


Heuristic Search


- Romania: Arad → Bucharest (for example)



7

Heuristic Search

- Romania:
 - Eyeballing it → certain cities first
 - They “look closer” to where we are going
- Can domain knowledge be captured in a **heuristic**?




8

Heuristics Examples

- 8-puzzle:
 - # of tiles in wrong place
- 8-puzzle (better):
 - Sum of distances from goal
 - Captures distance *and* number of nodes
- Romania:
 - Straight-line distance from start node to Bucharest
 - Captures “closer to Bucharest”

5	4	
6	1	8
7	3	2



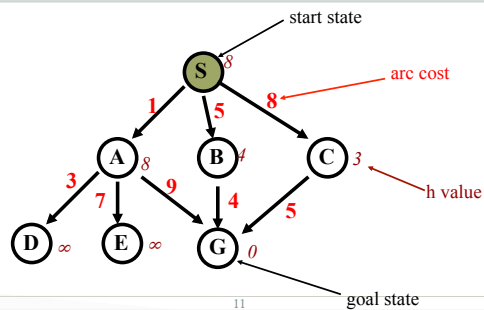
9

Heuristic Function

- All** domain-specific knowledge is encoded in heuristic function h
- h is some estimate of how desirable a move is
 - How “close” (we think) it gets us to our goal
- Usually:
 - $h(n) \geq 0$: for all nodes n
 - $h(n) = 0$: n is a goal node
 - $h(n) = \infty$: n is a dead end (no goal can be reached from n)

10

Example Search Space Revisited



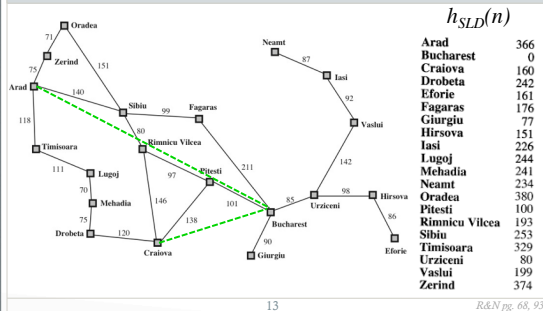
11

Informed Methods Add Domain-Specific Information

- Goal: **select** the best path to continue searching
- Define $h(n)$ to estimate the “goodness” of node n
 - $h(n)$ = **estimated cost** (or distance) of minimal cost path from n to a **goal state**
- Heuristic function is:
 - An estimate of how close we are to a goal
 - Based on domain-specific information
 - Computable from the current state description

12

Straight Lines to Budapest (km)



13

R&N pg. 68, 93

Admissible Heuristics

- Admissible heuristics never overestimate cost
 - They are *optimistic* – think goal is closer than it is
 - $h(n) \leq h^*(n)$
 - where $h^*(n)$ is **true** cost to reach goal from n
 - $h_{LSD}(\text{Lugoj}) = 244$
 - Can there be a shorter path?
- Using admissible heuristics guarantees that the first solution found will be optimal

14

Best-First Search

- A generic way of referring to informed methods
- Use an **evaluation function** $f(n)$ for each **node**
 - estimate of “desirability”
 - $f(n)$ incorporates domain-specific information
 - Different $f(n)$ → Different searches

15


Best-First Search (more)

- Order nodes on the list by
 - Increasing value of $f(n)$
- Expand **most desirable** unexpanded node
 - Implementation:
 - Order nodes in frontier in decreasing order of desirability
- Special cases:
 - Greedy best-first search
 - A* search

16

Greedy Best-First Search


- Idea: always choose “closest node” to goal
 - Most likely to lead to a solution quickly
- So, evaluate nodes based only on heuristic function
 - $f(n) = h(n)$
- Sort nodes by increasing values of f
- Select node believed to be **closest** to a goal node (hence “greedy”)
 - That is, select node with smallest f value



17


Greedy Best-First Search

- Admissible?
 - Why not?
- Example:
 - Greedy search will find: $a \rightarrow b \rightarrow c \rightarrow d \rightarrow e \rightarrow g$; cost = 5
 - Optimal solution: $a \rightarrow g \rightarrow h \rightarrow i$; cost = 3
- Not complete (why?)



18


Straight Lines to Budapest (km)



19

R&N pg. 68, 93

Greedy Best-First Search: Ex. 1




Greedy Best-First Search: Ex. 2



21

Greedy Best-First Search: Ex. 2




22

Greedy Best-First Search: Ex. 2



23

Greedy Best-First Search: Ex. 2



24

Beam Search

- Use an evaluation function $f(n) = h(n)$, but the maximum size of the nodes list is k , a fixed constant
- Only keeps k best nodes as candidates for expansion, and throws the rest away
- More space-efficient than greedy search, but may throw away a node that is on a solution path
- Not complete
- Not admissible

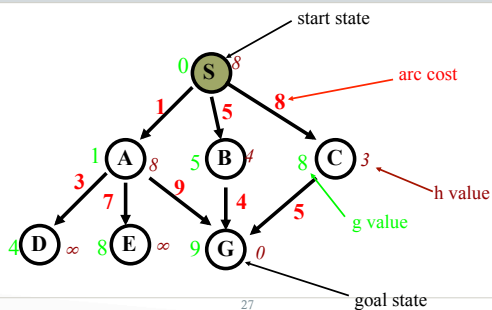
25

Algorithm A

- Use evaluation function $f(n) = g(n) + h(n)$
- $g(n)$ = minimal-cost path from any S to state n
- Ranks nodes on search frontier by *estimated* cost of solution
 - From start node, through given node, to goal
- Not complete if $h(n)$ can = ∞
- Not admissible


26

Example Search Space Revisited



27


Example Search Space Revisited



28

Algorithm A

- Use evaluation function $f(n) = g(n) + h(n)$
- $g(n)$ = minimal-cost path from any S to state n
- Ranks nodes on search frontier by *estimated* cost of solution
 - From start node, through given node, to goal
- Not complete if $h(n)$ can = ∞
- Not admissible



29

Algorithm A

1. Put start node S on the nodes list, called OPEN
2. If OPEN is empty, exit with failure
3. Select node in OPEN with minimal $f(n)$ and place on CLOSED
4. If n is a goal node, collect path back to start; terminate
5. Expand n , generating all its successors, and attach to them pointers back to n . For each successor n' of n
 1. If n' is not already on OPEN or CLOSED
 - put n' on OPEN
 - compute $h(n')$, $g(n') = g(n) + c(n, n')$, $f(n') = g(n') + h(n')$
 2. If n' is already on OPEN or CLOSED and if $g(n')$ is lower for the new version of n' , then:
 - Redirect pointers backward from n' along path yielding lower $g(n')$.
 - Put n' on OPEN.

30

Some Observations on A

- **Perfect heuristic:** If $h(n) = h^*(n)$ for all n :
 - Only nodes on the optimal solution path will be expanded
 - No extra work will be performed
- **Null heuristic:** If $h(n) = 0$ for all n :
 - This is an admissible heuristic
 - A* acts like Uniform-Cost Search

The closer h is to h^* , the fewer extra nodes will be expanded

31

Some Observations on A

- **Better heuristic:** If $h_1(n) < h_2(n) \leq h^*(n)$ for all non-goal nodes, h_2 is a better heuristic than h_1
- If A_1^* uses h_1 , A_2^* uses h_2 ,
 - every node expanded by A_2^* is also expanded by A_1^*
 - So A_1 expands at least as many nodes as A_2

We say that A_2^* is better informed than A_1^*

32

Quick Terminology Check

- What is $f(n)$?
 - An **evaluation function** that gives...
 - A cost estimate of...
 - The distance from n to G
- What is $h(n)$?
 - A **heuristic function** that...
 - Encodes domain knowledge about...
 - The search space
- What is $h^*(n)$?
 - A **heuristic function** that gives the...
 - **True** cost to reach goal from n
 - Why don't we just use that?
- What is $g(n)$?
 - The **path cost** of getting from S to n
 - describes the "spent" costs of the current search

A* Search


- Avoid expanding paths that are already expensive
 - Combines costs-so-far with expected-costs
- A* is **complete** iff
 - Branching factor is finite
 - Every operator has a fixed positive cost
- A* is **admissible** iff
 - $h(n)$ is admissible

34

A* Search

- **Idea:** Evaluate nodes by combining $g(n)$, the cost of reaching the node, with $h(n)$, the cost of getting from the node to the goal.

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




35

A* Example 1




36

A* Example 1




37

A* Example 1




38

A* Example 1




39

A* Example 1



40

A* Example 1




41

Algorithm A*

- Algorithm A with constraint that $h(n) \leq h^*(n)$
 - $h^*(n)$ = true cost of the minimal cost path from n to a goal.
- Therefore, $h(n)$ is an **underestimate** of the distance to the goal
- $h()$ is **admissible** when $h(n) \leq h^*(n)$
 - Guarantees optimality
- A* is **complete** whenever the branching factor is finite, and every operator has a fixed positive cost
- A* is **admissible**

42


Example Search Space Revisited



43

Example

n	$g(n)$	$h(n)$	$f(n)$	$h^*(n)$
S	0	8	8	9
A	1	8	9	9
B	5	4	9	4
C	8	3	11	5
D	∞	∞	∞	∞
E	8	∞	∞	∞
G	9	0	9	0




- $h^*(n)$ is the (hypothetical) perfect heuristic.
- Since $h(n) \leq h^*(n)$ for all n , h is admissible
- Optimal path = S B G with cost 9.

44

Greedy Search

$$f(n) = h(n)$$

Node expanded	Node list
	{ S(8) }
S	{ C(3) B(4) A(8) }
C	{ G(0) B(4) A(8) }
G	{ B(4) A(8) }




- Solution path found is S C G, 3 nodes expanded.
- Fast!! But NOT optimal.

45

A* Search

$$f(n) = g(n) + h(n)$$

node exp.	nodes list
	{ S(8) }
S	{ A(9) B(9) C(11) }
A	{ B(9) G(10) C(11) D(∞) E(∞) }
B	{ G(9) G(10) C(11) D(inf) E(∞) }
G	{ C(11) D(∞) E(∞) }



- Solution path found is S B G, 4 nodes expanded.
- Still pretty fast, and optimal

46

Proof of the Optimality of A*

- Assume that A* has selected G_2 , a goal state with a suboptimal solution ($g(G_2) > f^*$).
- We show that this is impossible.
 - Choose a node n on the optimal path to G.
 - Because $h(n)$ is admissible, $f(n) \leq f^*$.
 - If we choose G_2 instead of n for expansion, $f(G_2) \leq f(n)$.
 - This implies $f(G_2) \leq f^*$.
 - G_2 is a goal state: $h(G_2) = 0, f(G_2) = g(G_2)$.
 - Therefore $g(G_2) \leq f^*$
 - Contradiction.

47

Admissible heuristics

E.g., for the 8-puzzle:

$h_1(n)$ = number of misplaced tiles

$h_2(n)$ = total Manhattan distance
(i.e., # of squares each tile is from desired location)

7	2	4
5		6
8	3	1

Start

	1	2
3	4	5
6	7	8

Goal

- $h_1(S) = ?$
- $h_2(S) = ?$

48

Admissible heuristics

E.g., for the 8-puzzle:

$h_1(n)$ = number of misplaced tiles

$h_2(n)$ = total Manhattan distance
(i.e., # of squares each tile is from desired location)

- $h_1(S) = 8$
- $h_2(S) = 3+1+2+2+2+3+3+2 = 18$



49

Dealing with Hard Problems

- For large problems, A* often requires too much space.
- Two variations conserve memory: IDA* and SMA*
- IDA* – iterative deepening A*
 - uses successive iteration with growing limits on f . For example,
 - A* but don't consider any node n where $f(n) > 10$
 - A* but don't consider any node n where $f(n) > 20$
 - A* but don't consider any node n where $f(n) > 30, \dots$
- SMA* – Simplified Memory-Bounded A*
 - uses a queue of restricted size to limit memory use.
 - throws away the “oldest” worst solution.

50

What's a Good Heuristic?

- If $h_1(n) < h_2(n) \leq h^*(n)$ for all n , then:
 - Both are admissible
 - h_2 is strictly better than (**dominates**) h_1 .
- How do we find one?
 - Relaxing the problem:**
 - Remove constraints to create a (much) easier problem
 - Use the solution cost for this problem as the heuristic function
 - Combining heuristics:**
 - Take the max of several admissible heuristics
 - Still have an admissible heuristic, and it's better!

51

What's a Good Heuristic? (2)

- Use statistical estimates to compute h
 - May lose admissibility
- Identify good features, then use a learning algorithm to find a heuristic function
 - Also may lose admissibility
- Why are these a good idea, then?
 - Machine learning can give you answers you don't “think of”
 - Can be applied to new puzzles without human intervention
 - Often work

52

Some Examples of Heuristics?

- 8-puzzle?
 - Manhattan distance
- Driving directions?
 - Straight line distance
- Crossword puzzle?
- Making a medical diagnosis?

53

Summary: Informed Search

- Best-first search:** general search where the *minimum-cost nodes* (according to some measure) are expanded first.
- Greedy search:** uses *minimal estimated cost* $h(n)$ to the goal state as measure. Reduces search time but, is neither complete nor optimal.
- A* search:** combines UCS and greedy search
 - $f(n) = g(n) + h(n)$
 - A* is complete and optimal, but space complexity is high.
 - Time complexity depends on the quality of the heuristic function.
- IDA* and SMA* reduce the memory requirements of A*.

54

In-class Exercise: Creating Heuristics

8-Puzzle

Start State:

5	4	
6	1	8
7	3	2

Goal State:

1	2	3
8		4
7	6	5

Boat Problems

cabbage
wolf
sheep

Remove 5 Sticks

N-Queens

Water Jug Problem

5 2

Route Planning

In-Class Exercise

```

graph TD
    S((S)) -- 3 --> A((A))
    S -- 1 --> B((B))
    S -- 8 --> C((C))
    A -- 7 --> D((D))
    A -- 7 --> E((E))
    B -- 15 --> G((G))
    C -- 5 --> G
    D --- D_val[∞]
    E --- E_val[∞]
    G --- G_val[0]
    C --- C_h[3]
    
```

Apply the following to search this space. At each search step, show: the current node being expanded, $g(n)$ (path cost so far), $h(n)$ (heuristic estimate), $f(n)$ (evaluation function), and $h^*(n)$ (true goal distance).

Depth-first search Breadth-first search A* search
 Uniform-cost search Greedy search