

3010

ARTIFICIAL INTELLIGENCE

Lecture 3 A* search

Masashi Shimbo

2019-05-15

License:  CC BY-SA 4.0 except where noted

Today's agenda

- ▶ Heuristic evaluation function
- ▶ The A* algorithm

Heuristic search

ヒューリスティック探索

Heuristic search: Motivation

ヒューリスティック探索: 動機

4

Even if we have some “knowledge” about a given problem, Dijkstra’s algorithm doesn’t have a means to take advantage of it.

与えられた問題に対するなんらかの事前知識があっても, ダイクストラ法では活用することができない

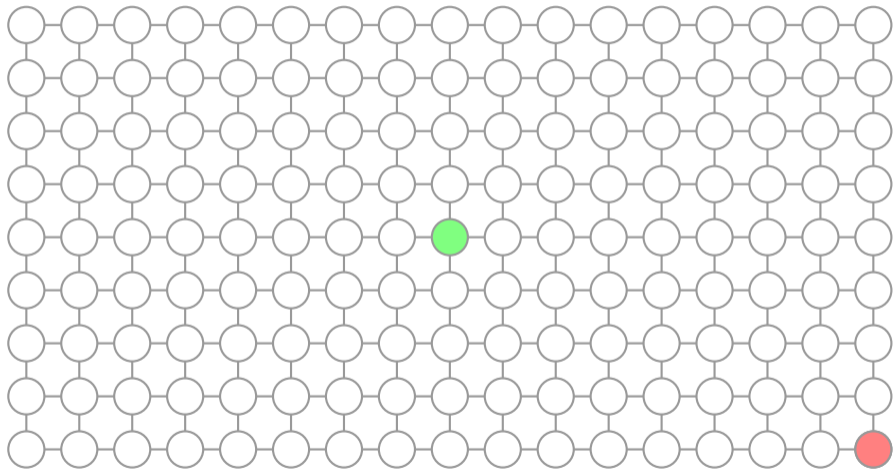
Here is a motivating example...

たとえば...

Even if you know the goal state is located at the lower right corner... 目標

節点が右下隅にあることがわかっても...

5




Dijkstra's algorithm does not take this information into account ダイクスト

Even if you know the goal state is located at the lower right corner... 目標

節点が右下隅にあることがわかっていても...

6



Dijkstra's algorithm does not take this information into account ダイクスト

Heuristic search ヒューリスティック探索

Also called “informed” search 「情報付き」探索とも呼ばれる

7

Search using domain- (problem-)specific **knowledge**, or **heuristics**

解きたい問題に関する「知識」 / 「ヒューリスティック」を活用した探索

But what kind of “knowledge”?

では, どんな「知識」が活用できる?

Heuristic evaluation function $h(v)$

It is assumed that the knowledge about the problem is given as a form of **heuristic evaluation function** $h(v)$

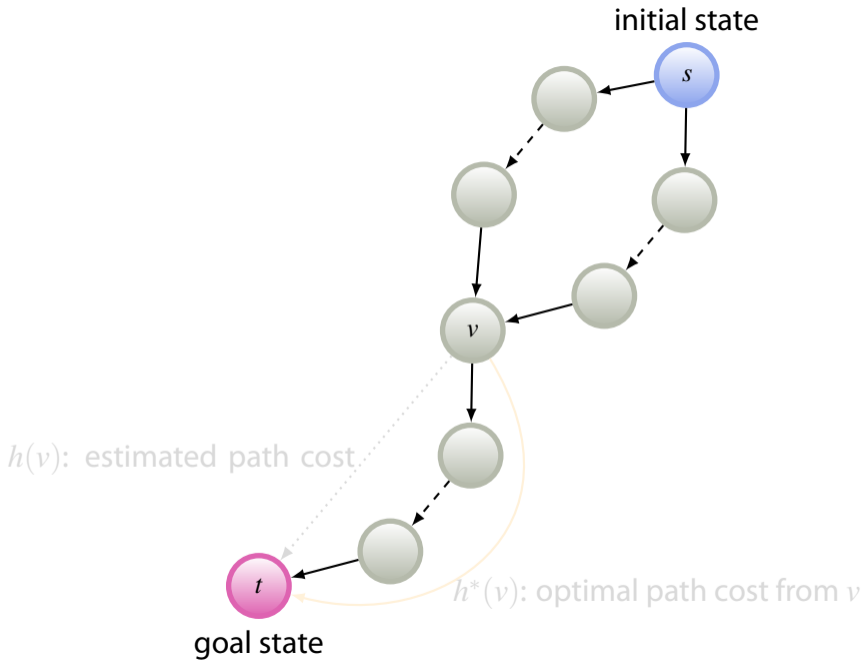
問題に関する知識は「ヒューリスティック評価関数」 $h(v)$ という形で与えられると仮定する

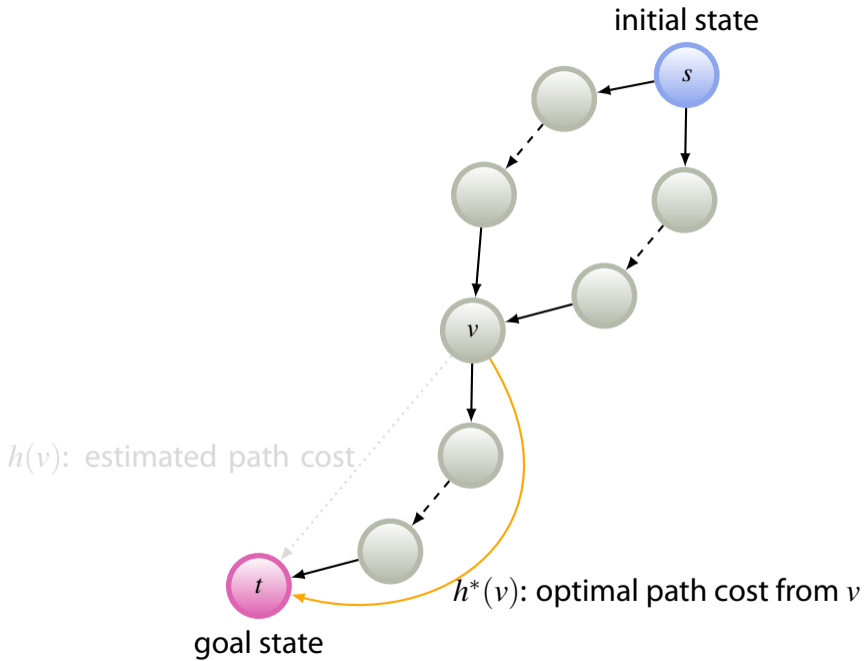
Sometimes simply called a **heuristic function**, or **heuristic**


単に「ヒューリスティック関数」とか「ヒューリスティック」と呼ばれることもある

$h(v)$ = **estimated** cost of the cheapest path from node v
to a terminal node

$h(v)$ の意味は「節点 v から一番近い (= コストが低い) 目標節点までの経路コストの見積もり」







Admissible heuristic 適格なヒューリスティック関数

An important class of heuristic function

A heuristic function h is said to be **admissible**



For every node v , $h(v)$ never overestimates the actual cheapest cost $h^*(v)$ to reach a goal from v



$$h(v) \leq h^*(v) \quad \text{for every node } v.$$



h gives optimistic estimates of actual cost h^* .

Admissible heuristics and the A* algorithm

適格なヒューリスティック関数と A*

13

Why admissible heuristics matter?

なぜヒューリスティック関数の適格性が重要か？

- ... Because the A* algorithm (described later) is admissible (= guaranteed to find a shortest path) if the used heuristic function h is admissible.

用いるヒューリスティック関数 h が適格なら A* は適格 (最短経路を発見することが保証される)

How do we build an admissible heuristic?


“Relaxed problems” approach:

- 1 Make an easier problem RP by removing constraints in the original problem OP .
- 2 Find the optimal solution for RP
- 3 Use the cost of the solution as heuristic h for OP .

The optimal solution (with cost h^*) for OP is also a solution of RP (but not necessarily optimal in RP ; better solutions may exist)

→ $h \leq h^*$


Gridworld



Gridworld

Many obstacles exist 障害物が多数存在

initial state




goal state

Gridworld

Many obstacles exist 障害物が多数存在

initial state




goal state

Gridworld: relaxed problem 緩和問題

Assume obstacles do not exist 障害物が一切存在しないと仮定

18

initial state




goal state

Gridworld: relaxed problem 緩和問題


Assume obstacles do not exist 障害物が一切存在しないと仮定

19



Gridworld: Manhattan-distance heuristic


initial state



goal state

Gridworld: Manhattan-distance heuristic


initial state



goal state

Gridworld: Manhattan-distance heuristic


initial state



goal state

Gridworld: Manhattan-distance heuristic

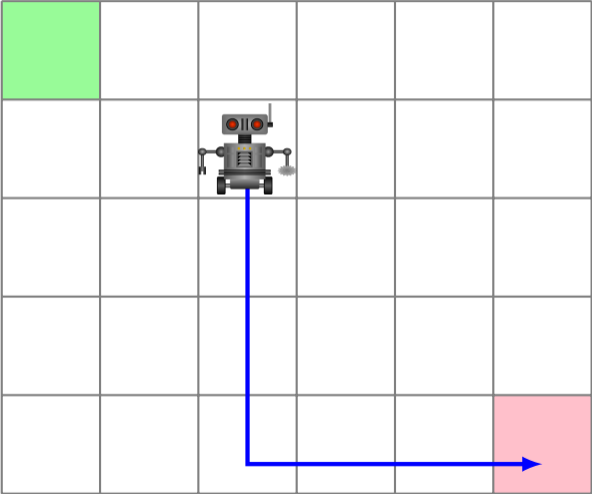
initial state



goal state

Gridworld: Manhattan-distance heuristic


initial state



goal state

Gridworld: Manhattan-distance heuristic

initial state



goal state

Gridworld: Manhattan-distance heuristic

Computation is easy

26

v : arbitrary state

Let

(x, y) coordinates of state v

(x_t, y_t) coordinates of the goal state

Then,

$$h(v) = |x - x_t| + |y - y_t|$$

Admissible heuristic via relaxed problems: Another example




Manhattan distance heuristic for $(n^2 - 1)$ -puzzles:

Sum of the Manhattan distance from each tile to its goal position.




Relaxed problem: tiles can overlap with each other.


Eight-puzzle: Manhattan distance heuristic




1 move




3 moves




0 moves




1 move




1 move



1 move



3 moves



2 moves

$$h = 12$$

Requirement for heuristic functions: They must be efficiently computable

ヒューリスティック関数には「簡単に計算できること」が求められる

No use if computing a heuristic function takes equal or more time and space than actually searching the state space, no matter how accurate its estimates are.

どんな正確なヒューリスティック関数でも、実際の探索を行う以上の時間やメモリが計算に必要ななら、そもそも使う意味がない

Heuristic evaluation function h : Summary

- ▶ h associates a non-negative real number $h(v)$ to each state v
- ▶ $h(v)$ is an **estimate** of the actual cheapest cost $h^*(v)$ necessary to reach a goal state from state v
- ▶ h must be efficiently computable
- ▶ h is said to be **admissible** if $h(v) \leq h^*(v)$ for every state v
- ▶ One way to construct an admissible h is to consider relaxed problems

The A* algorithm

A* アルゴリズム

The A* algorithm

[Hart, Nilsson & Raphael 1968]

Idea

Use the **sum** of

- ▶ the path cost from the initial state to state v
- ▶ the **estimated cost** from v to a goal state

to evaluate how “promising” it is to expand state v .

The A* algorithm

[Hart, Nilsson & Raphael 1968]

Idea

Use the **sum** of

- ▶ $g[v]$
- ▶ $h(v)$

to evaluate how “promising” it is to expand state v .

Evaluation function $f[v]$ of A^*

$$f[v] = g[v] + h(v)$$

where

$g[v]$ tentative minimum cost from the initial state to state v s から v への, これまで見つかった経路のなかで最小のコスト

$h(v)$ estimated cost from state v to the nearest goal state v から最も近い目標節点への経路コストの見積もり


Idea:

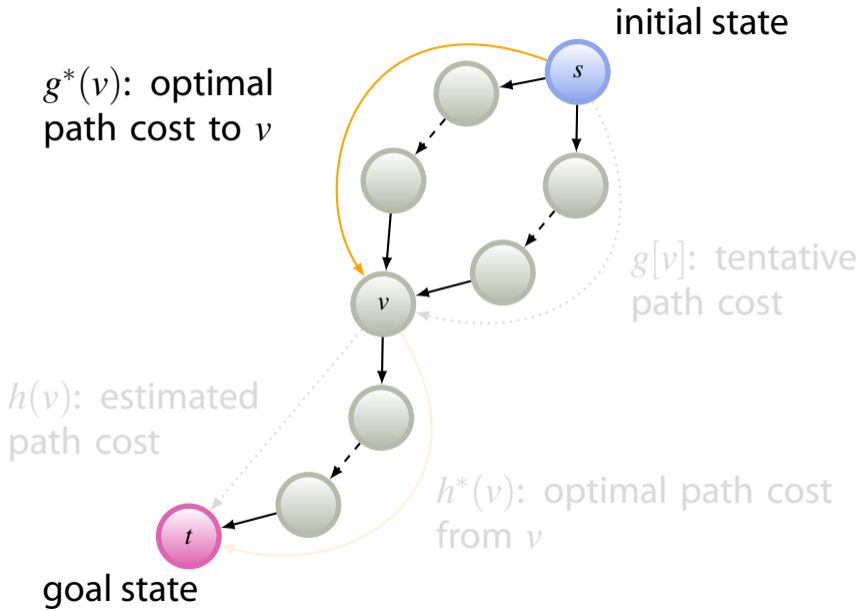
$f[v]$ smaller $\leftrightarrow v$ more promising


f が小さい節点ほど, より有望だ, とみなす

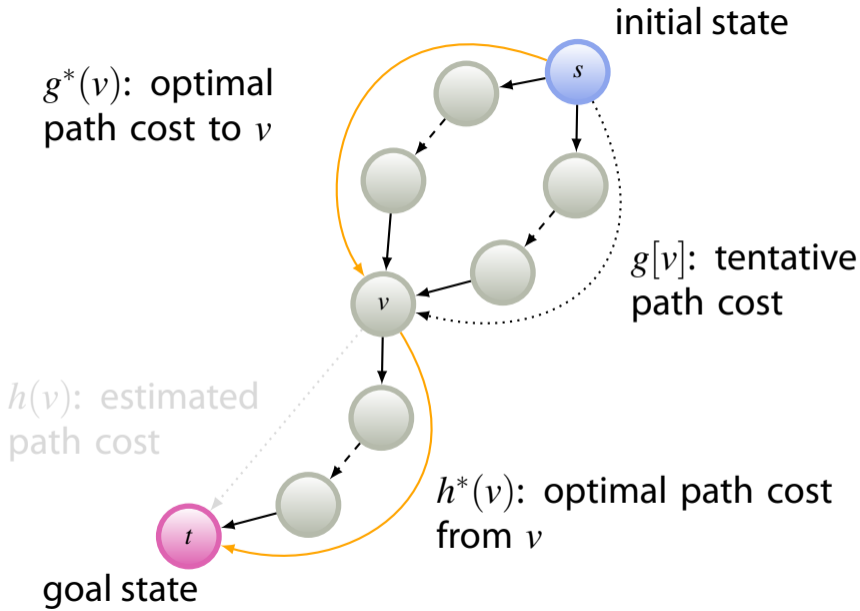
- ▶ $g[v]$: value may get updated if a better path from s to v is found later
- ▶ $h(v)$: once computed, the value will not change


- ▶ $g[v]$: upper-bound of the optimal cost $g^*(v)$ (from s to v)
- ▶ $h(v)$: lower-bound of the optimal cost $h^*(v)$ (from v to goal t) provided that $h(v)$ is admissible.











A* subsumes Dijkstra's shortest path algorithm as a special case

A* reduces to Dijkstra's algorithm if $h(v) = 0$ for every node v .

The A* algorithm

The algorithm is identical to Dijkstra's, except

- ▶ OPEN is a priority queue with priority function $f[v] = g[v] + h(v)$ (not $g[v]$).
- ▶ For each generated node v , $f[v]$ is recorded along with $g[v]$.
- ▶ Nodes in CLOSED can be **re-opened**.

The A* algorithm

The algorithm is identical to Dijkstra's, except

- ▶ OPEN is a priority queue with priority function $f[v] = g[v] + h(v)$ (not $g[v]$).
- ▶ For each generated node v , $f[v]$ is recorded along with $g[v]$.
- ▶ Nodes in CLOSED can be **re-opened**.

The A* algorithm

The algorithm is identical to Dijkstra's, except

- ▶ OPEN is a priority queue with priority function $f[v] = g[v] + h(v)$ (not $g[v]$).
- ▶ For each generated node v , $f[v]$ is recorded along with $g[v]$.
- ▶ Nodes in CLOSED can be **re-opened**.

Priority queue

Two functions for manipulating priority queue P_f :

$\text{Insert}_f(P_f, v)$

Put item v in P_f

$\text{DeleteMin}_f(P_f)$

Remove and return an item with the minimum f -value from P_f .

Thus, the returned item v is the one with

$$v = \underset{u \in P_f}{\operatorname{argmin}} f[u]$$

before removal

How a node changes its status in A*

Closed nodes can be re-opened

Status	Description
Unexplored	
↓	(when a parent node is expanded)
OPEN	the node is generated but not expanded
↓↑	(when the node itself is expanded)
CLOSED	the node is generated and expanded

Dijkstra's shortest path algorithm

Main routine

```
1 OPEN ← new PriorityQueueg
2 g[s] ← 0
3 Insertg(OPEN, s)
4 CLOSED ← ∅
5 loop do
6     if IsEmpty(OPEN) then return "failure"
7     v ← DeleteMing(OPEN)
8     CLOSED ← CLOSED ∪ {v}
9     if IsGoal(v) then return Solution(v, s)
10    Expand(v)
```

OPEN: set of states generated but not expanded

CLOSED: set of expanded states

choose a node with the smallest g

put v in CLOSED

The A* algorithm

Main routine

```
1 OPEN ← new PriorityQueuef # priority is based on  $f = g + h$ 
2  $g[s] \leftarrow 0$ ;  $f[s] \leftarrow h(s)$  #  $f[s] = g[s] + h(s) = 0 + h(s) = h(s)$ 
3 Insertf(OPEN,  $s$ ) # OPEN: set of states generated but not expanded
4 CLOSED ←  $\emptyset$  # CLOSED: set of expanded states
5 loop do
6   if IsEmpty(OPEN) then return "failure"
7    $v \leftarrow$  DeleteMinf(OPEN) # choose a node with the smallest  $f$ 
8   CLOSED ← CLOSED  $\cup$   $\{v\}$ 
9   if IsGoal( $v$ ) then return Solution( $v, s$ )
10  Expand( $v$ )
```


procedure Expand(v) for Dijkstra's algorithm

```
1  foreach  $u \in \text{Succ}(v)$  do
2      if  $u \notin \text{OPEN} \cup \text{CLOSED}$  then                                     # if  $u$  is a new state
3          Reserve memory for  $g[u], \text{Parent}[u]$ 
4           $g[u] \leftarrow g[v] + c(v, u)$                                      # memorize  $g[u]$ 
5           $\text{Parent}[u] \leftarrow v$                                          # also memorize Parent
6          Insert $_g(\text{OPEN}, u)$ 
7      else if  $u \in \text{OPEN}$  then                                           # "relax" edge  $(v, u)$  if  $u \in \text{OPEN}$ 
8          if  $g[v] + c(v, u) < g[u]$  then                                     # if it is better than the current path
9               $g[u] \leftarrow g[v] + c(v, u)$                              # update  $g$  if path through  $(v, u)$  is shorter
10              $\text{Parent}[u] \leftarrow v$                                      # update Parent, too
```


A* and Dijkstra's algorithm: Difference in $\text{Expand}(v)$

51

Case	Dijkstra	A*
$u \notin \text{OPEN}$ nor $u \notin \text{CLOSED}$	$g[u] \leftarrow g[v] + c(v, u)$ $\text{Insert}_g(\text{OPEN}, u)$ $\text{Parent}[u] \leftarrow v$	$g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ $\text{Insert}_f(\text{OPEN}, u)$ $\text{Parent}[u] \leftarrow v$
$u \in \text{OPEN}$	if $g[v] + c(v, u) < g[u]$ then $g[u] \leftarrow g[v] + c(v, u)$ $\text{Parent}[u] \leftarrow v$	if $g[v] + c(v, u) < g[u]$ then $g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ $\text{Parent}[u] \leftarrow v$
$u \in \text{CLOSED}$	Do nothing (always $g[u] \leq g[v] + c(v, u)$)	if $g[v] + c(v, u) < g[u]$ then $g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ $\text{CLOSED} \leftarrow \text{CLOSED} \setminus \{u\}$ $\text{Insert}_f(\text{OPEN}, u)$ $\text{Parent}[u] \leftarrow v$

v = node just expanded / u = a successor of v

Properties of A*

Assumptions



- ▶ At least one solution (path from the initial state to a goal state) exists
- ▶ function $h(\cdot)$ is admissible

Completeness

A* never fails to find a solution

Admissibility

The solution found by A* is optimal



Exercise

Trace the execution of the A* algorithm on this graph. In particular,

- 1 Trace which nodes are on OPEN and which are on CLOSED
- 2 Compute the g -value of each node at each stage
- 3 In what order are nodes expanded?
- 4 How many iterations are necessary before termination?

If you still have time left, trace the behavior of Dijkstra's algorithm (i.e., by setting $h = 0$ for all nodes) on the same graph

A* may reopen closed nodes

Cf. Dijkstra's algorithm never reopens a node.

Is there a class of heuristic functions h such that A* does not open a node more than once?


➡ **Monotone heuristic functions**

Monotone heuristic function

Heuristic function h is said to be **monotone** if

56

- ▶ $h(v) \leq c(v, u) + h(u)$ holds for **every** edge (v, u)
- ▶ $h(t) = 0$ for every goal node t



Monotonicity and admissibility

Monotonicity implies admissibility

h is monotone $\rightarrow h$ is admissible

A* guided by monotone heuristic function h **never** re-opens a node

A* algorithm

Main routine

```
1 OPEN ← new PriorityQueuef # priority is based on f
2 g[s] ← 0
3 f[s] ← h(s) # f[s] = g[s] + h(s) = h(s)
4 Insertf(OPEN, s) # OPEN: set of states generated but not expanded
5 CLOSED ← ∅ # CLOSED: set of expanded states
6 loop do
7   if IsEmpty(OPEN) then return "failure"
8   v ← DeleteMinf(OPEN) # choose a node with the smallest f
9   CLOSED ← CLOSED ∪ {v}
10  if IsGoal(v) then return Solution(v, s)
11  Expand(v)
```

A* algorithm with monotone h

Main routine — No change from the original A*

```
1 OPEN ← new PriorityQueuef # priority is based on  $f$ 
2  $g[s] \leftarrow 0$ 
3  $f[s] \leftarrow h(s)$  #  $f[s] = g[s] + h(s) = h(s)$ 
4 Insertf(OPEN,  $s$ ) # OPEN: set of states generated but not expanded
5 CLOSED ←  $\emptyset$  # CLOSED: set of expanded states
6 loop do
7   if IsEmpty(OPEN) then return "failure"
8    $v \leftarrow$  DeleteMinf(OPEN) # choose a node with the smallest  $f$ 
9   CLOSED ← CLOSED  $\cup$   $\{v\}$ 
10  if IsGoal( $v$ ) then return Solution( $v, s$ )
11  Expand( $v$ )
```

Procedure Expand(v)

for A* algorithm

61

```
1  foreach  $u \in \text{Succ}(v)$  do
2      if  $u \notin \text{OPEN} \cup \text{CLOSED}$  then                                     # if  $u$  is a new state
3          Reserve memory for  $g[u]$ ,  $f[u]$ , and  $\text{Parent}[u]$ 
4           $g[u] \leftarrow g[v] + c(v, u)$ ;  $f[u] \leftarrow g[u] + h(u)$            # memorize  $f[u]$  as well as  $g[u]$ 
5           $\text{Parent}[u] \leftarrow v$                                            # also memorize  $\text{Parent}$ 
6           $\text{Insert}_f(\text{OPEN}, u)$ 
7      else if  $u \in \text{OPEN}$  then                                           # "relax" edge  $(v, u)$  if  $u \in \text{OPEN}$ 
8          if  $g[v] + c(v, u) < g[u]$  then                                     # if it gives a better path than the current one
9               $g[u] \leftarrow g[v] + c(v, u)$ ;  $f[u] \leftarrow g[u] + h(u)$      # update  $g$  and  $f$ 
10              $\text{Parent}[u] \leftarrow v$                                        # update  $\text{Parent}$ , too
11      else                                                                 # if  $u \in \text{CLOSED}$ , "relax" edge  $(v, u)$  and re-open  $u$  if necessary
12          if  $g[v] + c(v, u) < g[u]$  then                                     # if a cheaper path is found
13               $g[u] \leftarrow g[v] + c(v, u)$ ;  $f[u] \leftarrow g[u] + h(u)$      # update  $g$  and  $f$ 
14               $\text{Parent}[u] \leftarrow v$                                        # update  $\text{Parent}$ 
15               $\text{CLOSED} \leftarrow \text{CLOSED} \setminus \{u\}$                    # take  $u$  out of  $\text{CLOSED}$ 
16               $\text{Insert}_f(\text{OPEN}, u)$                                        # and put it back in  $\text{OPEN}$ 
```

Procedure Expand(v)

for A* algorithm **when h is monotone**

62

```
1  foreach  $u \in \text{Succ}(v)$  do
2      if  $u \notin \text{OPEN} \cup \text{CLOSED}$  then                                     # if  $u$  is a new state
3          Reserve memory for  $g[u]$ ,  $f[u]$ , and  $\text{Parent}[u]$ 
4           $g[u] \leftarrow g[v] + c(v, u)$ ;  $f[u] \leftarrow g[u] + h(u)$            # memorize  $f[u]$  as well as  $g[u]$ 
5           $\text{Parent}[u] \leftarrow v$                                            # also memorize  $\text{Parent}$ 
6           $\text{Insert}_f(\text{OPEN}, u)$ 
7      else if  $u \in \text{OPEN}$  then                                           # "relax" edge  $(v, u)$  if  $u \in \text{OPEN}$ 
8          if  $g[v] + c(v, u) < g[u]$  then                                     # if it gives a better path than the current one
9               $g[u] \leftarrow g[v] + c(v, u)$ ;  $f[u] \leftarrow g[u] + h(u)$      # update  $g$  and  $f$ 
10              $\text{Parent}[u] \leftarrow v$                                        # update  $\text{Parent}$ , too
11     else                                                                 # if  $u \in \text{CLOSED}$ , "relax" edge  $(v, u)$  and re-open  $u$  if necessary
12         if  $g[v] + c(v, u) < g[u]$  then                                     # if a cheaper path is found
13              $g[u] \leftarrow g[v] + c(v, u)$ ;  $f[u] \leftarrow g[u] + h(u)$      # update  $g$  and  $f$ 
14              $\text{Parent}[u] \leftarrow v$                                        # update  $\text{Parent}$ 
15              $\text{CLOSED} \leftarrow \text{CLOSED} \setminus \{u\}$                    # take  $u$  out of  $\text{CLOSED}$ 
16              $\text{Insert}_f(\text{OPEN}, u)$                                        # and put it back in  $\text{OPEN}$ 
```

This test never succeeds if h is monotone

Procedure Expand(v)

for A* algorithm **when h is monotone**

63

```
1  foreach  $u \in \text{Succ}(v)$  do
2    if  $u \notin \text{OPEN} \cup \text{CLOSED}$  then                                     # if  $u$  is a new state
3      Reserve memory for  $g[u]$ ,  $f[u]$ , and  $\text{Parent}[u]$ 
4       $g[u] \leftarrow g[v] + c(v, u)$ ;  $f[u] \leftarrow g[u] + h(u)$            # memorize  $f[u]$  as well as  $g[u]$ 
5       $\text{Parent}[u] \leftarrow v$                                            # also memorize  $\text{Parent}$ 
6       $\text{Insert}_f(\text{OPEN}, u)$ 
7    else if  $u \in \text{OPEN}$  then                                           # "relax" edge  $(v, u)$  if  $u \in \text{OPEN}$ 
8      if  $g[v] + c(v, u) < g[u]$  then                                     # if it gives a better path than the current one
9         $g[u] \leftarrow g[v] + c(v, u)$ ;  $f[u] \leftarrow g[u] + h(u)$        # update  $g$  and  $f$ 
10        $\text{Parent}[u] \leftarrow v$                                          # update  $\text{Parent}$ , too
```

Thus, this part can be safely removed if we know h is monotone for sure

Dijkstra and A* : Difference in Expand(v)

Case	Dijkstra	A*
$u \notin \text{OPEN}$ nor $u \notin \text{CLOSED}$	$g[u] \leftarrow g[v] + c(v, u)$ $\text{Insert}_g(\text{OPEN}, u)$ $\text{Parent}[u] \leftarrow v$	$g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ $\text{Insert}_f(\text{OPEN}, u)$ $\text{Parent}[u] \leftarrow v$
$u \in \text{OPEN}$	if $g[v] + c(v, u) < g[u]$ then $g[u] \leftarrow g[v] + c(v, u)$ $\text{Parent}[u] \leftarrow v$	if $g[v] + c(v, u) < g[u]$ then $g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ $\text{Parent}[u] \leftarrow v$
$u \in \text{CLOSED}$	Do nothing (always $g[u] \leq g[v] + c(v, u)$)	if $g[v] + c(v, u) < g[u]$ then $g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ $\text{Parent}[u] \leftarrow v$ $\text{CLOSED} \leftarrow \text{CLOSED} \setminus \{u\}$ $\text{Insert}_f(\text{OPEN}, u)$

v = node just expanded / u = a successor of v

Dijkstra and A* with monotone h : Difference in Expand(v)


Case	Dijkstra	A*
$u \notin \text{OPEN}$ nor $u \notin \text{CLOSED}$	$g[u] \leftarrow g[v] + c(v, u)$ Insert $_g(\text{OPEN}, u)$ Parent $[u] \leftarrow v$	$g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ Insert $_f(\text{OPEN}, u)$ Parent $[u] \leftarrow v$
$u \in \text{OPEN}$	if $g[v] + c(v, u) < g[u]$ then $g[u] \leftarrow g[v] + c(v, u)$ Parent $[u] \leftarrow v$	if $g[v] + c(v, u) < g[u]$ then $g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ Parent $[u] \leftarrow v$
$u \in \text{CLOSED}$	Do nothing (always $g[u] \leq g[v] + c(v, u)$)	Do nothing (always $g[u] \leq g[v] + c(v, u)$)

v = node just expanded / u = a successor of v

How difficult is it to design a monotone heuristic function?

Good News!

Almost all well-known “natural” heuristics (e.g., those computed from relaxed problems) are monotone



Note: the heuristic used for the exercise was artificially constructed

It was

- ▶ admissible
- ▶ but **not** monotone