

3010

ARTIFICIAL INTELLIGENCE

Lecture 3 A* search

Masashi Shimbo

2019-05-15

License: © ⓘ CC BY-SA 4.0 except where noted

Today's agenda

2

- ▶ Heuristic evaluation function
- ▶ The A* algorithm

Heuristic search ヒューリスティック探索

3

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

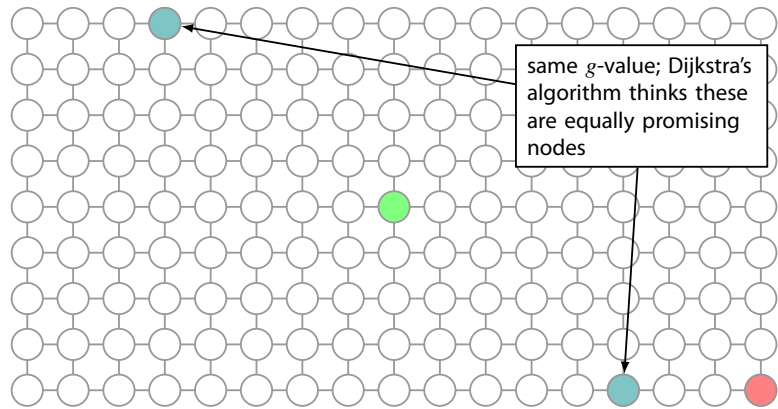
Here is a motivating example...

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いても...

5



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

ラ法はこの情報を活用した探索を行わない

Heuristic evaluation function $h(v)$

7

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」

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


6

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?



8

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^* .

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)

→ $h \leq h^*$

9

Admissible heuristics and the A* algorithm

適格なヒューリスティック関数と A*

10


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)

Gridworld




11

12

Gridworld

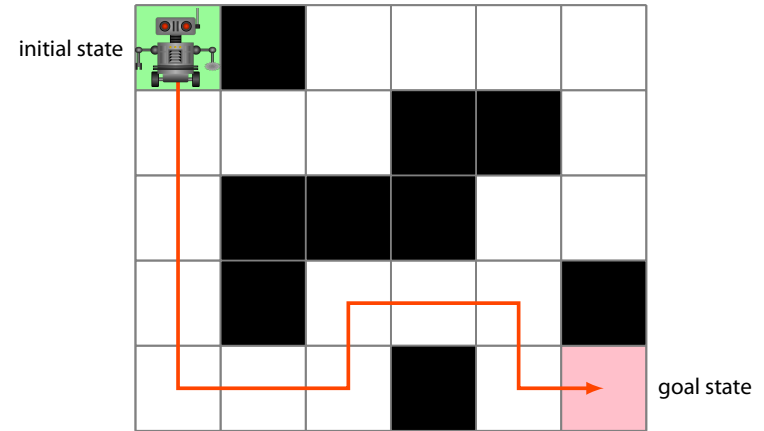
Many obstacles exist 障害物が多数存在



13

Gridworld

Many obstacles exist 障害物が多数存在



14

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



15


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



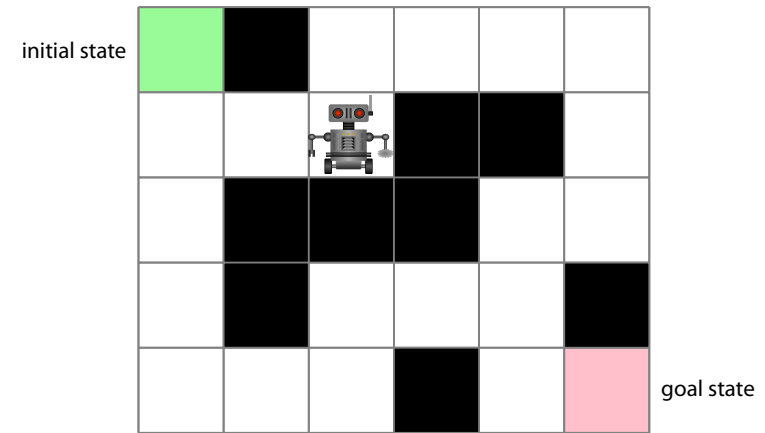
16

Gridworld: Manhattan-distance heuristic




17

Gridworld: Manhattan-distance heuristic



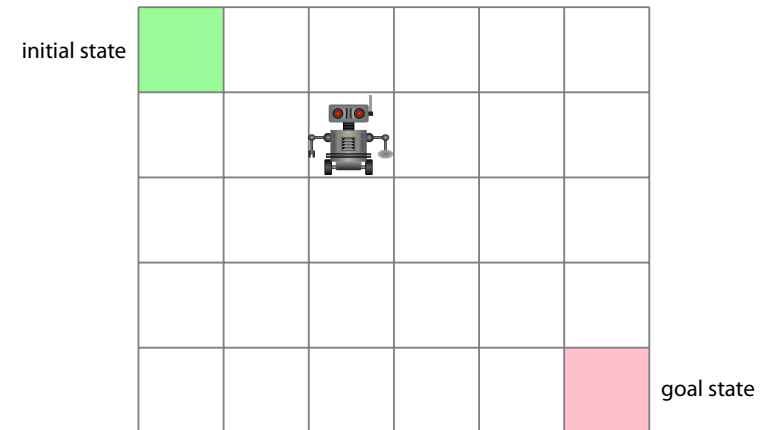
18

Gridworld: Manhattan-distance heuristic




19

Gridworld: Manhattan-distance heuristic



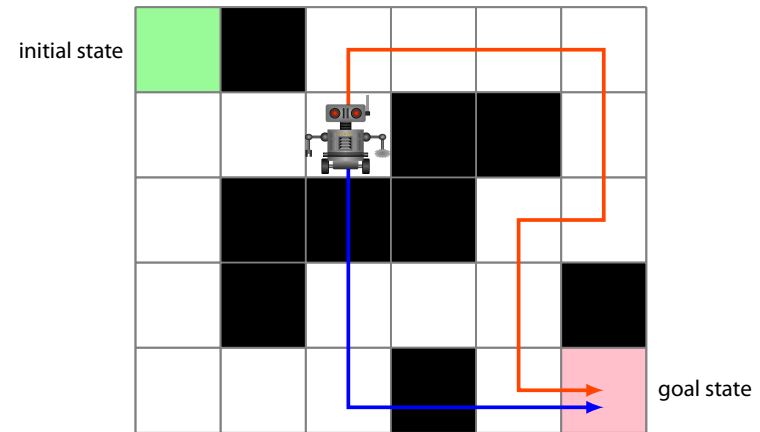
20

Gridworld: Manhattan-distance heuristic



21

Gridworld: Manhattan-distance heuristic



22

Gridworld: Manhattan-distance heuristic

Computation is easy

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.

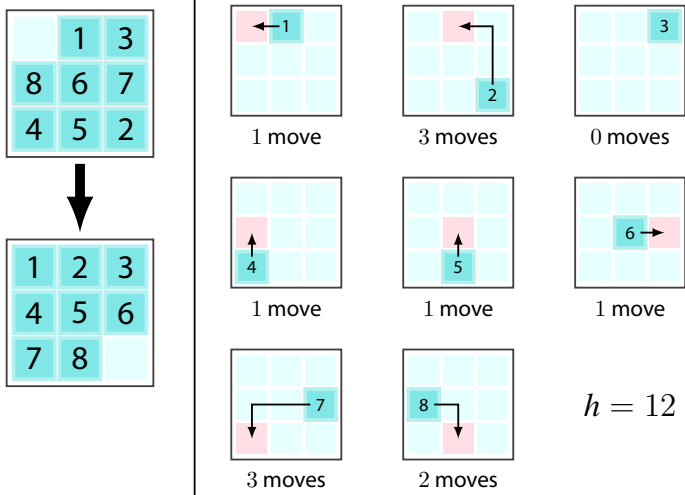


Relaxed problem: tiles can overlap with each other.

23

24

Eight-puzzle: Manhattan distance heuristic



25

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

26

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

27

The A* algorithm A* アルゴリズム

28

The A* algorithm

[Hart, Nilsson & Raphael 1968]

Idea

Use the **sum** of

- ▶ the path cost from the initial state to state v $g[v]$
- ▶ the **estimated cost** from v to a goal state $h(v)$

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

- ▶ $g[v]$: 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.

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

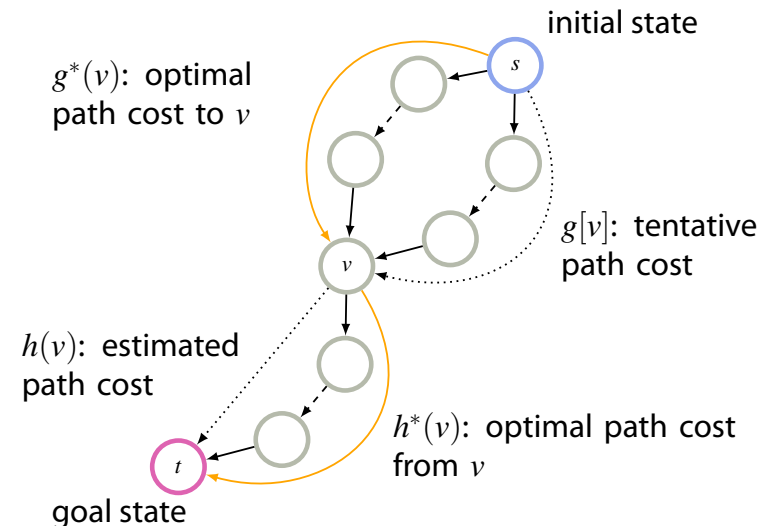
$h(v)$ estimated cost from state v to the nearest goal state v から最も近い目標節点への経路コストの見積もり

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

31



32

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

33

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

Priority queue

Two functions for manipulating priority queue P_f :

Insert _{f} (P_f, v)
Put item v in P_f

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

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

before removal

The A* algorithm

34

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**.

How a node changes its status in A*

Closed nodes can be re-opened

36

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) # OPEN: set of states generated but not expanded
4 CLOSED ← ∅ # CLOSED: set of expanded states
5 loop do
6   if IsEmpty(OPEN) then return "failure"
7   v ← DeleteMing(OPEN) # choose a node with the smallest g
8   CLOSED ← CLOSED ∪ {v} # put v in CLOSED
9   if IsGoal(v) then return Solution(v, s)
10  Expand(v)

```

procedure Expand(v) for Dijkstra's algorithm

```

1 foreach u ∈ Succ(v) do
2   if u ∉ OPEN ∪ CLOSED then # if u is a new state
3     Reserve memory for g[u], Parent[u]
4     g[u] ← g[v] + c(v, u) # memorize g[u]
5     Parent[u] ← v # also memorize Parent
6     Insertg(OPEN, u)
7   else if u ∈ OPEN then # "relax" edge (v, u) if u ∈ OPEN
8     if g[v] + c(v, u) < g[u] then # if it is better than the current path
9       g[u] ← g[v] + c(v, u) # update g if path through (v, u) is shorter
10      Parent[u] ← v # update Parent, too

```

The A* algorithm

Main routine

```

1 OPEN ← new PriorityQueuef # priority is based on f = g + h
2 g[s] ← 0; f[s] ← 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 ← ∅ # CLOSED: set of expanded states
5 loop do
6   if IsEmpty(OPEN) then return "failure"
7   v ← DeleteMinf(OPEN) # choose a node with the smallest f
8   CLOSED ← CLOSED ∪ {v}
9   if IsGoal(v) then return Solution(v, s)
10  Expand(v)

```

procedure Expand(v) for A* algorithm

```

1 foreach u ∈ Succ(v) do
2   if u ∉ OPEN ∪ CLOSED then # if u is a new state
3     Reserve memory for g[u], f[u], and Parent[u]
4     g[u] ← g[v] + c(v, u); f[u] ← g[u] + h(u) # memorize f[u] as well as g[u]
5     Parent[u] ← v # also memorize Parent
6     Insertf(OPEN, u)
7   else if u ∈ OPEN then # "relax" edge (v, u) if u ∈ OPEN
8     if g[v] + c(v, u) < g[u] then # if path through (v, u) is shorter
9       g[u] ← g[v] + c(v, u); f[u] ← g[u] + h(u) # update g and f
10      Parent[u] ← v # update Parent, too
11   else # if u ∈ 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] ← g[v] + c(v, u); f[u] ← g[u] + h(u) # update g and f
14       Parent[u] ← v # update Parent, too
15       CLOSED ← CLOSED \ {u} # then take u out of CLOSED
16       Insertf(OPEN, u) # and put it back into OPEN



```

A* and Dijkstra's algorithm: 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 $g[u] \leftarrow g[v] + c(v, u)$ $f[u] \leftarrow g[u] + h(u)$ $\text{CLOSED} \leftarrow \text{CLOSED} \setminus \{u\}$ $\text{Insert}_f(\text{OPEN}, u)$ $\text{Parent}[u] \leftarrow v$

41

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



43

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

42

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

44

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

Monotonicity and admissibility


Monotonicity implies admissibility

h is monotone $\rightarrow h$ is admissible

Monotone heuristic function

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

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



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

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] ← 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)

```

49

Procedure Expand(v)

for A* algorithm when h is monotone

```

1 foreach u ∈ Succ(v) do
2   if u ∉ OPEN ∪ CLOSED then # if u is a new state
3     Reserve memory for g[u], f[u], and Parent[u]
4     g[u] ← g[v] + c(v, u); f[u] ← g[u] + h(u) # memorize f[u] as well as g[u]
5     Parent[u] ← v # also memorize Parent
6     Insertf(OPEN, u)
7   else if u ∈ OPEN then # "relax" edge (v, u) if u ∈ OPEN
8     if g[v] + c(v, u) < g[u] then # if it gives a better path than the current one
9       g[u] ← g[v] + c(v, u); f[u] ← g[u] + h(u) # update g and f
10      Parent[u] ← v # update Parent, too
11    else # This test never succeeds if h is monotone # if u ∈ 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] ← g[v] + c(v, u); f[u] ← g[u] + h(u) # update g and f
14        Parent[u] ← v # update Parent
15        CLOSED ← CLOSED \ {u} # take u out of CLOSED
16        Insertf(OPEN, u) # and put it back in OPEN

```

50

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

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

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

51


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

How difficult is it to design a monotone heuristic function?

52

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