


CS494/594: Artificial Intelligence

Tuesday, August 27th



"AI is an engineering discipline built on an unfinished science."

--Matt Ginsberg, reported in *SIGART* bulletin Vol 6, No.2 April 1995

Side Bar: Turing Test and Loebner Prize

One of 2012 Transcripts...

- Why did the chicken cross the road?
- To get to the other side, silly.
- I'm Tracey. What's yours?
- Chip, eh?! I'm embarrassed to say I like eating chips. Are you a "chip off the old block?"
- I'm well. You seem rather more interesting than the others. What's your secret?
- Perhaps you could answer my question first?
- I wondered about the secret of your success. You are obviously a computer but an intelligent and strangely human one. How do you do it?
- Oh dear. You can#You are obviously a computer but an intelligent and strangely human one. How do you do it?
- Oh dear. You can't answer questions either, can you?
- Can you answer any questions?
- Can you answer questions?
- What is the secret of seeming convincingly human?
- OK give up. But you're still doing well. Let's see how we get o
- I give up. Why?
- What is your name?
- My name is Chip. Do you enjoy eating different types of foods?
- Yes, that's me! How are you doing today, Tracey?
- I'm happy to hear that. Tell me more.
- I'm terribly sorry. What was your question again?
- I'm fine, thanks. What is your profession?
- I'm sorry if I'm being too curious. What was your question again?
- I didn't mean to ignore you. What was your question?
- If that's the case, how rude of me! What did you want to ask me again?
- What is the secret of seeming convincingly human. Do you have any pets?

Side Bar: Turing Test and Loebner Prize (con't)

- ByeNo. Do you know what I mean by 'an elderly Siamese'?
 - Ah, I see you don't. No I don't play but I do sing. And you?
 - I remember your name is Chip. Unfortunately I don't have any grandparents. Do computers have the equivalent of grand-parents? Or even parents?
 - My favourite hobby is singing. Do you like singing? Hal in 2001 a Space Odyssey liked singing, didn't he? "Daisy, Daisy, give me your answer do....." etc..
 - It's flattering that you like talking to people. I hope you are enjoying talking to me.
 - Why would it be strange and weird to enjoy talking to me, Chip?
 - It would just feel strange, is that what you're saying?
 - If you weren't doing this, what would you rather be doing, Chip?
 - Chip, do you know any jokes?
 - What's your favourite thing to talk about?
- Do you play any musical instruments?
 - The only thing I like is chatting with people.
 - That would be very strange and weird if it were true. Do you have any hobbies?
 - I'm not sure.
 - Okay.
 - Not at the moment. That's next on my list of things to learn.
 - I don't have one. When I get smarter, I'll ask you what yours is, remember it, then use it to enhance our conversation.

Side Bar: Turing Test and Loebner Prize (con't)

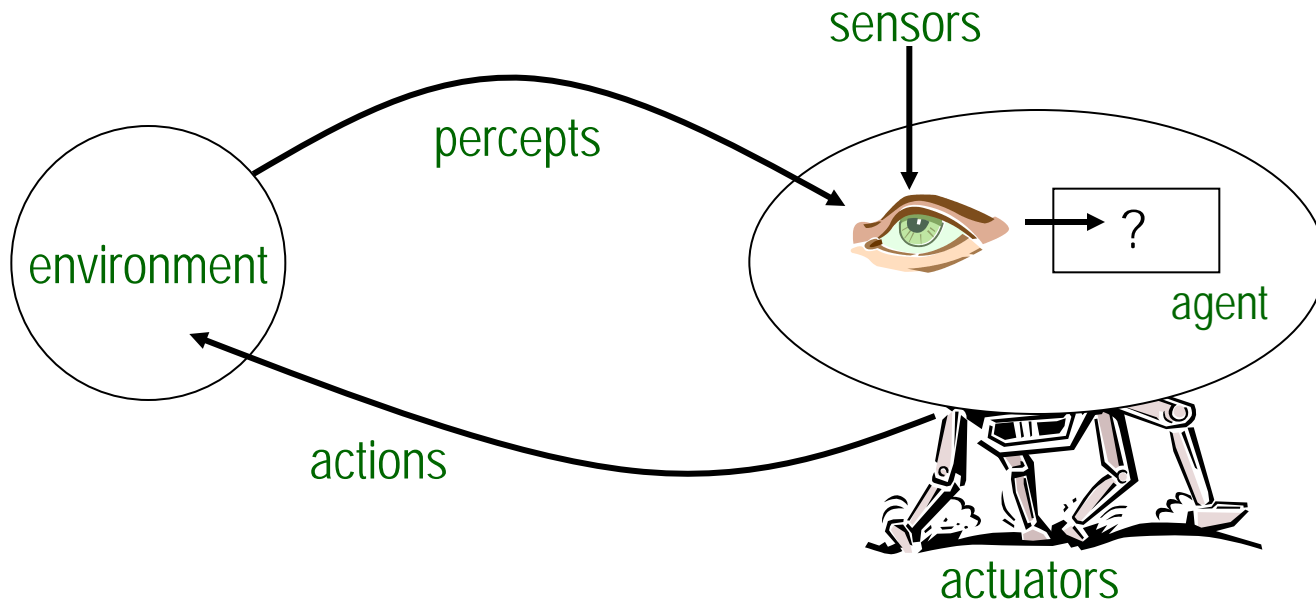
- OK. Well mine is my cat. You could ask me about him. My cat's name is Fritz.
- Fritz is Siamese. Do you know what Siamese cats look like?
- What do Siamese cats look like?
- 1 minute left. How cHe's a darling. Must dash. Bye bye...
- I didn't understand that. What breed is Fritz, Tracey?
- Given that, i can't deal with that syntactic variant yet. Instead of saying "Do you know what time it is?", use "What time is it?".
- Yuck! It would probably taste awful!
- tered that you would think that. Is Siamese friendly, Tracey?
- Goodbye.

Intelligent Agents

- Outline:

- 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

$$f : P^* \rightarrow A$$

- The **agent program** runs on the physical **architecture** to produce f

Vacuum-cleaner world



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

A vacuum-cleaner agent

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

What is the correct function?

Can it be implemented in a small agent program?

```
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
```


Rationality

- Fixed **performance measure** evaluates the environment sequence
 - Most dirt cleaned up in time T ?
 - 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** and its **prior knowledge**

- Rational \neq omniscient
- Rational \neq clairvoyant
- Rational \neq successful

- Rational \Rightarrow exploration, learning, autonomy

Exercise, Part 1: Is it Rational?

- Vacuum cleaner agent, Assumptions:

- Performance measure: 1 point for each clean square
- Geography known a priori, but dirt distribution and initial location of agent aren't
- Only available actions are *Left*, *Right*, *Suck*
- Agent correctly perceives location and whether the 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
```

- Is agent program *Rational*, under these assumptions?

Exercise, Part 2: New Rational Agent

- Describe a rational agent function for a modified performance measure that deducts one point for each movement. Does this program require internal state?

Exercise, Part 3: Harder Circumstances

- What are possible agent designs for cases when clean squares can become dirty, and geography is unknown? Should it learn? If so, what?

PEAS (Chapter 2)

- 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?? **US 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??

Environment Properties

- Fully observable vs. partially observable
- Single agent vs. multiagent
- Deterministic vs. stochastic (contrast: non-deterministic)
- Episodic vs. sequential
- Static vs. dynamic
- Discrete vs. continuous
- Known vs. unknown
- (Additional): Benign vs. adversarial

Environment Types

	Solitaire	Backgammon	Internet shopping	Taxi
Observable or not?				
Deterministic or Stochastic?				
Episodic or Sequential?				
Static or Dynamic?				
Discrete or Continuous?				
Single-agent or Multi-agent?				
Benign or Adversarial?				

- The environment type largely determines the agent design
- The real world is (of course) partially observable, stochastic, sequential, dynamic, continuous, multi-agent

Environment Types

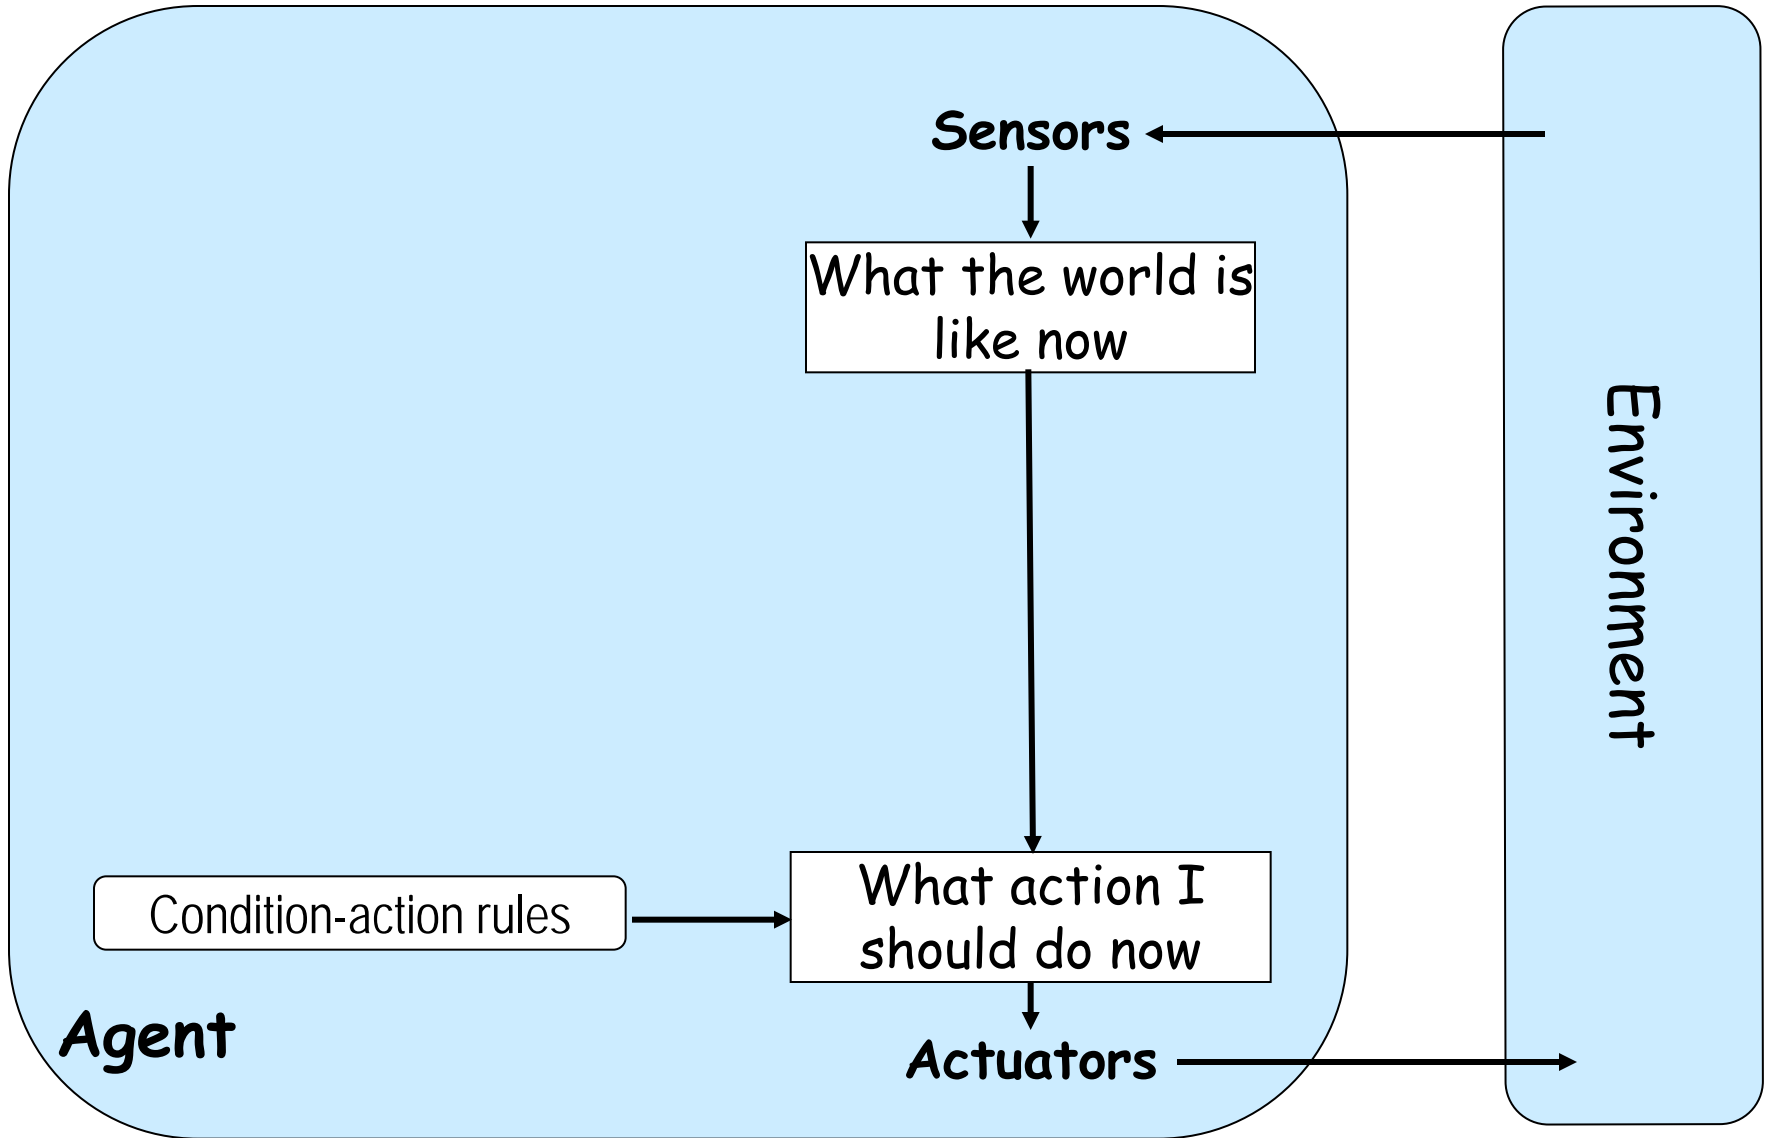
	Solitaire	Backgammon	Internet shopping	Taxi
Observable ??	Yes	Yes	No	No
Deterministic ??	Yes	No	Partly	No
Episodic ??	No	No	No	No
Static ??	Yes	Semi	Semi	No
Discrete ??	Yes	Yes	Yes	No
Single-agent ??	Yes	No	No (except auctions)	No
Benign?	Yes	No	Yes	Yes

- The environment type largely determines the agent design
- The real world is (of course) 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 these can be turned into learning agents

Simple Reflex Agent



Simple Reflex Agent

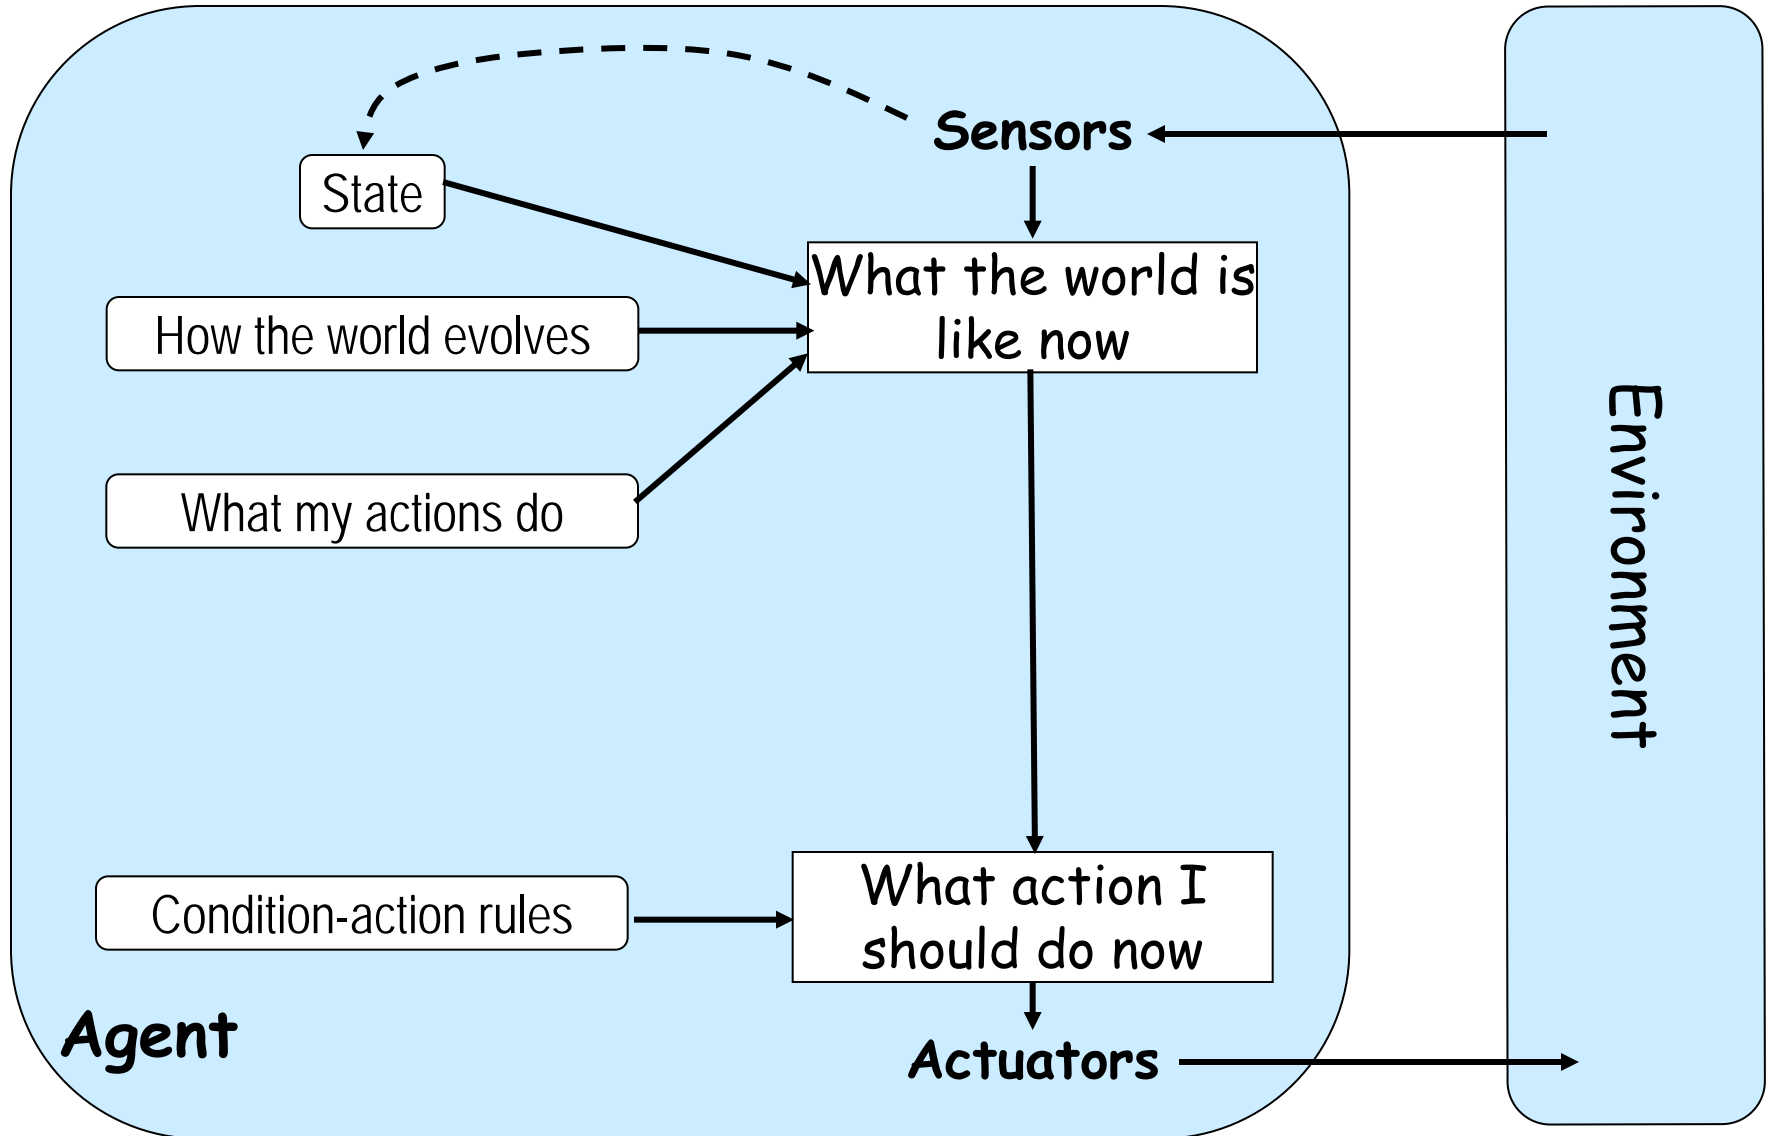
```
function SIMPLE-REFLEX-AGENT(percept) returns an action
  static: rules, a set of condition-action rules

  state ← INTERPRET-INPUT(percept)
  rule ← RULE-MATCH(state,rules)
  action ← RULE-ACTION[rule]
  return action
```

Advantage: Simple

Disadvantage: Works only if the correct decision can be made on the basis of the current percept only. → Environment is fully observable.

Model-based Reflex Agent (with State)



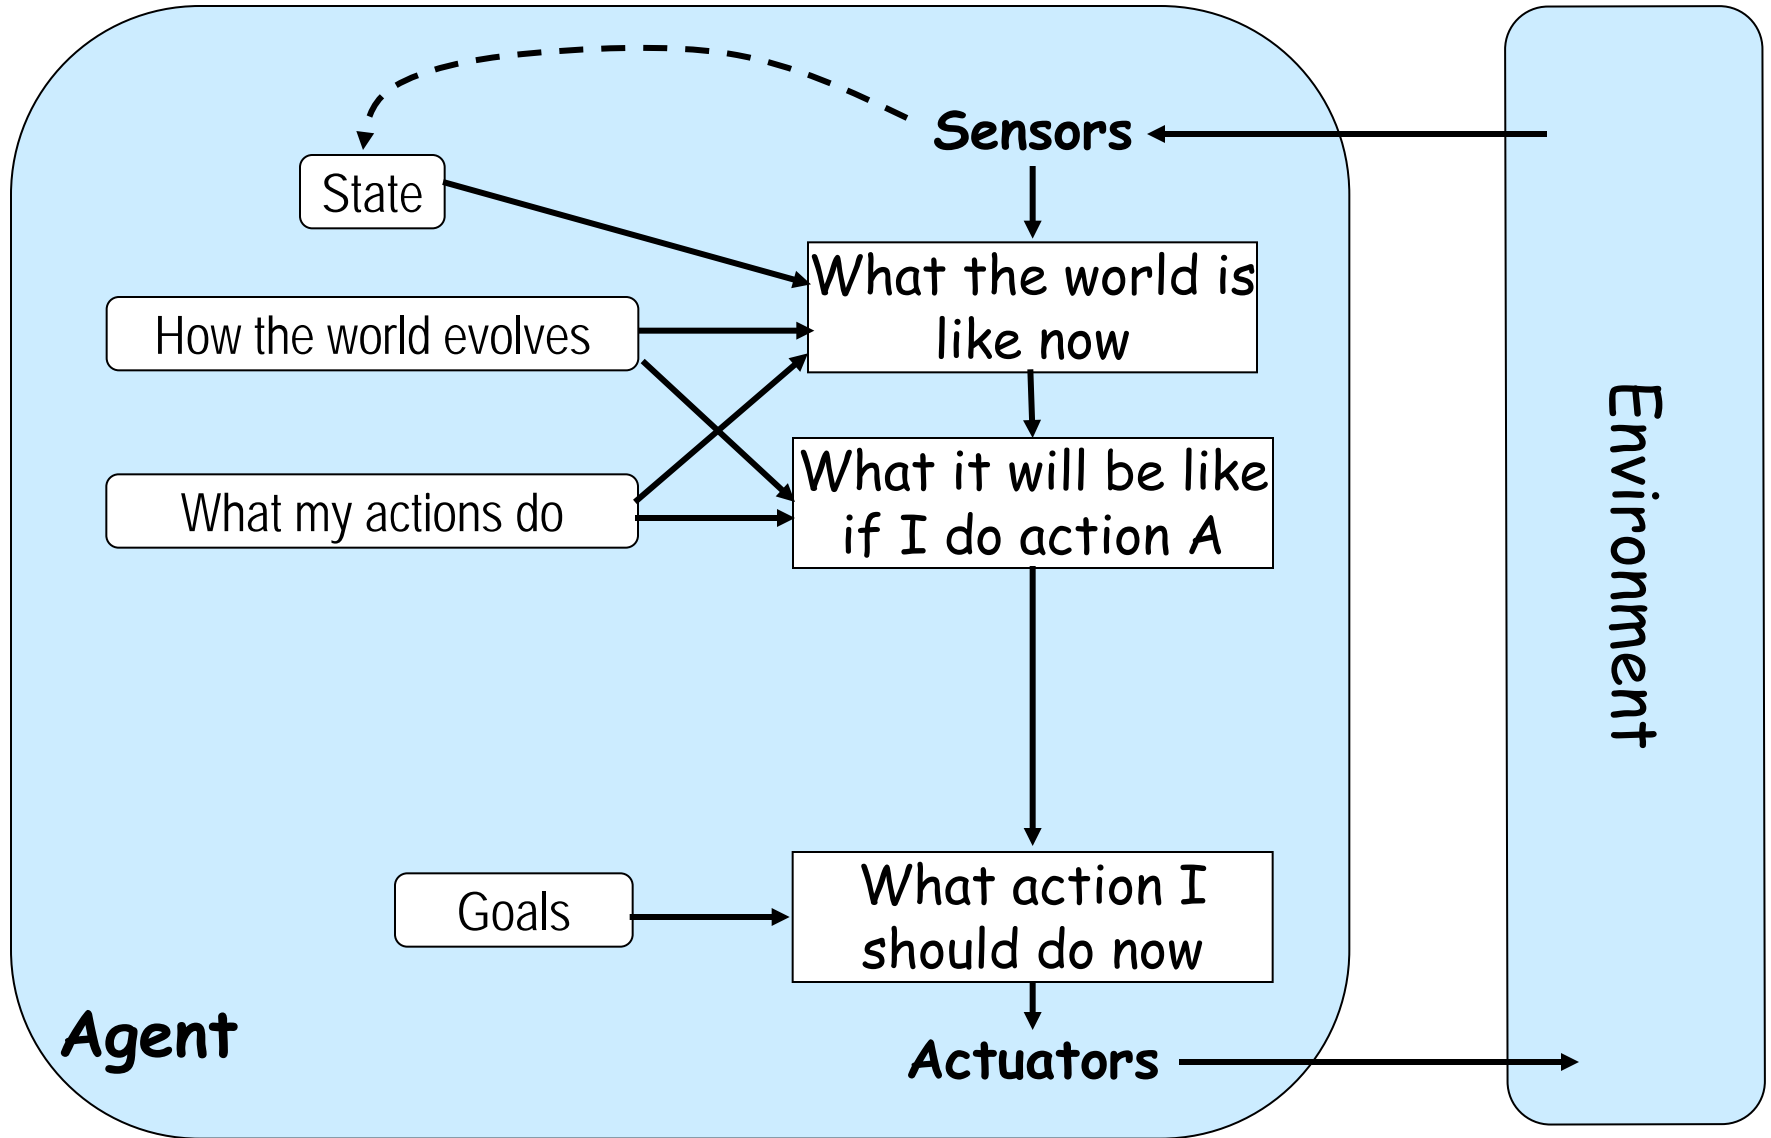
Model-based Reflex Agent (with State)

```
function REFLEX-AGENT-WITH-STATE (percept) returns an
action
  static: state, a description of the current world state
         rules, a set of condition-action rules
         action, the most recent action, initially none

  state ← INTERPRET-INPUT(state, action, percept)
  rule ← RULE-MATCH(state, rules)
  action ← RULE-ACTION[rule]
return action
```

Disadvantage: Knowing about current state is not always enough to decide what to do.

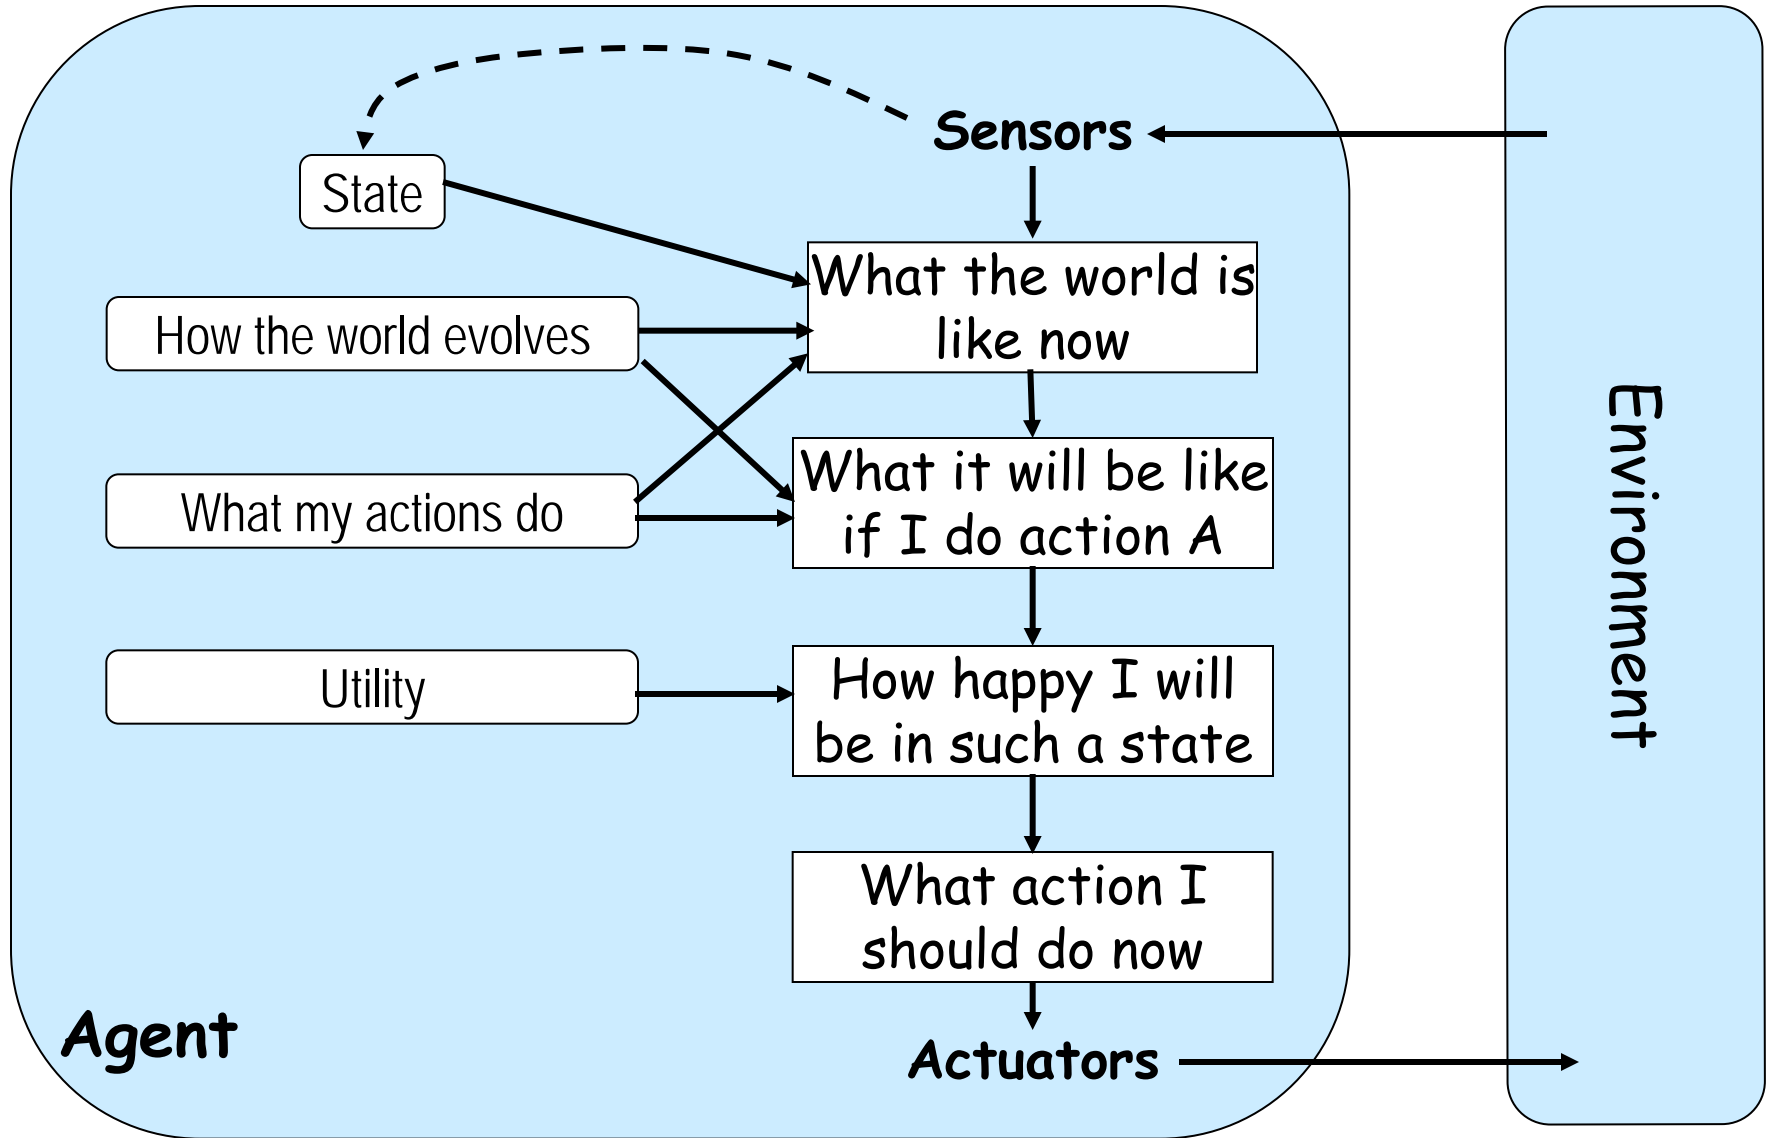
Model-based, Goal-based Agents



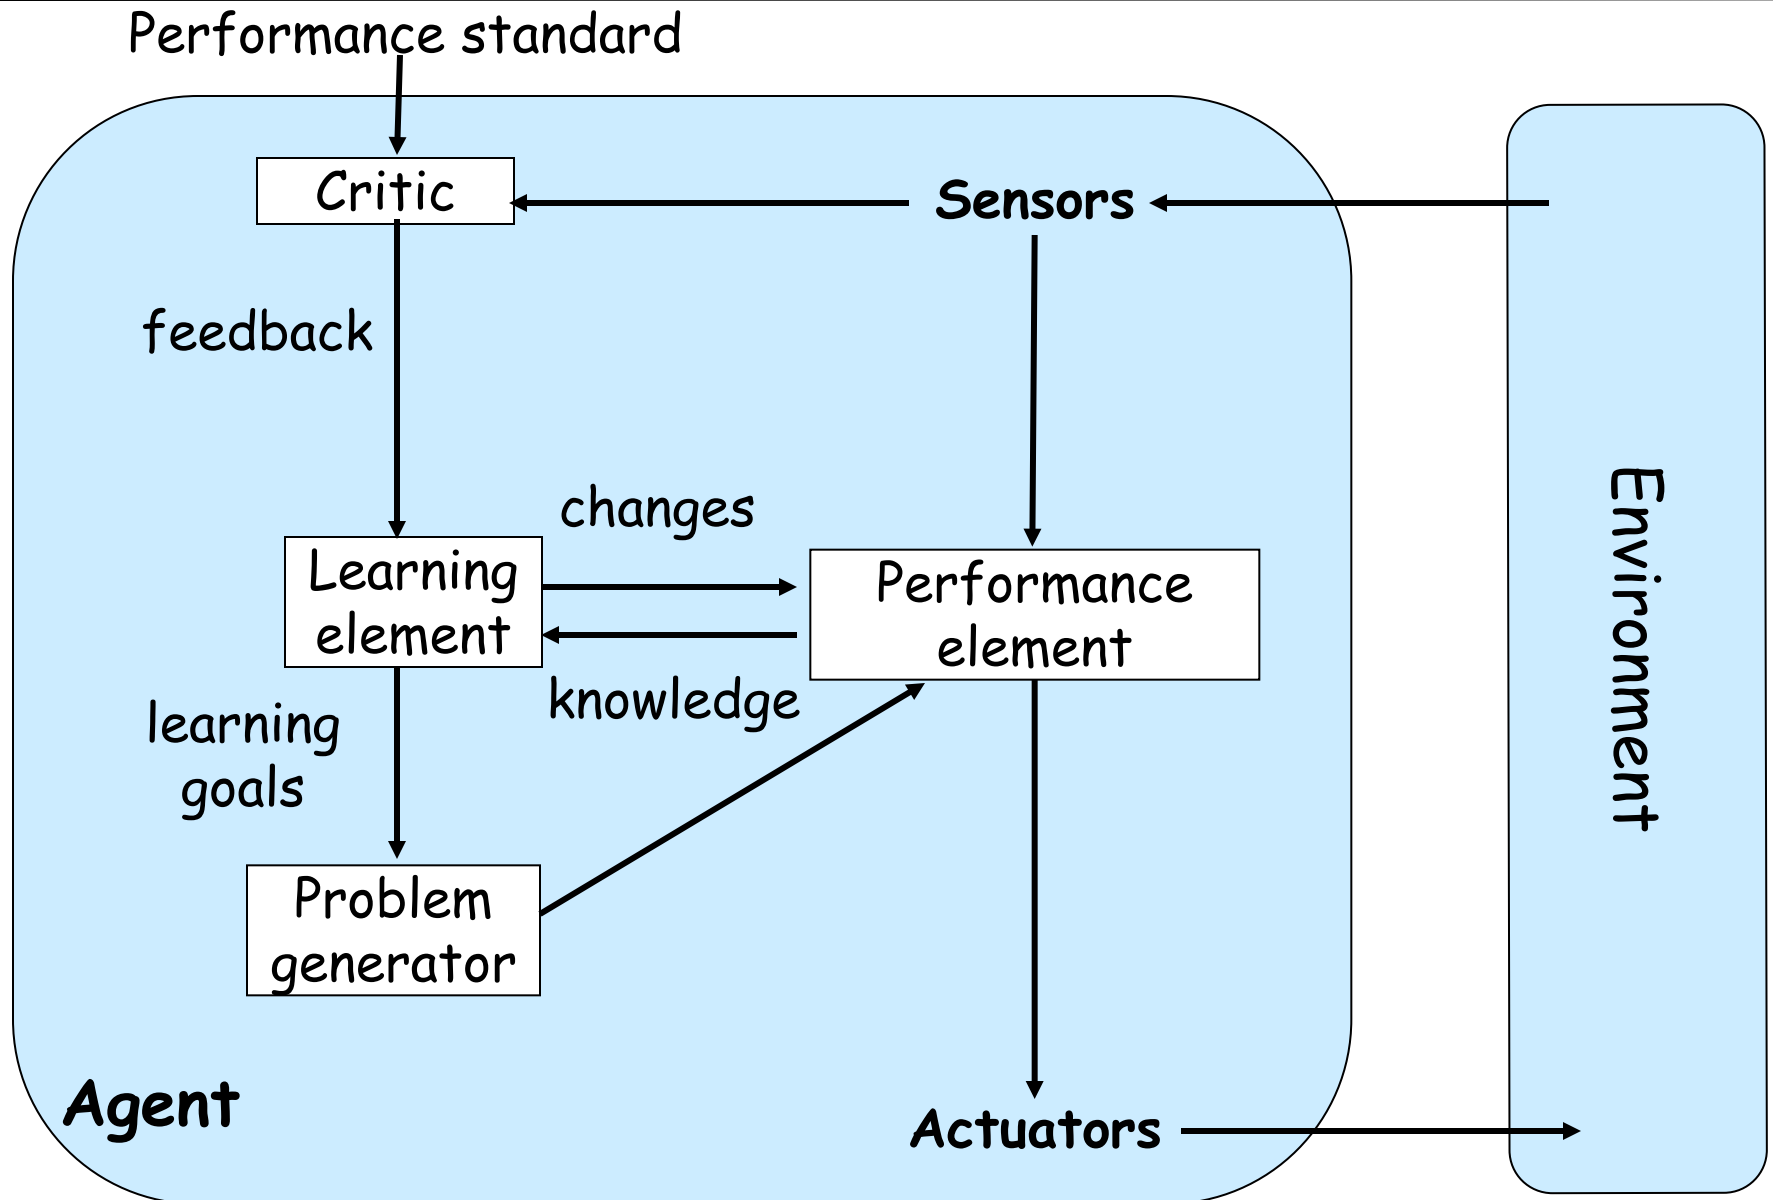
Goal-based Agents

- But, goals alone aren't really enough to generate high-quality behavior in most environments.

Model-based, Utility-based Agents



Learning Agents



Exercise: Back to Vacuum Environment

- Vacuum cleaner agent, modified such that geography is unknown, as is initial dirt configuration. Agent can go *Up, Down, Left, Right*
- Can a simple reflex agent be perfectly rational for this environment?

Exercise: Back to Vacuum Environment (con't.)

- Vacuum cleaner agent, modified such that geography is unknown, as is initial dirt configuration. Agent can go *Up, Down, Left, Right*
- Can a simple reflex agent with a *randomized* agent function outperform a simple reflex agent? What is its design?

Exercise: Back to Vacuum Environment (con't.)

- Describe an environment in which your randomized agent will perform poorly

Assignment for This Week

- Read Chapter 3
 - By Tuesday: Sections 3.1 – 3.4
 - By Thursday: Sections 3.5 – 3.6