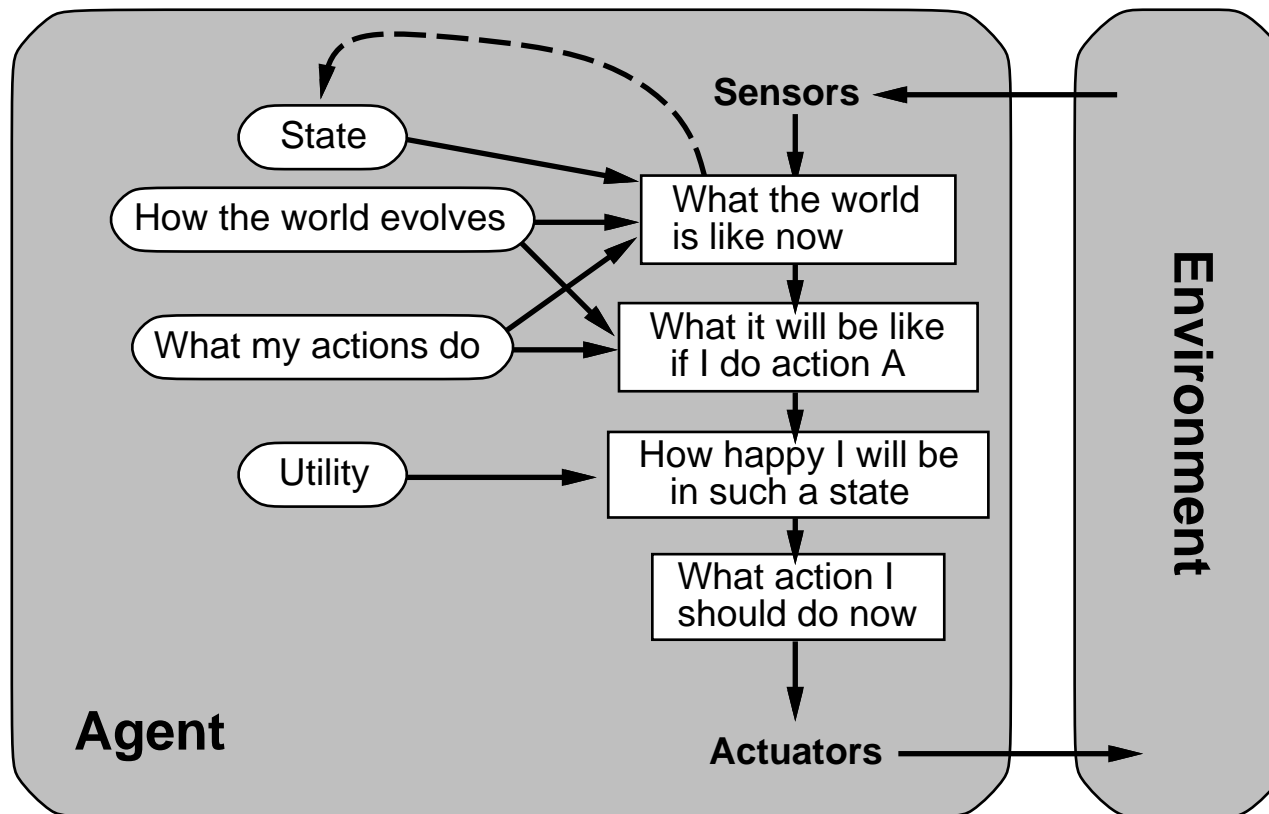
  if status = Dirty then return Suck
  else if location = A then
    location ← B
    return Right
  else if location = B then
    location ← A
    return Left
```

- ◇ The pseudocode assumes the agent always starts out in room *A*.
  - What if we didn't know this?

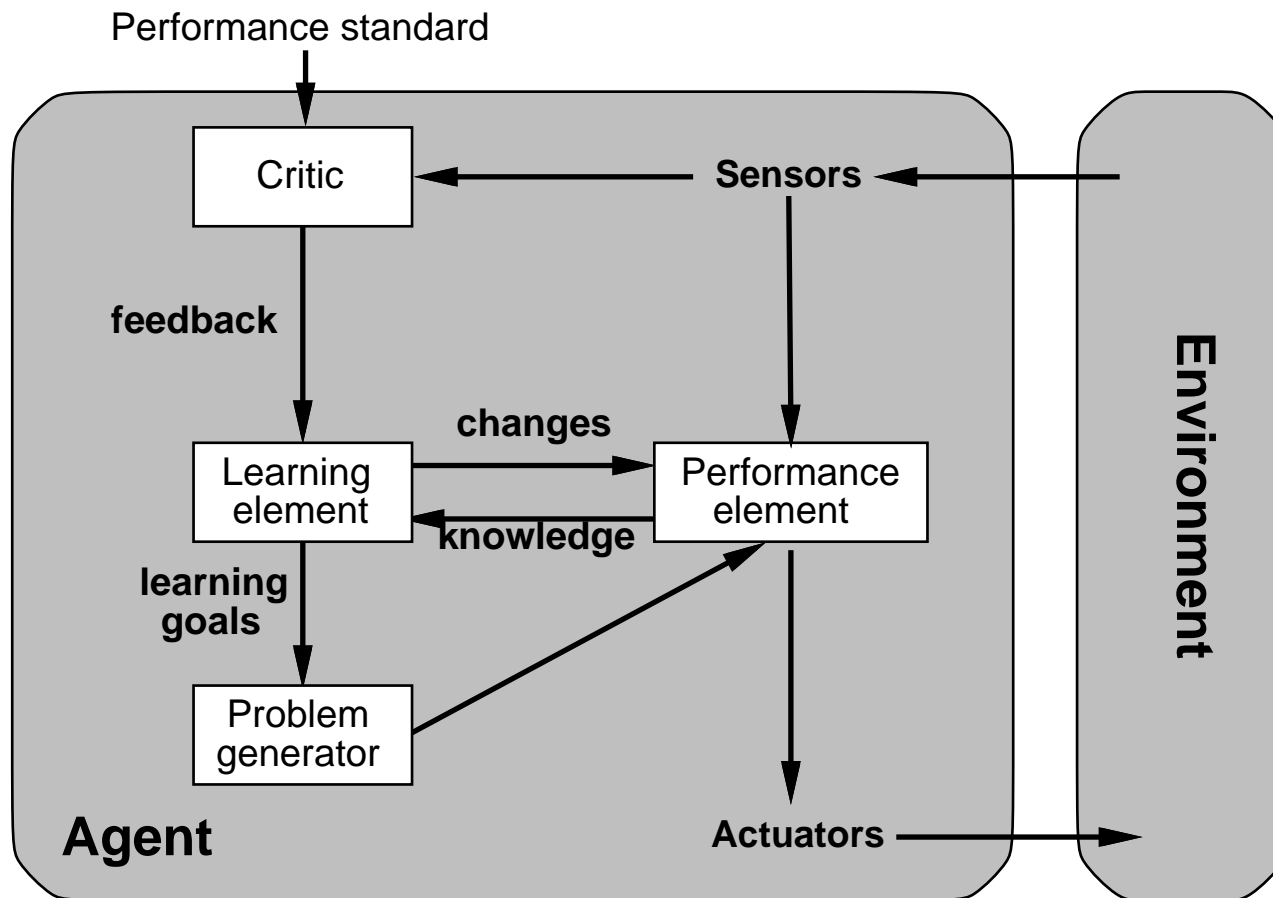
# Goal-based agent



# Utility-based agent



# Learning agent





# Summary

- ◇ *Agents* interact with *environments* through *actuators* and *sensors*
- ◇ The *agent function* describes what the agent does in all circumstances
- ◇ The *performance measure* evaluates the environment sequence
- ◇ A *perfectly rational* agent maximizes expected performance
- ◇ *Agent programs* implement (some) agent functions
- ◇ *PEAS* descriptions define task environments
- ◇ Environments are categorized along several dimensions:
  - *observable, deterministic, episodic, static, discrete, single-agent*
- ◇ Some basic agent architectures:
  - *reflex, reflex with state, goal-based, utility-based*