Artificial Intelligence

The Very Idea

Vasant G. Honavar
Dorothy Foehr Huck and J. Lloyd Huck Chair in Biomedical Data Sciences and Artificial Intelligence
Professor of Data Sciences, Informatics, Computer Science, Bioinformatics & Genomics and Neuroscience
Director, Artificial Intelligence Research Laboratory
Director, Center for Artificial Intelligence Foundations and Scientific Applications
Associate Director, Institute for Computational and Data Sciences
Pennsylvania State University

vhonavar@psu.edu
http://faculty.ist.psu.edu/vhonavar
http://ailab.ist.psu.edu

Intelligent agents – A practical avenue for AI advances

An agent can

- perceive its environment through its sensors
- represent aspects of its environment
- reason and predict consequences of its actions
- act on its environment through its effectors
- make rational choices
- learn from experience
- communicate with other agents
- interact with other agents
- display autonomy
- display creativity
- exhibit purposeful behavior

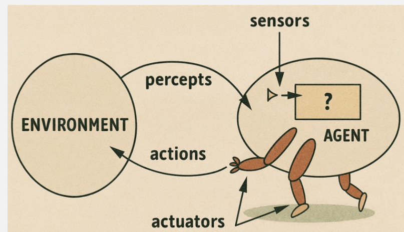

Agents and environments

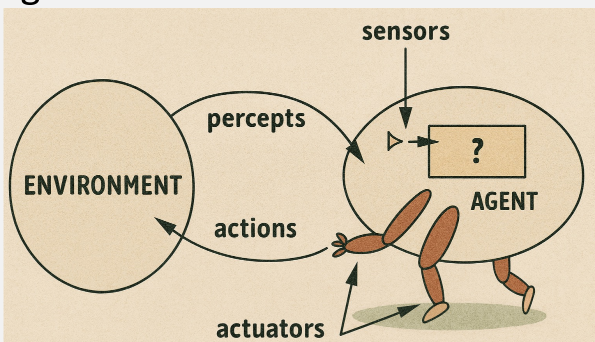

- Agents include humans, animals, robots, soft-bots, thermostats, etc.
- An agent senses the state of the environment through its sensors, obtaining *percepts*
- The *agent function* maps percept sequence to actions

Agents and environments

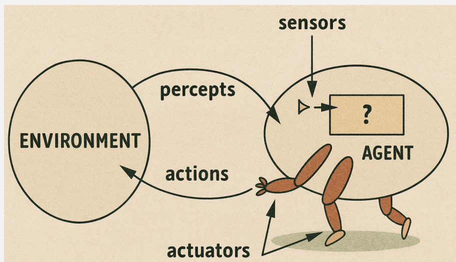

- The *agent function* f is encoded by the *agent program*.
- The agent program runs on a physical *architecture* (robot, human, etc.) to produce f.

AI Thermostat

- One of the simplest agents we can imagine is a simple thermostat
- The thermostat senses the ambient temperature using its temperature sensor
- If the sensed temperature is greater than the preset temperature, it turns on the air conditioner
- If the sensed temperature is less than equal to the preset temperature, it shuts off the air conditioner if it is on
- If the sensed temperature is less than the preset temperature, it turns on the heater
- If the sensed temperature is greater than equal to the preset temperature, it shuts off the heater if it is on

The vacuum-cleaner agent

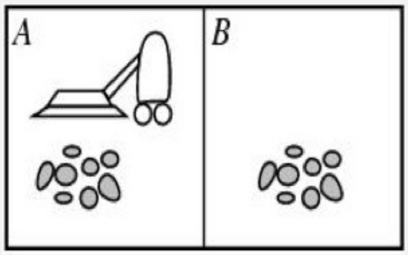

- Environment: rooms A and B
- Percepts: [location, content] e.g. [A, Dirty]
- Actions: left, right, cleanup, and no-op

The vacuum-cleaner agent

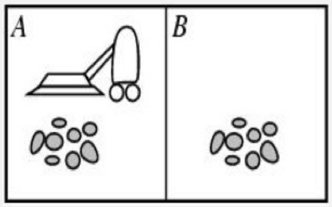

Percept sequence	Action
[A,Clean]	Right
[A, Dirty]	Cleanup
[B, Clean]	Left
[B, Dirty]	Cleanup
[A, Clean],[A, Clean]	Right
[A, Clean],[A, Dirty]	Cleanup
...	...

The vacuum-cleaner agent

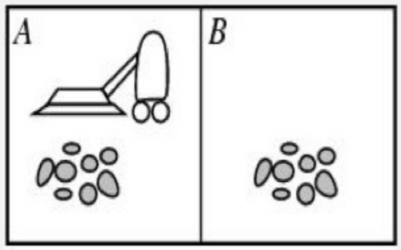

if *status is Dirty* then *Cleanup*
else if *location is A* then move *Right*
else if *location is B* then move *Left*

How an agent should behave

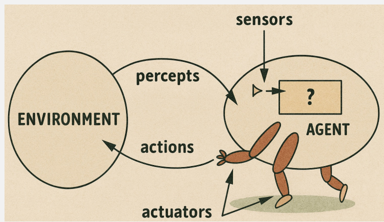

What an agent ought to do at any given time depends on:

- The performance measure that defines the criterion of success against which the agent is evaluated
- The agent's prior knowledge of its environment
- The actions that the agent has at its disposal
- The agent's percept sequence up to that time

Rational Agents

- 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
  - whatever built-in knowledge the agent has and
  - the actions at its disposal
- Examples of performance measures
  - the amount of dirt cleaned within a certain time
  - how clean the floor is
  - the amount of dirt cleaned per unit of electricity used
- What is <u>your</u> performance measure?
- Who decides what the performance measure should be?
  - Internal drives
  - External rewards

Rationality

- What is rational at a given time depends on:
  - The performance measure
  - What the agent knows
  - The actions that the agent can perform
  - What the agent has observed (through its sensors)

Rationality ≠ omniscience

- You are stopped at a red light at an intersection
- You watch it turn yellow and then green
- Being a rational driver who knows the traffic rules, you take your foot off the breaks and start driving across the intersection
- Meanwhile, at 30,000 feet in the air,
  - the cargo door falls off an airliner in flight,
  - comes crashing down, and
  - lands right on top of you,
  - pinning you down while
  - other vehicles crash into you causing a pileup
- Was your behavior rational?
  - If so, why?
  - If not, why not?

Rationality ≠ omniscience

- An omniscient agent knows everything there is to know
- A rational agent knows only
  - the current state of the world seen through its sensors and
  - the expected state of the world resulting from its action
- Your behavior at the intersection was rational based on what you knew
  - You are not omniscient
  - You couldn't possibly be expected to know that the cargo door of the airplane flying overhead would come crashing down on you

Rationality ≠ clairvoyance

- A clairvoyant agent knows the actual effects of each actions before it is performed
- A rational agent knows only the expected effect of its action

Rationality ≠ perfection

- Rational agent maximizes
  - expected performance
- Perfect agent maximizes
  - actual performance

Is the thermostat agent rational?

- The performance measure awards one point for each hour the temperature is maintained within ± 2 degrees Celsius from the preset temperature
- Suppose the agent knows:
  - The preset temperature in degrees Celsius.
  - If the air conditioner is turned on, it cools the room, causing a decrease in the temperature;
  - if the heater is turned on, it heats the room, causing an increase in the temperature;
  - The heater, once turned on, will remain on until it is turned off;
  - The air conditioner and the heater cannot both be on at the same time.
- The available actions are to turn the air conditioner on or off, and the heater on or off, or do nothing (leave things the way they are).

Is the thermostat agent rational?

- The thermometer or room temperature sensor is operational and provides the correct temperature readings
- An air conditioner status sensor tells the thermostat whether the air conditioner is on
- The heater status sensor tells the thermostat whether the heater is on
- If the sensed temperature is greater than the preset temperature, the thermostat turns on the air conditioner
- If the sensed temperature is less than equal to the preset temperature, it shuts off the air conditioner if it is on
- If the sensed temperature is less than the preset temperature, it turns on the heater
- If the sensed temperature is greater than equal to the preset temperature, it shuts off the heater if it is on

Is the new thermostat agent rational?

- Suppose the thermostat is additionally equipped with room occupancy sensor that tells it whether the room is occupied by people
- The performance measure awards the agent
  - one point for each hour the temperature is maintained within $\pm$ 2 degrees Celsius of the preset temperature whenever the room is occupied; and
  - -10 points if either the heater or air conditioner are found to be on when the room is unoccupied by people
- Under these circumstances, does the agent function described previously ensure that the agent is rational? Why or why not?
- How can you restore rationality?

Agents and their environments

- We need to specify
  - Performance measure
  - Environment
  - Sensors
  - Actuators or Effectors

Environments

For a fully self-driving taxi:

- Performance
  - Safety, reaching destination, profits, comfort
- Environment
  - Streets, other traffic, pedestrians, weather, …
- Actuators
  - Steering, accelerating, brake, horn, speaker, display,…
- Sensors
  - Video, sonar, speedometer, engine sensors, keyboard, GPS, …

Environment types

- Fully observable versus partially observable
- Deterministic versus non-deterministic
- Episodic vs non-episodic
- Static vs dynamic
- Discrete versus continuous
- Single-agent versus Multiagent
- Open versus closed

Environment types

Fully vs. partially observable: an environment is fully observable when the sensors can detect all aspects that are relevant to the choice of action

	Solitaire	Backgammon	Internet shopping	Taxi
Observable?? Deterministic?? Episodic?? Static?? Discrete?? Single-agent?? Open??				

Environment types

Fully vs. partially observable: an environment is fully observable when the sensors can detect all aspects that are relevant to the choice of action.

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??				
Episodic??				
Static??				
Discrete??				
Single-agent??				
Open??				

Is the environment fully observable in poker?

PennState Institute for Computational and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

PennState Clinical and Translational Science Institute

Environment types

Deterministic vs. stochastic: if the next environment state is completely determined by the current state the executed action then the environment is deterministic.

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??				
Episodic??				
Static??				
Discrete??				
Single-agent??				
Open??				

PennState College of Information Sciences And Technology

AI 100 Fall 2025

Vasant G Honavar

Dus het is beter om te kijken of de omgebing determinitisch of stochastisch is vanuit het standpunt van de agent.

Als de omgeving deterministisch is behalve de acties van de andere agenten dan is de omgeving strategisch.

Environment types

Deterministic vs. non-deterministic: if the next environment state is completely determined by the current state and the action executed by the agent, then the environment is deterministic

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??				
Static??				
Discrete??				
Single-agent??				
Open??				

Environment types

Episodic vs. non episodic: In an episodic environment the agent's experience can be divided into episodes – agent's actions in consecutive episodes are independent

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??				
Static??				
Discrete??				
Single-agent??				
Open??				

Environment types

Episodic vs. non episodic: In an episodic environment agent's experience can be divided into episodes – agent's actions in consecutive episodes are independent

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	MAYBE	YES
Static??				
Discrete??				
Single-agent??				
Open??				

Environment types

Static vs. dynamic: If the environment can change while the agent is choosing an action, the environment is dynamic

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	MAYBE	YES
Static??				
Discrete??				
Single-agent??				
Open??				

If the agent performance score changes ahen time passes, the environment is semi-dynamic.

Environment types

Static vs. dynamic: If the environment can change while the agent is choosing an action, the environment is dynamic

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	YES	MAYBE
Static??	YES	YES	MAYBE	NO
Discrete??				
Single-agent??				
Open??				

If the agent performance score changes ahen time passes, the environment is semi-dynamic.

Environment types

Discrete vs. continuous

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	YES	MAYBE
Static??	YES	YES	MAYBE	NO
Discrete??				
Single-agent??				
Open??				

If the agent performance score changes ahen time passes, the environment is semi-dynamic.

Environment types

Discrete vs. continuous

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	YES	MAYBE
Static??	YES	YES	MAYBE	NO
Discrete??	YES	YES	YES	NO
Single-agent??				
Open??				

If the agent performance score changes ahen time passes, the environment is semi-dynamic.

Environment types

Single vs. multi-agent: Does the environment contain other agents?

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	NO	MAYBE
Static??	YES	YES	MAYBE	NO
Discrete??	YES	YES	YES	NO
Single-agent??				
Open??				

If the agent performance score changes ahen time passes, the environment is semi-dynamic.

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

PennState Institute for Computational and Data Sciences

PennState Clinical and Translational Science Institute

Environment types

Single vs. multi-agent: Does the environment contain other agents?

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	YES	YES
Static??	YES	YES	MAYBE	NO
Discrete??	YES	YES	YES	NO
Single-agent??	YES	NO	NO	NO
Open??				

PennState College of Information Sciences And Technology

AI 100 Fall 2025

Vasant G Honavar

If the agent performance score changes ahen time passes, the environment is semi-dynamic.

Environment types

Open versus closed: Can entities enter and leave the environment?

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	YES	MAYBE
Static??	YES	YES	MAYBE	NO
Discrete??	YES	YES	YES	NO
Single-agent??	YES	NO	NO	NO
Open??				

If the agent performance score changes ahen time passes, the environment is semi-dynamic.

Center for Artificial Intelligence Foundations & Scientific Applications
Artificial Intelligence Research Laboratory

PennState Institute for Computational and Data Sciences

PennState Clinical and Translational Science Institute

Environment types

Open versus closed: Can agents enter and leave the environment?

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	FULL	FULL	PARTIAL	PARTIAL
Deterministic??	YES	NO	YES	NO
Episodic??	NO	NO	YES	YES
Static??	YES	YES	MAYBE	MAYBE
Discrete??	YES	YES	YES	NO
Single-agent??	YES	NO	NO	NO
Open??	NO	NO	YES	YES

PennState College of Information Sciences And Technology

AI 100 Fall 2025

Vasant G Honavar

If the agent performance score changes ahen time passes, the environment is semi-dynamic.

Environment types

- The simplest environment is
  - Fully observable, deterministic, episodic, static, discrete and single-agent
- Many real-world environments can be:
  - Partially observable, stochastic, non-episodic, dynamic, continuous and multi-agent

Agent types

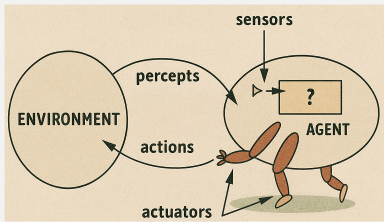

All agents have:

- Input = current percept
- Output = action
- Program = processes input to produce output

Types of Agents

- Simplest possible agent design makes use of a state-action table
- The state-action table lists for each state, the action to be taken by the agent
- The agent senses the state, and looks up the action to be taken from the state-action table and execute it
- Do you see a problem with this approach?
- It doomed to failure in all but the simplest cases
- Why?
- The number of states grows exponentially with the complexity of the percept
  - 1 bit sensed – 2 states (0, 1)
  - 2 bits sensed – 4 states (00, 01, 10, 11)
  - 3 bits sensed – 8 states (000, 001, 010, 011, 100, 101, 110, 111)

Agent types

Basic agent types

- Simple reflex agents or reactive agents
- Model-based reflex agents
- Deliberative agents (often goal-based)
- Utility-based agents

All these agent types can be

- Learning
- Communicative
- Interactive (competitive, collaborative)

Simple Reflex Agent

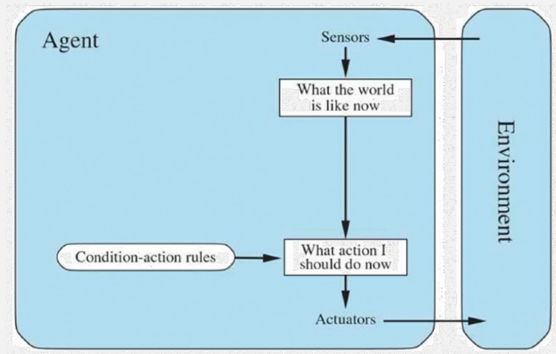

- Select action on the basis of only the current percept
  - E.g. the vacuum-cleaner agent
- Implemented through condition-action rules
  - Example: vacuum cleaner agent
  - If *status* is *dirty* then *cleanup*

The vacuum-cleaner agent

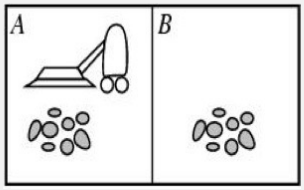

Agent function

if *status is Dirty* then *Cleanup*

else if *location is A* then move *Right*

else if *location is B* then move *Left*

- Note that this design works only for the room layout shown
- A more general design needs to use a map of the floor and plan to choose its moves

Simple reflex agent

- Has limited intelligence
- Has only a limited view of the state of the environment
- Has program specified by the designer
- The agent program must be modified if the environment changes
- Design can be modified to work in different environments or different tasks
  - simply by specifying a set of condition-action rules to be interpreted and applied in different contexts

The vacuum-cleaner agent

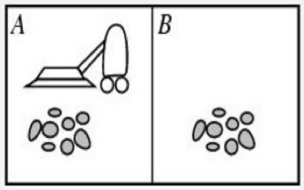

Agent function

if *status is Dirty* then *Cleanup*
else if *location is A* then move *Right*
else if *location is B* then move *Left*

- Suppose its location sensor fails and it is forced to operate with only its dirt sensor
- Now it has only two possible percepts, namely, *status* being *Dirty* or *Clean*
- Now if the *status* is *Dirty*, it will activate the first rule, triggering the action Cleanup
- But what about if the *status* is *Clean*?

The vacuum-cleaner agent

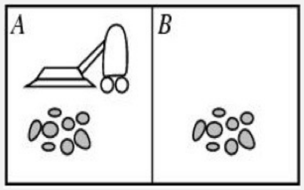

Agent function

if *status is Dirty* then *Cleanup*

else if *location is A* then move *Right*

else if *location is B* then move *Left*

- When the location sensor fails, the agent can't sense its location
- None of the rules match its current percept Clean, thereby effectively paralyzing the agent
- In this case, one can "fix" this problem by adding some randomness to the agent's behavior: when no rules match the current percept, randomly move to the other square
- Incomplete picture of the state of the environment makes the agent's task harder, and in some cases, impossible

Simple reflex agents

- Given a suitable agent function, a reflex agent succeeds only when the environment is fully observable
- What happens if the environment is only partially observable?
- The agent has no way to distinguish between two or more distinct environmental states that are indistinguishable from the agent's point of view
- This limits the applicability of simple reflex agents to fully observable envionments

Simple reflex agents and Dennett's Darwinian creatures

Simulated evolution

- Generate a population of agents with random agent functions
- Repeat:
  - Assess the fitness of agents in the environment
  - Replace a fraction of the population through fitness-proportionate selection and recombination

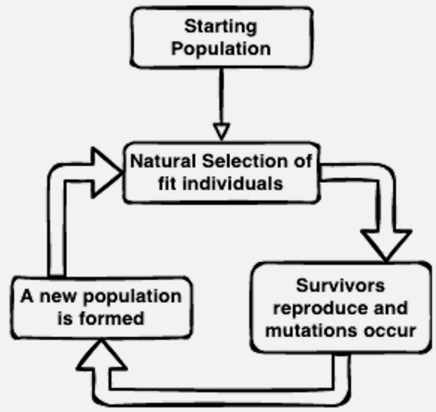

From simple reflex agents to reflex agents with state

- Consider a Mars rover agent designed to pick up rock samples of distinct types on the Mars surface
- A reflex agent would be able to pick up every rock it encounters (if *you see a rock*, then *pick it up*) by systematically exploring the Mars surface
- It can't to ensure that each rock sample it collects is of a different type than those that it has already collected
- Why not?
- It cannot keep track of the types of rocks it has collected
- Why? It has no memory

From simple reflex agents to reflex agents with state

- A reflex agent with state can encode as part of its internal state description, the types of rocks it has collected already
- It can decide whether to pick up the current rock sample based on whether it has already collected a rock of the same type or not

Reflex agent with state

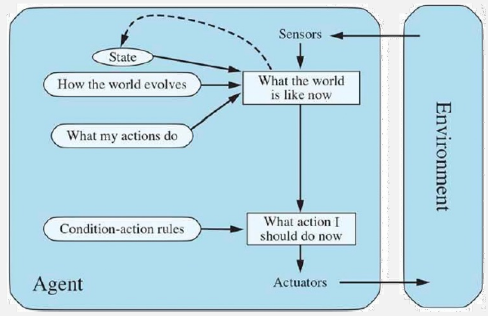

- Maintains internal state
- Over time update state using based on its knowledge and observation
  - How does the world change?
  - How do the actions affect the world?
- Can handle partially observable environments

Reflex agents with state

- Find a rule that matches current state and percept
- Apply the rule to choose action
- Perform action
- Update state
- Example: Agent controlling an elevator

Reflex agents with state

- State update requires
  - a state transition model that describes how the world changes over time
    - due to its intrinsic dynamics, independently of the agent's actions and
    - due to the agent's actions.
  - A sensor model that describes how aspects of state of the world are mapped to percepts.

- Example: in the case of the vacuum cleaner agent, the status of a location changes from *Dirty* to *Clean* when the Cleanup action is executed by the agent at that location

Example: A more sophisticated vacuum cleaner

- Imagine a robot vacuum cleaner operating in a house with many rooms
- It has sensors that tell it whether the current room is dirty or clean
- But it cannot see the *entire house* or know the dirt status of other rooms
- It can move left/right or clean
- Why does this robot need an internal state?
  - Without an internal state, it cannot remember which room it has cleaned already
- If it is given a map of the house, and can remember the rooms it has visited already, it can systematically visit all the rooms and clean the ones that are dirty until every room has been cleaned

Example: A more sophisticated vacuum cleaner

- Percepts: Current location, dirt status (dirty/clean).
- State: A record of what rooms have been visited and whether they were clean/dirty.
- Rules:
  - If current location is dirty → Clean.
  - Else if an adjacent room not yet visited → Move there.
  - Else if all known rooms visited → Return to base.
- Actions: Move or Clean.

Agents with goals

- Even perfect sensing of the current state of the environment is not necessarily sufficient for an agent to decide what it should do
- Consider a self-driving taxi at a 4-way intersection in Washington DC.
- The taxi has three choices: turn left, turn right, or go straight.
- The right thing to do depends on not just the state of the environment, but the taxi's destination!
- Without a destination to get to, all three actions are equally good so the the taxi would drive around aimlessly

Agents with goals

- Suppose the taxi has a goal that it needs to achieve, for example, get to the Union Station in Washington DC
- The taxi's decision as to what to do next can be informed by its goal of getting to the Union Station
- Example: Taxi should turn right if that puts it on the road that takes it towards the Union Station as opposed to turning left if it leads away from the Union Station

Goal-based agents

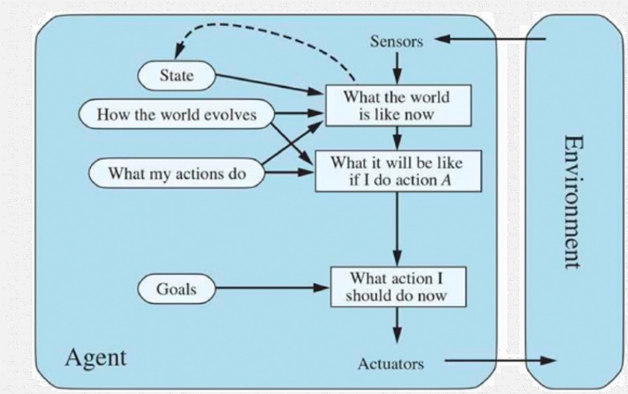

- Agent is provided a goal
- The agent seeks to achieve the specified goal
- Attaining a goal may require a long sequence of actions
- Needs a model (representation) of the world

Goal-based agents

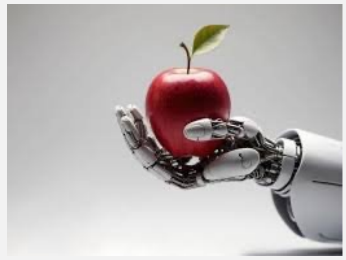

- In some scenarios, selecting an optimal action based on the goal is straightforward
- Example:
  - If the goal is to get an apple, and
  - the apple is within reach of the agent,
  - and one of the possible action choices is pickup the apple,
  - the goal can be achieved by simply by executing that action

Goal-based agents

- In general, achieving a goal may need an entire plan or a sequence of actions
- Example:
  - Get in the car → Drive to the grocery store → Park the car → Get out of the car → Walk into the grocery store → Find the aisle with fruits → Find apple → Pickup apple → Pay → Walk out of the store → Get in the car → Drive home → Park the car → Enter the home
- Finding such a plan requires knowledge of
  - The preconditions and effects of actions
    - To drive the car, the agent must be seated in the car, have car keys, and the license to drive, and so on.
  - the world e.g., that grocery stores sell apples …

Optimal plan seeking agents

- A goal-based self-driving taxi equipped with the appropriate knowledge and planning abilities can take you from your current location, say Ronald Reagan International Airport, to the Union Station in Washington DC.
- But it may take a longer route, or it may take congested route, when there are better options available.
- We may want the agent to find an optimal e.g., shortest or quickest route instead of simply finding any route
- More on this later

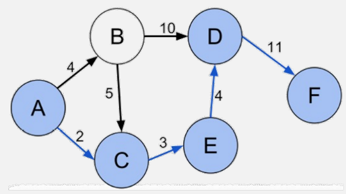

Utilitarian agent

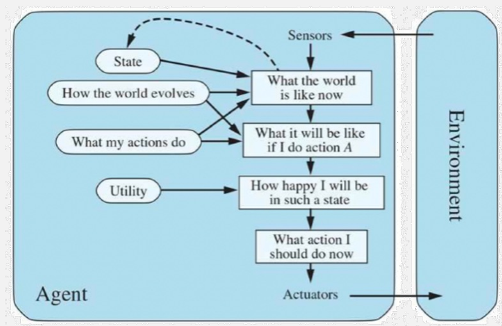

- Different goals have different utilities for the agent
- Utility function maps a (sequence of) state(s) onto a real number (utility)
- Utility can influence choice of action
- Goal-based agents are a special case (utility is 1 for the goal state, and 0 for all other states)

Learning Agents

- How do agents come to possess the knowledge that they need?
- The agent may be painstaking programmed by their designers with the knowledge needed
- Is there a better way?
  - Learning from their interactions with the environment, other agents, including humans.
  - Learning from examples, environmental reward/punishment, advice, imitation, active demonstration, among other means

Learning agents

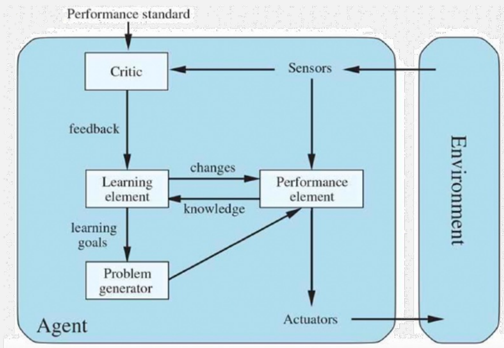

- Agents need knowledge (e.g., model of the environment)
- Where does this knowledge come from? – Learning

Learning agents

- Learning element improves performance
- Performance element realizes agent programs
- Problem generator suggests actions leading to informative experiences

Learning agents

- One of the simplest forms of learning is Skinnerian conditioning
  - A reflex agent randomly tries out different actions from its repertoire that match the percept
  - The environment provides reward / punishment
  - Eventually the agent figures out the optimal action to perform for each percept
- A major limitation of learning agents that learn from Skinnerian conditioning is that their success depends on being lucky enough to try out actions that are not dangerously bad.

Learning agents

- A major limitation of learning agents that learn from Skinnerian conditioning is that their success depends on being lucky enough to try out actions that are not dangerously bad
- Example:
    - A self-driving car that randomly tries out stepping on the accelerator as it approaches a red light.
    - The car will crash and get destroyed with no chance to learn from its experience.
- Skinnerian learning agents must start with agent programs that allow them to limit their choice of actions to those that are not downright dangerous

Learning Agents

- How can an agent limit its choice of actions to those that are not downright dangerous?
  - The agents could start out with some basic knowledge of the world and the ability to imagine and reason about the effects of its actions.
    - Example: The vacuum cleaner agent could know by design that it is dangerous to go into a flooded room because of the risk of dangerous short-circuit.
  - The agents could use virtual worlds for learning to limit their action choices
    - Driving simulator for training self-driving taxis
    - Flight simulator

Multi-agent Systems

- Inter-agent communication
- Multi-agent Interaction
  - Cooperation
  - Competition
- Multi-agent organizations