Artificial Intelligence

Fiona Yan Liu

Department of Computing

The Hong Kong Polytechnic University

Learning Outcomes of Lecture 01

- Introduction to COMP4431
  - Understand the motivation of the lecture teaching
    - Only focus on the discussion of the techniques that have been used widely in industry
  - Understand the motivation of the lab teaching
    - COMP4431 teaching teams spent a lot of time to design the content of the labs
    - TAs will teach the nine labs and one help session
    - Even more important than the lectures
    - Share the benefits brought by AI
  - Know the schedule of the lectures
    - Group demo is on Nov. 12 and the submission of the group project is on Nov. 7
  - Know the schedule of the labs
    - **CAN NOT** use any generative techniques in any assignment other than COMP4431
    - Please follow the instructions and the policy of the university strictly
- Feedback to the Quiz
  - Open book
  - Electronic device are allowed
  - The purpose of the quiz is to help you prepare the final exam
  - Final exam is close book
- History of artificial intelligence

Learning Outcomes of Lecture 01

- What is artificial intelligence
  - A science and a set of computational technologies that are inspired by—but typically operate quite differently from—the ways people use their neruo systems and bodies to sense, learn, reason, and take action
- Birth of artificial intelligence
  - Term coined by McCarthy in 1956 at Dartmouth Conference
  - Search technologies
  - Natural language processing
  - Computer vision
- Golden Time
  - The knowledge revolution
  - The rise of symbolicism that symbolic logic is claimed to explain the character of mind
- Winter
  - The difficult to reproduce general intelligence
- Struggle
  - The rise of machine learning
  - Support vector machines
- Revival
  - Deep learning
  - The success of connectionism

Agenda

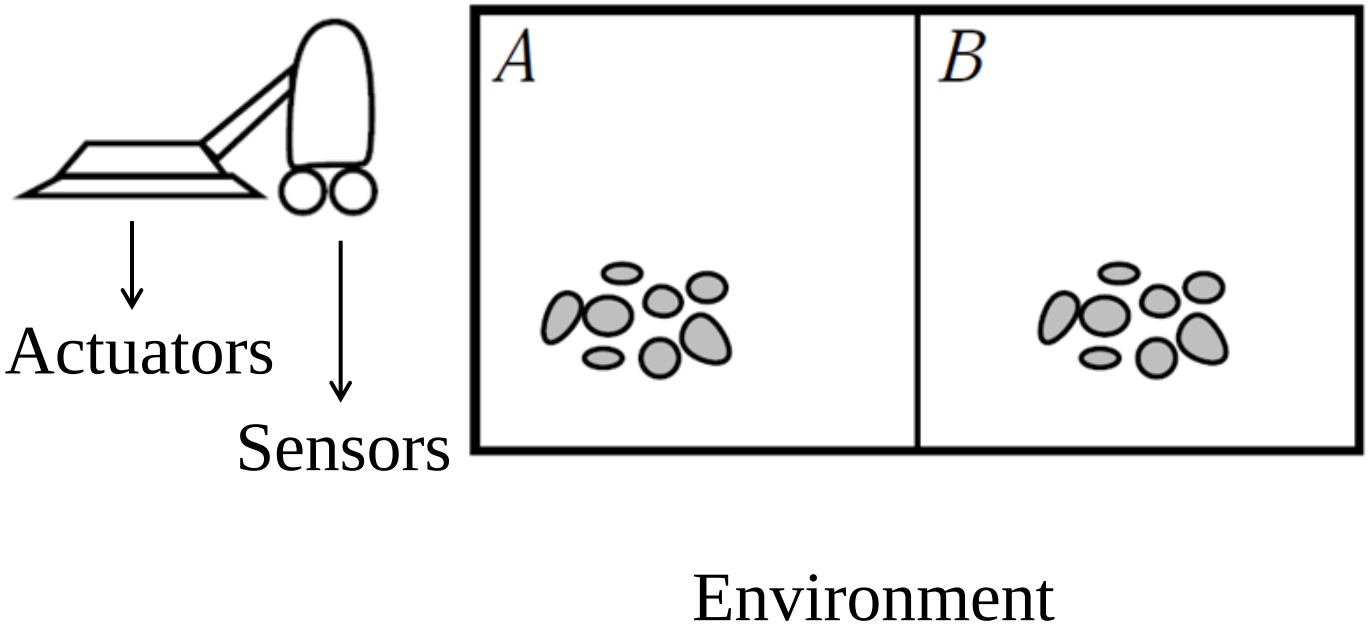
- Rational agent
- Search agent
  - Uninformed search and informed search
  - A* search
  - Heuristics
  - Domination of heuristic functions
- Technological trends of agent
  - Coze

Agenda

- Rational agent
- Search agent
  - Uninformed search and informed search
  - A* search
  - Heuristics
  - Domination of heuristic functions
- Technological trends of agent
  - Coze

Agent

- Agents in artificial intelligence
  - Agents denotes artificial entities that are capable of perceiving their surroundings using sensors, making decisions, and then taking actions in response using actuators
- Example
  - Vacuum-Cleaner
- How to perceives the environment
  - which room, clean or dirty
- How to make the decision
  - move right or left, suck the dust
- How to make the action



Rational Agent

- Rational Agent
  - An agent is an entity that perceives and acts
  - An agent function is to determine actions from percept histories:

    $[f: \mathcal{P}^* \rightarrow \mathcal{A}]$

  - For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Characteristics of rational agent
  - Distinct from omniscience
  - An agent is autonomous if its behavior is determined by its own experience
  - Agents can perform actions in order to modify future percepts so as to obtain useful information

Agenda

- Rational agent
- Search agent
  - Uninformed search and informed search
  - A* search
  - Heuristics
  - Domination of heuristic functions
- Technological trends of agent
  - Coze

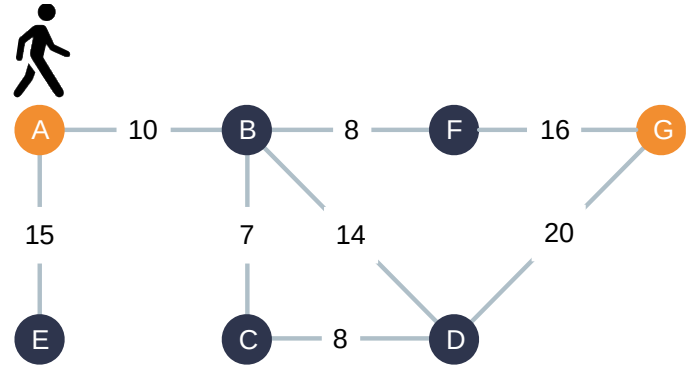
Scenario

- Students visit scientific park
  - Use intelligent information system to aid arrangement
- Real world problem
  - Students are in building A
  - Students want to go to building G
  - Find a path from building A to building G
  - Find a shortest path from building A to building G
  - Find a shortest path from any building to another building



Represent the Map using Graph

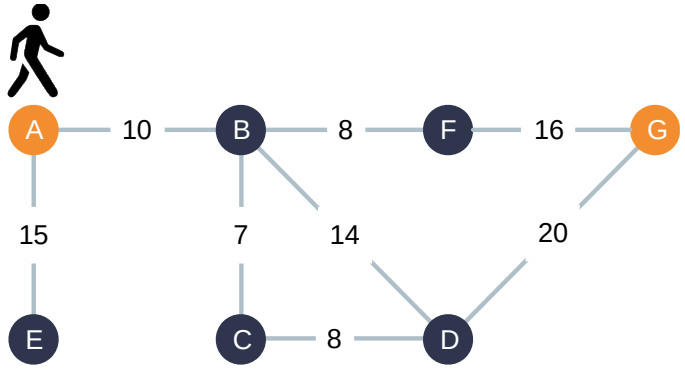
- Abstract the map to the graph
  - Ignore the detailed information
- Node
  - Location of the building
- Edge
  - The path connecting two buildings
- Value
  - The time walking from one building to another





Solve the Problem using Rational Agent

- How to perceive the environment
  - Which kind of information should be collected?
  - Which kind of information could be collected?
  - Which kind of information should be kept?
  - Which kind of information could be kept?
- How to make the action
  - Which kind of action could be made?
  - Which kind of action should be made?
- How to optimize the performance?
  - Could we find the solution?
  - Could we find the optimal solution?
- How to make the decision
  - Check the target
  - Check the cost

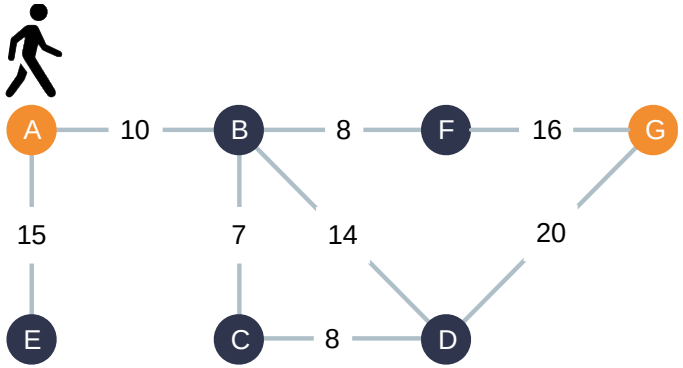


- An agent is an entity that perceives and acts
- An agent function is to determine actions from percept histories:

  $[f: \mathcal{P}^* \rightarrow \mathcal{A}]$
- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

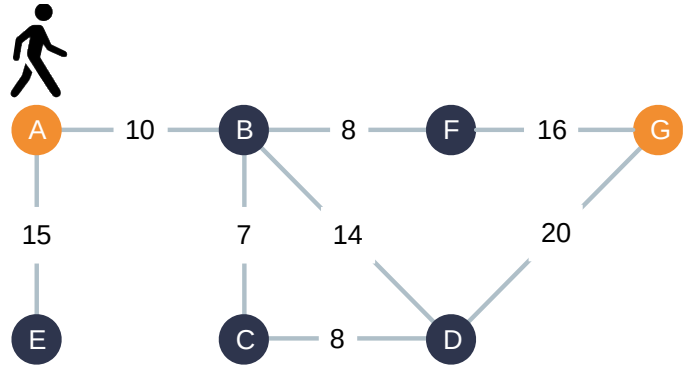
Problem Formulation

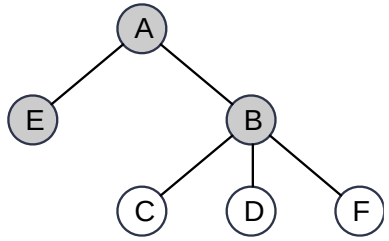
- Perception
  - State: In(B)
  - Initial state: In(A)
  - Goal state: In(G)
- Action
  - Decision result Go(B)
  - Solution: {Go(B),Go(F),Go(G)}
- Transition model
  - Action leads to the state transtition
  - Result (In(A), Go(B)) = In(B)
- Cost
  - If optimize the performance
  - c(In(A), {Go(B)}, In(B)) = 10
  - The sum of the costs of the individual actions along the path
  - c(In(A), {Go(B), Go(F)}, In(F)) = 10 + 8 = 18



Any Existing Methods to Model the Problem

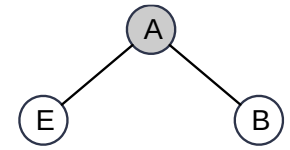
- Solution
  - An action sequence
- Search algorithm
  - Considering various possible action sequences
- Search tree
  - Nodes: states in the state space
  - Root: initial state
  - Branches: actions
- Expanding the current state
  - Apply each legal action to the current state, thereby generating a new set of states
  - Add branches from the parent node leading to child nodes
- Essence of search
  - Following up one option now and putting the others aside for later
  - In case the first choice does not lead to a solution

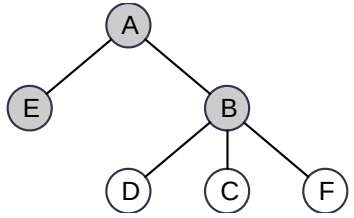


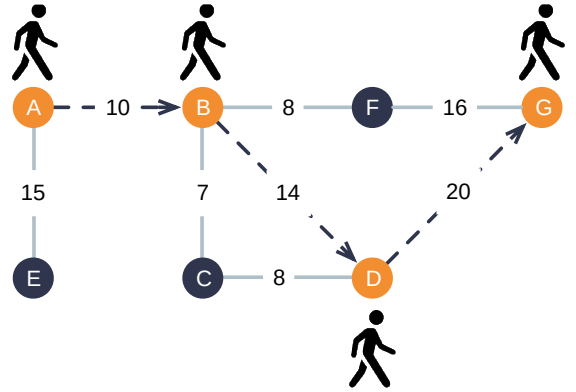


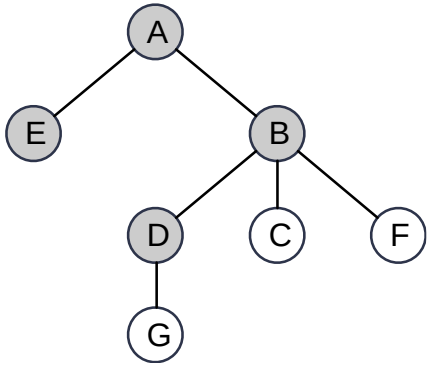
Breadth-first Search

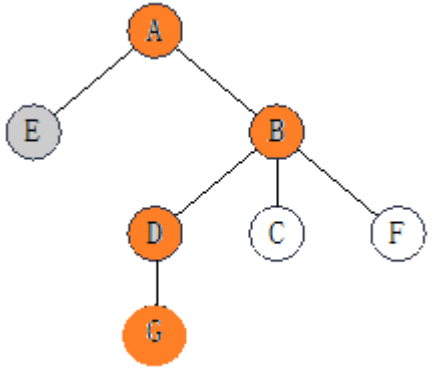
- Expand the shallowest unexpanded node
  - Cost = 10+14+20=44





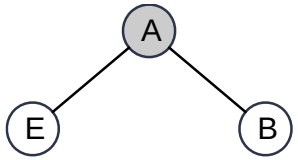


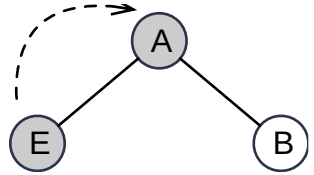


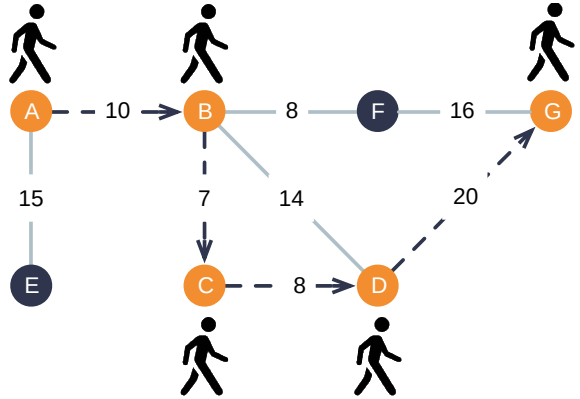


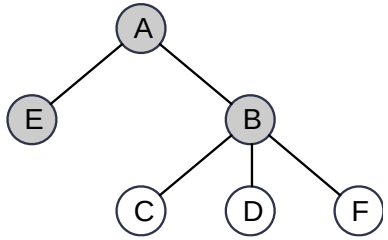
Depth-first Search

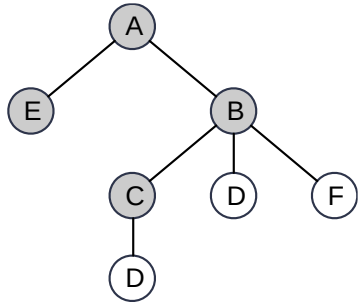
- Expand the deepest unexpanded node
  - Cost = 10+7+8+20=45

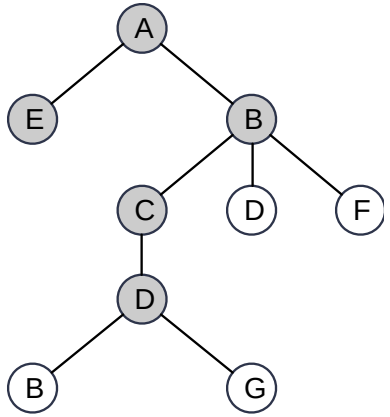


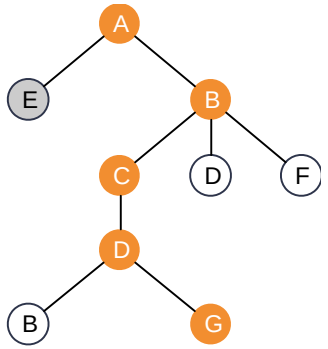






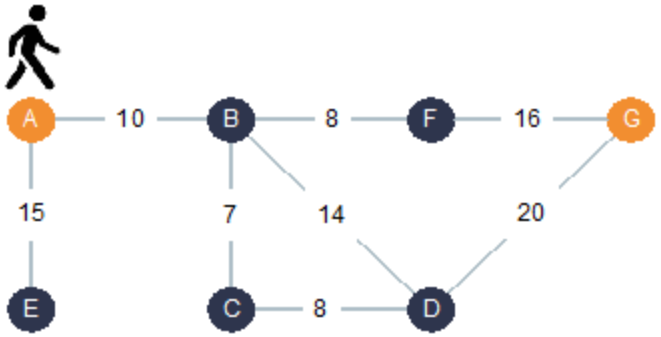






Uninformed Search and Informed Search

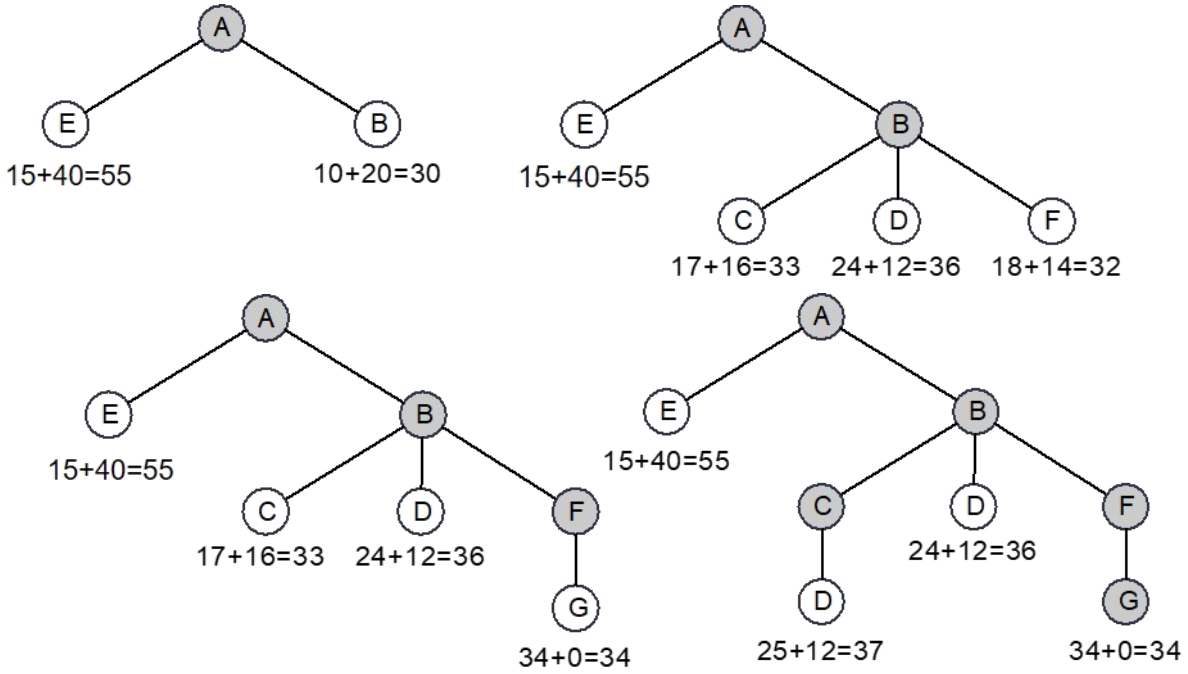
- Uninformed search
  - The strategies have no additional information about states beyond that provided in the problem definition
- Informed search
  - The strategies know whether one non-goal node is “more promising” than another
  - Heuristic information represents the relative suitability of a solution among its peers based on the intuition, experience or common sense

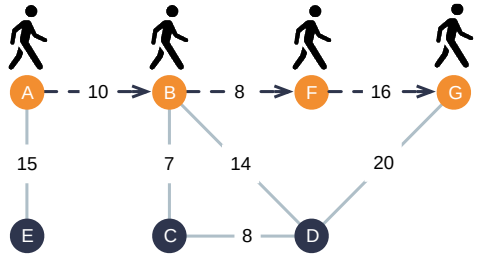


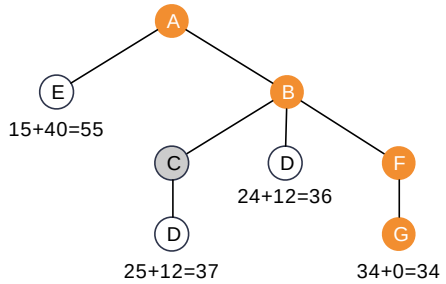
A	-	G :	28
B	-	G :	20
C	-	G :	16
D	-	G :	12
E	-	G :	40
F	-	G :	14
G	-	G :	0

A* Search

- Evaluation function: f(n) = g(n) + h(n).
  - g(n): the cost so far to reach the node n
  - h(n): the estimated cost to goal from the node n
  - f(n): the estimated total cost of path through the node n to goal



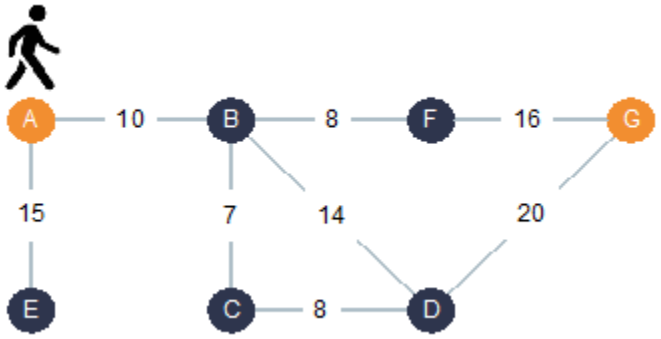




A - G : 28
B - G : 20
C - G : 16
D - G : 12
E - G : 40
F - G : 14
G - G : 0

Heuristics

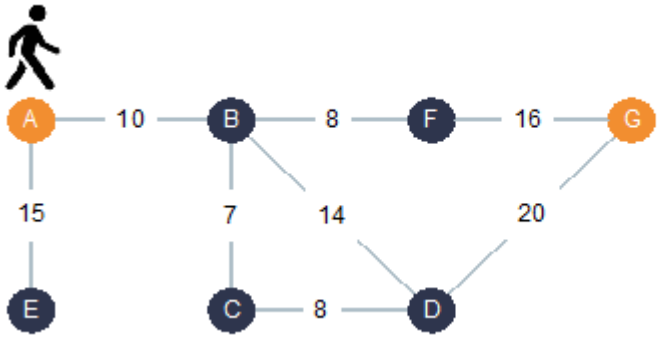
- Heuristics
  - A problem-solving method
  - Uses shortcuts to produce good-enough solutions given a limited time
  - Solution may not necessarily be optimal solution
- Heuristic function in search algorithms
  - A function that ranks alternatives based on available information
- An admissible heuristic never overestimates the cost to reach the goal
  - A heuristic h(n) is admissible if for every node n, h(n) ≤ h*(n), where h*(n) is the true cost to reach the goal state from n

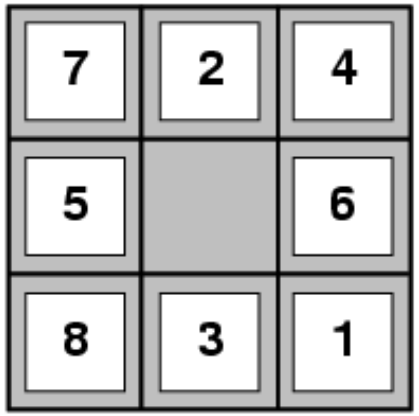


A	-	G	: 28
B	-	G	: 20
C	-	G	: 16
D	-	G	: 12
E	-	G	: 40
F	-	G	: 14
G	-	G	: 0

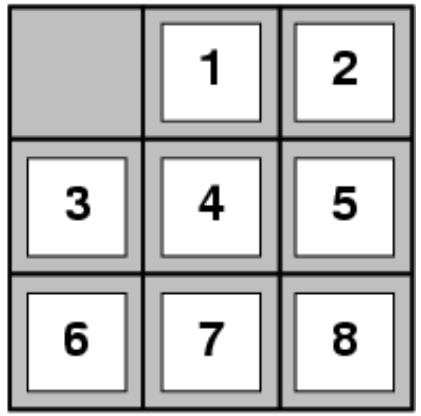
Domination of Heuristic Functions

- If $h_2(n) \geq h_1(n)$ for all n (both admissible)
  - then h_2 dominates h_1, which indicates that h_2 is better for search
- Domination translates directly into efficiency
  - A* using h_2 will never expand more nodes than A* using h_1
- $h_2(S)$ and $h_1(S)$ for all both admissible
  - $h_1(S) = 8$
  - $h_2(S) = 18$
- A* using h_2 will never expand more nodes than A* using h_1
  - $A^*(h_1)$ = 39,135 nodes
  - $A^*(h_2)$ = 1,641 nodes





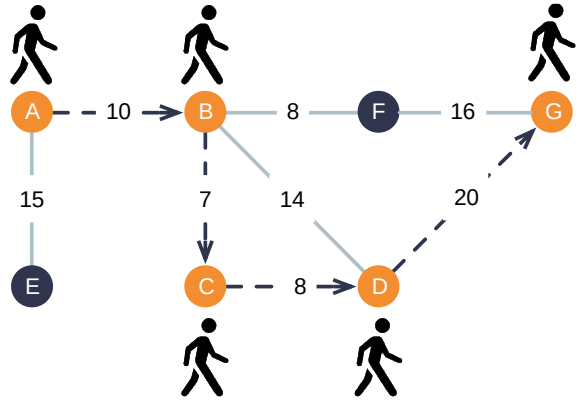
Start State



Goal State

Search Strategy

- Search strategy
  - Pick the order of node expansion
- Strategies are evaluated along the following dimensions:
  - completeness: does it always find a solution if one exists?
  - time complexity: number of nodes generated
  - space complexity: maximum number of nodes in memory
  - optimality: does it always find a least-cost solution?
- Informed search algorithms
  - Breadth-first search
  - Depth-first search
- Uninformed search algorithms
  - A* search

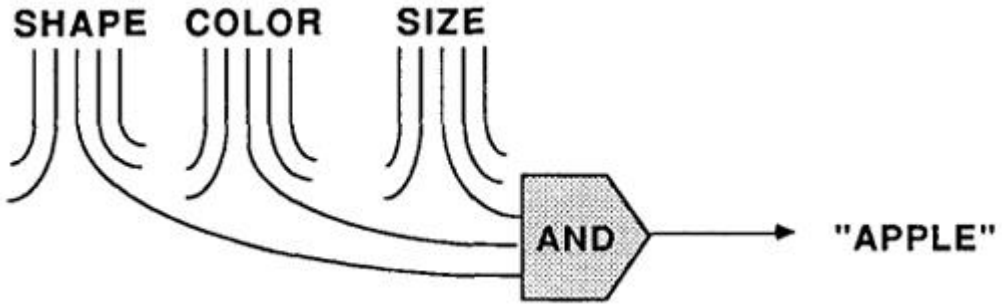


Agenda

- Rational agent
- Search agent
  - Uninformed search and informed search
  - A* search
  - Heuristics
  - Domination of heuristic functions
- Technological trends of agent
  - Coze

Technological Trends in Agent Research

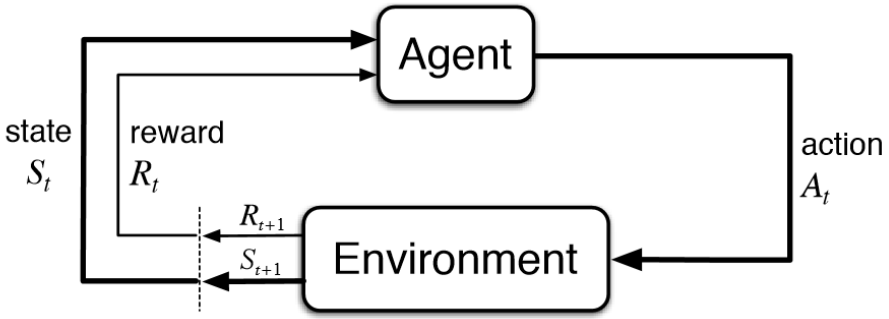
- The evolution of AI agents has undergone several stages
  - Symbolic agent stage: symbolic agent approaches employ logical rules and symbolic representations to encapsulate knowledge and facilitate reasoning processes.
  - A classic example of this approach is knowledge-based expert systems.



Minsky M. Society of mind[M]. Simon and Schuster, 1988.

Technological Trends in Agent Research

- The evolution of AI agents has undergone several stages
  - Reinforcement learning-based agent stage: reinforcement learning-based agents utilize reinforcement learning methods to train agents to learn through interactions with their environments, enabling them to achieve maximum cumulative rewards in specific tasks.

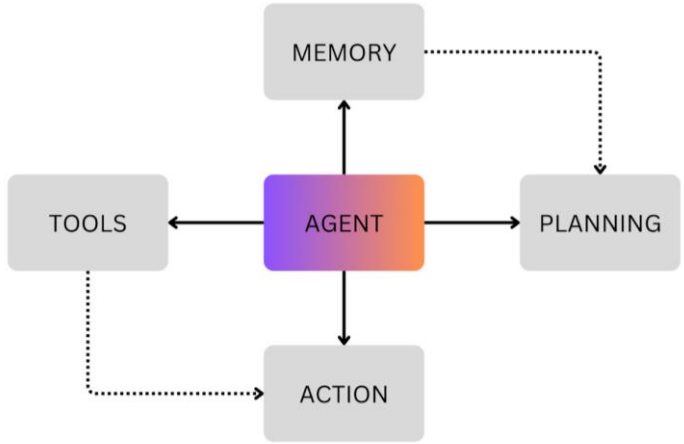


- The agent-environment interaction in reinforcement learning
  - An agent (decision maker) takes the action A_t in the given state s_t and receives the reward R_t.
  - After the reward is received, the agent update their parameters such that the total maximizing discounted expected reward.

Sutton R S, Barto A G. Reinforcement learning: An introduction[M]. MIT press, 2018.

Technological Trends in Agent Research

- The evolution of AI agents has undergone several stages
  - Large language model-based agent stage: employ large language model as the primary component of brain or controller of agents.
    - LLMs can exhibit reasoning and planning abilities comparable to symbolic agents
    - LLMs can acquire interactive capabilities with the environment, by learning with feedback
    - LLMs undergo pre-training and demonstrate the capacity for few- and zero-shot generalization
  - Several key components in LLM
    - Memory: The primary function of the memory system is to preserve and regulate knowledge, experiential data, and historical information.
    - Planning: The planning capability defines an LLM-based agent's ability to devise action sequences based on set objectives and existing environment constraints, ensuring goal fulfillment.
    - Tool use: encompassing APIs, calculators, code interpreters, or actions within a physical environment through text-based directives.



https://lilianweng.github.io/posts/2023-06-23-agent/

Coze

- Coze
  - Coze is a one-stop AI chat bot development platform. Regardless of your programming skills,
  - Coze allows you to swiftly create a wide range of chat bots powered by AI, capable of handling everything from simple questions to complex conversations.
- Coze facilitates and simplifies the process of AI bots development
  - Knowledge
    - Coze supports adding data from local files in the format of TXT, PDF, DOCX, Excel, and CXV, or obtaining online content and API JSON data via the specified URL.
  - Workflow
    - Coze empowers users to process intricate task flows with a set of workflow nodes such as LLMs, codes, and conditions. Regardless of your coding expertise, its intuitive drag-and-drop interface facilitates workflow creation
  - Plugins
    - Coze provides an extensive collection of plugins, enhances bot's functionality in areas such as information retrieval, travel assistance, productivity, image comprehension, and so on.

Arrangement of Following Classes

- Classes and labs on Sept. 17
  - Be cancelled due to the public holiday
  - I will be in the classroom of TU 103 on Sept. 17 3:30pm – 5:30pm
    - Additional office hour
    - If you have any question, you may come to the lecture room
- Class on Sept. 24
  - Quiz 2 covers lecture 2 page 5 – page 20
  - Constrains and games
- Lab on Sept. 24
  - AI-powered image generation
- Happy Mid-Autumn Festival !