Agents

CPSC 383: Explorations in Artificial Intelligence and Machine Learning
Fall 2025

Jonathan Hudson, Ph.D
Associate Professor (Teaching)
Department of Computer Science
University of Calgary

August 27, 2025

UNIVERSITY OF CALGARY

Outline

- Agents
- Agent types
- Learning Agents
- Task Environment
- Environment Types

Agents

Double O

Agent

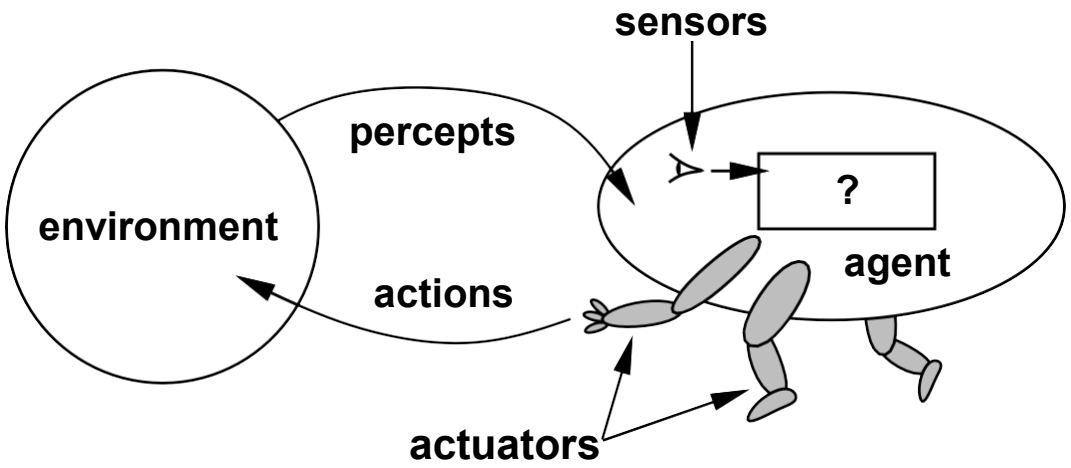

- Agents include humans, robots, softbots, thermostats, etc.
- An agent can be anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators
- The agent function maps from percept histories to actions

Agent Types

007?

Reflex Agents

- Reflex agents:
  - Choose action based on current perception

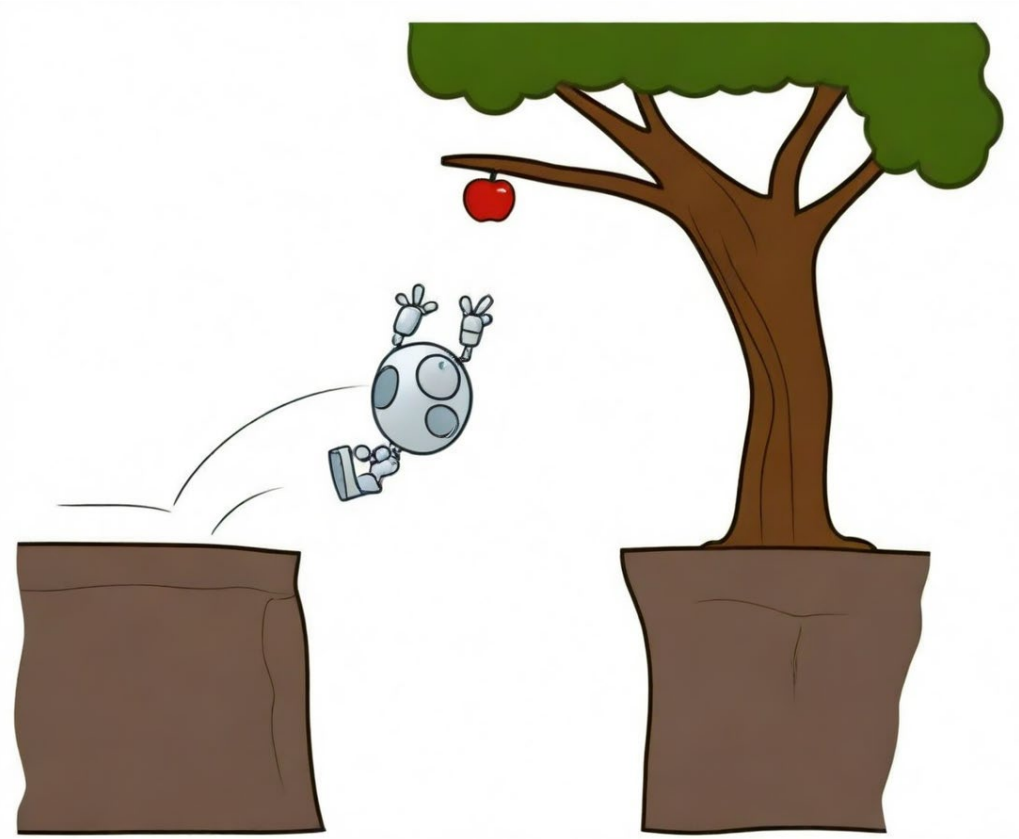

Reflex Agents

- Reflex agents:
  - **Choose action** based on current perception
- *Can a reflex agent be rational?*

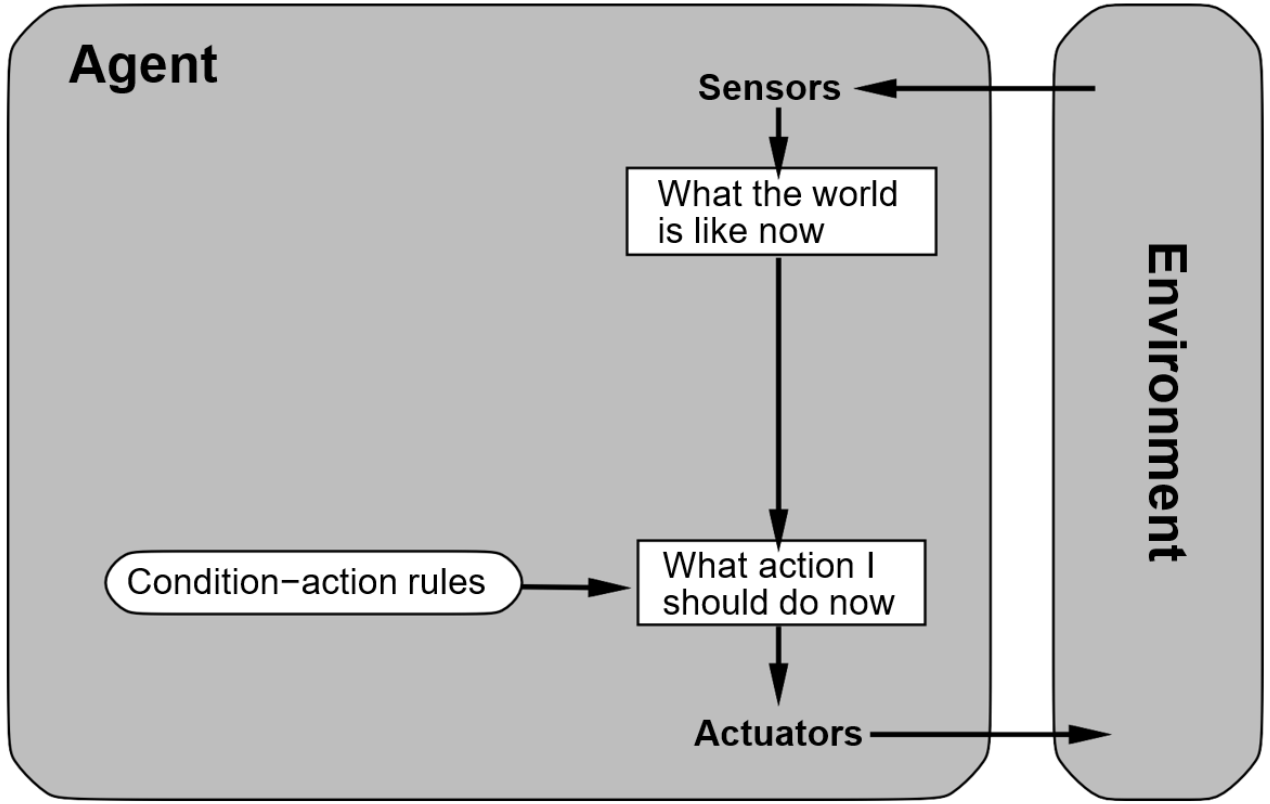

Reflex Agents

- Reflex agents:
  - **Choose action** based on current perception
- *Can a reflex agent be rational?*
  - *By itself, no*
  - *But the human designing itself could embed their rational decisions*

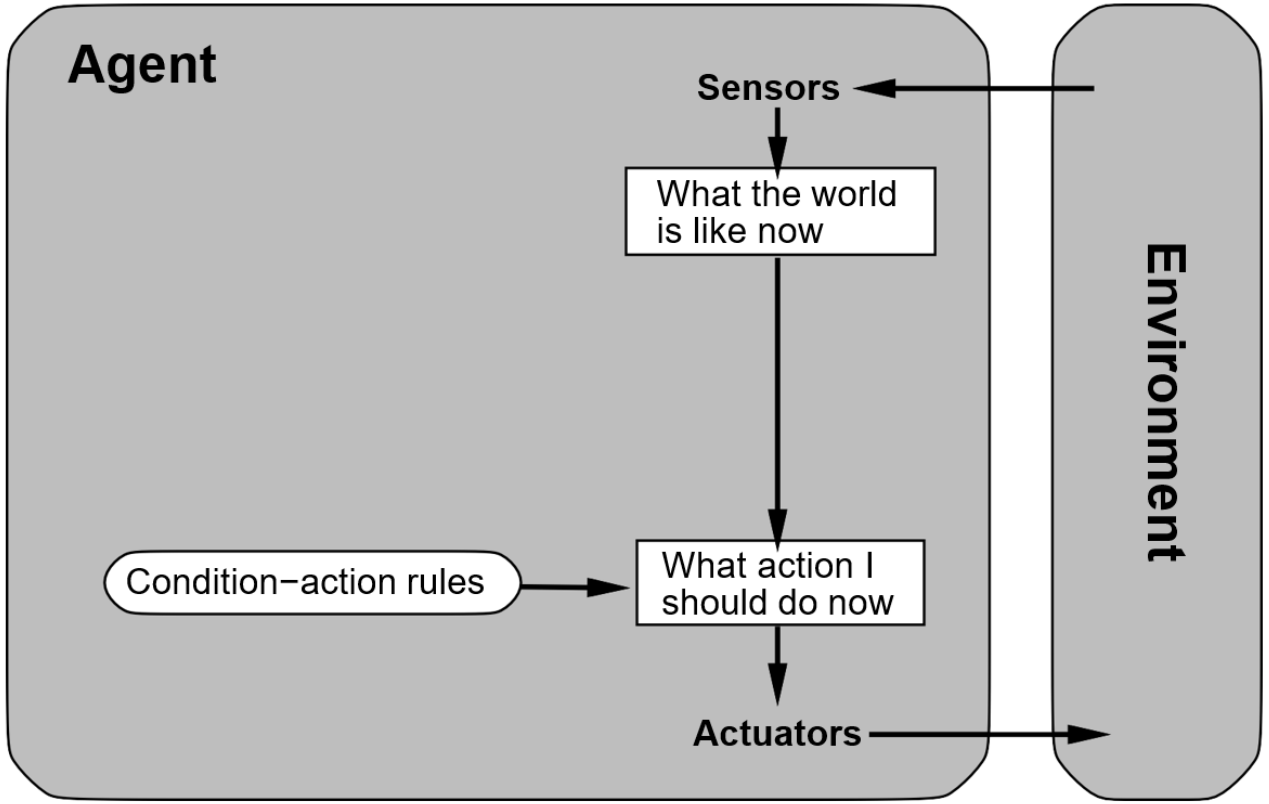

Reflex Agents with State

- Reflex agents with State:
  - Choose action based on current perception
  - Have **memory** or the world's state

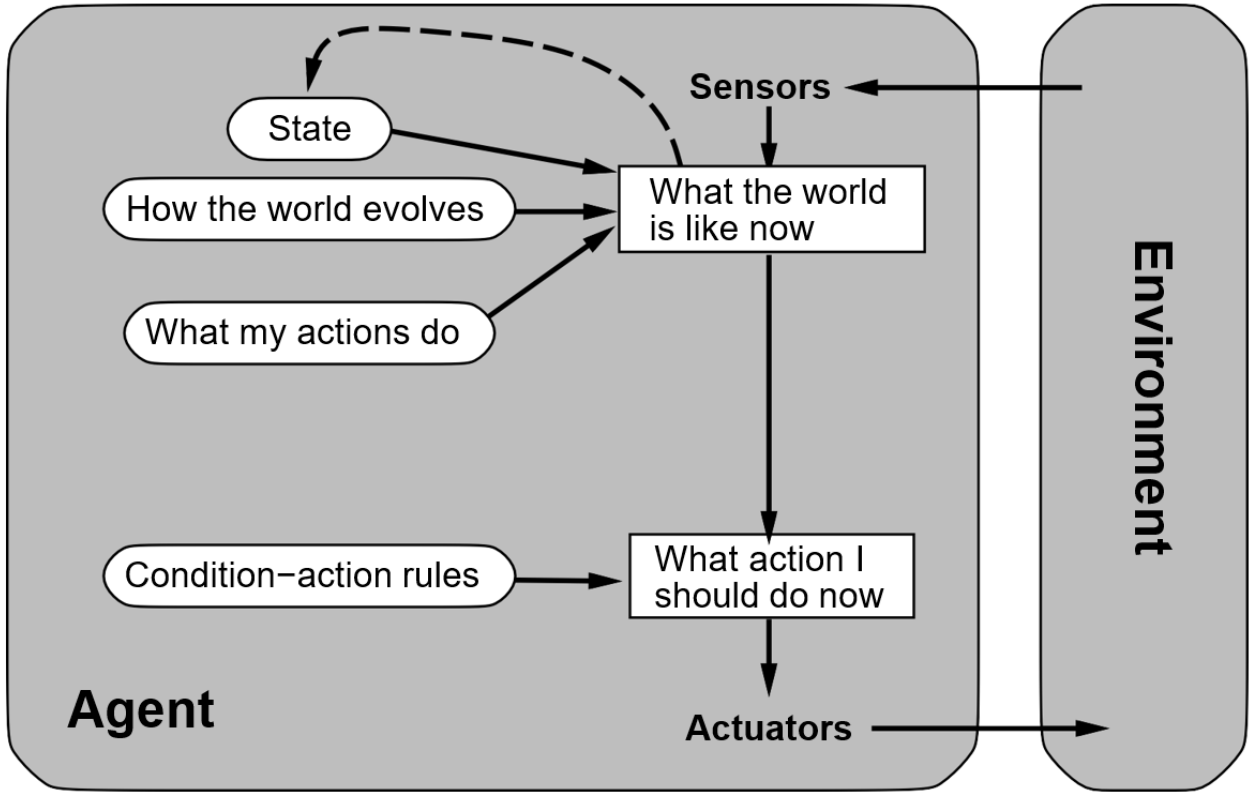

Goal-based Agents

- Goal-based Agents:
  - Choose action based on current perception
  - Have memory or the world's state
  - Have more **abstract definition of overall goal** that drives more complex choice of situation-action pairs

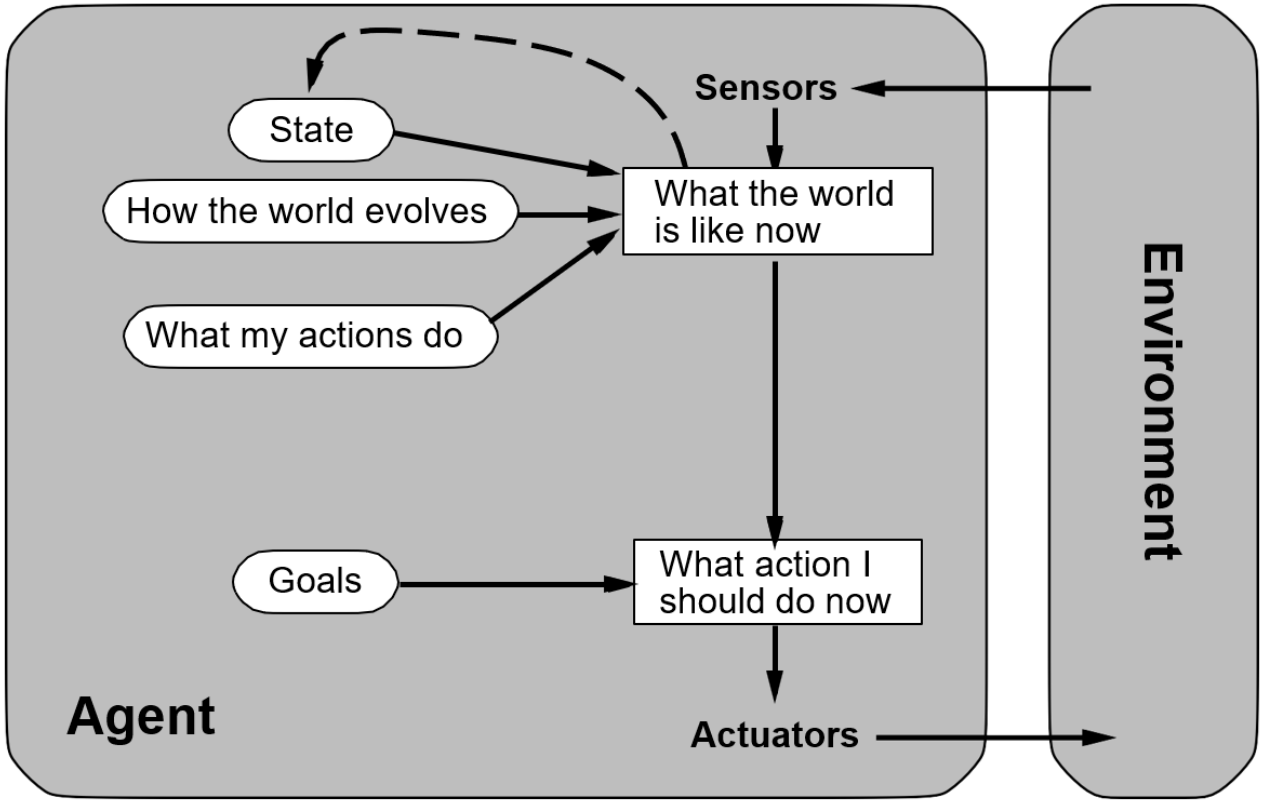

Model-based Agents

- Model-based Agents:
  - Choose action based on current perception
  - Have memory or the world's state
  - Have more abstract definition of overall goal that drives more complex choice of situation-action pairs
  - **Model** the **consequence** of action choice

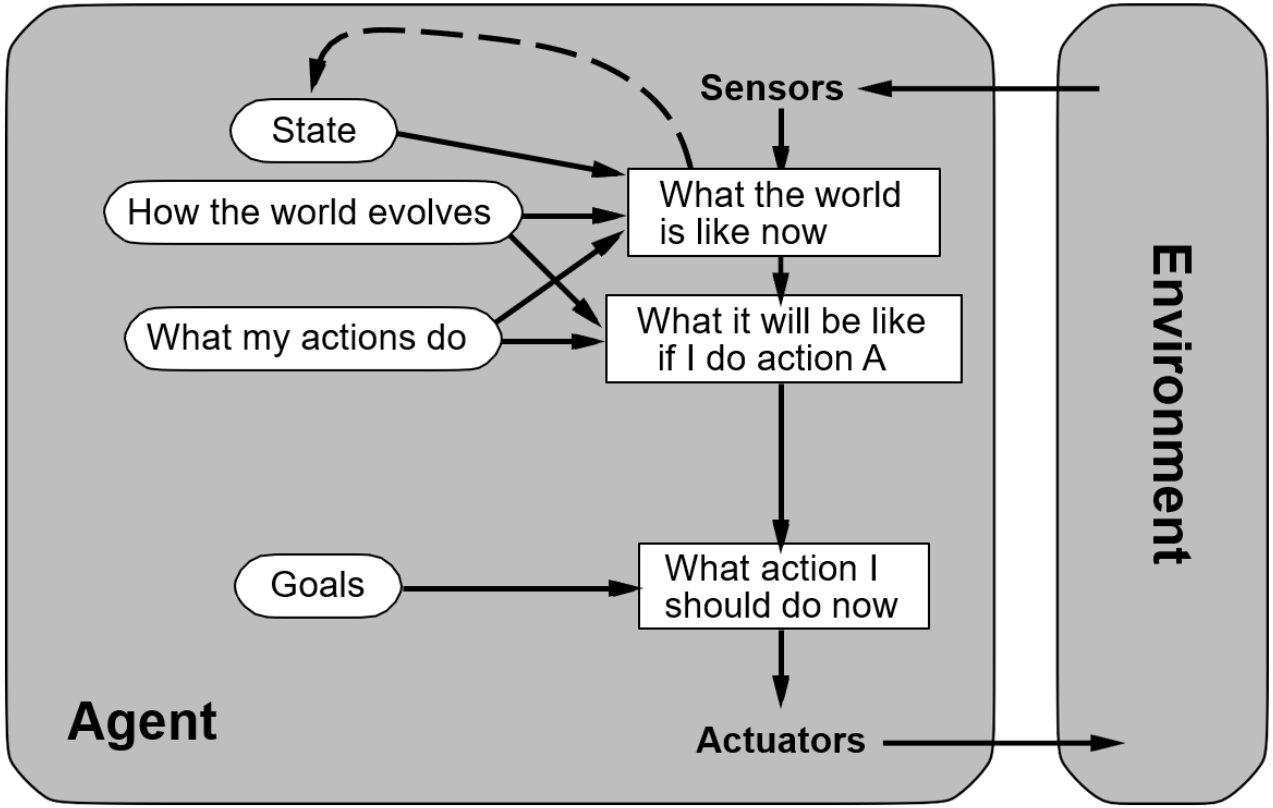

Planning Agents

- Planning agents:
  - Model-based agents that explore **multiple step** consequences of modeled actions

- Types
  - Optimal vs. complete planning
  - Planning vs. re-planning

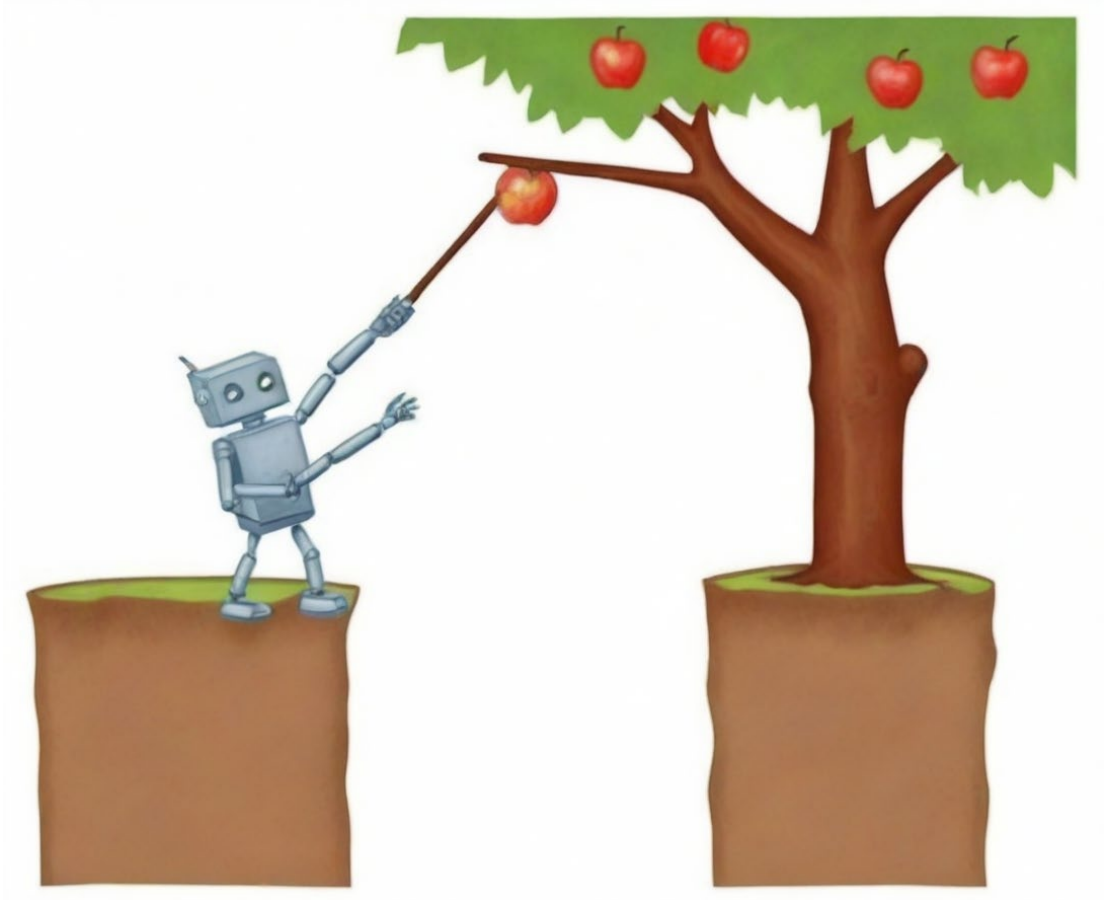

Utility-based Agents

- Utility-based Agents:
  - Choose action based on current perception
  - Have memory or the world's state
  - Have more abstract definition of overall goal that drives more complex choice of situation-action pairs
  - Model the consequence of action choice
  - Use **utility** function to judge consequences
- *A supplemented reflex agent into a utility-based agent is now rational!*
  - *existence of utility function we can maximize*

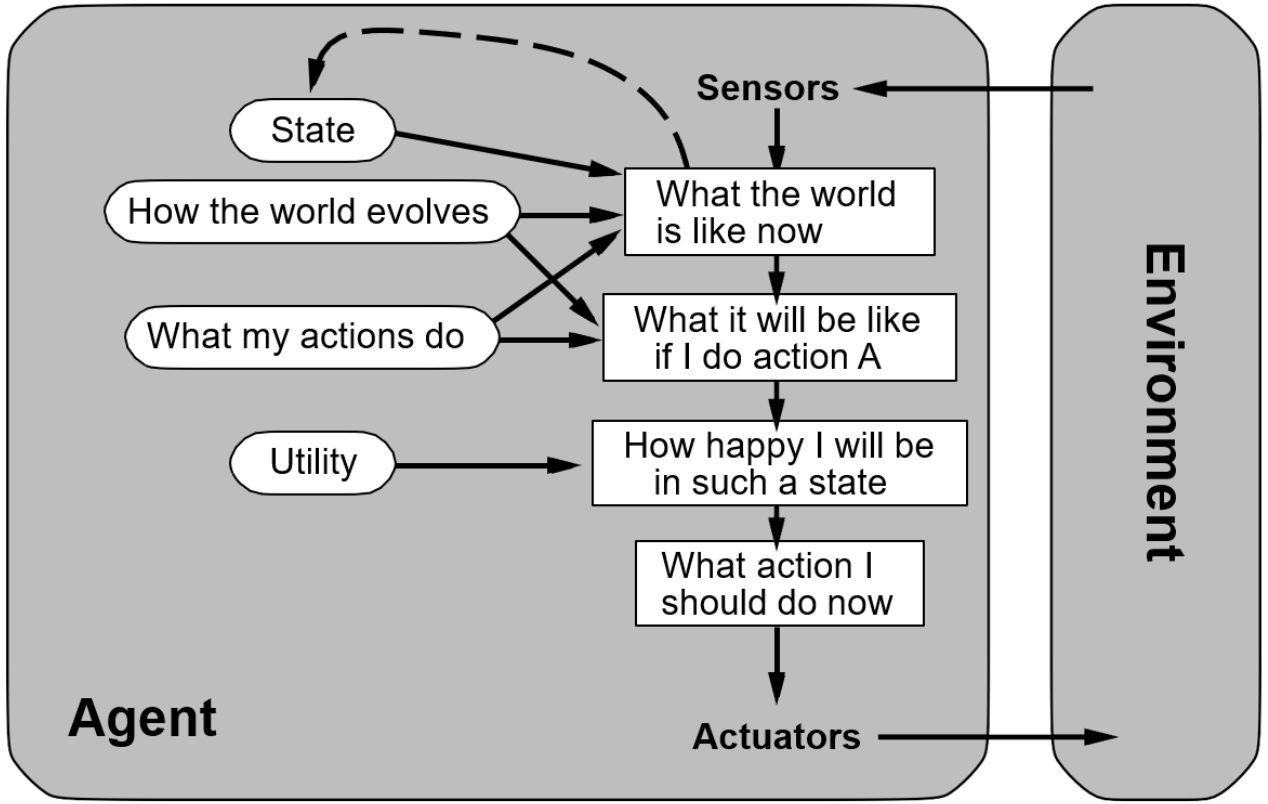

Learning Agents

Q

Learning Agents

- Learning Agents:
  - Choose action based on current perception
  - Have memory or the world's state
  - Have more abstract definition of overall goal that drives more complex choice of situation-action pairs
  - Model the consequence of action choice
  - Use utility function to judge consequences
  - Feedback loops to change future behaviour

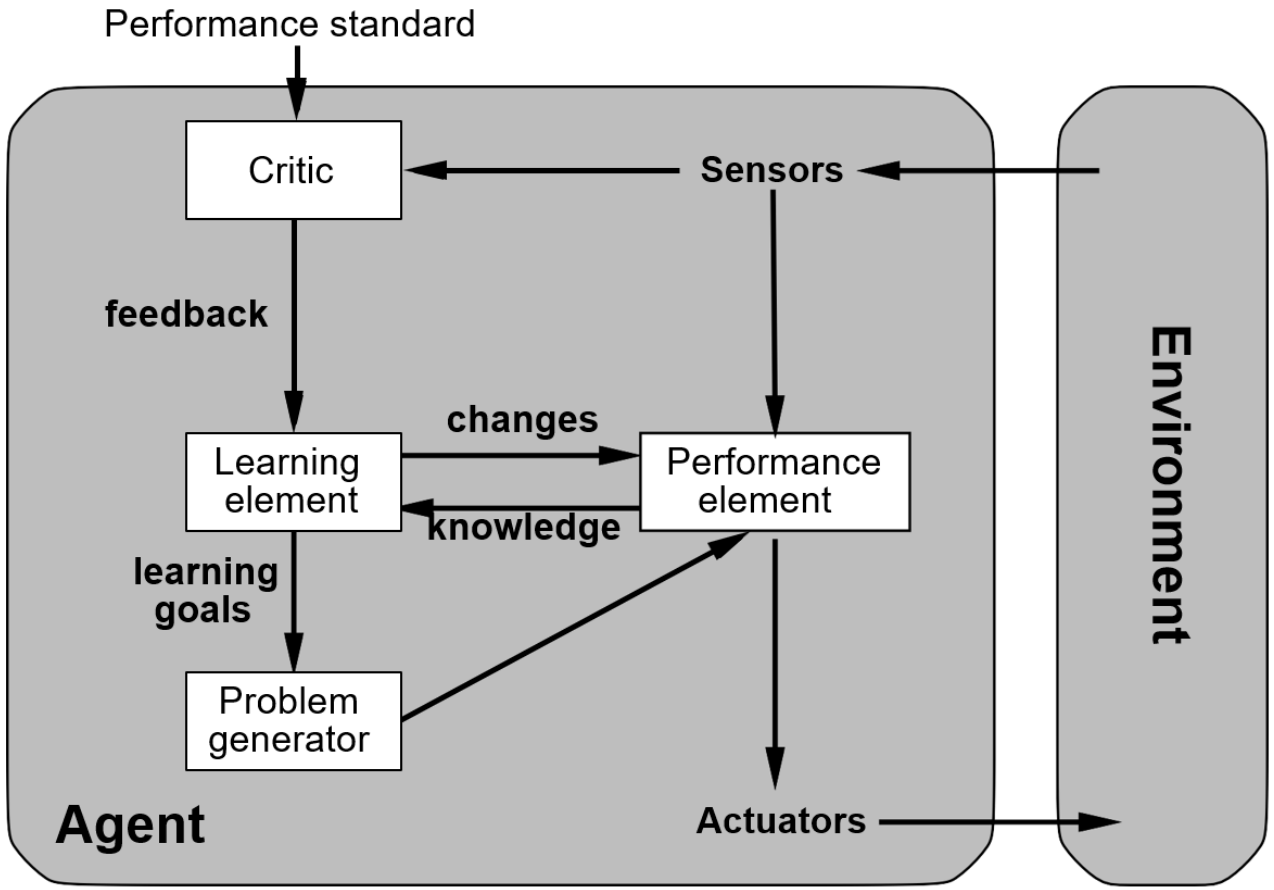

Task Environment

I have a mission for you if you choose to accept it

Task Environment

- To design a rational agent, we must specify the task environment
- Consider, e.g., the task of designing an automated taxi:

1. **Performance measure??**
2. **Environment??**
3. **Actuators??**
4. **Sensors??**

PEAS

- To design a rational agent, we must specify the task environment
- Consider, e.g., the task of designing an automated taxi:

1. **Performance measure??** safety, destination, profits, legality, comfort, . . .
2. **Environment??** US streets/freeways, traffic, pedestrians, weather, . . .
3. **Actuators??** steering, accelerator, brake, horn, speaker/display, . . .
4. **Sensors??** video, accelerometers, gauges, engine sensors, keyboard, GPS, . . .

Environment types

This message will self-destruct

Environment Types

- The environment type largely determines the agent design

	Solitaire	Backgammon	Internet Shopping	Taxi
Observable	Yes	Yes	No	No
Deterministic	Yes	No	Partially	No
Episodic	No	No	No	No
Static	Yes	Partially	Partially	No
Discrete	Yes	Yes	Yes	No
Single Agent	No	No	Yes (auctions?)	No

- The real world is (of course) partially observable, stochastic, sequential, dynamic, continuous, multi-agent

Summary

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
- Several basic agent architectures exist:
  - reflex, reflex with state, goal-based, utility-based
- PEAS descriptions define task environments
- Environments are categorized along several dimensions:
  - observable? deterministic? episodic? static? discrete? single-agent?

Next…search

Jonathan Hudson, Ph.D.
jwhudson@ucalgary.ca
https://cspages.ucalgary.ca/~jwhudson/

UNIVERSITY OF CALGARY