

Artificial Intelligence

Lecture 2: Uninformed Search

Prof. Antonio Khalil Moretti

Week 2

Spelman College

Problem Solving as Search

Search Algorithms Roadmap: Weeks 2-4

Week 2: Uninformed Search

- Breadth-First Search (BFS)
- Depth-First Search (DFS)
- Uniform-Cost Search (UCS)
- Iterative Deepening (IDDFS)

Week 3: Informed Search

- Greedy Best-First Search
- A* Search
- Heuristic Design
- IDA* (memory-efficient)

Week 4: Adversarial Search

- Minimax Algorithm
- Alpha-Beta Pruning
- Expectimax (stochastic games)
- Evaluation Functions

Week 5: Constraints

- CSP Formulation
- Backtracking Search
- Constraint Propagation

Today: Uninformed search - algorithms with no domain knowledge

Motivation: Why Search?

Many AI problems can be framed as search:

Classic Examples:

- Route finding (GPS)
- Puzzle solving (8-puzzle, Rubik's cube)
- Game playing (chess, checkers)
- Robot path planning

The Pattern:

- Start state
- Goal state
- Set of actions
- Find sequence of actions

Key Insight: Search is about exploring possibilities systematically

Note: This is "classical" AI - elegant but limited. Modern AI uses learning for most problems.

Problem Formulation

Components of a search problem

A search problem consists of:

1. **Initial state:** Where we start
2. **Actions:** What we can do in each state
3. **Transition model:** What each action does
4. **Goal test:** Are we done?
5. **Path cost:** Cost of a sequence of actions

Solution: A sequence of actions from initial state to goal

Optimal solution: Solution with lowest path cost

Interactive Visualization: Search Algorithms

Watch the algorithms explore

<https://qiao.github.io/PathFinding.js/visual/>

What you'll see:

- Maze pathfinding with multiple algorithms
- Visual exploration: blue = exploring, yellow = visited
- Compare BFS, A*, Dijkstra, Best-First, and more

Try this:

1. Click and drag to draw walls (create a maze)
2. Select "Breadth-First Search" from the dropdown
3. Click "Start Search" and watch it explore evenly
4. Try "A* Search" - notice how it's more focused toward the goal

Great for comparing BFS with informed search algorithms!

Example: Maze Solving - Concrete Representation

Problem Components:

Initial State:

- Position (row, col): (0, 0)

Actions:

- {UP, DOWN, LEFT, RIGHT}

Transition Model:


- UP: $(r, c) \rightarrow (r-1, c)$
- Must be valid: in bounds, not wall

Goal Test:

- Current position == goal position

Path Cost:

- 1 per move



How to represent in code:

- **Grid:** 2D array (0=open, 1=wall)
- **State:** (row, col) tuple
- **Check valid:** `grid[r][c] == 0`

Interactive Visualization: DFS vs BFS

Watch the algorithms explore

<https://visualgo.net/en/dfsdfs>

What you'll see:

- Graph nodes colored as they're explored
- Step-by-step execution of each algorithm
- Clear difference between DFS (goes deep) and BFS (explores level by level)

Try this:

1. Select "BFS" from the menu and click "Go"
2. Watch how it explores all neighbors before going deeper
3. Now select "DFS" and click "Go"
4. Notice how it follows one path all the way down before backtracking

Example: 8-Puzzle

Initial State:

7	2	4
5		6
8	3	1

Goal State:

	1	2
3	4	5
6	7	8

Formulation:

- **States:** 3×3 grid configurations
- **Actions:** Move blank up/down/left/right
- **Transition:** Swap blank with adjacent tile
- **Goal test:** Match goal configuration
- **Cost:** Number of moves (1 per move)


State space: $9!/2 = 181,440$ reachable states

How do we find the solution efficiently?

Search Trees vs Search Graphs


Search Tree:

- Nodes can repeat
- Shows exploration path
- May be infinite
- Simple to implement



Search Graph:

- Track visited states
- Avoid revisiting
- More efficient
- Requires more memory



In practice, always use graph search to avoid redundant work!

Uninformed Search Strategies

Uninformed Search

No domain knowledge, just brute force

Also called "blind search"

Characteristics:

- No knowledge about goal location
- Only know: goal vs non-goal
- Differ in order of exploration

Key data structure:

- **Frontier (fringe):** nodes to explore
- Order determines strategy

We'll cover:

1. Breadth-First Search (BFS)
2. Depth-First Search (DFS)
3. Uniform-Cost Search (UCS)

Evaluate on:

- Completeness
- Optimality
- Time complexity
- Space complexity

Search Algorithm Template

```
1: frontier  $\leftarrow$  {initial state}
2: explored  $\leftarrow$  {}
3: while frontier not empty do
4:   node  $\leftarrow$  frontier.pop()
5:   if node is goal then
6:     return solution
7:   end if
8:   explored.add(node)
9:   for child in expand(node) do
10:    if child not in explored then
11:      frontier.add(child)
12:    end if
13:  end for
14: end while
15: return failure
```

Key variations:

- **BFS:** frontier is queue (FIFO)
- **DFS:** frontier is stack (LIFO)
- **UCS:** frontier is priority queue

The only difference is how we choose which node to expand next!

Breadth-First Search

Breadth-First Search (BFS)

Explore level by level


Strategy:

- Expand shallowest unexpanded node
- Use FIFO queue
- Explores by depth level

Properties:

- ✓ Complete
- ✓ Optimal (if uniform cost)
- ✗ High memory

Exploration Order:



Level 0: 1

Level 1: 2, 3

Level 2: 4, 5, 6, 7

BFS Implementation

```
from collections import deque
def bfs(problem):
    frontier = deque([problem.initial_state])
    explored = set()
    while frontier:
        node = frontier.popleft() # FIFO - leftmost
        if problem.is_goal(node):
            return solution(node)
        explored.add(node)
        for child in problem.expand(node):
            if child not in explored and child not in frontier:
                frontier.append(child)
    return None # No solution found
```

Key: Use deque for efficient FIFO operations

BFS Analysis

Let:

- b = branching factor (children per node)
- d = depth of shallowest solution

Time Complexity:

- Level 0: 1 node
- Level 1: b nodes
- Level 2: b^2 nodes
- Level d : b^d nodes

Total: $O(b^d)$

Space Complexity:

- Must store all nodes at current level
- Frontier size at level d : b^d

Total: $O(b^d)$

Memory is the main limitation!

Maze Example: Designing Breadth-First Search

To Implement BFS, What Do We Need?

Designing the algorithm before writing code

Before we write *any code*, we must answer a few key design questions:

- What is a **state**?
- What states should we explore next?
- When do we mark a state as **visited**?
- What does it mean to **expand** a state?
- How do we check for **success**?
- How do we **reconstruct the solution**?

Breadth-First Search is defined entirely by how we answer these questions.

What Is a State?

Defining the state space

Question: What information fully describes a situation?

What Is a State?

Defining the state space

Question: What information fully describes a situation?

Maze example:

- A state is the agent's current position
- Represented as a tuple: (row, col)

State space:

- All valid (non-wall) grid cells
- Each cell is a node in a graph

Search problems are graphs — states are nodes, actions are edges.

What States Should We Explore Next?

The frontier

Question: How do we choose which state to explore next?

What States Should We Explore Next?

The frontier

Question: How do we choose which state to explore next?

BFS answer:

- Always explore the **shallowest** unexpanded states
- Use a **FIFO queue**

Frontier:

- Contains states that have been discovered
- But not yet expanded

The frontier data structure defines the search strategy.

When Do We Mark a State as Visited?

Avoiding repeated work

Question: When should a state be marked as visited?

When Do We Mark a State as Visited?

Avoiding repeated work

Question: When should a state be marked as visited?

BFS rule:

- Mark a state as visited **when it is first discovered**
- That is, when it is added to the frontier

Why?

- Prevents adding the same state multiple times
- Guarantees the first path found is the shortest

Visited \neq explored. Visited means “seen before.”

What Does It Mean to Expand a State?

Generating successors

Question: What happens when we expand a state?

What Does It Mean to Expand a State?

Generating successors

Question: What happens when we expand a state?

To **expand** a state:

1. Remove it from the frontier
2. Generate all valid successor states
3. Add unvisited successors to the frontier

In a maze:

- Up, down, left, right
- Must be in bounds and not a wall

Expansion defines how the search graph grows.

How Do We Check for Success?

Goal testing

Question: When do we check if we are done?

How Do We Check for Success?

Goal testing

Question: When do we check if we are done?

BFS approach:

- Check the goal when a state is expanded
- The first time we reach the goal:
 - We have found the shortest path

Why this works:

- BFS explores states in increasing depth order

This is why BFS is optimal for uniform-cost problems.

How Do We Reconstruct the Solution?

Remembering how we got there

Question: How do we recover the path after reaching the goal?

How Do We Reconstruct the Solution?

Remembering how we got there

Question: How do we recover the path after reaching the goal?

Idea: parent pointers

- When adding a state, record where it came from
- Store: `parent[child] = parent`

At the goal:

- Follow parent pointers backward
- Reverse the list to get start \rightarrow goal

Search finds the goal; parent pointers give the solution.

Maze for BFS Example (with Coordinates)

(S)	(0,1)	(0,2)		(0,4)	(0,5)	(0,6)	(0,7)	(0,8)	(0,9)
(1,0)		(1,2)		(1,4)					(1,9)
(2,0)		(2,2)	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)		(2,9)
(3,0)							(3,7)		(3,9)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)	(4,8)	(4,9)
(5,0)									(5,9)
(6,0)	(6,1)	(6,2)	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)	(6,8)	(6,9)
			(7,3)		(7,5)				(7,9)
(8,0)	(8,1)	(8,2)	(8,3)	(8,4)	(8,5)	(8,6)	(8,7)	(8,8)	(8,9)
(9,0)							(9,7)	(9,8)	(G)

Start = (0,0), Goal = (9,9)

Breadth-First Search: Setup

Maze coordinates: (row, col) **Neighbor order:** Up, Down, Left, Right

Start: $(0, 0)$

Goal: $(9, 9)$

Data structures:

- Queue (FIFO)
- `came_from` dictionary
- explored set

BFS Step 0

Initial state

Queue:

[_____]

Expand: (0, 0)

Which neighbors are added (in order)?

_____, _____

BFS Step 0 — Answer

Queue:

$[(0, 0)]$

Neighbors of $(0, 0)$:

$(1, 0), (0, 1)$

Queue after expansion:

$[(1, 0), (0, 1)]$

BFS Step 1

Expand: (1, 0)

Which new node is added?

Queue after expansion:

[_____]

BFS Step 1 — Answer

New node added:

$(2, 0)$

Queue:

$[(0, 1), (2, 0)]$

BFS Step 2

Expand: (0, 1)

Which neighbor is valid and new?

Queue becomes:

[_____]

BFS Step 2 — Answer

New node added:

$(0, 2)$

Queue:

$[(2, 0), (0, 2)]$

BFS Step 3

Expand: (2, 0)

Which node is added?

Queue after expansion:

[_____]

BFS Step 3 — Answer

New node added:

$(3, 0)$

Queue:

$[(0, 2), (3, 0)]$

BFS Concept Check

Q1. Why is $(0, 0)$ never added again?

Answer: It is already recorded in `came_from`.

Q2. Why is $(3, 0)$ expanded after $(0, 2)$?

Answer: FIFO queue — $(0, 2)$ was discovered earlier.

Q3. Why does BFS guarantee the shortest path?

Answer: Nodes are expanded in order of increasing path length.

Q4. What changes if we reverse neighbor order?

Answer: The path shape may change, but its length remains minimal.

Depth-First Search: Setup

Maze coordinates: (row, col) **Neighbor order:** Up, Down, Left, Right

Start: $(0, 0)$

Goal: $(9, 9)$

Data structures:

- Stack (LIFO)
- `came_from` dictionary
- explored set

Depth-First Search (DFS): First Steps

(S)	(0,1)	(0,2)		(0,4)	(0,5)	(0,6)	(0,7)	(0,8)	(0,9)
(1,0)		(1,2)		(1,4)					(1,9)
(2,0)		(2,2)	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)		(2,9)
(3,0)							(3,7)		(3,9)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)	(4,8)	(4,9)
(5,0)									(5,9)
(6,0)	(6,1)	(6,2)	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)	(6,8)	(6,9)
			(7,3)		(7,5)				(7,9)
(8,0)	(8,1)	(8,2)	(8,3)	(8,4)	(8,5)	(8,6)	(8,7)	(8,8)	(8,9)
(9,0)							(9,7)	(9,8)	(G)

Start (0, 0) Goal (9, 9)

Algorithm: Depth-First Search **Structure:** Stack (LIFO) **Neighbor order:** Up, Down, Left, Right

Step 0

Stack:

[(0, 0)]

Expand (0, 0)

Valid neighbors:

(1, 0), (0, 1)

Push in order:

[(1, 0), (0, 1)]

Step 1

Pop top:

(0, 1)

New neighbor:

(0, 2)

Stack:

[(1, 0), (0, 2)]

Key idea: DFS follows one path as deep as possible before backtracking.

DFS Step 0

Initial state

Stack:

[_____]

Expand: (0, 0)

Which neighbors are pushed (in order)?

_____, _____

DFS Step 0 — Answer

Stack:

[(0, 0)]

Neighbors of (0, 0) (Up blocked):

(1, 0), (0, 1)

Stack after push:

[(1, 0), (0, 1)]

Next expanded: (0, 1) (top of stack)

DFS Step 1

Expand: (0, 1)

Which neighbor is valid and new?

Stack after expansion:

[_____]

DFS Step 1 — Answer

New node pushed:

$(0, 2)$

Stack:

$[(1, 0), (0, 2)]$

Next expanded: $(0, 2)$

DFS Step 2

Expand: (0, 2)

Which neighbor is pushed next?

Stack becomes:

[_____]

DFS Step 2 — Answer

New node pushed:

$(1, 2)$

Stack:

$[(1, 0), (1, 2)]$

DFS Step 3

Expand: (1, 2)

Which direction does DFS go next?

DFS Step 3 — Answer

DFS continues **downward**:

(2, 2)

Key observation:

DFS commits to a direction until it hits a dead end.

DFS Concept Check

Q1. Why does DFS go deep instead of wide?

Answer: Because it uses a LIFO stack.

Q2. Can DFS find the goal faster than BFS?

Answer: Sometimes — if it goes in the right direction early.

Q3. Does DFS guarantee the shortest path?

Answer: No. It may find a much longer path first.

Q4. What happens if DFS reaches a dead end?

Answer: It backtracks by popping from the stack.

BFS vs DFS: Key Difference

	BFS	DFS
Data structure	Queue (FIFO)	Stack (LIFO)
Exploration	Level by level	Deep first
Shortest path	Yes	No
Memory use	Higher	Lower

Implementation: What the BFS Code Is Doing

BFS: From Design to Code

Mapping ideas to implementation

Every part of BFS design corresponds to a concrete data structure:

Concept	Implementation
State	(row, col) tuple
Frontier	deque (FIFO queue)
Visited	Dictionary / set
Expansion	Neighbor generation
Goal test	Equality check
Path	Parent pointers

Frontier and Visited in Code

Initialization

```
frontier = deque([start])  
came_from = {start: None}
```

- `frontier`: states waiting to be expanded
- `came_from`: marks states as visited
- Also stores parent pointers

A state is visited as soon as it enters the frontier.

Expansion Loop in Code

The core of BFS

```
current = frontier.popleft()
for neighbor in get_neighbors(grid, current):
    if neighbor not in came_from:
        came_from[neighbor] = current
        frontier.append(neighbor)
```

- FIFO queue enforces breadth-first order
- Neighbors are added exactly once
- Parent pointers recorded immediately

This loop is BFS. Everything else is bookkeeping.

Goal Test and Path Reconstruction

Finishing the search

```
if current == goal:
    path = []
    while current is not None:
        path.append(current)
        current = came_from[current]
    path.reverse()
```

- Goal test during expansion
- Parent pointers recover the solution
- Path length equals depth of goal

Search explores the graph; reconstruction tells the story.

Key Takeaway: BFS as Design Choices

Breadth-First Search

BFS is not a single algorithm — it is a set of design decisions:

- States define the graph
- Frontier ordering defines behavior
- Visited policy ensures correctness
- Parent pointers recover solutions

Change one design choice, and you get DFS, UCS, or A*.

Depth-First Search

Depth-First Search (DFS)

Explore as deep as possible


Strategy:

- Expand deepest unexpanded node
- Use LIFO stack
- Go deep, then backtrack

Properties:

- ✗ Not complete (can get stuck)
- ✗ Not optimal
- ✓ Low memory

Exploration Order:



Order: 1 → 2 → 3, 4 → 5, 6, 7

DFS Implementation

```
def dfs(problem):
    frontier = [problem.initial_state] # Stack (list)
    explored = set()
    while frontier:
        node = frontier.pop() # LIFO - rightmost
        if problem.is_goal(node):
            return solution(node)
        explored.add(node)
        for child in problem.expand(node):
            if child not in explored and child not in frontier:
                frontier.append(child)
    return None
```

Key: Use list with pop() for LIFO (stack) operations

DFS Analysis

Let:

- b = branching factor
- m = maximum depth of tree

Time Complexity:

- Worst case: explore entire tree
- Could explore all nodes up to depth m

Total: $O(b^m)$

Problem: If $m \gg d$, very wasteful!

Example: $b = 10$, $m = 100 \Rightarrow$ only 1,000 nodes in memory

Space Complexity:

- Only store path from root to current node
- Plus siblings of nodes on path

Total: $O(bm)$

Advantage: Much less memory than BFS!

DFS Variants

Depth-Limited Search (DLS):

- DFS with depth limit ℓ
- Prevents infinite loops
- Not complete if $\ell < d$

Iterative Deepening DFS (IDDFS):

- Run DLS with increasing limits: $\ell = 0, 1, 2, \dots$
- Combines benefits of BFS and DFS
- ✓ Complete like BFS
- ✓ Low memory like DFS
- ✓ Optimal (if uniform cost)

IDDFS is often the best uninformed search strategy!

Uniform-Cost Search

Uniform-Cost Search (UCS)


Expand lowest-cost node

Strategy:

- Expand node with lowest path cost
- Use priority queue
- Priority = $g(n)$ (cost so far)

Properties:

- ✓ Complete
- ✓ Optimal
- ✗ Can be slow



Expansion order: $A(0) \rightarrow C(2) \rightarrow E(4) \rightarrow B(5)$

UCS Implementation

```
import heapq
def ucs(problem):
    frontier = [(0, problem.initial_state)] # (cost, node)
    explored = set()
    while frontier:
        cost, node = heapq.heappop(frontier) # Get min cost
        if problem.is_goal(node):
            return solution(node)
        explored.add(node)
        for child, action_cost in problem.expand(node):
            new_cost = cost + action_cost
            if child not in explored:
                heapq.heappush(frontier, (new_cost, child))
    return None
```

Key: Use heapq for efficient priority queue

UCS Analysis

Let:

- C^* = cost of optimal solution
- ϵ = minimum step cost

Time Complexity:

- Expands all nodes with cost $< C^*$
- Number of such nodes: $O(b^{C^*/\epsilon})$

Can be much worse than BFS!

Space Complexity:

- Must store all generated nodes
- Same as time: $O(b^{C^*/\epsilon})$

Also high memory usage

When to use: When actions have different costs and you need optimal solution

Uniform Cost Search: Setup

Maze coordinates: (row, col) **Neighbor order:** Up, Down, Left, Right **Step cost:** 1 per move

Start: $(0, 0)$

Goal: $(9, 9)$

Data structures:

- Priority queue ordered by $g(n)$
- `g_cost` dictionary
- `came_from` dictionary
- explored set

Uniform Cost Search (UCS): First Steps

(S)	(0,1)	(0,2)		(0,4)	(0,5)	(0,6)	(0,7)	(0,8)	(0,9)
(1,0)		(1,2)		(1,4)					(1,9)
(2,0)		(2,2)	(2,3)	(2,4)	(2,5)	(2,6)	(2,7)		(2,9)
(3,0)							(3,7)		(3,9)
(4,0)	(4,1)	(4,2)	(4,3)	(4,4)	(4,5)	(4,6)	(4,7)	(4,8)	(4,9)
(5,0)									(5,9)
(6,0)	(6,1)	(6,2)	(6,3)	(6,4)	(6,5)	(6,6)	(6,7)	(6,8)	(6,9)
			(7,3)		(7,5)				(7,9)
(8,0)	(8,1)	(8,2)	(8,3)	(8,4)	(8,5)	(8,6)	(8,7)	(8,8)	(8,9)
(9,0)							(9,7)	(9,8)	(G)

Start (0, 0) Goal (9, 9)

Data structure: Priority Queue (min g) **Edge cost:** 1 per move

Step 0

Priority Queue:

$[(0, (0, 0))]$

Expand (0, 0)

Valid neighbors:

(1, 0), (0, 1)

Add with cost:

(1, (1, 0)), (1, (0, 1))

Queue:

$[(1, (1, 0)), (1, (0, 1))]$

Step 1

Remove lowest-cost node:

(1, (1, 0))

New neighbor:

(2, 0) with $g = 2$

Queue:

$[(1, (0, 1)), (2, (2, 0))]$

Key idea: UCS always expands the lowest-cost path so far.

Uniform Cost Search: Step 0

Initial state

Priority queue:

[(____, _____)]

Expand: (0, 0)

Which neighbors are inserted (with cost)?

_____, _____

Uniform Cost Search: Step 0 — Answer

Priority queue:

$[(0, (0, 0))]$

Neighbors of $(0, 0)$:

$(1, 0), (0, 1)$

Each has cost:

$$g = 1$$

Priority queue:

$[(1, (1, 0)), (1, (0, 1))]$

Uniform Cost Search: Step 1

Which node is removed next?

Why?

Uniform Cost Search: Step 1 — Answer

Node removed:

(1, 0)

Reason:

It has the smallest $g(n)$

New neighbor added:

(2, 0) with $g = 2$

Priority queue:

[(1, (0, 1)), (2, (2, 0))]

Uniform Cost Search: Step 2

Next node removed:

New node added:

Uniform Cost Search: Step 2 — Answer

Node removed:

$(0, 1)$

New node added:

$(0, 2)$ with $g = 2$

Priority queue:

$[(2, (2, 0)), (2, (0, 2))]$

Important Observation

- UCS always expands the node with **lowest path cost**
- When all step costs are equal:
 - UCS behaves exactly like BFS
- When step costs differ:
 - UCS may expand deeper nodes first

UCS guarantees:

First time goal is removed \Rightarrow optimal path

Uniform Cost Search: Concept Check

Q1. Why doesn't UCS use a FIFO queue?

Answer: Because expansion order depends on cost, not discovery time.

Q2. Why is UCS optimal?

Answer: It never expands a higher-cost path before a lower-cost one.

Q3. When does UCS reduce to BFS?

Answer: When all step costs are equal.

Q4. What does A* add on top of UCS?

Answer: A heuristic estimate of remaining cost.

Comparison and Summary

Algorithm Comparison

Algorithm	Complete	Optimal	Time	Space
BFS	✓	✓*	$O(b^d)$	$O(b^d)$
DFS	✗	✗	$O(b^m)$	$O(bm)$
IDDFS	✓	✓*	$O(b^d)$	$O(bd)$
UCS	✓	✓	$O(b^{C^*/\epsilon})$	$O(b^{C^*/\epsilon})$

*Optimal only if all step costs are equal

Use BFS when:

- Uniform costs
- Shallow solution
- Have memory

Use DFS/IDDFS when:

- Limited memory
- Deep trees
- Many solutions

The Limits of Uninformed Search

Why these algorithms are limited:

- Exponential complexity: $O(b^d)$ grows **very** fast
- No domain knowledge used
- Wastes time on unpromising paths

Example: 8-puzzle

- State space: 181,440 states
- Average solution depth: 22 moves
- BFS would explore thousands of states

The solution: **Informed search** (next lecture!)

- Use problem-specific knowledge
- Heuristics guide the search
- Much more efficient in practice

Key Takeaways

Search as Problem Solving

Many AI problems can be solved by searching through a state space

BFS vs DFS

BFS: Complete and optimal but memory-hungry

DFS: Memory-efficient but incomplete

UCS for Weighted Paths

Finds optimal solution when actions have different costs

Exponential Complexity

All uninformed searches have exponential time complexity - we need heuristics!

Next lecture: Informed search with heuristics (A* algorithm)