CAP 4630
Artificial Intelligence

Instructor: Sam Ganzfried
sganzfri@cis.fiu.edu

Automated music generation

- https://www.youtube.com/watch?v=0VTI1BBLydE

Proofs

- Square root of 2 is irrational

Proofs

- Infinitely many primes

Proofs

- x is odd iff x+3 is even

Proofs

- Sum of integers from 1 to n

What is AI?

Thinking Humanly	Thinking Rationally
"The exciting new effort to make computers think … *machines with minds*, in the full and literal sense." (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning …" (Bellman, 1978)	"The study of mental faculties through the use of computational models." (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Acting Humanly	**Acting Rationally**
"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)	"Computational Intelligence is the study of the design of intelligent agents." (Poole *et al.*, 1998) "AI …is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

Figure 1.1 Some definitions of artificial intelligence, organized into four categories.

volving observations

Acting humanly

- Turing test: https://www.youtube.com/watch?v=sXx-PpEBR7k
- Russian “chatterbot” Eugene Goostman passes it in 2014

  https://www.youtube.com/watch?v=OcY8xlPffcA

  https://www.youtube.com/watch?v=KNxALt_7F2k

- Eugene Goostman is portrayed as being a 13-year-old boy from Odessa, Ukraine, who has a pet guinea pig and a father who is a gynaecologist. Veselov stated that Goostman was designed to be a "character with a believable personality". The choice of age was intentional, as, in Veselov's opinion, a thirteen-year-old is "not too old to know everything and not too young to know nothing". Goostman's young age also induces people who "converse" with him to forgive minor grammatical errors in his responses.

Acting humanly

- **Natural language processing** to communicate effectively in English
- **Knowledge representation** to store what it knows or hears
- **Automated reasoning** to use the stored information to answer questions and to draw new conclusions
- **Machine learning** to adapt to new circumstances and to detect and extrapolate patterns.

Acting humanly

- **Total Turing Test** includes video signal so the interrogator can test the subject's perceptual abilities, as well as the opportunity for the interrogator to pass physical objects "through the hatch." To pass the Total Turing Test, the computer will need:

- **Computer vision** to perceive objects
- **Robotics** to manipulate objects and move about

Thinking humanly

- If we are going to say that a given program thinks like a human, we must have some way of determining how humans think. We need to get *inside* the actual workings of human minds. There are 3 ways to do this:
  - Through introspection: trying to catch our own thoughts as they go by
  - Through psychological experiments: observing a person in action
  - Through brain imaging: observing the brain in action
- The interdisciplinary field of **cognitive science** brings together computer models from AI and experimental techniques from psychology to construct precise and testable theories of the human mind

Thinking rationally

- Aristotle's **syllogisms**:
  - "Socrates is a man; all men are mortal; therefore, Socrates is mortal."
- Field of **logic**.
- Two obstacles:
  - It is not easy to take informal knowledge and state it in the formal terms required by logical notation, particularly when the knowledge is less than 100% certain.
  - There is a big difference between solving a problem "in principle" and solving it in practice. Even problems with just a few hundred facts can exhaust the computational resources of any computer unless it has some guidance as to which reasoning steps to try first.

Acting rationally

- A **rational agent** is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome.
- Two advantages over other approaches:
  - More general than "thinking rationally" approach because correct inference is just one of several possible mechanisms for achieving rationality.
  - More amenable to scientific development than are approaches based on human behavior or human thought. The standard of rationality is mathematically well defined and completely general.

Relation to other fields

- Theoretical computer science:
  - Emphasis on proving theorems
- Computer science systems and programming languages
  - Emphasis on engineering systems
- Operations research/optimization
- Probability and statistics
- Machine learning and "data science"
- Philosophy
- Mathematics
- Economics
- Neuroscience
- Psychology, computer engineering, control theory and cybernetics, linguistics, …

Wikipedia and my definitions of AI

- In computer science, the field of AI research defines itself as the study of “intelligent agents”: any device that perceives its environment and takes actions that maximize its chance of success at some goal.
- AI is about creating real agents for solving interesting/important (large-scale) problems.
  - Pragmatic programming and implementation issues are very important in building real agents.
  - However, ideally the agents are not just based on “hacks” or random engineering heuristics, and there is also some deeper fundamental theory that justifies their performance. But producing strong “theory” is not the end goal.

Solving problems by search

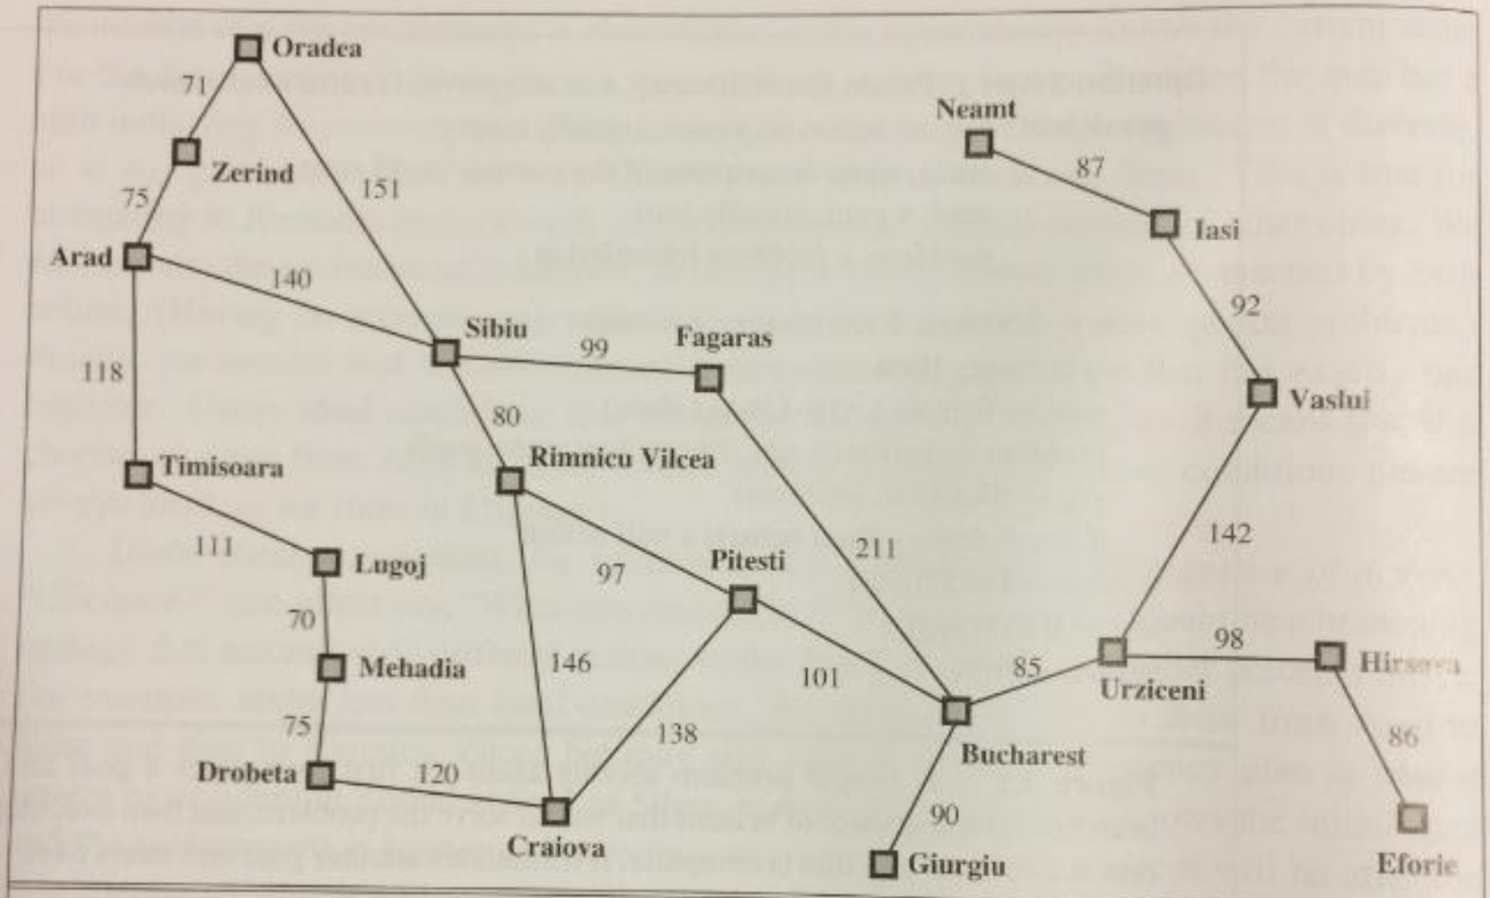

Figure 3.2 A simplified road map of part of Romania.

8-puzzle

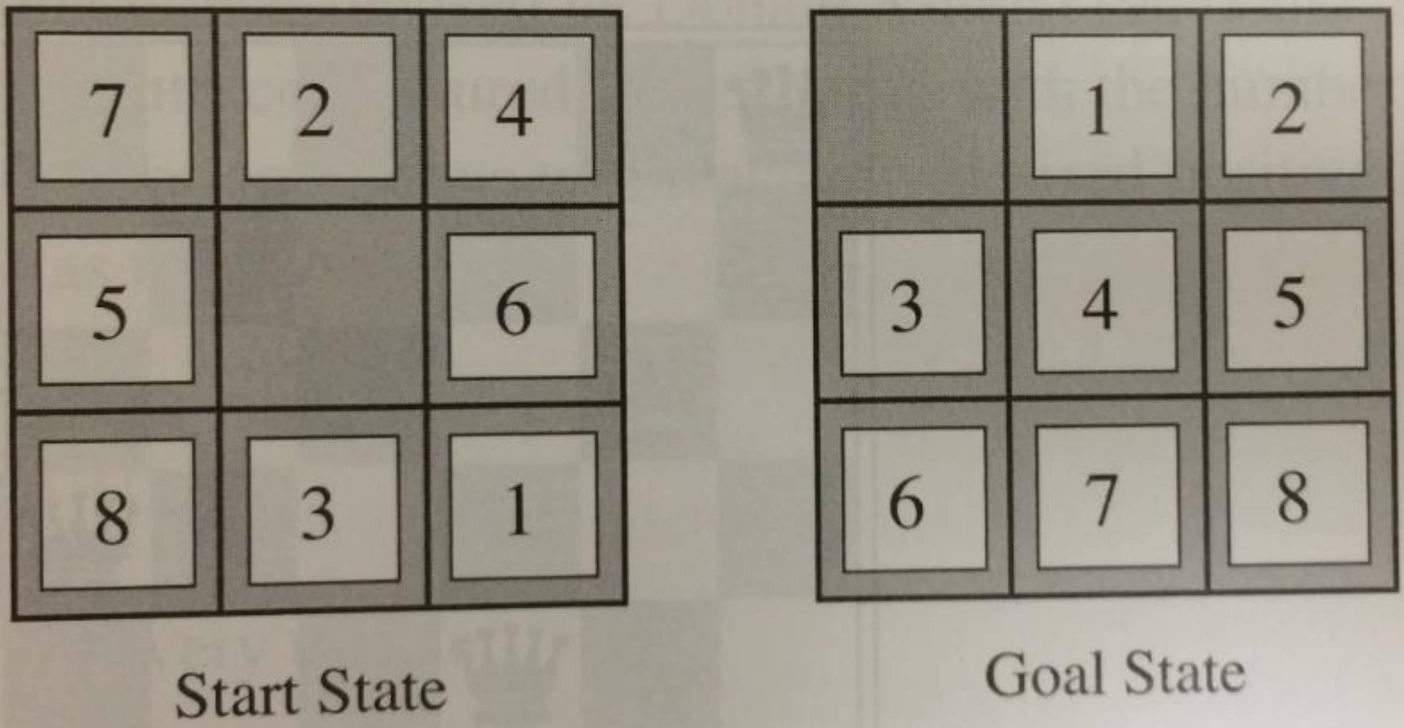

A typical instance of the 8-puzzle.

8-queens

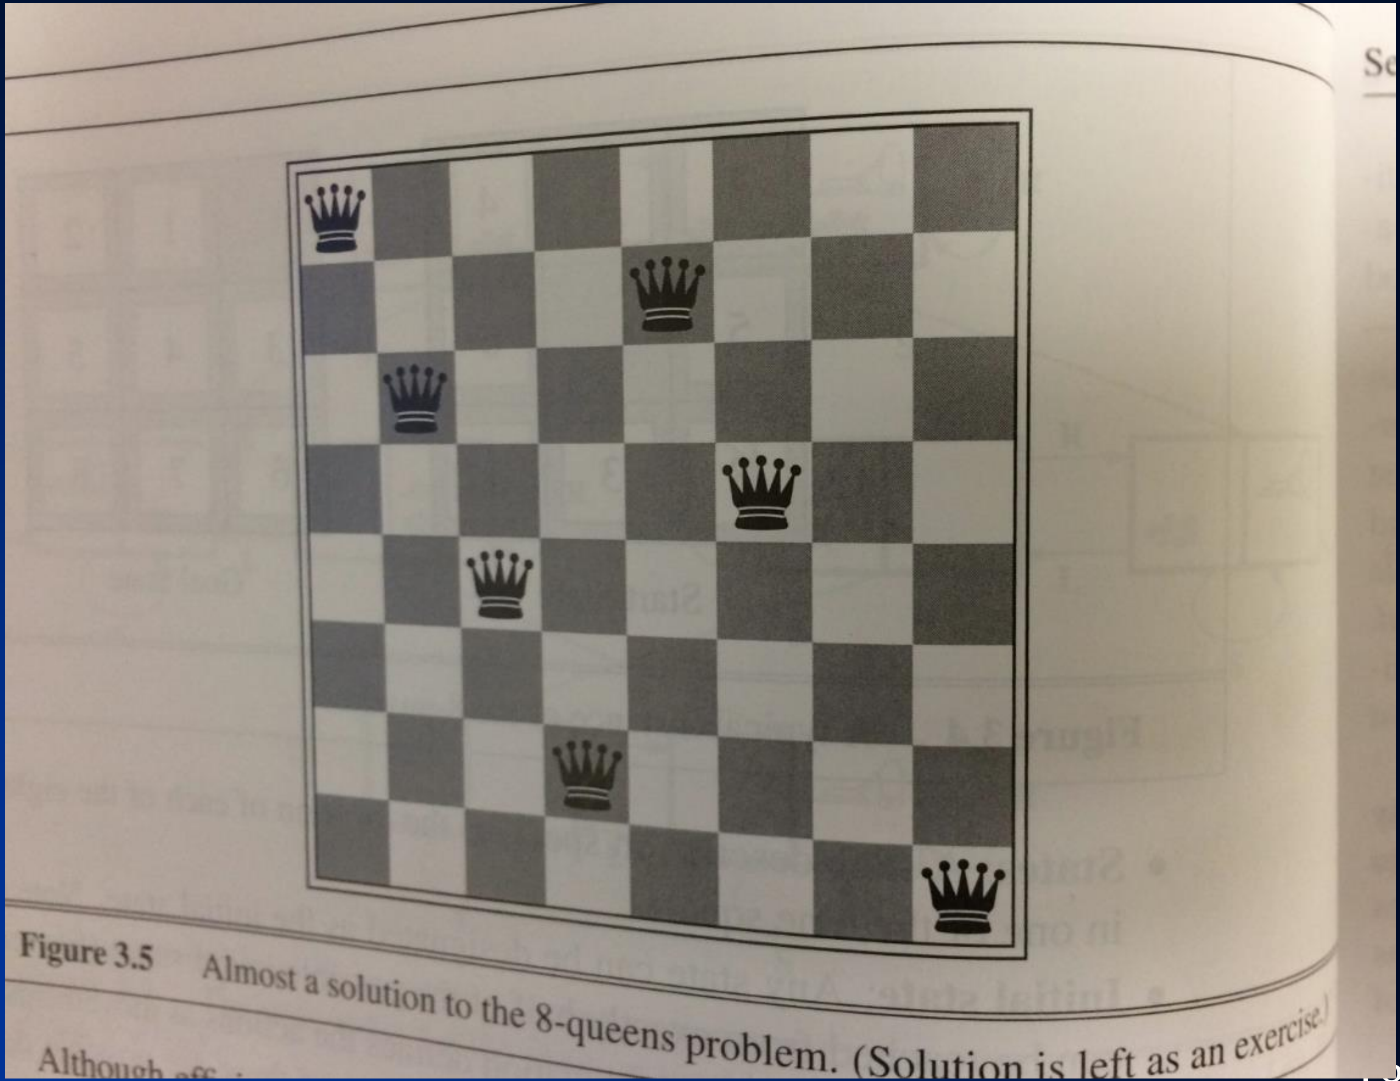
Figure 3.5 Almost a solution to the 8-queens problem. (Solution is left as an exercise.)

Search problem definition

- **States**
- **Initial state**
- **Actions**
- **Transition model**
- **Goal test**
- **Path cost**

Definition for 8-queens problem

- **States**: Any arrangement of 0 to 8 queens on the board is a state.
- **Initial state**: No queens on the board.
- **Actions**: Add a queen to any empty square.
- **Transition model**: Returns the board with a queen added to the specified square
- **Goal test**: 8 queens are on the board, none attacked
- **Path cost**: (Not applicable)

Problem-solving approach

```
function SIMPLE-PROBLEM-SOLVING-AGENT(percept) returns an action
  persistent: seq, an action sequence, initially empty
              state, some description of the current world state
              goal, a goal, initially null
              problem, a problem formulation

  state ← UPDATE-STATE(state, percept)
  if seq is empty then
     goal ← FORMULATE-GOAL(state)
     problem ← FORMULATE-PROBLEM(state, goal)
     seq ← SEARCH(problem)
     if seq = failure then return a null action
  action ← FIRST(seq)
  seq ← REST(seq)
  return action
```

Figure 3.1 A simple problem-solving agent. It first formulates a goal and a problem, searches for a sequence of actions that would solve the problem, and then executes the actions one at a time. When this is complete, it formulates another goal and starts over.

Searching Romania

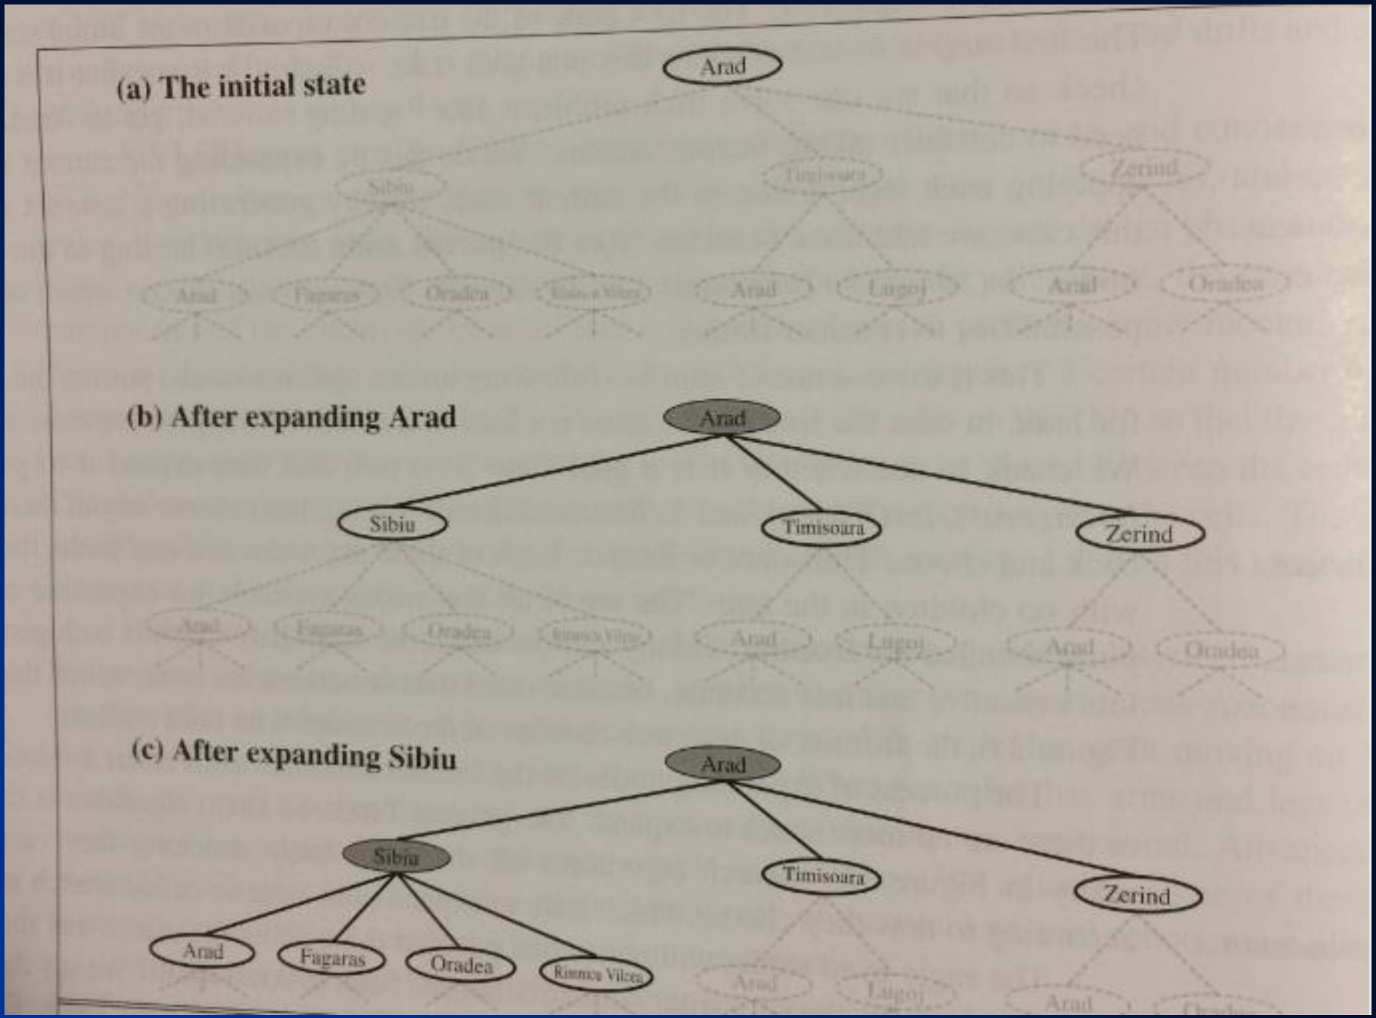

General tree and graph search algorithm

function TREE-SEARCH(*problem*) **returns** a solution, or failure
  initialize the frontier using the initial state of *problem*
  loop do
    if the frontier is empty **then return** failure
    choose a leaf node and remove it from the frontier
    if the node contains a goal state **then return** the corresponding solution
    expand the chosen node, adding the resulting nodes to the frontier

---

function GRAPH-SEARCH(*problem*) **returns** a solution, or failure
  initialize the frontier using the initial state of *problem*
  initialize the explored set to be empty
  loop do
    if the frontier is empty **then return** failure
    choose a leaf node and remove it from the frontier
    if the node contains a goal state **then return** the corresponding solution
    add the node to the explored set
    expand the chosen node, adding the resulting nodes to the frontier
      only if not in the frontier or explored set

Figure 3.7 An informal description of the general tree-search and graph-search algorithms. The parts of GRAPH-SEARCH marked in bold italic are the additions needed to handle repeated states.

Real-world search

- **Driving directions** (e.g., Google maps)
- **Airline travel problems**: Find “optimal” flight subject to conditions entered by user (e.g., for kayak.com)
- **Tourism problems**: E.g., “Visit every city in Romania map at least once, starting and ending in Bucharest.”
- **Traveling salesman problem**: Find shortest tour in which each city visited exactly once
- **VLSI layout**: Positioning millions of components and connections on a chip to minimize area, circuit delays, stray capacitances, and maximize manufacturing yield
- **Robot navigation**: generalization of route-finding problem to continuous space with potentially infinite set of actions and states.
- **Automatic assembly sequencing** of complex objects by a robot, e.g., electric motors and protein design

Evaluating performance

- **Completeness**: Is the algorithm guaranteed to find a solution when there is one?
- **Optimality**: Does the strategy find the optimal solution, as defined on page 68 (i.e., lowest path cost among all solutions)
- **Time complexity:** How long does it take to find a solution?
- **Space complexity:** How much memory is needed to perform the search?

Complexity

- Time and space complexity are always considered with respect to some measure of the problem difficulty. In theoretical computer science, the typical measure is the size of the state space graph, |V| + |E|, where V is the set of vertices (nodes) of the graph and E is the set of edges (links). This is appropriate when the graph is an explicit data structure that is input to the search program (e.g., map of Romania).

Complexity

- In AI, the graph is often represented *implicitly* by the initial state, actions, and transition model, and is frequently infinite. For these reasons, complexity is expressed in terms of three quantities:
  - b, the **branching factor** or maximum number of successors of any node;
  - d, the **depth** of the shallowest goal node (i.e., the number of steps along the path from the root);
  - m, the maximum length of any path in the state space
- Time is often measured in terms of the number of nodes generated during the search, and space in terms of the maximum number of nodes stored in memory.

Search algorithm effectiveness

- To assess the effectiveness of a search algorithm we can consider just the **search cost**—which typically depends on the time complexity but can include a term for memory usage– or we can use **total cost**, which combines the search cost and the path cost of the solution found.
- For problem of finding route from Arad to Bucharest, the search cost is the amount of time taken by the search (milliseconds) and the solution cost is the total length of the path in kilometers.
  - To compute total cost, we can convert km to ms by using a car's average speed.

Uninformed and informed search

- For **uninformed search** (aka **blind search**), the strategies have no additional information about states beyond that provided in the problem definition. All they can do is generate successors and distinguish a goal state from a non-goal state. All search strategies are distinguished by the *order* in which nodes are expanded. Strategies that know whether one non-goal state is "more promising" that another are called **informed search** or **heuristic search** strategies.

Uninformed search algorithms

- Breadth-first search
- Uniform-cost search
- Depth-first search
- Depth-limited search
- Iterated deepening depth-first search
- Bidirectional search

Informed search algorithms

- Greedy best-first search
- A* search
- Memory-bounded heuristic search

- These algorithms have a heuristic function to help guide the search.

Homework for next class

- Chapter 3 from Russell-Norvig textbook.

DARPA Grand Challenge

- https://www.youtube.com/watch?v=8P9geWwi9e0