Artificial Intelligence

Fiona Yan Liu

Department of Computing

The Hong Kong Polytechnic University

Learning Outputs of Lecture 01

- Artificial Intelligence
  - Is concerned with the design of intelligence in an artificial device
- Acting humanly
  - Can machines behave intelligently as the human
  - The Turing test
- Thinking humanly
  - Scientific theories of internal activities of the brain
  - Cognitive science and Neuroscience
- Thinking rationally
  - Correct argument/thought processes
  - Little widely accepted conclusion has been made
- Acting rationally
  - Doing the right thing
  - Which is expected to maximize goal achievement, given the available information

The History of Artificial Intelligence

- Artificial Intelligence
  - Term coined by McCarthy in 1956
- 1956 – 1974
  - Search technology
  - Natural language processing
  - Computer vision
- 1980 – 1987
  - Artificial neural networks
  - Expert systems
  - Industry robots
- 1993 – 2003
  - Support vector machine
  - Machine learning
  - Automatic cars
- 2012 –

Feedback of Questionnaire

- Why choose this course?
  - Be interested in this course
  - Useful for further study
- Expectation of this course
  - Learn solid theory of artificial intelligence
  - Learn some useful skills of developing intelligent system
  - More chance to discuss and practice
- Please choose the course you have taken
  - Data structure
- Please choose the techniques you want to learned
  - Various requirements
- How many courses do you have this semester
  - 5 – 6
  - I have final year project
  - My final year project is related with artificial intelligence

General Information

- Course web page
  - http://www.comp.polyu.edu.hk/~csyliu/course/comp406/main.html
- Text book
  - Stuart Russell and Peter Norvig, "Artificial Intelligence A Modern Approach"
- Lecture of our class
  - Tue. 15:30 – 17:20
  - Y410
  - Contact with Fiona csyliu@comp.polyu.edu.hk
- Lab of our class
  - Tue. 17:30 – 18:20
  - QT402
  - Contact with Songtao Wu csstwu@comp.polyu.edu.hk

Course Presentation

- Introduction to a movie related with artificial intelligence
  - Group work with 1 – 4 person(s) each group
  - Every group should work on different movies
- 10 minutes presentation
  - No requirement of report
- Send email to TA including the following information
  - Group member name and student ID
  - Movie name
  - Before Oct. 13 2015
- Presentation date
  - Oct. 27 2015 at 15: 30
  - The confirmed presentation order is announced on Oct. 20
- 100 points and 10% for the final grading
  - For detail information, check notes of lecture 01

A Room Clean Job

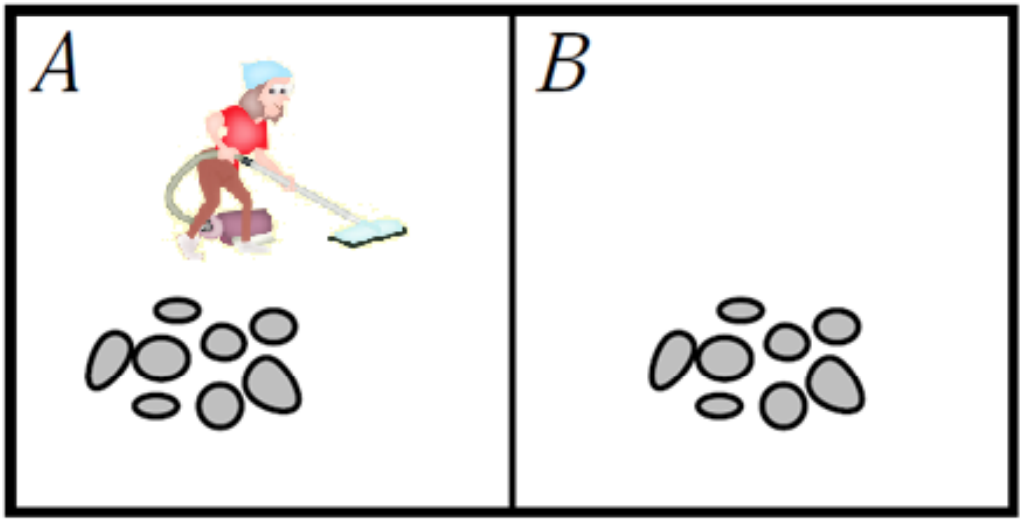

- The human perceives the environment
  - which room, clean or dirty
- Decides what to do
  - move right or left, suck the dust
- And then acts

A Vacuum-Cleaner Agent

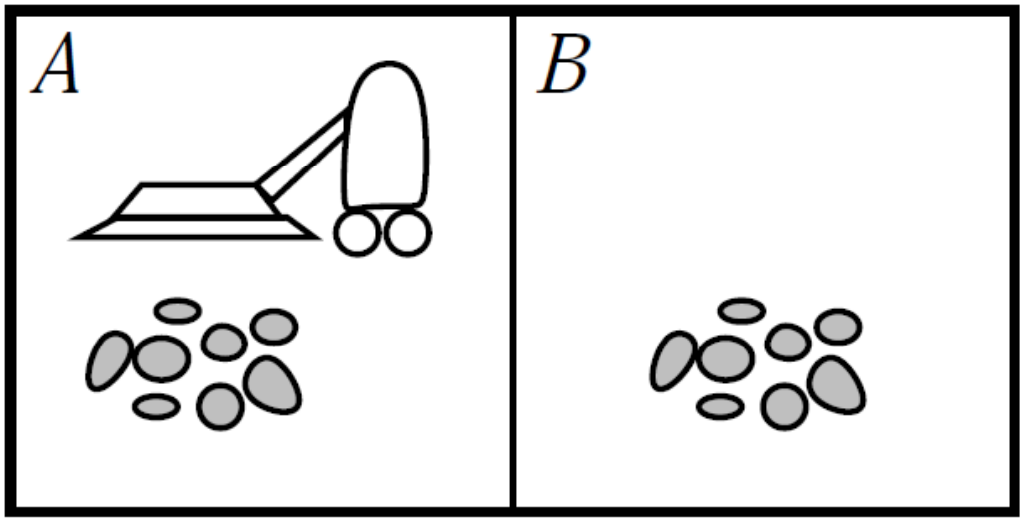

- How to perceives the environment
  - which room, clean or dirty
- How to make the decision
  - move right or left, suck the dust
- How to make the action

Agents

- An agent is an entity that perceives and acts
  - Perceiving its environment through sensors
  - Acting upon that environment through actuators

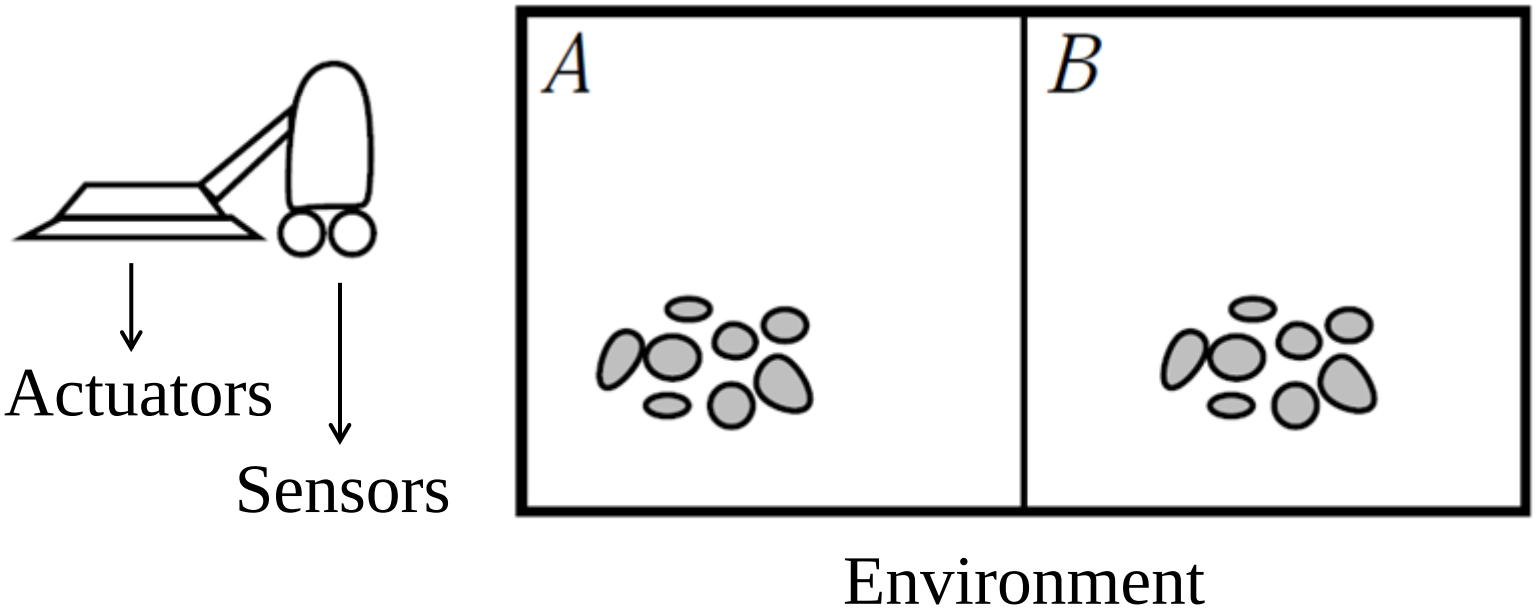

Rational Agent

- 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
  - Agents can perform actions in order to modify future percepts so as to obtain useful information
  - An agent is autonomous if its behavior is determined by its own experience

PEAS of Intelligent Agent Design

- Intelligent agent design
  - Under the assumption that there exists rational agent
  - Aim: find a way to implement the rational agent
- PEAS must first specify the setting for intelligent agent design
  - Performance measure
  - Environment
  - Actuators
  - Sensors
- Design an automated taxi driver
  - Performance measure
    - Safe, fast, legal, comfortable trip, maximize profits
  - Environment
    - Roads, other traffic, customers
  - Actuators
    - Steering wheel, accelerator, brake, signal, horn
  - Sensors
    - Cameras, sonar, speedometer, GPS

Intelligent Agent Types

- Simple reflex agents
- Model-based reflex agents
- Goal-based agents
- Utility-based agents

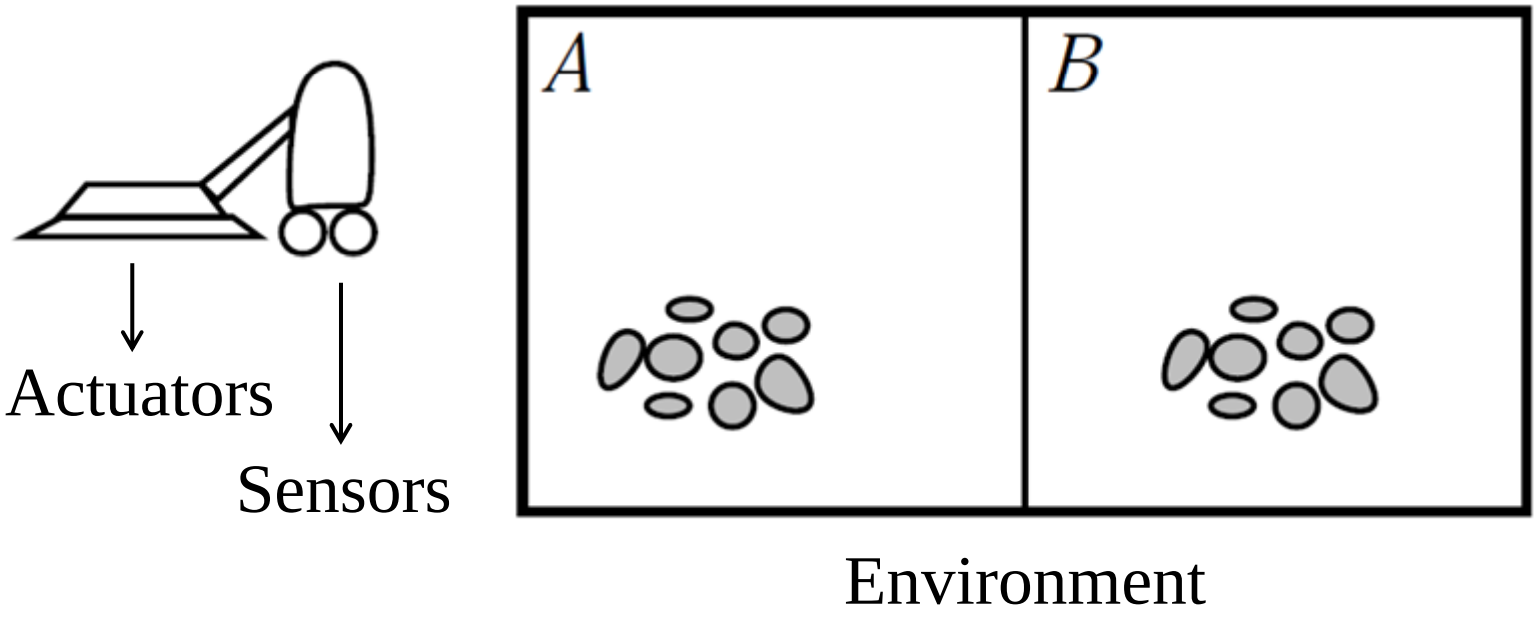

Simple Reflex Agents

- Select actions only based on current perception
  - Infinite loops are often unavoidable

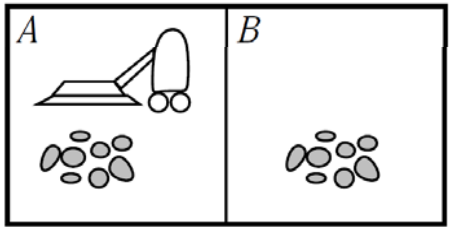

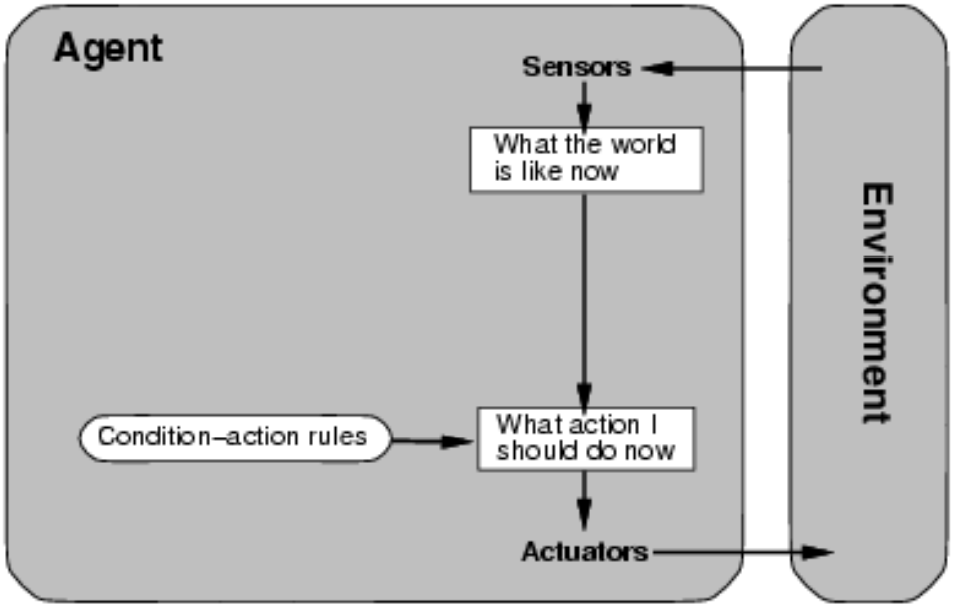

Model-based Reflex Agents

- Maintain internal states
  - A model of the world

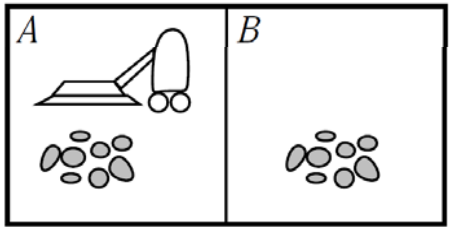

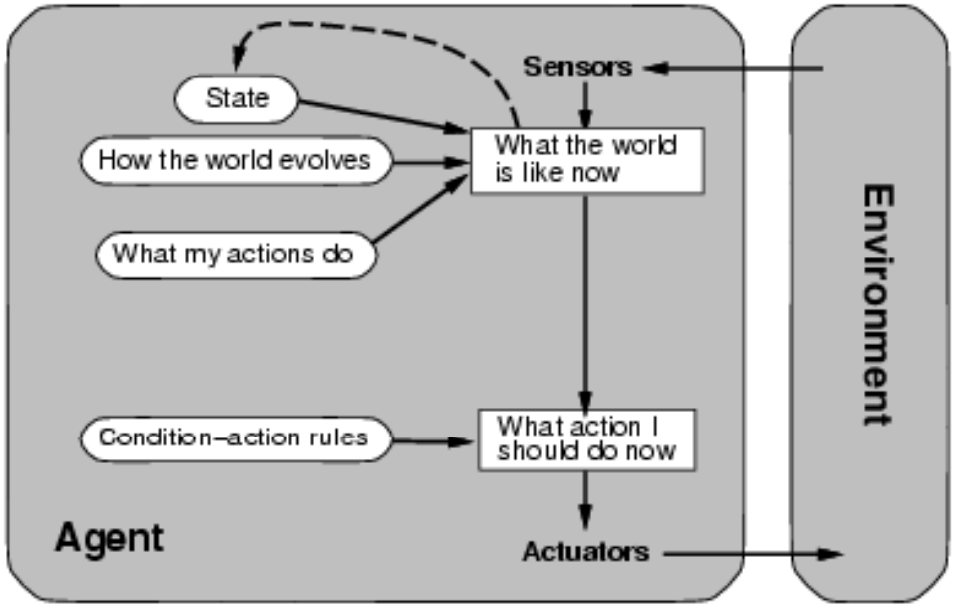

Goal-based Agents

- Future is taken into account
  - The agent can choose one from among multiple possible solutions

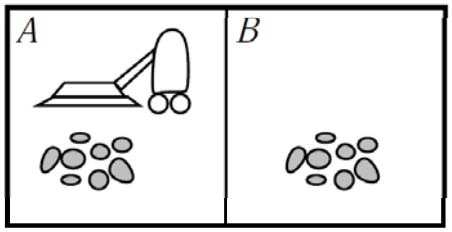

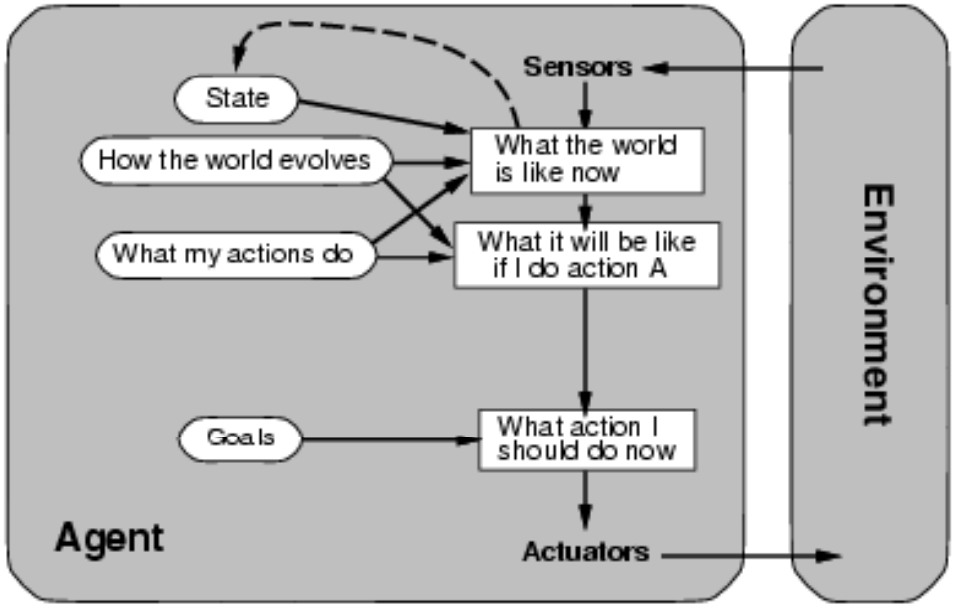

Utility-based Agents

- Cost will be considered to achieve the goal
  - If save the cost is a kind of goal, it can be classified to a kind of goal-based agent

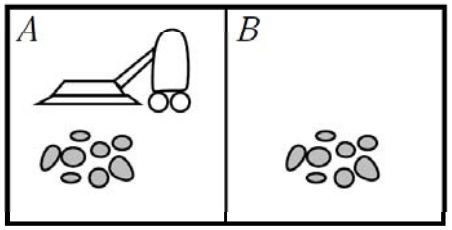

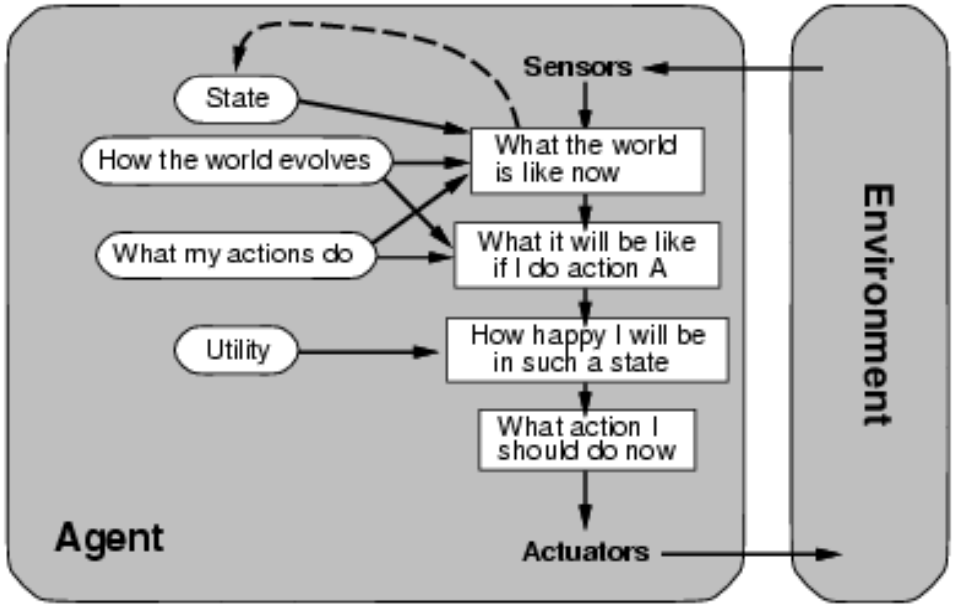

Example of An Intelligent Agent

- Drive from Arad to Bucharest
- Find the best way

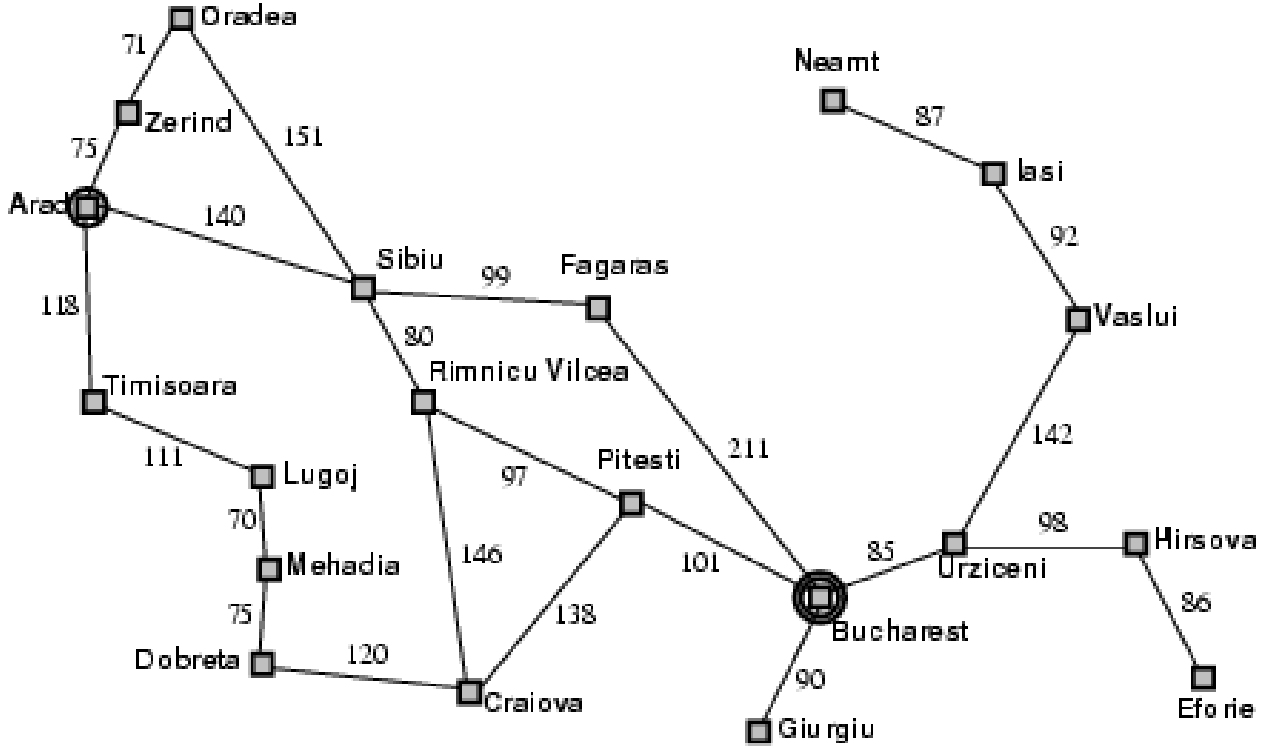

Solve the Problem by Search Agent

- Formulate the perception
  - Initial state: In(Arad)
- Formulate goal
  - Be in Bucharest
  - In(Bucharest)
- Formulate the environment
  - States: various cities with known distances
  - Path cost: The sum of the costs of the individual actions along the path
- Formulate the actions
  - Drive between cities
  - Actions: Go(Sibiu), Go(Timisoara), Go(Zerind)
- Find solution
  - Sequence of cities
  - Transition model: RESULT(In(Arad), Go(Sibiu)) = In(Sibiu)

Searching for Solution

- 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

Example of Tree Search

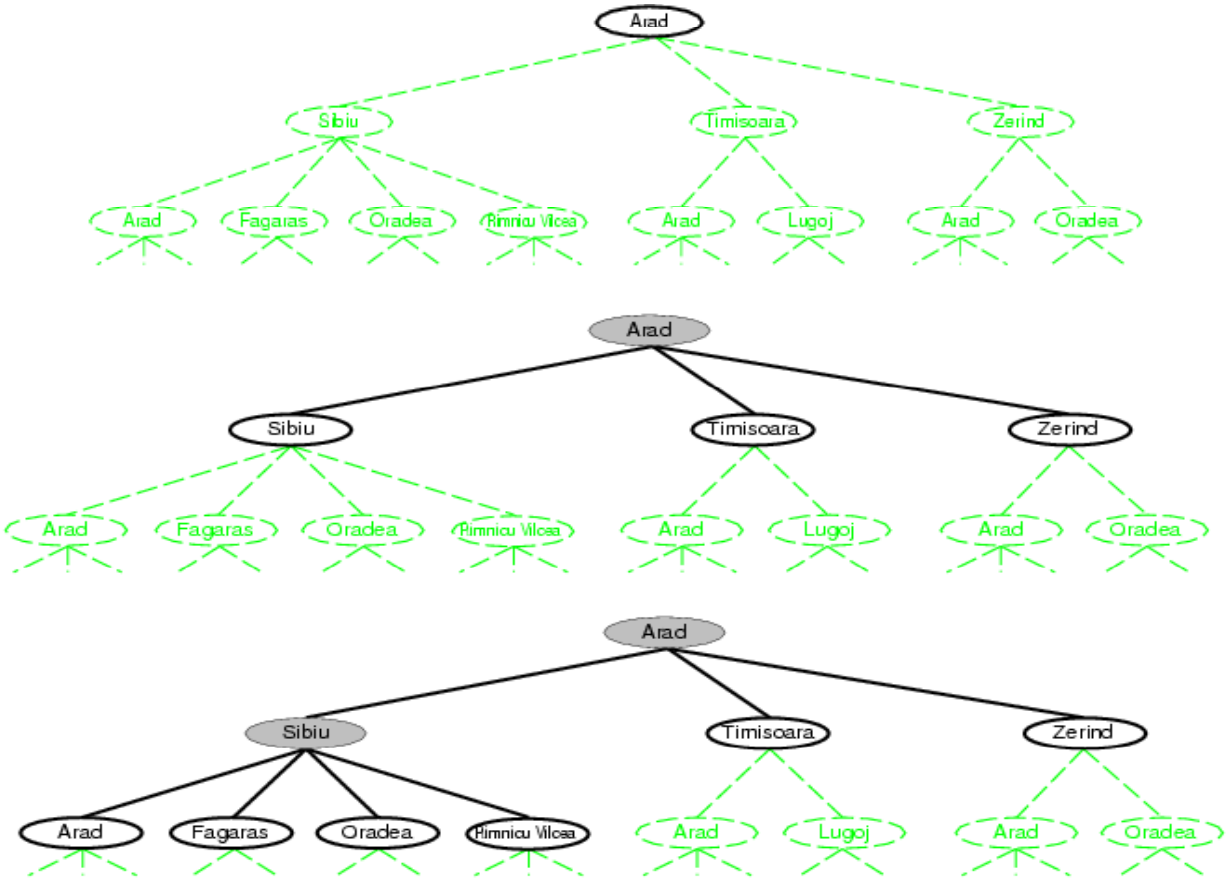

Loopy Path

- Redundant Paths
  - Loopy Path

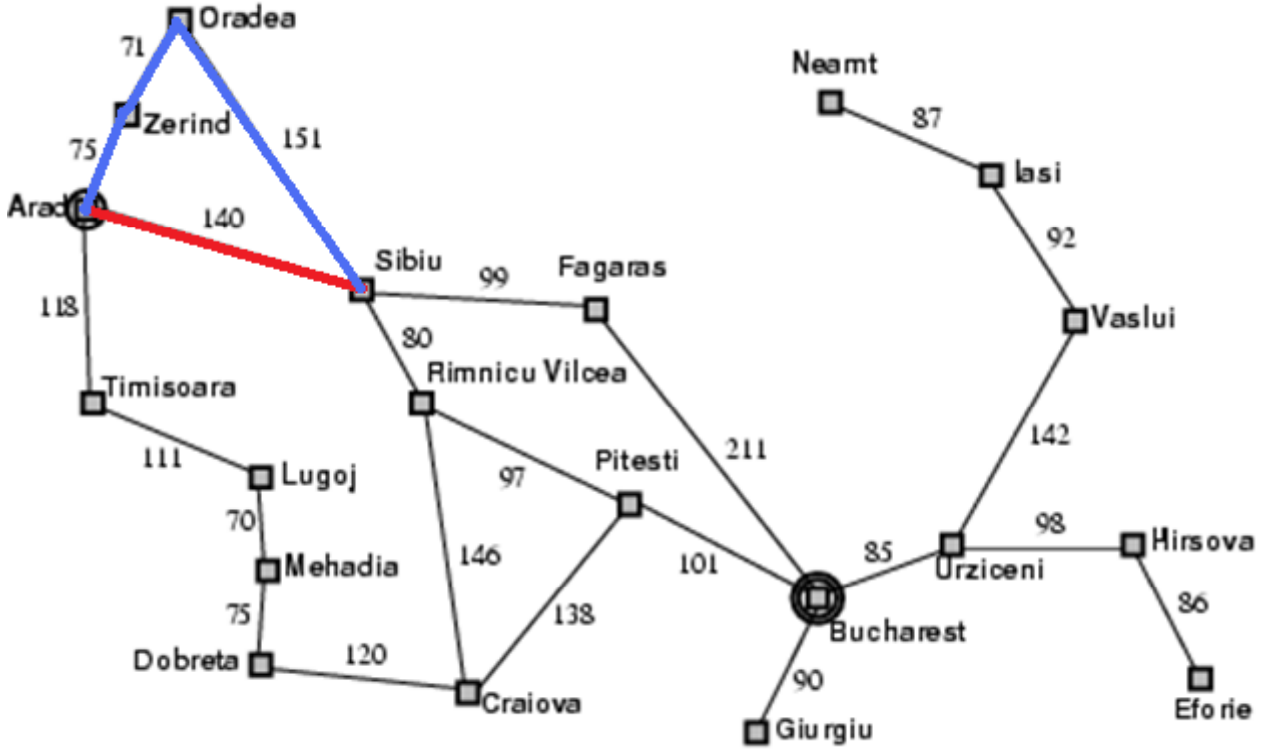

Search Strategy

- 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?
- Time and space complexity are measured in terms of
  - *b:* maximum branching factor of the search tree
  - *d:* depth of the least-cost solution
  - *m*: maximum depth of the state space (may be ∞)

Uninformed Search

- Breadth-first search
  - Expand the shallowest unexpanded node
- Depth-first search
  - Expand the deepest unexpanded node
- Uniform-cost search
  - Expand least-cost unexpanded node
- Uninformed search
  - Also called blind search
  - The strategies have no additional information about states beyond that provided in the problem definition

Informed Search

Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

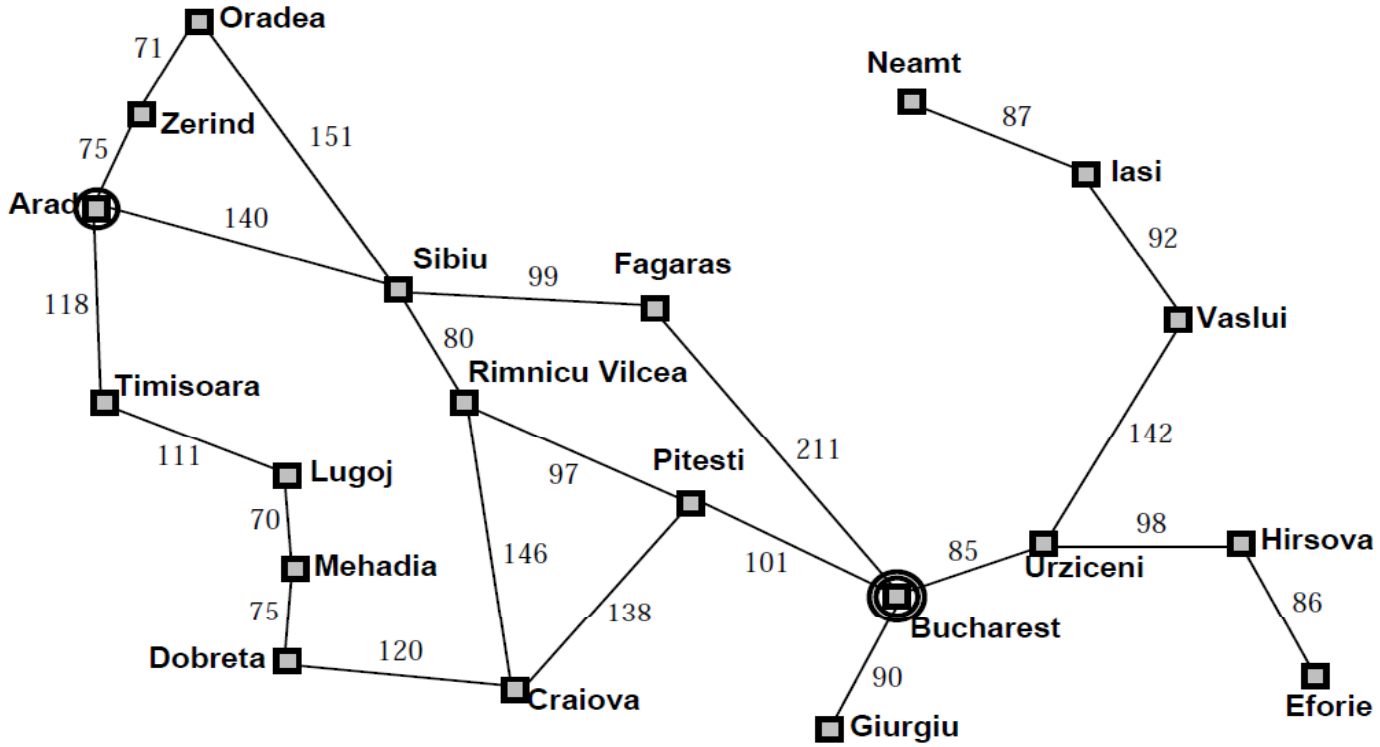

Informed Search

- Informed search
  - Also called heuristic search
  - The strategies know whether one non-goal node is “more promising” than another
  - The “desirability” of a node is estimated by a evaluation function *f(n)*
- Greedy best-first search
  - expands the node that is closest to the goal, which is evaluated by the heuristic function $h(n)$
- A* search
  - Avoid expanding paths that are already expensive
  - 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