

Artificial Intelligence

Informed Search

LESSON 5

prof. Antonino Staiano

M.Sc. In "Machine Learning e Big Data" - University Parthenope of Naples

Informed Search

- Uninformed search is based on systematically exploring the search space, and does not exploit any information (if any) about what nodes are more “promising” than others toward the solution
- When **some knowledge is available**, it can be exploited to improve the effectiveness and efficiency of tree search
- Idea
 - Use the available problem-specific knowledge to identify the **best node** to expand at each step of the general tree-search algorithm
 - This general approach is named **best-first** search

Best-First Search

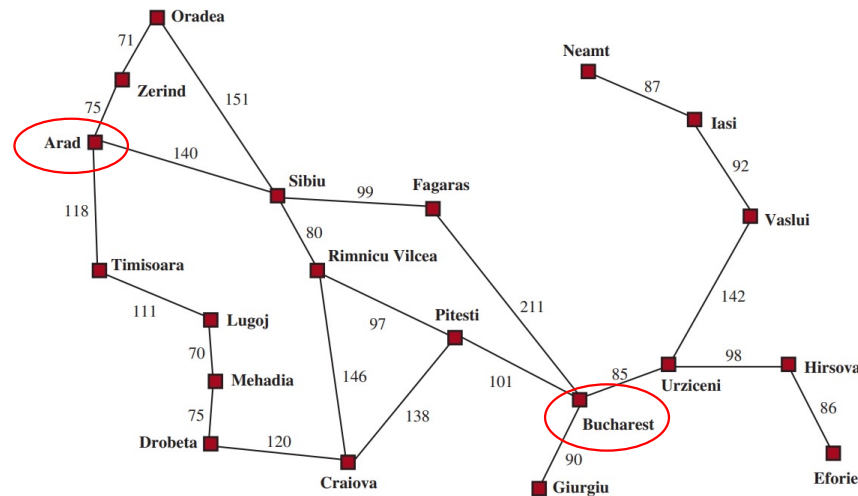
- It is based on quantitatively evaluating how promising a given node n is toward a solution
 - It uses a suitable node evaluation function $f(n)$
- Different definitions of $f(n)$ correspond to different specific best-first strategies, i.e.,
 - Greedy search
 - A*-search and its many variants
- Given a suitable $f(n)$ the search algorithm can be implemented based on the [general tree-search algorithm](#)
- Best-first search can be implemented by sorting nodes in the **frontier** for increasing values of $f(n)$
 - This means that the node n with the lowest $f(n)$ will be selected for expansion at each iteration

Best-First Search

- To define $f(n)$, the cost of the actions that will lead from any given node n to a goal state might be used
- The **exact** cost is usually unknown, thus an **estimate** can be easily computed
 - The estimated cost, a function of the nodes, is denoted with $h(n)$
 - Note that, by definition, $h(n) = 0$ if n contains a goal state (this is the only case in which the cost is **exactly** known)
- For historical reasons $h(n)$ is named the **heuristic function**, and search strategies are named **heuristic search**
- Heuristic search is one of the earlier achievements of AI (50's) and is still widely used in real-world problems and investigated by researchers in AI

Heuristic Function: Example

- Let's consider the shortest route findings on maps
 - From Arad to Bucharest, using the information on the map



- The actions' cost is evaluated as the route length
- A heuristic function is defined as estimating the distance between any given city and the destination

Heuristic Function: Example

- An easy-to-compute estimate is a straight-line distance (using the geographical coordinates of each city)
- If the destination is Bucharest, the heuristic function $h(n)$ can be defined as the straight-line distance from the node n city to Bucharest
- The values of $h(n)$ are

Arad	366	Mehadia	241
Bucharest	0	Neamt	234
Craiova	160	Oradea	380
Drobeta	242	Pitesti	100
Eforie	161	Rimnicu Vilcea	193
Fagaras	176	Sibiu	253
Giurgiu	77	Timisoara	329
Hirsova	151	Urziceni	80
Iasi	226	Vaslui	199
Lugoj	244	Zerind	374

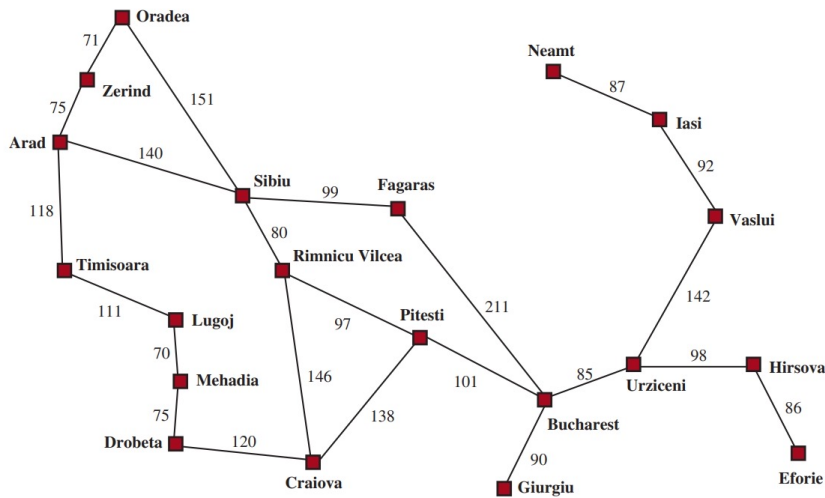

Greedy Best-First Search

- It is the simplest best-first search strategy
 - Expanding the node closest to the solution
- Since the exact cost is unknown, we use an estimated cost, the heuristic function $h(n)$
- It's a greedy strategy since it favors partial solutions that seem the closest to the actual solution
 - However, that's not an optimal choice

Greedy Best-First Search: Example

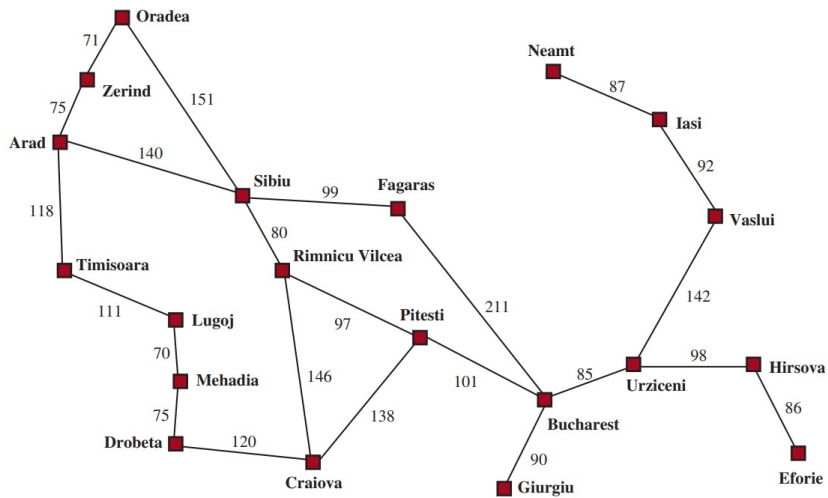

- Consider again the problem of finding the shortest route from Arad to Bucharest, using the straight-line distance heuristic
- In the following, a greedy strategy is used to build the search tree
 - It is shown the value of $f(n)$
 - An arrow denotes the node chosen for the expansion

Greedy search example




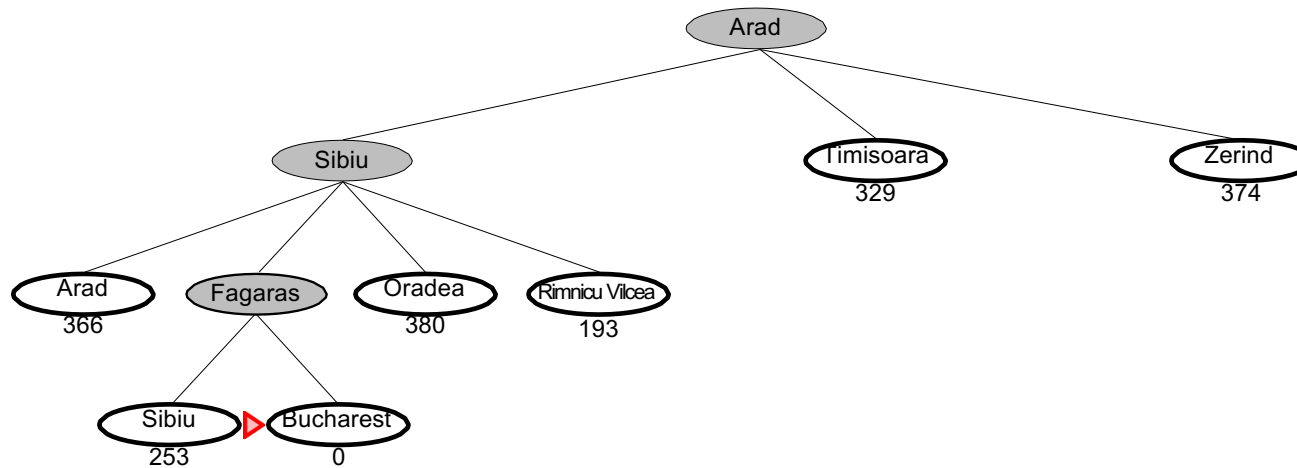
Arad	366	Mehadia	241
Bucharest	0	Neamt	234
Craiova	160	Oradea	380
Drobeta	242	Pitesti	100
Eforie	161	Rimnicu Vilcea	193
Fagaras	176	Sibiu	253
Giurgiu	77	Timisoara	329
Hirsova	151	Urziceni	80
Iasi	226	Vaslui	199
Lugoj	244	Zerind	374

Greedy search example



Arad	366	Mehadia	241
Bucharest	0	Neamt	234
Craiova	160	Oradea	380
Drobeta	242	Pitesti	100
Eforie	161	Rimnicu Vilcea	193
Fagaras	176	Sibiu	253
Giurgiu	77	Timisoara	329
Hirsova	151	Urziceni	80
Iasi	226	Vaslui	199
Lugoj	244	Zerind	374

Greedy search example





Arad	366	Mehadia	241
Bucharest	0	Neamt	234
Craiova	160	Oradea	380
Drobeta	242	Pitesti	100
Eforie	161	Rimnicu Vilcea	193
Fagaras	176	Sibiu	253
Giurgiu	77	Timisoara	329
Hirsova	151	Urziceni	80
Iasi	226	Vaslui	199
Lugoj	244	Zerind	374

Properties of greedy best-first search

- Complete
 - Unless there are infinite paths
- Non-optimal
 - For instance, a shorter route exists between Arad and Bucharest (through Sibiu and Rimnicu Vilcea)
- Exponential worst-case time and space complexity
 - $O(b^m)$, with m depth of the shallowest solution and b branching factor

A* Search

- A* is the most relevant best-first search strategy
 - Devised for robot navigation in '60s



- Many variants proposed to tune the trade-off between its effectiveness and efficiency


A* Search

- A greedy search estimates the costs of the actions from a node n to expand to a solution and chooses the closest one
- A* estimates the **total** cost of the action sequence from the root to the goal state through n
 - Sum of the path cost to n and the estimated cost from n to a solution
 - The node evaluation function is defined as


$$f(n) = \underbrace{g(n)}_{\text{path cost from root to } n} + \underbrace{h(n)}_{\text{estimated cost of the shortest path from node } n \text{ to a goal}}$$

A* Search: Example


- The following search tree built by A* is shown from the previous example (Arad-Bucharest, using the straight-line distance heuristic)
 - The value of $f(n) = g(n) + h(n)$ is also shown for each node



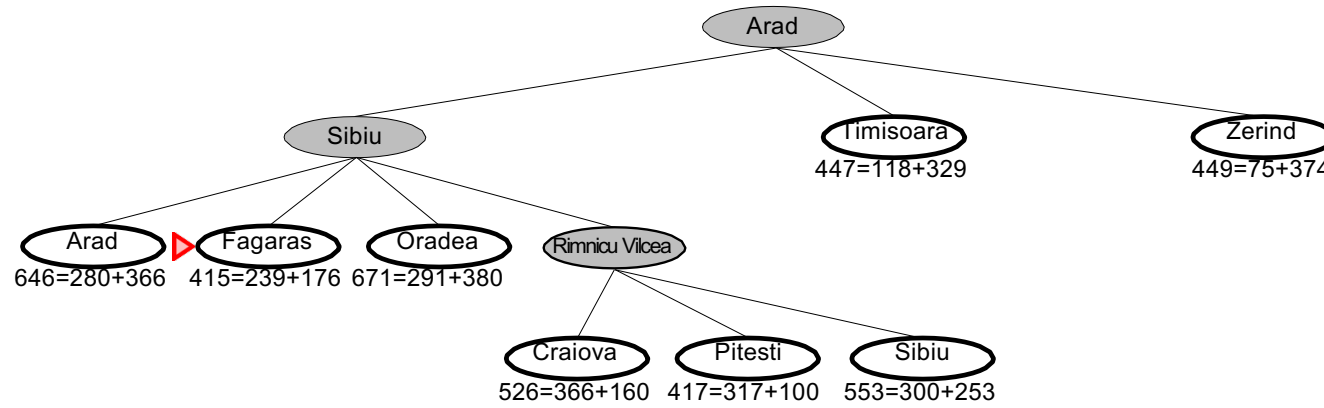
A* Search: Example



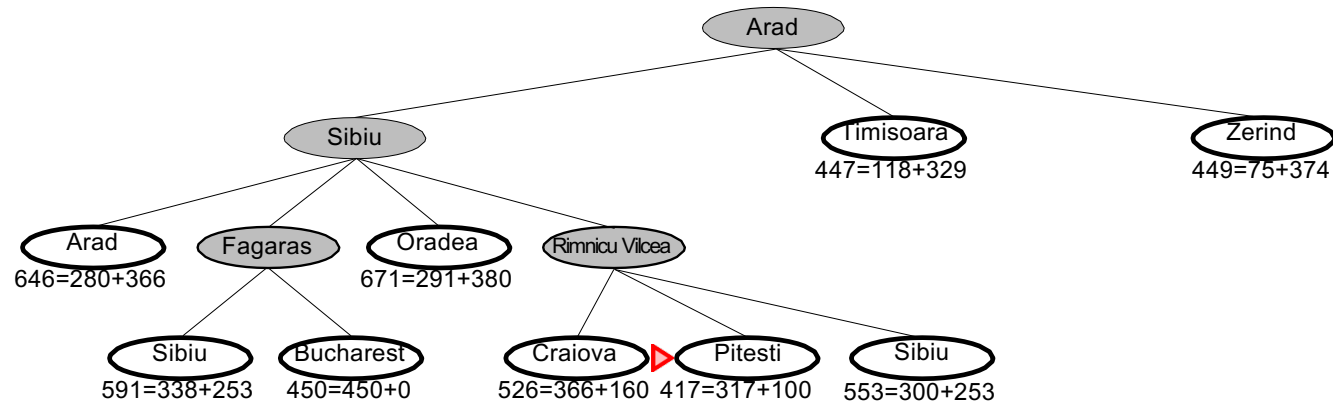
A* Search: Example



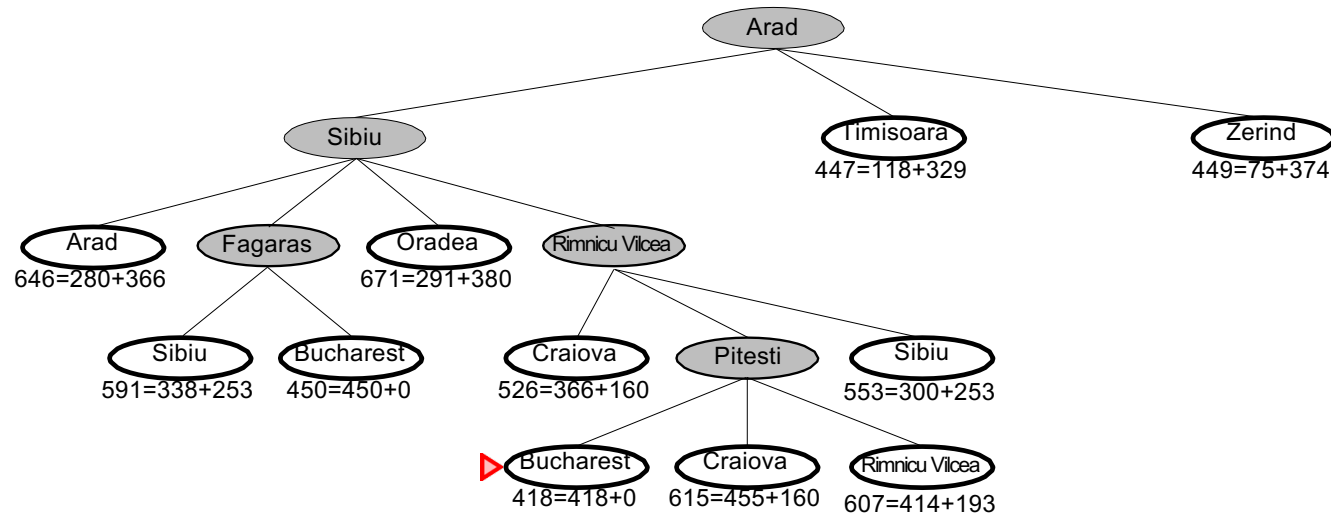
A* Search: Example



A* Search: Example



A* Search: Example



Properties of A* Search

- Optimal
 - Provided that the heuristic is **admissible**
 - That is, it never overestimates the cost of the solution
- Complete
 - If costs greater than $\epsilon > 0$ and the state-space is finite
- Optimally efficient
 - expands the minimum number of nodes for any admissible heuristic
- Exponential worst-case time and space complexity
 - However, A* is much more efficient (generates a much smaller number of nodes) than other uninformed and informed search strategies

Proof (by contradiction) of Optimality of A*

- Let's pretend the optimal path cost is C^* , but algorithm A* returns a path with cost $C > C^*$ (a [suboptimal path](#))
 - There must exist some node n which is on the optimal path and unexpanded

$$f(n) > C^* \quad (\text{otherwise } n \text{ would have been expanded})$$

$$f(n) = g(n) + h(n) \quad (\text{by definition})$$

$$f(n) = g^*(n) + h(n) \quad (\text{because } n \text{ is on an optimal path})$$

$$f(n) \leq g^*(n) + h^*(n) \quad (\text{because of admissibility, } h(n) \leq h^*(n))$$

$$f(n) \leq C^* \quad (\text{by definition, } C^* = g^*(n) + h^*(n))$$

A* Conditions for Optimality

- A slightly stronger property is called **consistency**
 - A heuristic $h(n)$ is consistent if, for every node n and every successor n' of n generated by an action a , we have: $h(n) \leq c(n, a, n') + h(n')$
- This is a form of the triangle inequality




Figure 3.19 Triangle inequality: If the heuristic h is **consistent**, then the single number $h(n)$ will be less than the sum of the cost $c(n, a, a')$ of the action from n to n' plus the heuristic estimate $h(n')$.

- An example of a consistent heuristic is the straight-line distance we have seen earlier for getting to Bucharest
 - Every consistent heuristic is admissible (but not vice versa), so with a consistent heuristic A* is cost-optimal

Improving A* Search

- Good heuristics can lower time and memory demand, particularly w.r.t. uninformed search
- Nonetheless, in many practical problems A* may result unfeasible
 - Alternative approaches
 - Non-optimal A* variants (i.e., find quickly suboptimal solutions)
 - Optimal A* variants with reduced memory requirements and a small increase in execution time

Defining Heuristic Functions

- Intuitively, the more accurate the estimate of the cost to the solution from a given node provided by the heuristic function, the more efficient a best-first algorithm is
- Defining a **good** (i.e., accurate) heuristics is therefore crucial for informed search
- Moreover, heuristics must be **admissible** to guarantee the optimality of A^*

Defining Heuristic Functions: Example

- We have seen that a possible heuristic for route finding in maps is the straight-line distance
- Consider now the 8-puzzle problem
 - Remember that about 3.1×10^{10} nodes are generated on average by breadth-first (uninformed) search -> a good heuristic can be of great practical help also in this toy problem

7	2	4
5		6
8	3	1


Start State

	1	2
3	4	5
6	7	8

Goal State

Defining Heuristic Functions: Example

- Admissible heuristics for k-puzzle:
 - Number of misplaced tiles (h_1)
 - Sum of the distances of each tile from its goal position (h_2)
 - City block or Manhattan
- For instance, the value of h_1 and h_2 for the state (left) w.r.t. the goal state (right):
 - h_1 (start state): 8 (all 8 tiles are misplaced)
 - h_2 (start state): $3+1+2+2+2+3+3+2 = 18$ (tiles 1 to 8)



Start State

Goal State

Defining Heuristic Functions

- Generally, not straightforward to define a heuristic function
 - The approach is to set $h(n)$ to the exact cost of a relaxed version of a problem at hand
- Examples
 - **k-puzzle**: by relaxing the constraint that tiles can move only to a free adjacent square, and allowing them to move to any adjacent square, one obtains $h_2(n)$
 - **k-puzzle**: similarly, allowing tiles to move to any square (even non-adjacent and occupied ones), one obtains $h_1(n)$
 - **route finding on maps**: by relaxing the constraint that an adjacent city can be reached only through the corresponding route, and allowing one to move straight to it, one obtains the straight-line distance heuristic

Choosing Heuristic Functions

- On the other hand, for some problems, it can be possible to define several admissible heuristics h_1, \dots, h_p (e.g., h_1 or h_2 for 8-puzzle)
- In this case, one can choose or define a single heuristic h which dominates all the other ones, i.e.:
 - For each node n , $h(n) \geq h_i(n)$, $i=1, \dots, p$
- It is easy to see that such a heuristic is admissible, and provides a more accurate estimate of the cost to the solution than h_1, \dots, h_p
- To this aim, h can be defined as follows
 - If there is a dominating heuristic among h_1, \dots, h_p , choose it as the heuristic for the problem at hand
 - Otherwise, for a given node n use the following heuristic
 - $h(n) = \max \{h_1(n), \dots, h_p(n)\}$, which dominates by definition h_1, \dots, h_p

Evaluating Heuristic Functions

- To evaluate the quality of heuristic functions the concept of **effective branching factor** (denoted as b^*) is used:
 - let N be the number of nodes generated by A^* for a given problem, and d be the depth of the (optimal) solution
 - b^* is defined as the branching factor of a **uniform** tree of depth d containing N nodes, which is the solution of the equation:
 - $N = 1 + b^* + (b^*)^2 + \dots + (b^*)^d$
- The **lower** the value of b^* , the better the heuristic
- Since b^* depends on the problem **instance**, it is usually evaluated **empirically** as the **average** over a set of instances

Evaluating Heuristic Functions

- Example: Empirical evaluation of the effective branching factor of heuristics h_1 and h_2 for the 8-puzzle (used in A^*), and, for comparison, of an uninformed search strategy, BFS

d	Search Cost (nodes generated)			Effective Branching Factor		
	BFS	$A^*(h_1)$	$A^*(h_2)$	BFS	$A^*(h_1)$	$A^*(h_2)$
6	128	24	19	2.01	1.42	1.34
8	368	48	31	1.91	1.40	1.30
10	1033	116	48	1.85	1.43	1.27
12	2672	279	84	1.80	1.45	1.28
14	6783	678	174	1.77	1.47	1.31
16	17270	1683	364	1.74	1.48	1.32
18	41558	4102	751	1.72	1.49	1.34
20	91493	9905	1318	1.69	1.50	1.34
22	175921	22955	2548	1.66	1.50	1.34
24	290082	53039	5733	1.62	1.50	1.36
26	395355	110372	10080	1.58	1.50	1.35
28	463234	202565	22055	1.53	1.49	1.36

Figure 3.26 Comparison of the search costs and effective branching factors for 8-puzzle problems using breadth-first search, A^* with h_1 (misplaced tiles), and A^* with h_2 (Manhattan distance). Data are averaged over 100 puzzles for each solution length d from 6 to 28.