

Homework 4 (Generic Version)

Assigned on: Friday, September 30, 2022

Due: Friday, October 14, 2022

Programming assignment should be submitted with **handin**.
Do not hesitate to seek help during recitation and office hours.

This homework has two parts:

1. A regular ‘exercise’ on AI games.
2. A programming assignment. The goal of the programming assignment is to implement both uninformed and heuristic search.
 - (a) Romanian Holidays is mandatory for all students. The implementation is explained in great details.
 - (b) Eight-sliding puzzle is a bonus assignment for curious students. Rough explanations about the implementation are provided.

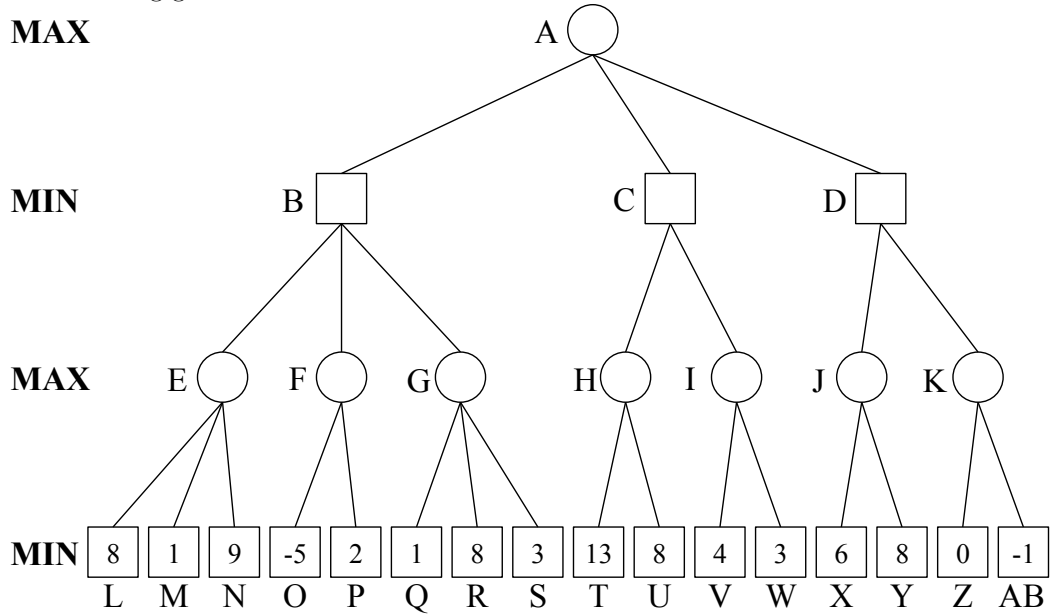
AIMA exercises are available online : ”<https://aimacode.github.io/aima-exercises>”

Contents

1 Adversarial Search (Total 5 points)	2
2 Alpha-beta Pruning (Total 5 points)	3
3 Chapter 5, Exercise 8, Source: AIMA online site. (Total 10 points)	3
4 Romanian Holidays (Total 100 points)	4
4.1 Data structures (50 points)	4
4.1.1 Tasks	4
4.2 Implementing Search (50 points)	5
4.2.1 Results to report	6
4.2.2 Some indications	6
5 Eight-Piece Sliding Puzzle (Bonus 80 points)	7
5.1 Requirements	7
5.2 Indications	7

1 Adversarial Search (Total 5 points)

Consider the following game tree:



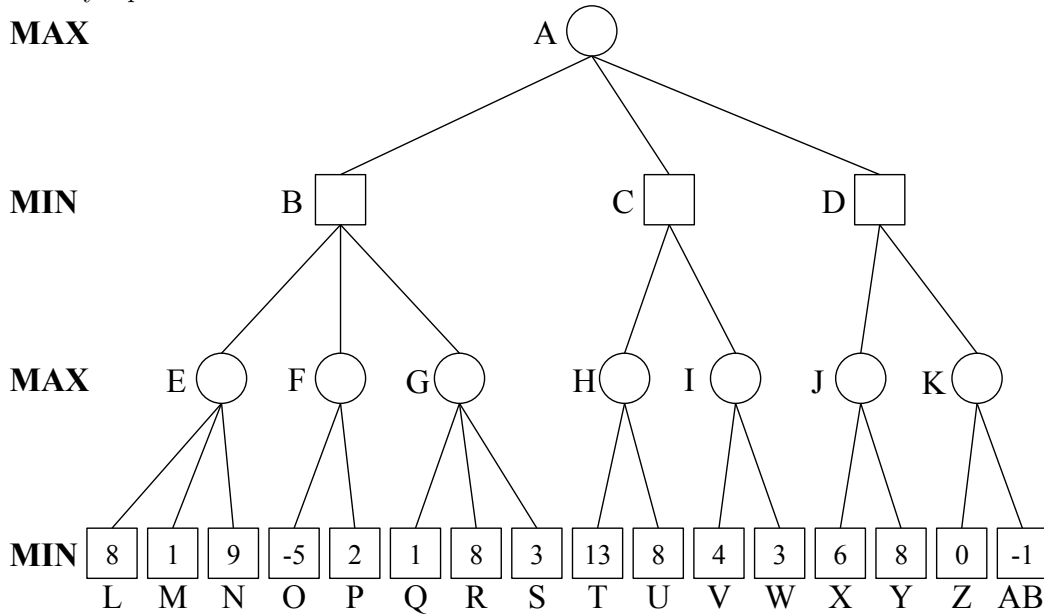
1. Compute the minimax decision. Show your answer by writing the values at the appropriate nodes in the above tree. 4 points

2. What move should Max choose?

1 point

2 Alpha-beta Pruning (Total 5 points)

Using the *alpha-beta pruning method*, with standard left-to-right evaluation of nodes, show what nodes are *not* examined by alpha-beta.



3 Chapter 5, Exercise 8, Source: AIMA online site. (Total 10 points)

Consider the two-player game described in Chapter 5, Exercise 7.

1. Draw the complete game tree, using the following conventions:
 - Write each state as (s_A, s_B) , where s_A and s_B denote the token locations.
 - Put each terminal state in a square box and write its game value in a circle.
 - Put *loop states* (states that already appear on the path to the root) in double square boxes. Since their value is unclear, annotate each with a “?” in a circle.
2. Now mark each node with its backed-up minimax value (also in a circle). Explain how you handled the “?” values and why.
3. Explain why the standard minimax algorithm would fail on this game tree and briefly sketch how you might fix it, drawing on your answer to part (2). Does your modified algorithm give optimal decisions for all games with loops?
4. This 4-square game can be generalized to n squares for any $n > 2$. Prove that A wins if n is even and loses if n is odd.

4 Romanian Holidays (Total 100 points)

This exercise will guide you, step by step, to implement the data structures representing Romania's map and the search algorithms for conducting search. It is mandatory to all students.

Note that in the following sections, we refer to variable and function names that may not be accepted by the programming language you choose (e.g., C, C++, and Java variables cannot contain hyphens). In these cases, just use names that are similar but acceptable by the language by, for example, using camel case instead of hyphens. Additionally, an association list is merely a list of key-value pairs.

It is not necessary to implement these structures *exactly*, but your implementation should be similar enough that the graders can clearly see the analogous variables and structures.

4.1 Data structures (50 points)

Create the data structures required to represent the map of Romania. Include the information about the distances between two cities linked by a road as well as the distance from any given city to Bucharest as indicated in Figure 1.

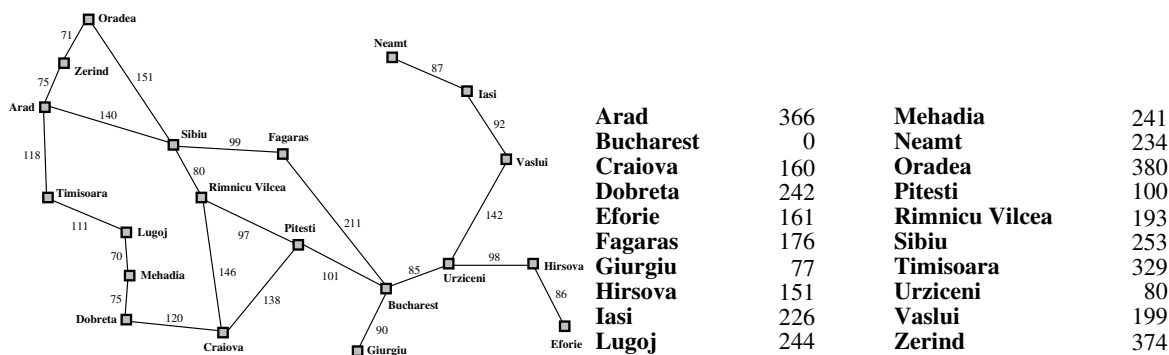


Figure 1: Map of Romania with road distances in kilometers and straight-line distances to Bucharest.

Indications (follow illustration in Figure 2):

- Create a data structure for a city.
- Include an attribute `name` to store the name of the city.
- Include an attribute `neighbors` to store the neighboring cities.
- Include an attribute `h` that provides the value of the straight-line distance to Bucharest.
- Create two ways to access the cities:
 1. a list called `allCitiesList` and
 2. a hash-table `allCitiesHash`
- After creating structures for all the cities, loop through them again in order to include, in the relevant attribute of a city, a reference its neighboring cities. Store these neighbors as an association list of the structure of a neighbor and the distance between the two.

4.1.1 Tasks

1. (10 points) Design, implement and test your map.
2. (5 points) Write a function `allCitiesFromList` that takes a global variable, `allCitiesList`, and returns a list of all *names* of cities on the map.

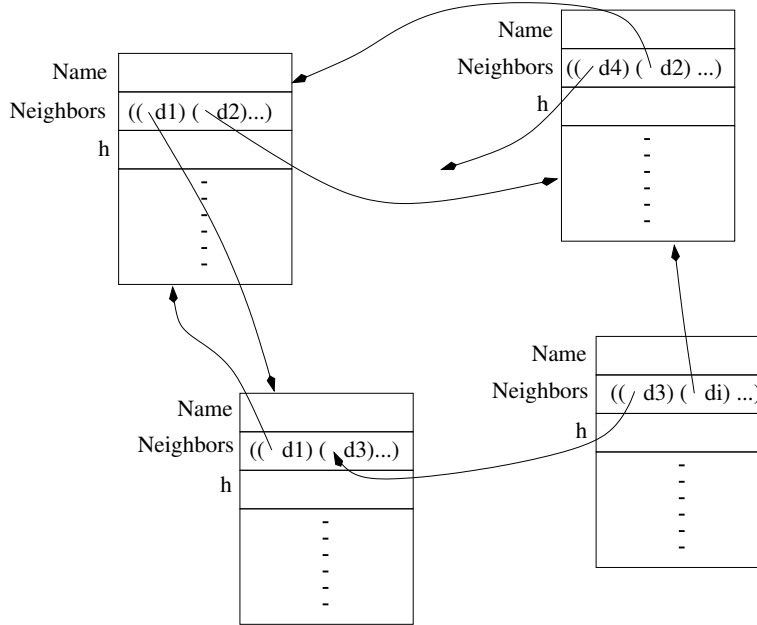


Figure 2: Data structures.

3. (5 points) Write a function `allCitiesFromHtable` that takes a variable, `allCitiesHtable` and returns a list of all the *structures* of cities on the map.
4. (5 point) Write two functions `getCityFromList` and `getCityFromHtable` that take the *name* of a city as input and return the corresponding structure (by accessing a variable, `allCitiesList` and `allCitiesHtable`, respectively).
5. (5 points) Write two functions `neighborsUsingList` and `neighborsUsingHtable` that take the *name* of a city as input and return the list of structures of its direct neighbors. `neighborsUsingList` and `neighborsUsingHtable` should use `getCityFromList` and `getCityFromHtable`, respectively.
6. (10 points) Using `allCitiesHtable`, write a function `neighborsWithinD` that takes the name of a city `myCity` and a number `distance`, then returns, for all direct neighbors within `distance` from `myCity` (\leq), an association list of the structures of the neighbors of `myCity` and their distance to `myCity`.
7. (10 points) Using `allCitiesHtable`, write a function `neighborsP` that takes the *name* of two cities `cityOne` and `cityTwo`, and returns the distance between them if they are directly connected or nil if they are not.

4.2 Implementing Search (50 points)

You are asked to implement search the following search strategies, first as a TREE-SEARCH then as a GRAPH-SEARCH:

- Any uninformed search strategy of your choice, 10 points
- A Greedy search strategy, and 10 points
- An A* search strategy. 10 points

for the ‘Romanian Holiday’ problem. Needless to say, you should first get Section 1 to work. Write `Search` that take as input the name of any city on the map, the name of a search strategy, and outputs:

1. The path to Bucharest,

2. The number of nodes generated/visited by the search process,
3. The cost of the path found (even when the function $g(n)$ is not used to choose the node to expand),
4. The running time spent on the search.

Hints:

- You may choose to write one search function and give it the strategy as an argument.

4.2.1 Results to report

In addition to your code, report the results of your two functions applied to *each* city in Romania as indicated in the table below. Report two tables:

1. One result table for TREE-SEARCH. 10 points
2. One result table for GRAPH-SEARCH. 10 points

Uninformed search of your choice				
City name	#nodes visited	Path to Bucharest	Total cost of path	CPU time
Arad				
Bucharest				
⋮				
Vaslui				
Zerind				

Greedy Search				
City name	#nodes visited	Path to Bucharest	Total cost of path	CPU time
Arad				
Bucharest				
⋮				
Vaslui				
Zerind				

A* Search				
City name	#nodes visited	Path to Bucharest	Total cost of path	CPU time
Arad				
Bucharest				
⋮				
Vaslui				
Zerind				

4.2.2 Some indications

Follow the requirements below:

1. For GRAPH-SEARCH, modify the data structure of a city that you implemented in Section 4.1 to add one more field `visited`, initialized to `Nil`. Use this attribute for loop control during search: when a city is visited, set this field to `T`.
2. Create a new data structure to represent a node in the search tree. The structure should have attributes that *point* to the structures of its parent (when applicable), its children (list), the city it represents. Other attributes may be necessary, such as path value at the node.
3. Implement a function `expand-node` that takes a node in the search tree and generates its children, which, for GRAPH-SEARCH, should correspond to cities not *yet* visited. It needs to generate one node data-structure per child.

4. Implement a function `evaluate-node` that takes a node and a search strategy and returns the value of the node (e.g., $g(n)$, $h(n)$ or $f(n)$).
5. Implement a function that takes a fringe (i.e., a list of nodes to be expanded) and returns the node to expand. As a refinement, you can provide the name of the search strategy as an optional second argument (check `:key` in the list of arguments of a function).
6. If you separate the implementation search strategy from the evaluation functions cleverly enough, you may be able to use the same search function for all search strategies you implement.
7. Implement the search strategies iteratively, not recursively.
8. Declare a global variable `*nnv*` for storing the *number of nodes visited*. The search function should set up its value and the function `expand-node` should increment this value at every expansion (technically, every instantiation of a search-node structure).
9. Load the information about the cities from the file `all-cities.lisp` which is on the course website (it is also acceptable to hard-code this information if you don't want to parse a lisp file).

5 Eight-Piece Sliding Puzzle (Bonus 80 points)

The goal is to implement A* search for solving the Eight-Piece Sliding Puzzle Problem with the two admissible heuristics: the displaced tile and the Manhattan distance heuristics.

5.1 Requirements

Below is the list of basic requirements:

1. Implement a generator of random states, to be used to generate an initial and a goal state.
2. Implement A* that
 - takes as arguments an initial state, a goal state, and the name of a heuristic function and
 - returns the list of moves of the empty tile from the initial state to the goal state, the cost of the path found, and the CPU time.
3. Count the movement of each tile as unity cost.
4. Generate 100 combinations of random initial and goal states (more if possible).
5. Run A* with each heuristic on each combination initial and goal states, reporting the following:

Instance	Displaced tile			Manhattan Distance		
	#NV	Path Cost	CPU time	#NV	Path Cost	CPU time
Combination 1						
Combination 2						
⋮						
Combination 99						
Combination 100						

5.2 Indications

In contrast to the previous problem, you do not need to generate a state space. Use to the extent possible the explanations provided above for the Romanian Holidays and the mechanisms

1. Implement a state as a 2 dimensional array (Chapter 17 in CITI).
2. Implement a generator of random states, to be used to generate an initial and a goal state.

3. Implement a function that takes as input a state and computes the value of the heuristic function (there are two of them).
4. Implement functions that correspond to the actions of moving the *empty tile* north, south, east and west (similar to the Farmer's dilemma).
5. Implement functions that determine whether a move is legal or not (similar to the Farmer's dilemma). A move north is not legal when the empty tile is in the first row.