

Search

CSCI 4511/6511

Joe Goldfrank

Good Afternoon

- Good afternoon

Announcements

- Homework 1 is due on 7 February at 11:55 PM
 - Automatic extensions
 - Pay attention!


Why Are We Here?

- We're designing rational agents!
 - Perception
 - Logic
 - Action

In Practice


- Environment
 - What happens next
- Perception
 - What agent can see
- Action
 - What agent can do
- Measure/Reward
 - Encoded utility function

Reframed

- Building a model of the real world
 - Model is based on sensor inputs
 - Model is flawed
- Solve problems *on the model*
 - Take actions based on solution
- Model close to reality → solution useful
 - Else: 



Search: Why?

- Fully-observed problem
- Deterministic actions and state
- Well defined *start* and *goal*



State

What is the state space?



State



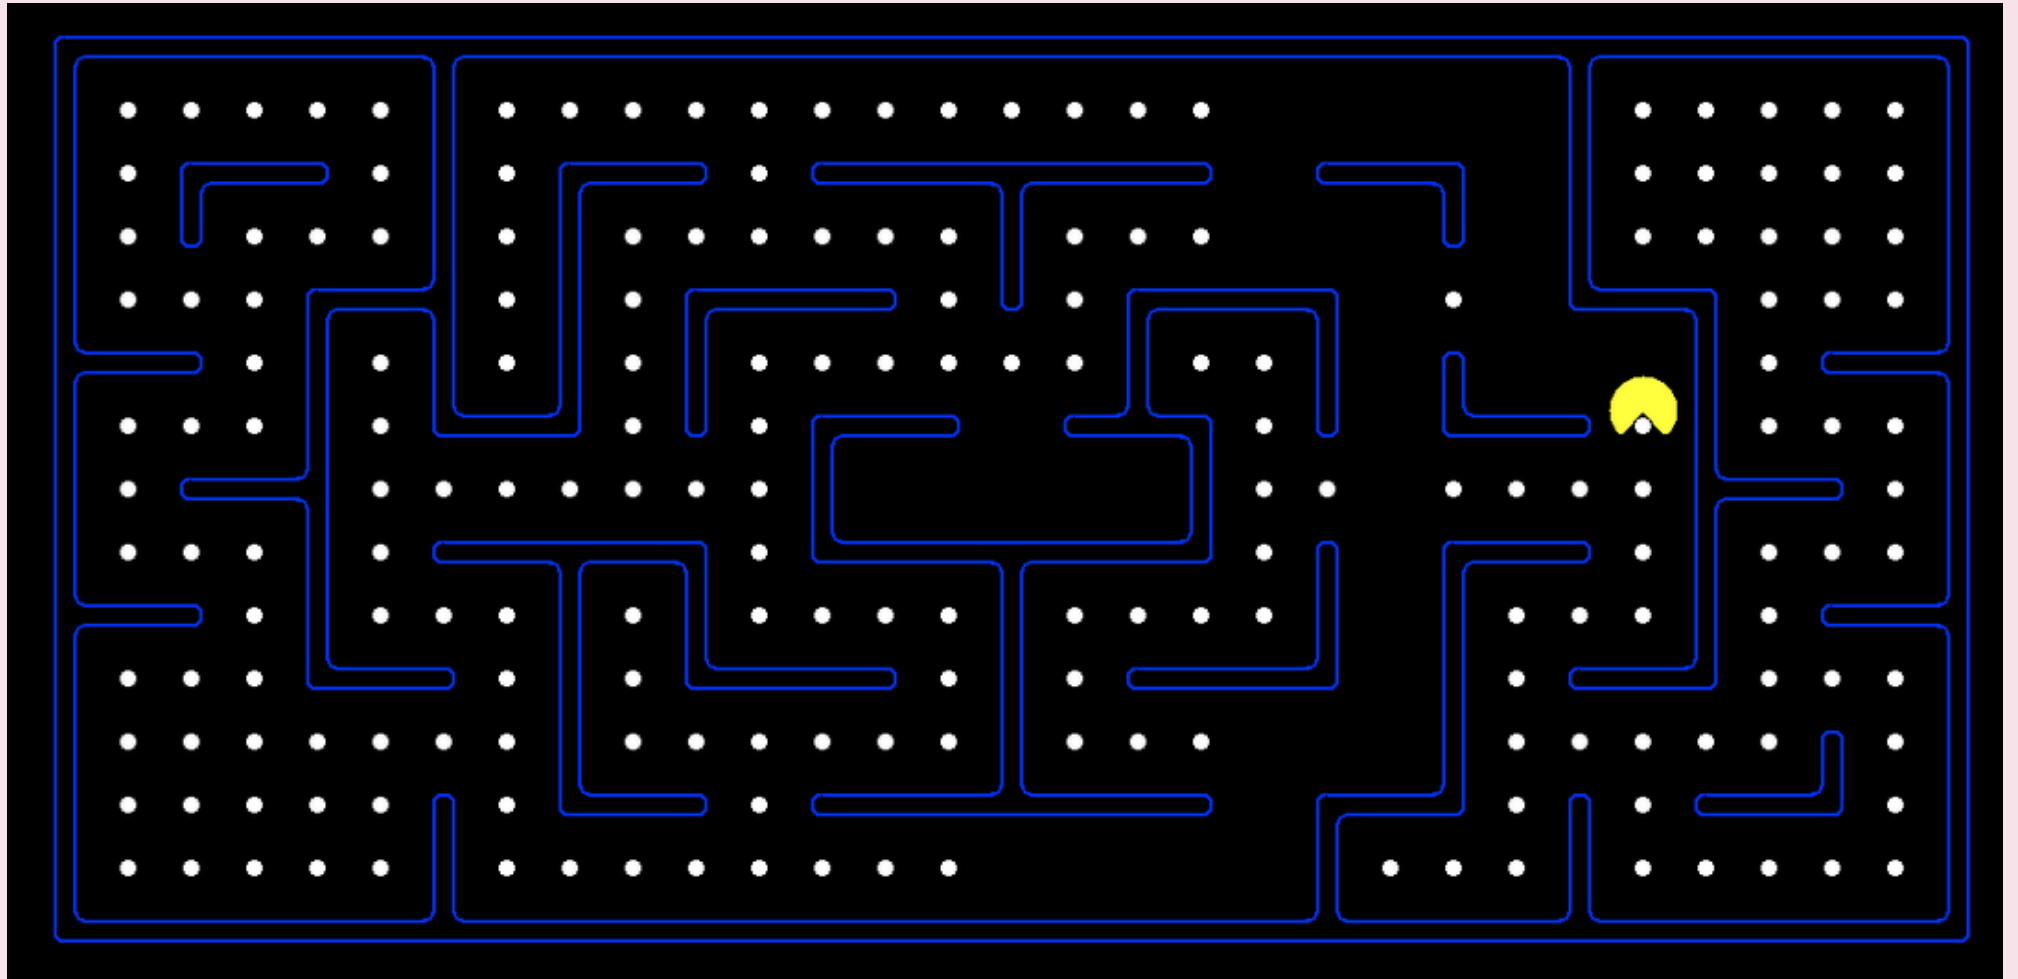
Other Applications

- Route planning
- Protein design
- Robotic navigation
- Scheduling
 - Science
 - Manufacturing

Not Included

- Uncertainty
 - State transitions known
- Adversary
 - Nobody wants us to lose
- Cooperation
- Continuous state

Who Is The Pac-Man?




Search Problem

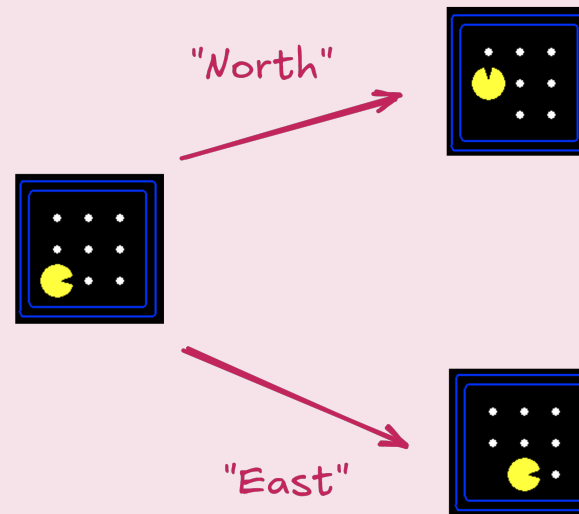
Search problem includes:

- Start State
- State Space
- State Transitions
- Goal Test

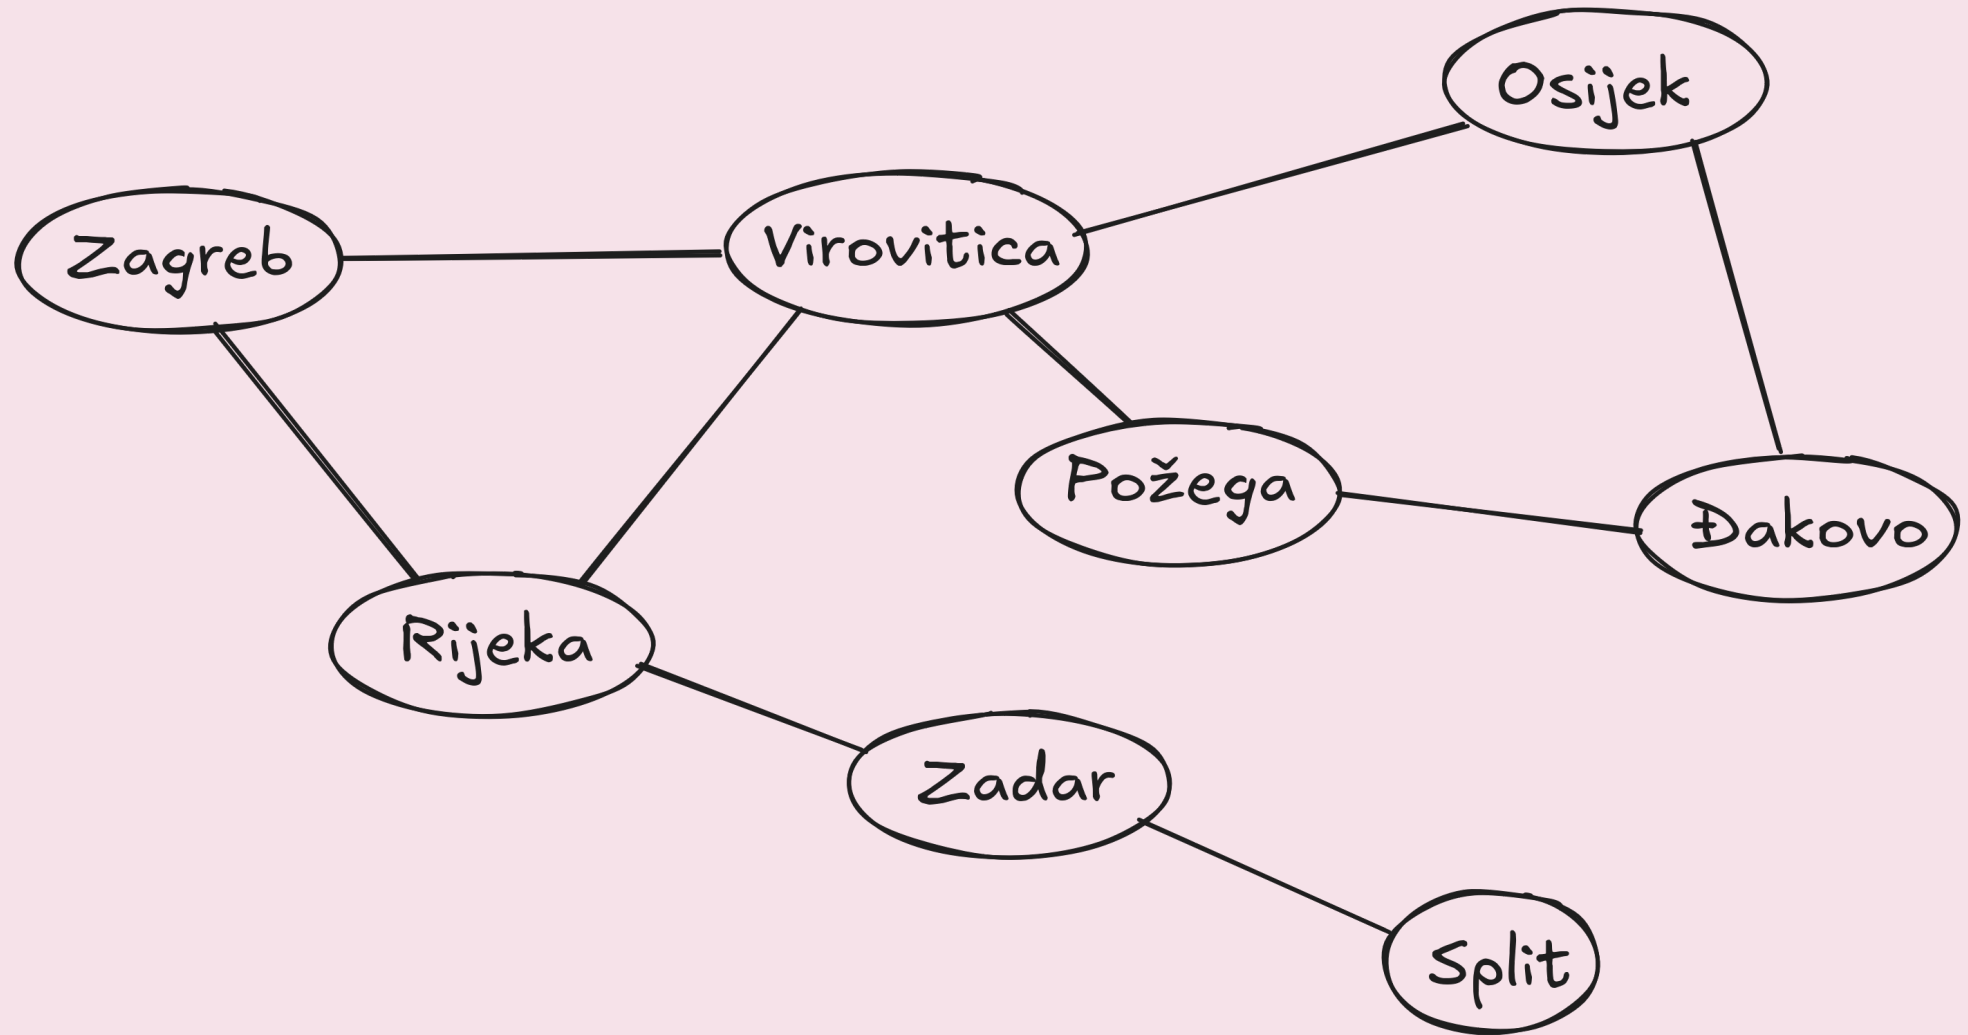
State Space:



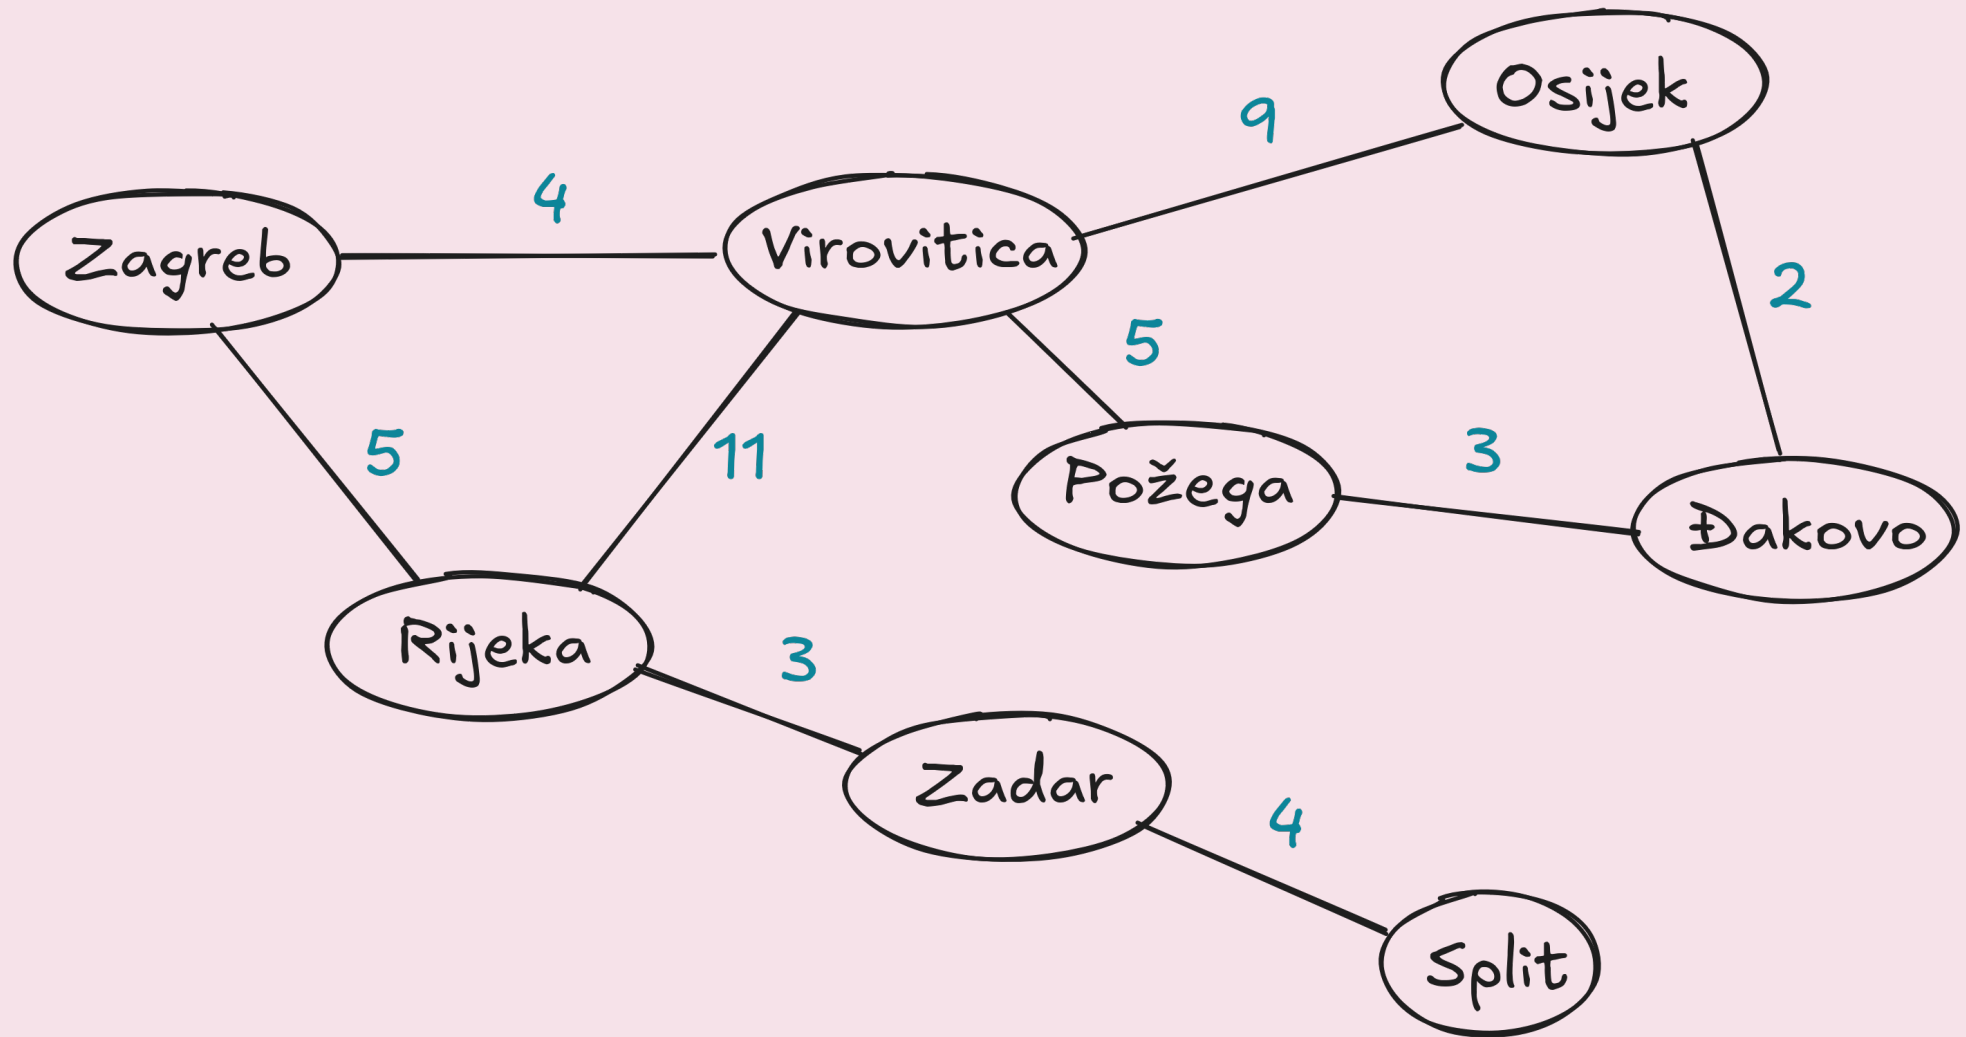
Actions & Successor States:



Tour of Croatia




Tour of Croatia

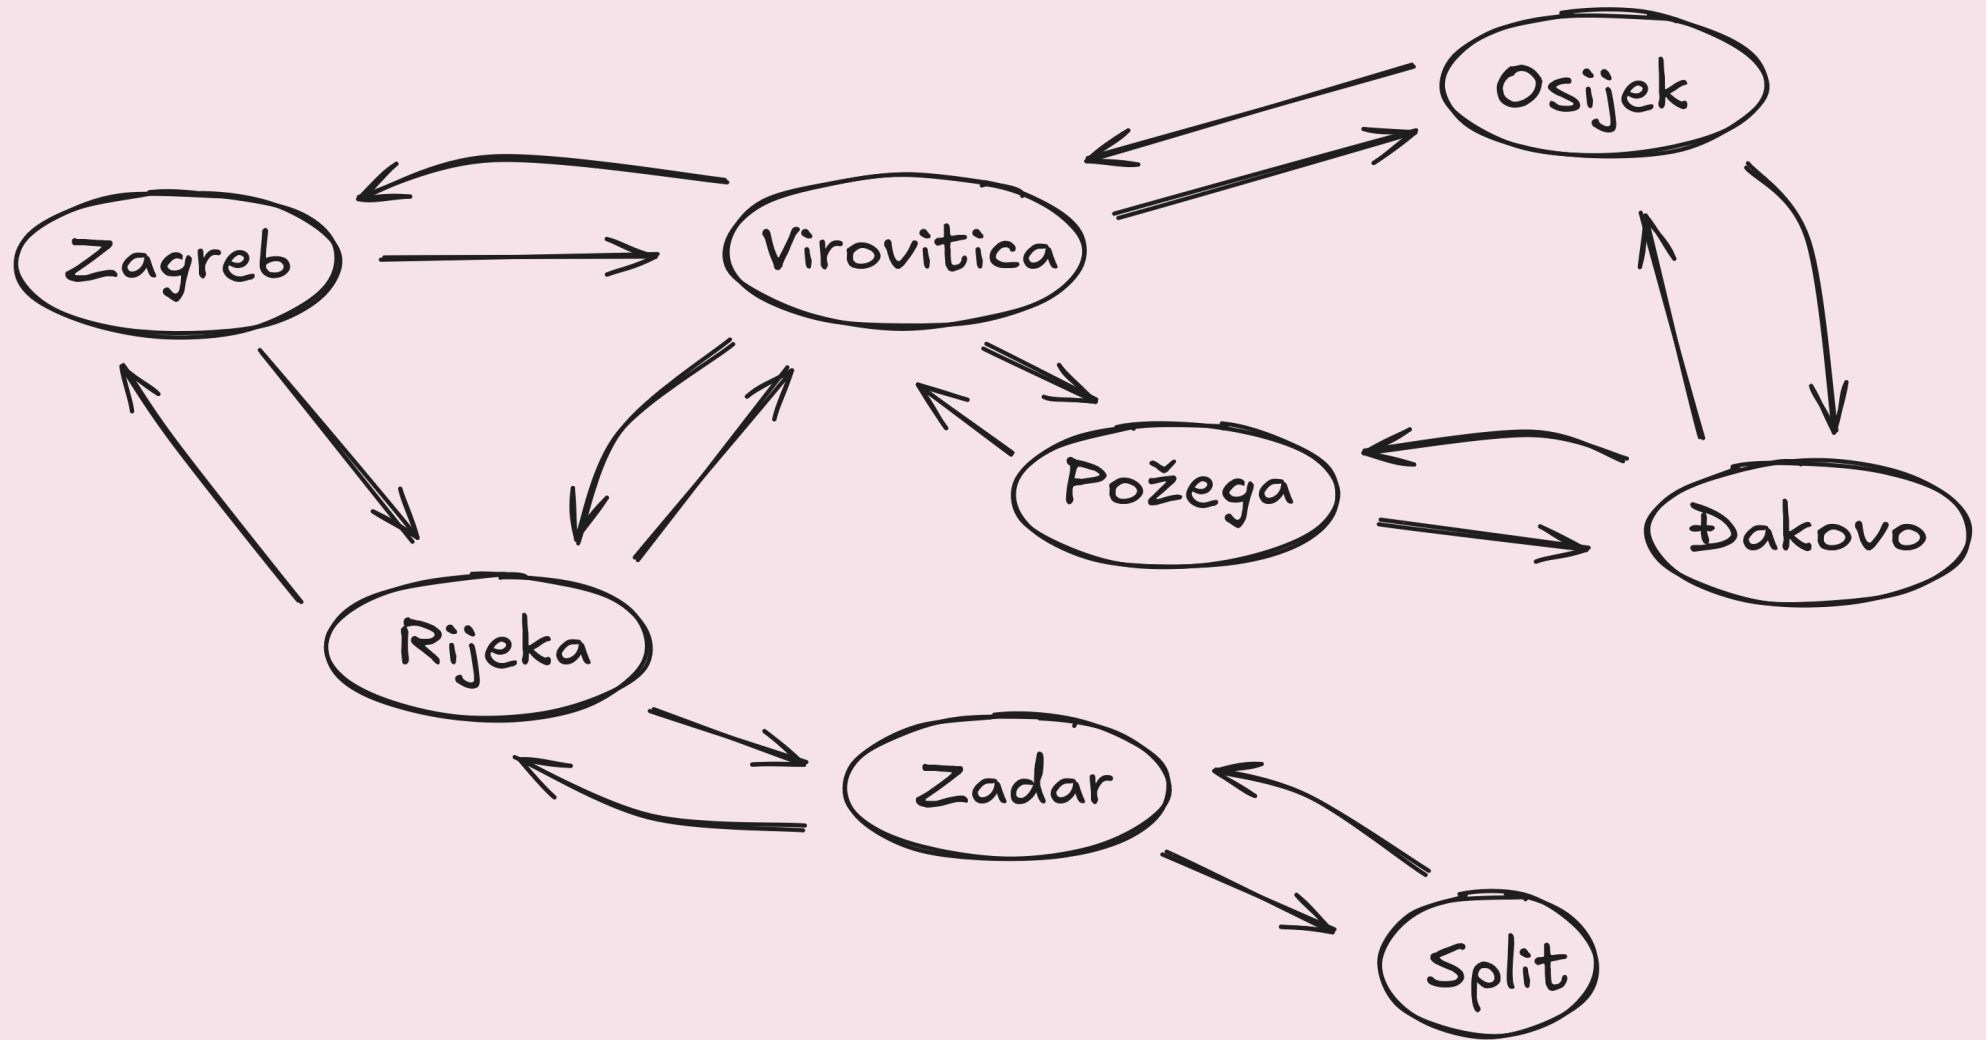


State Space Size?

- Pacman positions, Wall Positions
- Food positions, Food Status?
- Ghost positions, Ghost Status?

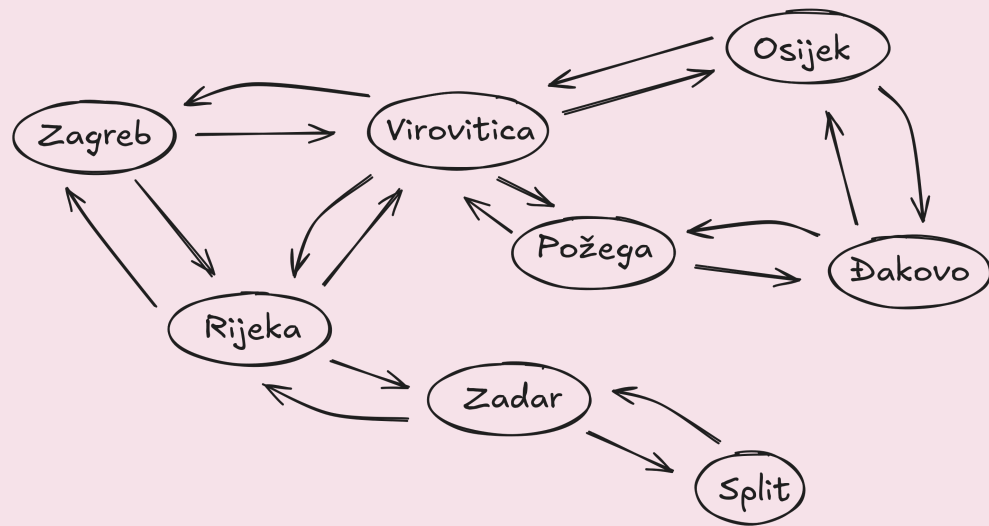


State Space Graph




Search Trees

Graph:

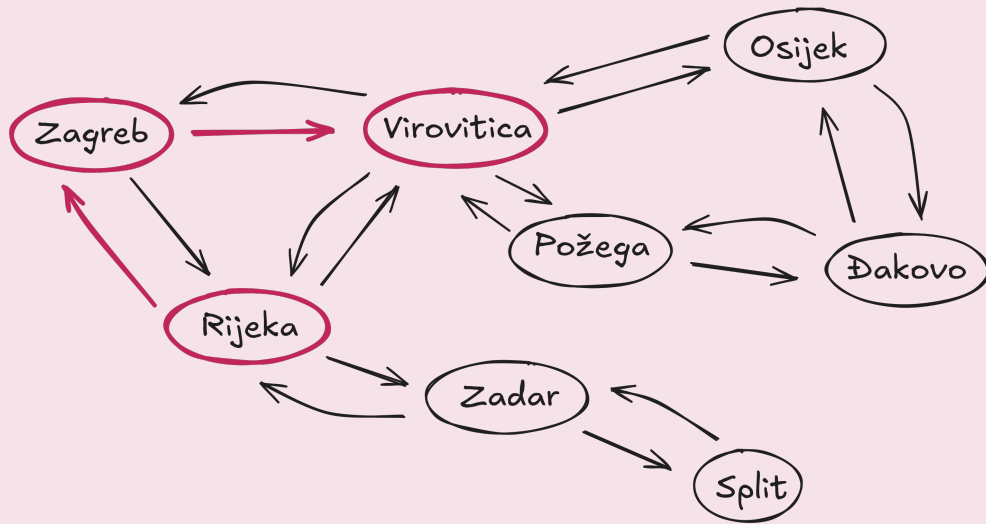


Tree:

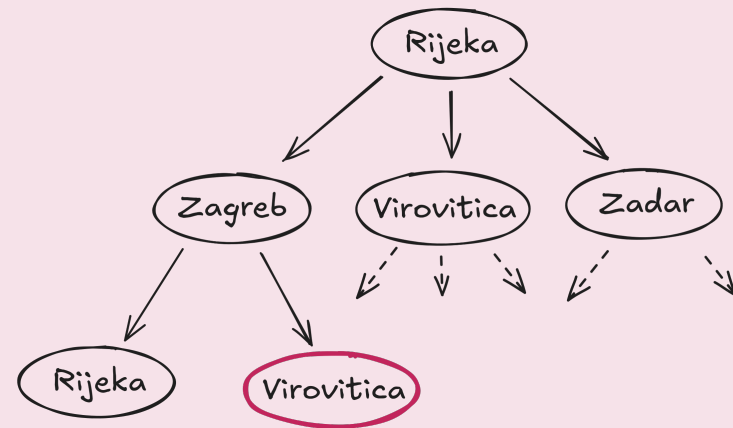


Node Representation

Graph:



Tree:



Let's Talk About Trees

- For any non-trivial problem, they're *big*
 - (Effective) branching factor
 - Depth
- Graph and tree both too large for memory
 - Successor function (graph)
 - Expansion function (tree)

How To Solve It

Given:

- Starting node
- Goal test
- Expansion

Do:

- Expand nodes from start
- Test each new node for goal
 - If goal, success
- Expand new nodes
 - If nothing left to expand, failure

Best-First Search

Algorithm Best-First Search

```
1: function BEST-FIRST-SEARCH(problem, f)
2:   node ← NODE(STATE=problem.INITIAL)
3:   frontier ← priority queue ordered by f
4:   frontier.ADD(node)
5:   reached ← lookup table
6:   reached[node] ← problem.INITIAL
7:   while not IS-EMPTY(frontier) do
8:     node ← POP(frontier)
9:     if problem.IS-GOAL(node.STATE) then
10:      return node
11:     for each child in EXPAND(problem,node) do
12:       s ← child.STATE
13:       if not s ∈ reached or child.PATH-COST < reached[s].PATH-COST then
14:         reached[s] ← child
15:         frontier.ADD(child)
16:   return failure
17:
18: function EXPAND(problem, node)
19:   s ← node.STATE
20:   for each action in problem.ACTIONS(s) do
21:     s' ← problem.RESULT(s, action)
22:     cost ← node.PATH-COST + problem.ACTION-COST(s, action, s')
23:     yield NODE(STATE= s', PARENT=node, ACTION=action, PATH-COST=cost)
```

Frontier Expansion



Frontier Expansion

- Frontier: nodes “currently” expanded
 - If no frontier node is goal, need to add to frontier
 - How?
- Can we have cycles?
 - How do we deal with cycles?

Queues & Searches

- Priority Queues
 - Best-First Search
 - Uniform-Cost Search¹
- FIFO Queues
 - Breadth-First Search
- LIFO Queues²
 - Depth-First Search

1. Also known as “Dijkstra’s Algorithm,” because it is Dijkstra’s Algorithm

2. Also known as “stacks.” because they are stacks.

Search Features

- Completeness
 - If there is a solution, will we find it?
- Optimality
 - Will we find the *best* solution?
- Time complexity
- Memory complexity

Breadth-First Search

- FIFO Queue
- Complete
- Optimal
- $O(b^d)$
- Nice features for equal-weight arcs:
 - Lowest-cost path first
 - *reached* collection can be a set

Breadth-First Search

Algorithm Breadth-First Search

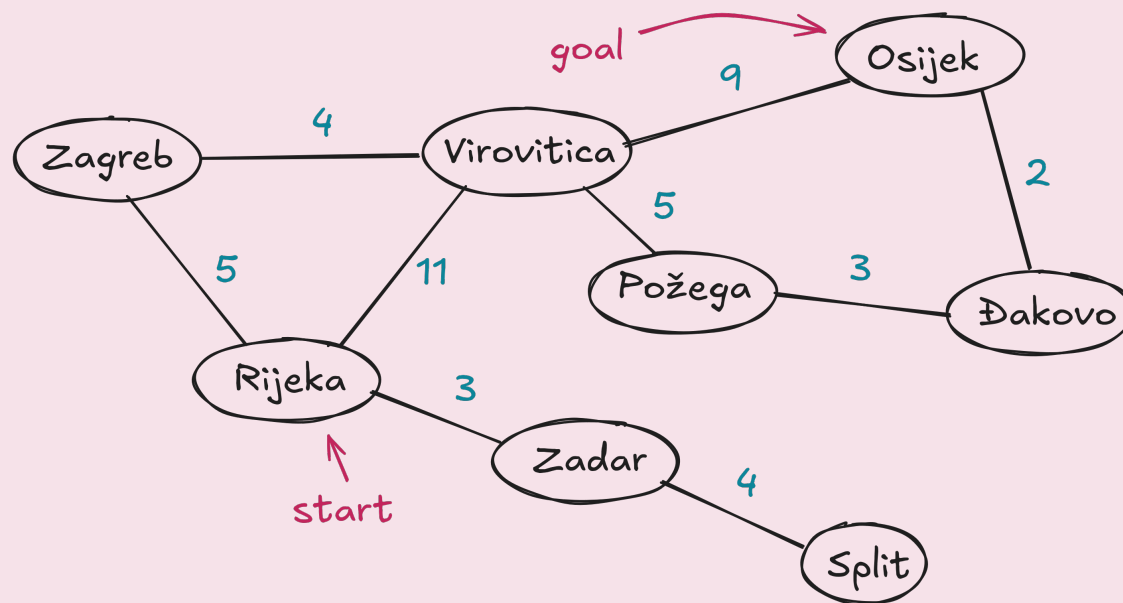
```
1: function BREADTH-FIRST-SEARCH(problem)
2:   node ← NODE(STATE=problem.INITIAL)
3:   if problem.IS-GOAL(node.STATE) then
4:     return node
5:   frontier ← FIFO queue
6:   frontier.ADD(node)
7:   reached ← set
8:   reached ← problem.INITIAL
9:   while not IS-EMPTY(frontier) do
10:    node ← POP(frontier)
11:    for each child in EXPAND(problem,node) do
12:      s ← child.STATE
13:      if problem.IS-GOAL(s) then
14:        return child
15:      if not s ∈ reached then
16:        reached.ADD(child)
17:        frontier.ADD(child)
18:   return failure
```

Uniform-Cost Search

Non-uniform costs \rightarrow BFS inappropriate.

Algorithm Uniform-Cost Search

```
1: function UNIFORM-COST-SEARCH(problem)  
2:   return BEST-FIRST-SEARCH(problem, PATH-COST)
```



Depth-First Search

- “Family” of searches
- LIFO stack
- Problems?

Algorithm Depth-First Search

```
1: function DEPTH-FIRST-SEARCH(problem)
2:   node ← NODE(STATE=problem.INITIAL)
3:   frontier ← LIFO stack
4:   frontier.PUSH(node)
5:   while not IS-EMPTY(frontier) do
6:     node ← POP(frontier)
7:     if problem.IS-GOAL(node.STATE) then
8:       return node
9:     else if not IS-CYCLE(node) then
10:      for each child in EXPAND(problem,node) do
11:        frontier.PUSH(child)
12:   return failure
```

Uninformed Search Variants

- Depth-Limited Search
 - Fail if depth limit reached (why?)
- Iterative deepening
 - vs. Breadth-First Search
- Bidirectional Search

How to Choose?

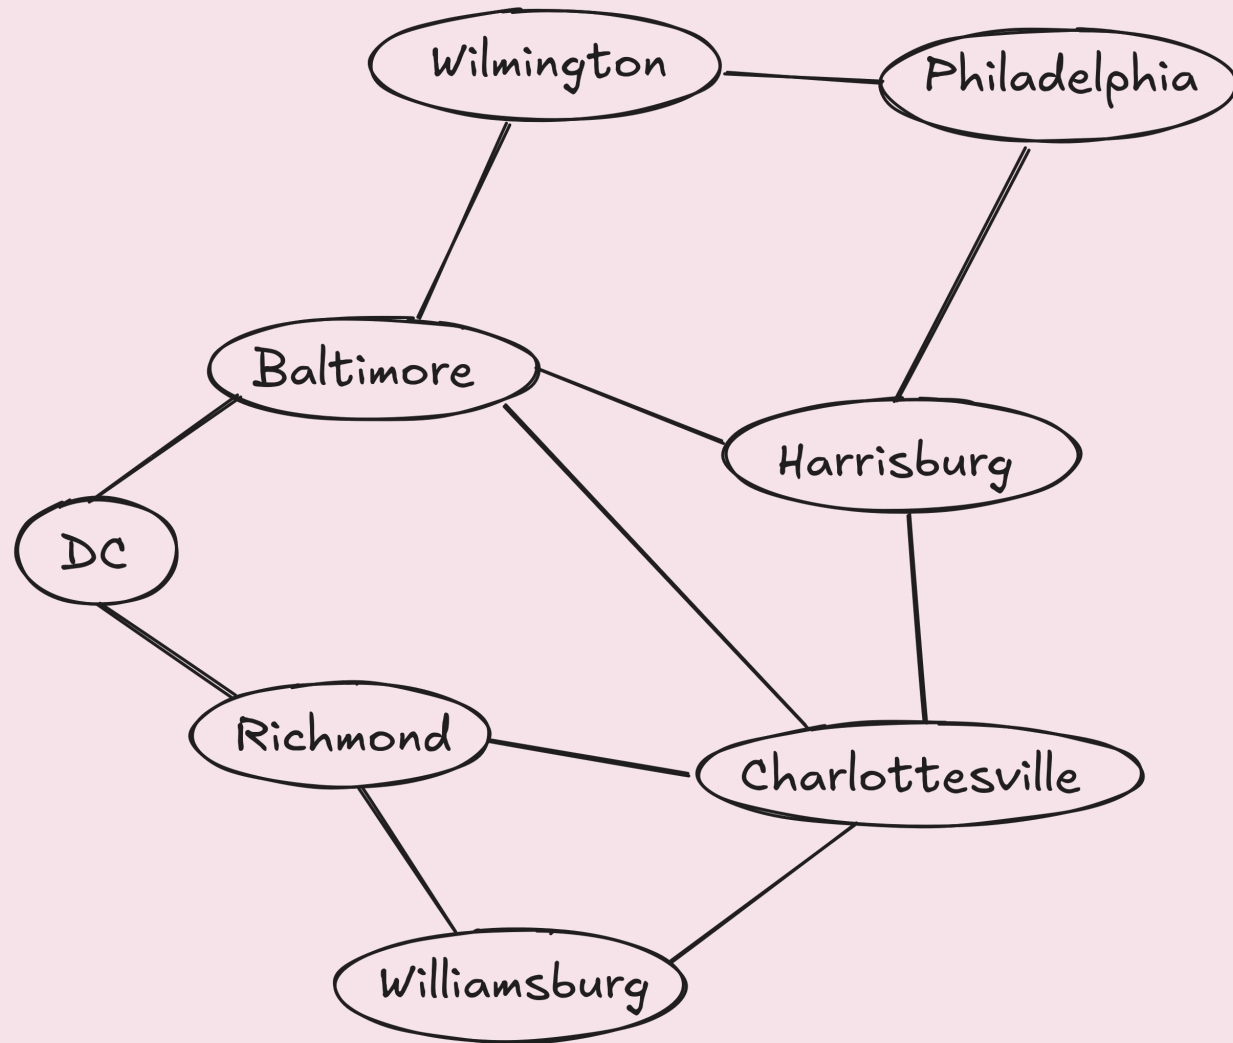
- Think about when the searches “fail”
- Think about complexity
- Do we need an optimal solution?
 - Are we looking for “any” solution

Informed Search

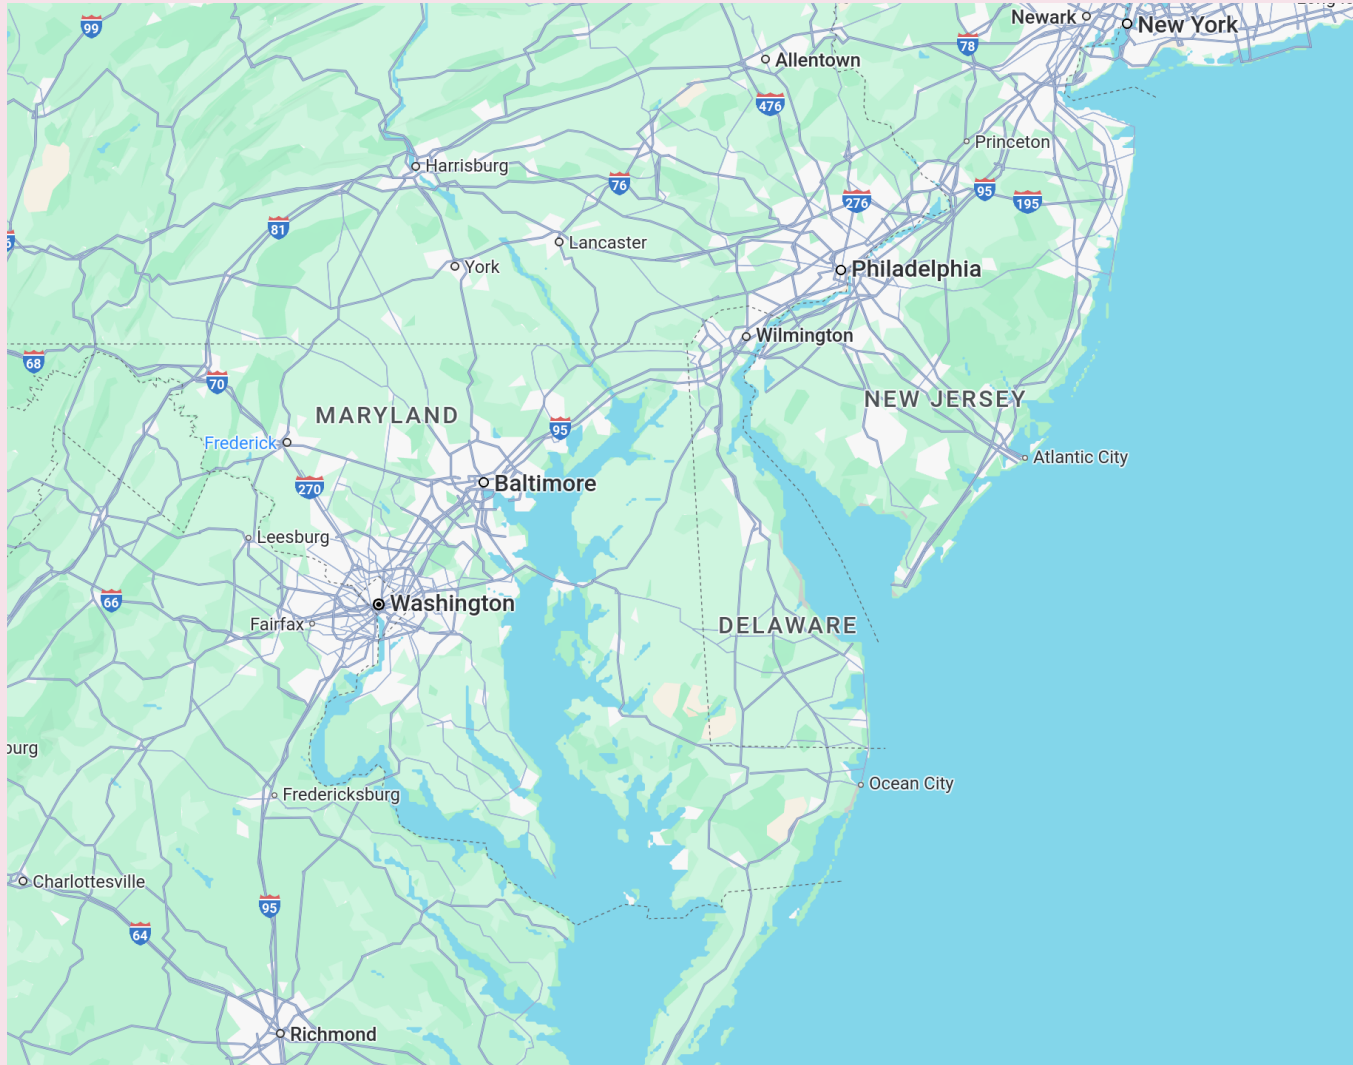
It Is Possible To Know Things




It Is Possible To Know Things



Mid-Atlantic



DC Metro Area



Heuristics

heuristic - *adj* - Serving to discover or find out.¹

- We know things about the problem
- These things are external to the graph/tree structure
 - We could model the problem differently
 - We can use the information directly

1. Webster's. 1913

Best-First Search (reprise)

Algorithm Best-First Search

```
1: function BEST-FIRST-SEARCH(problem, f)
2:   node ← NODE(STATE=problem.INITIAL)
3:   frontier ← priority queue ordered by f
4:   frontier.ADD(node)
5:   reached ← lookup table
6:   reached[node] ← problem.INITIAL
7:   while not IS-EMPTY(frontier) do
8:     node ← POP(frontier)
9:     if problem.IS-GOAL(node.STATE) then
10:      return node
11:     for each child in EXPAND(problem,node) do
12:       s ← child.STATE
13:       if not s ∈ reached or child.PATH-COST < reached[s].PATH-COST then
14:         reached[s] ← child
15:         frontier.ADD(child)
16:   return failure
17:
18: function EXPAND(problem, node)
19:   s ← node.STATE
20:   for each action in problem.ACTIONS(s) do
21:     s' ← problem.RESULT(s, action)
22:     cost ← node.PATH-COST + problem.ACTION-COST(s, action, s')
23:     yield NODE(STATE= s', PARENT=node, ACTION=action, PATH-COST=cost)
```

Greedy Best-First Search

- Heuristic $h(n)$
 - n is the search-tree node
 - $h(n)$ estimates cost from n to goal
- Best-first search: $f(n)$ orders priority queue
 - Use $f(n) = h(n)$
- Complete
- No optimality guarantee
 - (expected)

A* Search

- Include path-cost $g(n)$
 - $f(n) = g(n) + h(n)$


Algorithm A* Search

```
1: function A*-SEARCH(problem)  
2:   return BEST-FIRST-SEARCH(problem,  $g(n) + h(n)$ )
```

- Complete (always)
- Optimal (sometimes)
- Painful $O(b^m)$ time and space complexity

Choosing Heuristics

- Recall: $h(n)$ estimates cost from n to goal



- Admissibility
- Consistency

Choosing Heuristics

- Admissibility
 - *Never* overestimates cost from n to goal
 - Cost-optimal!
- Consistency
 - $h(n) \leq c(n, a, n') + h(n')$
 - n' successors of n
 - $c(n, a, n')$ cost from n to n' given action a

Consistency

- Consistent heuristics are admissible
 - Inverse not necessarily true
- Always reach each state on optimal path
- Implications for inconsistent heuristic?

Is Optimality Desirable?

Is Optimality Desirable?

- Yes

Is Optimality Desirable?

- Yes, but it isn't always *feasible*
 - A* search still exponentially complex in solution length
 - Optimality is never guaranteed “inexpensively”
- We need strategies for “good enough” solutions

Satisficing

satisfy - *verb* - To give satisfaction; to afford gratification; to leave nothing to be desired.¹

suffice - *verb* - To be enough, or sufficient; to meet the need (of anything)²

1. Webster's, 1913

2. Webster's, 1913

Weighted A* Search

- Greedy: $f(n) = h(n)$
- A*: $f(n) = h(n) + g(n)$
- Uniform-Cost Search: $f(n) = g(n)$
- ...
- Weighted A* Search: $f(n) = W \cdot h(n) + g(n)$
 - Weight $W > 1$

Reducing Complexity

- Frontier Management
- Elimination of *reached* collection
 - Reference counts
 - How else?
- Other searches

Iterative-Deepening A* Search

“IDA*” Search

- Similar to Iterative Deepening with Depth-First Search
 - DFS uses depth cutoff
 - IDA* uses $h(n) + g(n)$ cutoff *with DFS*
 - Once cutoff breached, new cutoff:
 - Typically next-largest $h(n) + g(n)$
 - $O(b^m)$ time complexity 😞
 - $O(d)$ space complexity¹ 😊

1. This is slightly complicated based on heuristic branching factor b_h .

Beam Search

Best-First Search:

- Frontier is all expanded nodes

Beam Search:

- k “best” nodes are kept on frontier
 - Others discarded
- Alt: all nodes within δ of best node
- Not Optimal
- Not Complete

Recursive Best-First Search (RBFS)

- No *reached* table is kept
- Second-best node $f(n)$ retained
 - Search from each node cannot exceed this limit
 - If exceeded, recursion “backs up” to previous node
- Memory-efficient
 - Can “cycle” between branches

Recursive Best-First Search (RBFS)

Algorithm Recursive Best-First Search

```
1: function RECURSIVE-BEST-FIRST-SEARCH(problem)
2:   solution, f_value  $\leftarrow$  RBFS(problem, NODE(problem.INITIAL),  $\infty$ )
3:   return solution
4:
5: function RBFS(problem, node, f_limit)
6:   if problem.Is-GOAL(node.STATE) then
7:     return node
8:   successors  $\leftarrow$  LIST(EXPAND(node))
9:   if IS-EMPTY(successors) then
10:    return failure,  $\infty$ 
11:   for each s in successors do
12:     s.f  $\leftarrow$  MAX(s.PATH-COST + h(s), node.f)
13:   while True do
14:     best  $\leftarrow$  node in successors with lowest f
15:     if best.f > f_limit then
16:       return failure, best.f
17:     alternative  $\leftarrow$  node in successors with second-lowest f
18:     result, best.f  $\leftarrow$  RBFS(problem, best, MIN(f_limit, alternative))
19:     if result  $\neq$  failure then
20:       return result, best.f
```

Heuristic Characteristics

- What makes a “good” heuristic?
 - We know about admissability and consistency
 - What about performance?
- Effective branching factor
- Effective depth
- # of nodes expanded

Where Do Heuristics Come From?

- Intuition
 - “Just Be Really Smart”
- Relaxation
 - The problem is constrained
 - Remove the constraint
- Pre-computation
 - Sub problems
- Learning

References

- Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 4th Edition, 2020.
- Stanford CS231
- UC Berkeley CS188