

Q3. [9 pts] Search

Suppose we have a connected graph with N nodes, where N is finite but large. Assume that every node in the graph has exactly D neighbors. **All edges are undirected.** We have exactly one start node, S , and exactly one goal node, G .

Suppose we know that the shortest path in the graph from S to G has length L . That is, it takes at least L edge-traversals to get from S to G or from G to S (and perhaps there are other, longer paths).

We'll consider various algorithms for searching for paths from S to G .

(a) [2 pts] Uninformed Search

Using the information above, give the tightest possible bounds, using big \mathcal{O} notation, on **both the absolute best case and the absolute worst case number of node expansions** for each algorithm. Your answer should be a function in terms of variables from the set $\{N, D, L\}$. You may not need to use every variable.

(i) [1 pt] DFS **Graph** Search

Best case: $\mathcal{O}(L)$. If we are lucky, DFS could send us directly on the shortest path to the goal without expanding anything else. Worst case: $\mathcal{O}(N)$. Worst Case is we expand every node in the graph before expanding G ; because this is graph search, we can't expand anything more than once.

(ii) [1 pt] BFS **Tree** Search

Best case: $\mathcal{O}(D^{L-1})$ Worst case: $\mathcal{O}(D^L)$

In the best case, G is the first node expanded at depth L of the tree (expanded immediately after all nodes of depth $L - 1$ are expanded). The structure of the graph gives that there are no more than D^{L-1} nodes of depth $L - 1$, and since we can ignore this one extra node at depth L in the asymptotic bound, we have $\mathcal{O}(D^{L-1})$. In the worst case, BFS needs to expand all paths with depth $\leq L$ (i.e. G is the last node of depth L expanded), and so needs to expand $\mathcal{O}(D^L)$ nodes.

(b) [2 pts] Bidirectional Search

Notice that because the graph is undirected, finding a path from S to G is equivalent to finding a path from G to S , since reversing a path gives us a path from the other direction of the same length.

This fact inspired **bidirectional search**. As the name implies, bidirectional search consists of two simultaneous searches which both use the same algorithm; one from S towards G , and another from G towards S . When these searches meet in the middle, they can construct a path from S to G .

More concretely, in bidirectional search:

- We start Search 1 from S and Search 2 from G .
- The searches take turns popping nodes off of their separate fringes. First Search 1 expands a node, then Search 2 expands a node, then Search 1 again, etc.
- This continues until one of the searches expands some node X which the other search has also expanded.
- At that point, Search 1 knows a path from S to X , and Search 2 knows a path from G to X , which provides us with a path from X to G . We concatenate those two paths and return our path from S to G .

Don't stress about further implementation details here!

Repeat part (a) with the bidirectional versions of the algorithms from before. Give the tightest possible bounds, using big \mathcal{O} notation, on both the absolute best and worst case number of node expansions by the bidirectional search algorithm. Your bound should still be a function of variables from the set $\{N, D, L\}$.


(i) [1 pt] Bidirectional DFS **Graph** Search

Best case: $\mathcal{O}(L)$. Bidirectional Search does not meaningfully change the number of nodes visited for DFS. If we are lucky, Bidi-DFS could send us directly on the shortest path in both directions without expanding anything else. Worst case: $\mathcal{O}(N)$. Worst Case is our two searches expands every node in the graph before meeting at some X ; because this is graph search, we can't expand anything more than once.

(ii) [1 pt] Bidirectional BFS **Tree** Search

Best case: $\mathcal{O}(D^{\frac{L}{2}-1})$. Bidirectional Search improves BFS. Each search will expand half of the optimal path to the goal before meeting in the middle, at some node at depth $L/2$ for both searches. In the best case, this node is the first one expanded at that depth for both searches, so the number of node expansions is $\mathcal{O}(D^{\frac{L}{2}-1})$ for the same reason as in part a(ii). Worst case: $\mathcal{O}(D^{\frac{L}{2}})$. In the worst case the searches both need to expand at depths up to and including $D^{\frac{L}{2}}$.

In parts (c)-(e) below, consider the following graph, with start state S and goal state G . Edge costs are labeled on the edges, and heuristic values are given by the h values next to each state.



(c) [1 pt] Greedy Graph Search

What is the path returned by greedy graph search, using the given heuristic?

- ☒ $S \rightarrow A \rightarrow G$
- ☐ $S \rightarrow A \rightarrow C \rightarrow G$
- ☐ $S \rightarrow B \rightarrow A \rightarrow C \rightarrow G$
- ☐ $S \rightarrow B \rightarrow A \rightarrow G$
- ☐ $S \rightarrow B \rightarrow C \rightarrow G$

(d) A* Graph Search

(i) [1 pt] List the nodes in the order they are expanded by A* graph search:

Order: S, A, C, B, G

(ii) [1 pt] What is the path returned by A* graph search?

- ☐ $S \rightarrow A \rightarrow G$
- ☒ $S \rightarrow A \rightarrow C \rightarrow G$
- ☐ $S \rightarrow B \rightarrow A \rightarrow C \rightarrow G$
- ☐ $S \rightarrow B \rightarrow A \rightarrow G$
- ☐ $S \rightarrow B \rightarrow C \rightarrow G$

(e) Heuristic Properties

- (i) [1 pt] Is this heuristic *admissible*? If so, mark *Already admissible*. If not, find a minimal set of nodes that would need to have their values changed to make the heuristic admissible, and mark them below.

☒ Already admissible

☐ Change $h(S)$ ☐ Change $h(A)$ ☐ Change $h(B)$

☐ Change $h(C)$ ☐ Change $h(D)$ ☐ Change $h(G)$

- (ii) [1 pt] Is this heuristic *consistent*? If so, mark *Already consistent*. If not, find the minimal set of nodes that would need to have their values changed to make the heuristic consistent, and mark them below.

☐ Already consistent

☐ Change $h(S)$ ☐ Change $h(A)$ ☒ Change $h(B)$

☐ Change $h(C)$ ☐ Change $h(D)$ ☐ Change $h(G)$


Q4. [8 pts] CSPs

Four people, A, B, C, and D, are all looking to rent space in an apartment building. There are three floors in the building, 1, 2, and 3 (where 1 is the lowest floor and 3 is the highest). Each person must be assigned to some floor, but it's ok if more than one person is living on a floor. We have the following constraints on assignments:

- A and B must not live together on the same floor.
- If A and C live on *the same* floor, they must both be living on floor 2.
- If A and C live on *different* floors, one of them must be living on floor 3.
- D must not live on the same floor as anyone else.
- D must live on a higher floor than C.

We will formulate this as a CSP, where each person has a variable and the variable values are floors.

- (a) [1 pt] Draw the edges for the constraint graph representing this problem. Use binary constraints only. You do not need to label the edges.



- (b) [2 pts] Suppose we have assigned $C = 2$. Apply forward checking to the CSP, filling in the boxes next to the values for each variable that are eliminated:

A	<input checked="" type="checkbox"/>	1	<input type="checkbox"/>	2	<input type="checkbox"/>	3
B	<input type="checkbox"/>	1	<input type="checkbox"/>	2	<input type="checkbox"/>	3
C	<input type="checkbox"/>		<input type="checkbox"/>	2		
D	<input checked="" type="checkbox"/>	1	<input checked="" type="checkbox"/>	2	<input type="checkbox"/>	3

- (c) [3 pts] Starting from the original CSP with full domains (i.e. without assigning any variables or doing the forward checking in the previous part), enforce arc consistency for the entire CSP graph, filling in the boxes next to the values that are eliminated for each variable:

A	<input checked="" type="checkbox"/>	1	<input type="checkbox"/>	2	<input type="checkbox"/>	3
B	<input type="checkbox"/>	1	<input type="checkbox"/>	2	<input type="checkbox"/>	3
C	<input type="checkbox"/>	1	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	3
D	<input checked="" type="checkbox"/>	1	<input type="checkbox"/>	2	<input type="checkbox"/>	3


- (d) [2 pts] Suppose that we were running local search with the min-conflicts algorithm for this CSP, and currently have the following variable assignments.

A	3
B	1
C	2
D	3

Which variable would be reassigned, and which value would it be reassigned to? Assume that any ties are broken alphabetically for variables and in numerical order for values.



The variable ☒ A will be assigned the new value ☐ 1
☐ B ☒ 2
☐ C ☐ 3
☐ D

Q10. [15 pts] Neural Networks: Representation





For each of the piecewise-linear functions below, mark all networks from the list above that can represent the function **exactly** on the range $x \in (-\infty, \infty)$. In the networks above, *relu* denotes the element-wise ReLU nonlinearity: $relu(z) = \max(0, z)$. The networks G_i use 1-dimensional layers, while the networks H_i have some 2-dimensional intermediate layers.

(a) [5 pts]



The networks G_3, G_4, G_5 include a ReLU nonlinearity on a scalar quantity, so it is impossible for their output to represent a non-horizontal straight line. On the other hand, H_3, H_4, H_5 have a 2-dimensional hidden layer, which allows two ReLU elements facing in opposite directions to be added together to form a straight line. The second subpart requires a bias term because the line does not pass through the origin.

(b) [5 pts]





These functions include multiple non-horizontal linear regions, so they cannot be represented by any of the networks G_i which apply ReLU no more than once to a scalar quantity.

The first subpart can be represented by any of the networks with 2-dimensional ReLU nodes. The point of nonlinearity occurs at the origin, so nonzero bias terms are not required.

The second subpart has 3 points where the slope changes, but the networks H_i only have a single 2-dimensional ReLU node. Each application of ReLU to one element can only introduce a change of slope for a single value of x .

(c) [5 pts]



Both functions have two points where the slope changes, so none of the networks $G_i; H_1, H_2$ can represent them.

An output bias term is required for the first subpart because one of the flat regions must be generated by the flat part of a ReLU function, but neither one of them is at $y = 0$.


The second subpart doesn't require a bias term at the output: it can be represented as $-\text{relu}(\frac{-x+1}{2}) - \text{relu}(x+1)$. Note how if the segment at $x > 2$ were to be extended to cross the x axis, it would cross exactly at $x = -1$, the location of the other slope change. A similar statement is true for the segment at $x < -1$.

Q11. [9 pts] Backpropagation

In this question we will perform the backward pass algorithm on the formula

$$f = \frac{1}{2} \|\mathbf{Ax}\|^2$$

Here, $\mathbf{A} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}$, $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$, $\mathbf{b} = \mathbf{Ax} = \begin{bmatrix} A_{11}x_1 + A_{12}x_2 \\ A_{21}x_1 + A_{22}x_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}$, and $f = \frac{1}{2} \|\mathbf{b}\|^2 = \frac{1}{2} (b_1^2 + b_2^2)$ is a scalar.



(a) [1 pt] Calculate the following partial derivatives of f .

(i) [1 pt] Find $\frac{\partial f}{\partial \mathbf{b}} = \begin{bmatrix} \frac{\partial f}{\partial b_1} \\ \frac{\partial f}{\partial b_2} \end{bmatrix}$.

- ☐ $\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$
☒ $\begin{bmatrix} b_1 \\ b_2 \end{bmatrix}$
☐ $\begin{bmatrix} b_2 \\ b_1 \end{bmatrix}$
☐ $\begin{bmatrix} f(b_1) \\ f(b_2) \end{bmatrix}$
☐ $\begin{bmatrix} A_{11} \\ A_{22} \end{bmatrix}$
☐ $\begin{bmatrix} b_1 + b_2 \\ b_1 - b_2 \end{bmatrix}$

(b) [3 pts] Calculate the following partial derivatives of b_1 .

(i) [1 pt] $\left(\frac{\partial b_1}{\partial A_{11}}, \frac{\partial b_1}{\partial A_{12}} \right)$

- ☐ (A_{11}, A_{12})
☐ $(0, 0)$
☐ (x_2, x_1)
☐ $(A_{11}x_1, A_{12}x_2)$
☒ (x_1, x_2)

(ii) [1 pt] $\left(\frac{\partial b_1}{\partial A_{21}}, \frac{\partial b_1}{\partial A_{22}} \right)$

- ☐ (A_{21}, A_{22})
☐ (x_1, x_2)
☐ $(1, 1)$
☒ $(0, 0)$
☐ $(A_{21}x_1, A_{22}x_2)$

(iii) [1 pt] $\left(\frac{\partial b_1}{\partial x_1}, \frac{\partial b_1}{\partial x_2} \right)$

- ☒ (A_{11}, A_{12})
☐ (A_{21}, A_{22})
☐ $(0, 0)$
☐ (b_1, b_2)
☐ $(A_{21}x_1, A_{22}x_2)$

(c) [3 pts] Calculate the following partial derivatives of f .

(i) [1 pt] $\left(\frac{\partial f}{\partial A_{11}}, \frac{\partial f}{\partial A_{12}} \right)$

- ☐ (A_{11}, A_{12})
☐ $(A_{11}b_1, A_{12}b_2)$
☐ $(A_{11}x_1, A_{12}x_2)$
☒ (x_1b_1, x_2b_1)
☐ (x_1b_2, x_2b_2)
☐ (x_1b_1, x_2b_2)

(ii) [1 pt] $\left(\frac{\partial f}{\partial A_{21}}, \frac{\partial f}{\partial A_{22}} \right)$

- ☐ (A_{21}, A_{22})
☐ $(A_{21}b_1, A_{22}b_2)$
☐ $(A_{21}x_1, A_{22}x_2)$
☐ (x_1b_1, x_2b_1)
☒ (x_1b_2, x_2b_2)
☐ (x_1b_1, x_2b_2)

(iii) [1 pt] $\left(\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2} \right)$

- ☐ $(A_{11}b_1 + A_{12}b_2, A_{21}b_1 + A_{22}b_2)$
☒ $(A_{11}b_1 + A_{21}b_2, A_{12}b_1 + A_{22}b_2)$
☐ $(A_{11}b_1 + A_{12}b_1, A_{21}b_2 + A_{22}b_2)$
☐ $(A_{11}b_1 + A_{21}b_1, A_{12}b_2 + A_{22}b_2)$

(d) [2 pts] Now we consider the general case where \mathbf{A} is an $n \times d$ matrix, and \mathbf{x} is a $d \times 1$ vector. As before, $f = \frac{1}{2} \|\mathbf{Ax}\|^2$.

(i) [1 pt] Find $\frac{\partial f}{\partial \mathbf{A}}$ in terms of \mathbf{A} and \mathbf{x} only.

- ☐ $\mathbf{x}^\top \mathbf{A}^\top \mathbf{Ax}$
☒ \mathbf{Axx}^\top
☐ $\mathbf{A}(\mathbf{A}^\top \mathbf{A})^{-1}$
☐ $\mathbf{AA}^\top \mathbf{Ax}$
☐ \mathbf{A}

(ii) [1 pt] Find $\frac{\partial f}{\partial \mathbf{x}}$ in terms of \mathbf{A} and \mathbf{x} only.

- ☐ \mathbf{x}
☐ $(\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{x}$
☐ $\mathbf{xx}^\top \mathbf{x}$
☐ $\mathbf{x}^\top \mathbf{A}^\top \mathbf{Ax}$
☒ $\mathbf{A}^\top \mathbf{Ax}$