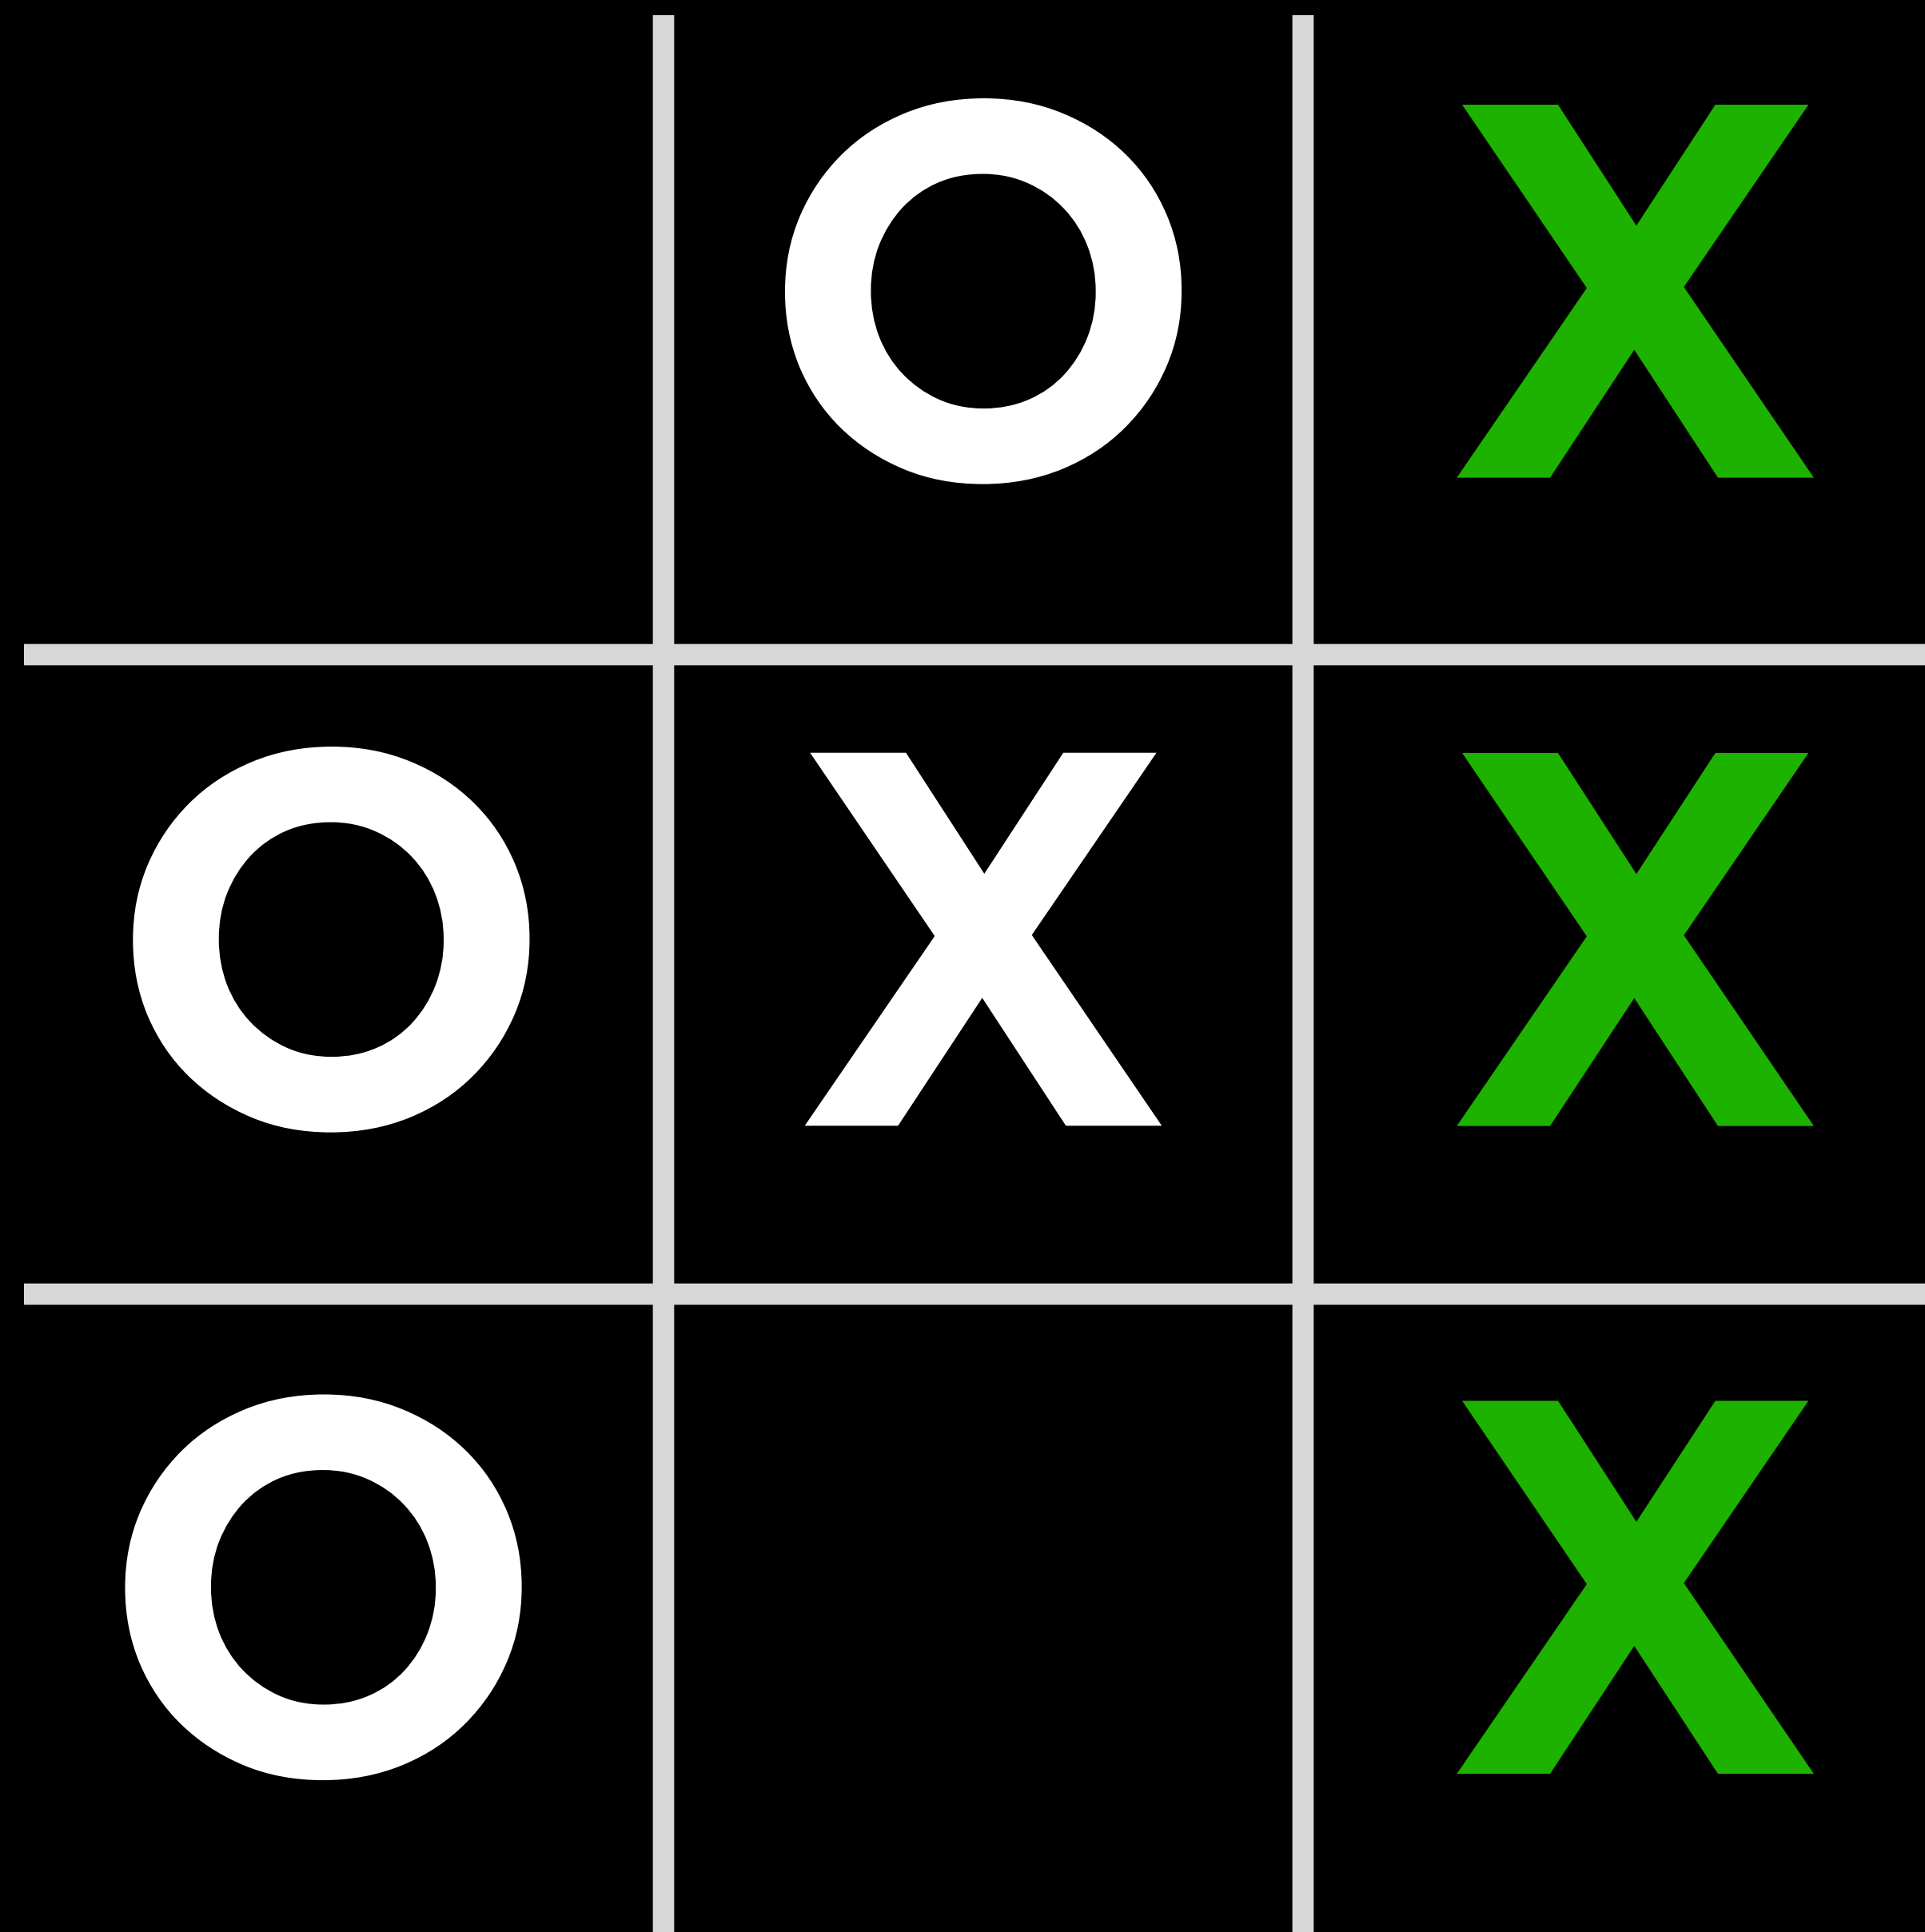
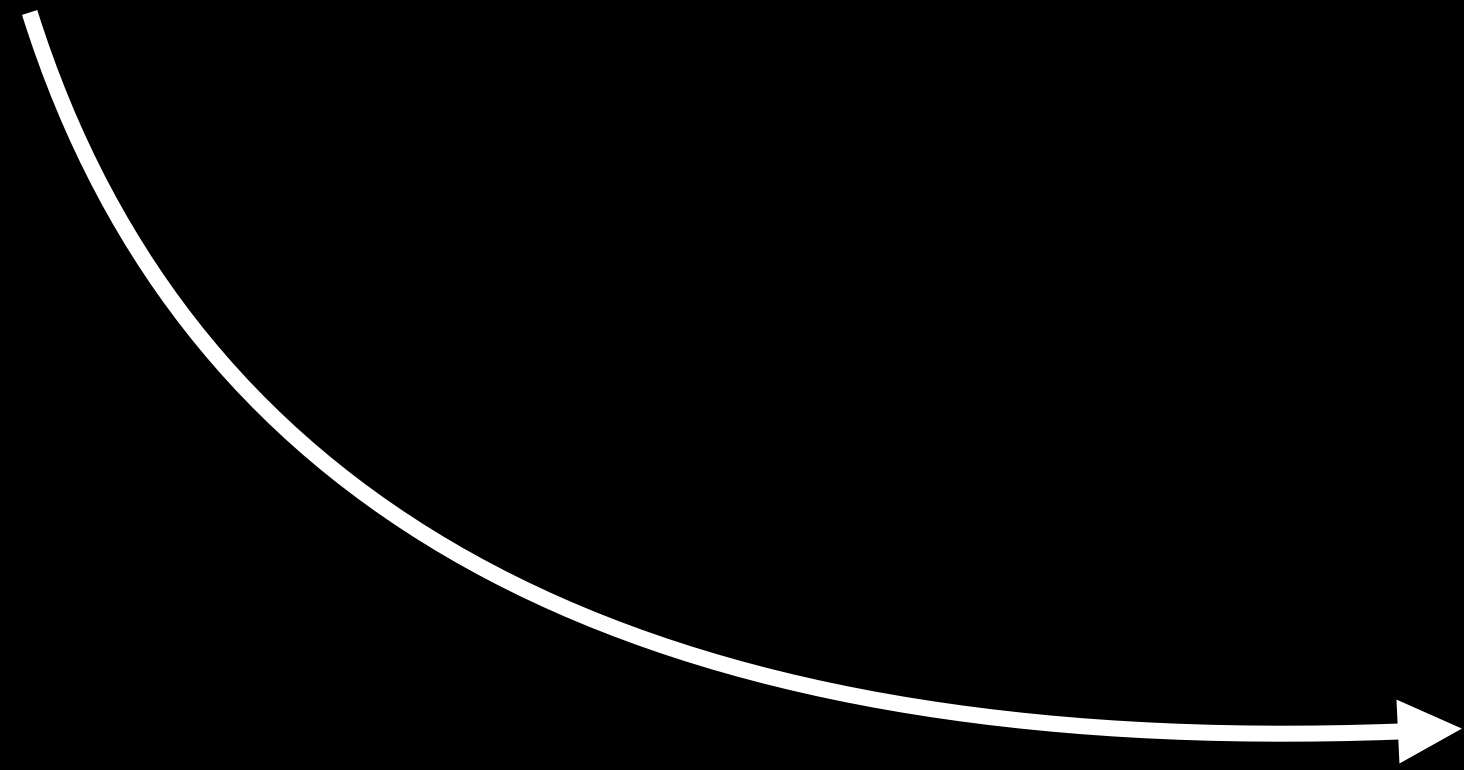


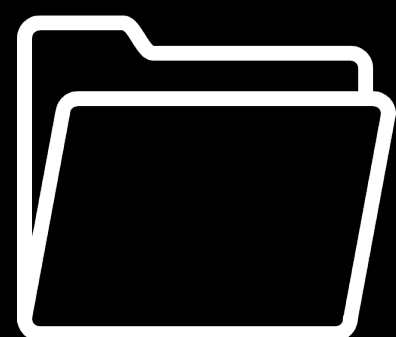
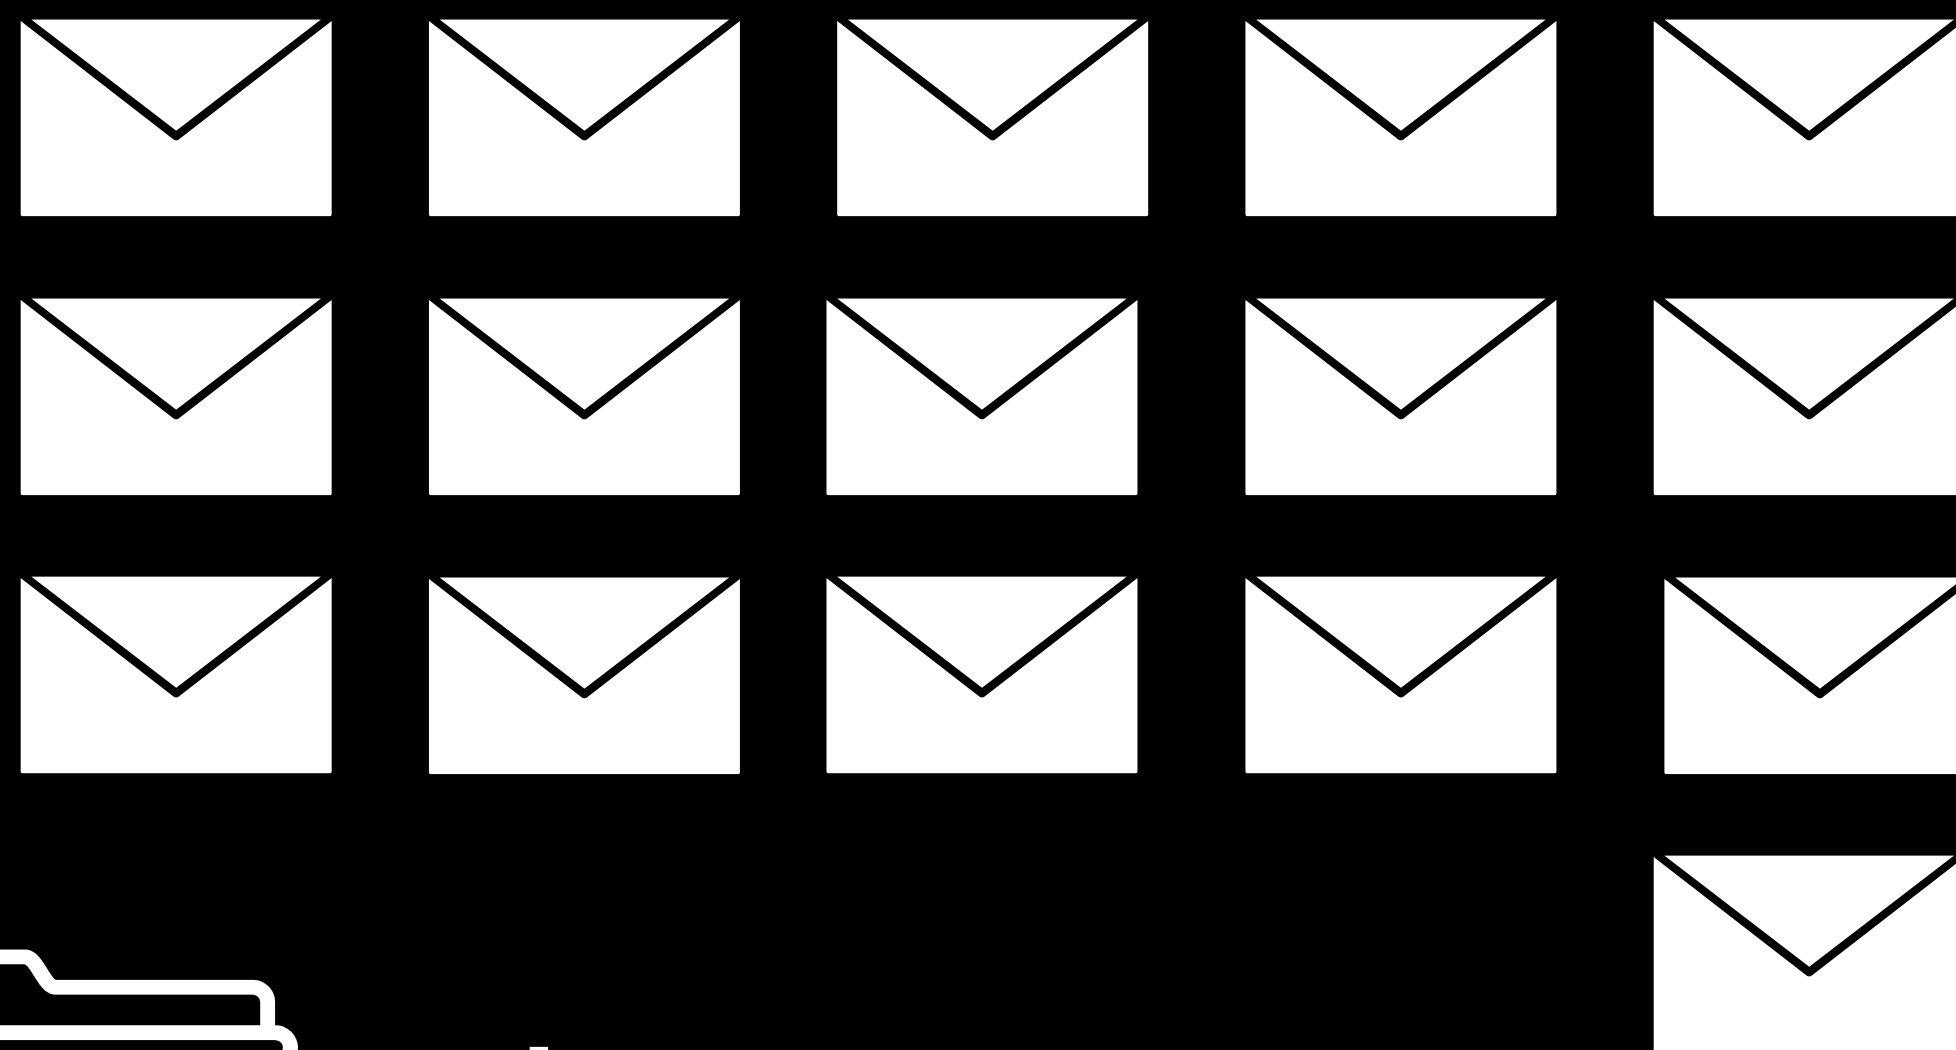
Artificial Intelligence



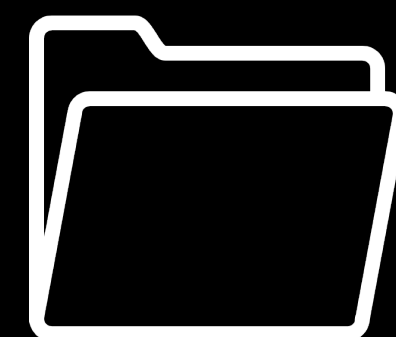
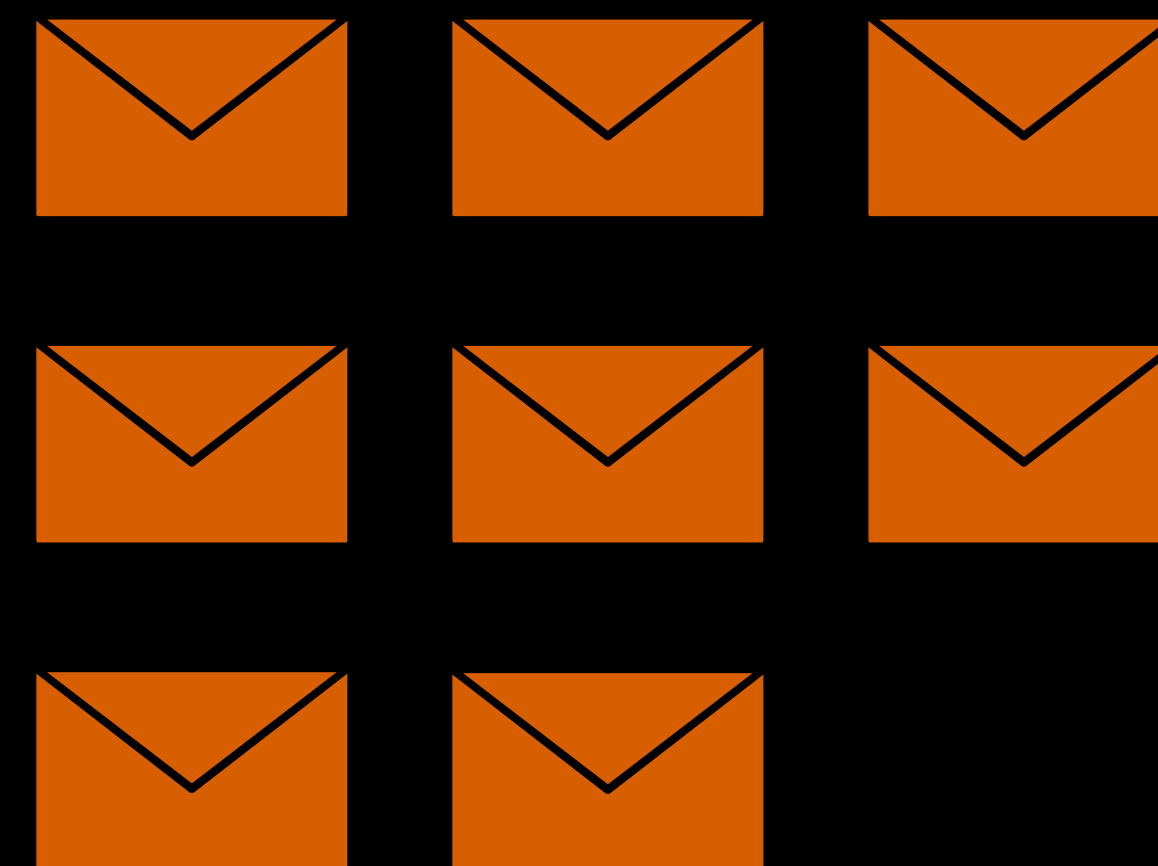
handwriting



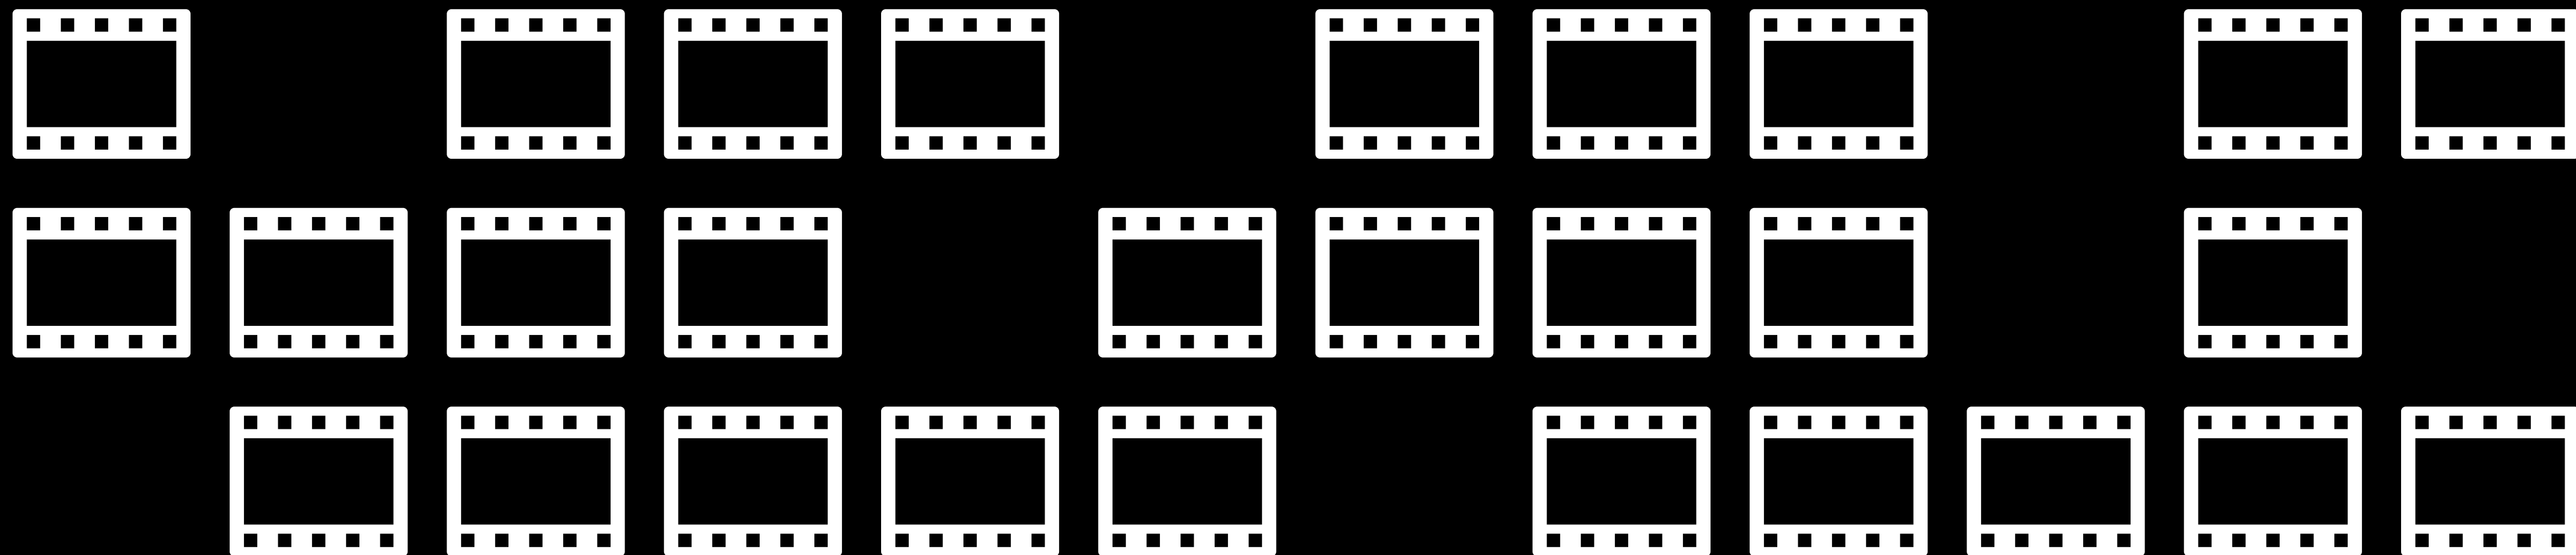
handwriting



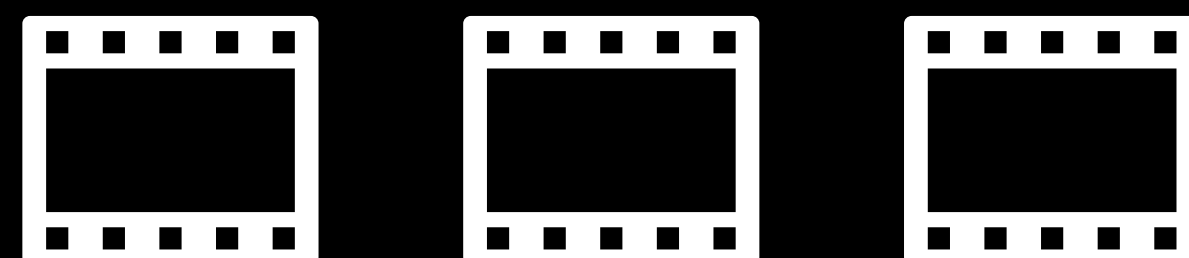
Inbox



Spam



Watch History



Recommended

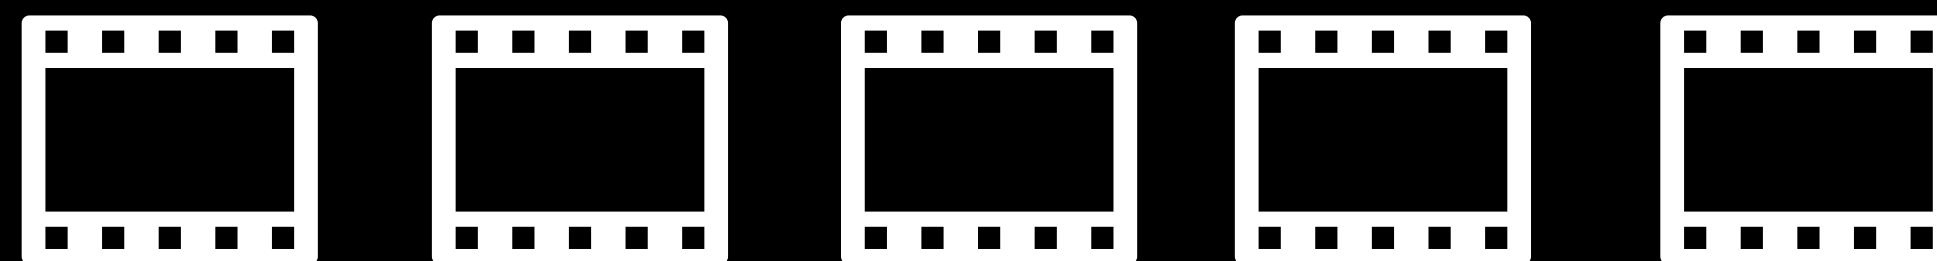


Image 1

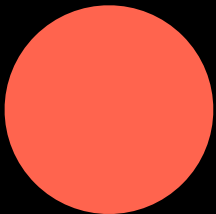
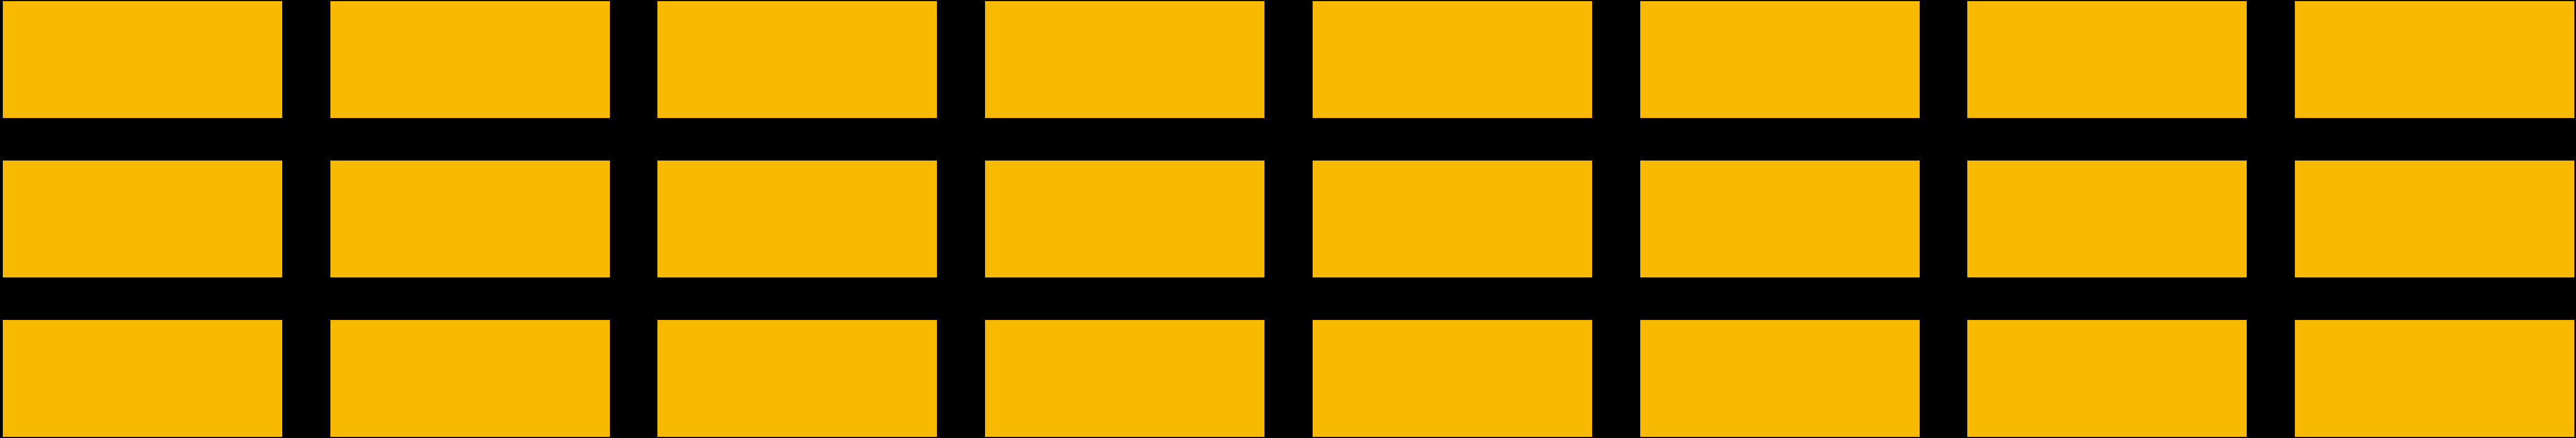


Image 2



Artificial Intelligence

Decision-Making



Decision Trees

Is ball left of paddle?

Yes

No

Move paddle left.

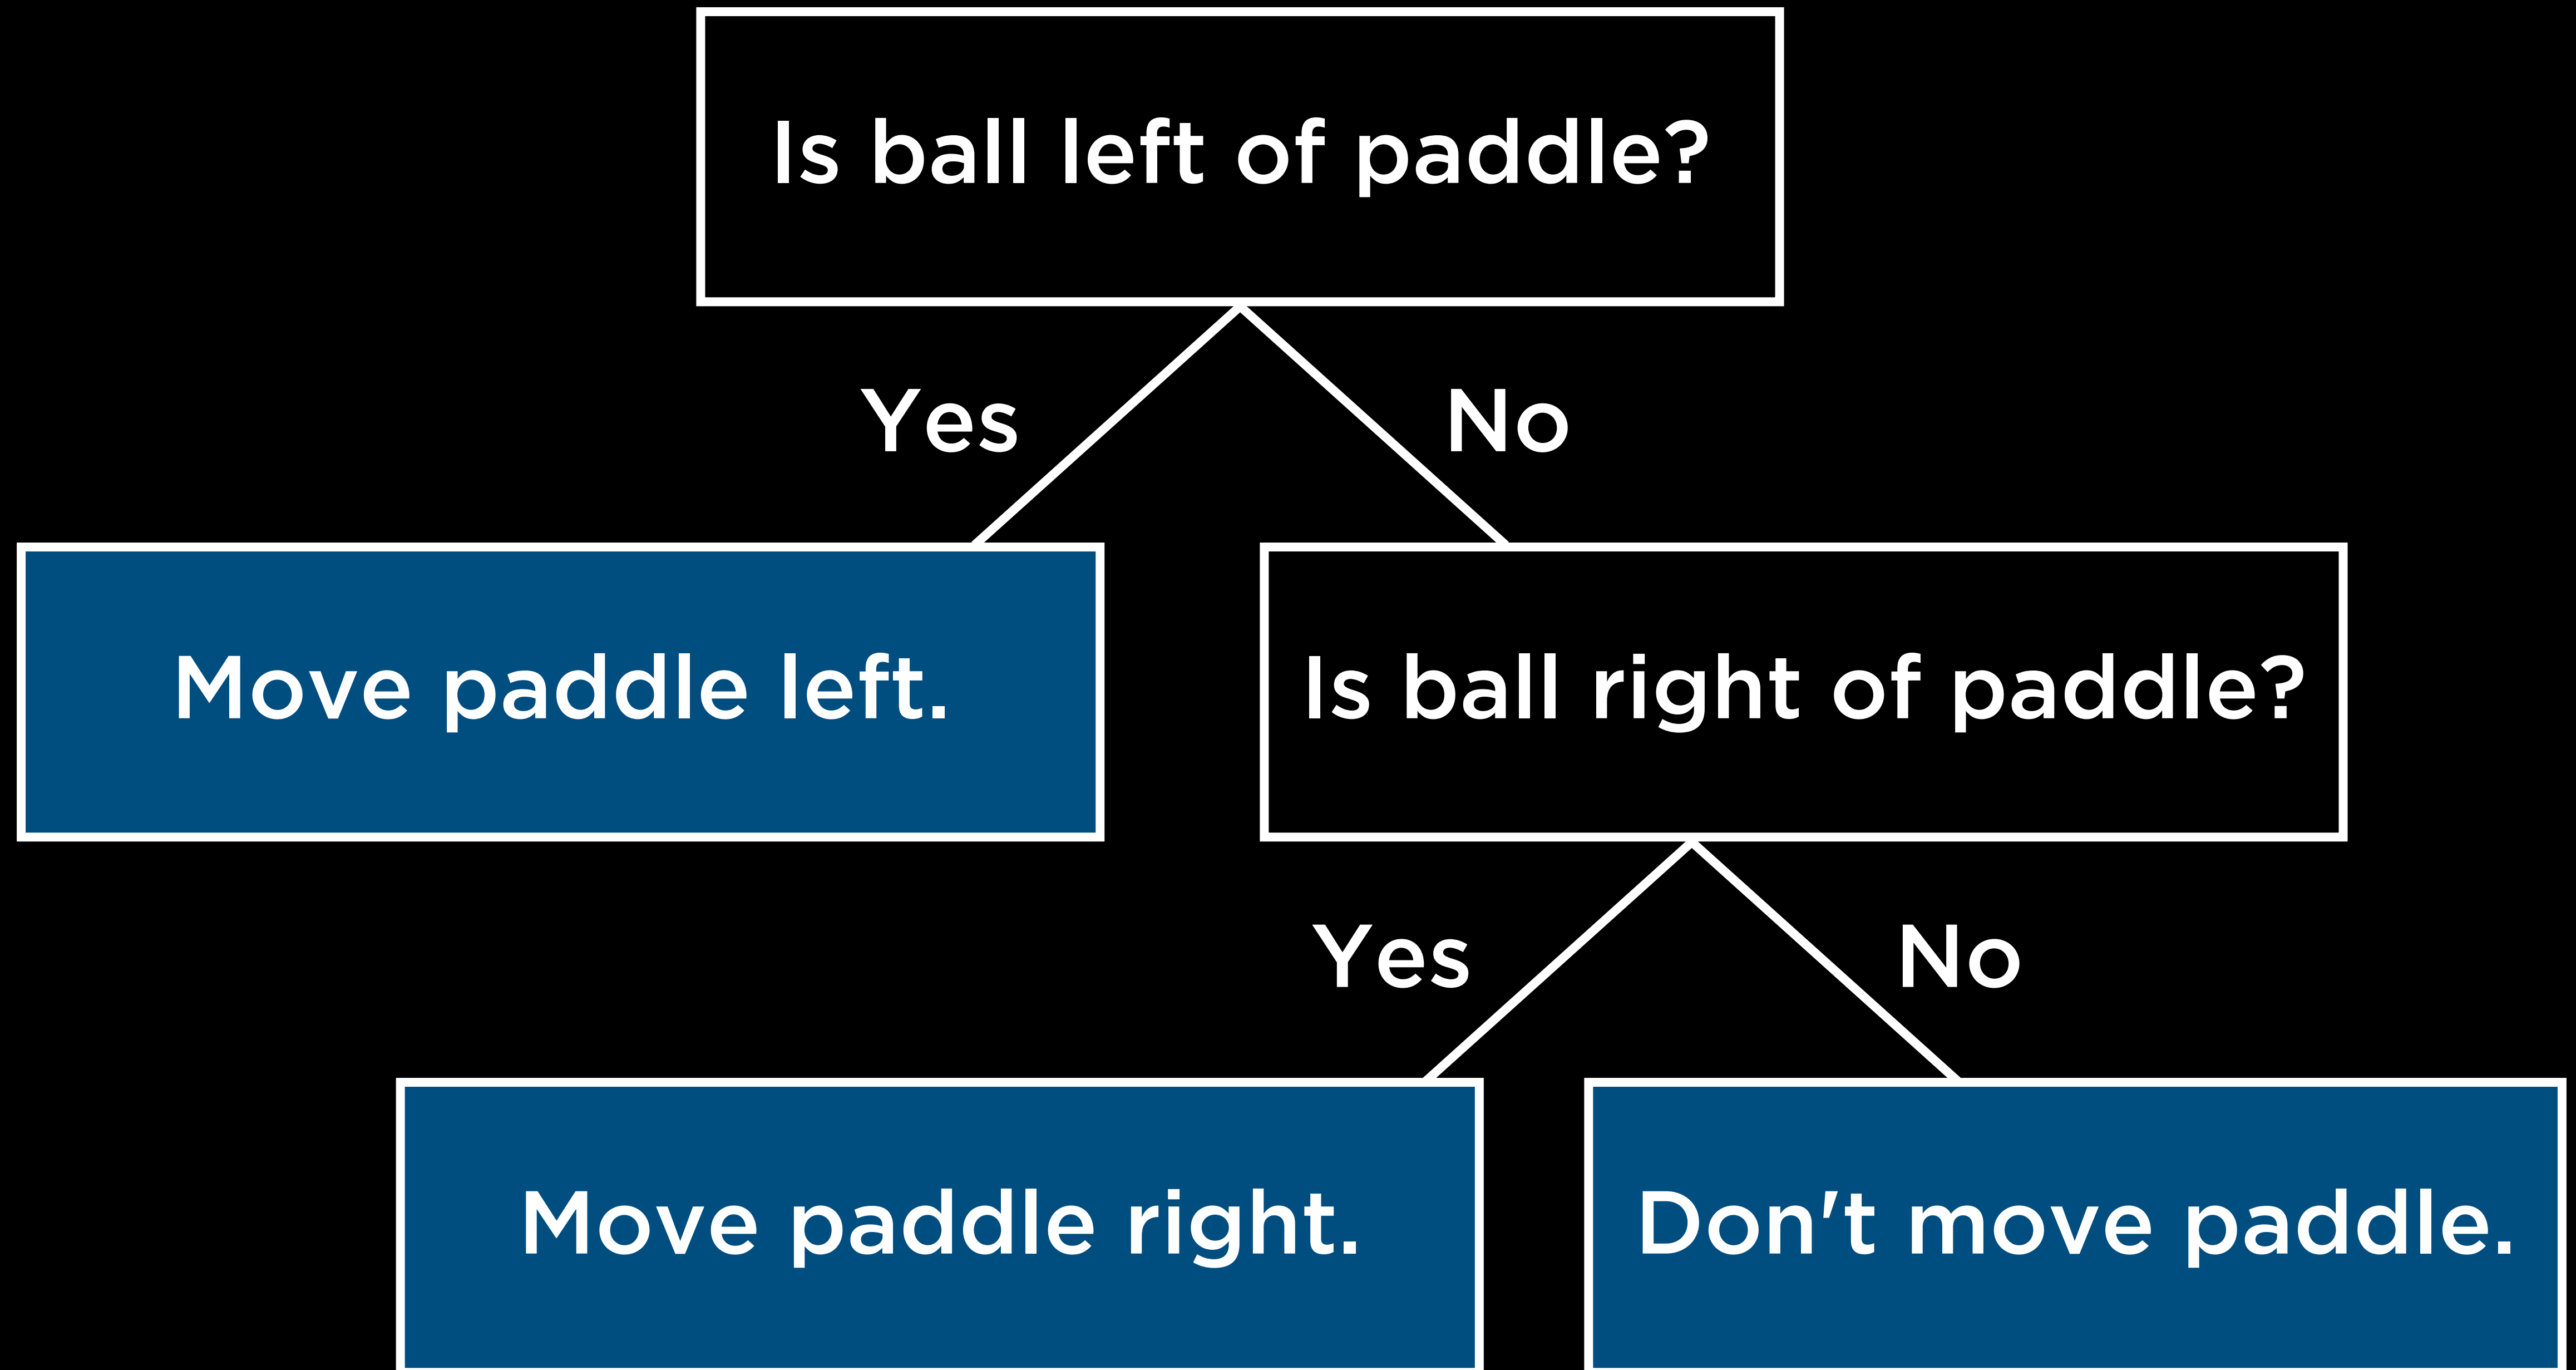
Is ball right of paddle?

Yes

No

Move paddle right.

Don't move paddle.



```
while game is ongoing:  
    if ball left of paddle:  
        move paddle left  
    else if ball right of paddle:  
        move paddle right  
    else:  
        don't move paddle
```


		O
	X	
X		O

Can I get 3 in a row on this turn?

Yes

No

Play in square to get 3 in a row.

Can my opponent get 3 in a row on next turn?

Yes

No

Play in square to block opponent's 3 in a row.

?

Minimax

- MAX (X) aims to maximize score.
- MIN (O) aims to minimize score.

O	X	X
O	O	
O	X	X

-1

X	O	X
O	O	X
X	X	O

0

O		X
	X	O
X	O	X

1

O	X	O
O	X	X
X	X	O

VALUE: 1

TURN = 0

VALUE:
0

	X	O
O	X	X
X		O

VALUE:
1

O	X	O
O	X	X
X		O

VALUE:
0

	X	O
O	X	X
X	O	O

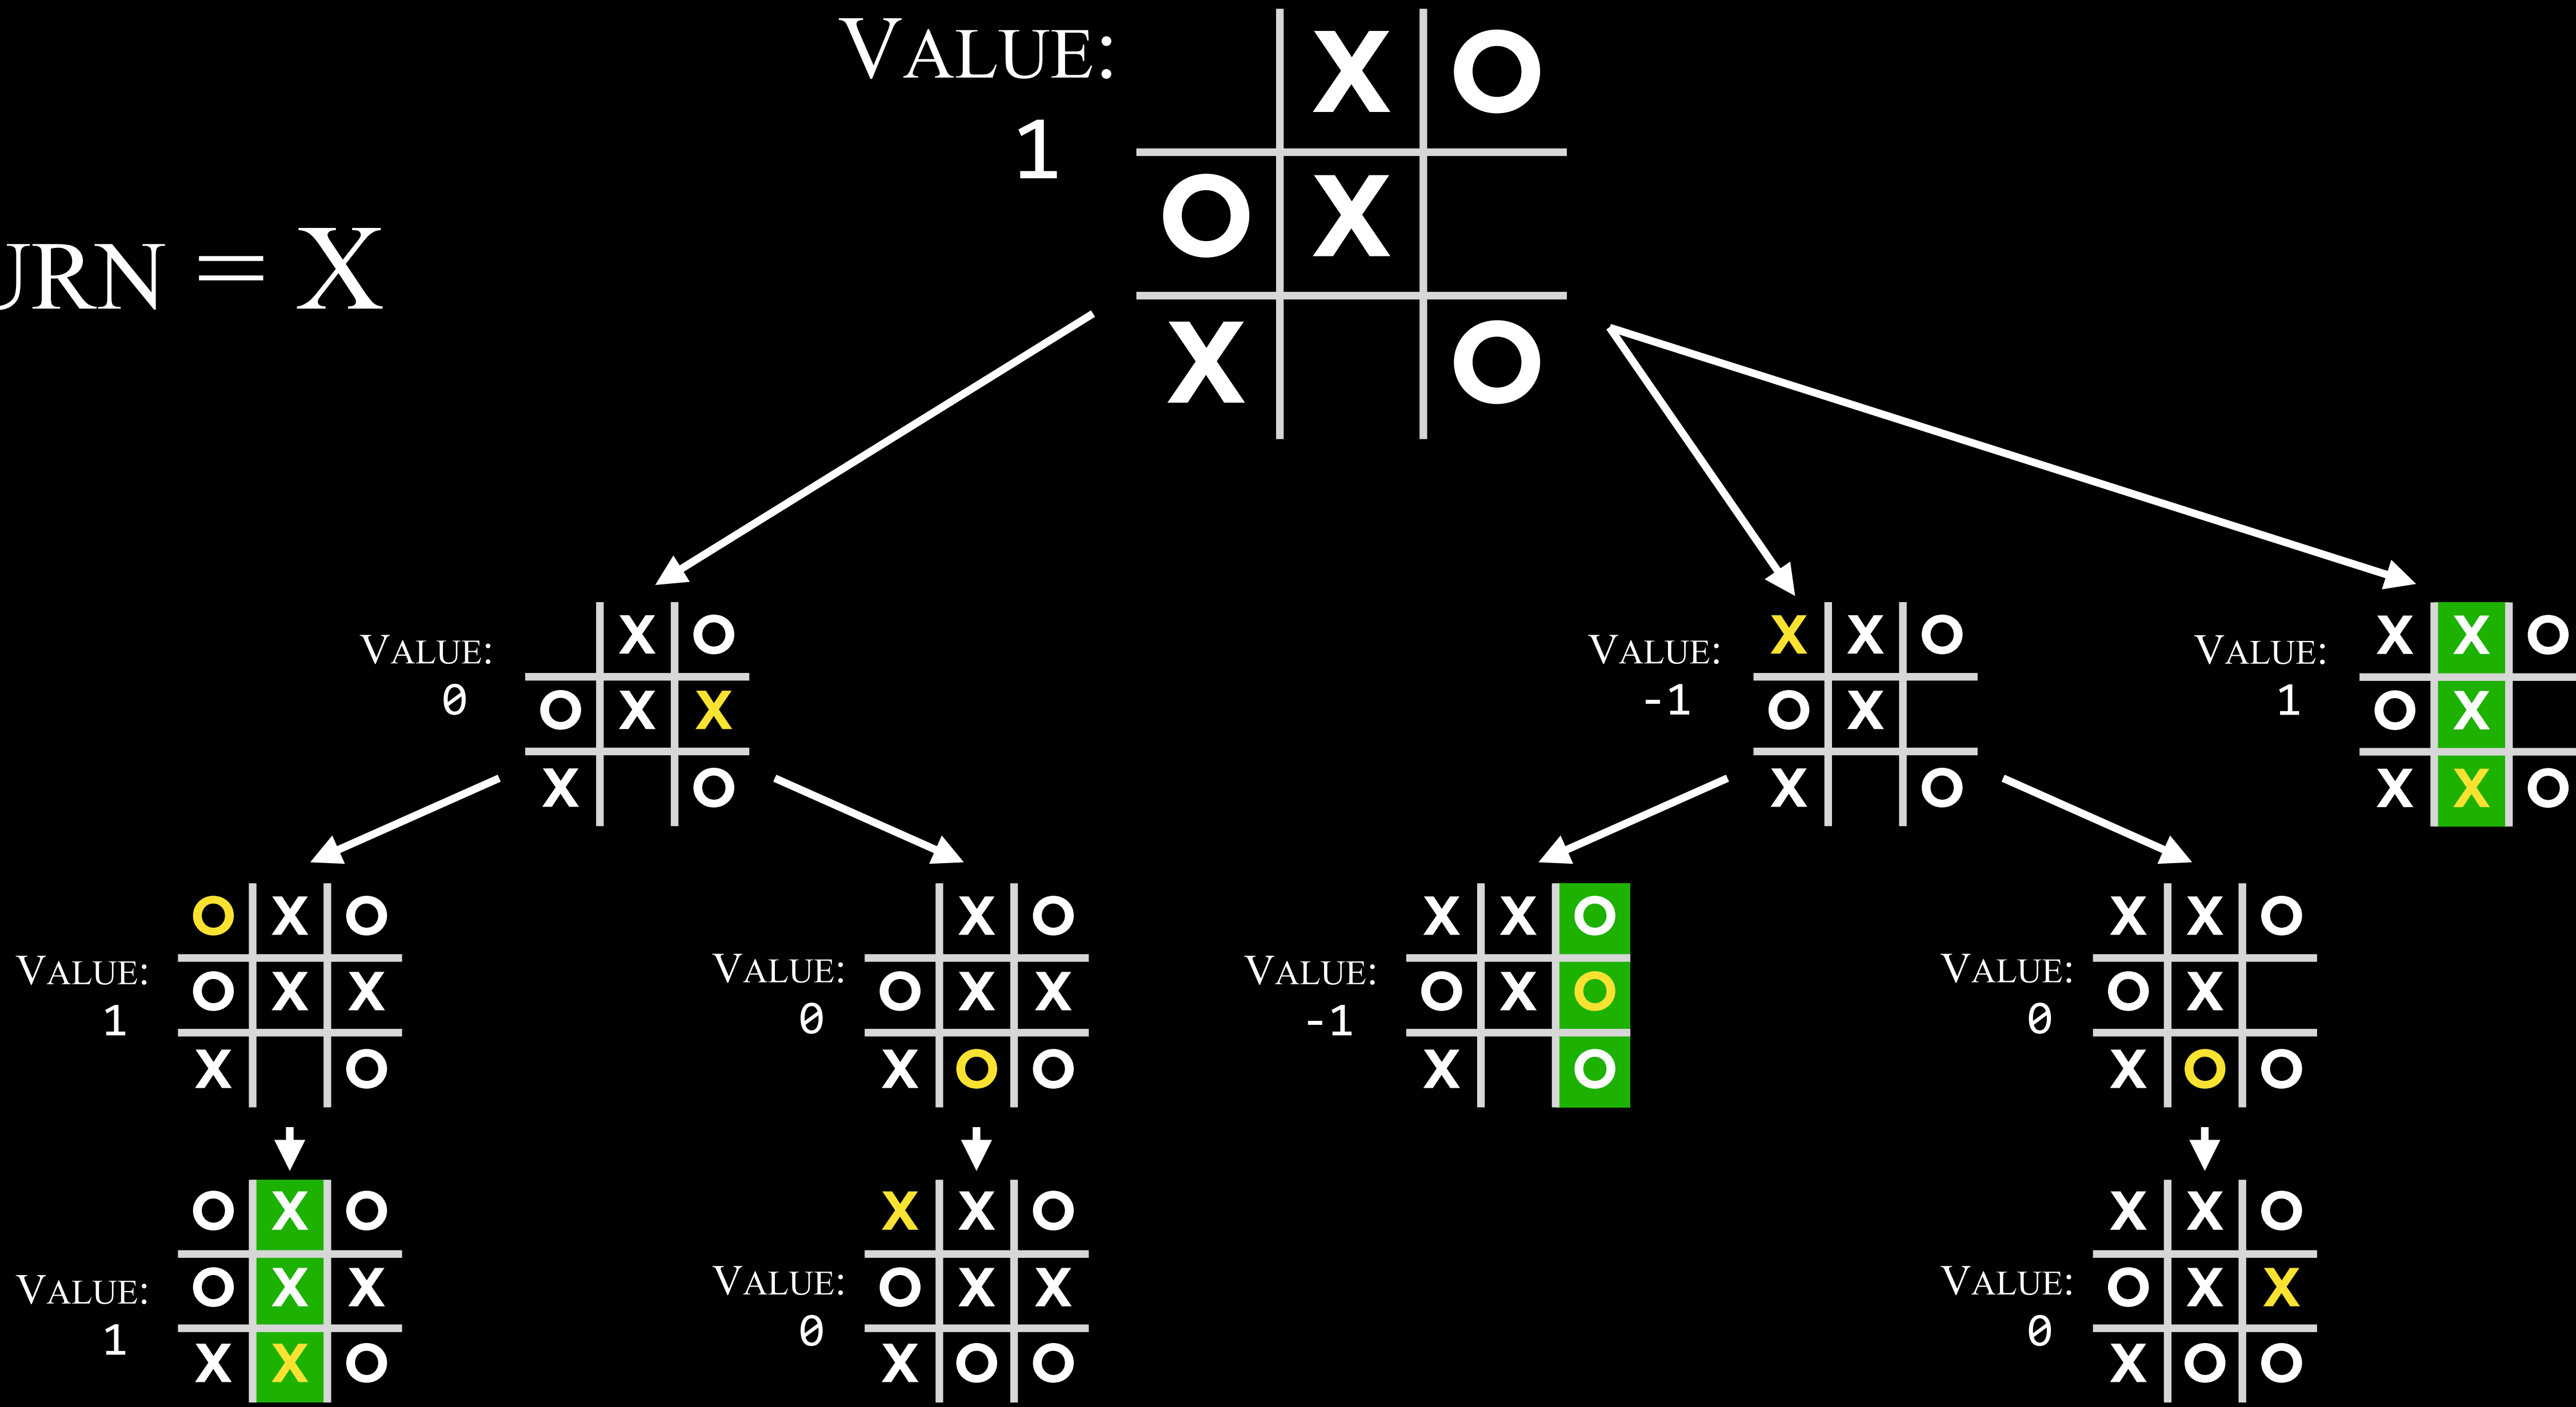
VALUE:
1

O	X	O
O	X	X
X	X	O

VALUE:
0

X	X	O
O	X	X
X	O	O

TURN = X



Minimax

```
if player is X:  
    for all possible moves:  
        calculate score for board  
    choose move with highest score  
  
else:  
    for all possible moves:  
        calculate score for board  
    choose move with lowest score
```


255,168

total possible Tic-Tac-Toe games



288,000,000,000

total possible chess games
for the first four moves

10^{29241}

total possible chess games
(lower bound)

Depth-Limited Minimax

evaluation function

function that estimates the expected utility of the game from a given state

Search

1

2

3

4

5

6

7

8

9

10

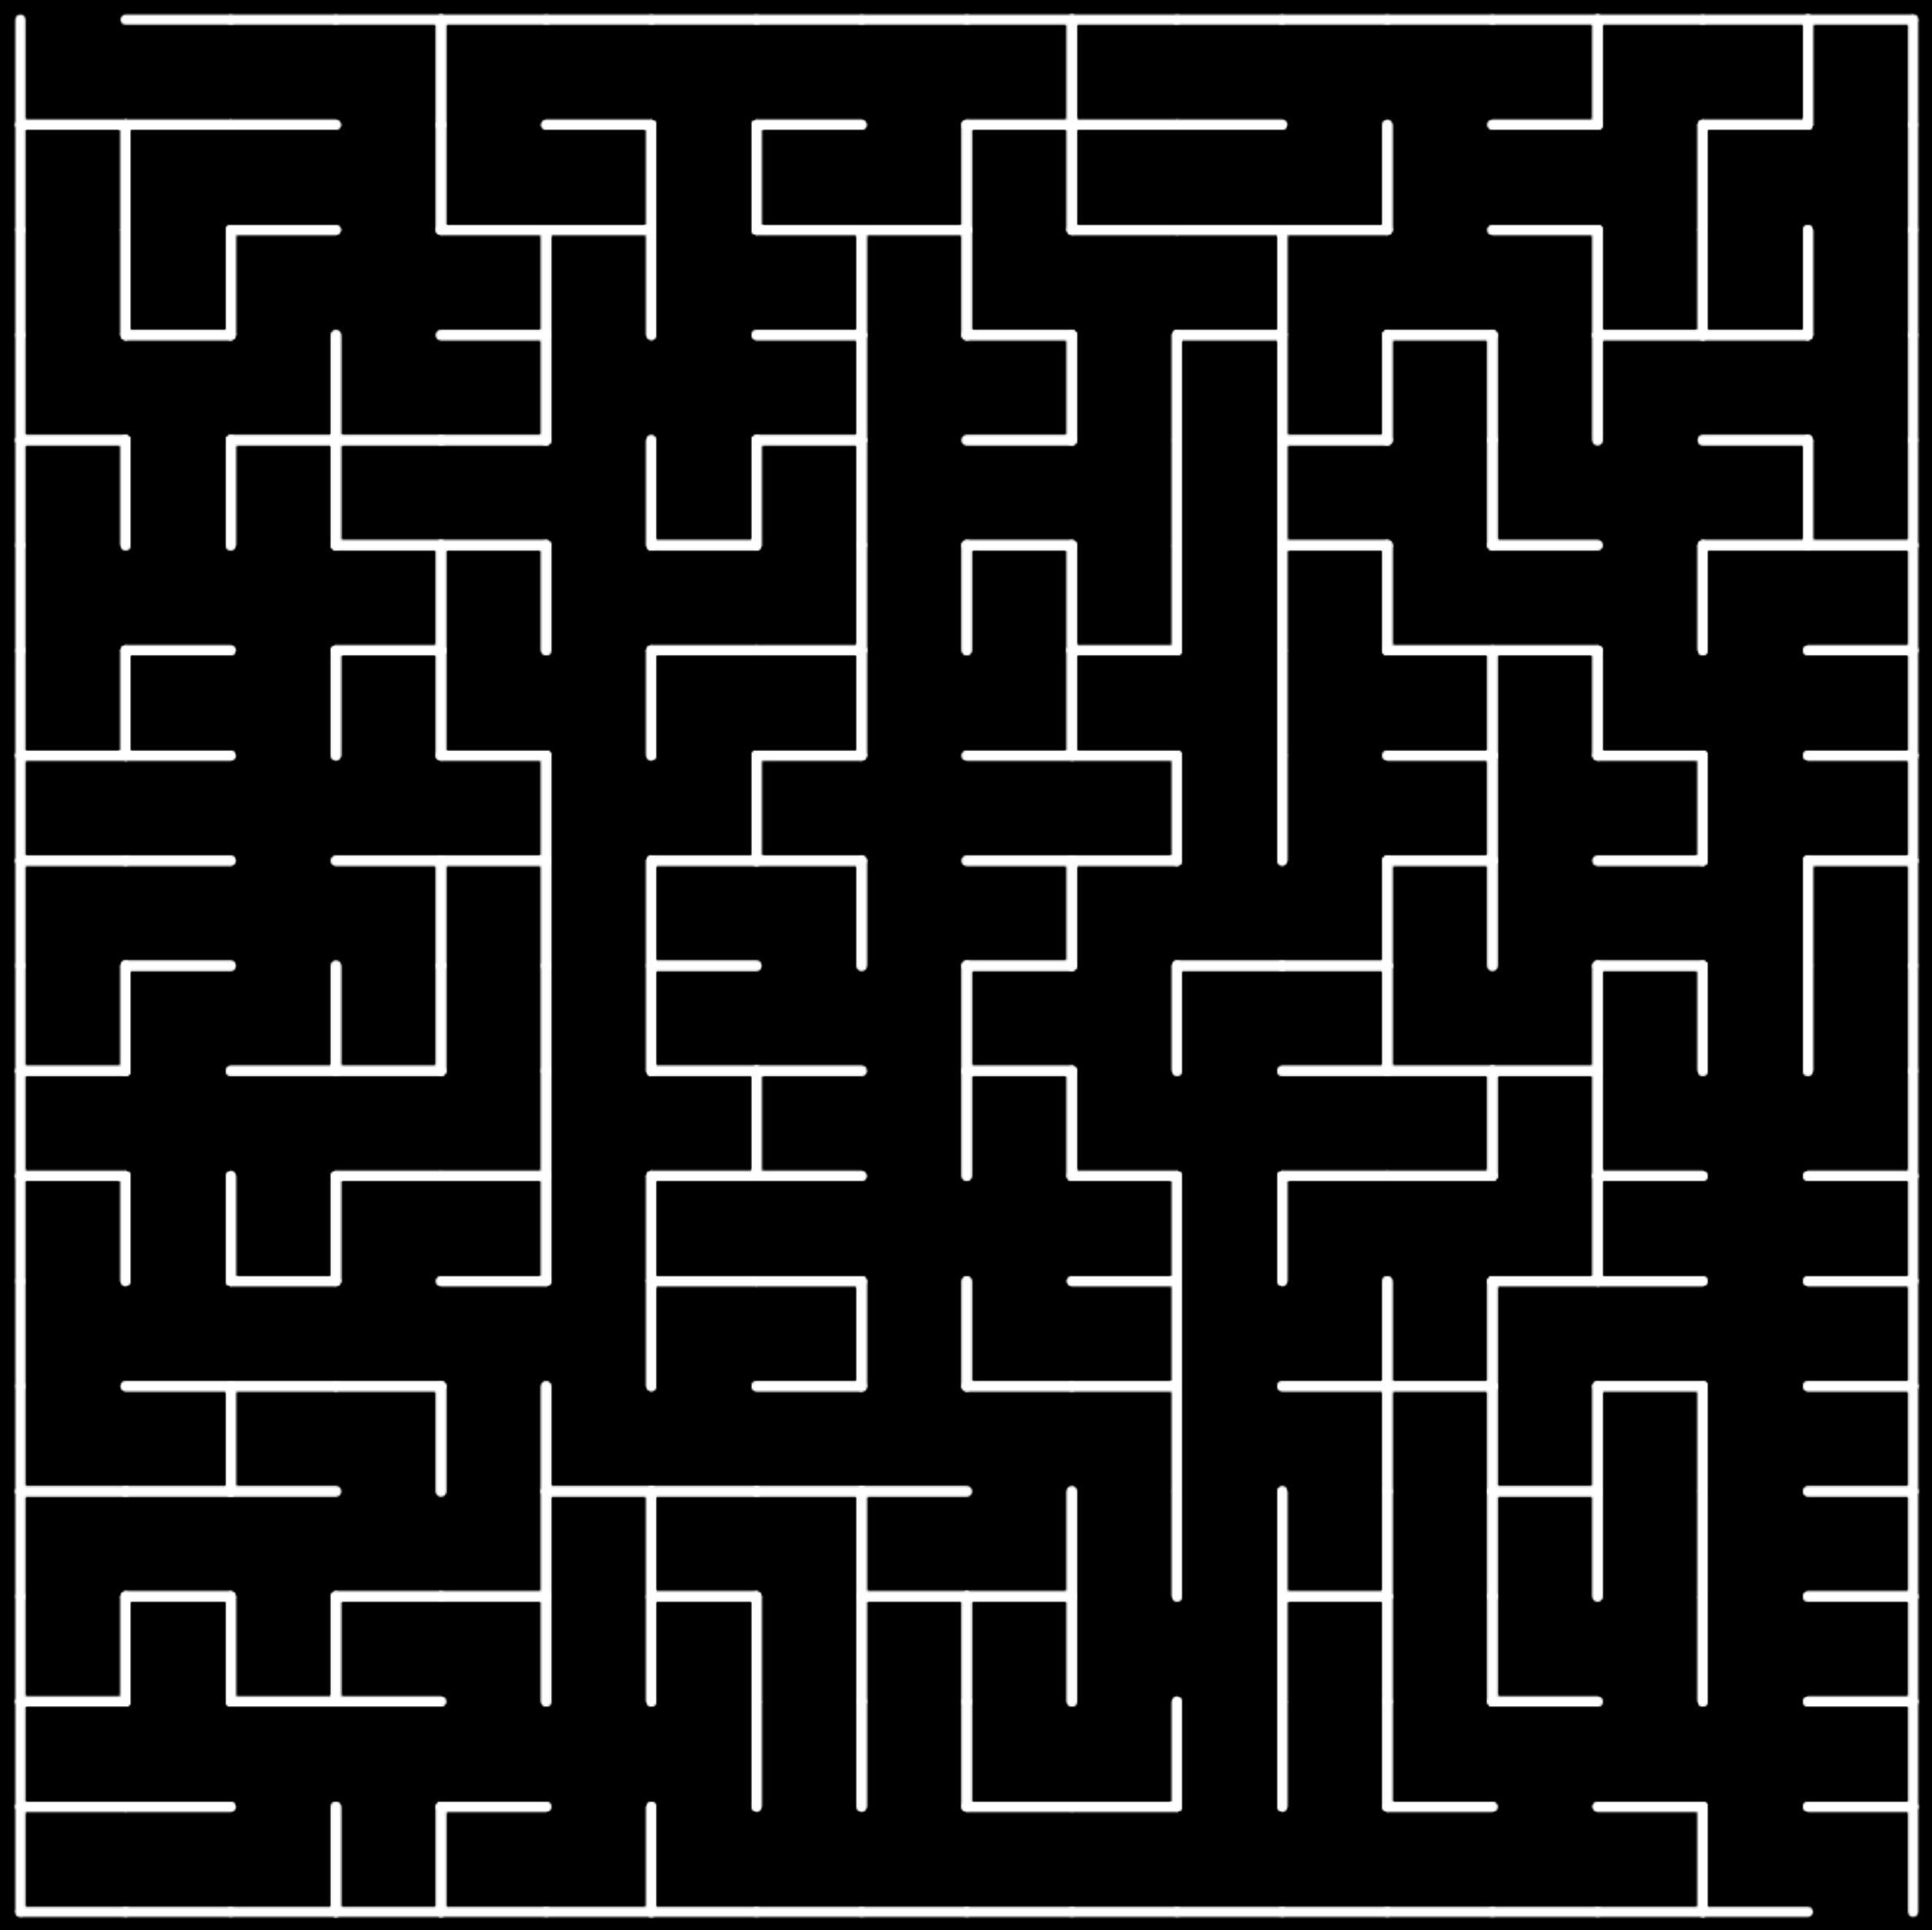
11

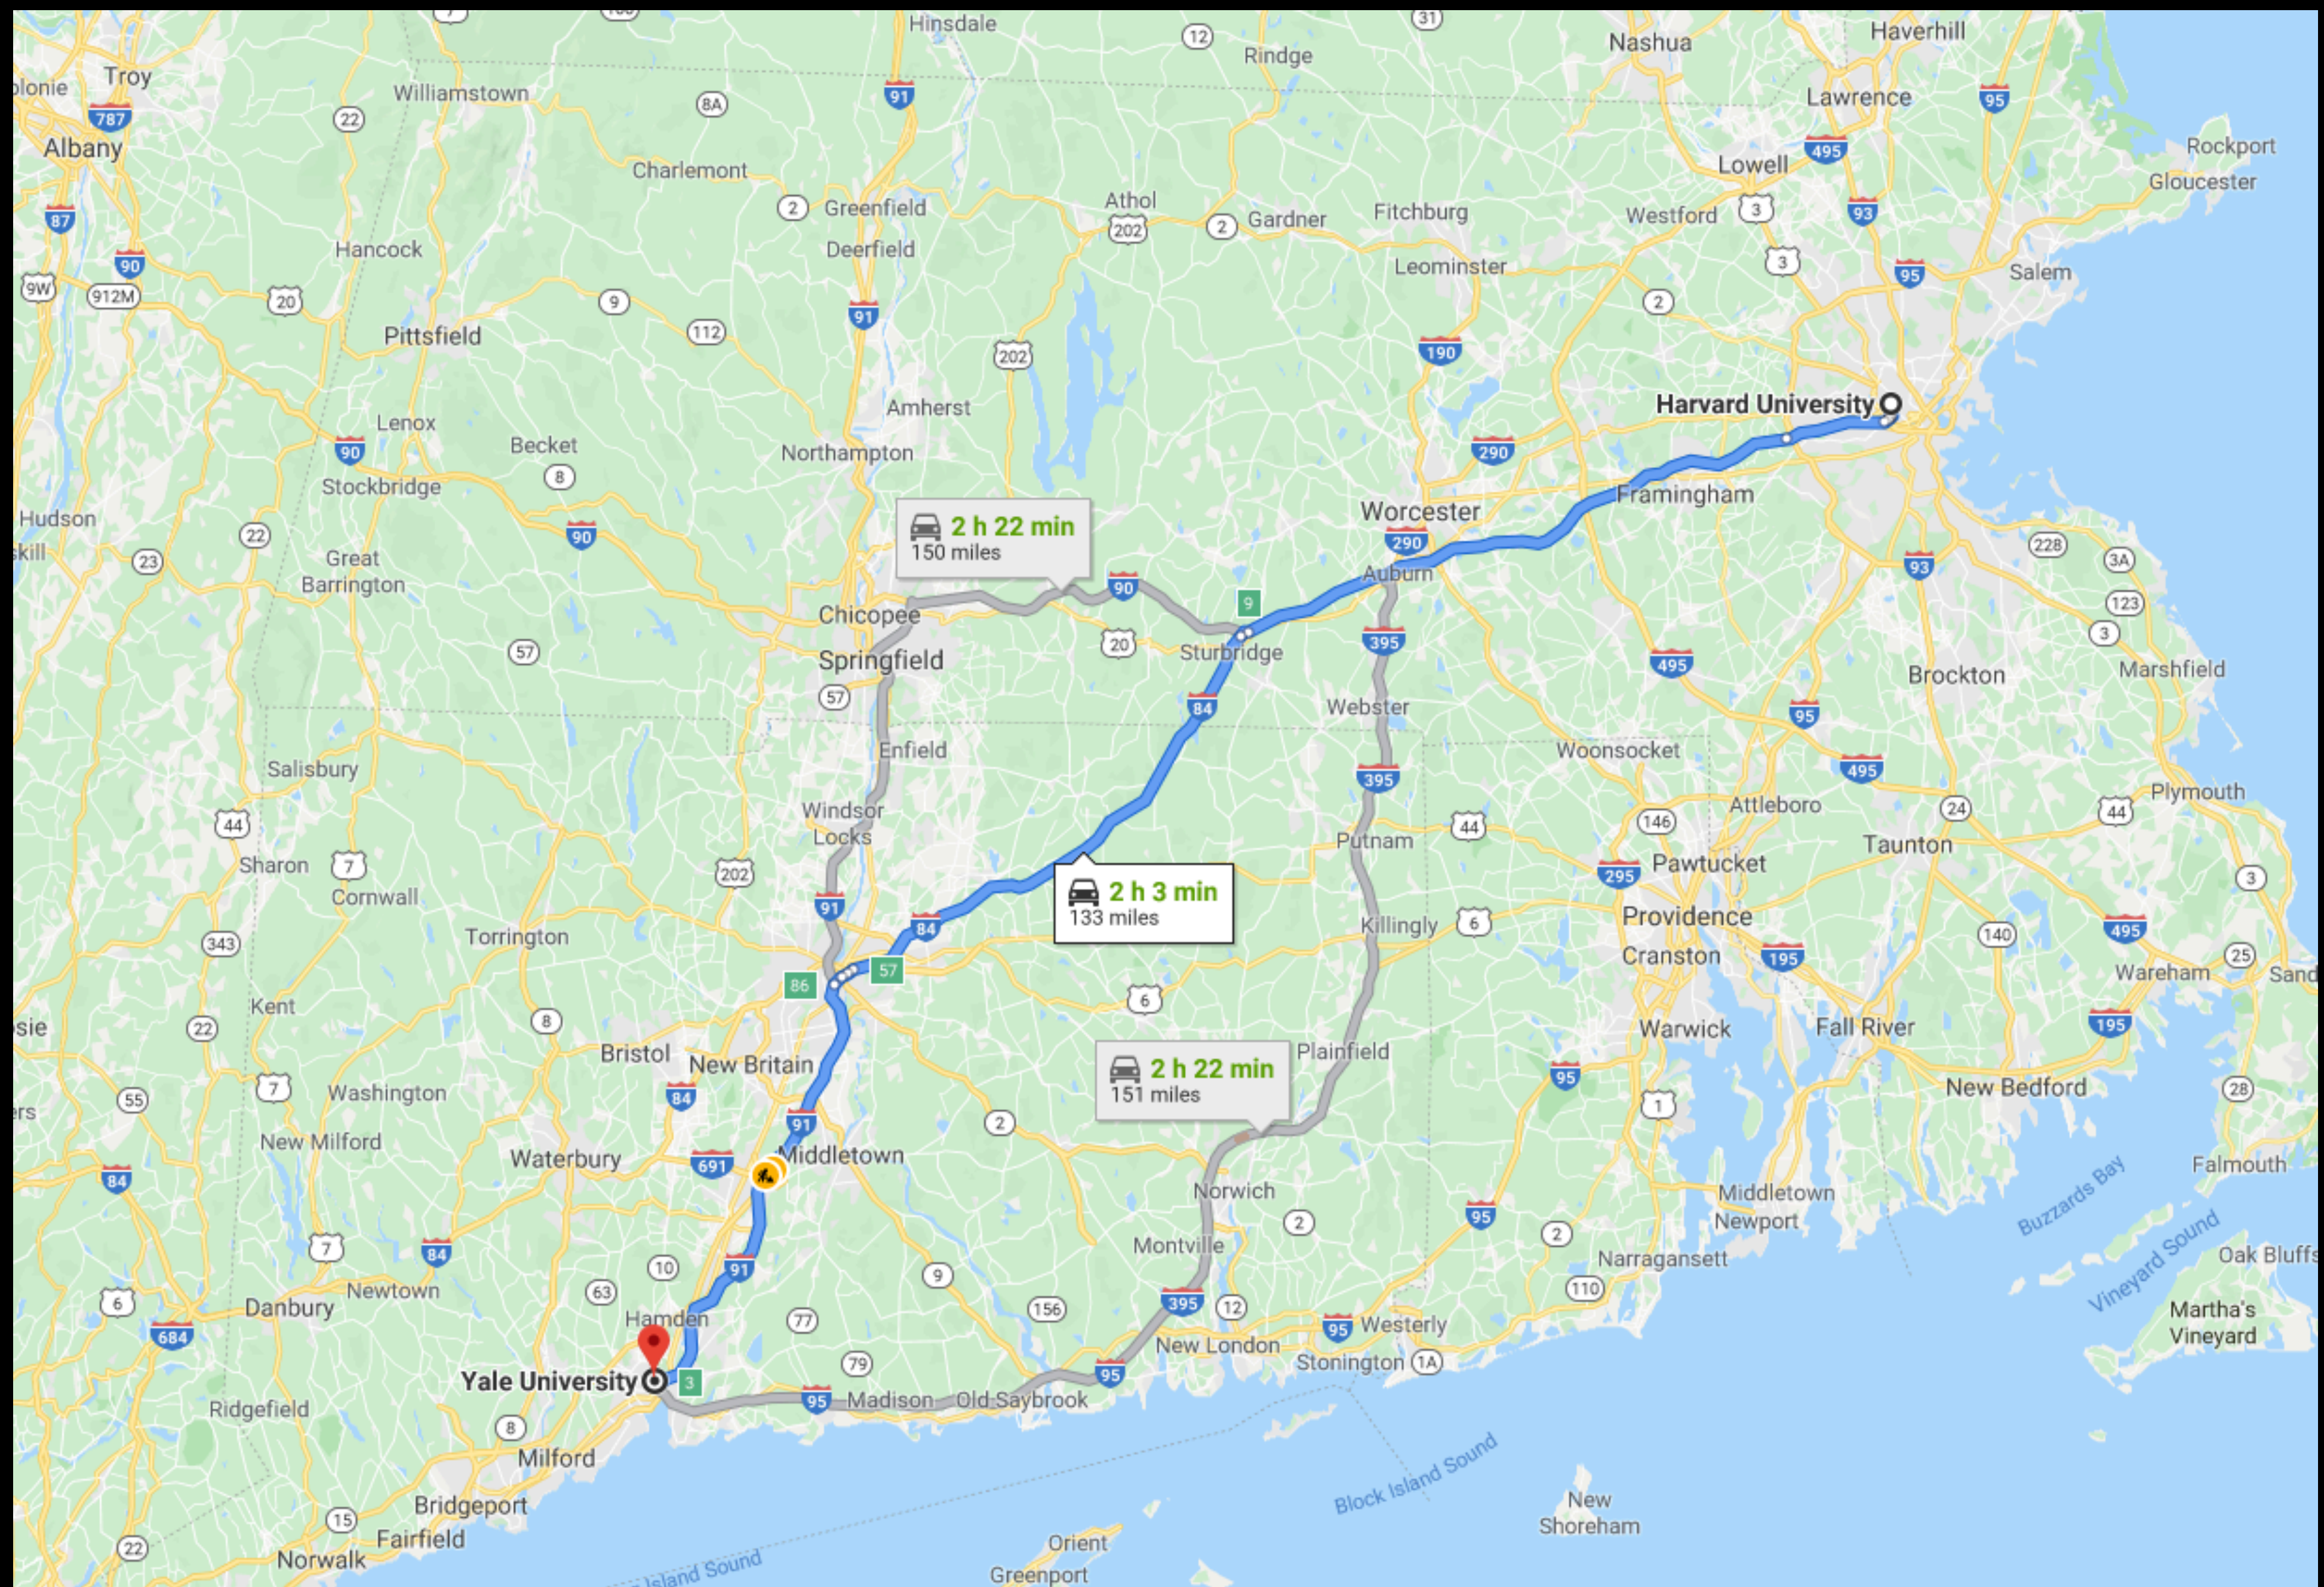
12

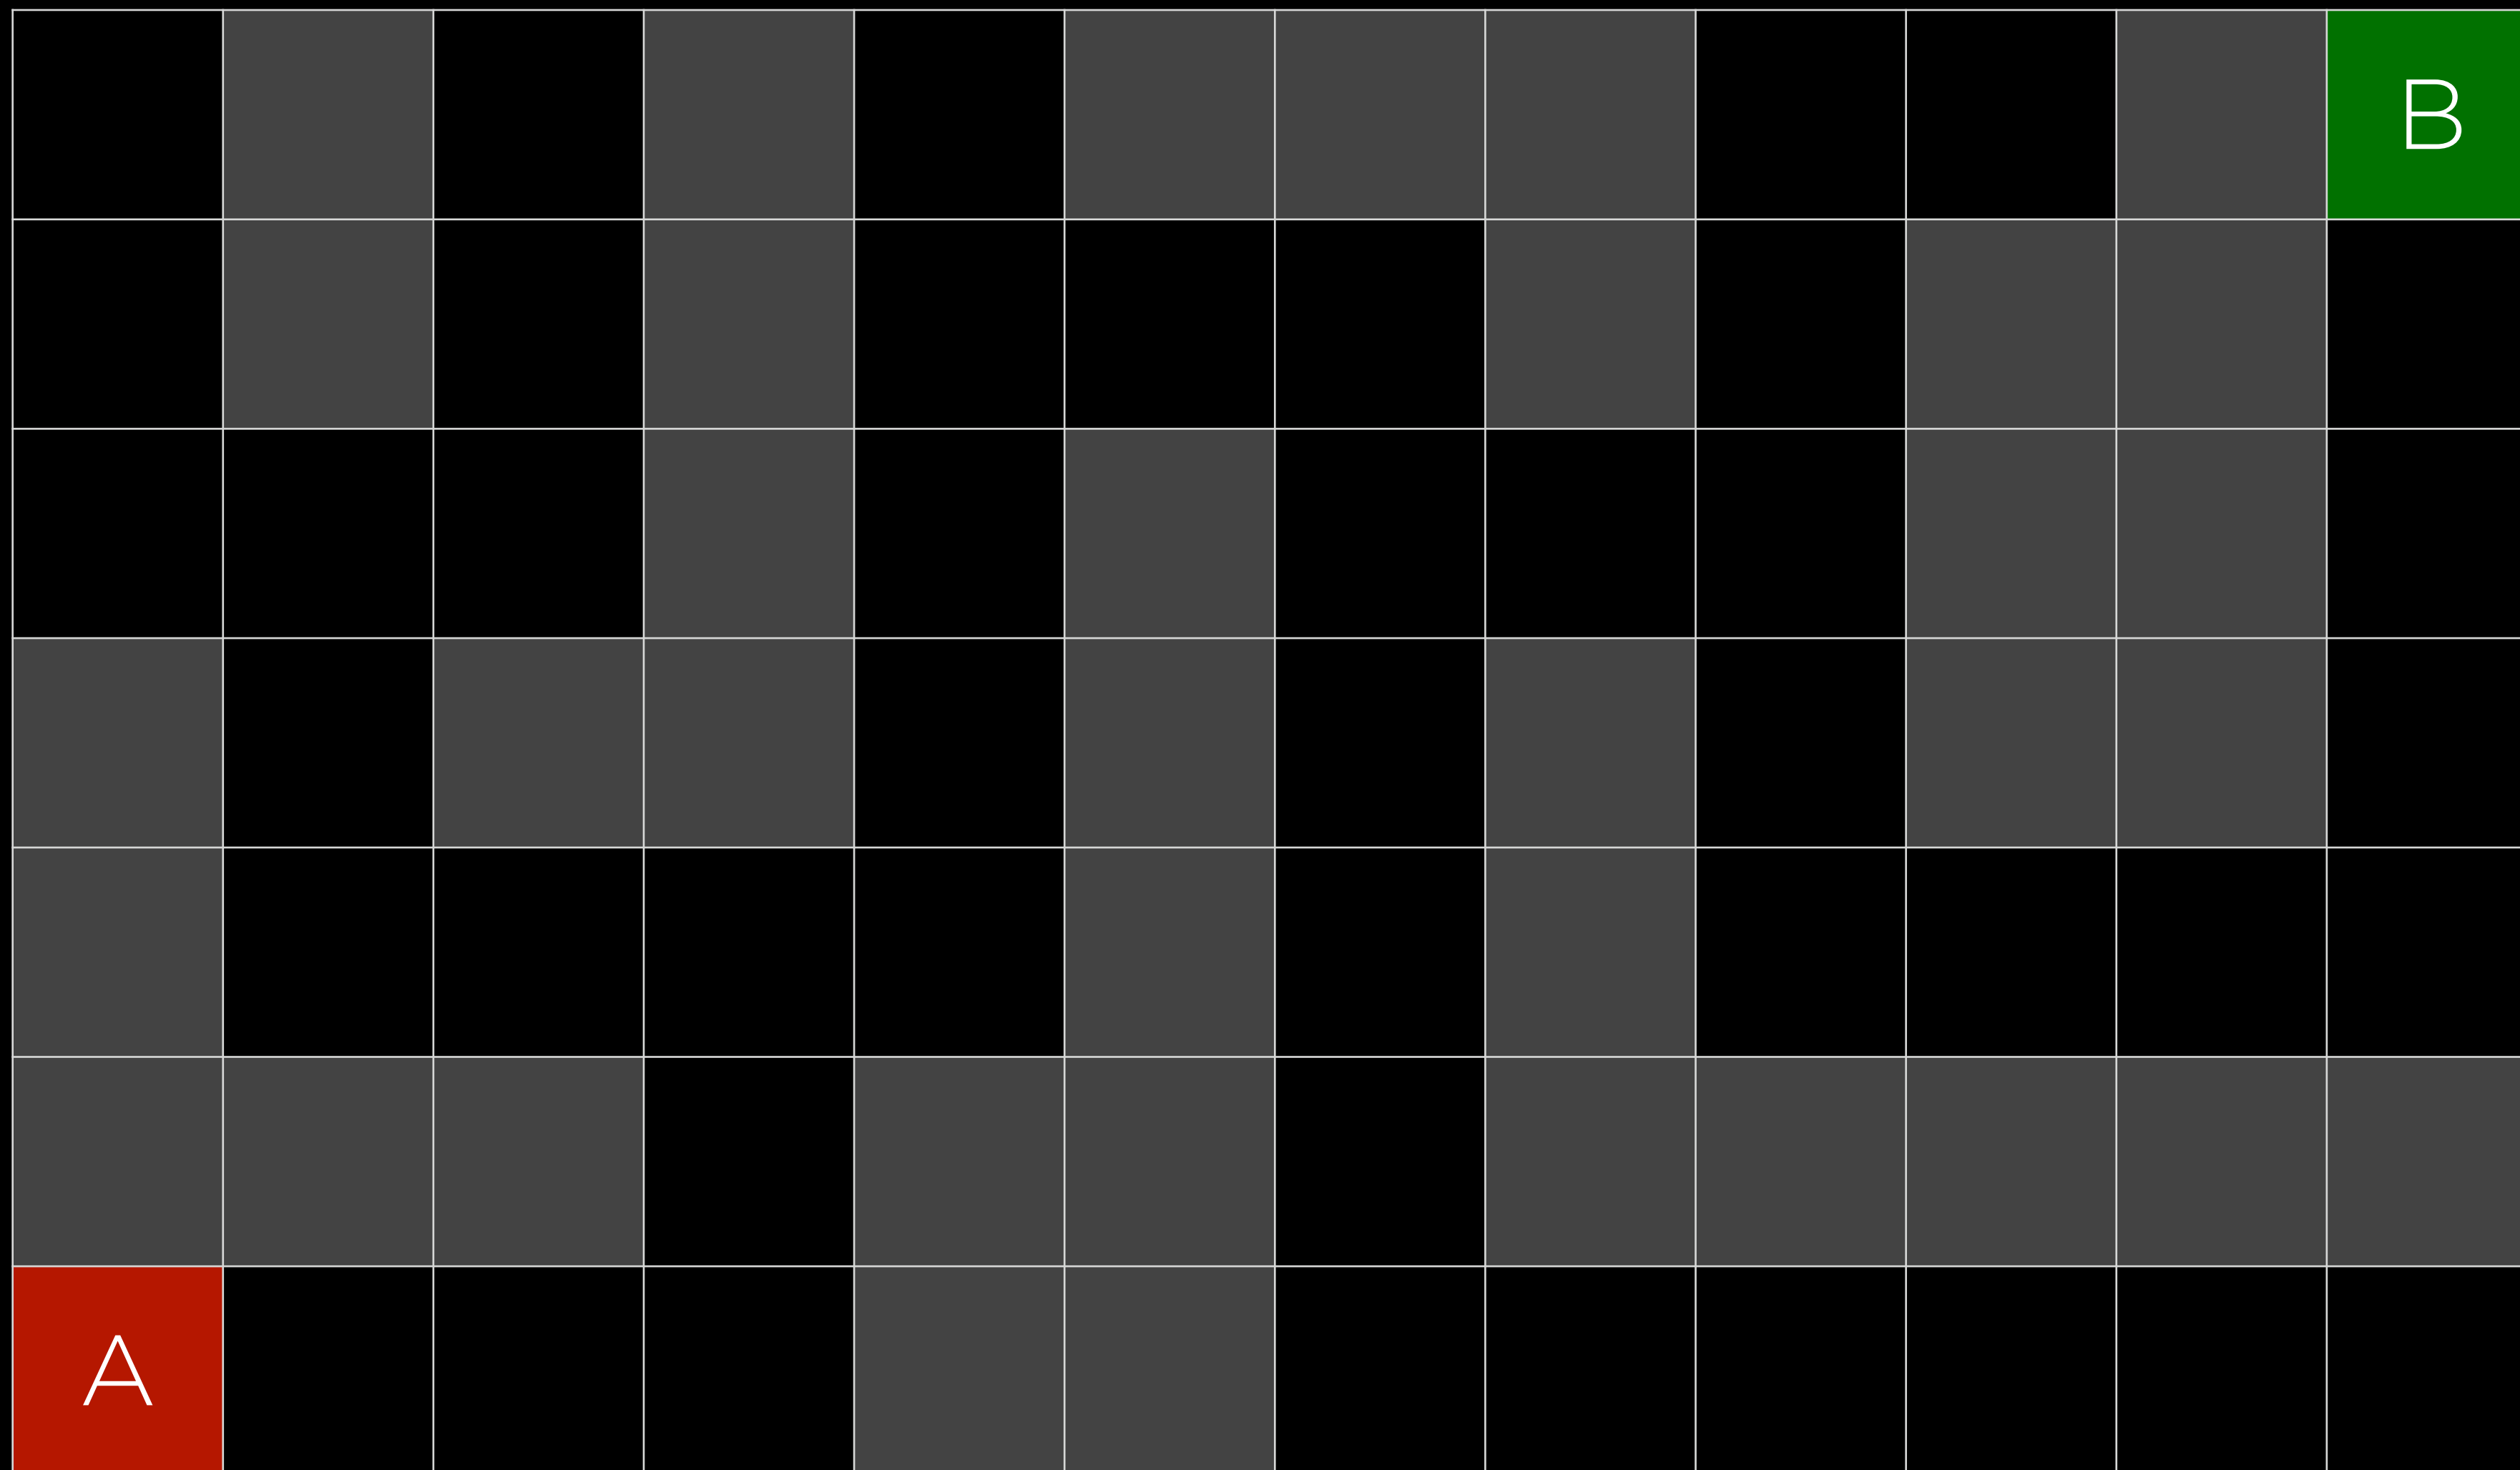
13

14

15

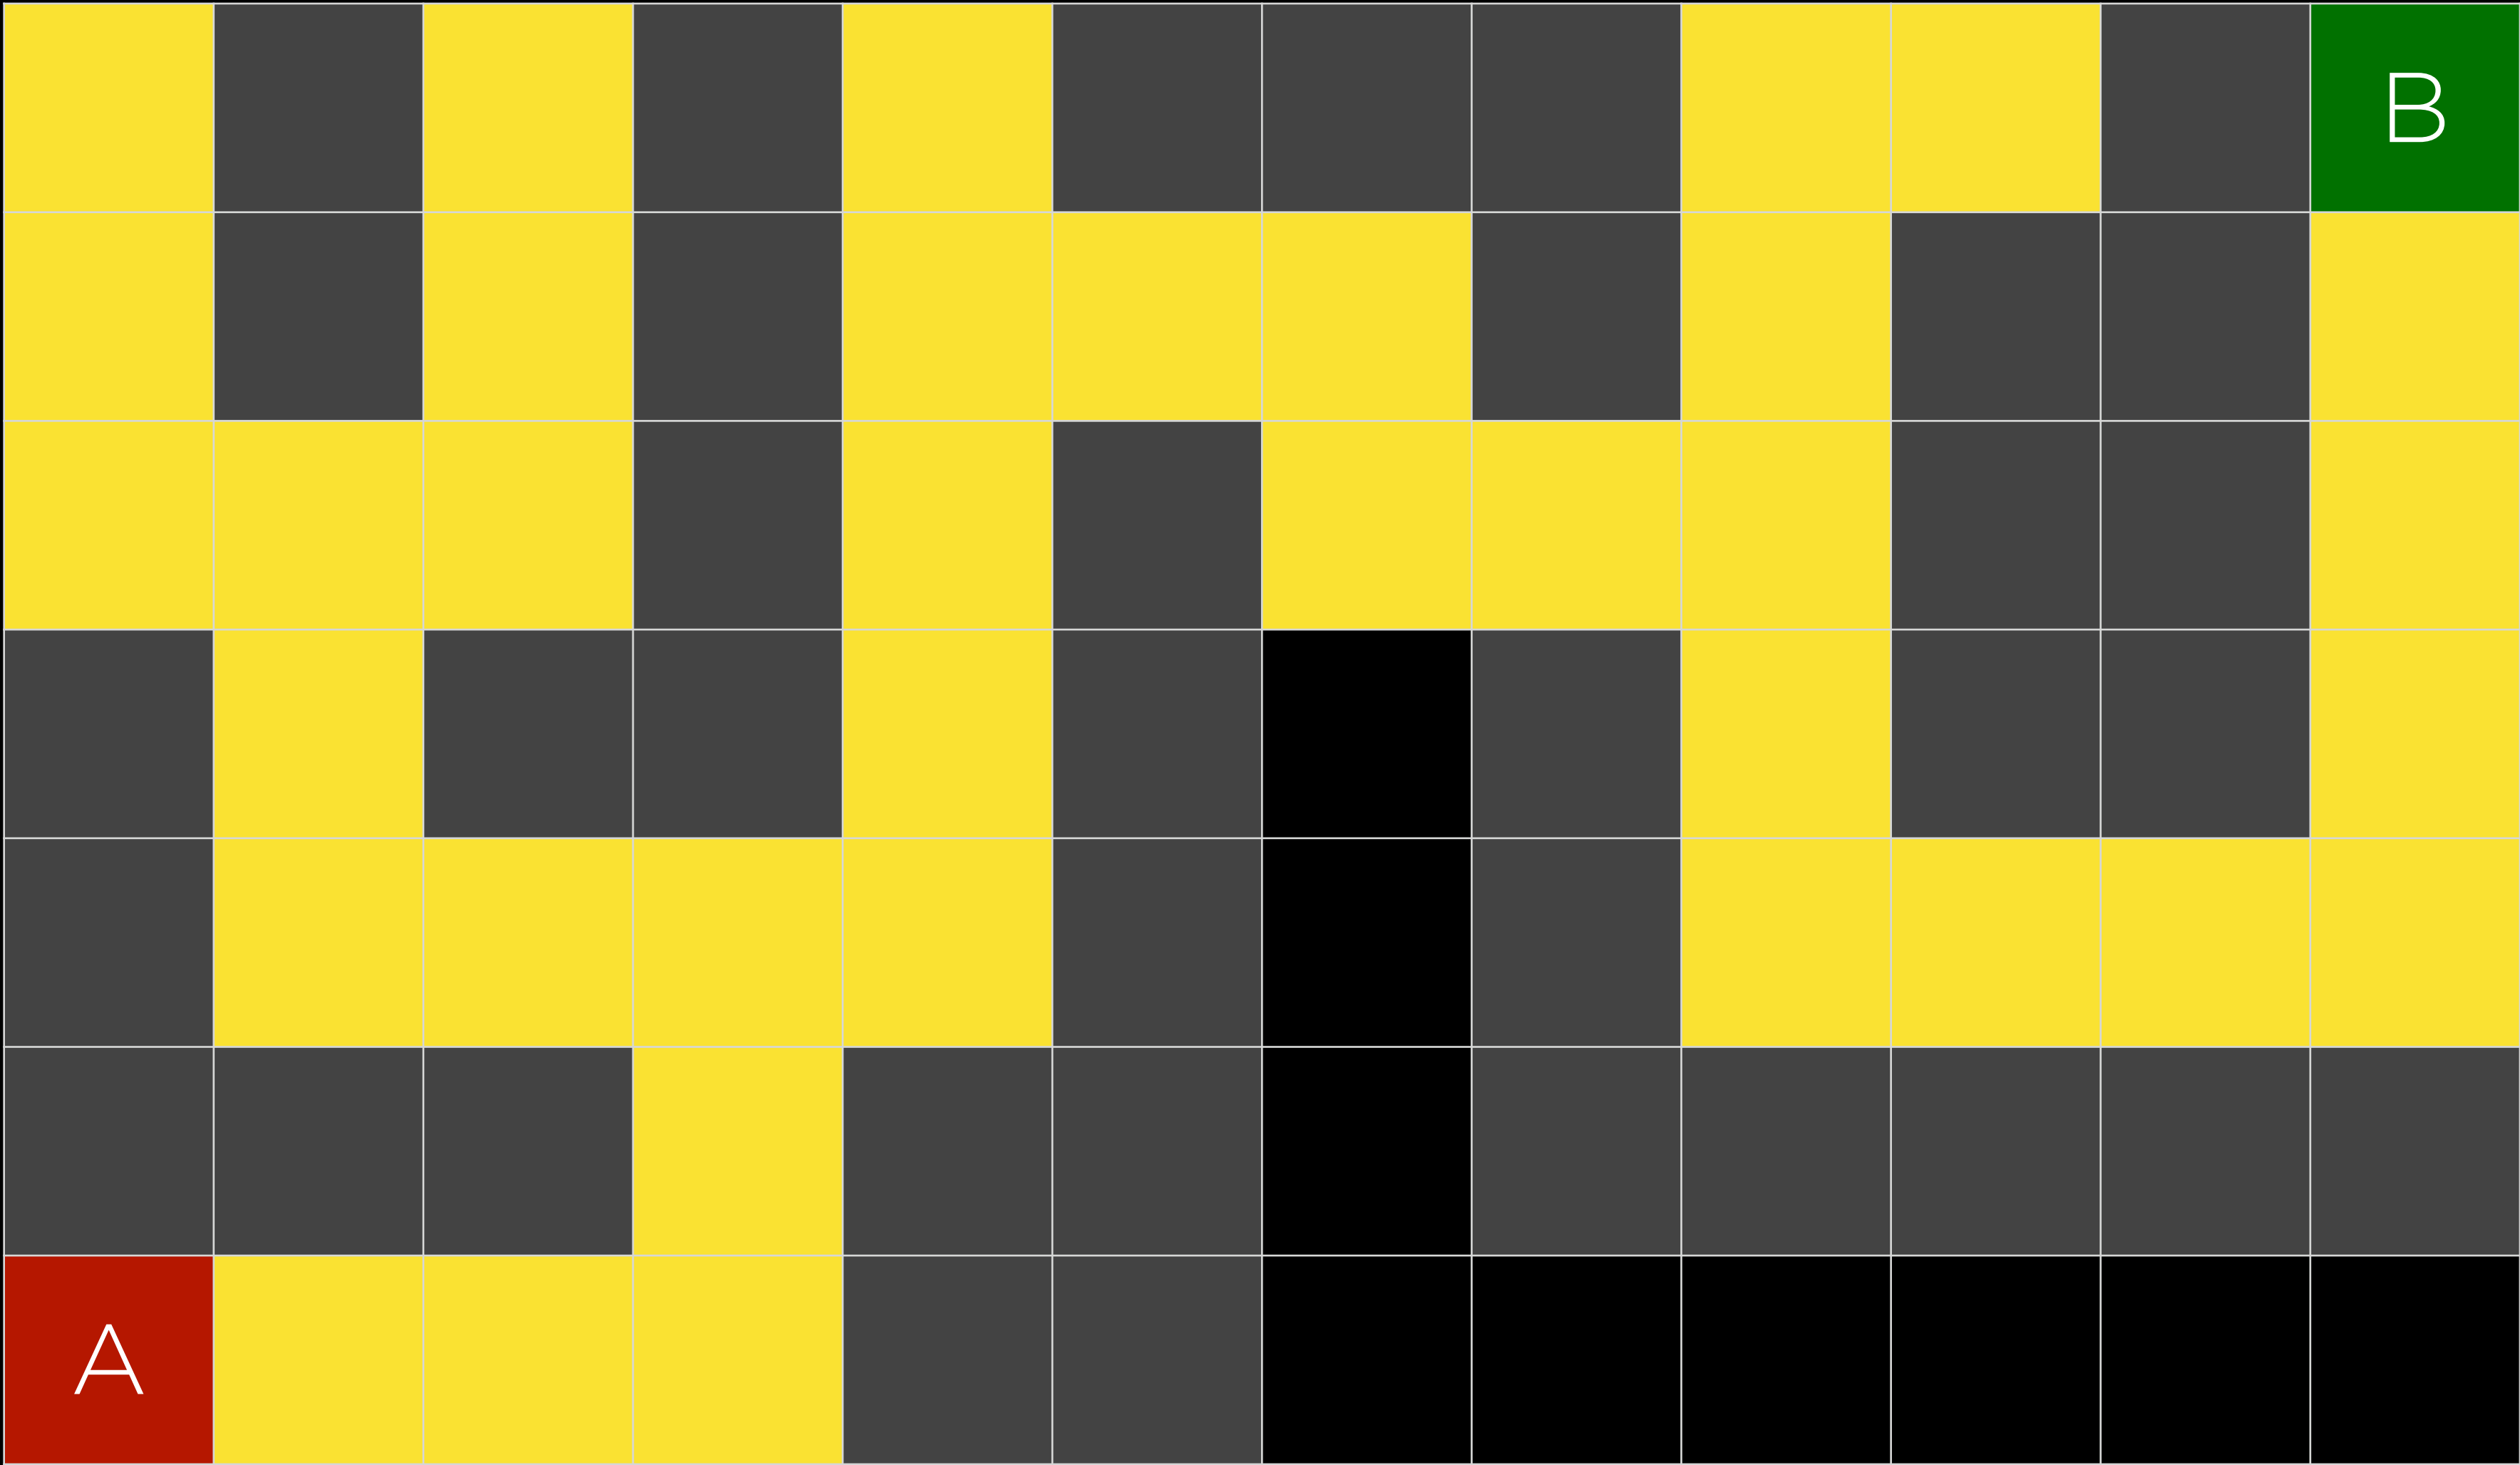




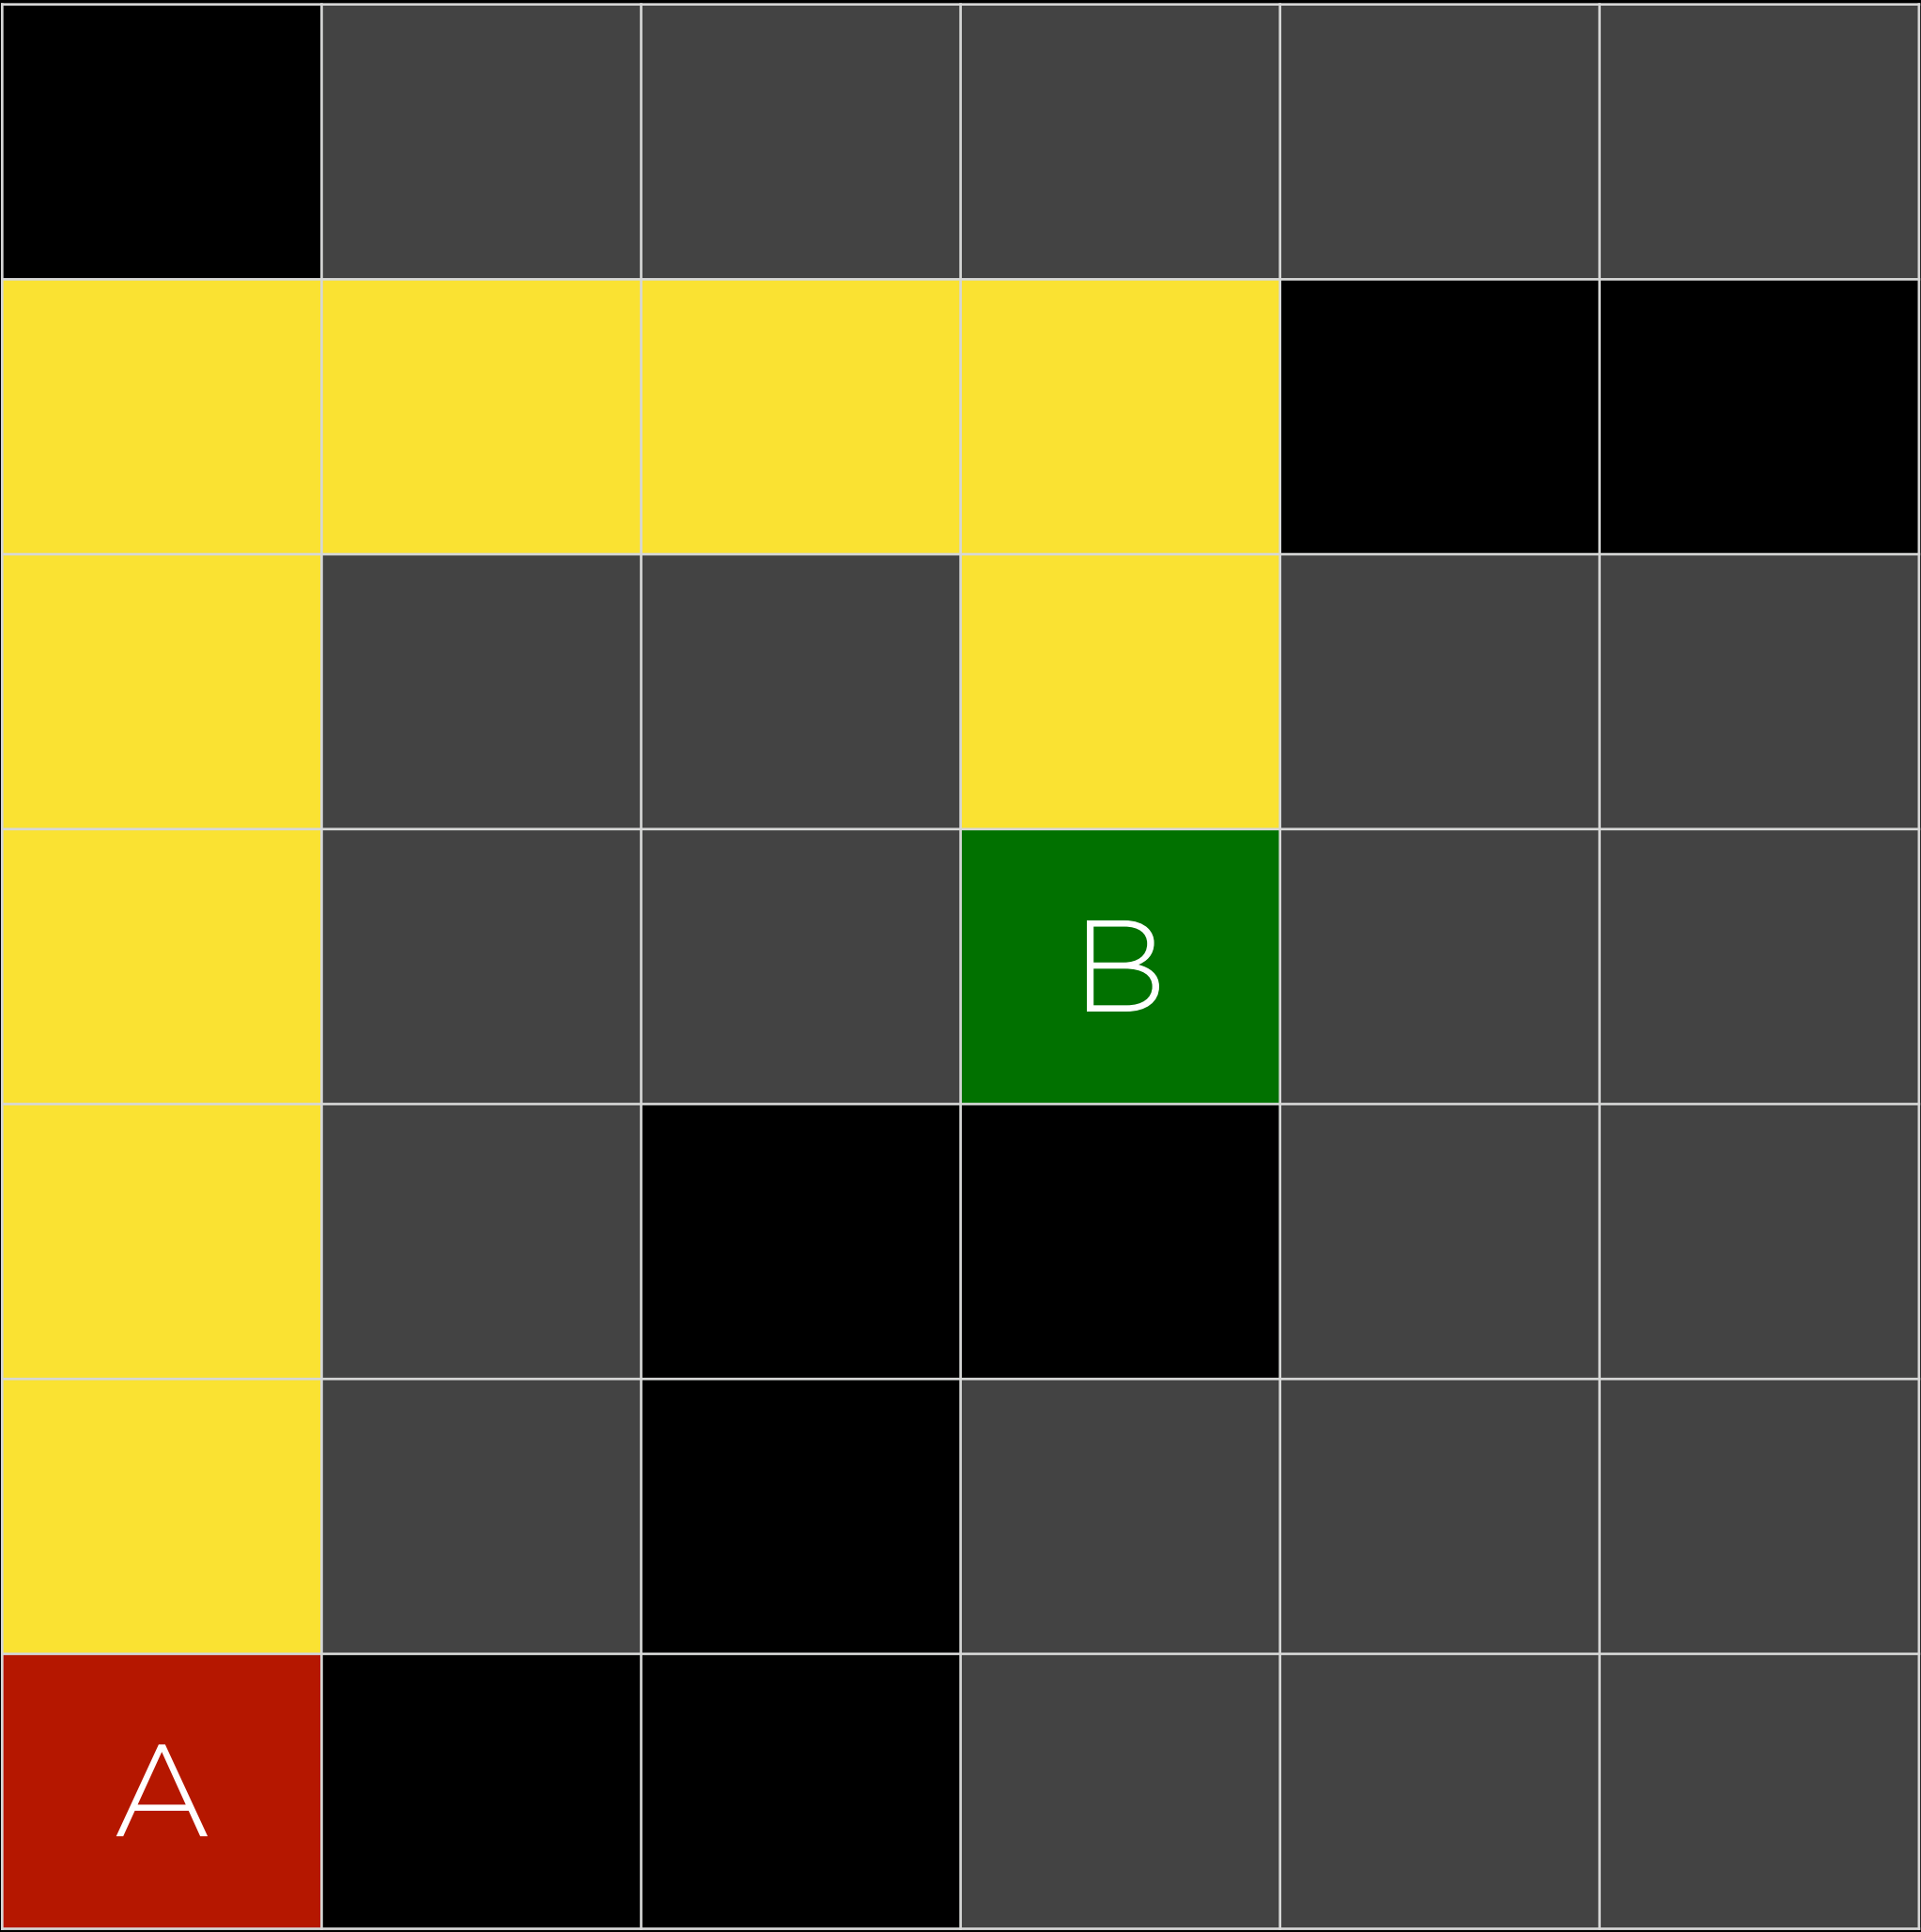


Depth-First Search

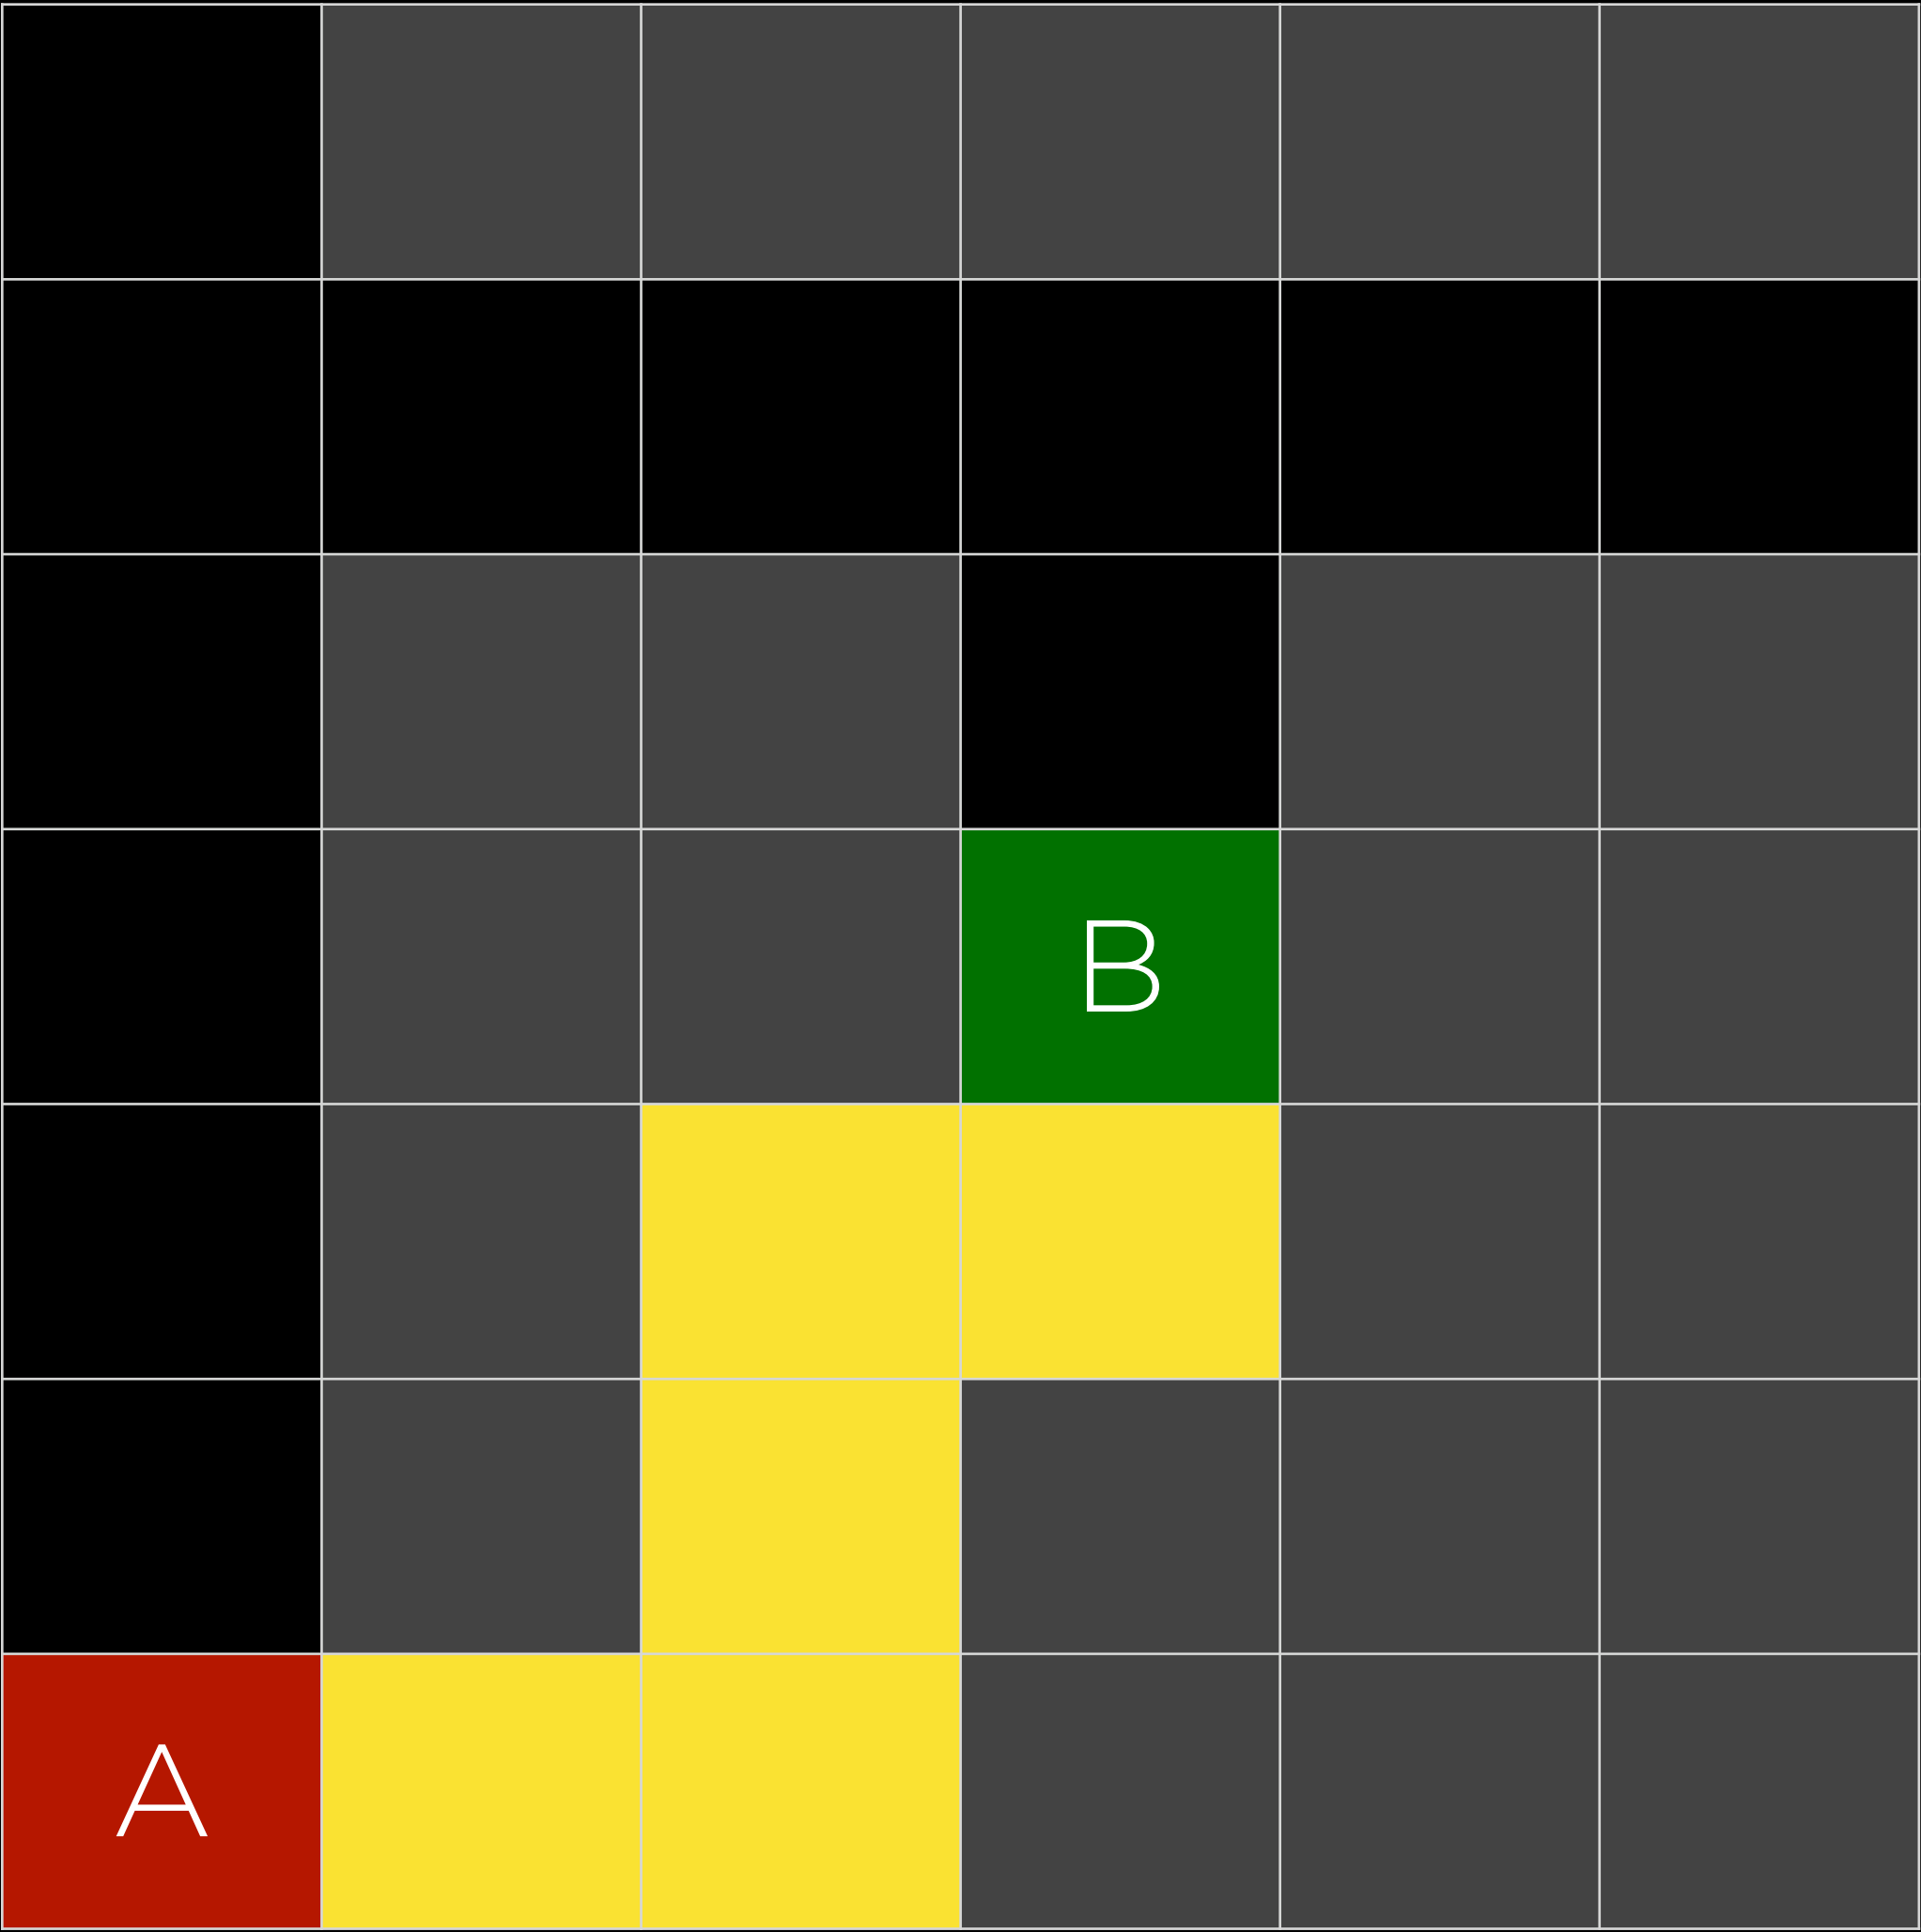
Depth-First Search



Depth-First Search

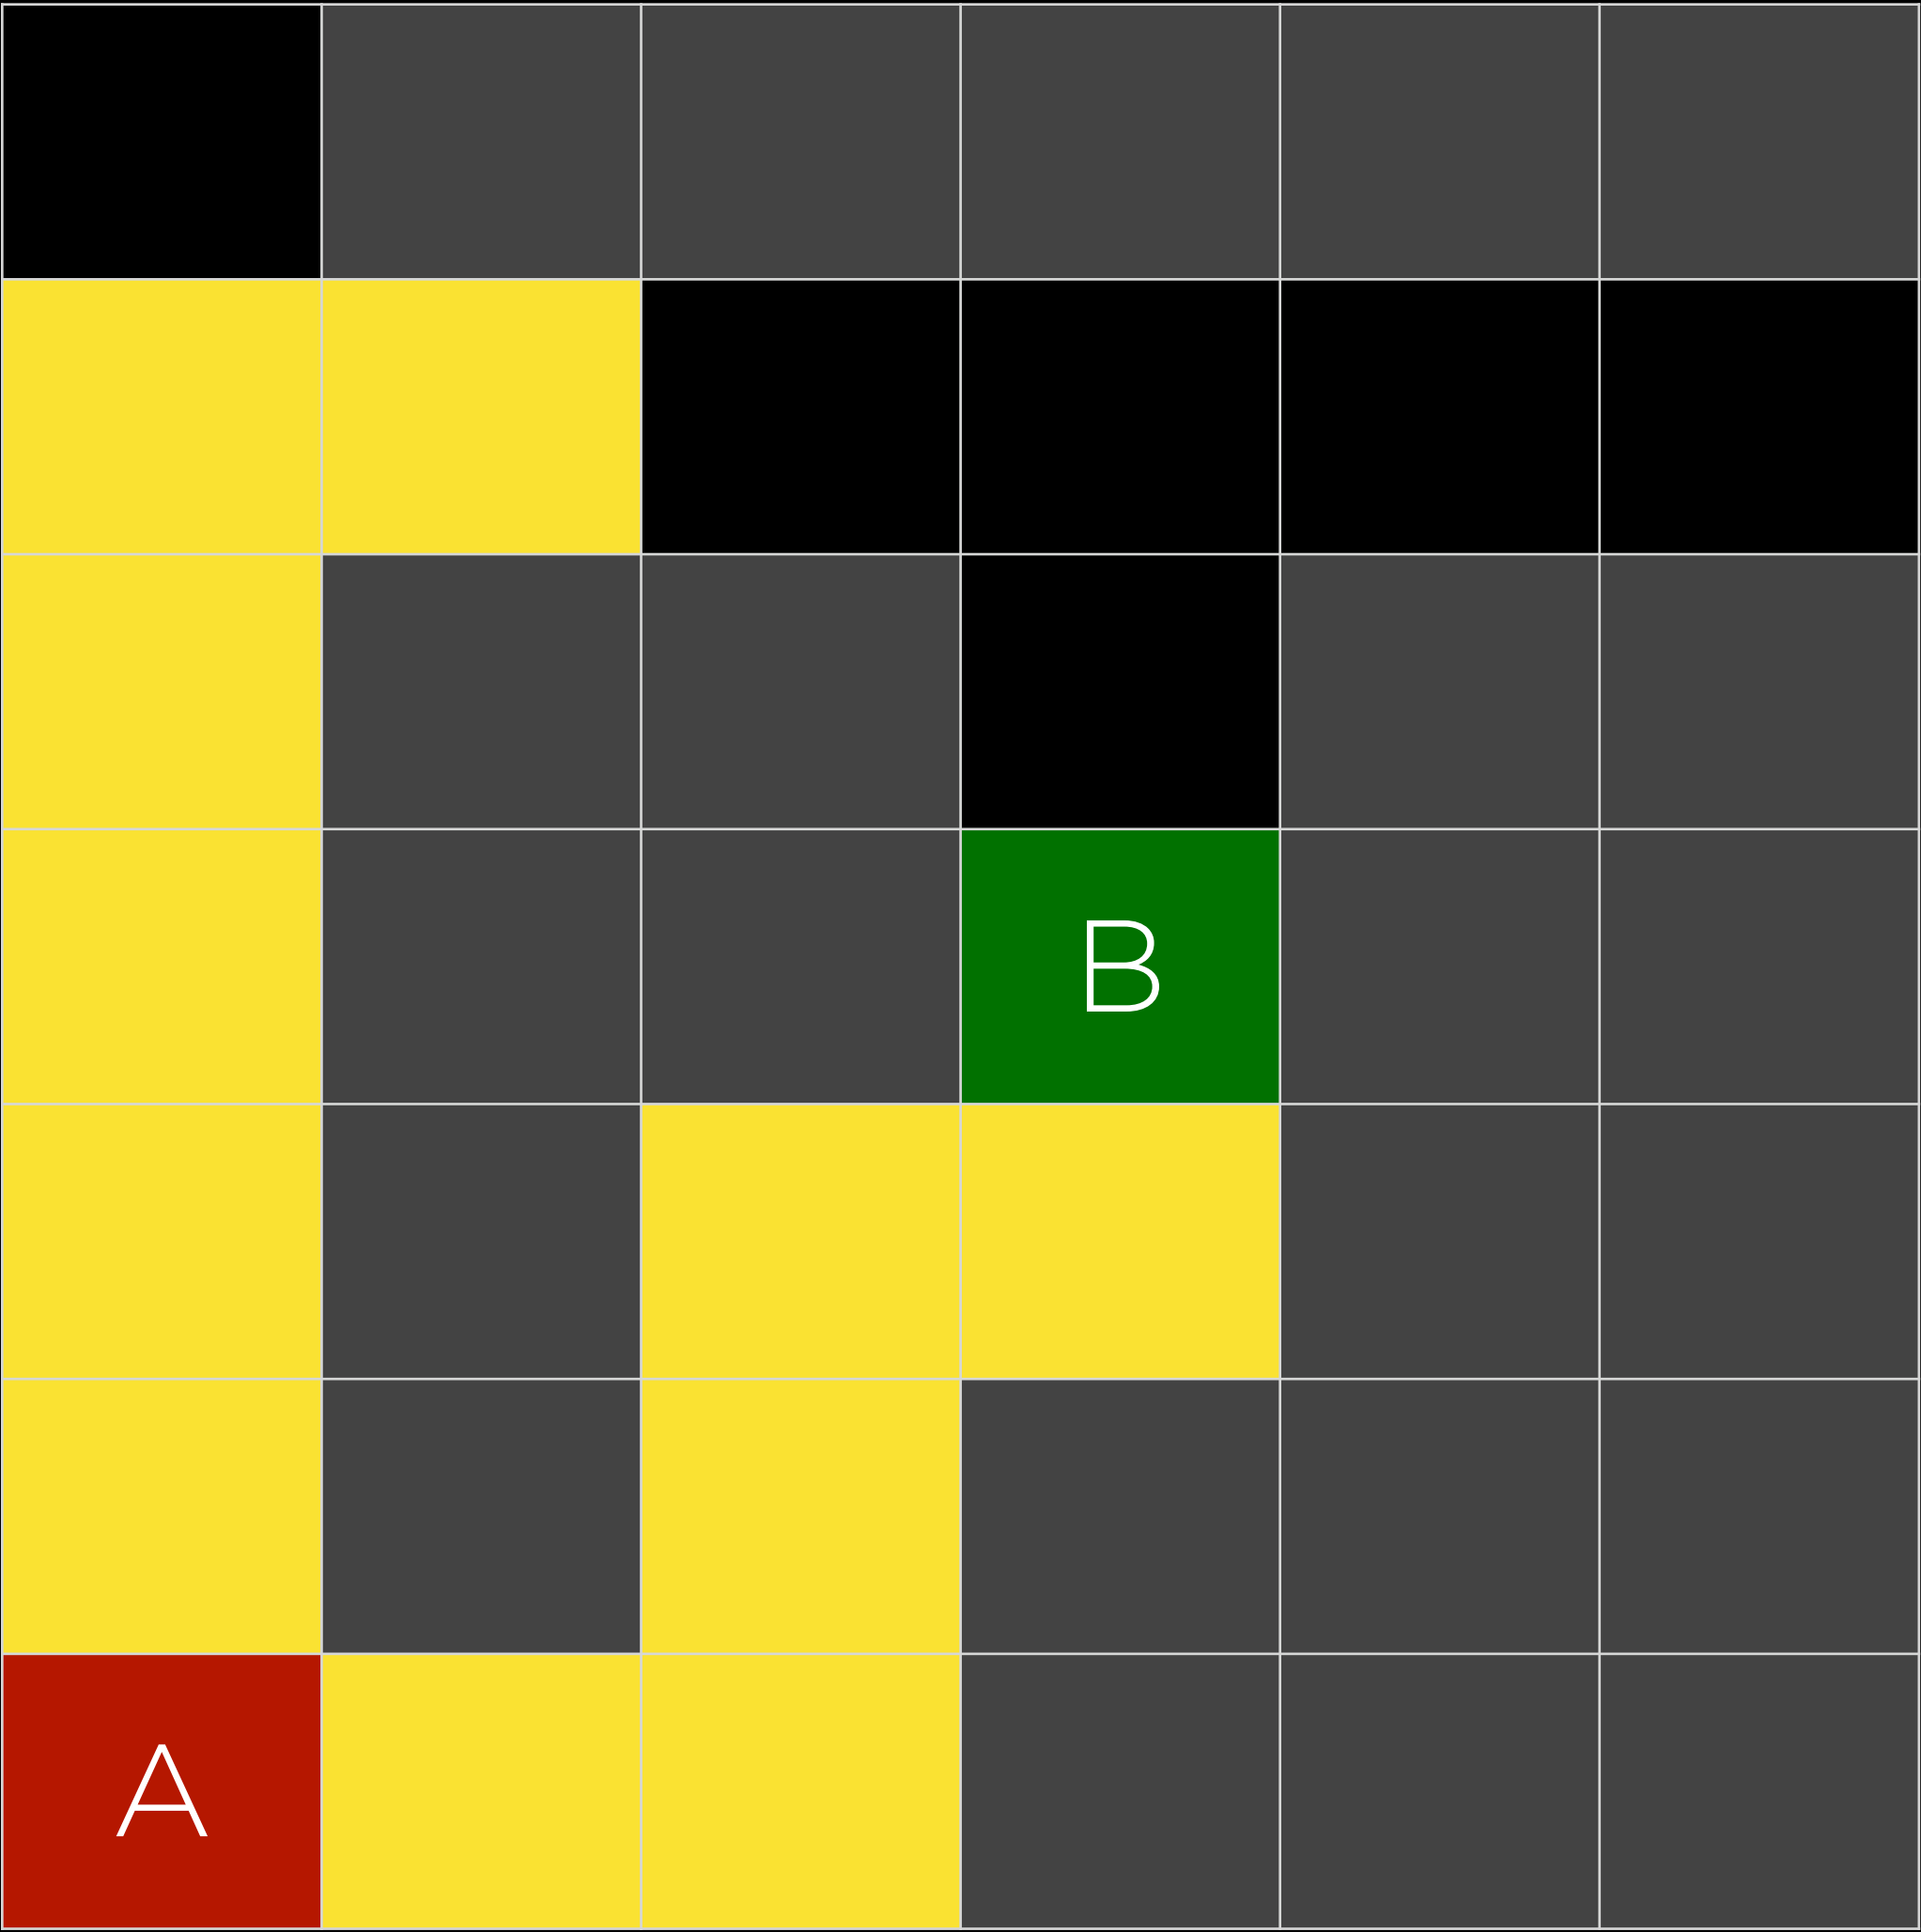


Depth-First Search

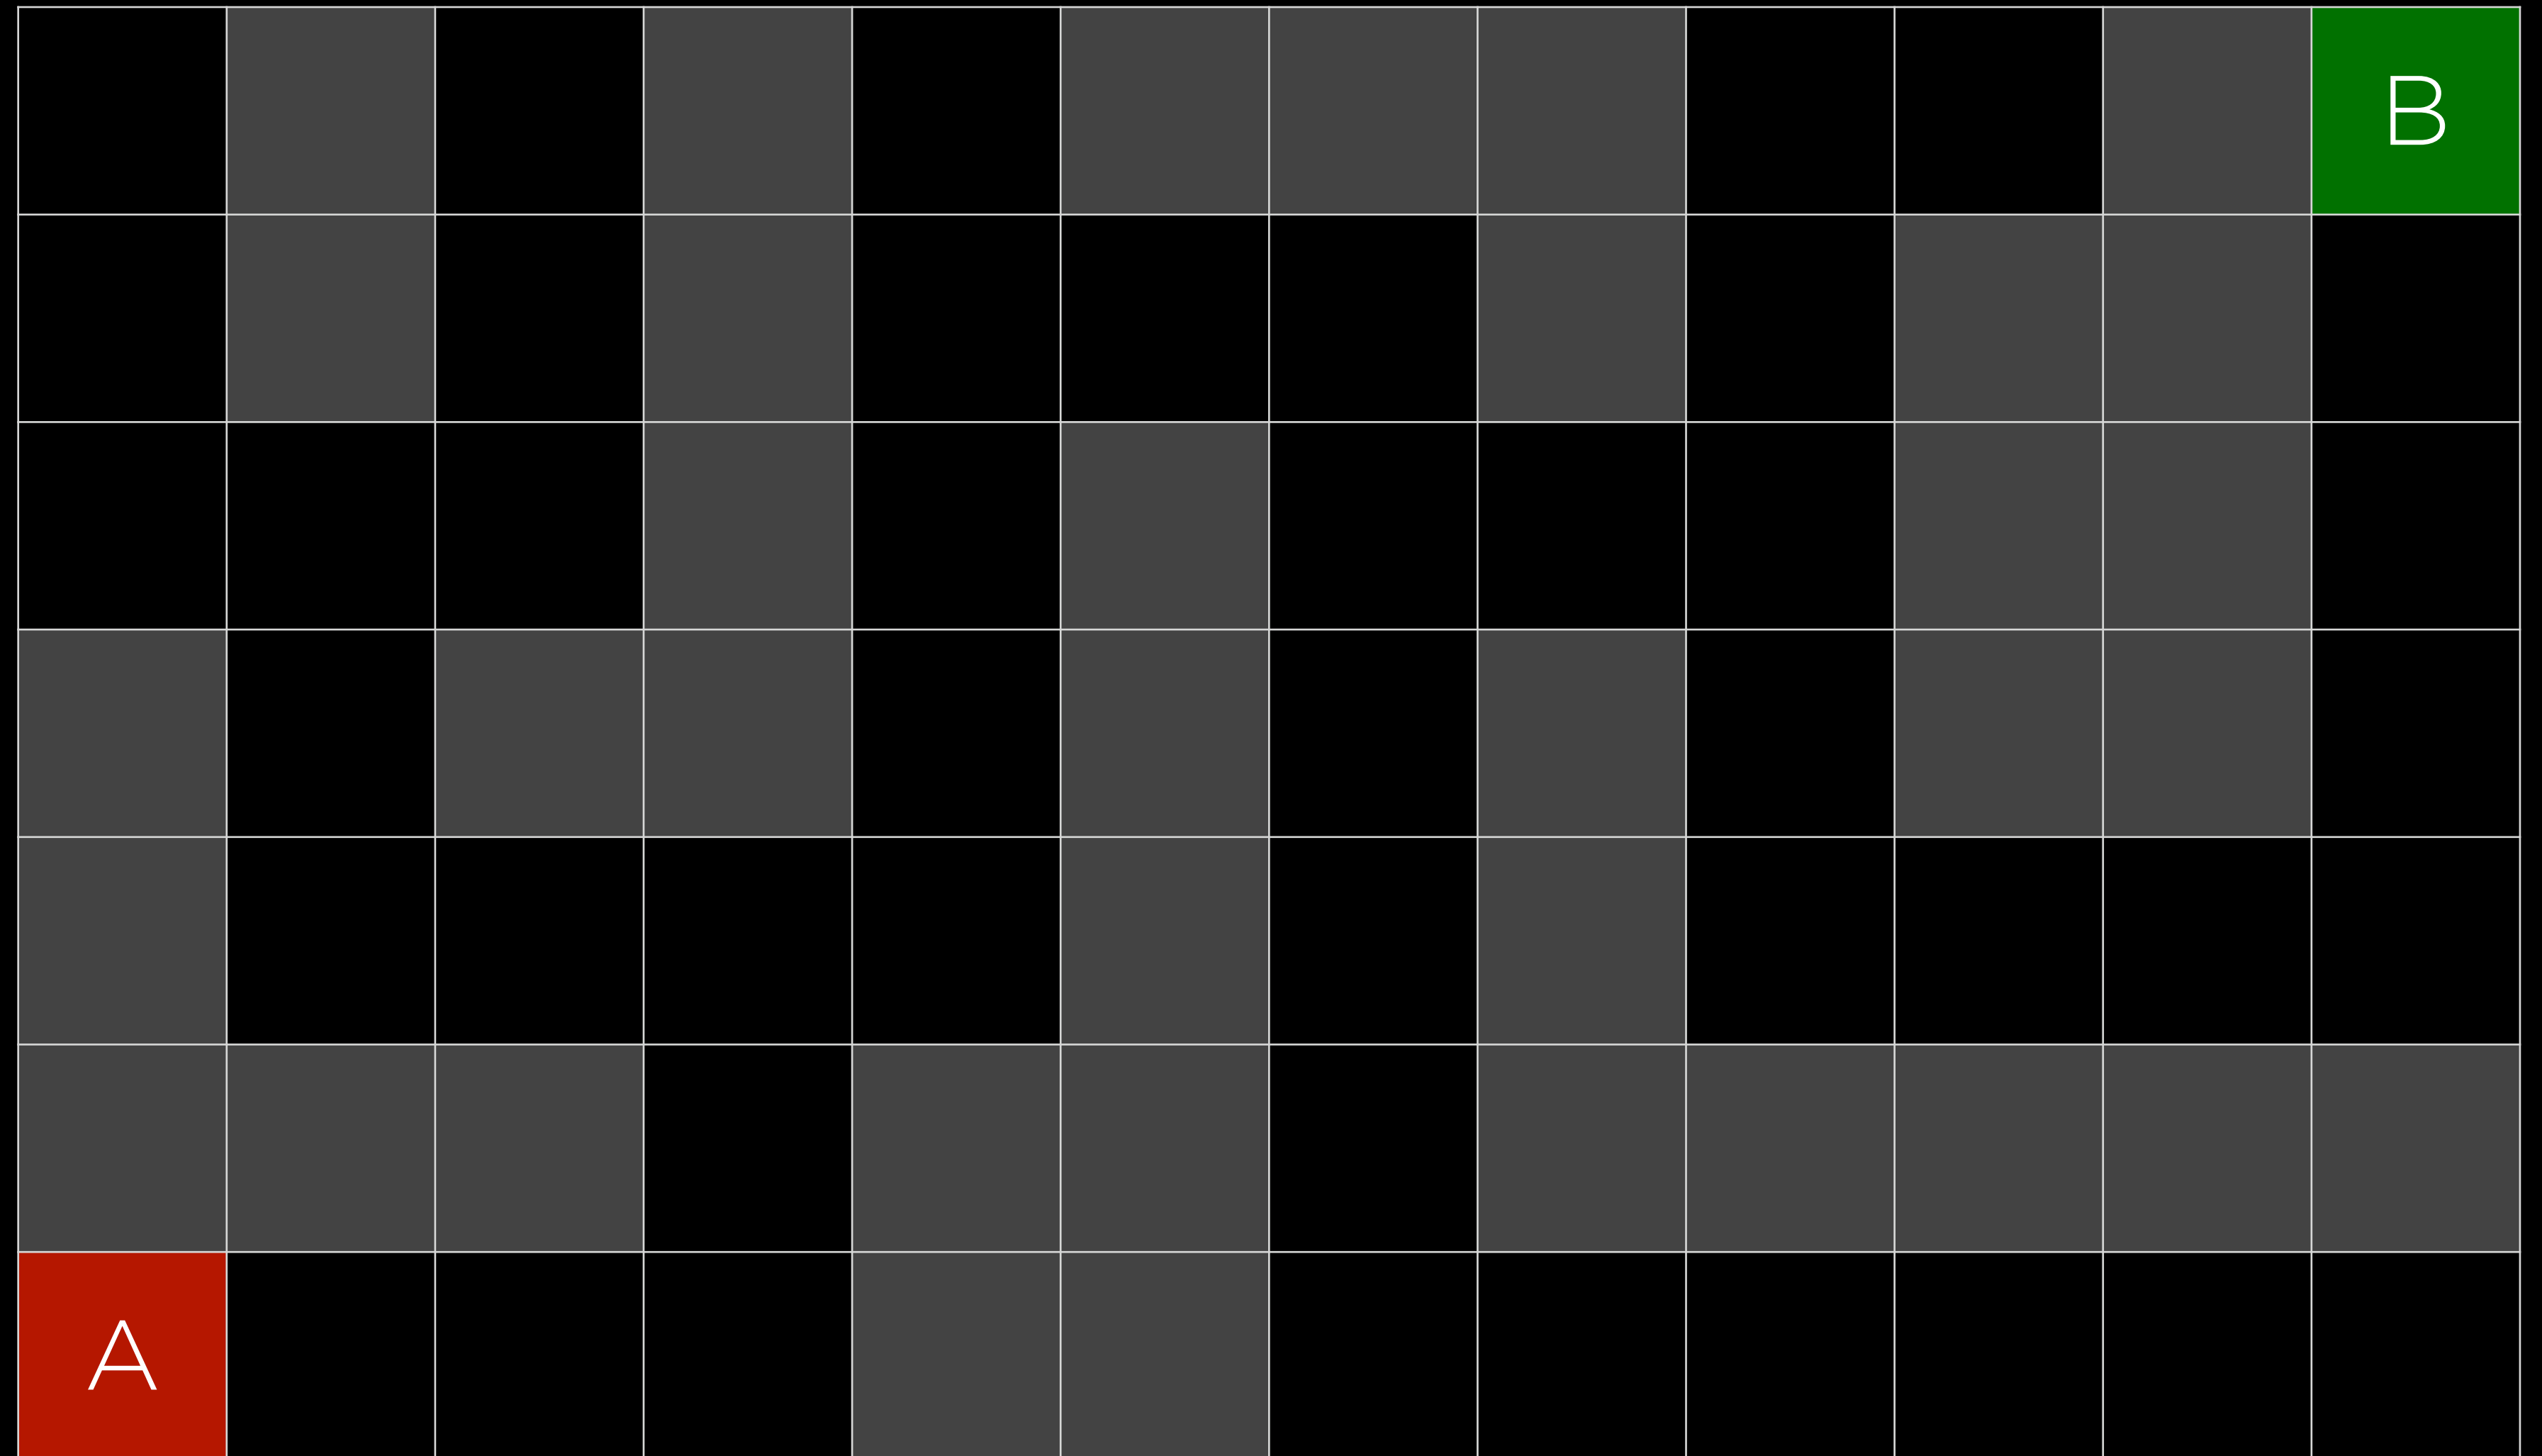


Breadth-First Search

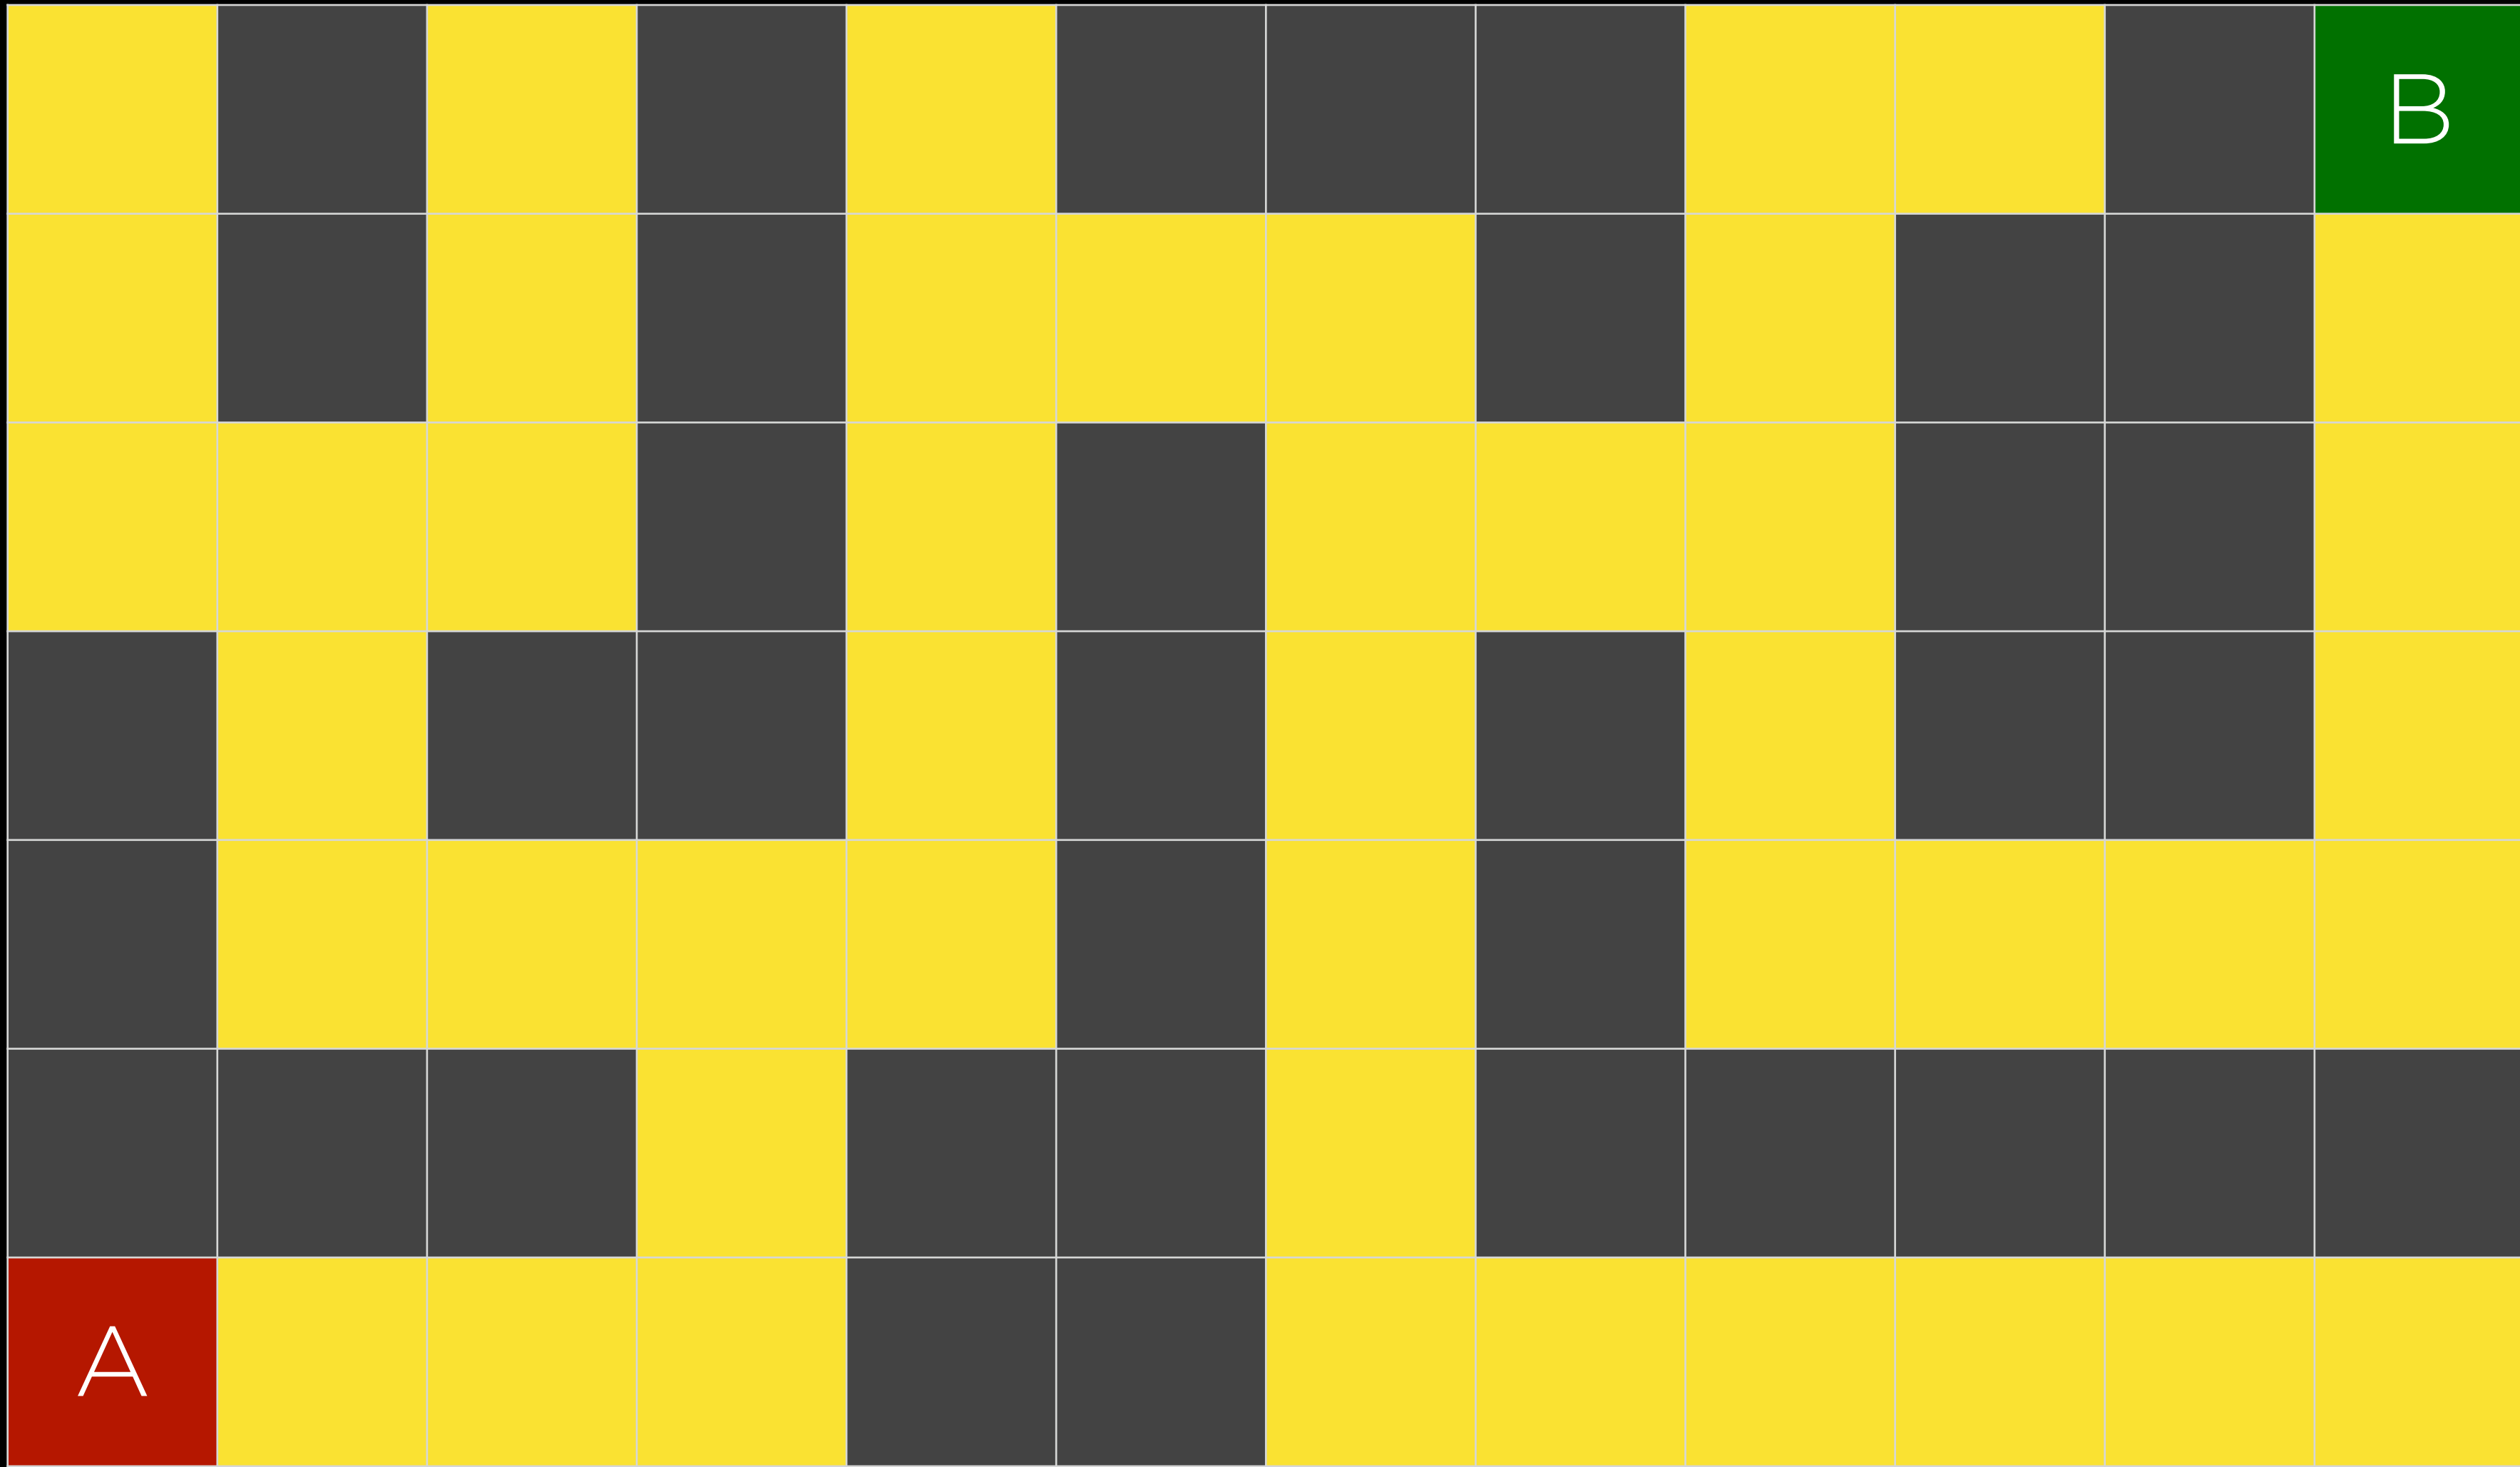
Breadth-First Search



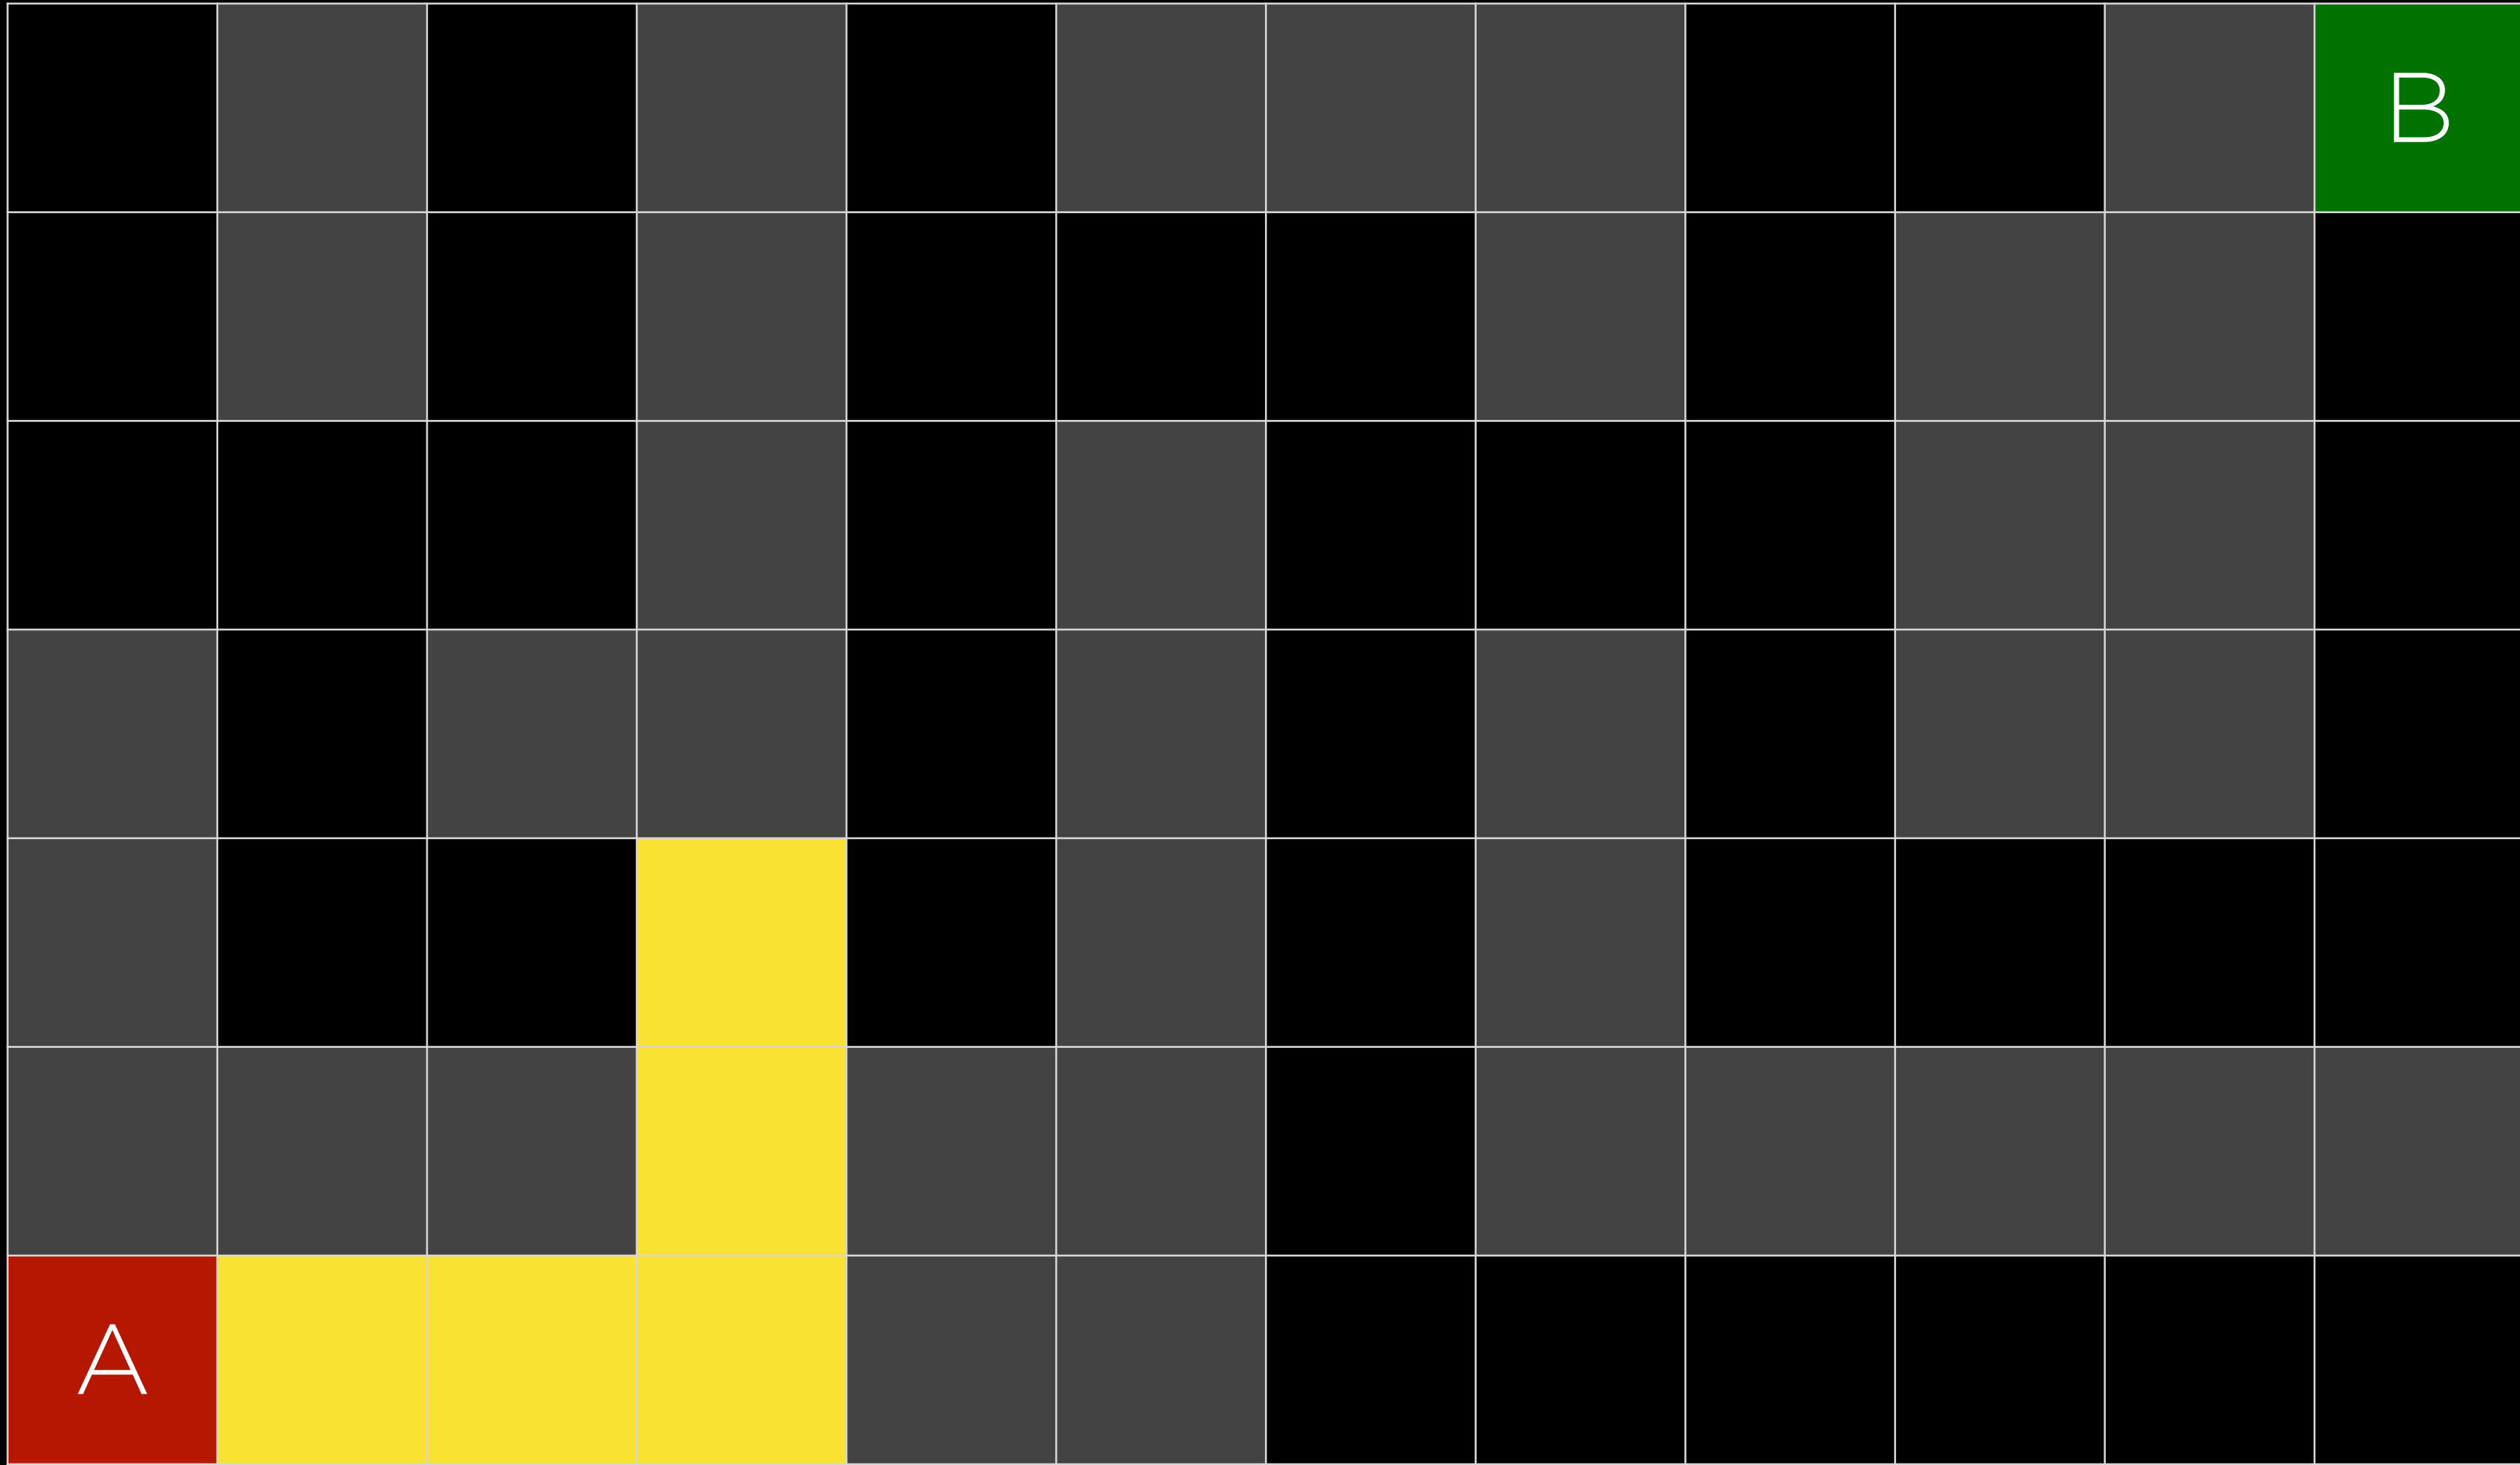
Breadth-First Search



Breadth-First Search



Breadth-First Search



uninformed search

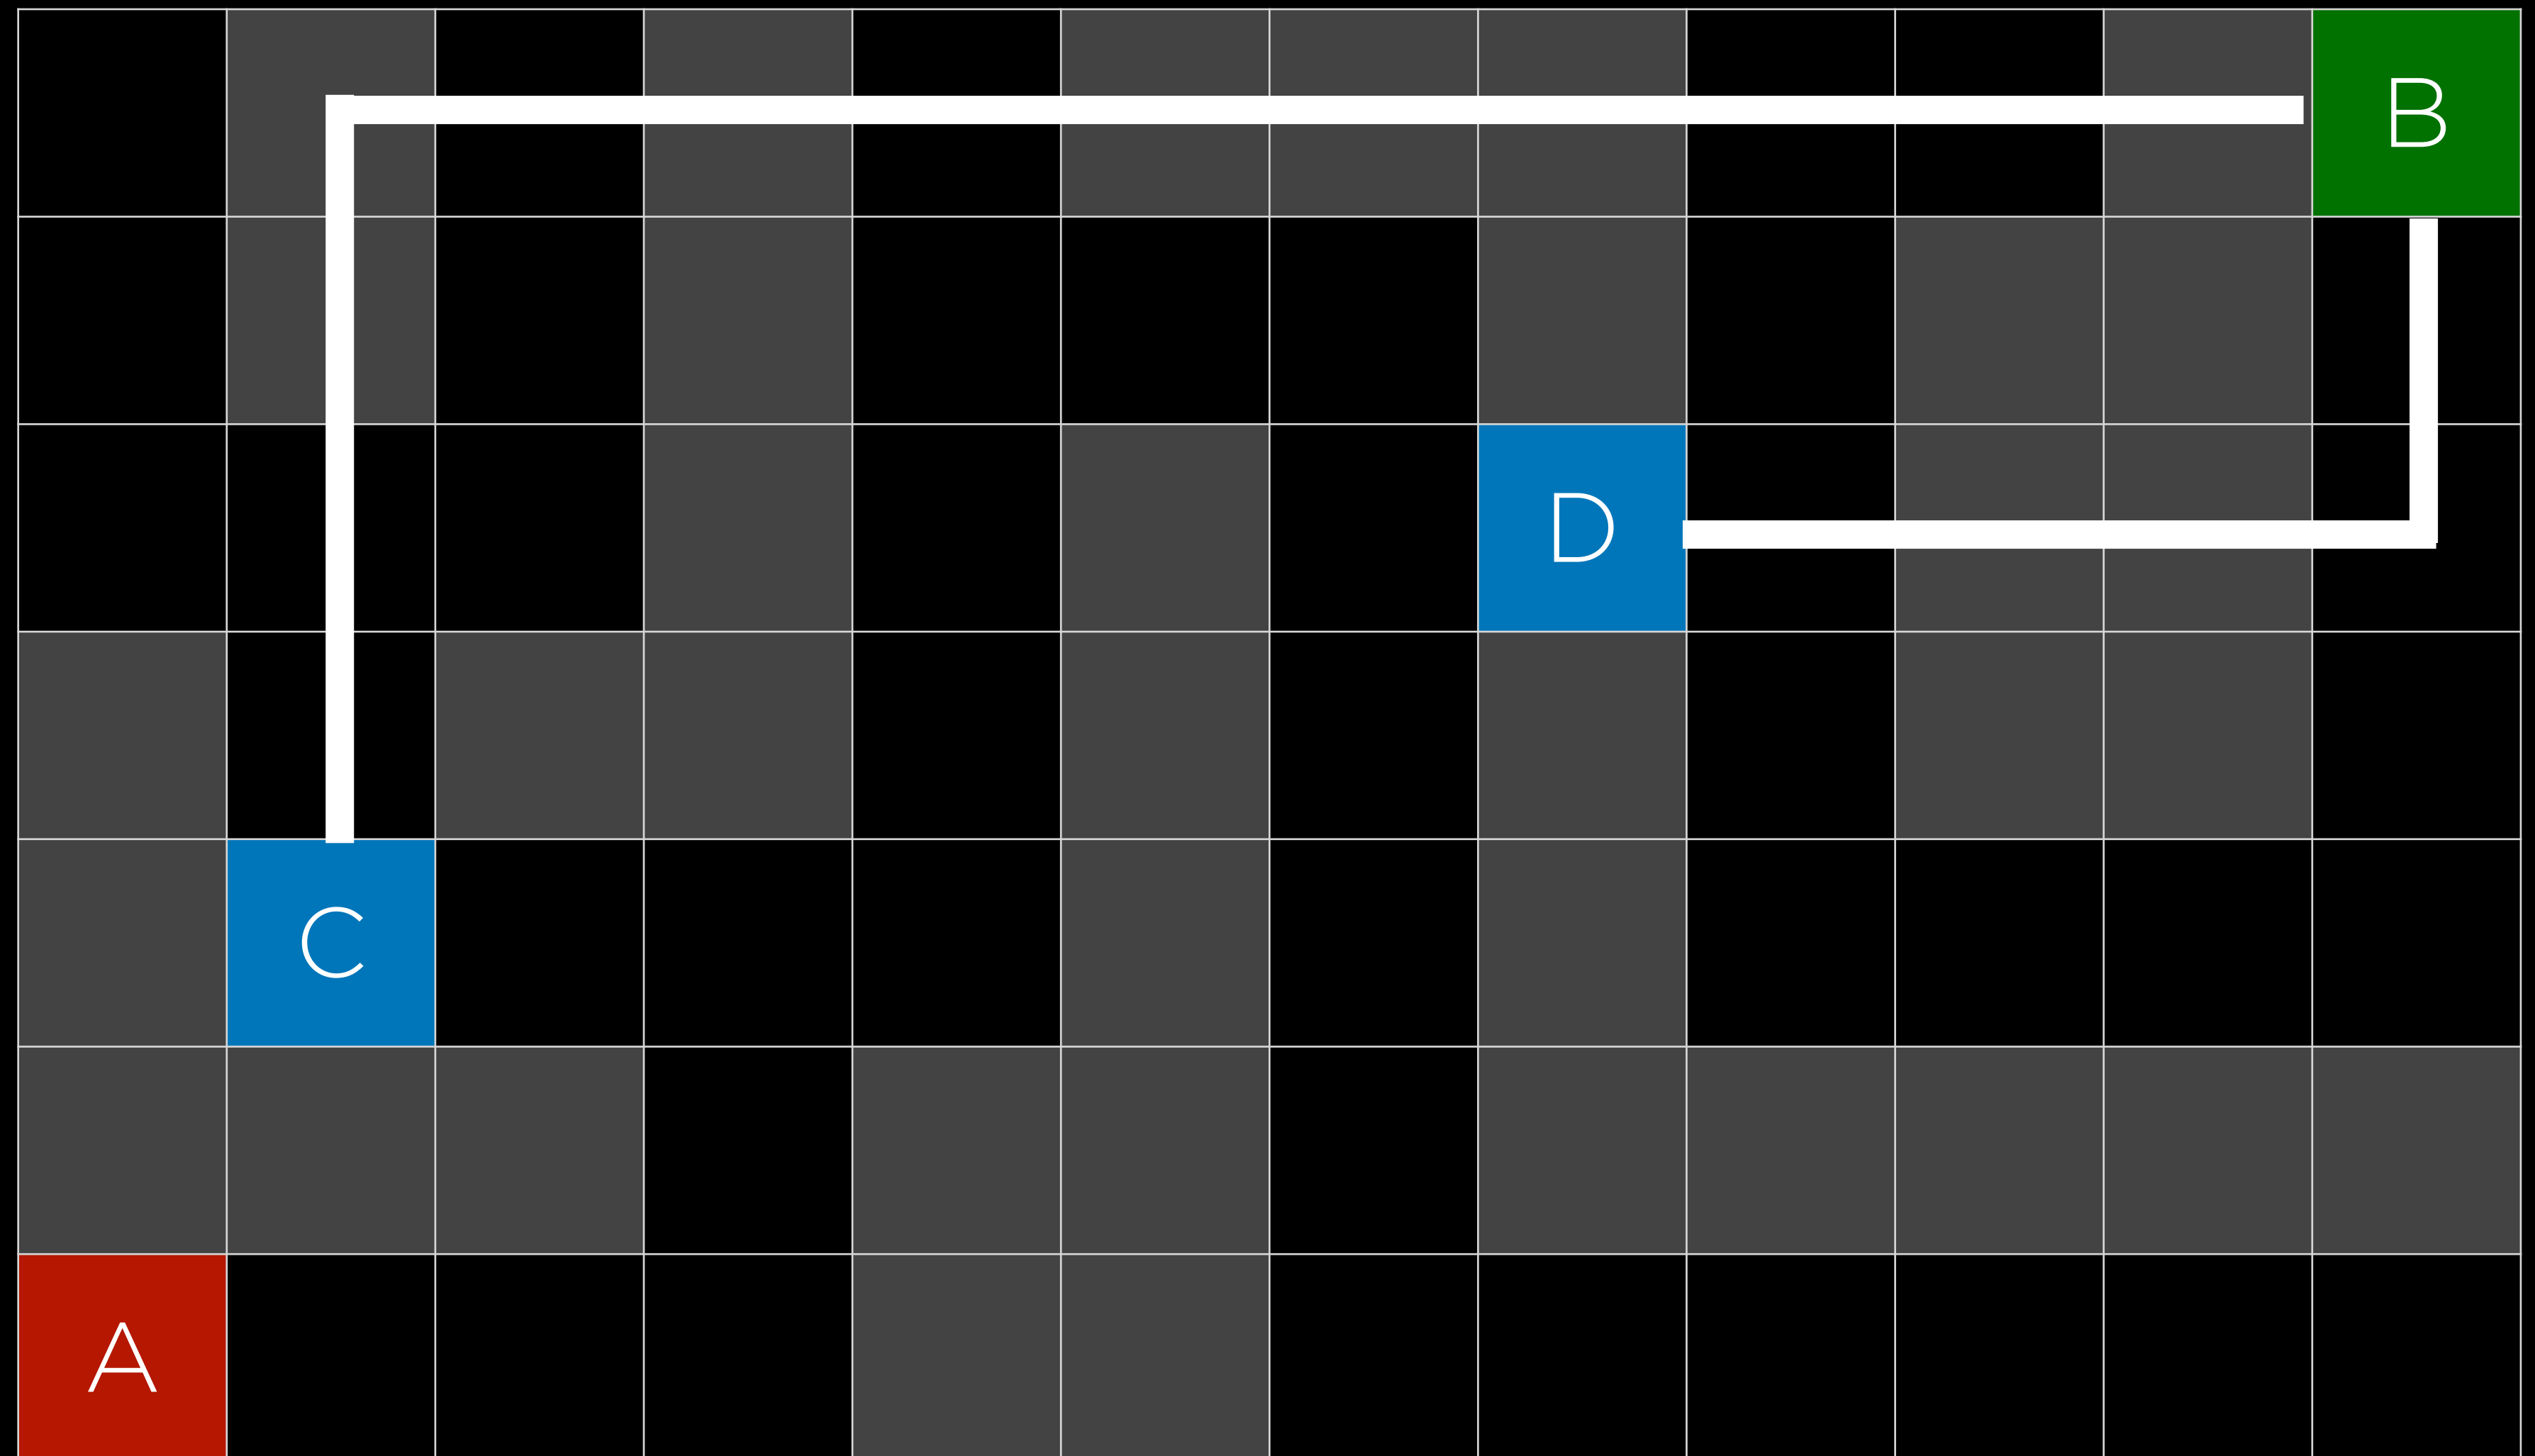
informed search

search strategy that uses problem-specific knowledge to find solutions more efficiently

greedy best-first search

search algorithm that chooses what to explore based on a heuristic

Heuristic function? Manhattan distance.



Greedy Best-First Search

11		9		7				3	2		B
12		10		8	7	6		4			1
13	12	11		9		7	6	5			2
	13			10		8		6			3
	14	13	12	11		9		7	6	5	4
			13			10					
A	16	15	14			11	10	9	8	7	6

Greedy Best-First Search

11		9		7				3	2		B
12		10		8	7	6		4			1
13	12	11		9		7	6	5			2
	13			10		8		6			3
	14	13	12	11		9		7	6	5	4
			13			10					
A	16	15	14			11	10	9	8	7	6

Greedy Best-First Search

	10	9	8	7	6	5	4	3	2	1	B
	11										1
	12		10	9	8	7	6	5	4		2
	13		11						5		3
	14	13	12		10	9	8	7	6		4
			13		11						5
A	16	15	14		12	11	10	9	8	7	6

Greedy Best-First Search

	10	9	8	7	6	5	4	3	2	1	B
	11										1
	12		10	9	8	7	6	5	4		2
	13		11						5		3
	14	13	12		10	9	8	7	6		4
			13		11						5
A	16	15	14		12	11	10	9	8	7	6

Greedy Best-First Search

	10	9	8	7	6	5	4	3	2	1	B
	11										1
	12		10	9	8	7	6	5	4		2
	13		11						5		3
	14	13	12		10	9	8	7	6		4
			13		11						5
A	16	15	14		12	11	10	9	8	7	6

A* search

search algorithm that chooses what to explore based both on cost so far and a heuristic

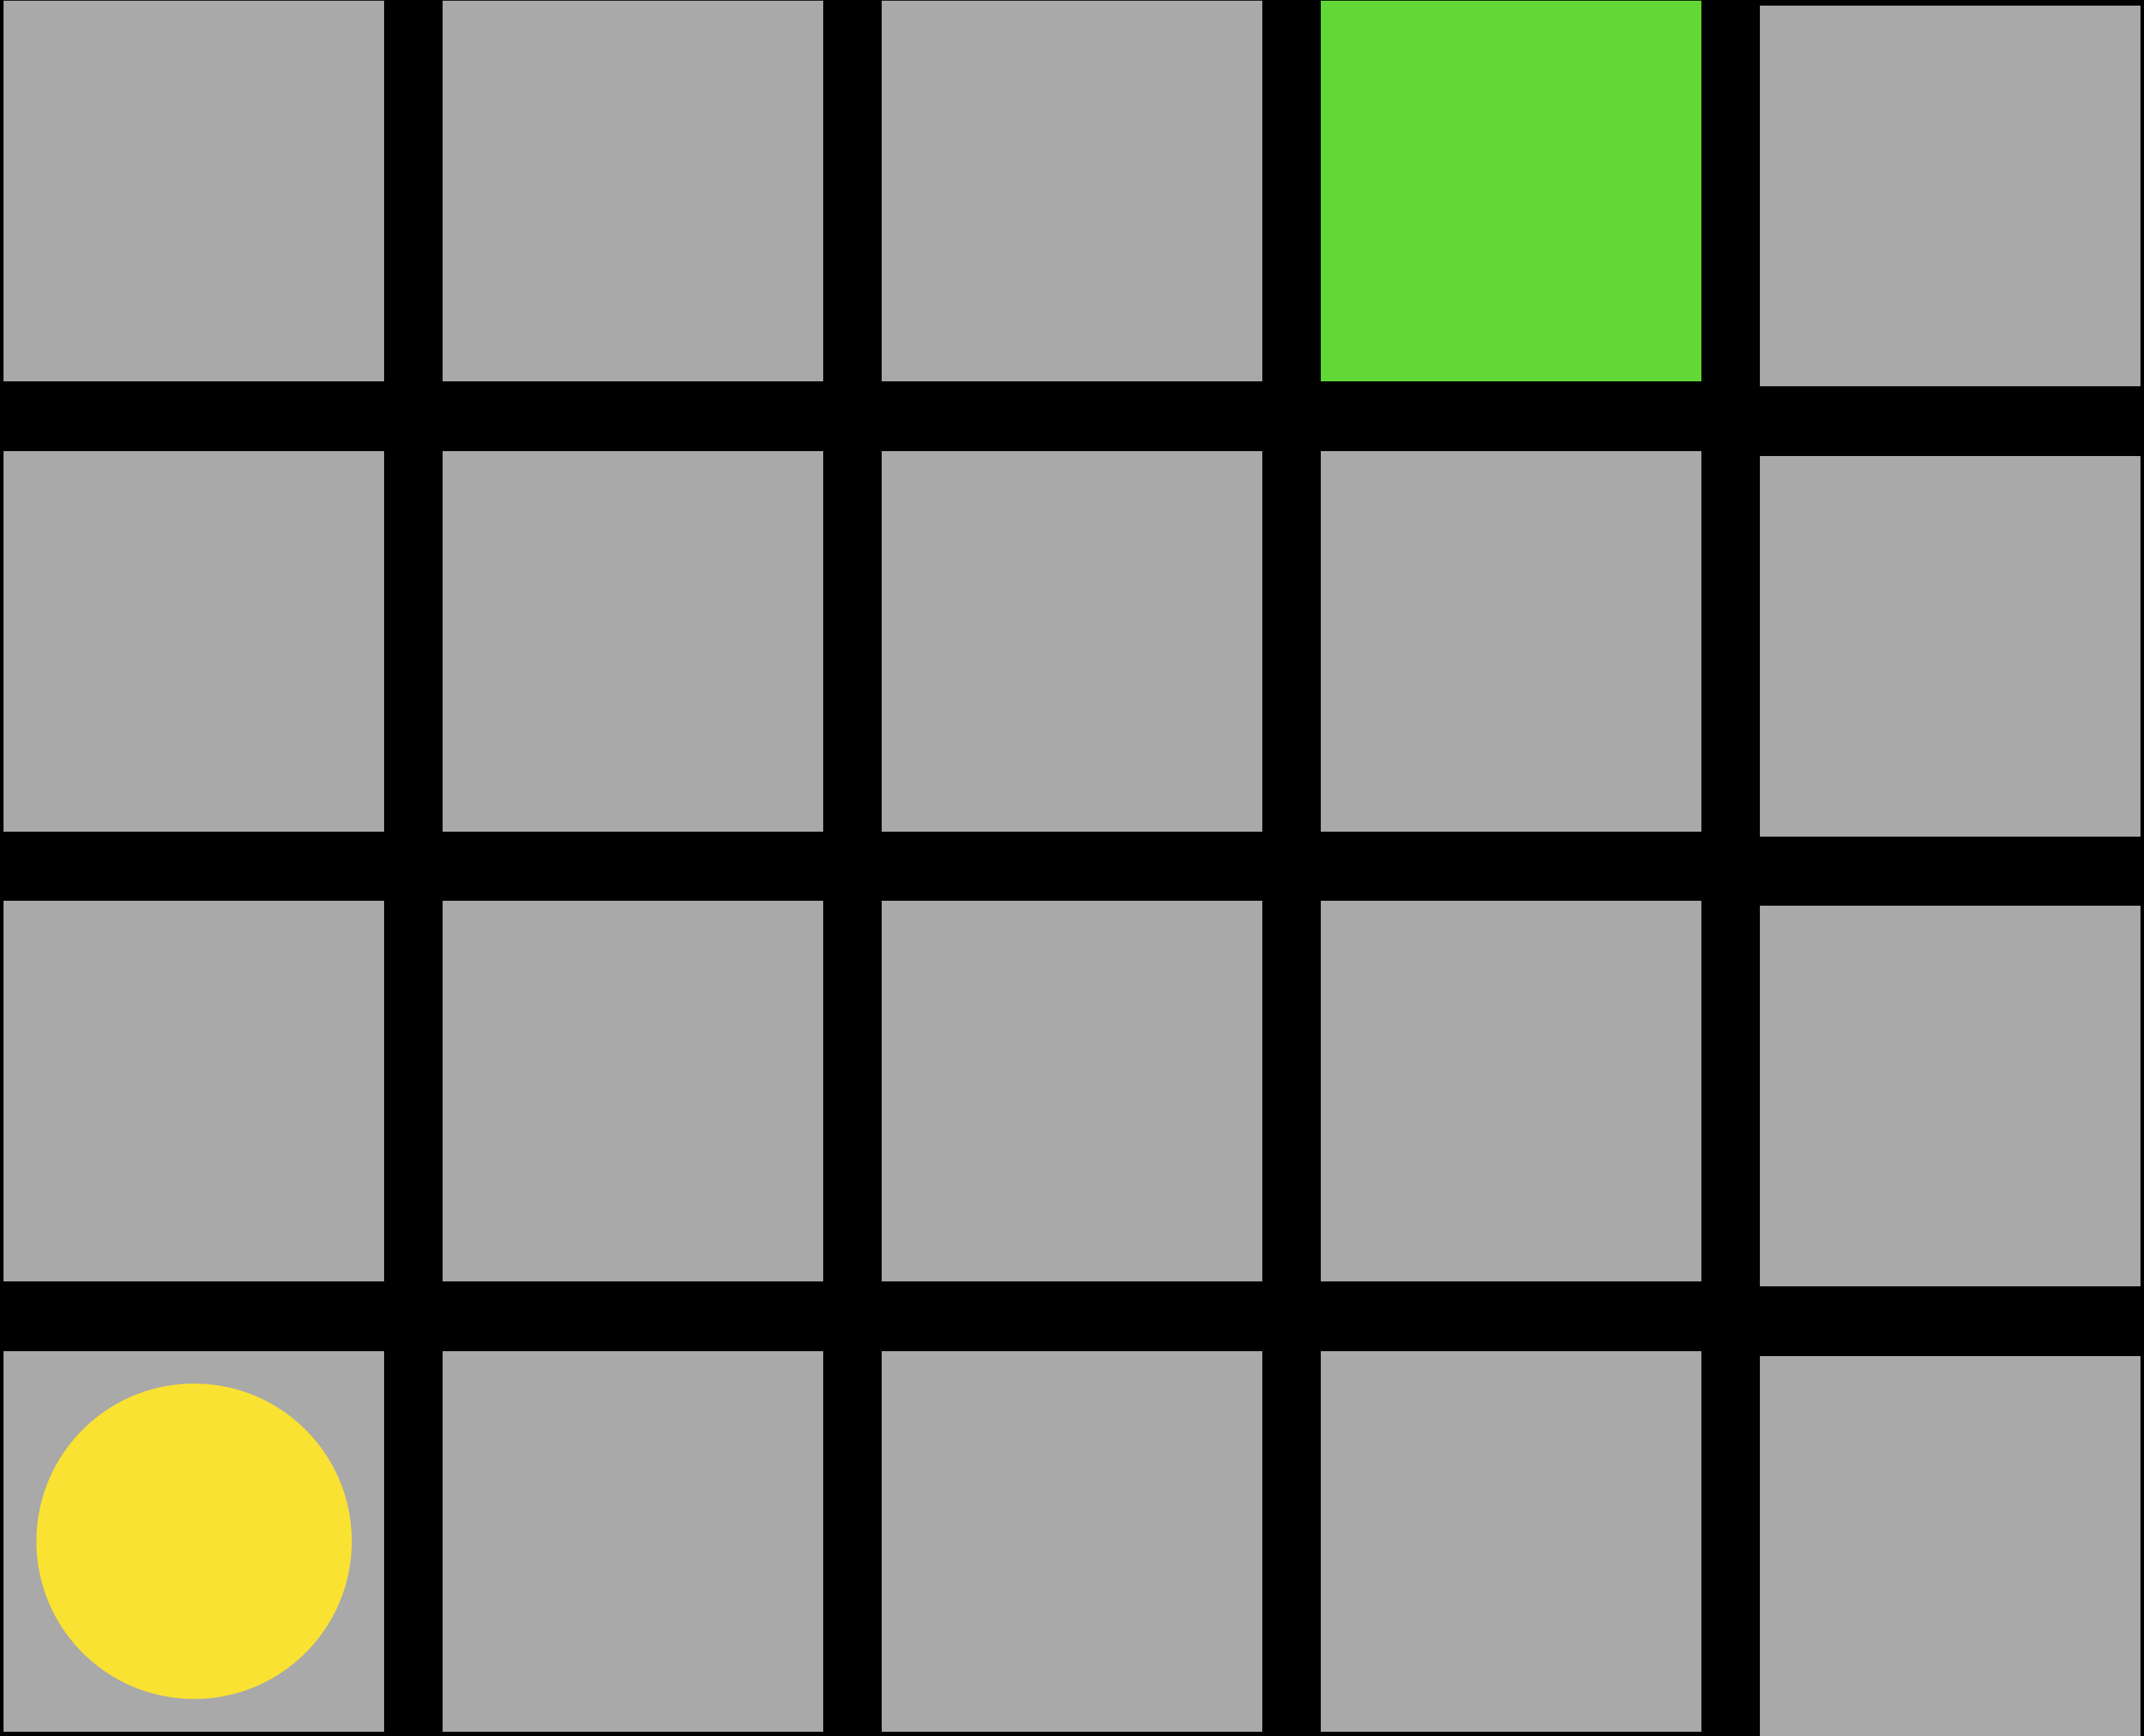
A* Search

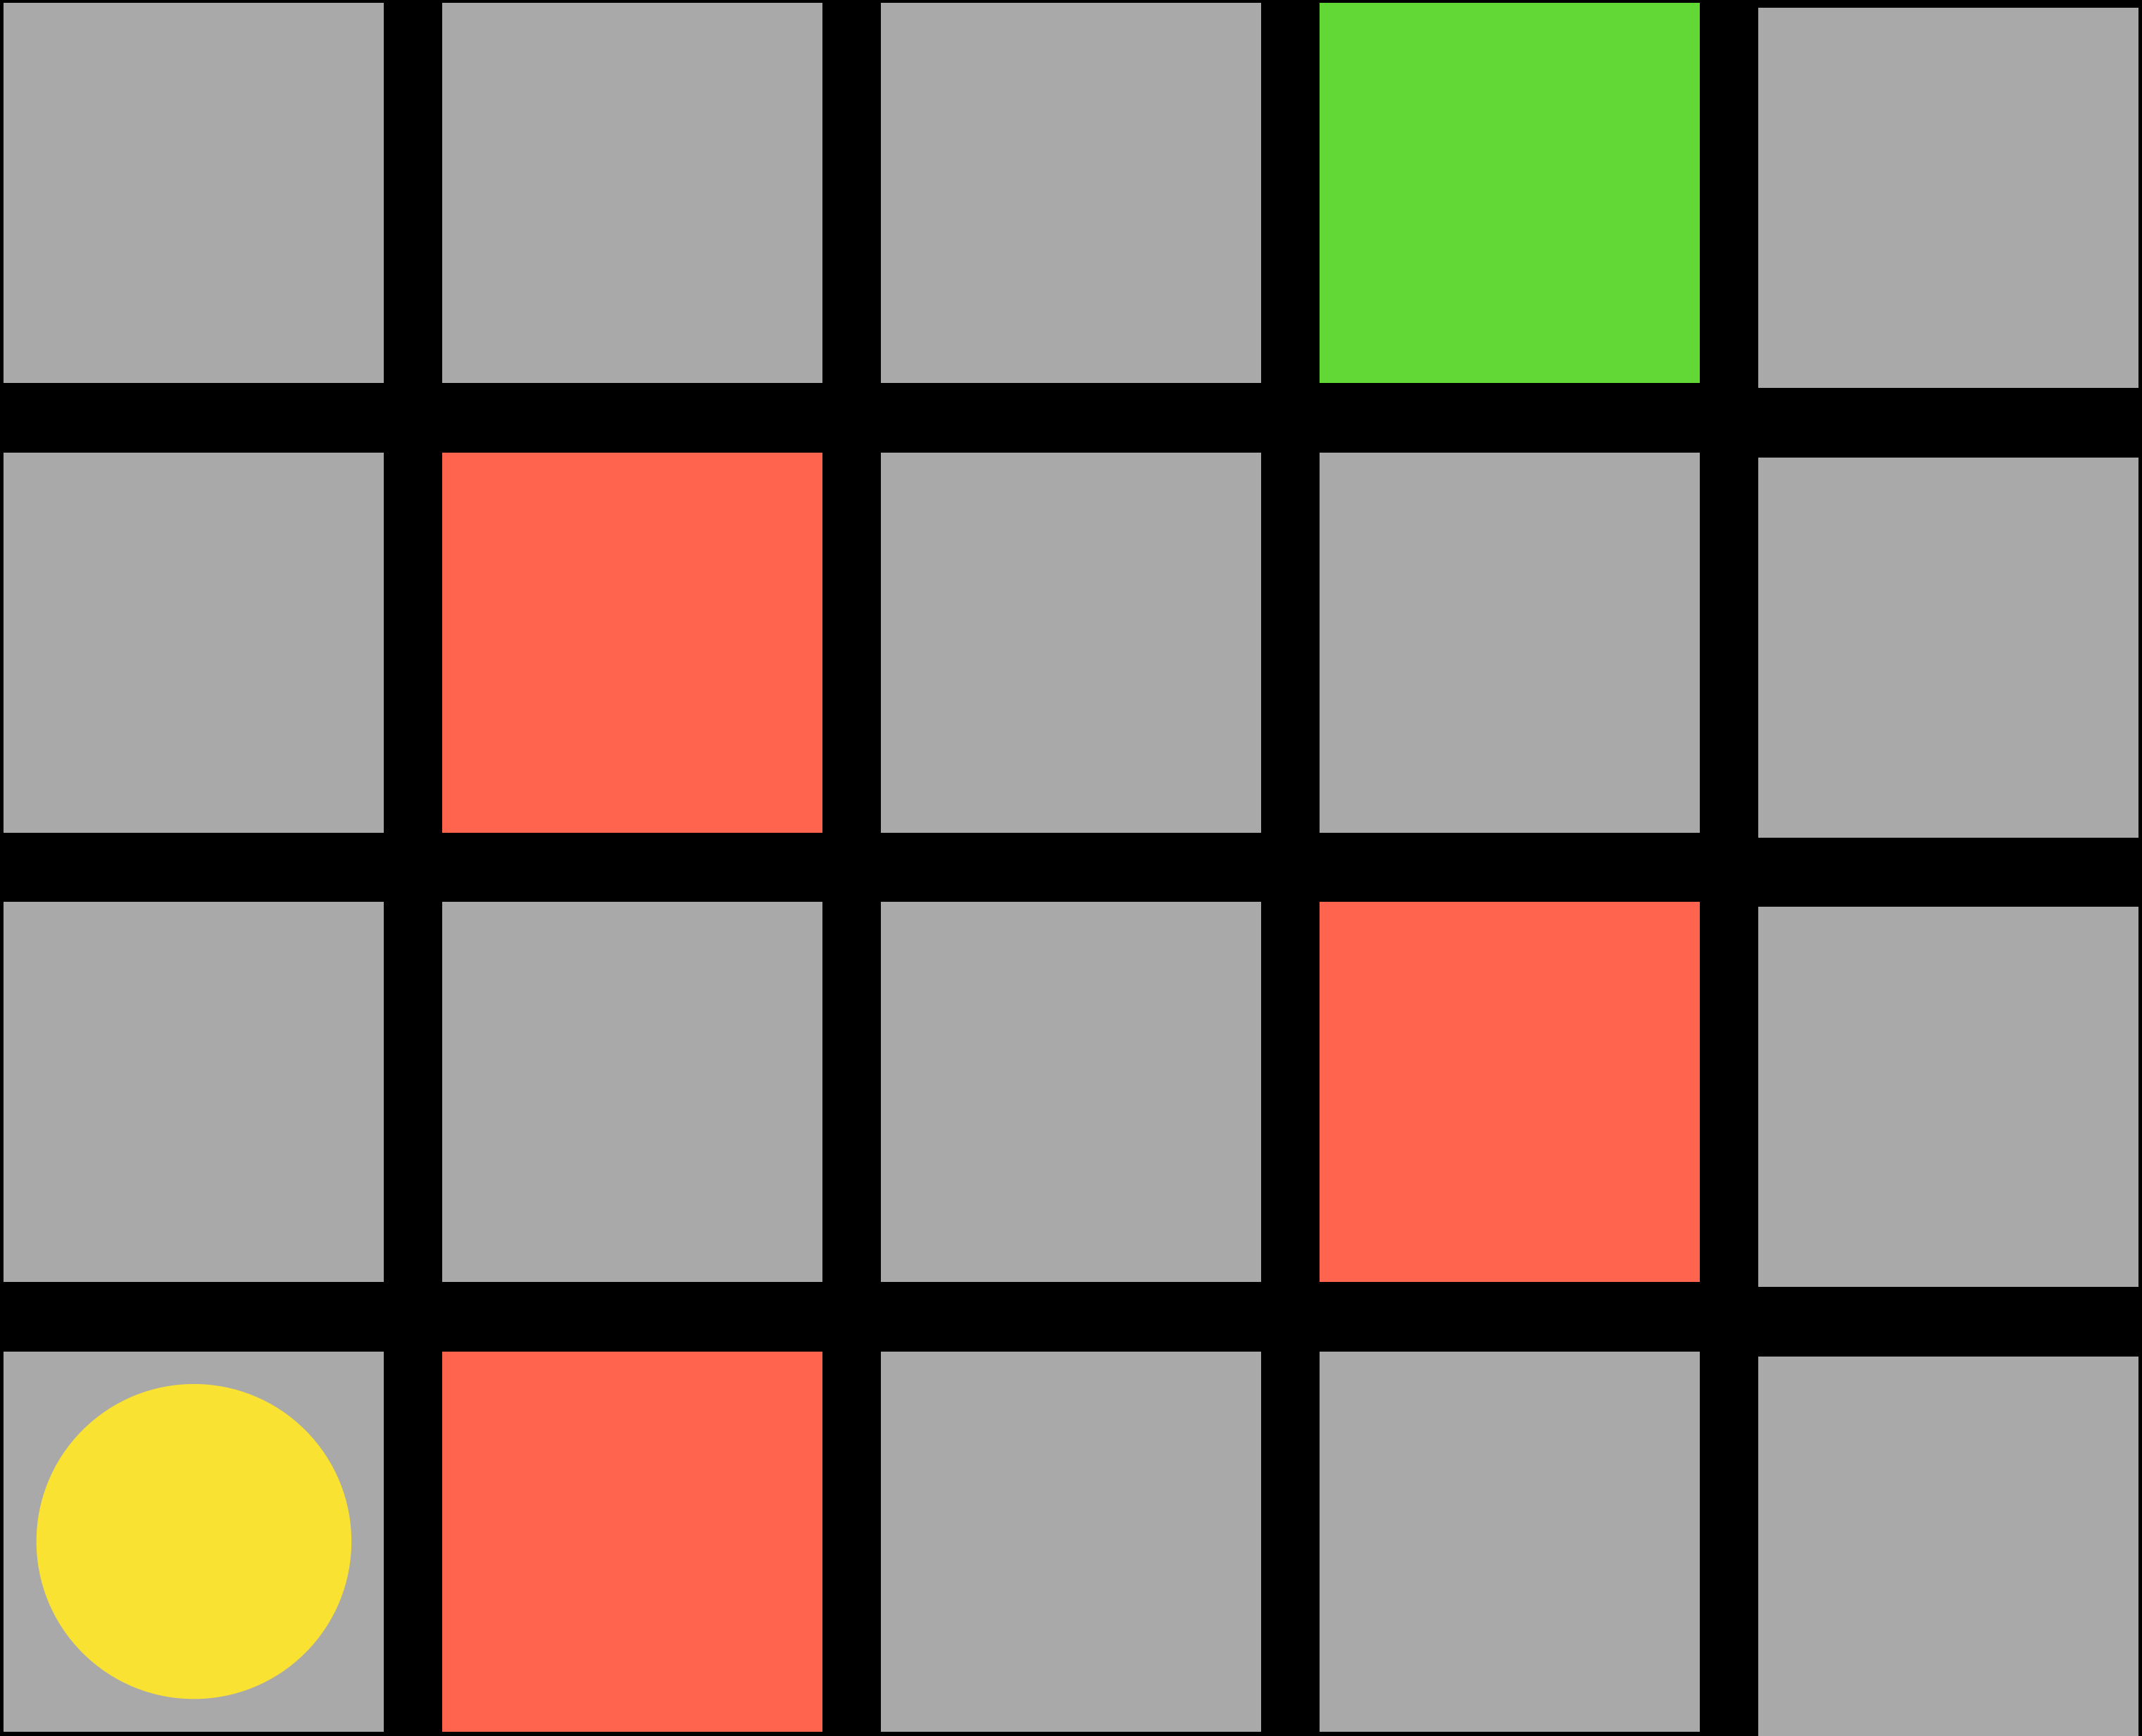
	10	9	8	7	6	5	4	3	2	1	B
	11										1
	12		10	9	8	7	6	5	4		2
	13		11						5		3
	14	13	12		10	9	8	7	6		4
			13		11						5
A	16	15	14		12	11	10	9	8	7	6

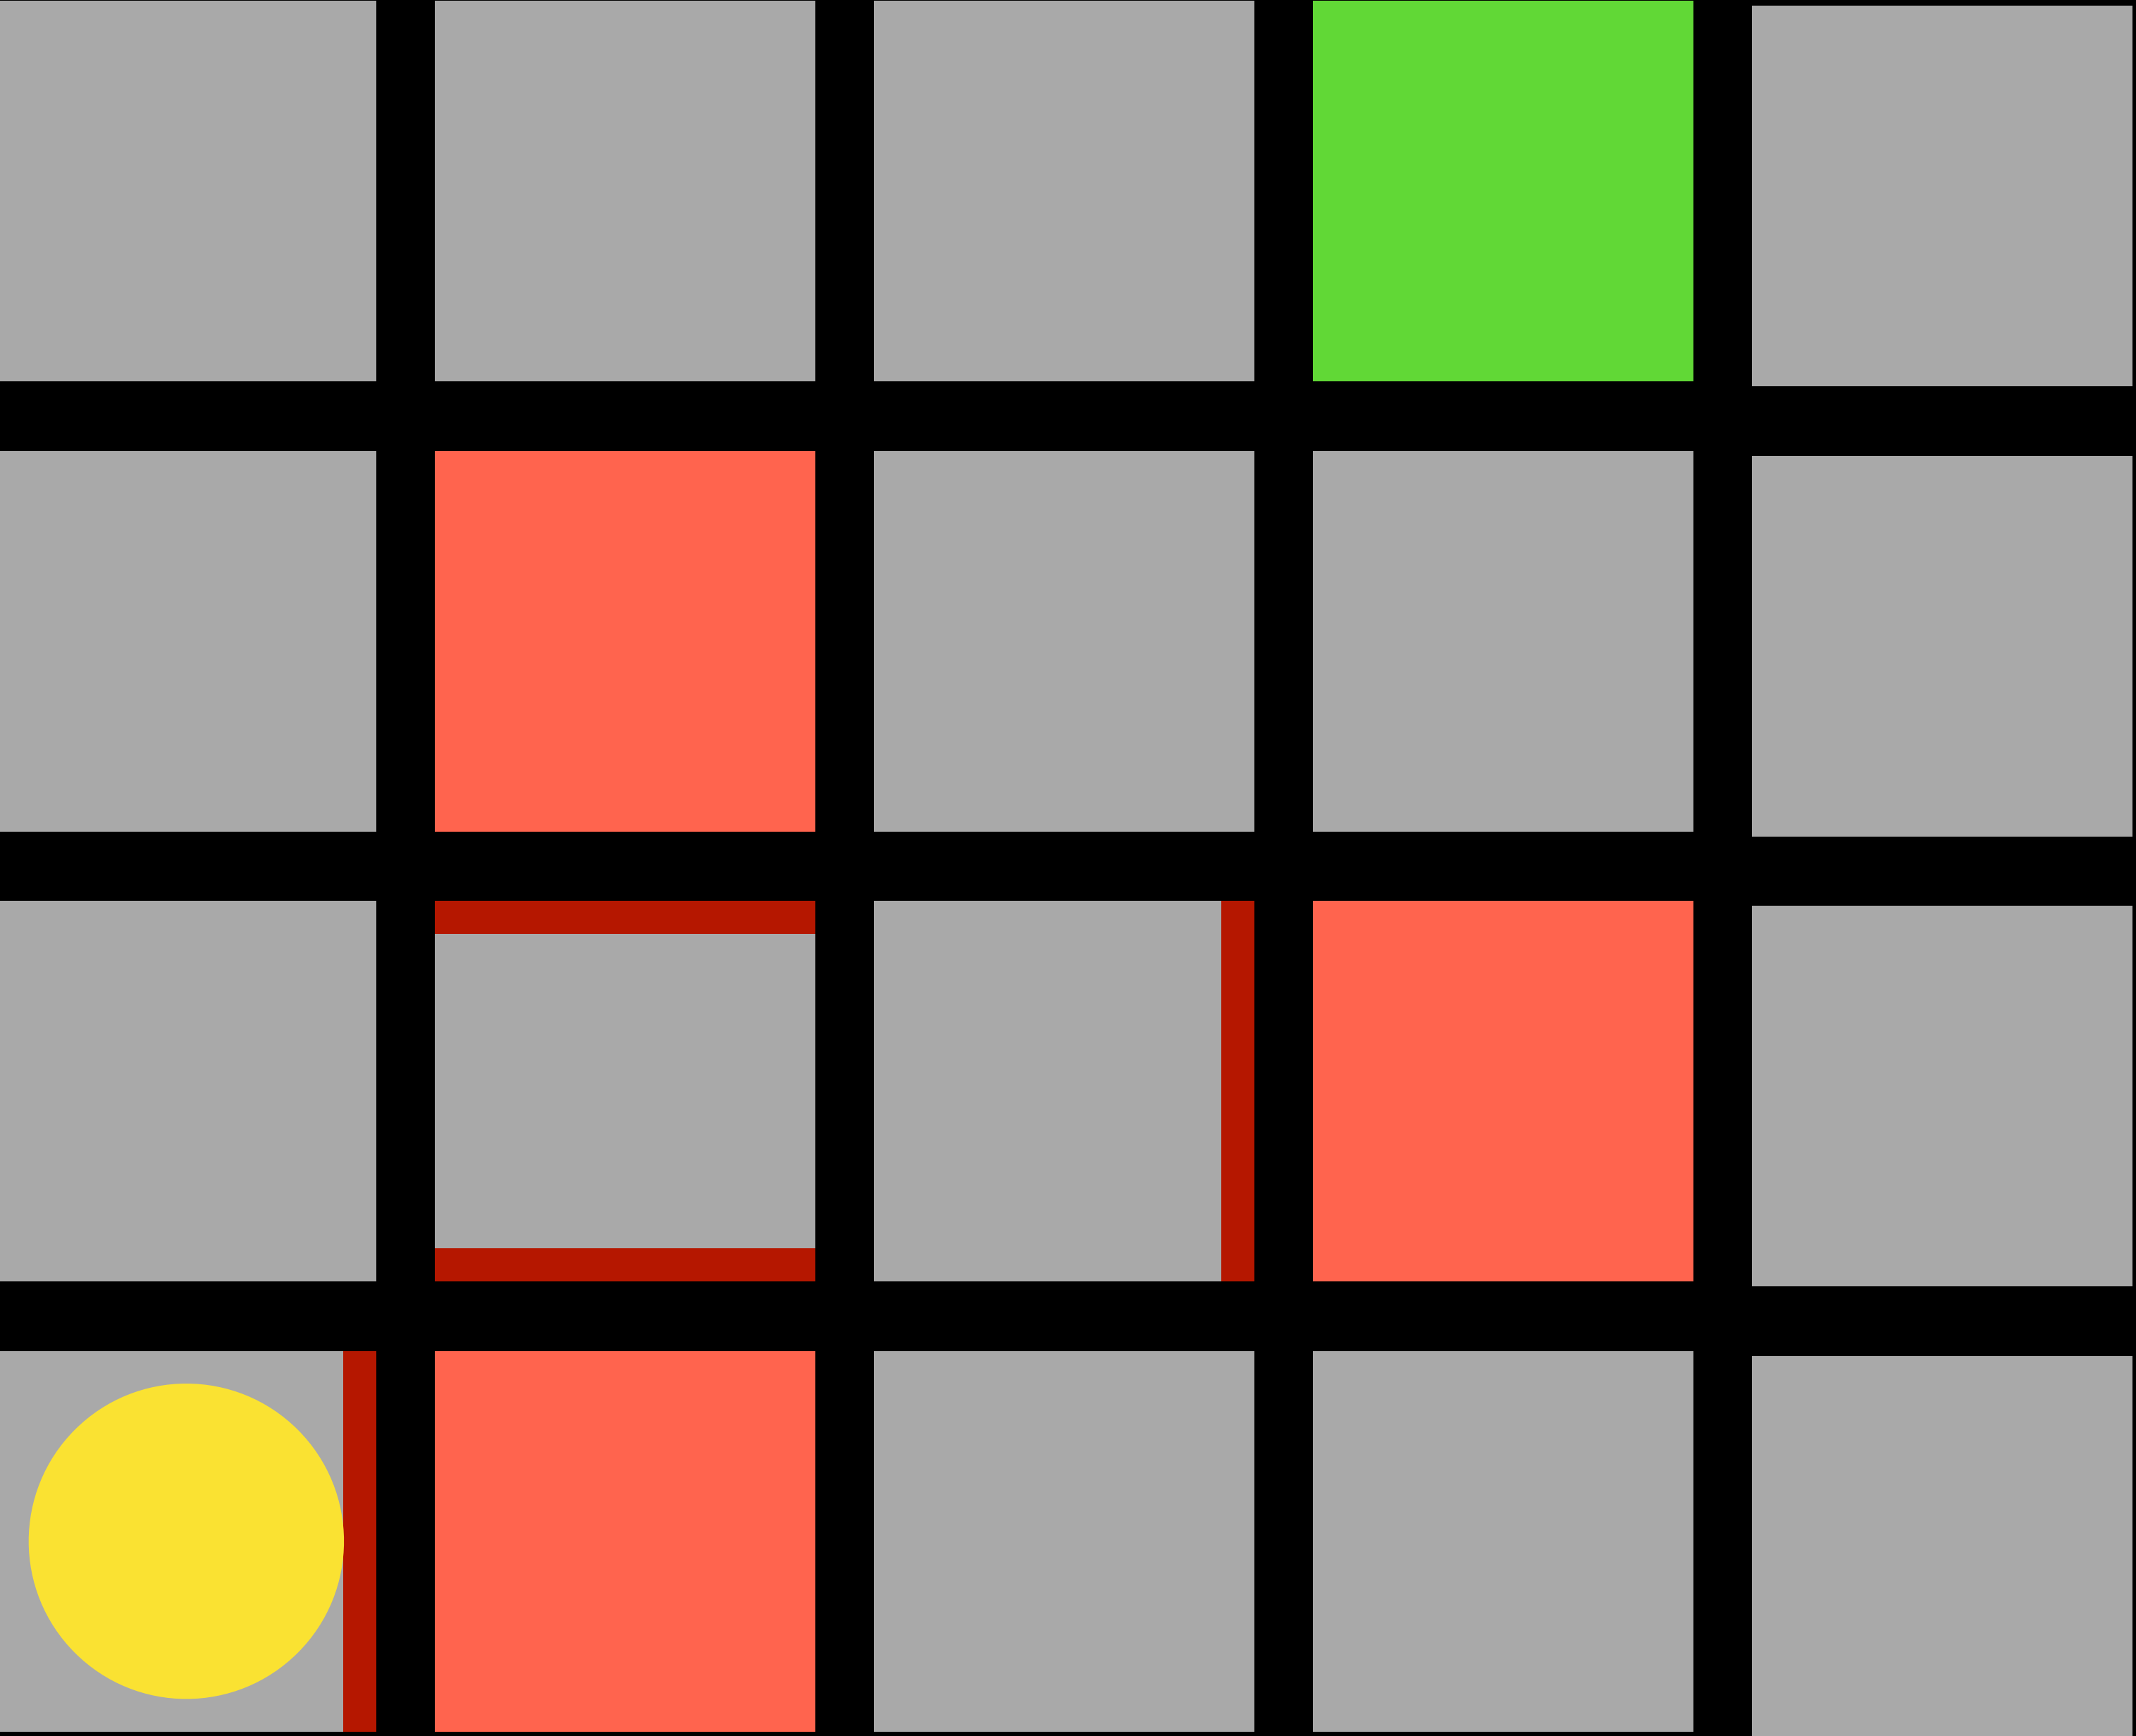
A* Search

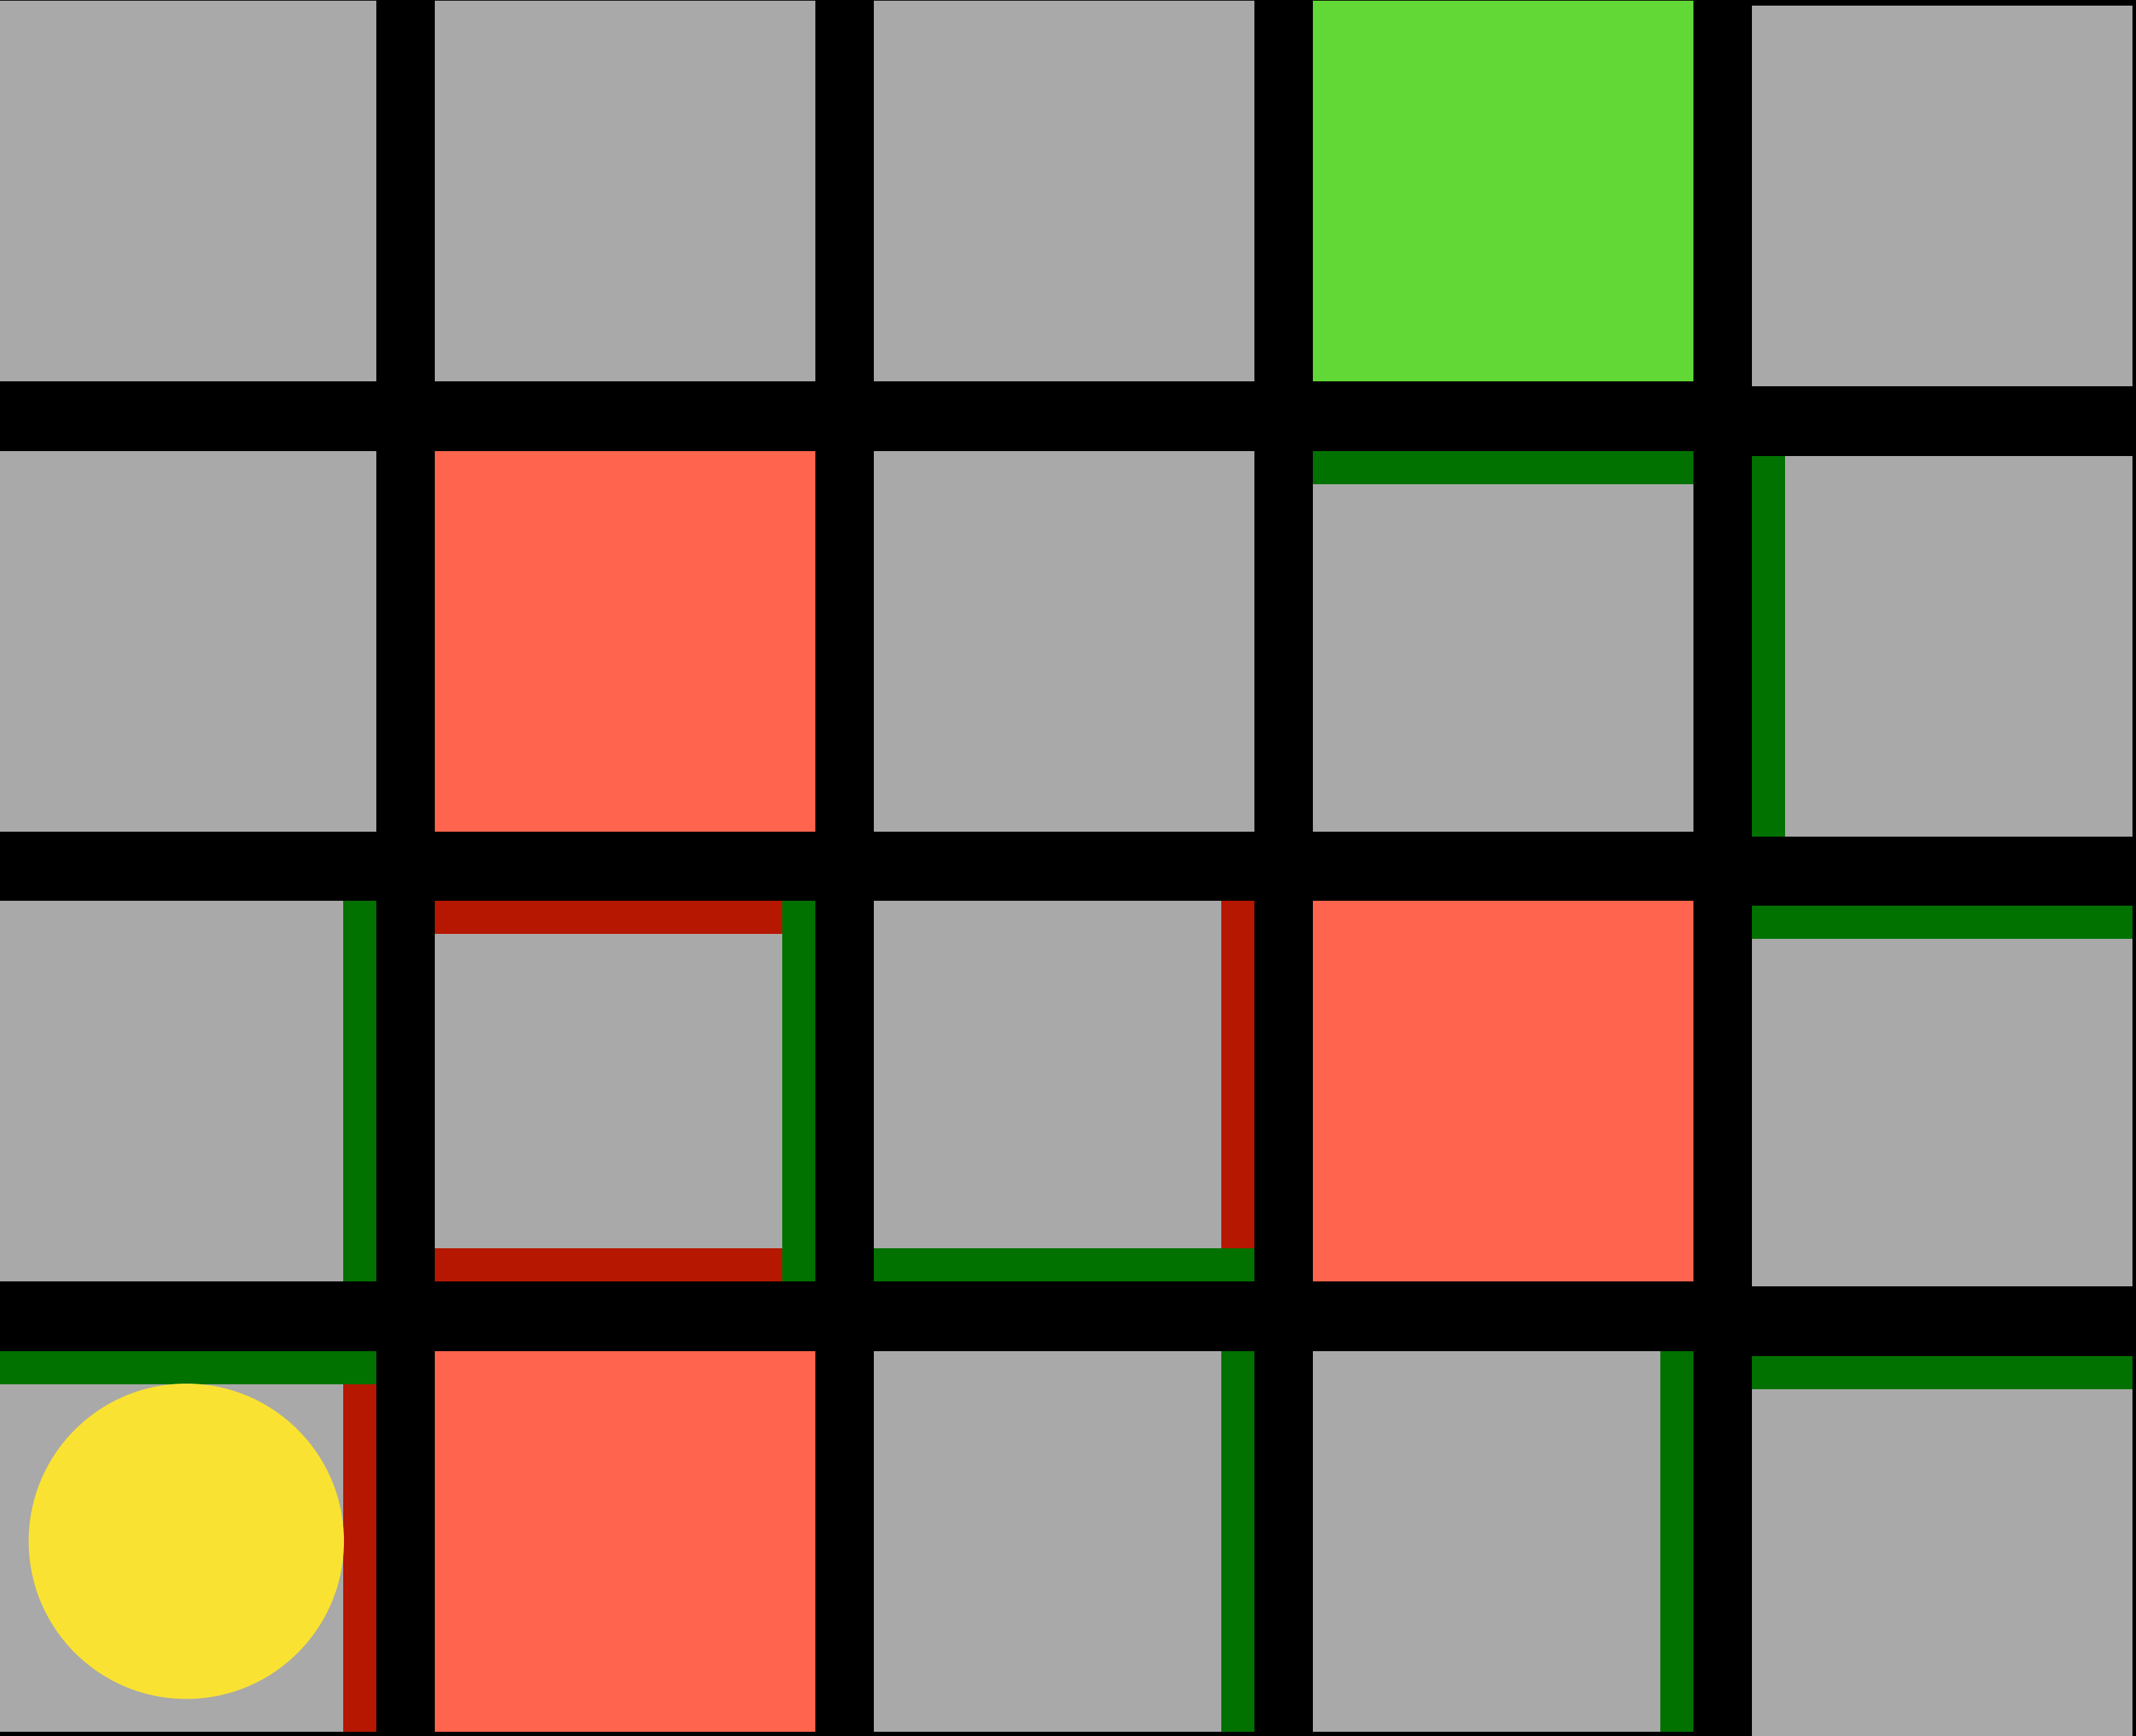
	11+10	12+9	13+8	14+7	15+6	16+5	17+4	18+3	19+2	20+1	B
	10+11										1
	9+12		7+10	8+9	9+8	10+7	11+6	12+5	13+4		2
	8+13		6+11						14+5		3
	7+14	6+13	5+12		10	9	8	7	15+6		4
			4+13		11						5
A	1+16	2+15	3+14		12	11	10	9	8	7	6

Reinforcement Learning









Explore vs. Exploit

Explore vs. Exploit Strategy

```
epsilon = 0.10
```

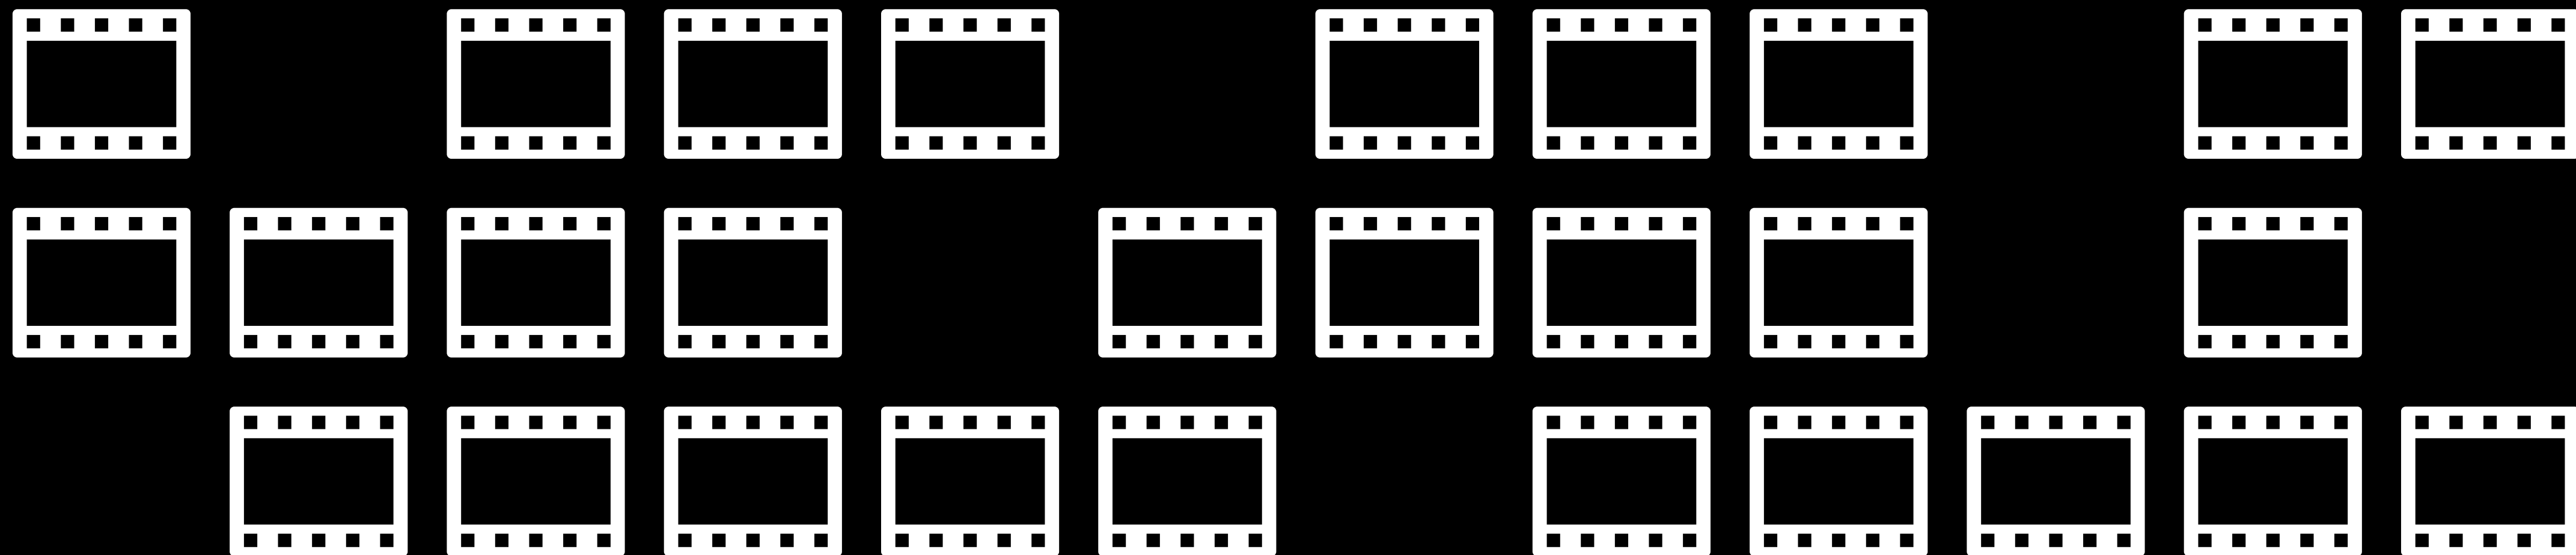
```
if random() < epsilon:  
    make a random move
```

```
else:  
    make the move with the highest value
```

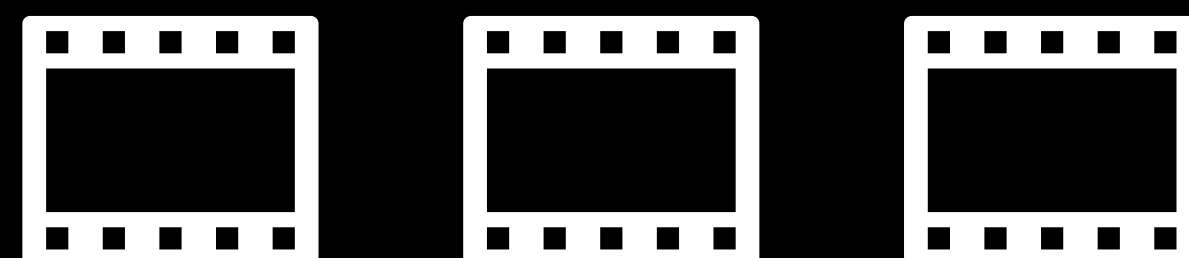
Robot Motor Skill Coordination with EM-based Reinforcement Learning

**Petar Kormushev, Sylvain Calinon,
and Darwin G. Caldwell**

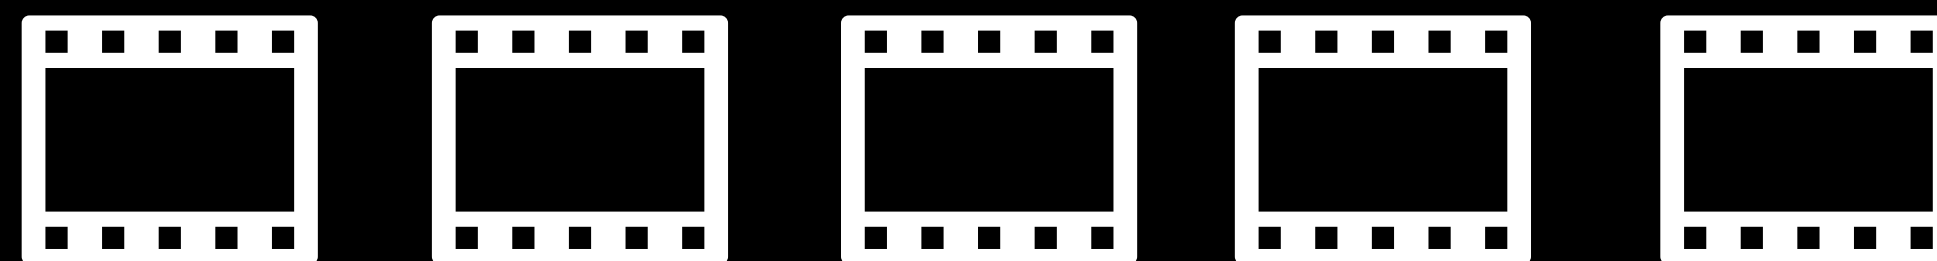
Italian Institute of Technology

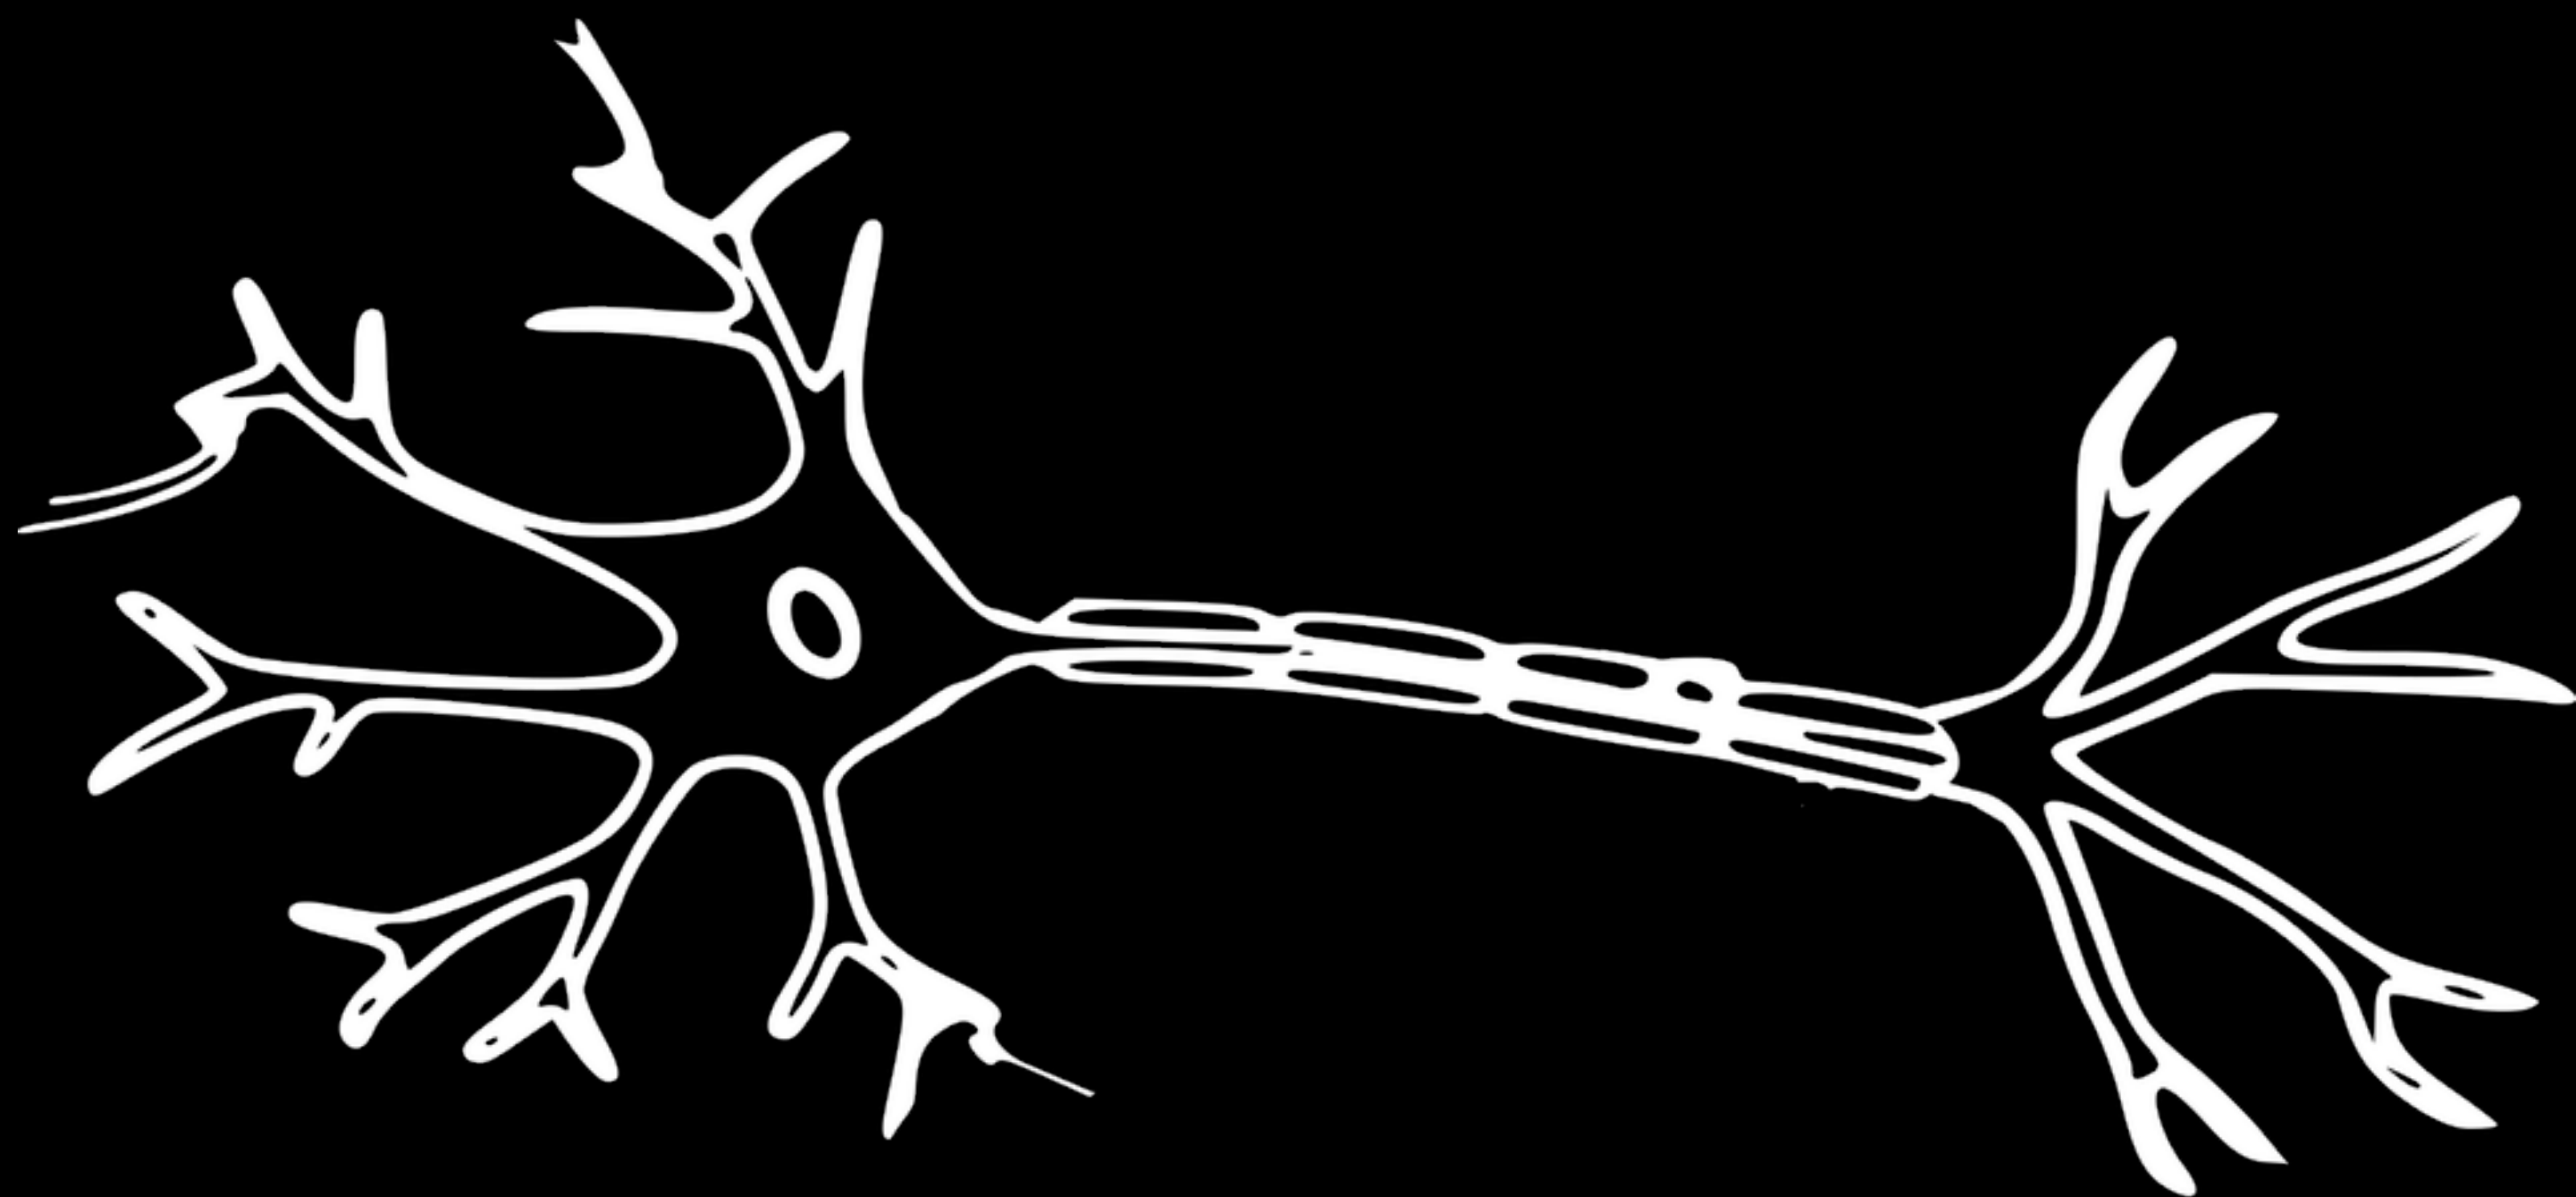


Watch History



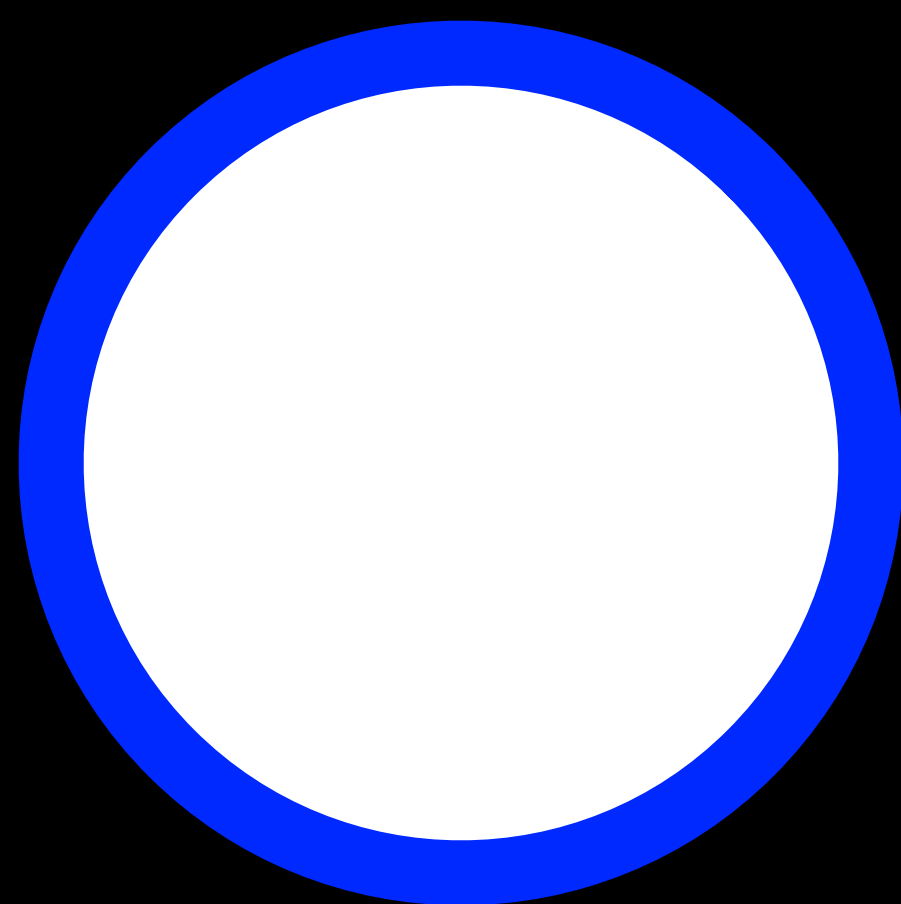
Recommended

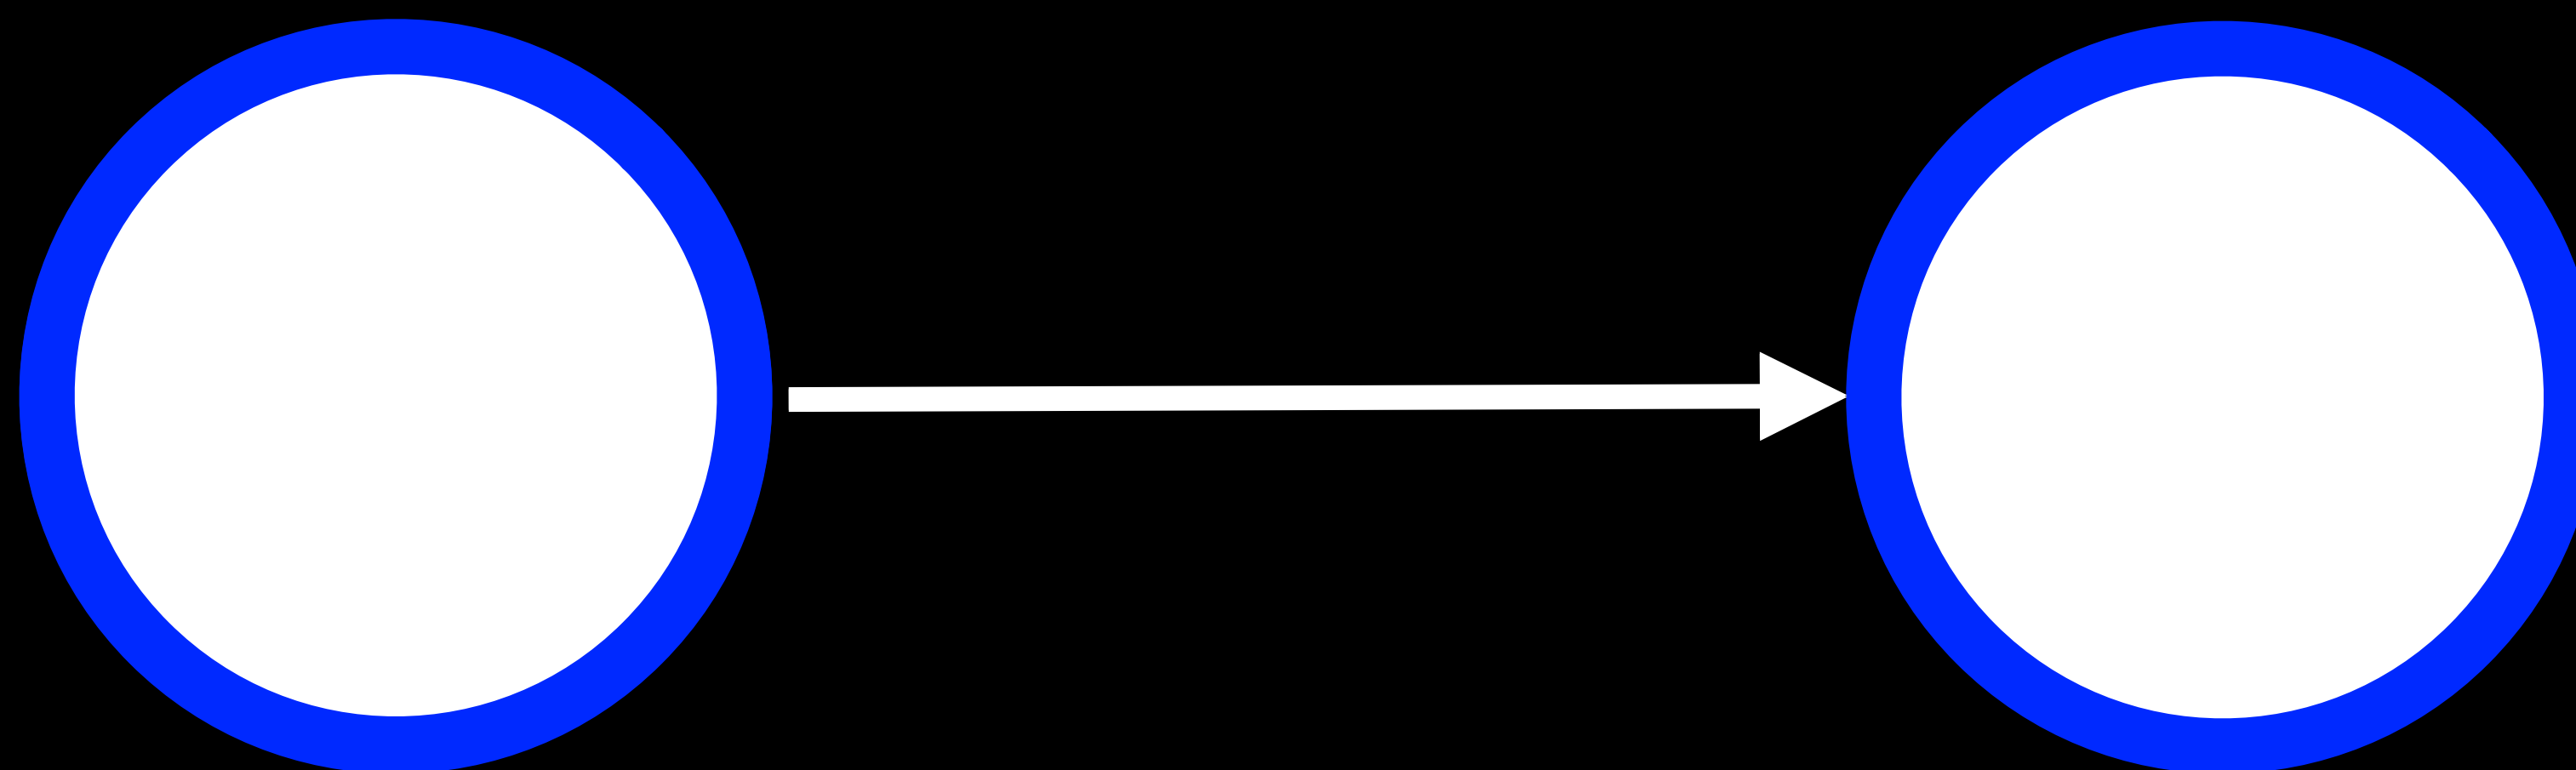


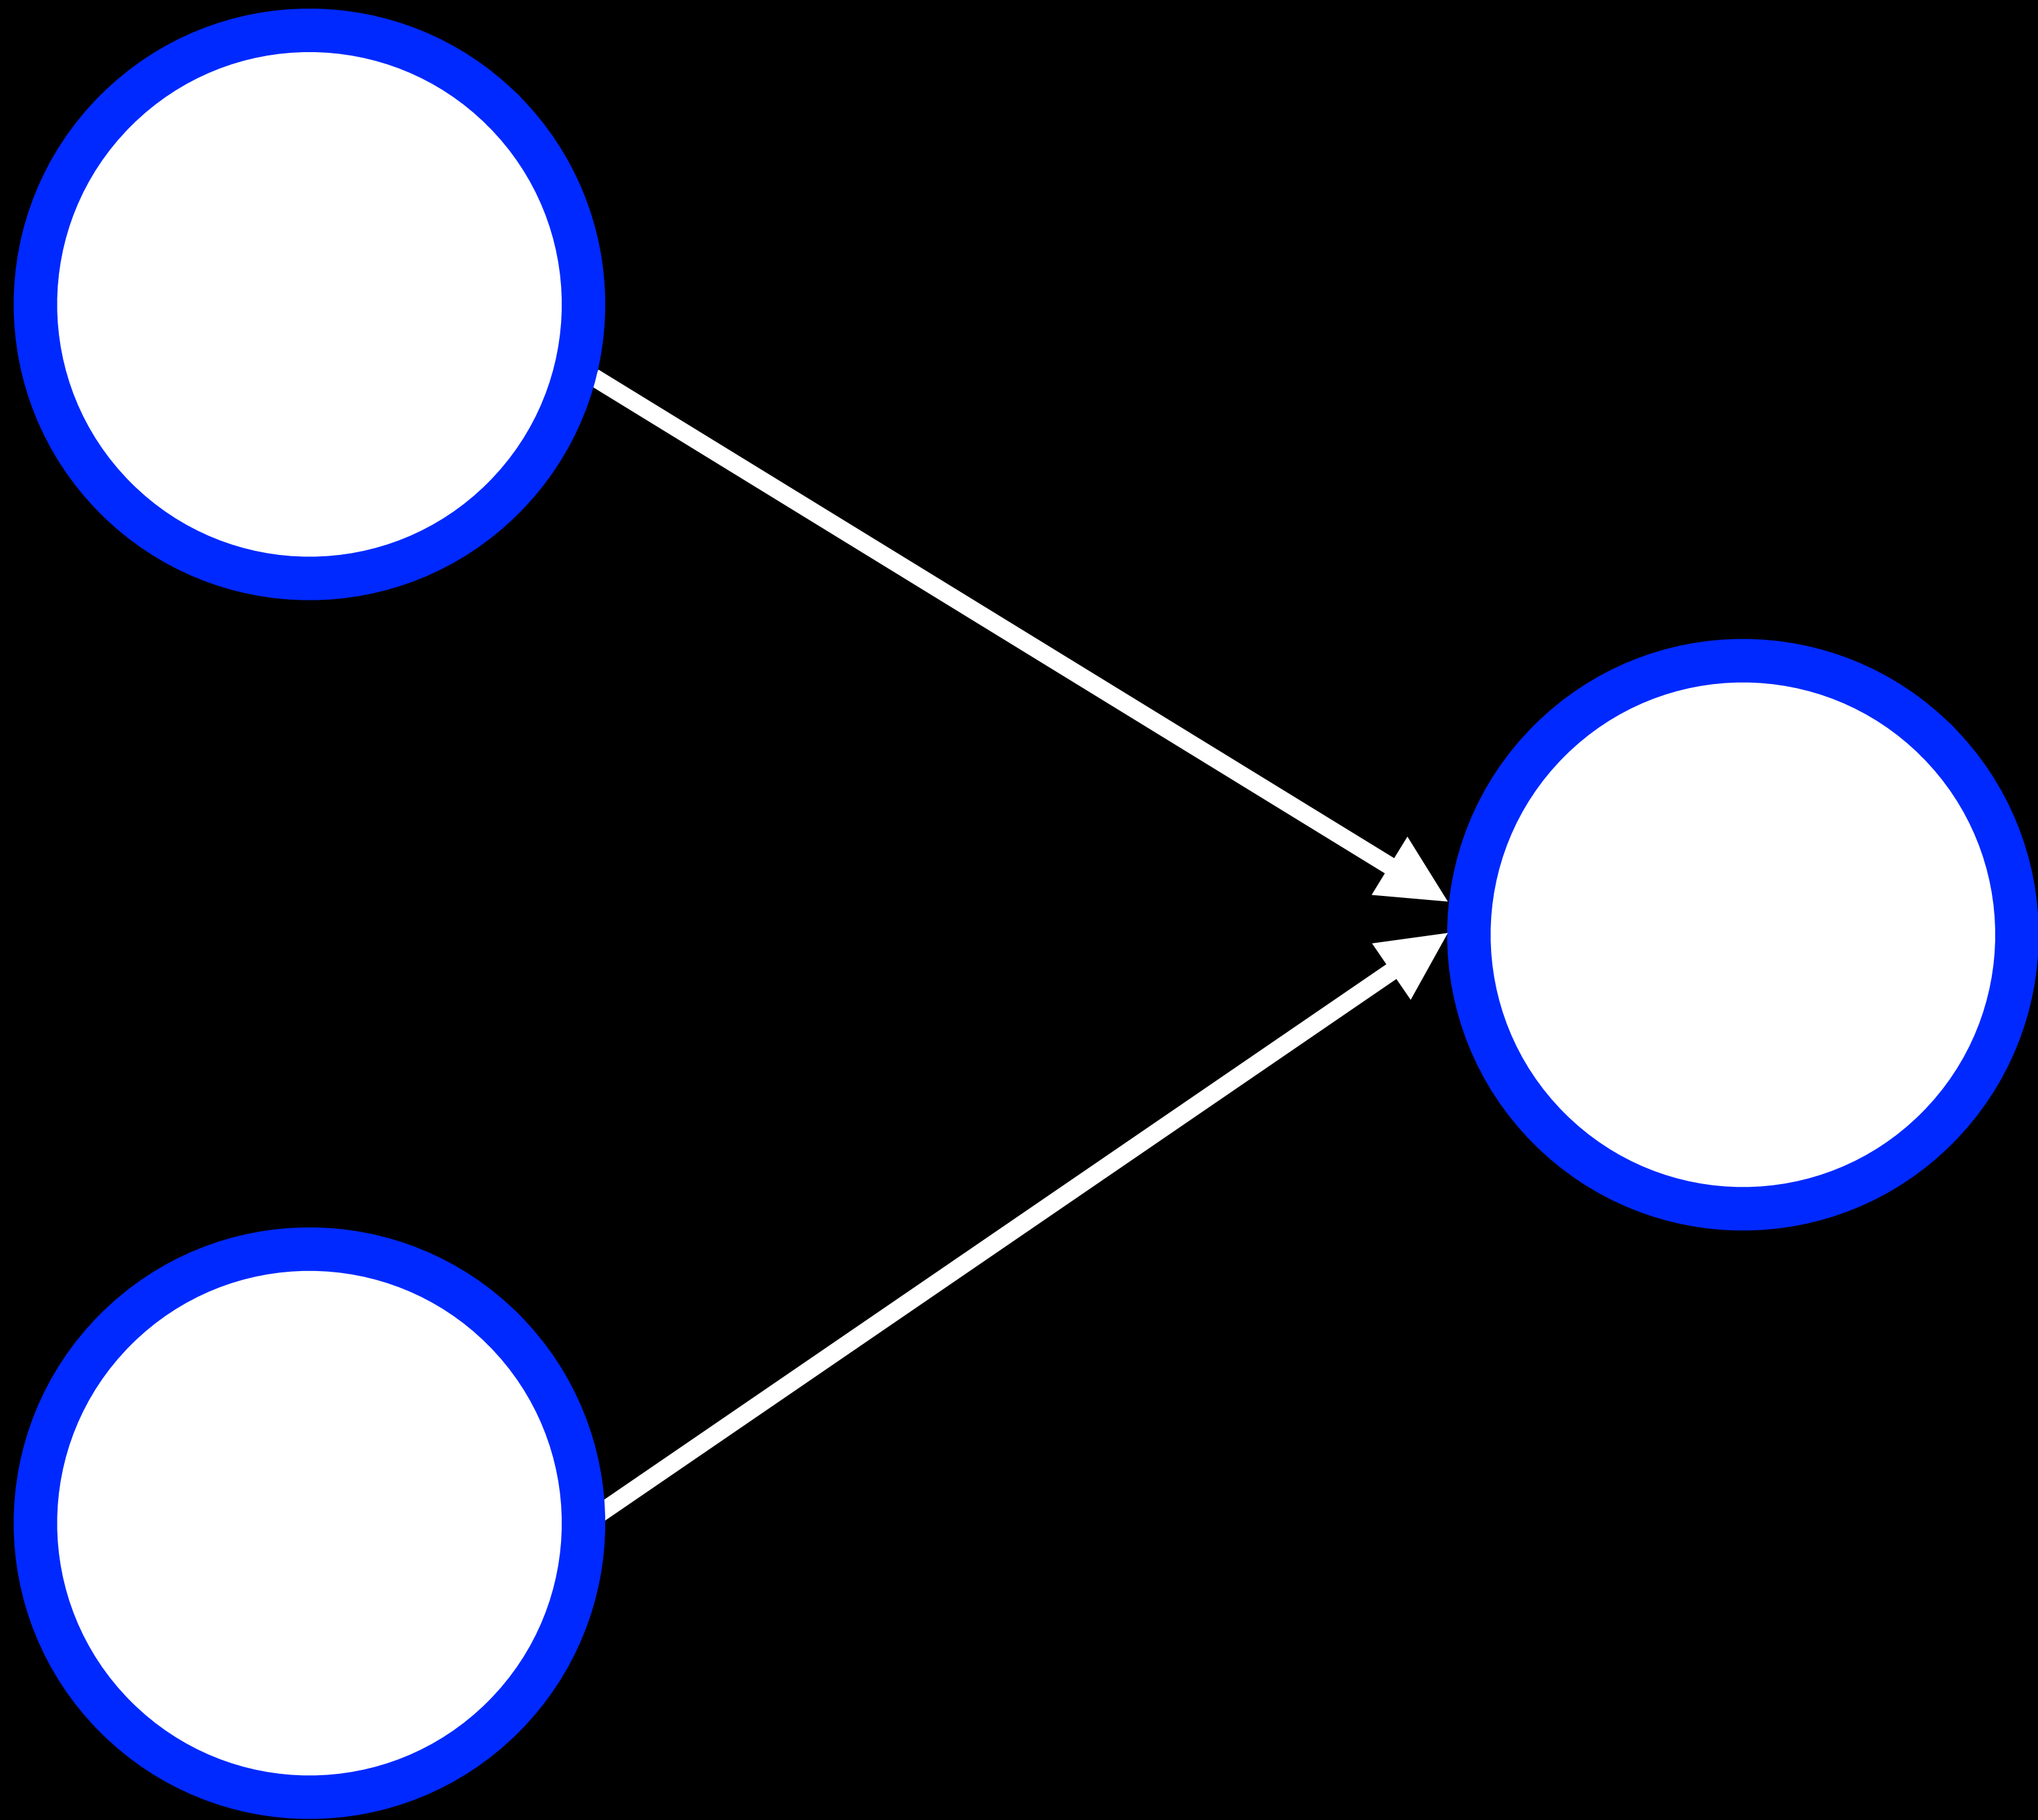


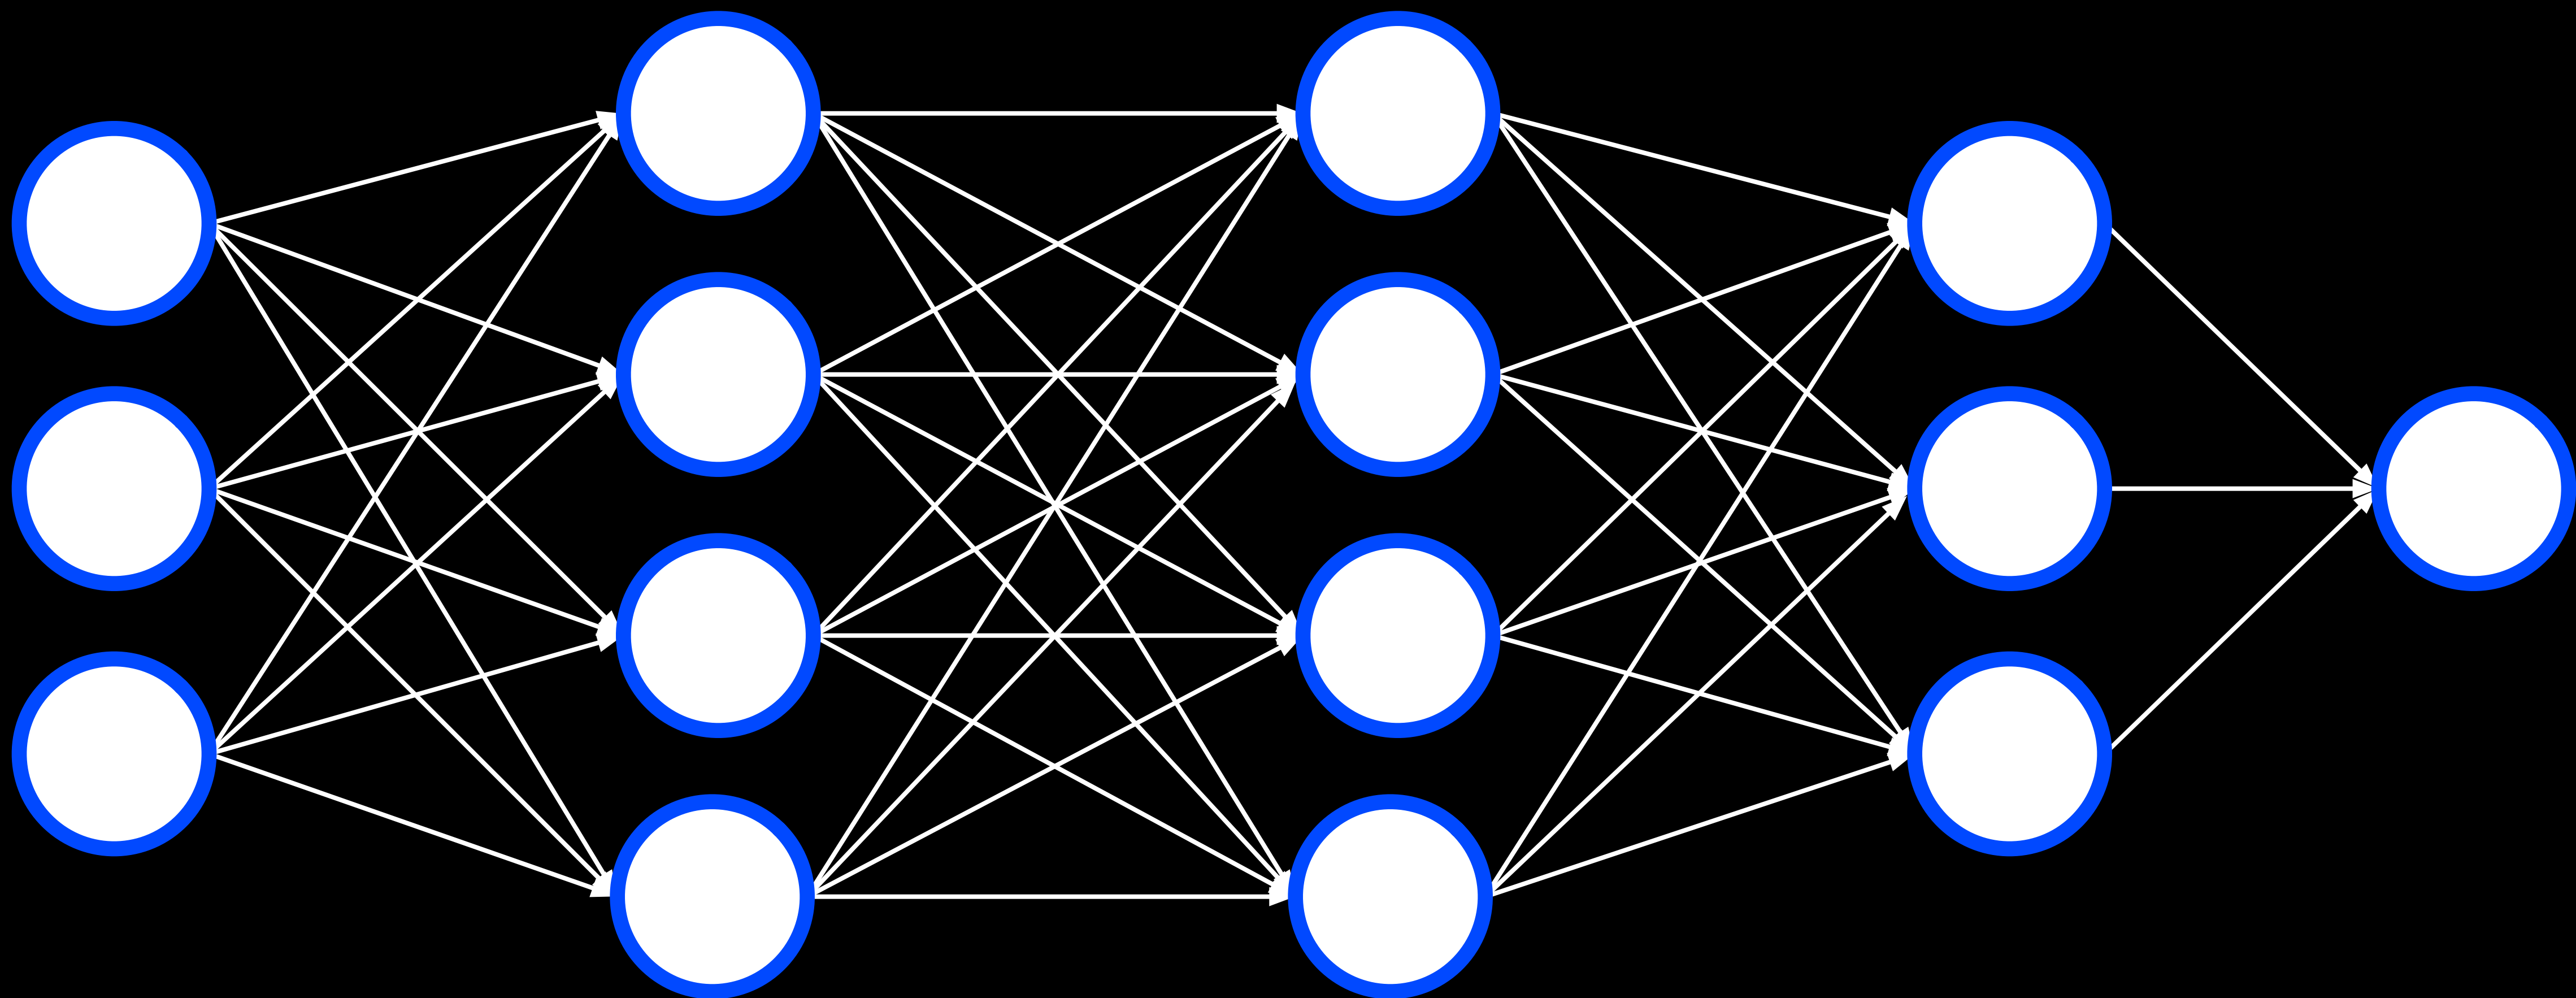


Neural Networks







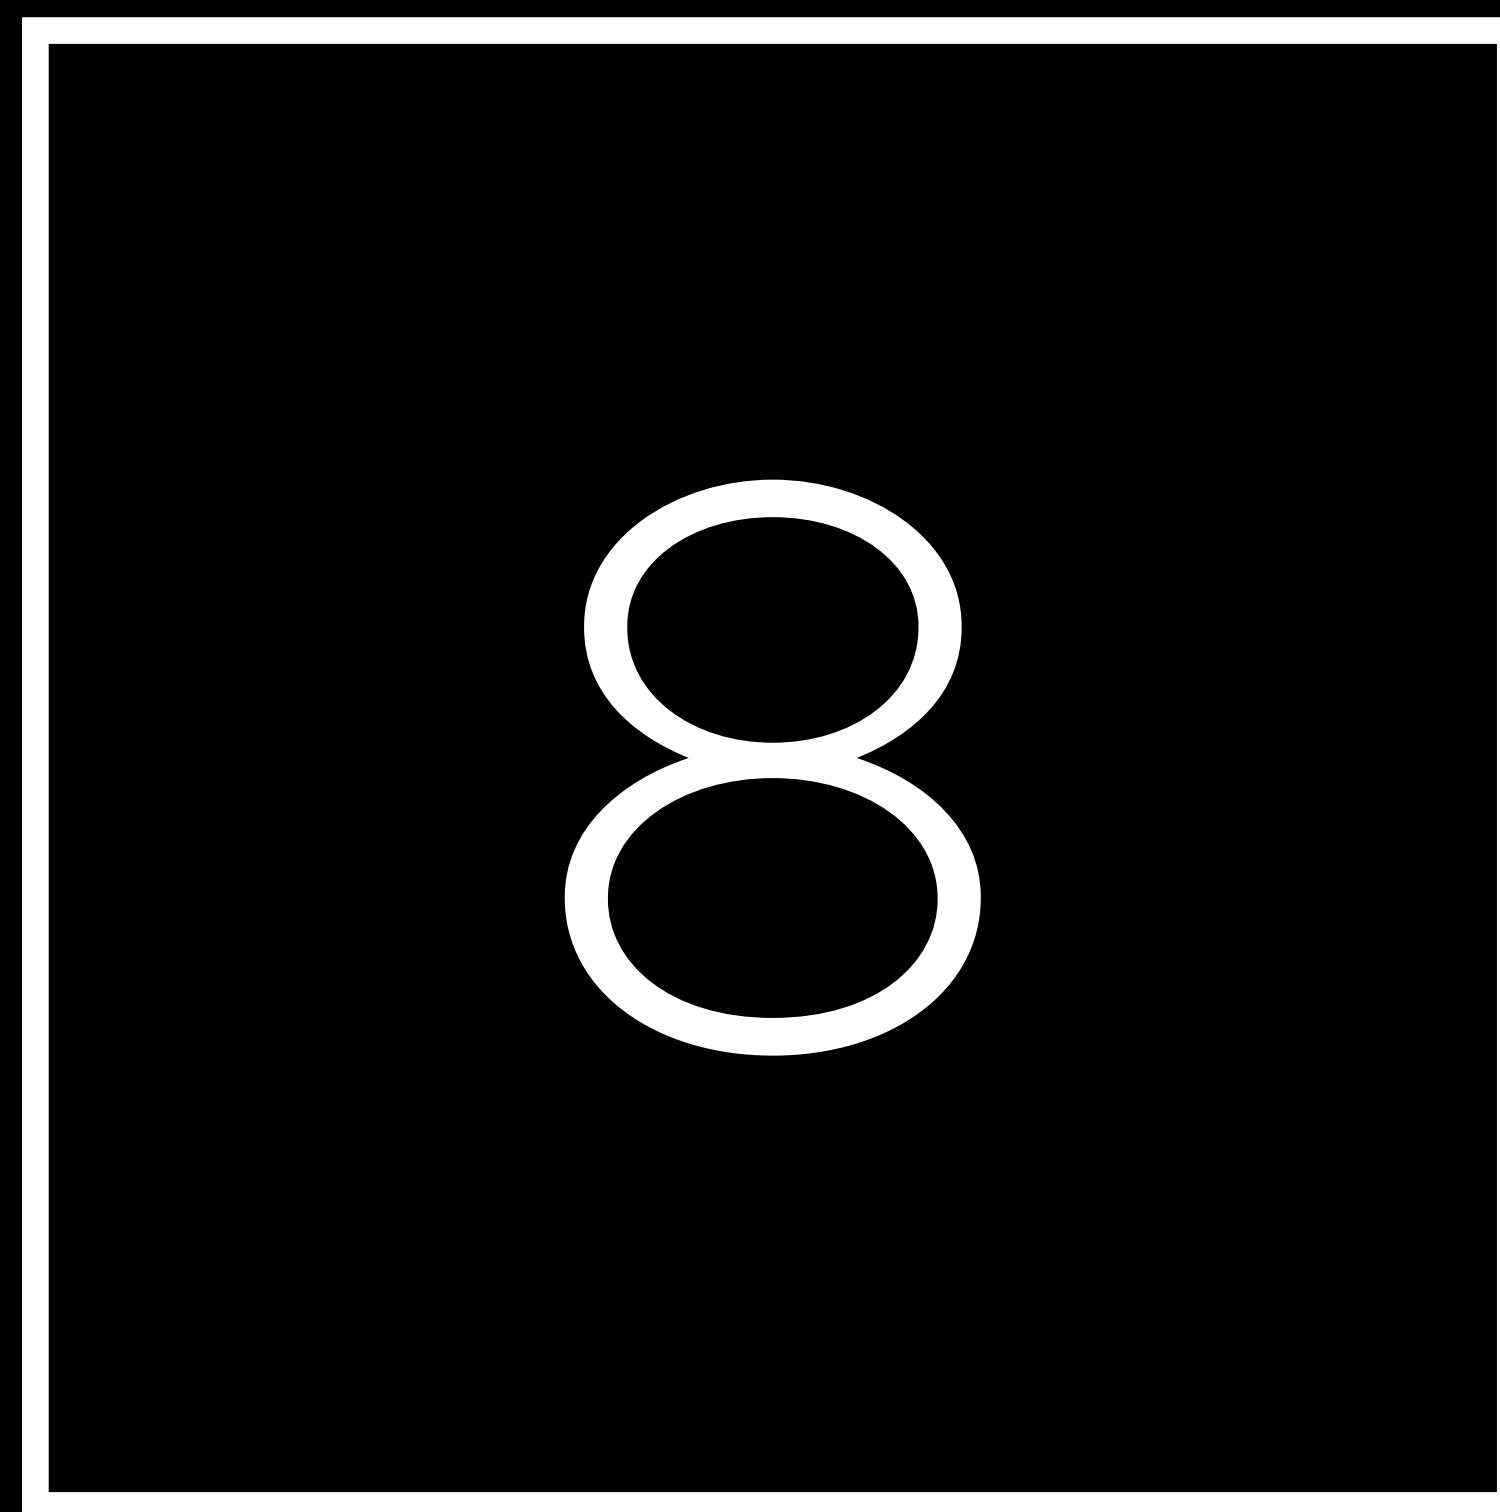


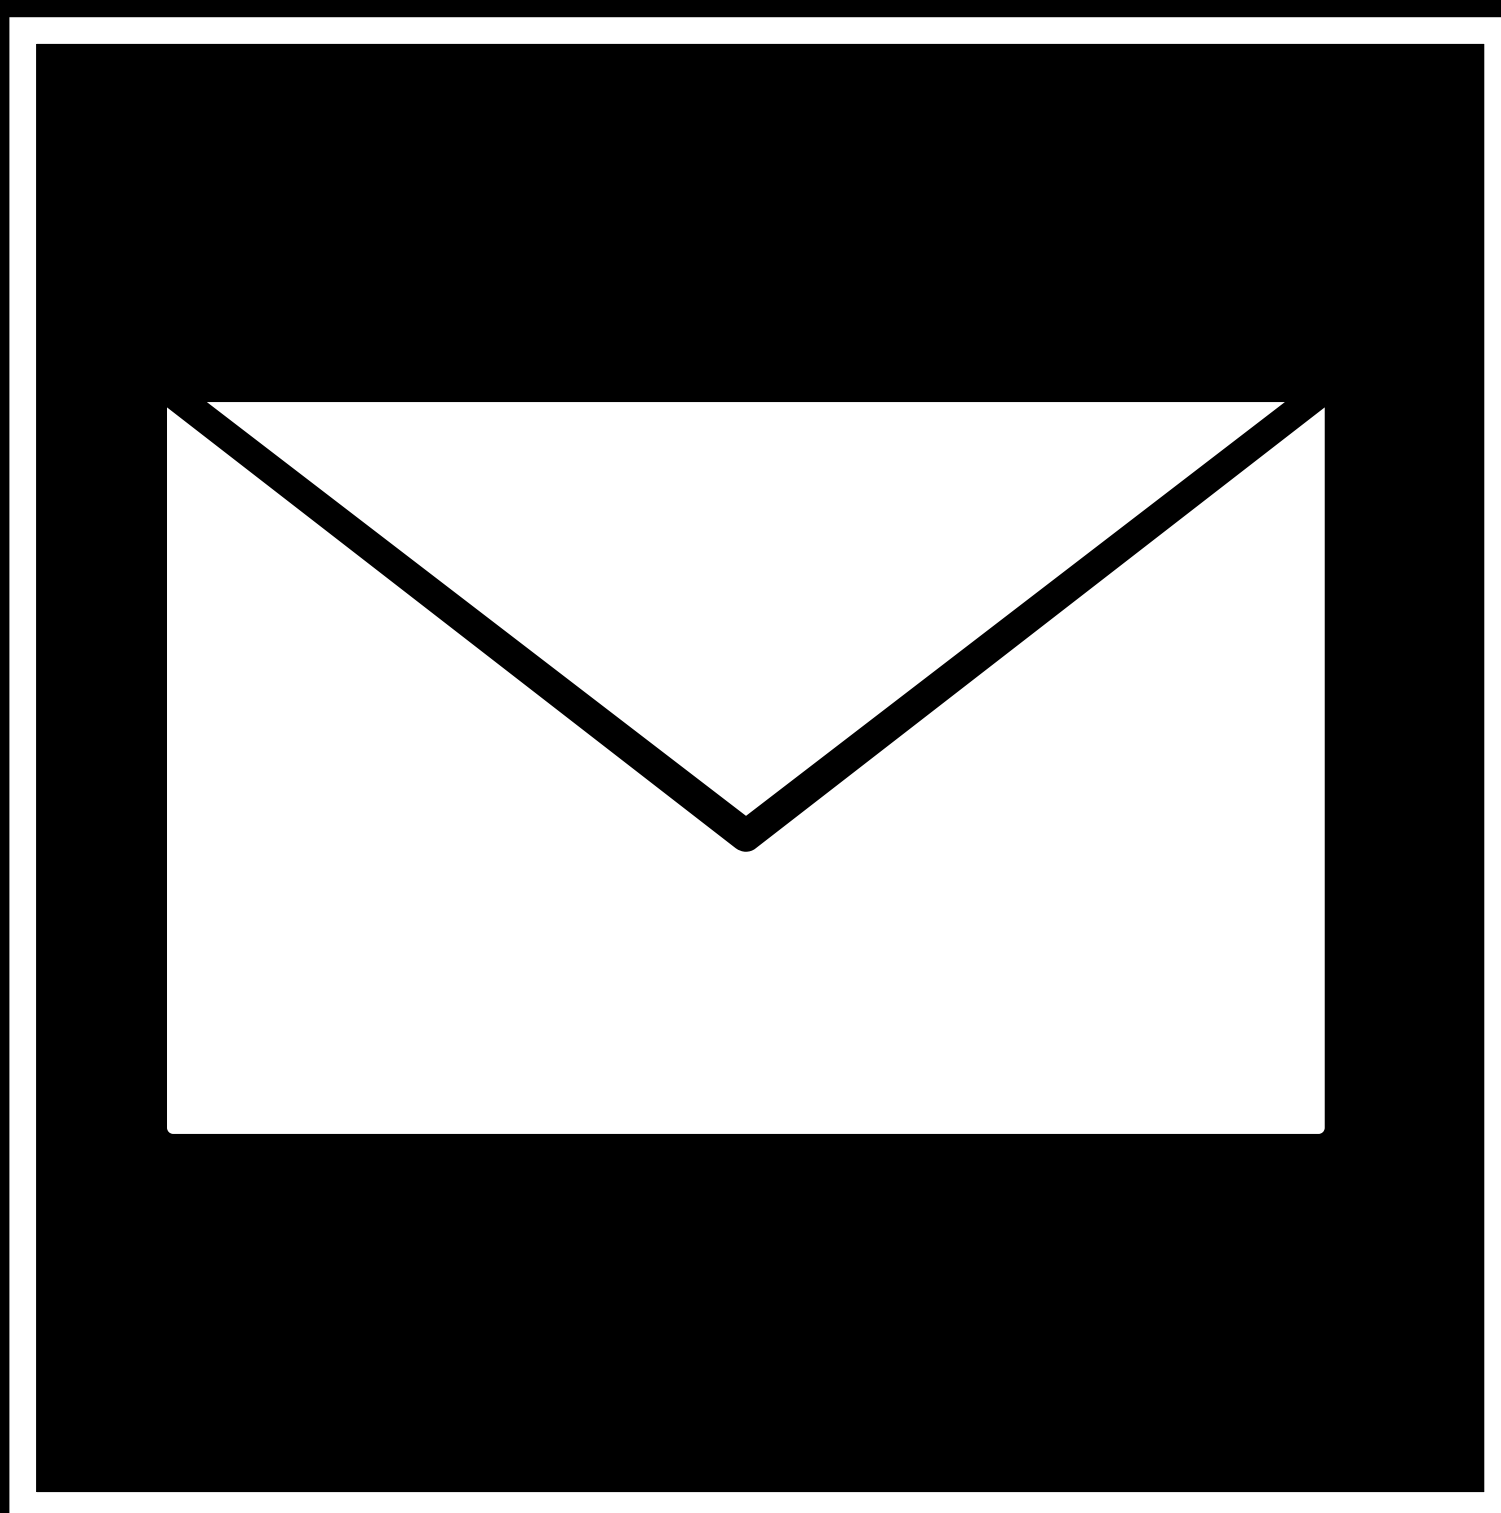
input



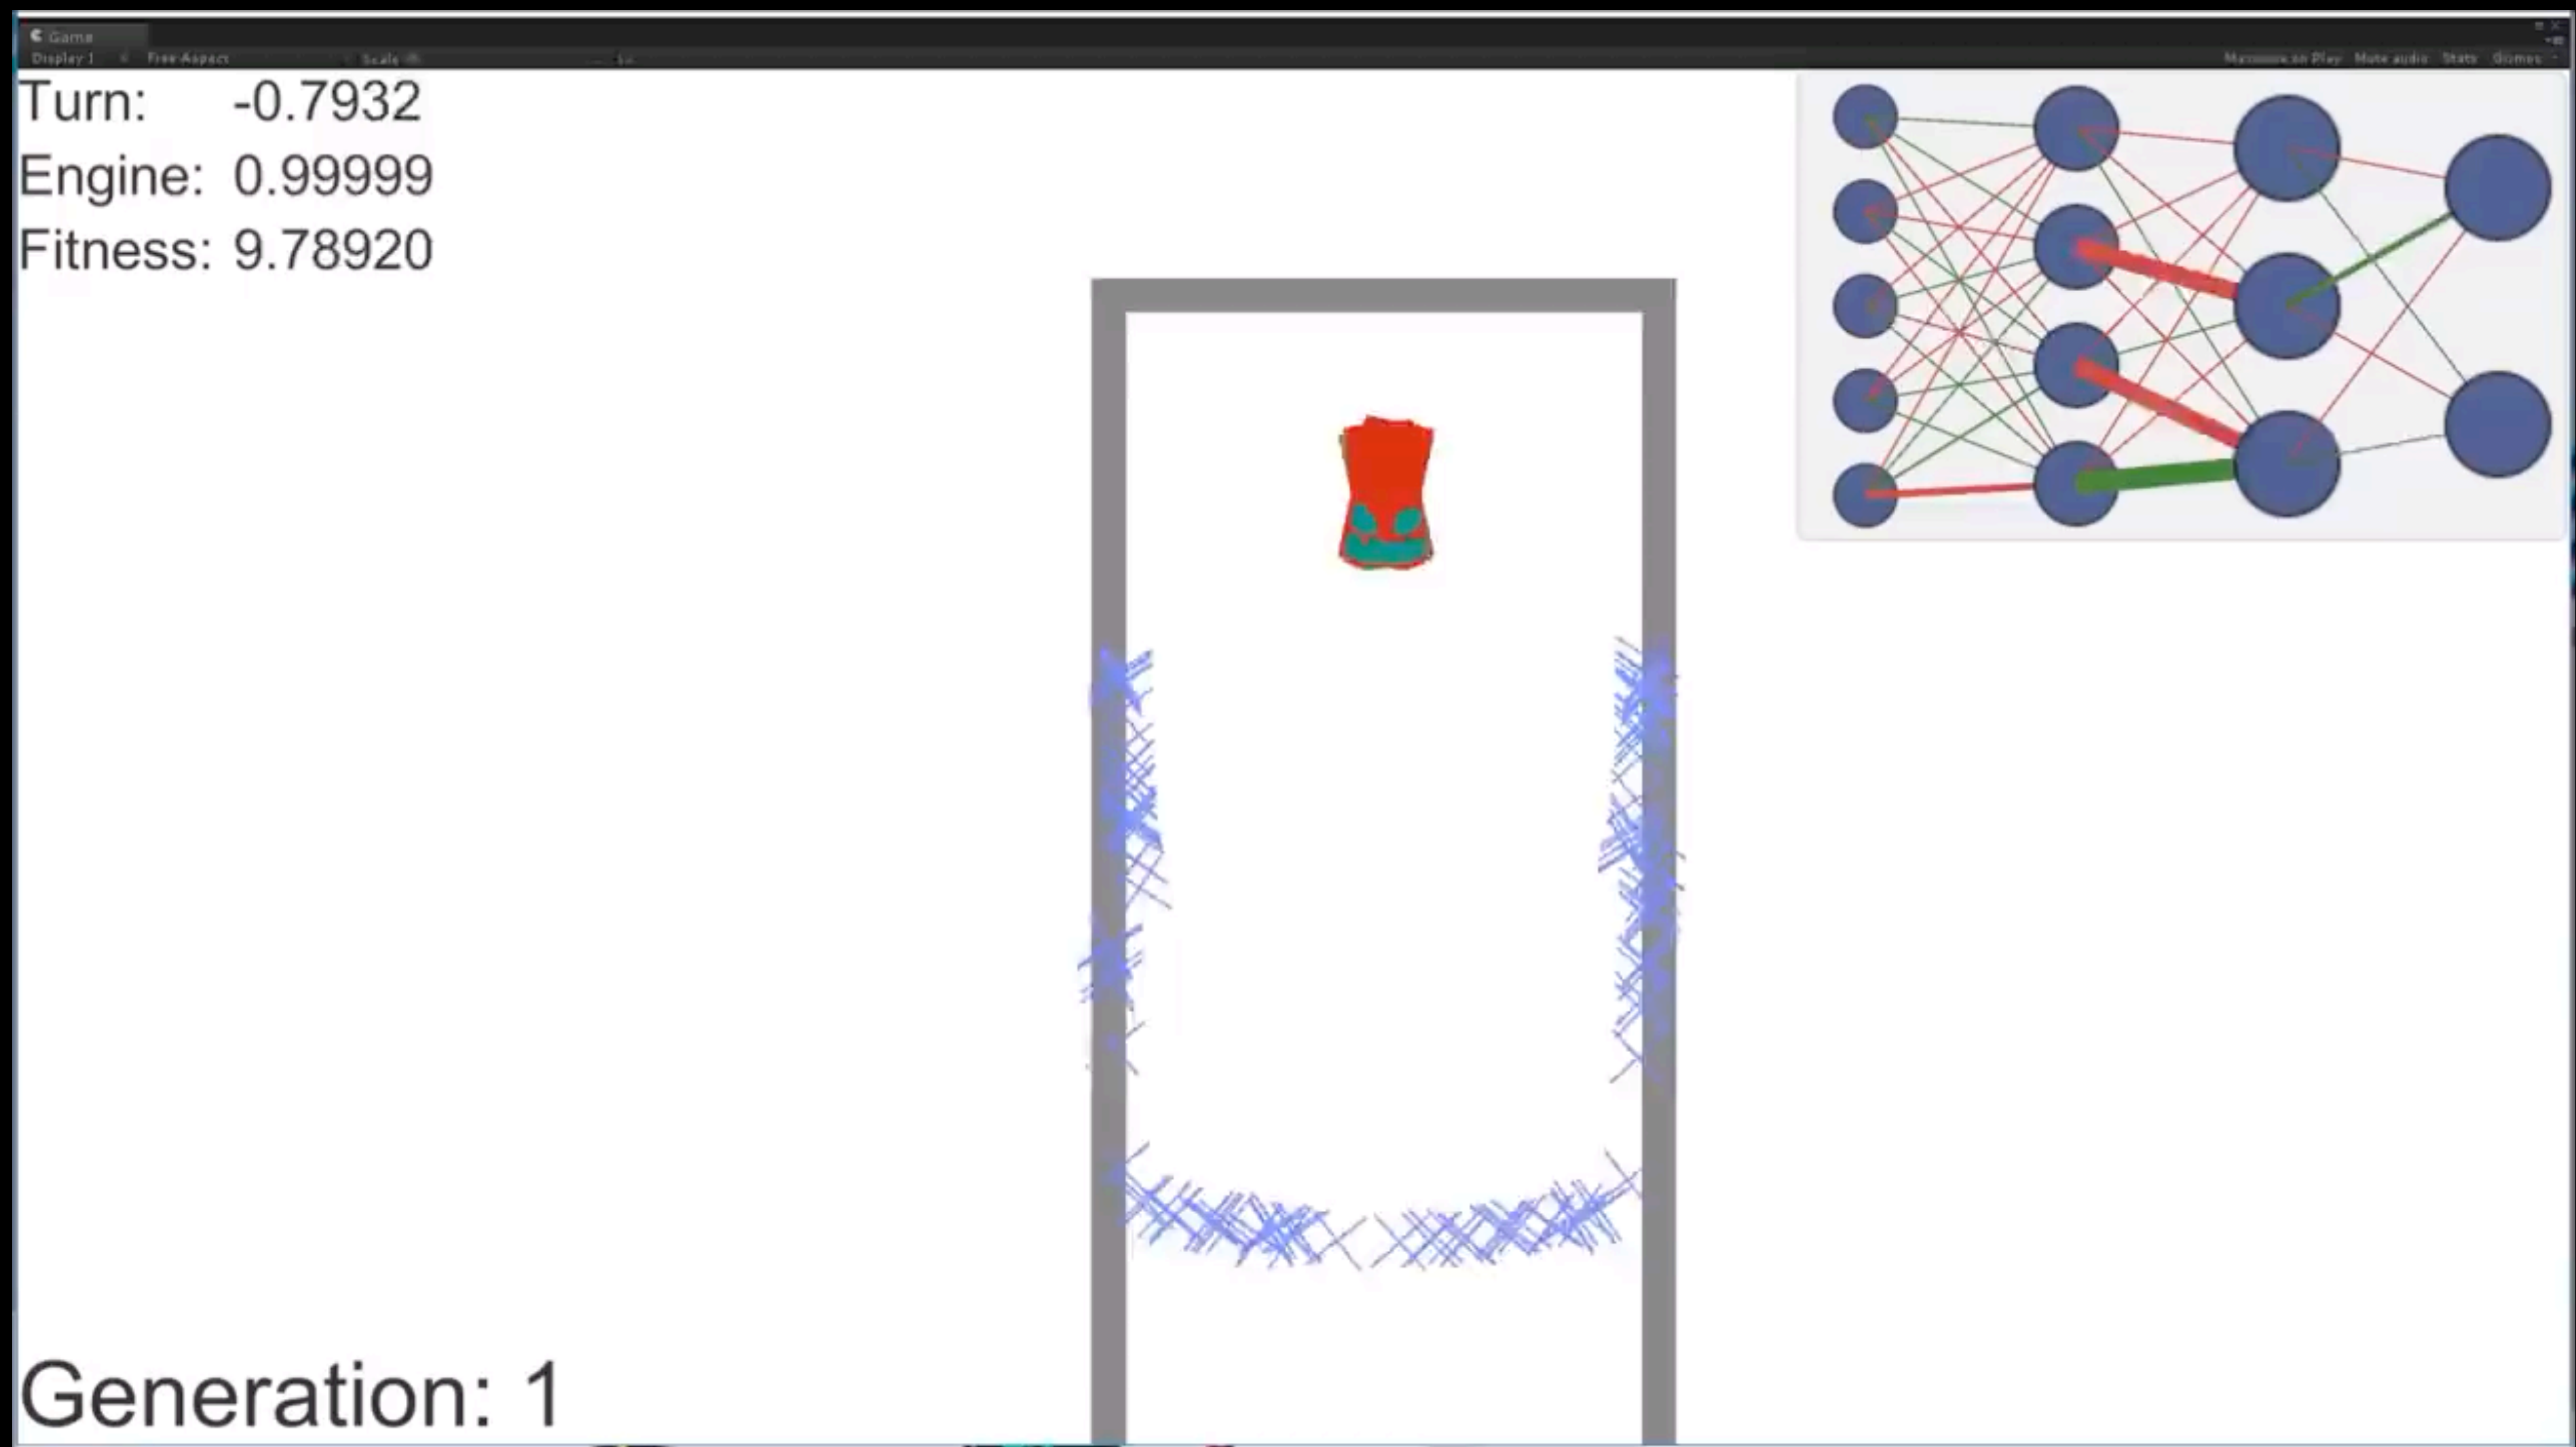
output

[illegible]





spam











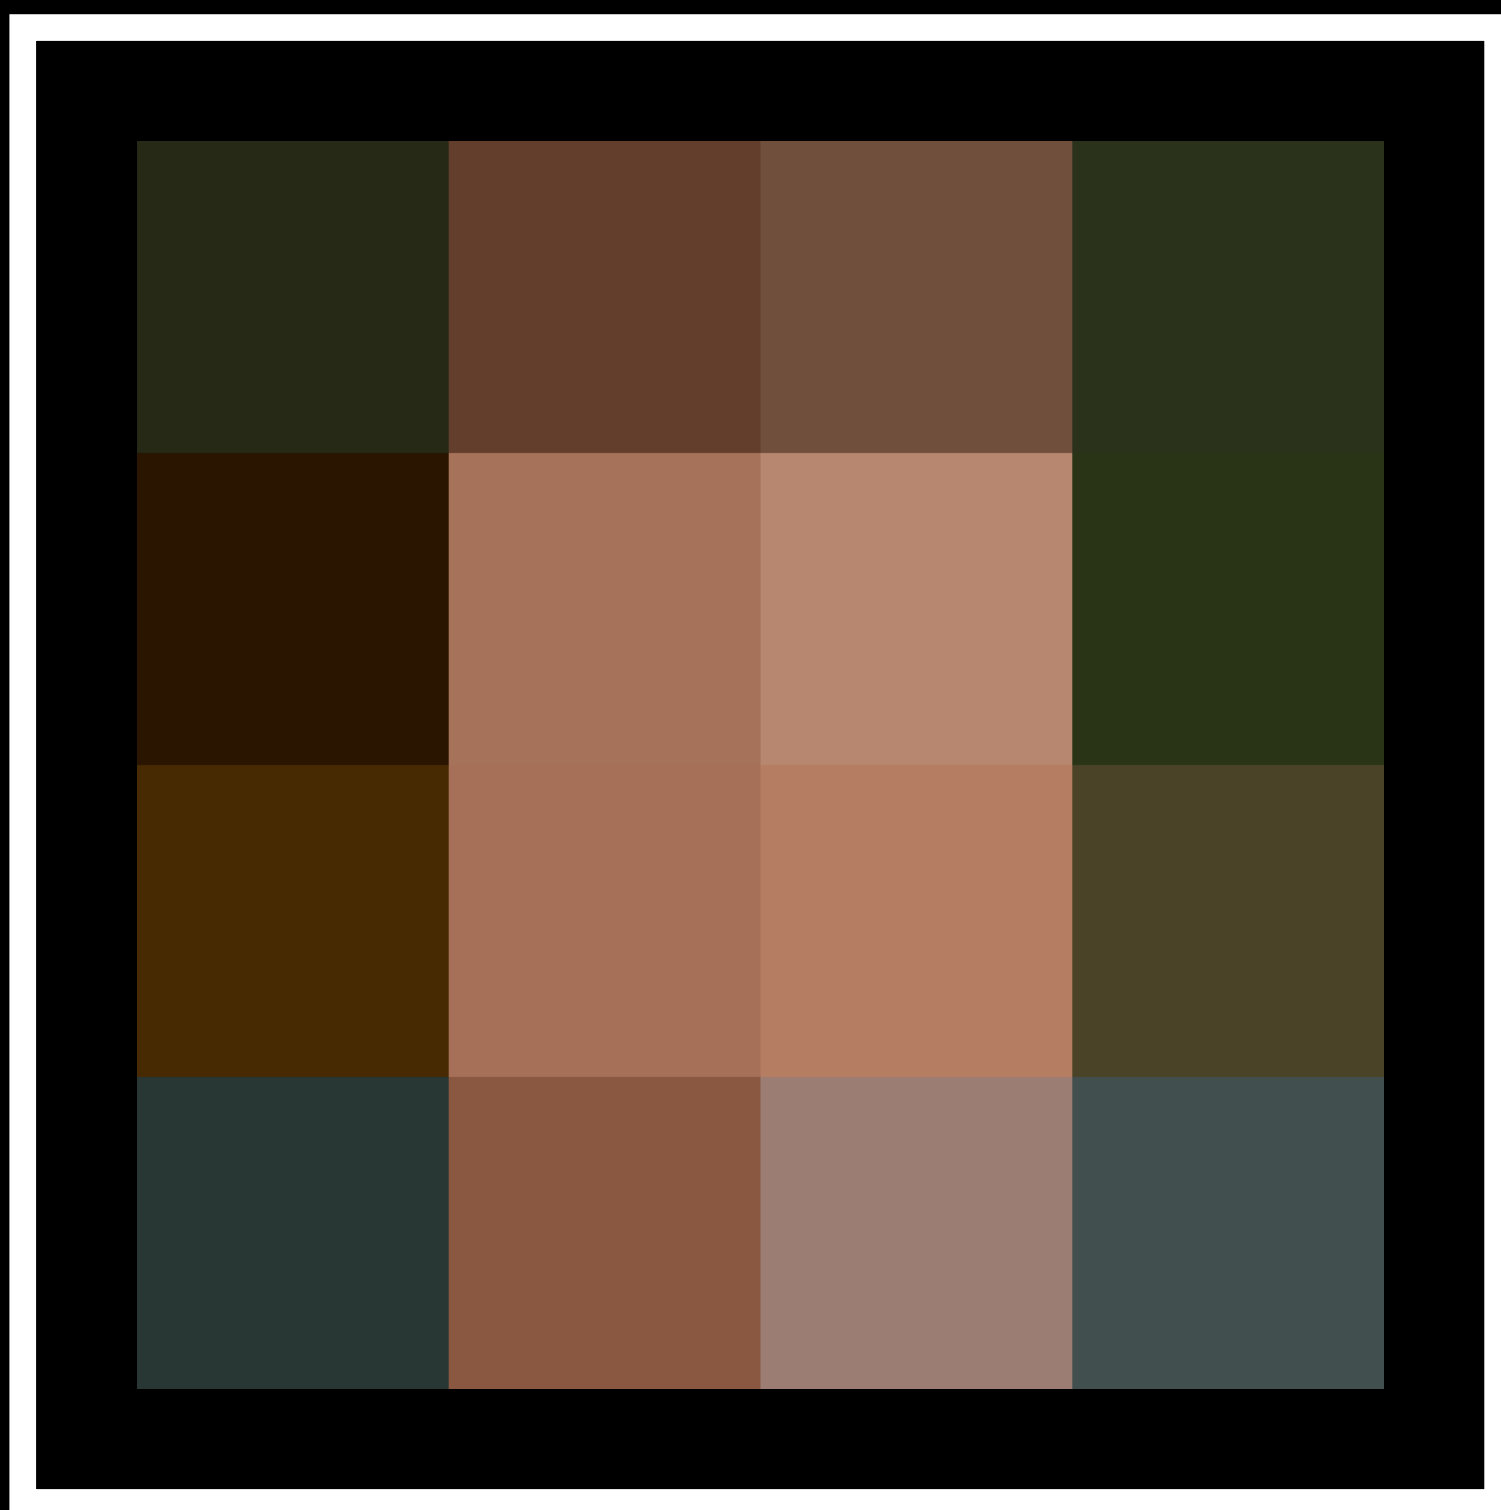


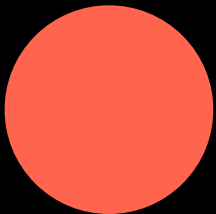
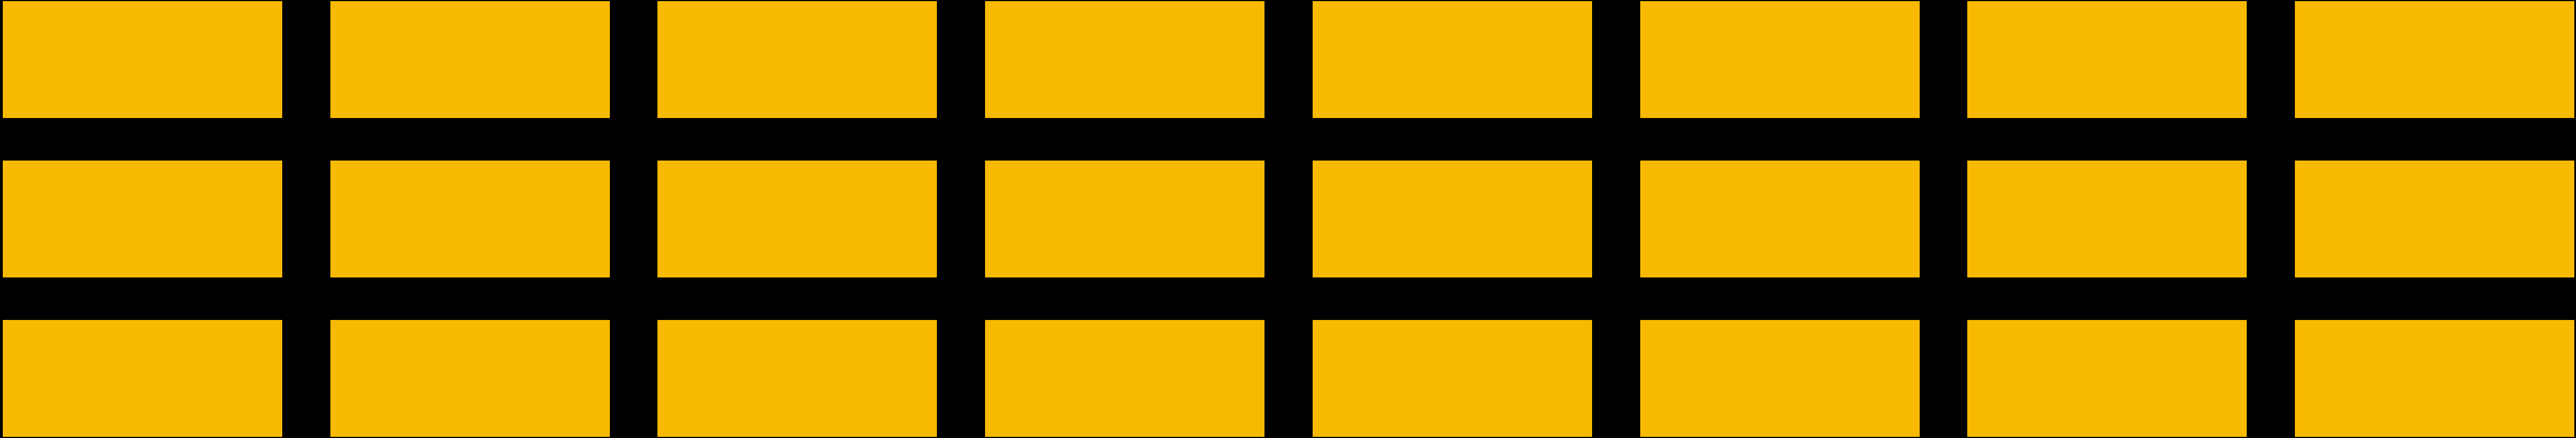












Google DeepMind

Training Time: 10 minutes



Google DeepMind

Training Time: 2 hours



Google DeepMind

Training Time: 4 hours



Artificial Intelligence