

CS540 Introduction to Artificial Intelligence

Lecture 17

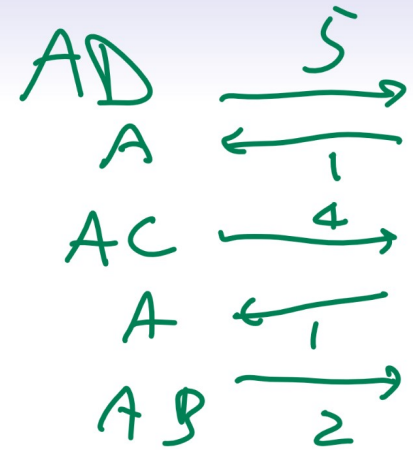
Young Wu

Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

July 26, 2021

Bridge and Torch Game

Motivation



- Four people with one flashlight (torch) want to go across a river. The bridge can hold two people at a time, and they must cross with the flashlight. The time it takes for each person to cross the river:

A	B	C	D
1	2	4	5

- What is the minimum total time required for everyone to cross the river?
- A: 10, B: 11, C: 12, D: 13, E: 14

Remind Me to Start Recording

Admin

- The messages you send in chat will be recorded: you can change your Zoom name now before I start recording.

Reduced Space Example 2

Quiz

Go through on Friday

- $\hat{\Sigma} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 3 \end{bmatrix}$. If one original data is $x = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$. What is

the new representation using only the first two principal components?

- A: $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$, B: $\begin{bmatrix} 2 \\ 1 \end{bmatrix}$, C: $\begin{bmatrix} 1 \\ 3 \end{bmatrix}$, D: $\begin{bmatrix} 3 \\ 1 \end{bmatrix}$, E: $\begin{bmatrix} 2 \\ 3 \end{bmatrix}$


Reduced Space Example 3

Quiz

- $\hat{\Sigma} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 3 \end{bmatrix}$. If one original data is $x = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$. What is the reconstructed vector using only the first two principal components?

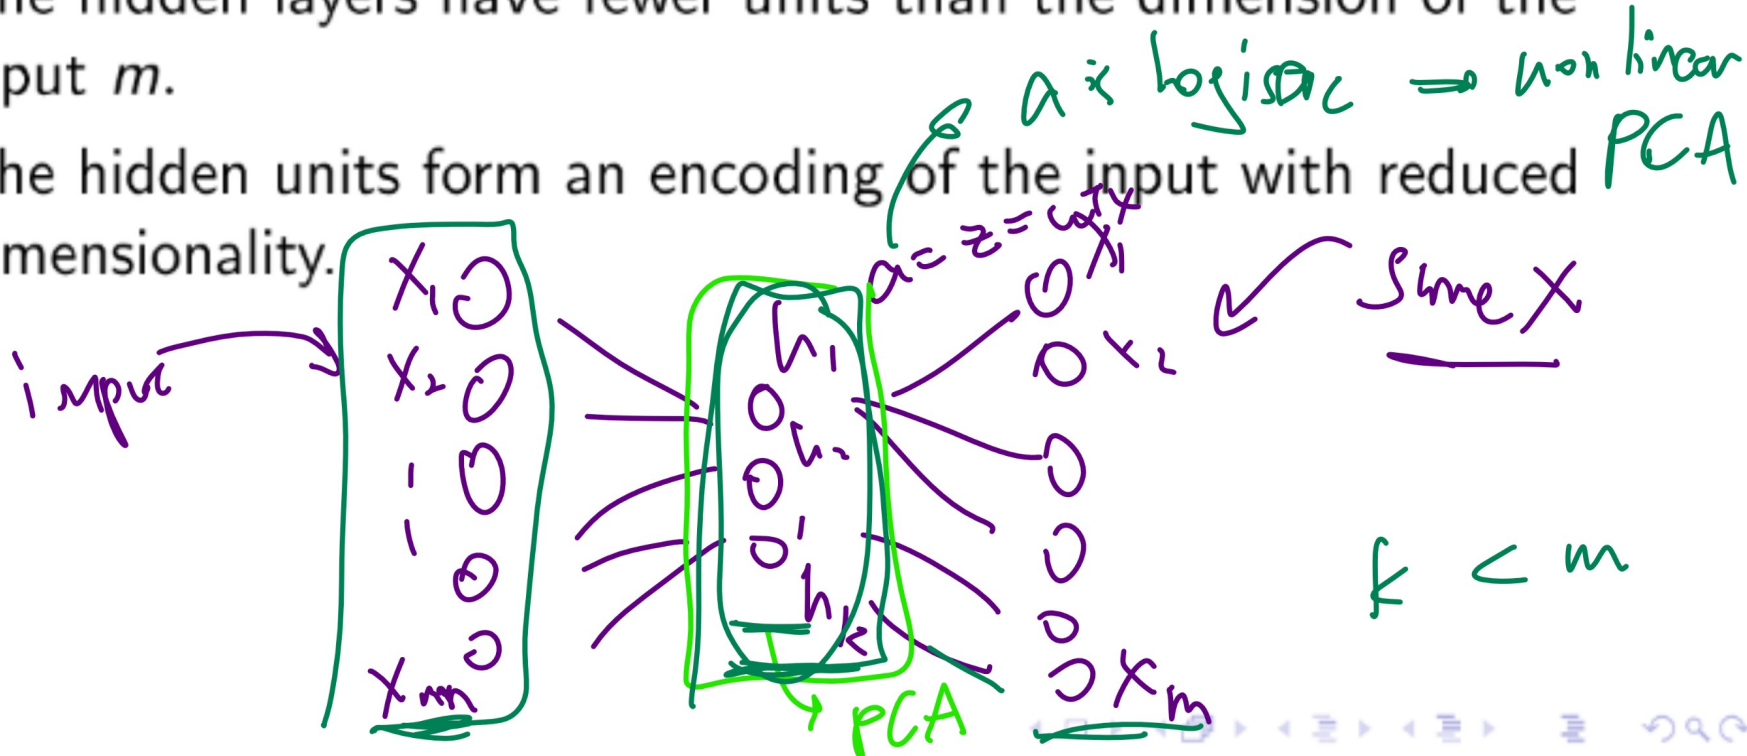
- A: $\begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}$, B: $\begin{bmatrix} 2 \\ 3 \\ 0 \end{bmatrix}$, C: $\begin{bmatrix} 0 \\ 2 \\ 3 \end{bmatrix}$, D: $\begin{bmatrix} 2 \\ 3 \\ 1 \end{bmatrix}$, E: $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$

Autoencoder Discussion



non-linear
PCA

- A multi-layer neural network with the same input and output $y_i = x_i$ is called an autoencoder.
- The hidden layers have fewer units than the dimension of the input m .
- The hidden units form an encoding of the input with reduced dimensionality.



Kernel PCA

Discussion

SVM

- A kernel can be applied before finding the principal components.

$$\hat{\Sigma} = \frac{1}{n-1} \sum_{i=1}^n \varphi(x_i) \varphi(x_i)^T$$

K Kernel matrix

var for transformed feature

- The principal components can be found without explicitly computing $\varphi(x_i)$, similar to the kernel trick for support vector machines.
- Kernel PCA is a non-linear dimensionality reduction method.

Learning vs Search

Motivation



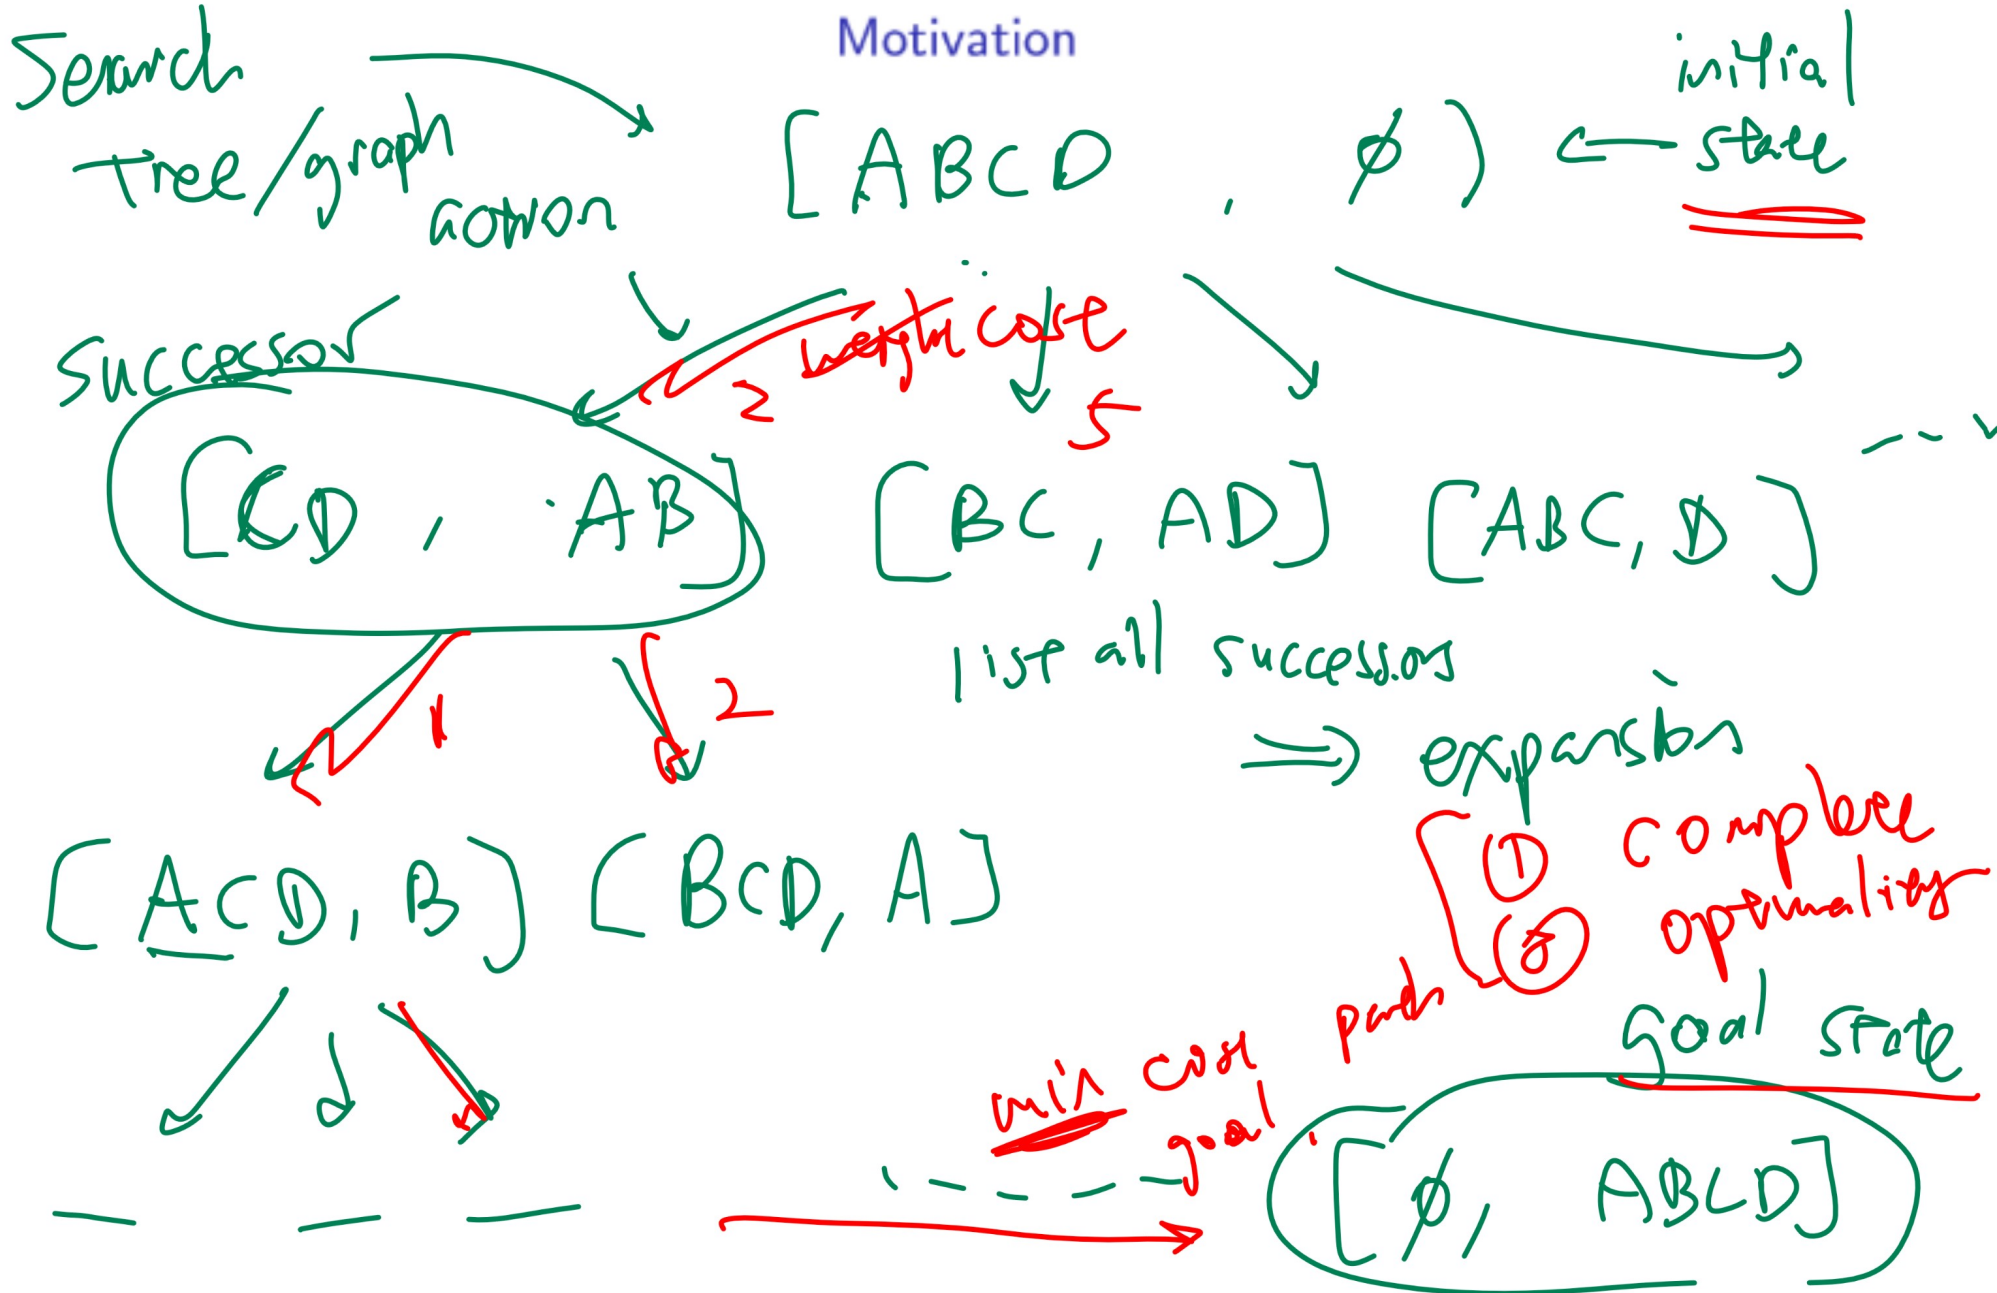
- In reinforcement learning, the reward and state transition need to be learned by taking actions.
- In search problems, the reward and state transitions are given.
- The problem is to find a sequence of actions that lead to the goal with minimum cost.

Q table



Bridge and Torch Game States 1

Motivation



Bridge and Torch Game States 2

Motivation

Complexity

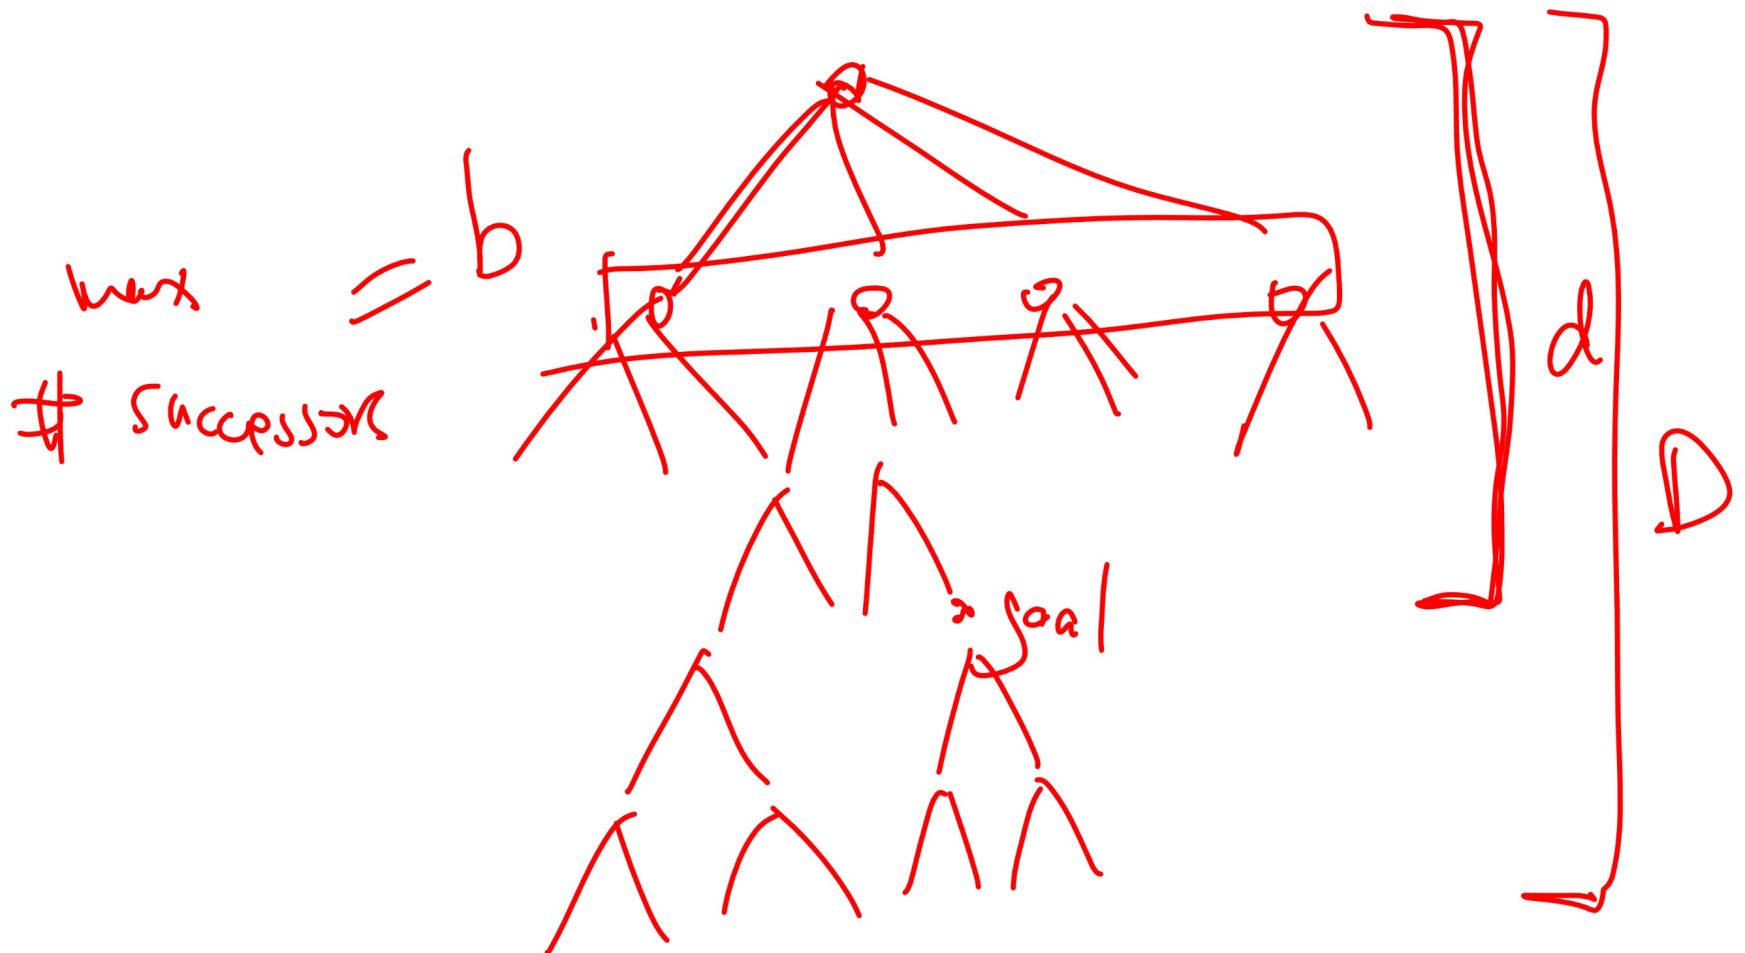
Definition

- The time complexity of a search strategy is the worst case maximum number of vertices expanded. *nodes*
- The space complexity of a search strategy is the worst case maximum number of states stored in the frontier at a single time. *search tree/graph*
- Notation: the goals are d edges away from the initial state. This means assuming a constant cost of 1, the optimal solution has cost d . The maximum depth of the graph is D .
- Notation: the branching factor is b , the maximum number of actions associated with a state.

$$b = \max_{s \in V} |s'(s)|$$

Search Tree Diagram

Definition



BFS Example 1

Quiz

- Fall 2018 Midterm Q2, Fall 2017 Midterm Q13, Fall 2010 Final Q2

- Suppose the states are positive integers between 1 and 10, initial state is 1, goal state is 9, successors of i is $2i$ and $2i + 1$ (if exist). What a BFS expansion sequence?

may get a


shorter path

Costs = 1

Queue


remove a state
add successors

~~1~~ 2 ~~3~~ 4 ~~5~~ ~~6~~ ~~7~~ ~~8~~ 9 10
stop when we expand goal (9)



BFS Example 1 Diagram

Quiz



BFS Example 2


Quiz

- Suppose the states are integers between 1 and $2^{10} = 1024$. The initial state is 1, and the goal state is 1024. The successors of a state i are $2i$ and $2i + 1$, if exist. How many states are expanded during a BFS search?

Q3

- A: 10
- B: 11
- C: 12
- D: 1023
- E: 1024

search length



~~1 2 3 4 5 6 7~~

BFS Example 3

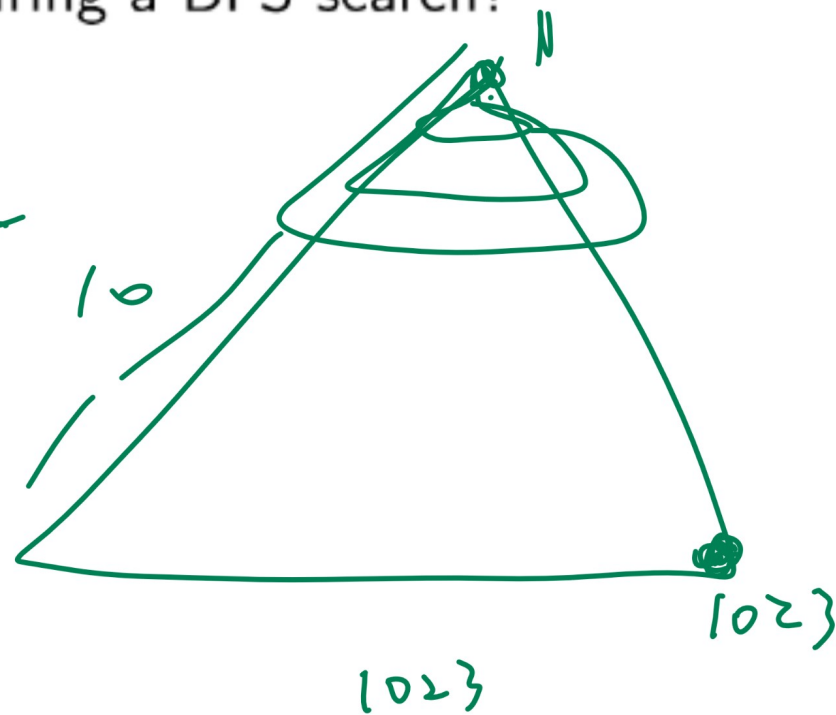
Quiz

Q 3

• Suppose the states are integers between 1 and $2^{10} - 1 = 1023$. The initial state is 1, and the goal state is 1023. The successors of a state i are $2i$ and $2i + 1$, if exist. How many states are expanded during a BFS search?

- A: 10
- B: 11
- C: 12
- D: 1023
- E: 1024

*include
expand goal*



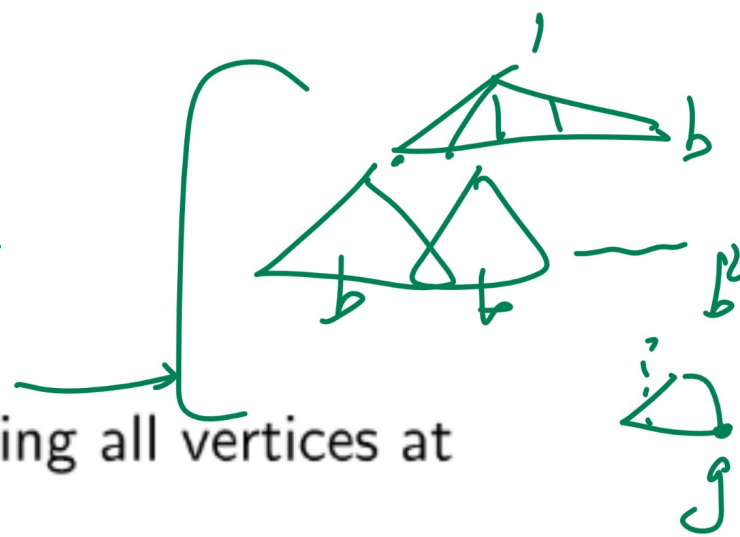
Breadth First Search Complexity

Discussion

- Time complexity: the worst case occurs when the goal is the last vertex at depth d .


$$T \approx b + b^2 + \dots + b^d$$

$O(b^d)$



- Space complexity: the worst case is storing all vertices at depth d is in the frontier.

$$S = b^d$$



DFS Example 1

Quiz

Space: 3, 5, 8, 9, Time: 1, 2, 4, 8, 9

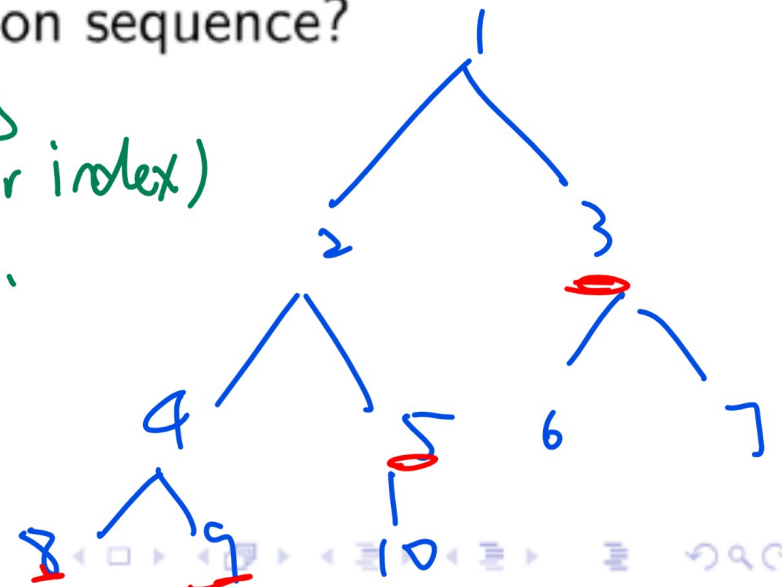

- Fall 2018 Midterm Q2, Fall 2017 Midterm Q13, Fall 2010 Final Q2
- Suppose the states are positive integers between 1 and 10, initial state is 1, goal state is 9, successors of i is $2i$ and $2i + 1$ (if exist). What a DFS expansion sequence?

⑤

④

convention put nodes on the right (larger index) first.

Stack



DFS Example 1 Diagram

Quiz

DFS Example 2

Quiz


Q4

Suppose the states are integers between 1 and $2^{10} = 1024$. The initial state is 1, and the goal state is 1024. The successors of a state i are $2i$ and $2i + 1$, if exist. How many states are expanded during a DFS search?

- A: 10
- B: 11
- C: 12
- D: 1023
- E: 1024

Goal is expanded

1024



DFS Example 3

Quiz

Q5

• Suppose the states are integers between 1 and $2^{10} - 1 = 1023$. The initial state is 1, and the goal state is 1023. The successors of a state i are $2i$ and $2i + 1$, if exist. How many states are expanded during a DFS search?


- A: 10
- B: 11
- C: 12
- D: 1023
- E: 1024



Depth First Search Performance

Discussion

- DFS is incomplete if $D = \infty$.
- DFS is not optimal.



Depth First Search Complexity

Discussion

- Time complexity: the worst case occurs when the goal is the root of the last subtree expanded in the whole graph.

$$T = b^{D-d+1} \dots + b^{D-1} + b^D$$

- Space complexity: the worst case is storing all vertices sharing the parents with vertices in the current path.

$$S = (b - 1) D + 1$$

IDS Example 1

Quiz

- Fall 2018 Midterm Q2, Fall 2017 Midterm Q13, Fall 2010 Final Q2
- Suppose the states are positive integers between 1 and 10, initial state is 1, goal state is 9, successors of i is $2i$ and $2i + 1$ (if exist). What a IDS expansion sequence?

Iterative Deepening Search

Algorithm

- Input: a weighted digraph (V, E, c) , initial states I and goal states G .
- Output: a path from I to G .
- Perform DFS on the digraph restricted to vertices with depth ≤ 1 from the initial state.
- Perform DFS on the digraph restricted to vertices with depth ≤ 2 from the initial state.
- Repeat until the goal is deQueued.

Iterative Deepening Search Performance

Discussion

- IDS is complete.
- IDS is optimal with $c = 1$.

Iterative Deepening Search Complexity

Discussion

- Time complexity: the worst case occurs when the goal is the last vertex at depth d .

$$T = db + (d - 1) b^2 + \dots + 3b^{d-2} + 2b^{d-1} + 1b^d$$

- Space complexity: it has the same space complexity as DFS.

$$S = (b - 1) d$$

Configuration Space

Discussion