

(version without (potentially) copyrighted images)

CS494/594: Artificial Intelligence

Tuesday, August 25th

“AI is an engineering discipline built on an unfinished science.”

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

“Thought Discussion” State of the art

Which of the following can be done at present?

- Play a decent game of table tennis
 - Yes. A reasonable level of proficiency was achieved by Andersson’s robot, 1988
- Drive along a curving mountain road
 - Pretty close, at least for relatively empty roads. Some significant progress has been made in following roads, even dirt roads, when there are no other vehicles in the area to worry about. The DARPA contest (“DARPA Grand Challenge”) showed the ability of robots to follow unimproved roads in difficult terrains (132 miles!)
- Drive in the center of Cairo
 - Not really. Although there has been a lot of progress in automated driving, and the DARPA Urban Challenge has addressed some issues, all such systems currently rely on certain relatively constant clues: that the road has shoulders and a center line, that the car ahead will travel a predictable course, that cars will keep to their side of the road (most of the time), etc. Driving in downtown Cairo is still too unpredictable for any of these to work, although progress is being made.

“Thought Discussion” State of the art

- Buy a week’s worth of groceries at Kroger
 - No. No robot can currently put together the tasks of moving in a crowded environment, using vision to identify a wide variety of objects, and grasping the objects (including squishable vegetables) without damaging them. The component pieces are nearly able to handle the individual tasks, but it would take a major integration effort to put it all together.
- Buy a week’s worth of groceries on the web
 - Yes. Software robots are capable of handling such tasks, particularly if the design of the web grocery shopping site does not change radically over time.
- Play a decent game of bridge
 - Yes. Programs such as GIB now play at a solid level.
- Discover and prove a new mathematical theorem
 - Yes. For example, the proof of Robbins algebra described in our text on page 309.
- Write an intentionally funny story
 - No. While some computer-generated prose and poetry is hysterically funny, this is invariably unintentional, except in the case of programs that echo back prose that they have memorized.

“Thought Discussion” State of the art (con’t.)

- Give a competent legal advice in a specialized area of law
 - Yes, in some cases. AI has a long history of research into applications of automated legal reasoning. Two outstanding examples are the Prolog-based expert systems used in the UK to guide members of the public in dealing with the intricacies of the social security and nationality laws. The social security system is said to have saved the UK government approximately \$150 million in its first year of operation. However, extension into more complex areas such as contract law awaits a satisfactory encoding of the vast web of common-sense knowledge pertaining to commercial transactions and agreement and business practices.
- Translate spoken English into spoken Swedish in real time
 - Yes. In a limited way, this is already being done. See Kay, Gawron, and Novig (1994) and Wahlster (2000) for an overview of the field of speech translation, and some limitations on the current state of the art.
- Perform a complex surgical operation
 - Yes. Robots are increasingly being used for surgery, although always under the command of a doctor.

Next time...

- “Thought Discussion” for next time:

Consider: “Intelligence is in the eye of the beholder.”

Do you agree? Disagree?

Why/why not?

What are the implications to either position?

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

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 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?				

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

Environment Types

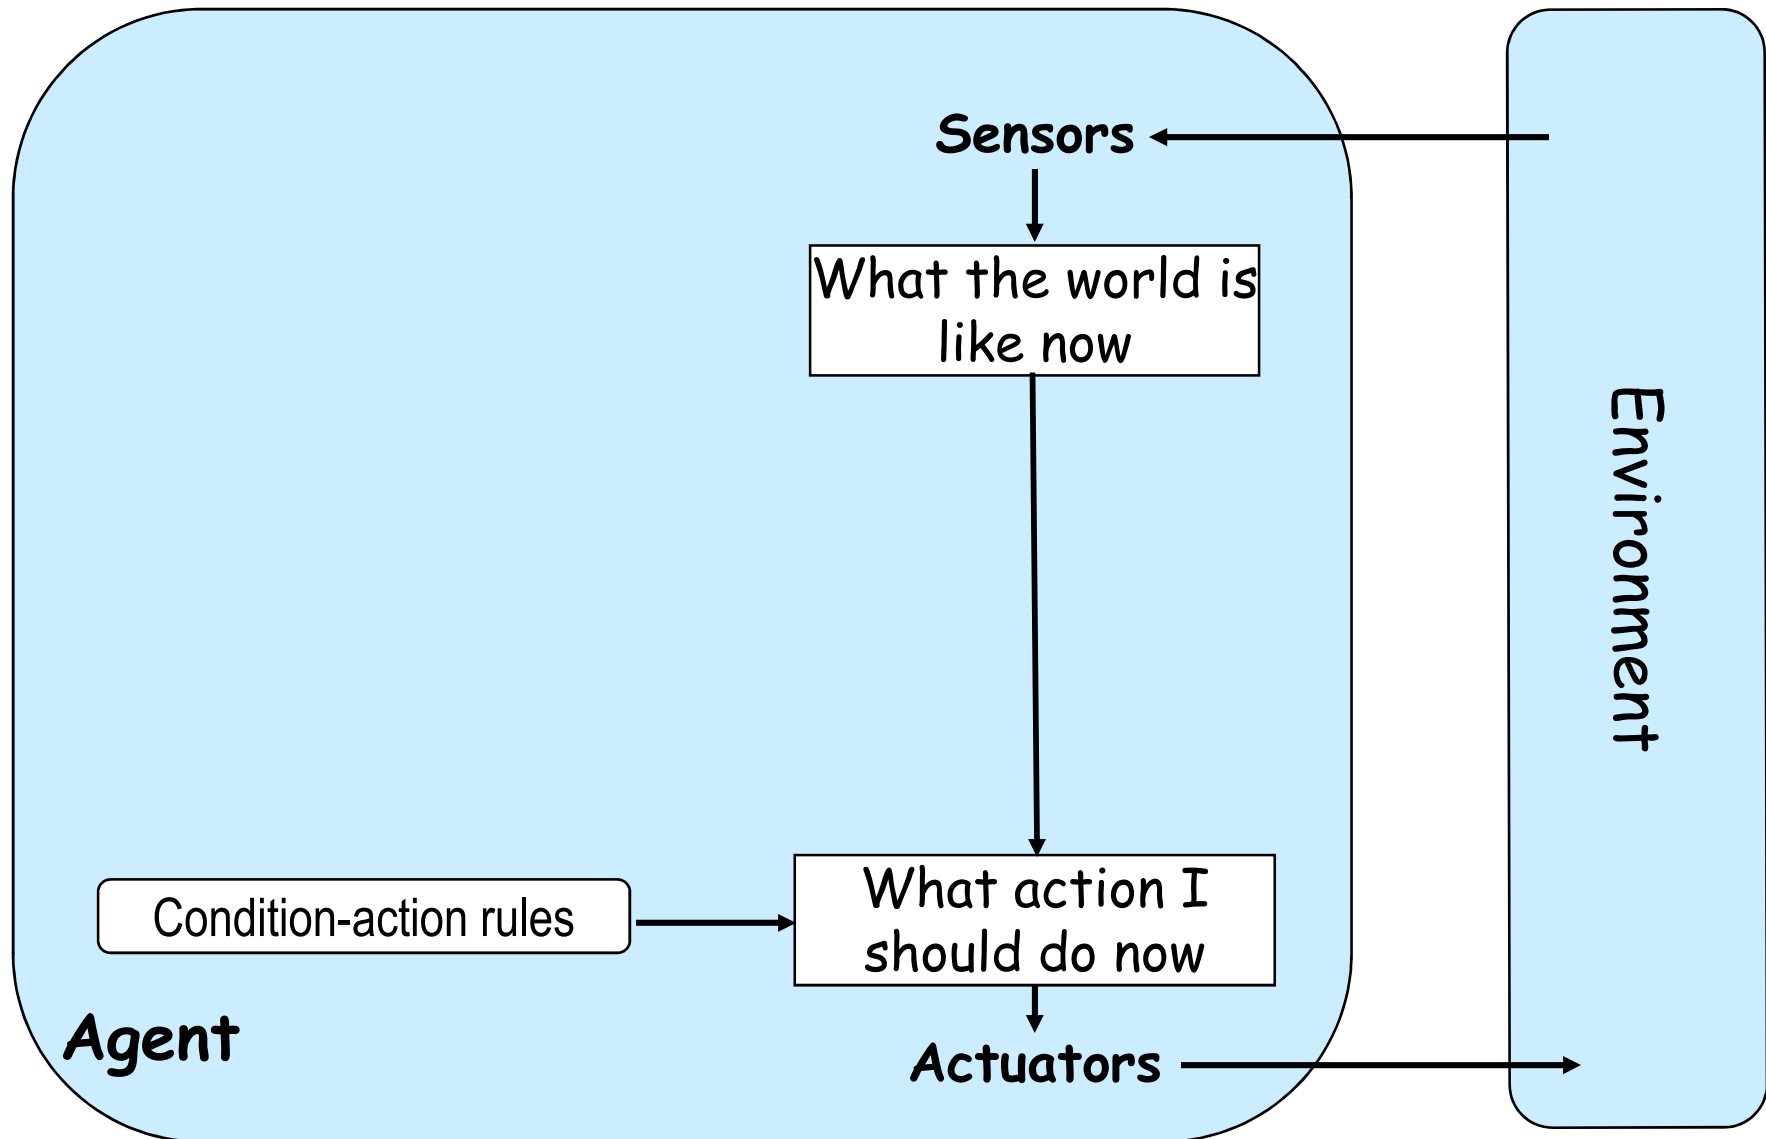
	Solitaire	Backgammon	Internet shopping	Taxi
Observable ??	Yes	Yes	No	No
Deterministic ??	Yes	No	Partly	No
Episodic ??	No	No	No	No
Static ??	Yes	Semi	Semi	No
Discrete ??	Yes	Yes	Yes	No
Single-agent ??	Yes	No	No (except auctions)	No

- 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

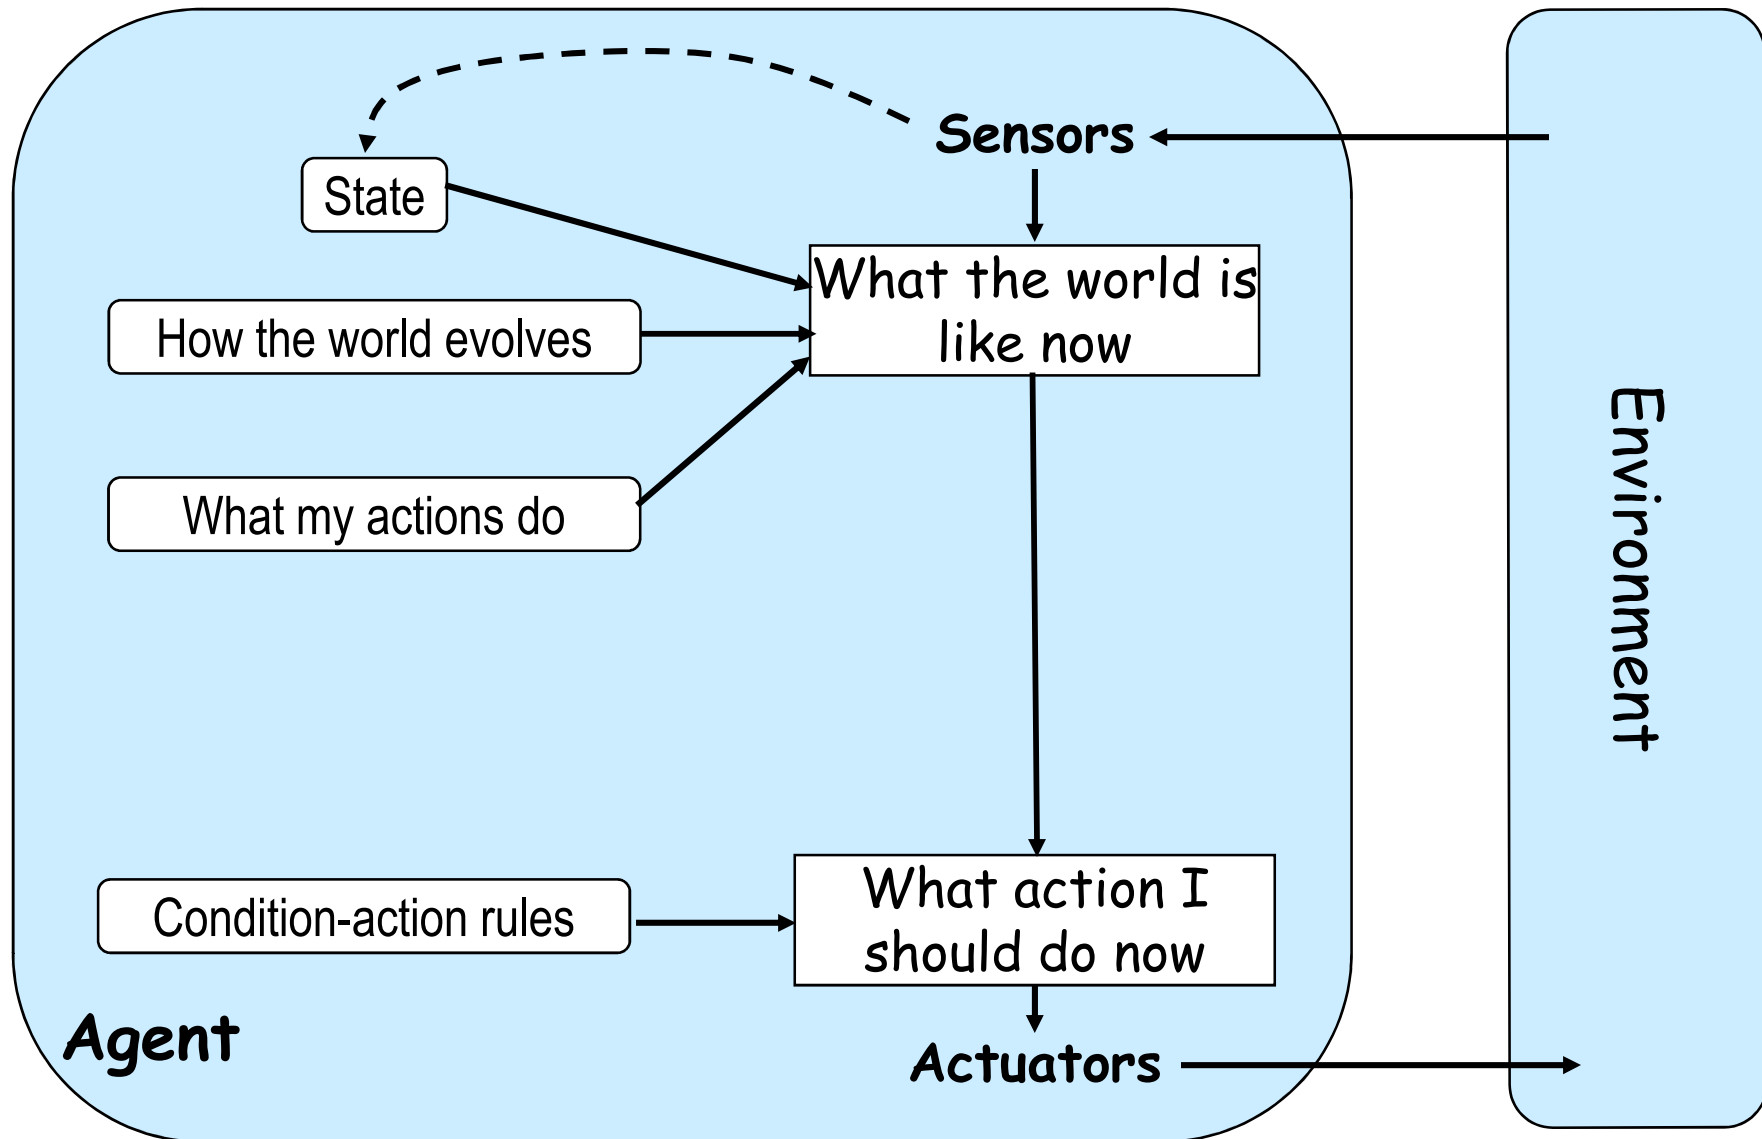
```
function SIMPLE-REFLEX-AGENT(percept) returns an action
  static: rules, a set of condition-action rules

  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 *note*

state ← INTERPRET-INPUT(*state*, *action*, *percept*)

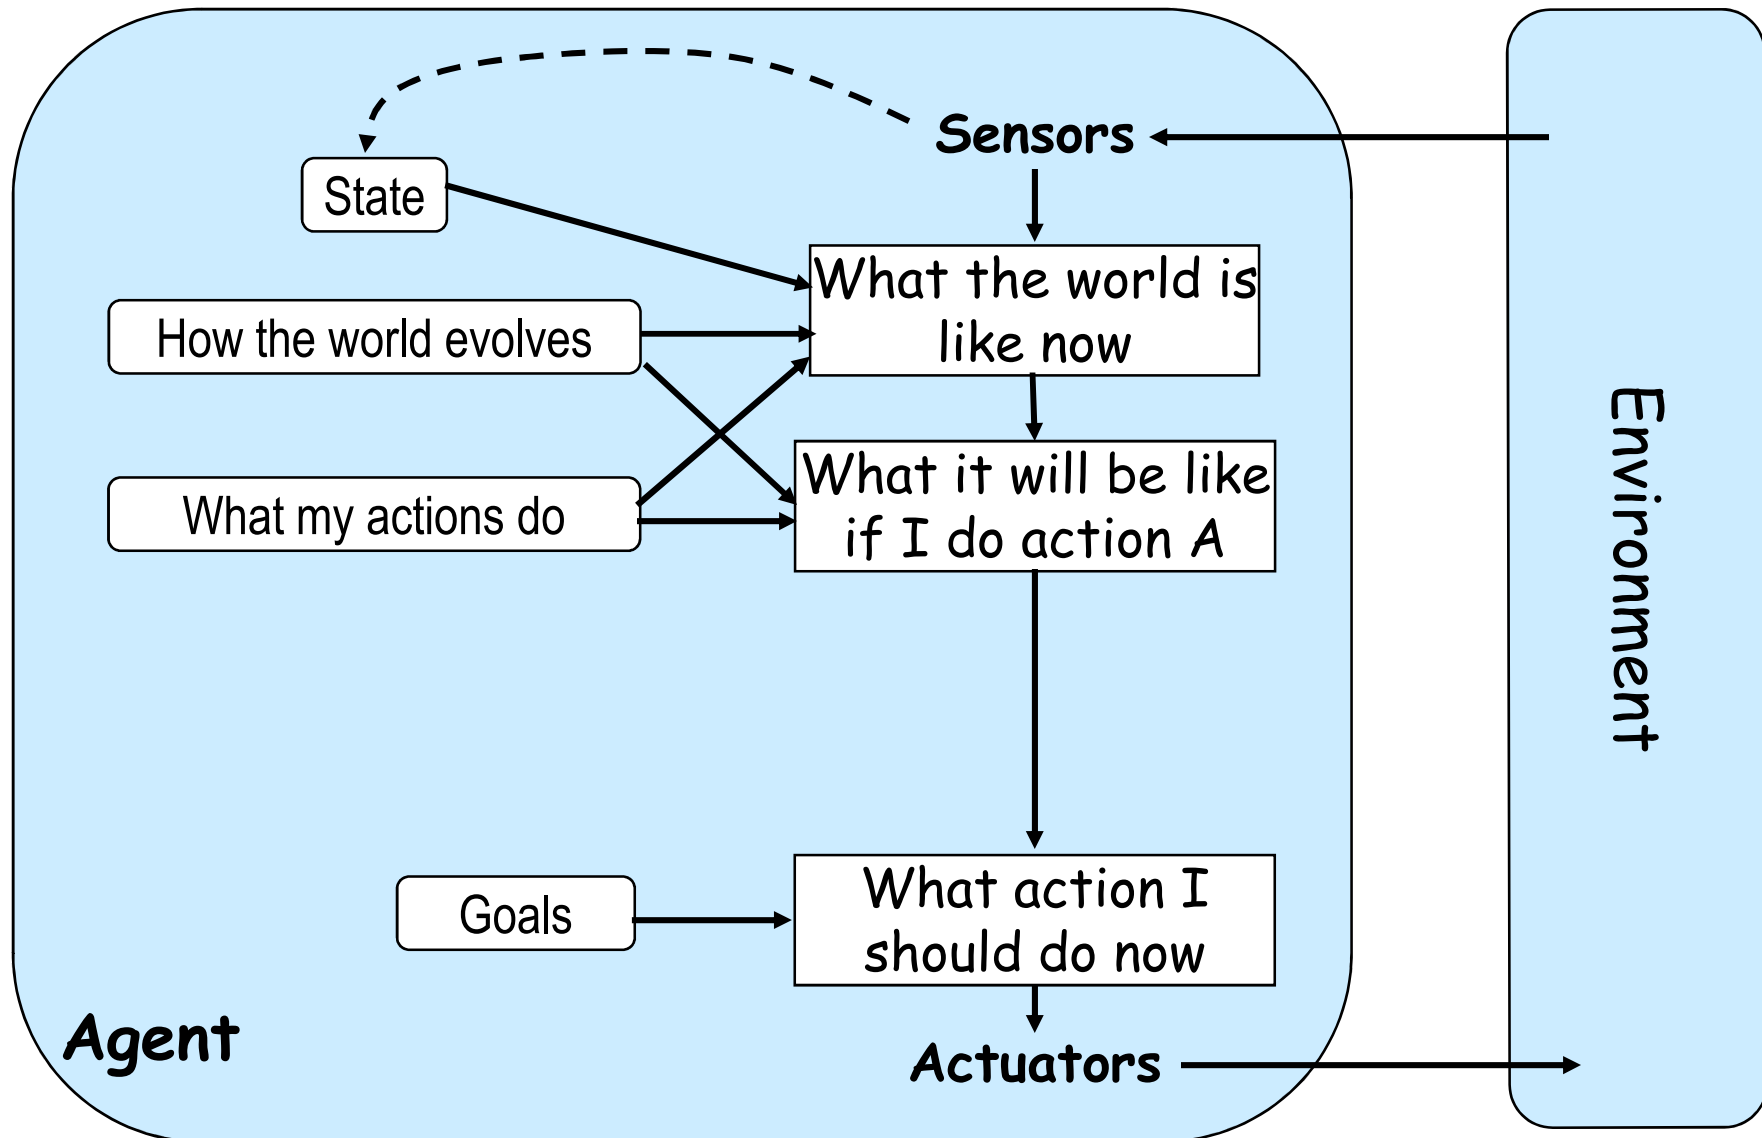
rule ← RULE-MATCH(*state*, *rules*)

action ← RULE-ACTION[*rule*]

return *action*

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

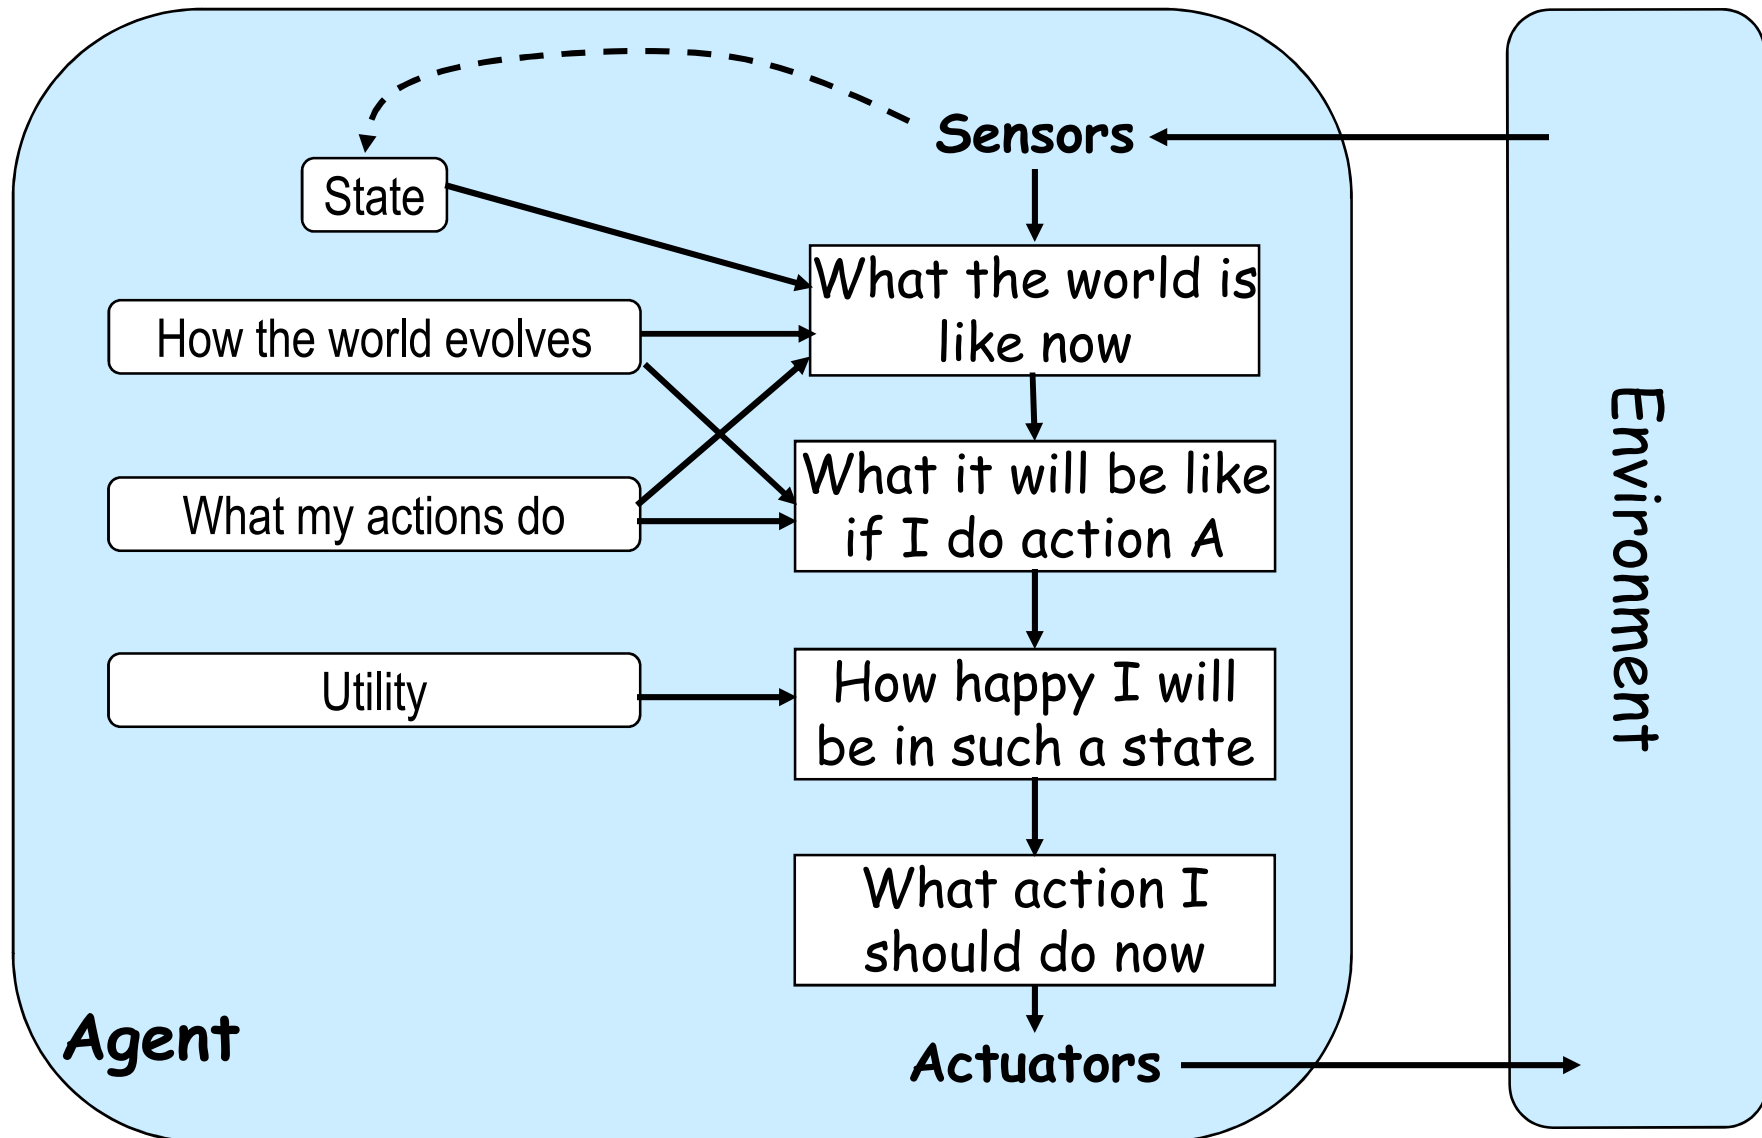
Model-based, Goal-based Agents



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

