

Informed Search

Dr. Neil T. Dantam

CSCI-534, Colorado School of Mines

Spring 2020

Introduction

Uninformed vs. Informed Search

Definition: Uninformed Search

Search that has no additional information about states beyond the problem definition, i.e., can only distinguish between goal and non-goal states. Also called *blind search*.

(last lecture)

Definition: Informed Search

Search that can consider whether some non-goal states are “more promising” than others. Also called *heuristic search*.

(this lecture)

Outcomes

- ▶ Understand common informed search algorithms / algorithm families:
 - ▶ Dijkstra's algorithm
 - ▶ Greedy Search
 - ▶ A*
- ▶ Evaluate the applicability and trade-offs of and common informed search methods:
 - ▶ Can we apply a particular algorithm to our problem?
(do we have the right inputs / given information?)
 - ▶ Will a particular algorithm be optimal?
 - ▶ Can we expect a particular algorithm perform “well?”

Outline

Planning and Search Problems

Dijkstra's Algorithm

Optimality and Heuristics

Greedy Search


A* Search

The Planning / Search Problem

- Given:
1. State space: \mathcal{Q}
 2. Transition function $\delta : \mathcal{Q} \mapsto \mathcal{P}(\mathcal{Q})$
 3. Start state: $q_0 \in \mathcal{Q}$
 4. Goal set: $A \subseteq \mathcal{Q}$

Find: Path $p = (p_0, \dots, p_n)$ from start to goal such that:

- ▶ $p_0 = q_0$ is the start state
- ▶ $p_n \in A$ is a goal state
- ▶ Subsequent states are valid transitions: $p_{k+1} \in \delta(p_k)$



Forward Search

Procedure search($f_{\text{ins}}, f_{\text{rem}}, q_0, \delta, A$)

```

1   $T[q_0] \leftarrow \text{NIL};$  // Search Tree
2   $W \leftarrow f_{\text{ins}}(q_0, \text{NIL});$  // Frontier
3  while  $W$  do
4      let  $q \leftarrow f_{\text{rem}}(W)$  in
5          if  $q \in A$  then
6              return tree-path( $T, q$ );
7          else
8              foreach  $q' \in \delta(q)$  do
9                  if  $\neg \text{contains}(T, q')$  then
10                      $T[q'] \leftarrow q$ ;
11                      $W \leftarrow f_{\text{ins}}(q', W)$ ;
12 return NIL;
```

Procedure depth-first-search(q_0, δ, A)

```
1 return search(push, pop,  $q_0, \delta, A$ );
```

Procedure breadth-first-search(q_0, δ, A)

```
1 return search(enqueue, dequeue,  $q_0, \delta, A$ );
```

Planning Properties

Correctness: Do we get a right answer?

Completeness: Do we always get an answer?

Optimality: Do we get the best answer?

Optimality

Definition (Optimality)

A planning algorithm is optimal if it produces the lowest cost (/ highest reward) plan.

State Cost

▶ $\delta : \mathcal{Q} \mapsto \mathcal{P}(\mathcal{Q})$

▶ $q_{i+1} \in \delta(q_i)$

▶ $C : \mathcal{Q} \mapsto \mathbb{R}$

$$p = \operatorname{argmin}_{p_0, \dots, p_n} \sum_{i=0}^n C(p_i)$$

Transition Cost

▶ $\delta : \mathcal{Q} \mapsto \mathcal{P}(\mathcal{Q})$

▶ $q_{i+1} \in \delta(q_i)$

▶ $C : \mathcal{Q} \times \mathcal{Q} \mapsto \mathbb{R}$

$$p = \operatorname{argmin}_{p_0, \dots, p_n} \sum_{i=0}^{n-1} C(p_i, p_{i+1})$$

Action Cost

▶ $\delta : \overbrace{\mathcal{Q}}^{\text{state}} \times \overbrace{\mathcal{U}}^{\text{action}} \mapsto \mathcal{Q}$

▶ $q_{i+1} = \delta(q_i, u_i)$

▶ $C : \mathcal{U} \mapsto \mathbb{R}$

$$u = \operatorname{argmin}_{u_0, \dots, u_n} \sum_{i=0}^n C(u_i)$$

Outline

Planning and Search Problems

Dijkstra's Algorithm

Optimality and Heuristics

Greedy Search

A* Search

Dijkstra's Algorithm Overview

Steps

1. Store node work list in a priority queue
2. Order priority queue by cost to reach node
3. Each iteration, visit the least-cost unexplored node

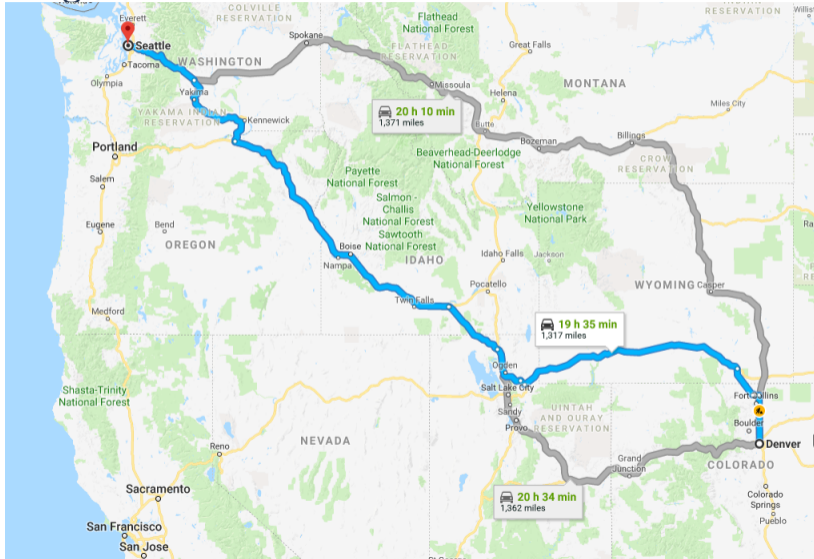
Costs

Given: Cost of an individual node/transition: $C(q)$

Computed: Total cost to reach a node from the start:

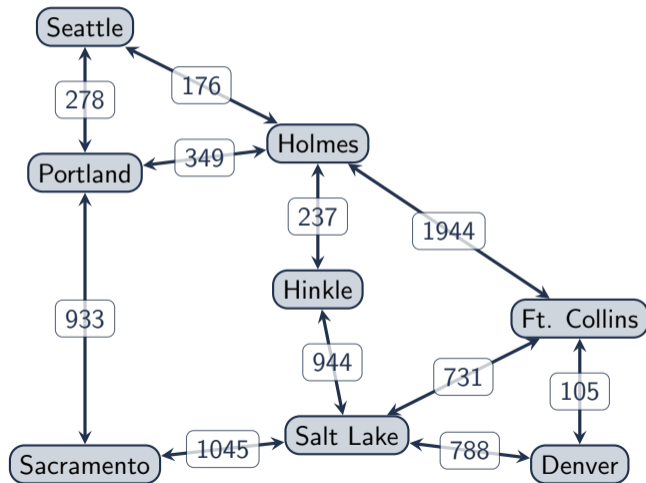
$$D(q) = C(q) + D(\text{parent}(q))$$

Example: Navigation



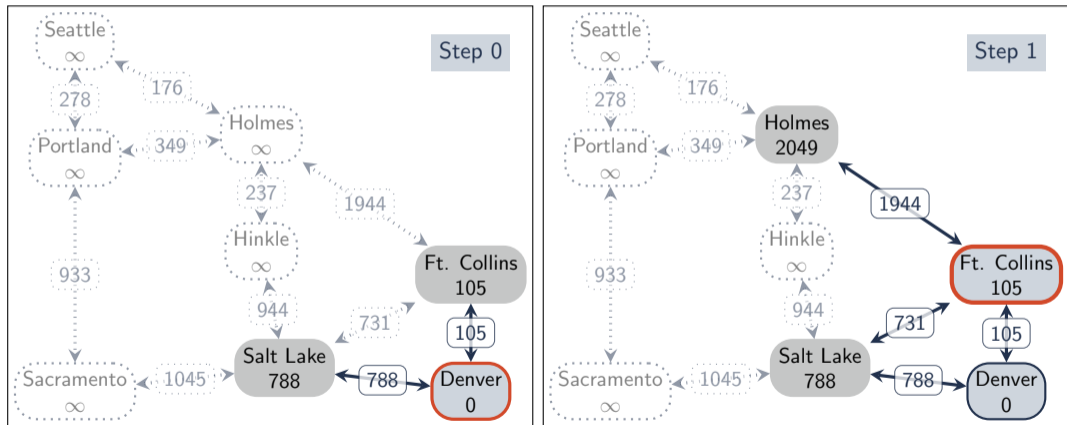
Example: Dijkstra's Algorithm

Graph



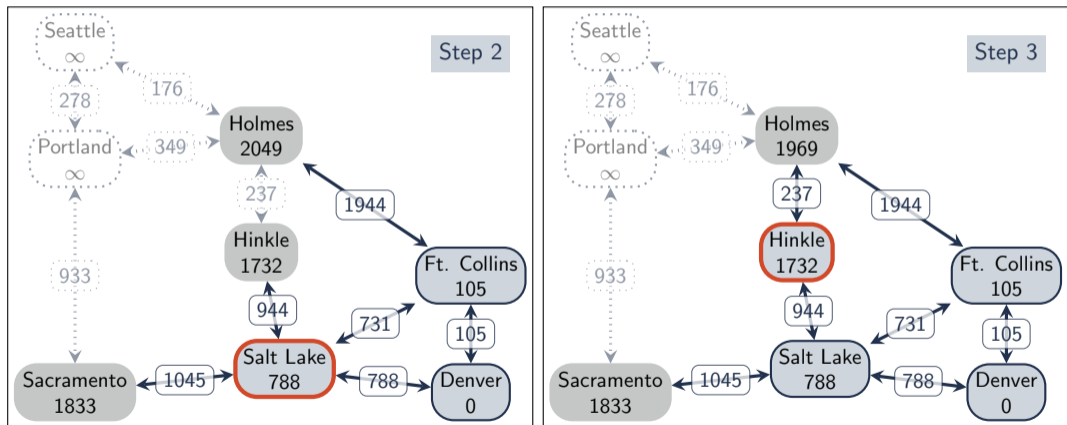
Example: Dijkstra's Algorithm

continued – 1



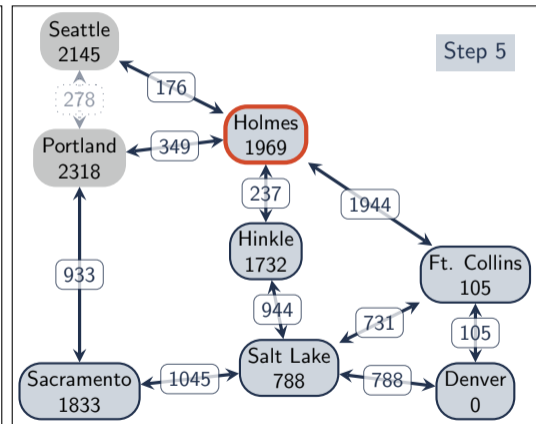
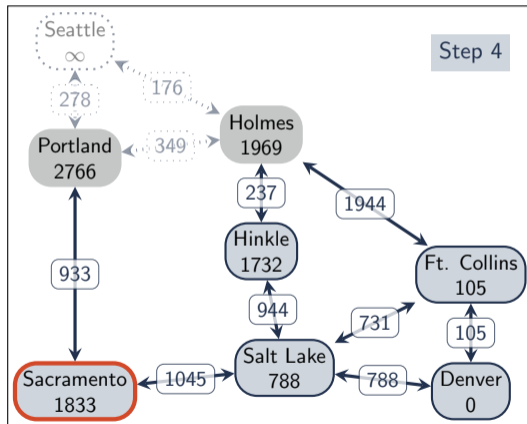
Example: Dijkstra's Algorithm

continued – 2



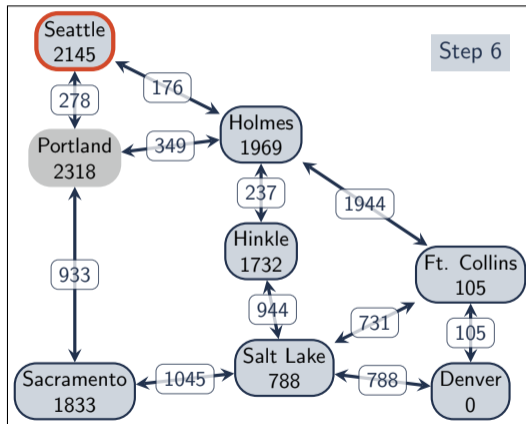
Example: Dijkstra's Algorithm

continued – 2



Example: Dijkstra's Algorithm

continued – 3



Dijkstra's Algorithm – Transition Cost


Procedure dijkstra-transition(C, q_0, δ, A)

```

1   $T[q_0] \leftarrow \text{NIL}$ ; // Search Tree: node  $\mapsto$  parent
2   $D[q_0] \leftarrow 0$ ; // Cost: node  $\mapsto$  path cost
3   $W \leftarrow \text{insert}(q_0, 0)$ ; // Priority Queue
4  while  $W$  do
5      let  $q \leftarrow \text{remove-min}(W)$  in
6          if  $q \in A$  then return tree-path( $T, q$ );
7          else
8              foreach  $q' \in \delta(q)$  do
9                  let  $d' \leftarrow D[q] + C(q, q')$  in
10                     if  $\neg \text{contains}(T, q') \vee (d' < D[q'])$  then
11                          $T[q'] \leftarrow q$ ;
12                          $D[q'] \leftarrow d'$ ;
13                          $W \leftarrow \text{insert}(q', d')$ ;
14 return NIL;
```

Exercise: Dijkstra's Algorithm

2	3	4	☆2
1	2	1	2
0	1	3	1
	0	1	2



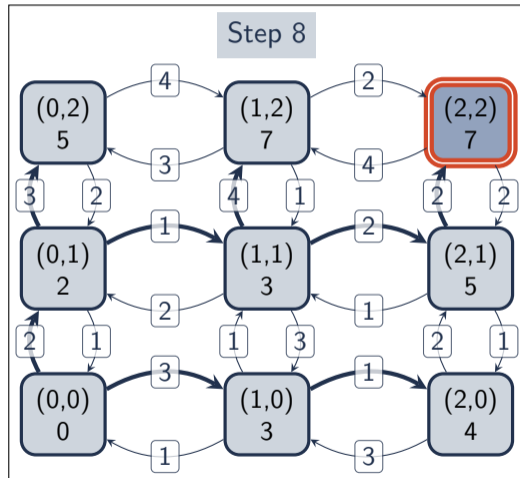
Exercise: Dijkstra's Algorithm

Exercise: Dijkstra's Algorithm

Exercise: Dijkstra's Algorithm

Exercise: Dijkstra's Algorithm

Exercise: Dijkstra's Algorithm



Dijkstra's Algorithm – Transition vs. State Cost

Procedure djikstra-transition(C, q_0, δ, A)

```

1  $T[q_0] \leftarrow \text{NIL}$ ; // Search Tree: node  $\mapsto$  parent
2  $D[q_0] \leftarrow 0$ ; // Cost: node  $\mapsto$  path cost
3  $W \leftarrow \text{insert}(q_0, 0)$ ; // Priority Queue
4 while  $W$  do
5   let  $q \leftarrow \text{remove-min}(W)$  in
6   if  $q \in A$  then return tree-path( $T, q$ );
7   else
8     foreach  $q' \in \delta(q)$  do
9       let  $d' \leftarrow D[q] + \overbrace{C(q, q')}^{\text{transition cost}}$  in
10      if
11         $\neg \text{contains}(T, q') \vee (d' < D[q'])$ 
12      then
13         $T[q'] \leftarrow q$ ;
14         $D[q'] \leftarrow d'$ ;
15         $W \leftarrow \text{insert}(q', d')$ ;
16 return NIL;
```

Procedure djikstra-state(C, q_0, δ, A)

```

1  $T[q_0] \leftarrow \text{NIL}$ ; // Search Tree: node  $\mapsto$  parent
2  $D[q_0] \leftarrow 0$ ; // Cost: node  $\mapsto$  path cost
3  $W \leftarrow \text{insert}(q_0, C(q_0))$ ; // Priority Queue
4 while  $W$  do
5   let  $q \leftarrow \text{remove-min}(W)$  in
6   if  $q \in A$  then return tree-path( $T, q$ );
7   else
8     foreach  $q' \in \delta(q)$  do
9       let  $d' \leftarrow D[q] + \overbrace{C(q')}^{\text{state cost}}$  in
10      if
11         $\neg \text{contains}(T, q') \vee (d' < D[q'])$ 
12      then
13         $T[q'] \leftarrow q$ ;
14         $D[q'] \leftarrow d'$ ;
15         $W \leftarrow \text{insert}(q', d')$ ;
16 return NIL;
```


Exercise: Dijkstra's Algorithm – Action Cost

Transition Function:

$$\delta : \underbrace{Q}_{\text{predecessor}} \times \underbrace{U}_{\text{action}} \mapsto \underbrace{Q}_{\text{sucessor}}$$

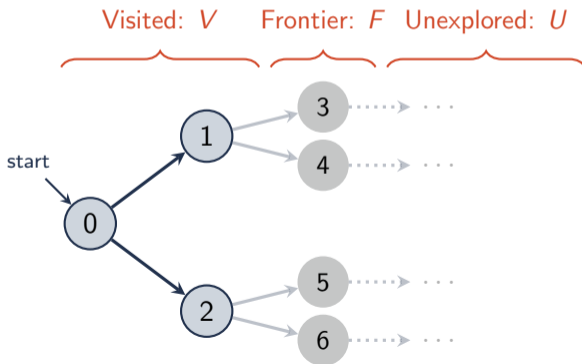
Cost Function:

$$C : U \mapsto \mathbb{R}$$

Dijkstra's Algorithm is Optimal

Proof Outline

Proof by induction: For each iteration of the loop, we visit the next unexplored node with least cost from q_0 . □



$$\blacktriangleright \tilde{C}(q_i, q_j) = \min \left(\sum_{i=0}^{n-1} C(p_i, p_{i+1}) \right),$$

where:

- $\blacktriangleright p_0 = q_i$
- $\blacktriangleright p_n = q_j$
- $\blacktriangleright p_{i+1} \in \delta p_i$
- $\blacktriangleright \forall q_v \in V, \forall q_f \in F, \tilde{C}(q_0, q_v) \leq \tilde{C}(q_0, q_f)$
- $\blacktriangleright \forall q_v \in V, \forall q_u \in U, \tilde{C}(q_0, q_v) \leq \tilde{C}(q_0, q_u)$

Edsger Wybe Dijkstra

“The question of whether a computer can think is no more interesting than the question of whether a submarine can swim.”

–Edsger W. Dijkstra



Outline

Planning and Search Problems

Dijkstra's Algorithm

Optimality and Heuristics

Greedy Search

A* Search

Bellman's Principle of Optimality

Definition (Principle of Optimality)

"An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision."



Richard Bellman

$$\text{Goal State: } \forall p_a \in A, \quad \overbrace{f_{\text{opt}}(p_a)}^{\text{zero cost at goal}} = 0$$

$$\text{Non-Goal State: } f_{\text{opt}}(p_i) = \underbrace{C(p_i, \underbrace{\delta_{\text{opt}}(p_i)}_{\text{optimal decision}})}_{\text{cost of optimal decision}} + \underbrace{f_{\text{opt}}(\delta_{\text{opt}}(p_i))}_{\text{cost of remaining decisions}}$$

Optimal Cost: As hard to find as optimal decision/policy

Heuristic

Definition (Heuristic)

An approximation used to quickly find a solution.



In search, an approximate ranking of branches at each step:

$$h : \mathcal{Q} \mapsto \mathbb{R}$$

From the Greek *εὕρισκω* “find” or “discover”.

Heuristic approximates cost-to-go $f_{\text{opt}}(p_i)$

Example: Euclidean distance heuristic




$$d_{\text{euclidean}} = \sqrt{x^2 + y^2}$$

- ▶ $h(a) = \|e - a\|_2 = \sqrt{12^2 + 0^2} = 12$
- ▶ $h(b) = \|e - b\|_2 = \sqrt{5^2 + 12^2} = 13$
- ▶ $h(c) = \|e - c\|_2 = \sqrt{6^2 + 8^2} = 10$
- ▶ $h(d) = \|e - d\|_2 = \sqrt{5^2 + 0^2} = 5$
- ▶ $h(e) = \|e - e\|_2 = \sqrt{0^2 + 0^2} = 0$


Exercise: Manhattan Distance Heuristic

3	4	☆2
2	1	2
①	3	1



$$d_{\text{euclidean}} = \sqrt{x^2 + y^2}$$

$$d_{\text{manhattan}} = |x| + |y|$$



Historical Note



Solomon Lefschetz



John McCarthy



Richard Bellman

Outline

Planning and Search Problems

Dijkstra's Algorithm

Optimality and Heuristics

Greedy Search

A* Search

Greedy Search Overview

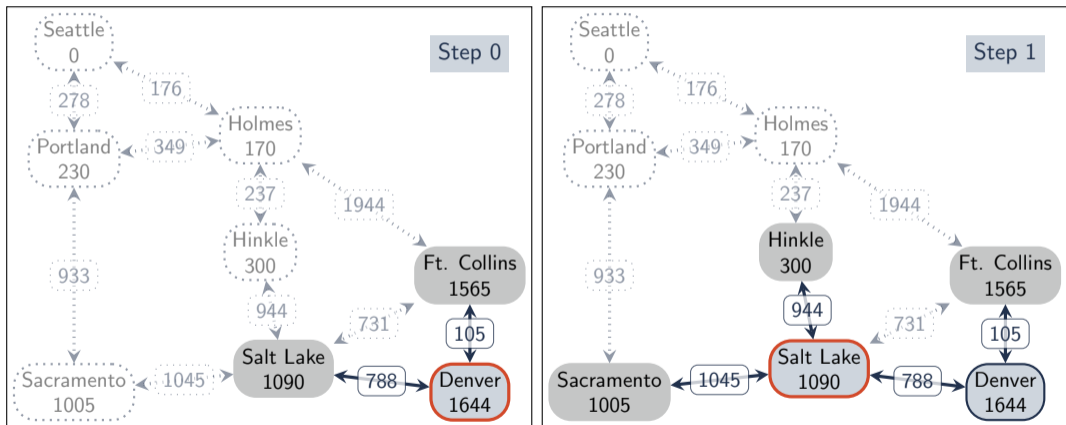
Steps

1. Store node work list in a priority queue
2. Order priority queue by heuristic cost-to-go
3. Each iteration, visit the least-heuristic-cost-to-go unexplored node

Costs

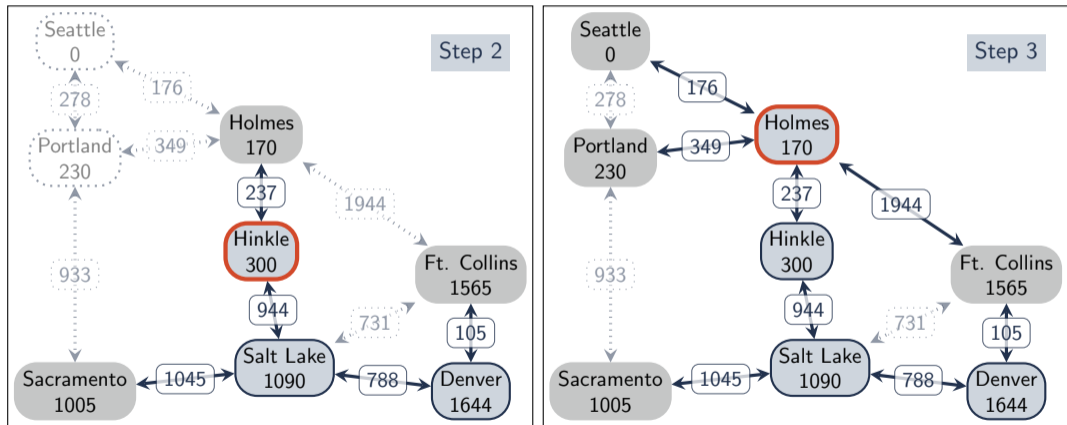
Given: Heuristic cost-to-go: $h(q)$

Example: Greedy Search



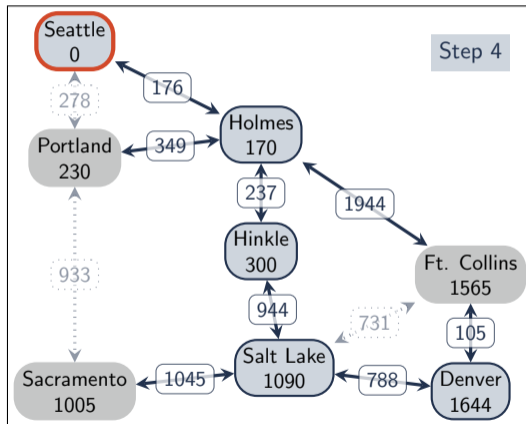
Example: Greedy Search

continued – 1



Example: Greedy Search

continued – 2



Greedy Search Algorithm


Procedure greedy-search(h, q_0, δ, A)

```

1   $T[q_0] \leftarrow \text{NIL}$ ; // Search Tree: node  $\mapsto$  parent
2   $W \leftarrow \text{insert}(q_0, h(q_0))$ ; // Priority Queue
3  while  $W$  do
4      let  $q \leftarrow \text{remove-min}(W)$  in
5          if  $q \in A$  then return tree-path( $T, q$ );
6          else
7              foreach  $q' \in \delta(q)$  do
8                  if  $\neg \text{contains}(T, q')$  then
9                       $T[q'] \leftarrow q$ ;
10                      $W \leftarrow \text{insert}(q', h(q'))$ ;
11 return NIL;
```


Exercise: Greedy Search

3	4	☆2
2	1	2
①	3	1



$$d_{\text{euclidean}} = \sqrt{x^2 + y^2}$$

$$d_{\text{manhattan}} = |x| + |y|$$




Exercise: Greedy Search

Exercise: Greedy Search

Exercise: Greedy Search

Greedy Search is Not Optimal

Proof by Counterexample



Heuristic:
 ▶ $h(a) = 12$
 ▶ $h(b) = 13$
 ▶ $h(c) = 10$
 ▶ $h(d) = 0$

Optimal Path:
 ▶ $p = (a, b, d)$
 ▶ cost = 18

Greedy Path:
 ▶ (a, c, d)
 ▶ cost = 20

Outline

Planning and Search Problems

Dijkstra's Algorithm

Optimality and Heuristics

Greedy Search

A* Search

A* Search Overview

Steps

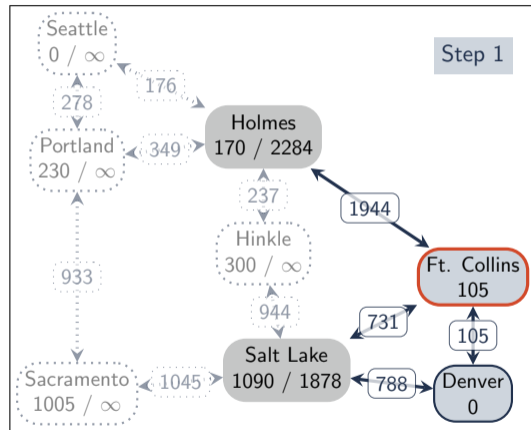
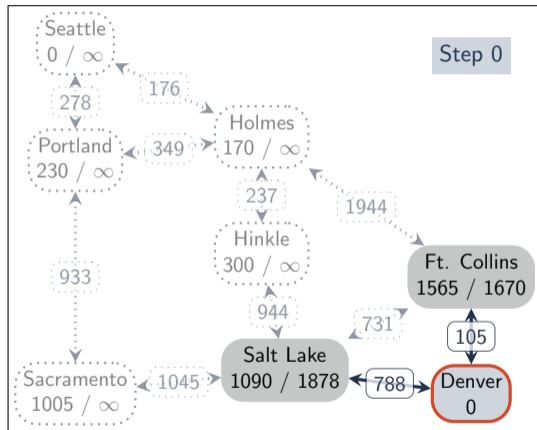
1. Store node work list in a priority queue
2. Order priority queue by sum of cost to reach node and heuristic cost-to-go
3. Each iteration, visit the least expected-total-cost unexplored node

Costs

- Given:
- ▶ Cost of an individual node/transition: $C(q)$
 - ▶ Heuristic cost-to-go: $h(q)$

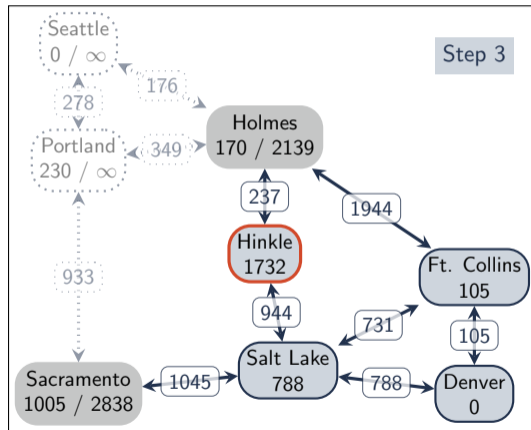

- Computed:
- ▶ Total cost to reach a node from the start:
 $D(q) = C(q) + D(\text{parent}(q))$
 - ▶ Expected total cost:
 $E(q) = h(q) + D(q)$

Example: A* Search



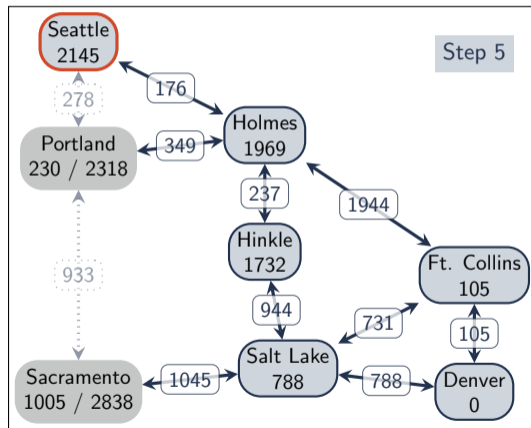
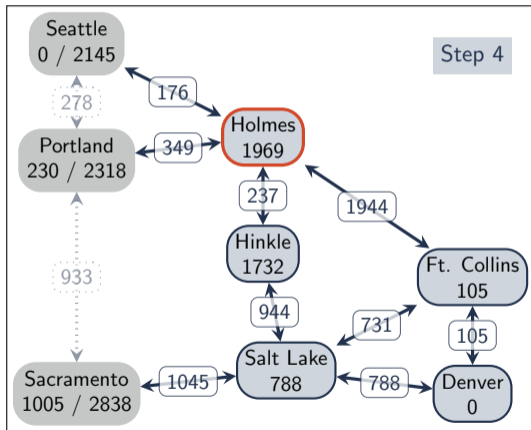
Example: A* Search

continued - 1



Example: A* Search

continued – 2



A* Search Algorithm

Procedure $A^*(C, h, q_0, \delta, A)$


```

1  $T[q_0] \leftarrow \text{NIL}$ ; // Search Tree: node  $\mapsto$  parent
2  $D[q_0] \leftarrow 0$ ; // Cost: node  $\mapsto$  path cost
3  $W \leftarrow \text{insert}(q_0, h(q_0))$ ; // Priority Queue
4 while  $W$  do
5   let  $q \leftarrow \text{remove-min}(W)$  in
6   if  $q \in A$  then return tree-path( $T, q$ );
7   else
8     foreach  $q' \in \delta(q)$  do
9       let  $d' \leftarrow D[q] + C(q, q')$  in
10      if  $\neg \text{contains}(T, q') \vee (d' < D[q'])$  then
11         $T[q'] \leftarrow q$ ;
12         $D[q'] \leftarrow d'$ ;
13         $W \leftarrow \text{insert} \left( q', \underbrace{d' + h(q')}_{\text{total cost estimate}} \right)$ ;
14 return NIL;
```


Exercise: A* Search

3	4	☆2
2	1	2
①	3	1

$$d_{\text{euclidean}} = \sqrt{x^2 + y^2}$$



$$d_{\text{manhattan}} = |x| + |y|$$



Exercise: A* Search

Exercise: A* Search

Exercise: A* Search

Exercise: A* Search

Admissible Heuristic

Definition (Admissible Heuristic)

An admissible heuristic never over-estimates the optimal cost to the goal:

$$\underbrace{h_{\text{admiss}}(q)}_{\text{admissible heuristic}} \leq \underbrace{f_{\text{opt}}(q)}_{\text{optimal cost}}$$

Optimality of A^*

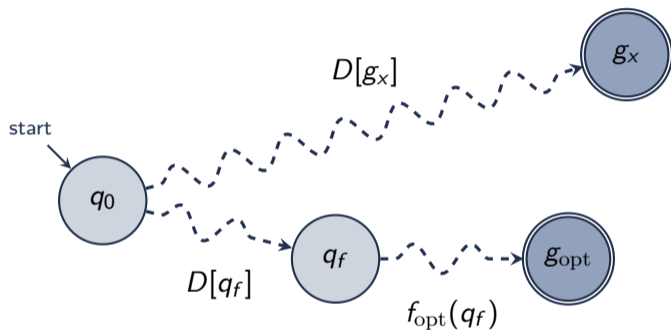
Theorem

When using an admissible heuristic, A^ is optimal.*

Proof by Contradiction

1. Assume we visit goal node g_x (remove from W), where g_x is suboptimal: $D[g_x] > f_{\text{opt}}(q_0)$ and $h(g_x) = 0$
2. If g_x is suboptimal, then the following must hold:
 - 2.1 There exists another path to a goal g_{opt} of cost $f_{\text{opt}}(q_0)$, where $f_{\text{opt}}(q_0) < D[g_x]$
 - 2.2 There is some frontier node q_f in W on the path to g_{opt} where $(D[q_f] + h(q_f)) \leq f_{\text{opt}}(q_0)$
3. But, since q_f has lower value than g_x , we would visit q_f before visiting g_x
4. Contradiction

A* Optimality Illustration



- ▶ $D[g_x] > f_{opt}(q_0)$
- ▶ $D[q_f] + f_{opt}(q_f) = f_{opt}(q_0)$
- ▶ $D[q_f] + h(q_f) \leq f_{opt}(q_0)$

Must visit q_f along optimal path before suboptimal g_x

Summary

Informed Search: Store frontier nodes in a priority queue

Dijkstra's Algorithm: Order queue by cost-to-reach node

- ▶ Optimal

Greedy Search: Order queue heuristic cost-to-go

- ▶ Not optimal

A* Search: Order queue cost-to-reach plus heuristic cost-to-go

- ▶ Optimal (with admissible heuristic)

References

Textbook Russell & Norvig.

- ▶ Ch 3.5. Informed (Heuristic) Search

