CMPSCI 383: Artificial Intelligence
Lecture 2, September 8, 2011

Philip Thomas (TA)

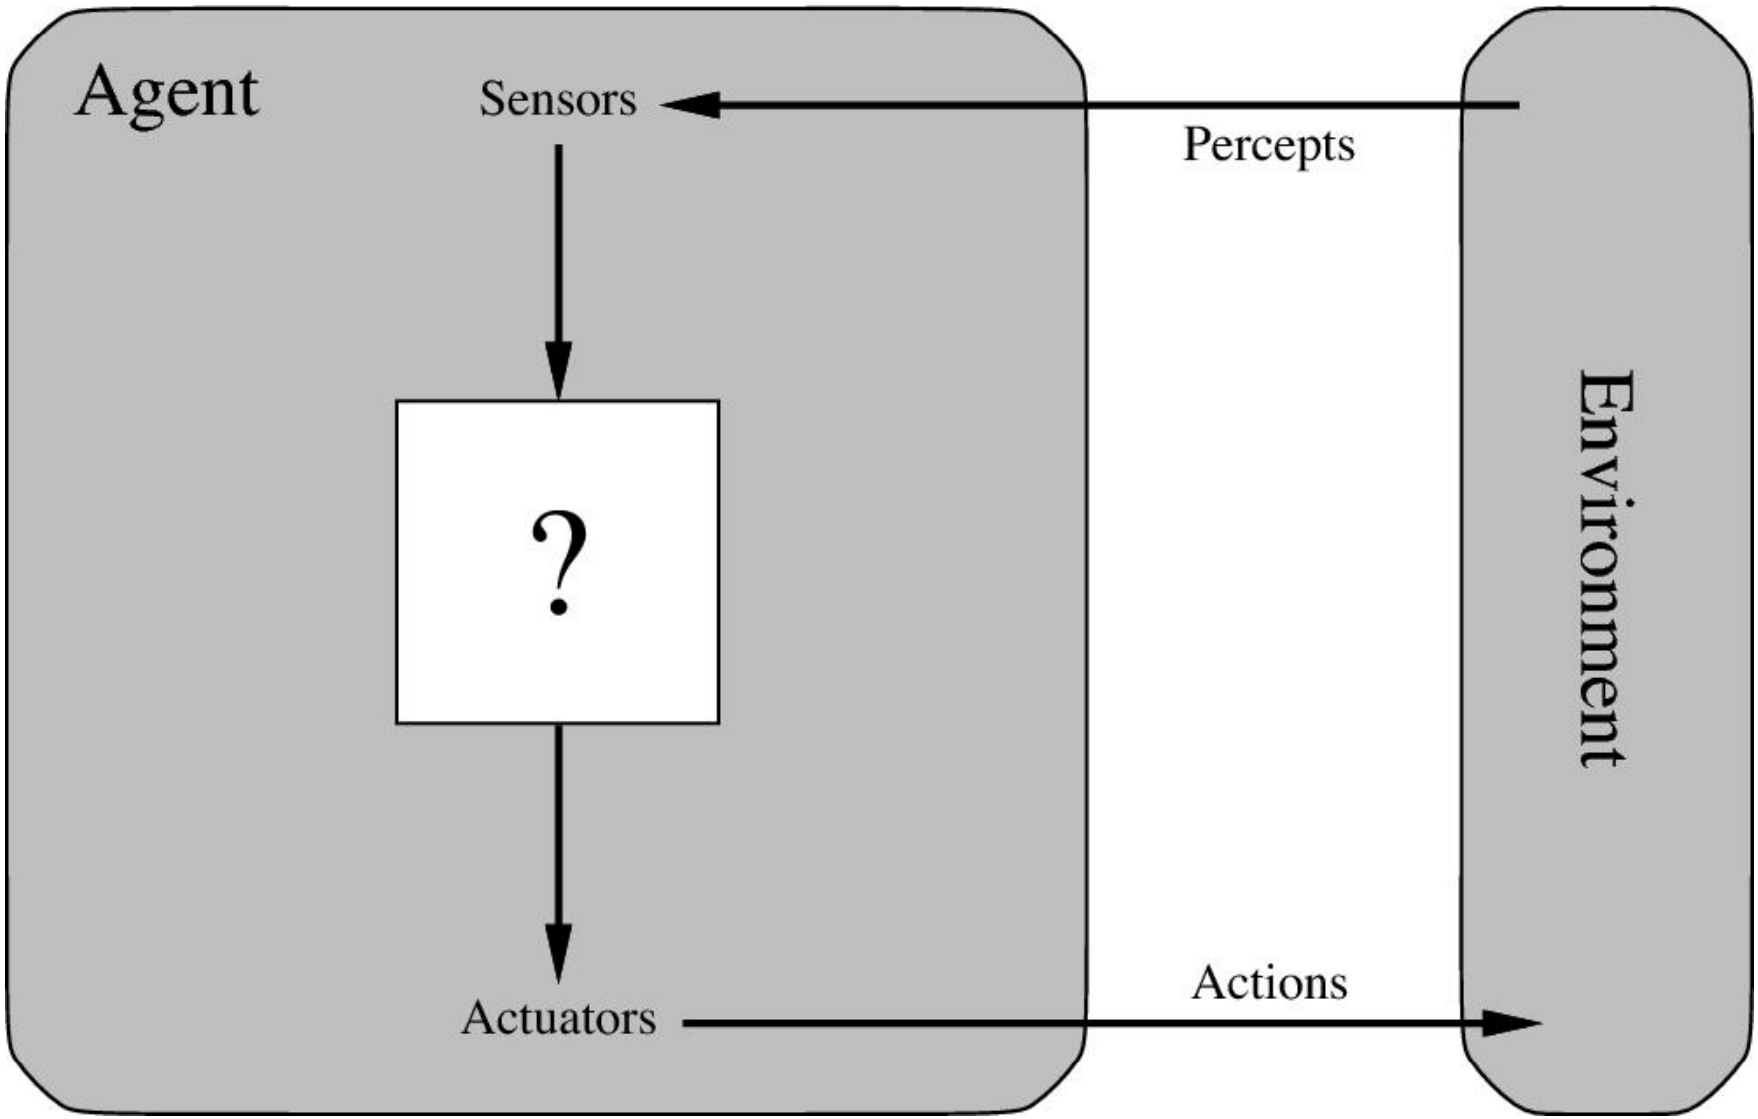
Agent
Sensors
Percepts
Environment
?
Actuators
Actions

Agent Function vs. Agent Program

- Agent function: Maps any percept sequence to an action
- Agent program: Internal implementation of the agent function
- Tabular?

Vacuum-Cleaner World

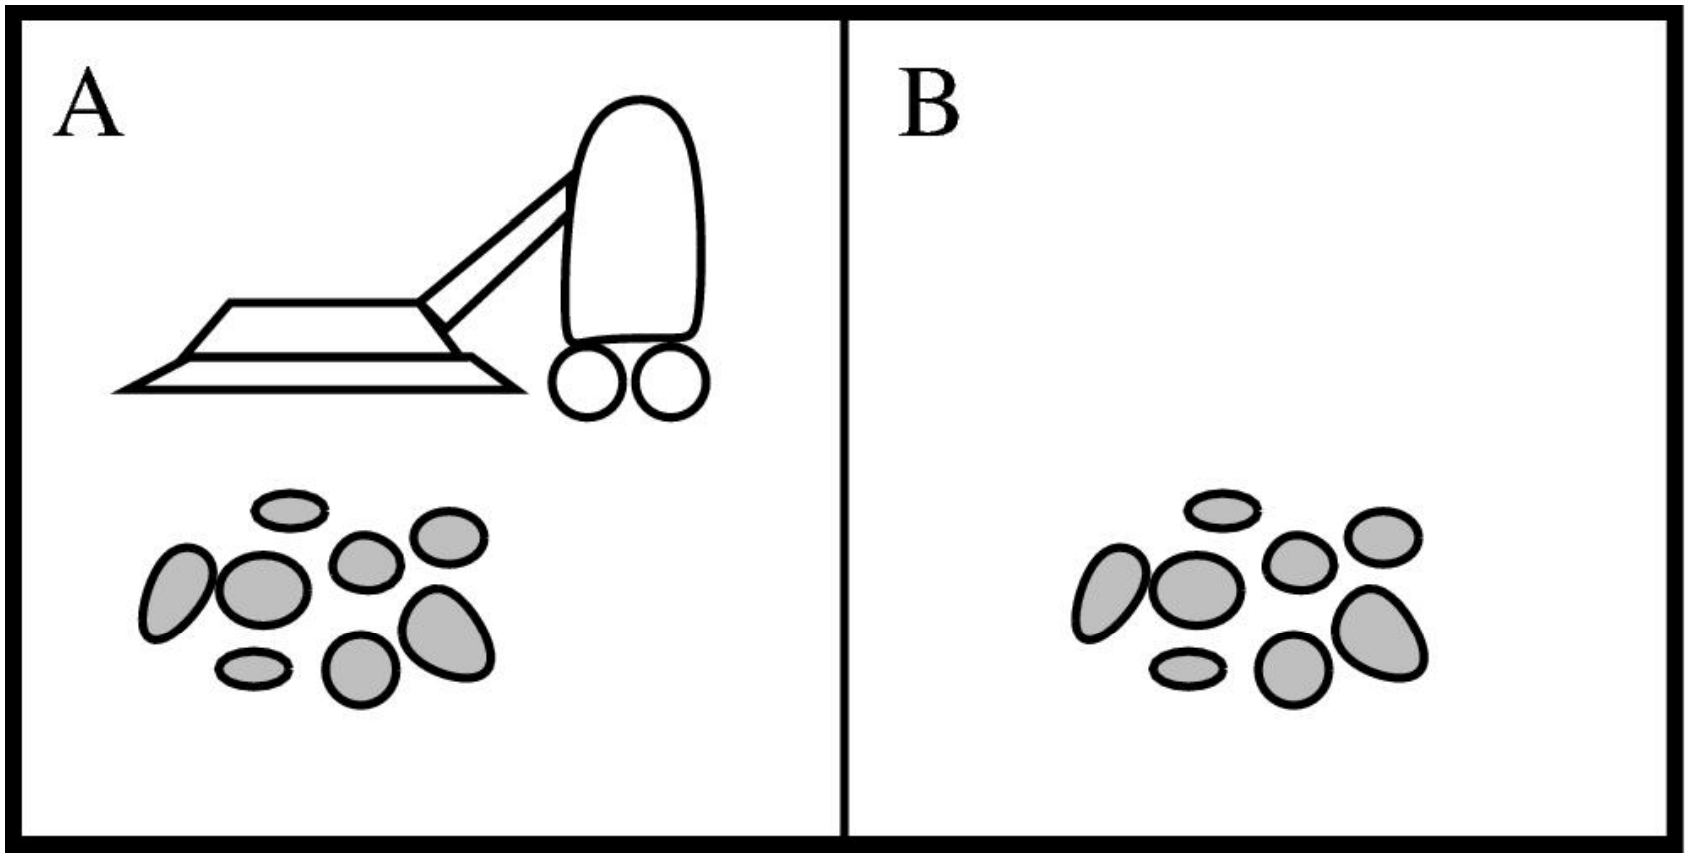

Tabular Agent Function

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
…	…
[A, Clean], [A, Clean], [A, Clean]	Right
…	…

Tabular Agent Program?

- Huge, or infinite!
- How big would it be for chess?
  - Estimated number of possible percept histories is 10^(10^50)

    Hardy, G. H. Ramanujan: Twelve Lectures on Subjects Suggested by His Life and Work, 3rd ed. New York: Chelsea, 1999.
- Can be represented more succinctly:
  - If dirty, suck, otherwise move to other location

Performance Measure

- “Pick up as much dirt as possible”
  - What might a rational agent do?
- “Keep the floor as clean as possible”

- Human arm movements
  - “Minimize error in hand location”
  - “Minimize error in joint angles”
  - “Minimize error in ^^^^^^^^, with penalty for using a lot of force”
  - “Maximize the amount of time spent moving toward the goal, while minimizing the time spent moving away”

Examples

- Give me examples of
  - Agent
    - Sensors and Percepts
    - Actuators and Actions
  - Environment
  - Performance measure

Rational Agent

- 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 whatever built-in knowledge the agent has.

Rationality vs. Perfection

- Crossing the street example
- Rationality: Maximize expected performance
- Perfection: Maximize actual performance

- Is it rational to cross a road without looking both ways?
- Information gathering / exploration

Black Sheep Wall

Chess	Poker
Checkers	Blackjack
Tag	Marco Polo

Fully Observable

Partially Observable

Other Examples?

Piano	Poker
Running	Tag
Balancing	Marco Polo
Sudoku	Chess

Single Agent

Multi Agent

- Is driving single or multi-agent?

Driving (street)	Poker
Taxiing in Airplane	Chess
Ninja Turtles arcade game	Missile Guidance
Fire Truck Allocation	Racing (car)

Cooperative

Competitive

Chess	Fire Truck Allocation
Robot Control (ideally)	Robot Control (in reality)
Vending Machine	Monopoly
Poker	D&D

Deterministic Environment

Stochastic Environment

- Stochastic = Random

Exploration vs. Exploitation

- You're new to a town, and you find a way to work, it takes 10 minutes.
- You have 5 minutes before you have to be at work for a meeting.
- Do you
  - 1) Take the route you know? (exploit)
  - 2) Try to take a shortcut that could waste time? (explore)

Episodic vs. Sequential

- Episodic: Experience divided into atomic episodes. In each episode the agent receives a percept and then performs a single action.
  - The next episode does not depend on the actions taken in previous episodes.
- Sequential: Current decision influences future decisions.

Chess	Driving
Poker	Robot Control
Games using the direction pad	Games using the joystick

Discrete

Continuous

- State, time, action

Reflex Agent

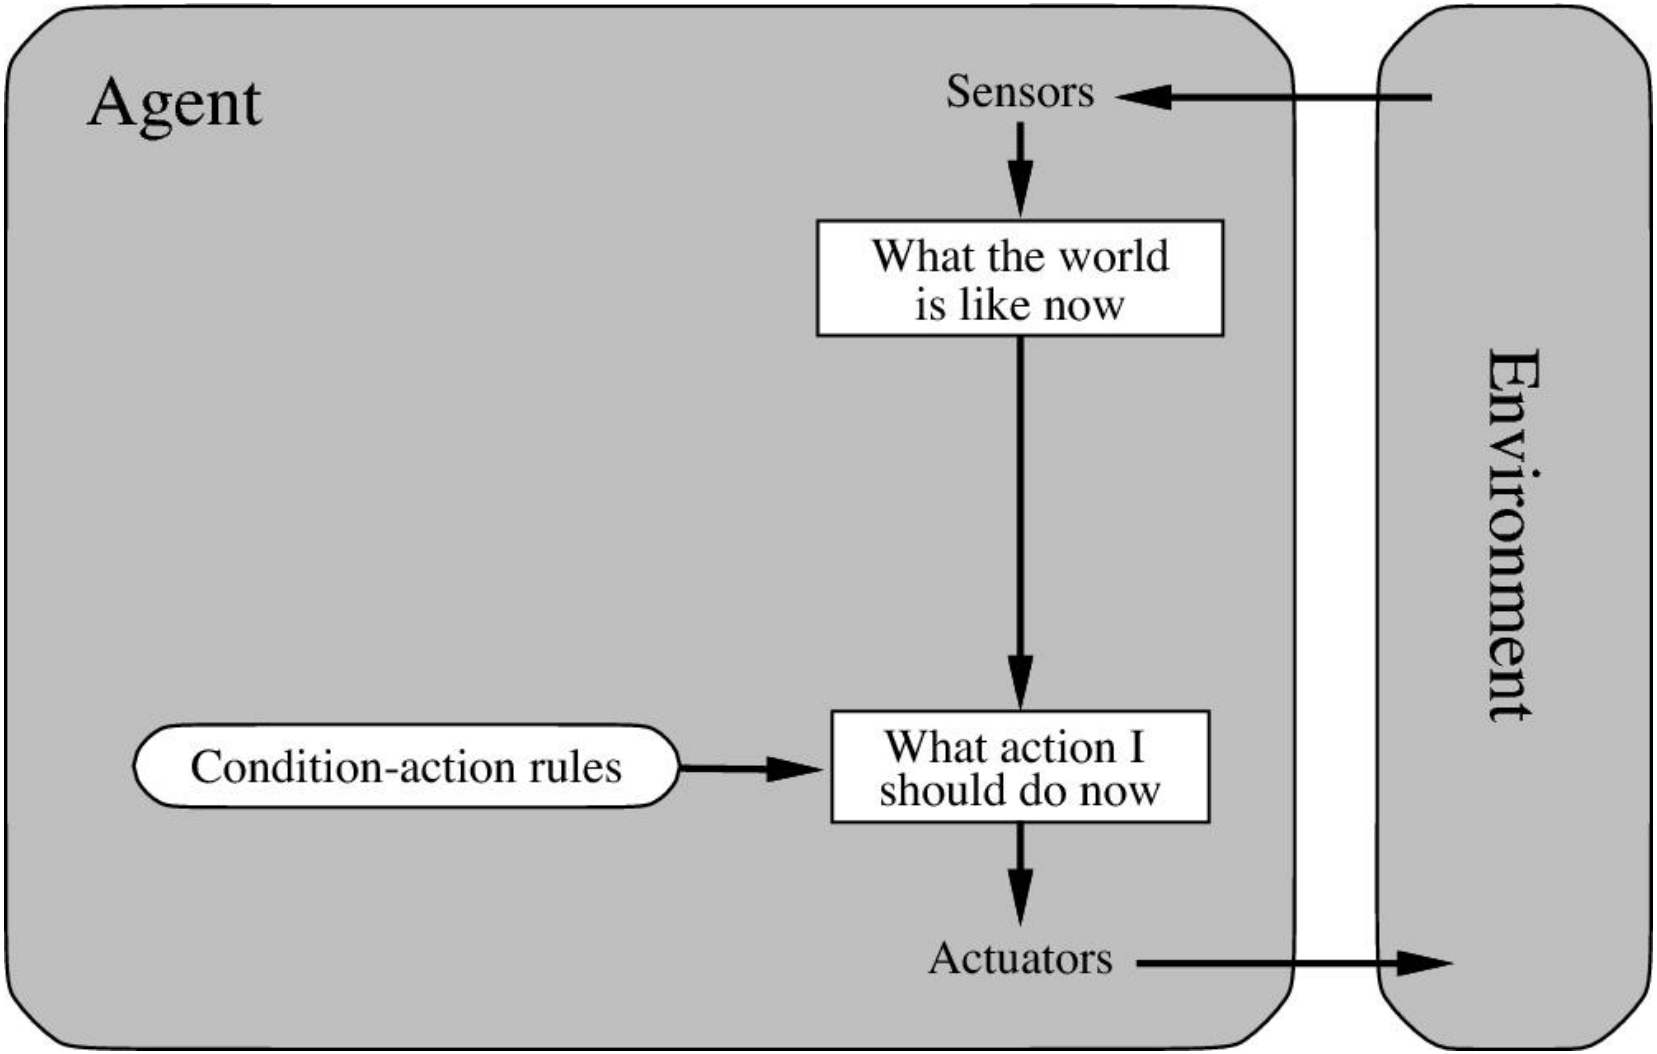

Model-Based Reflex Agent

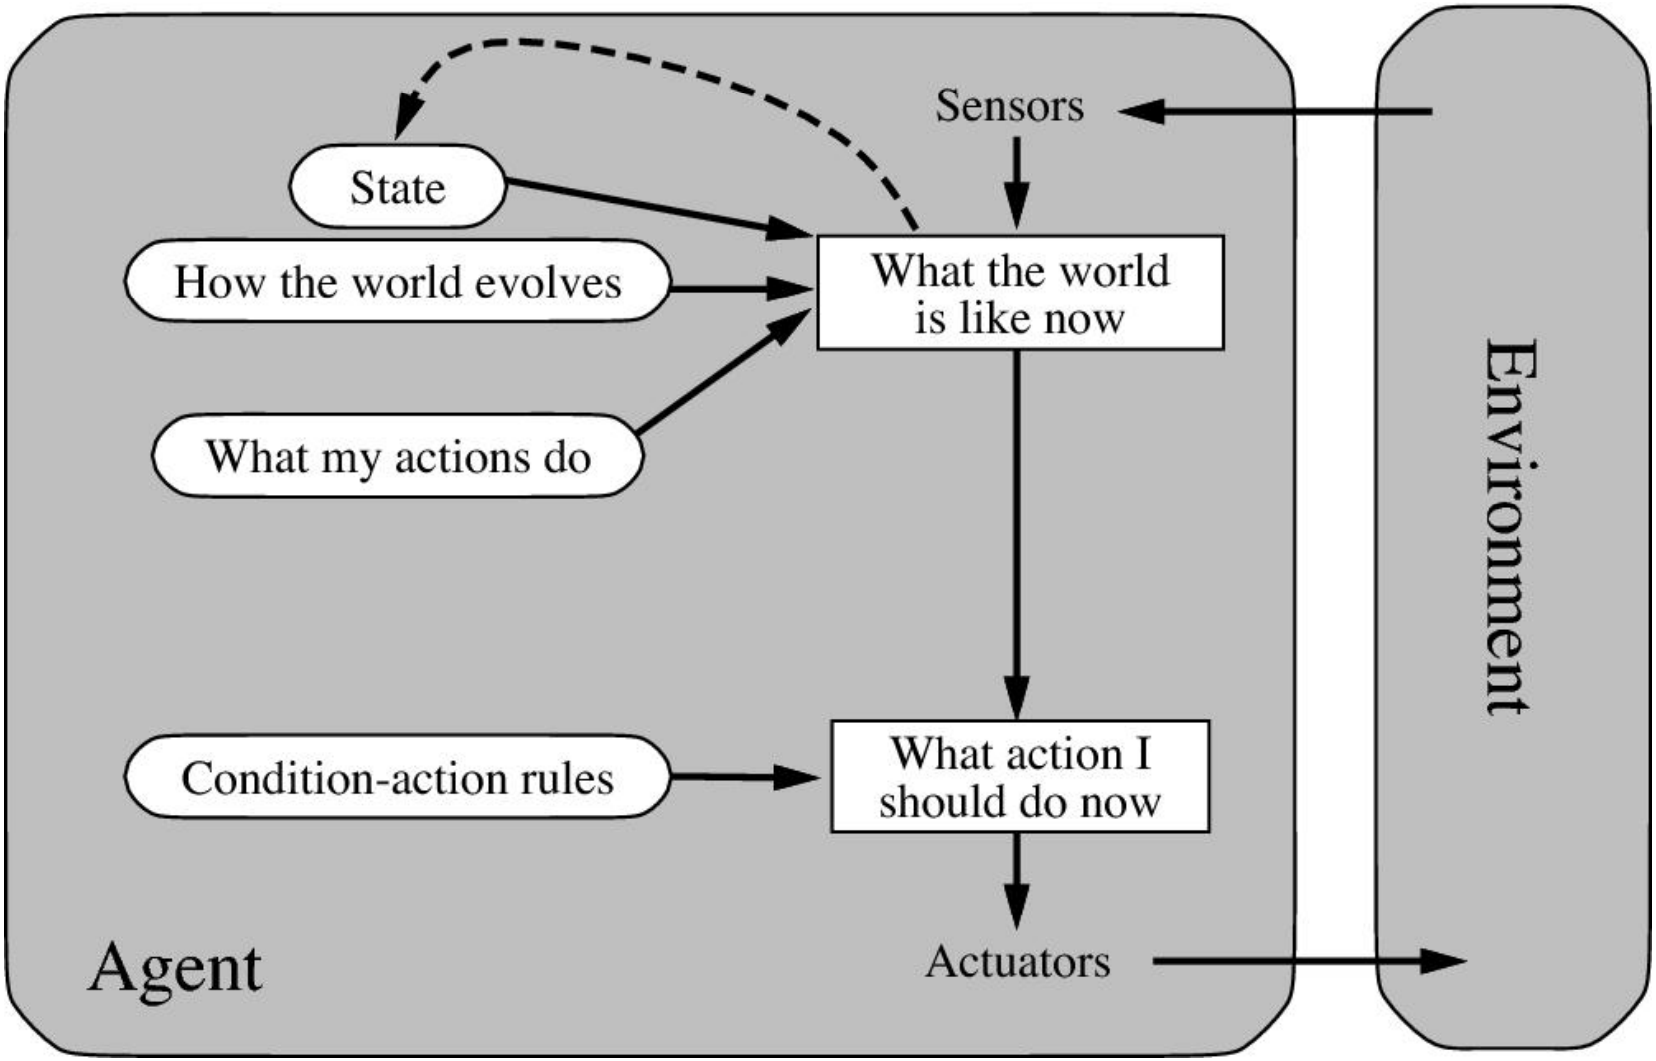

Model-Based Goal-Based

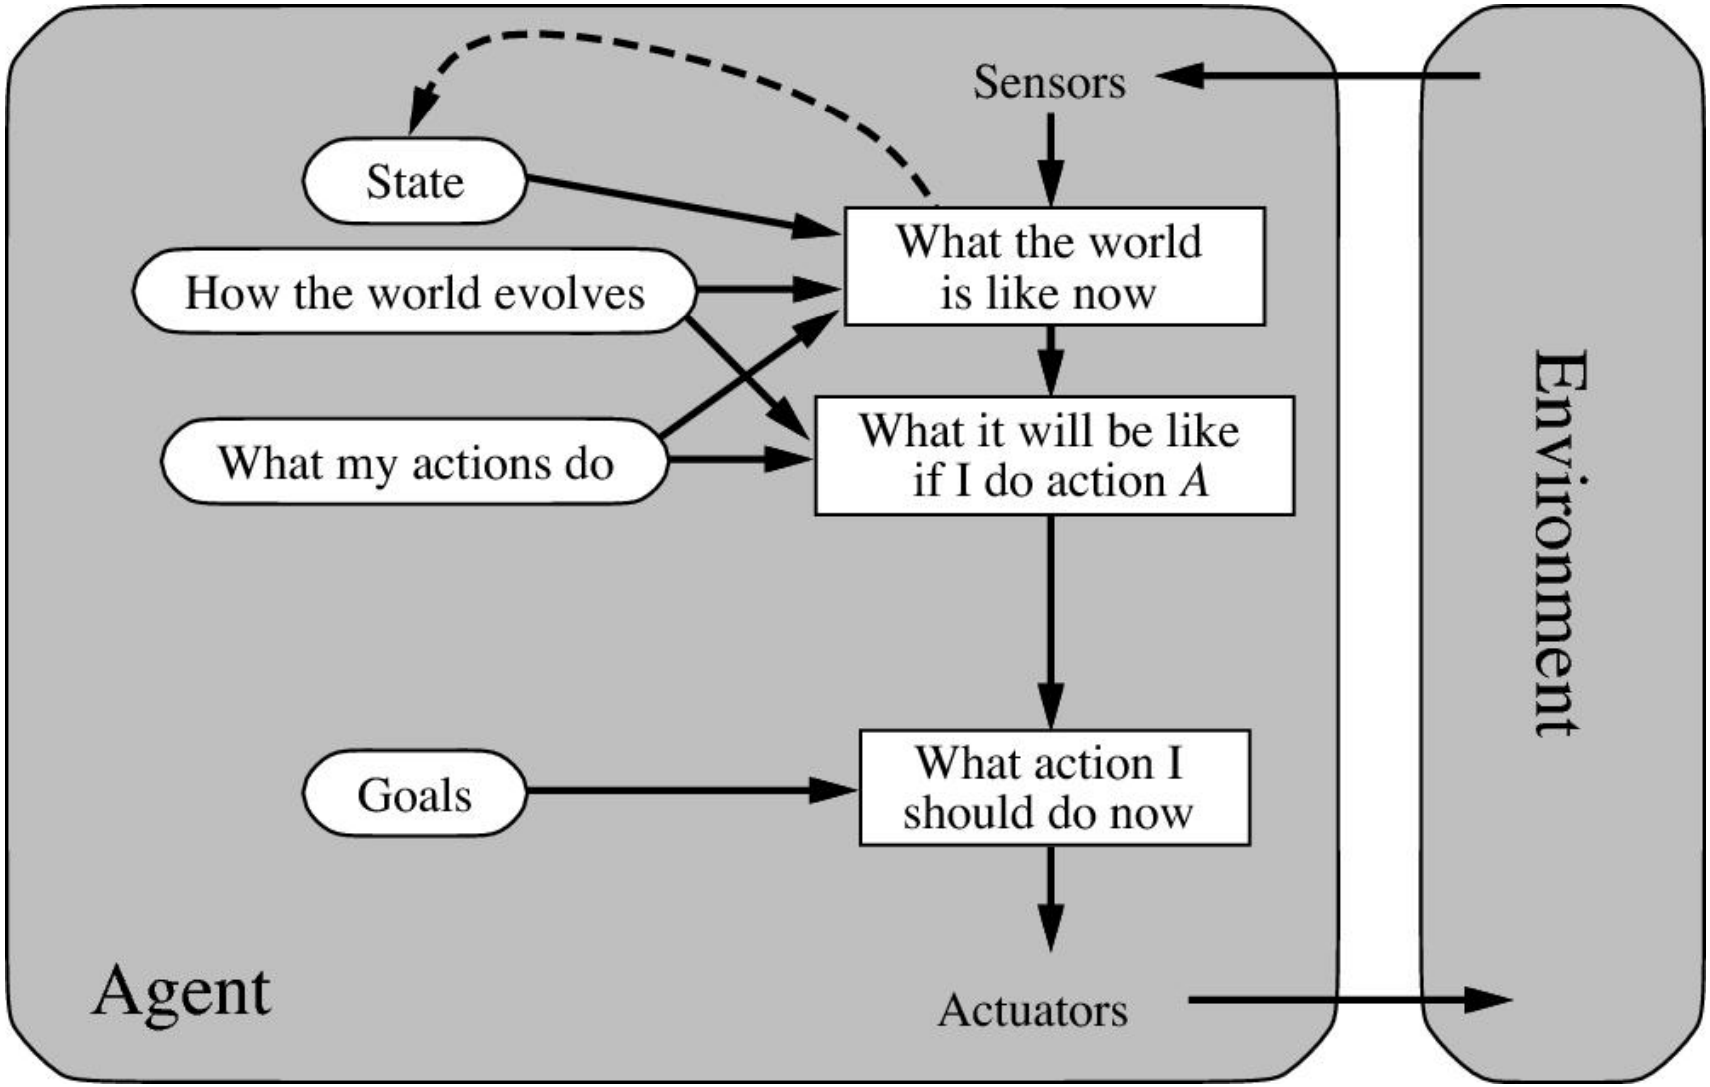

Model-Based Utility-Based

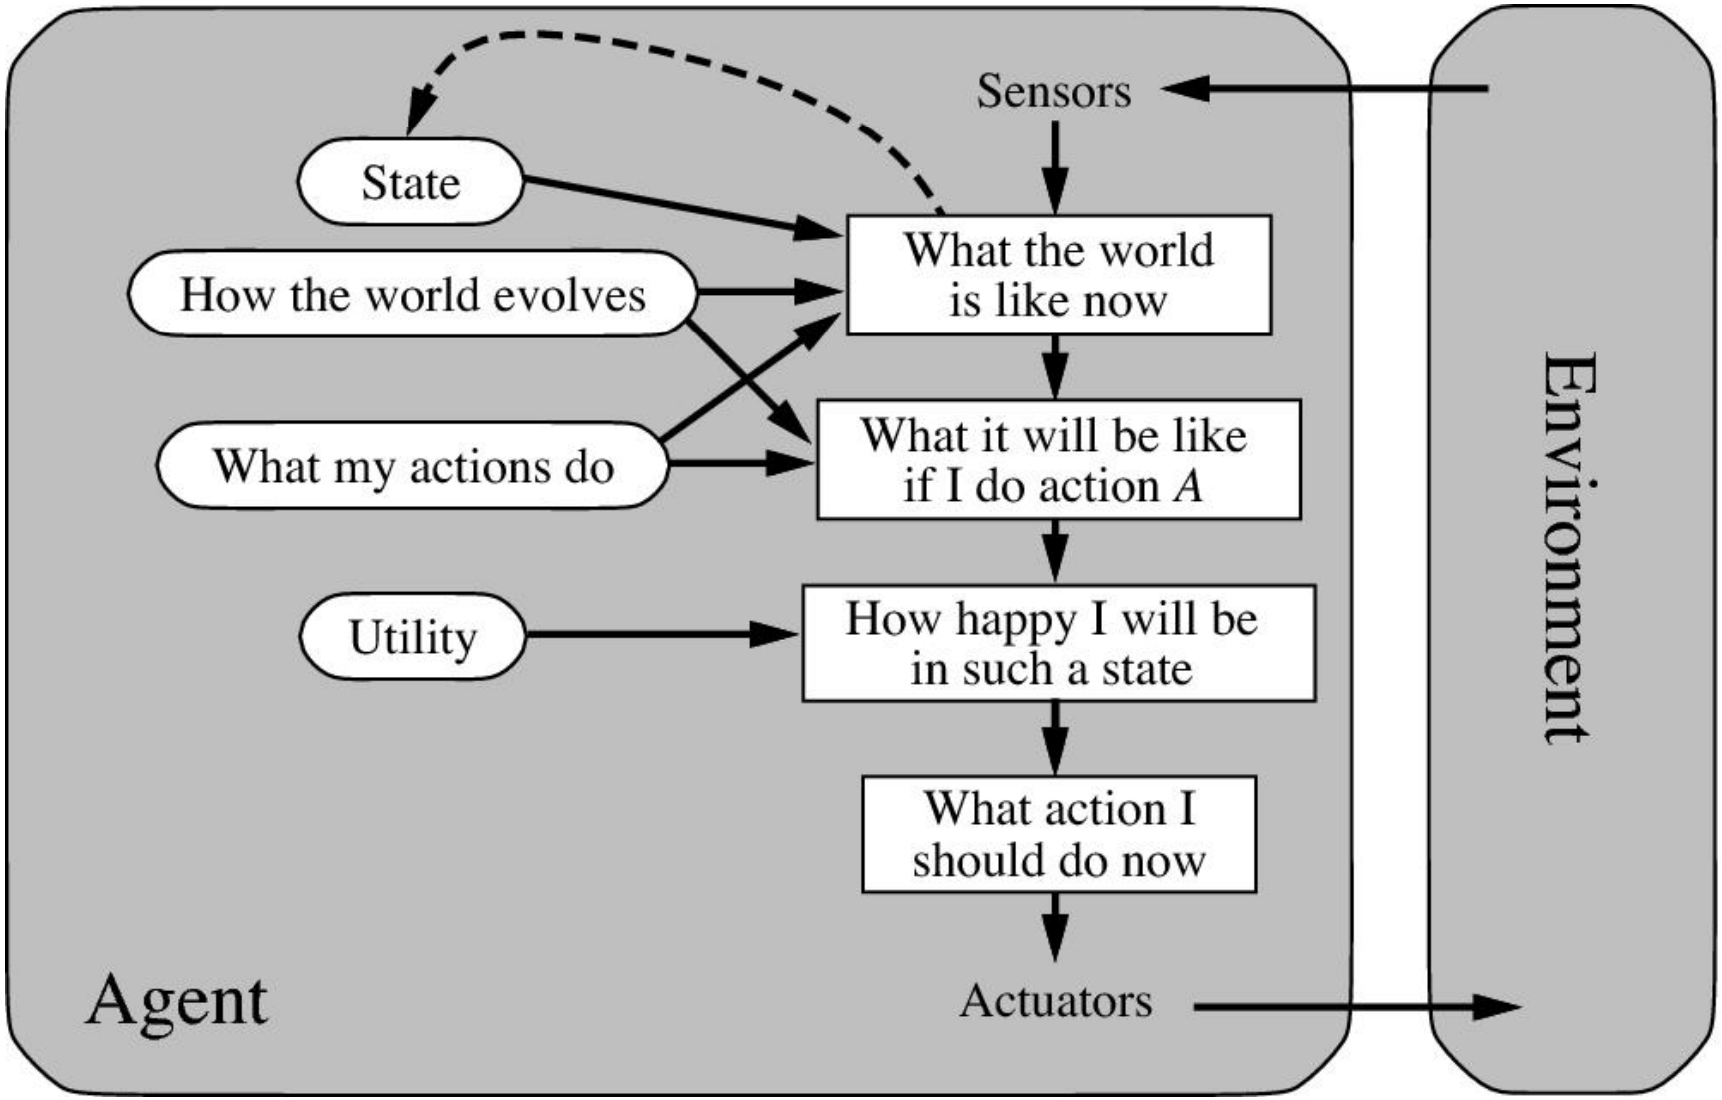

Learning Agent

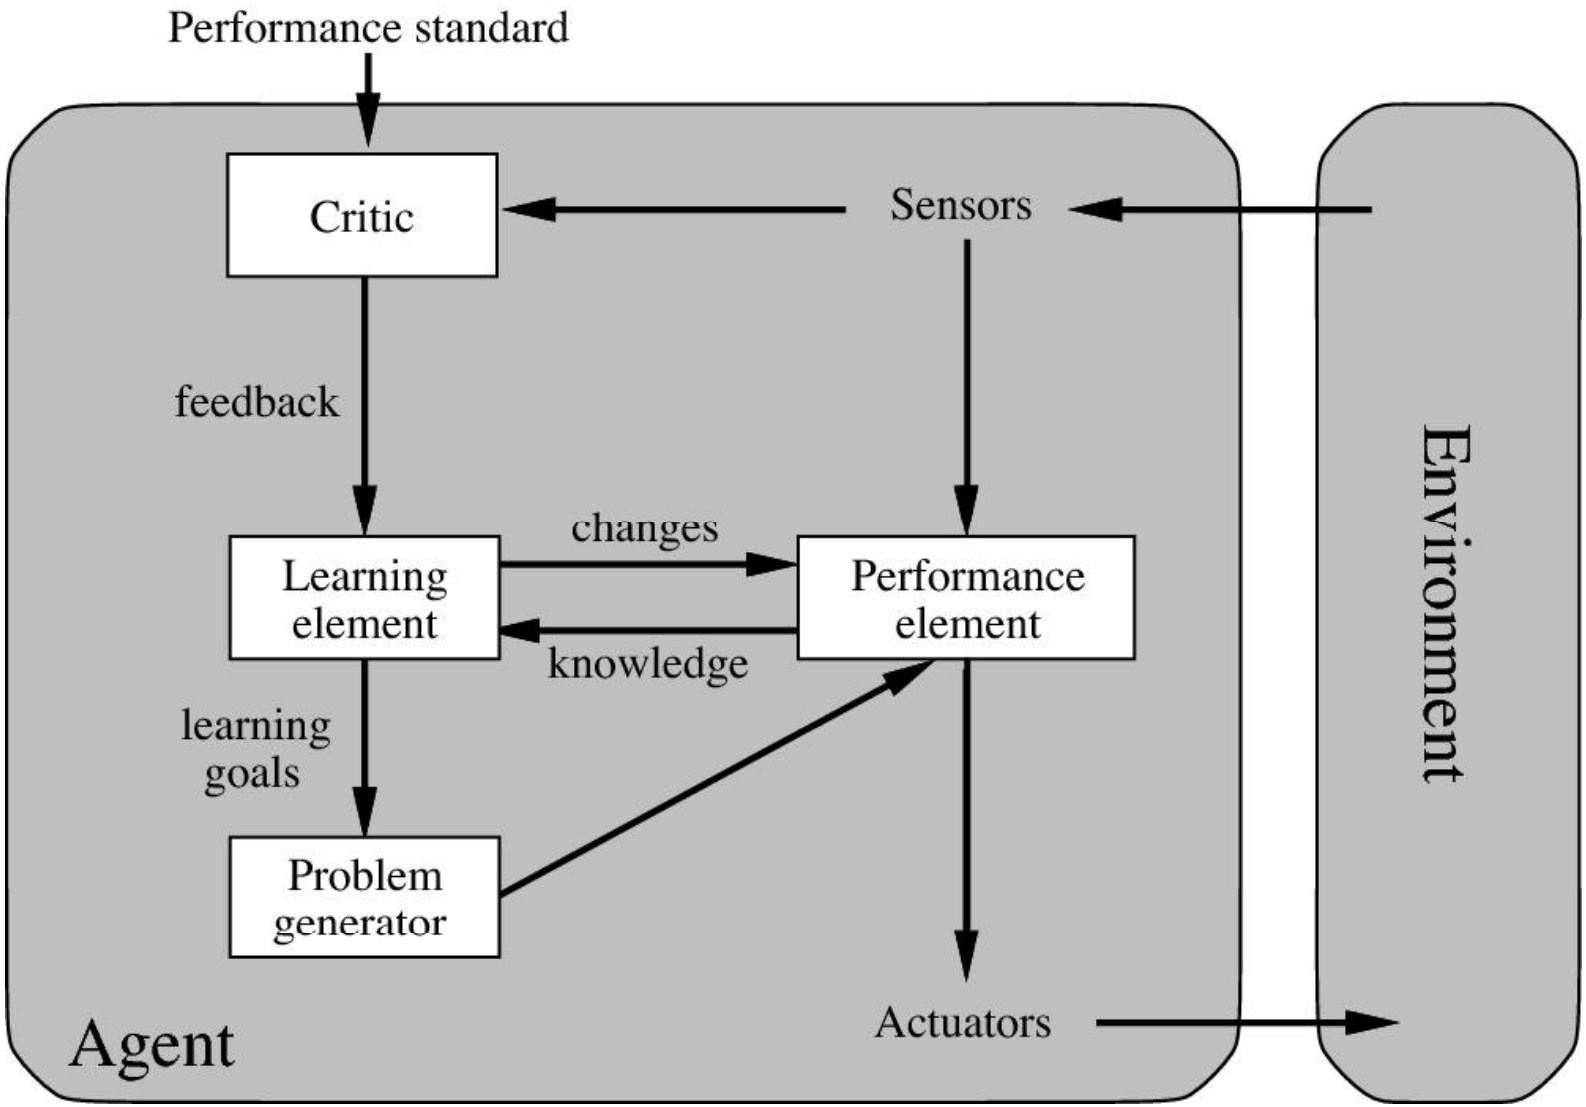

Atomic/Factored/Structured