6.390: Final Exam, Fall 2022

Do not tear exam booklet apart!

- This is a closed book exam. Two pages (8 1/2 in. by 11 in) of notes, front and back, are permitted. Calculators are not permitted.

- The total exam time is 3 hours.

- The problems are not necessarily in any order of difficulty.

- Record all your answers in the places provided. If you run out of room for an answer, continue on a blank page and mark it clearly.

- If a question seems vague or under-specified to you, make an assumption, write it down, and solve the problem given your assumption.

- If you absolutely *have* to ask a question, come to the front.

- **Write your name on every piece of paper.**

Name: _____   MIT Email: _____

Question	Points	Score
1	12	
2	10	
3	10	
4	8	
5	14	
6	18	
7	18	
8	10	
Total:	100	

Mountain Ascent

1. (12 points) Randy got back from hiking his favorite mountain. At various points of time x (in hours relative to before and after reaching the summit) he recorded his altitude y (in miles) to form a dataset $\mathcal{D}_n = \{(x^{(i)}, y^{(i)})\}_{i=1}^{n}$. The dataset is plotted below.


(a) Having learned about linear regression, he wonders if he can fit a piecewise linear function to the data. For each of the following ranges, give the formula for a linear function which perfectly fits the data in the range.

$$
y = \begin{cases}
\underline{\hspace{3cm}}, & x < -1, \\[2mm]
\underline{\hspace{3cm}}, & -1 \le x < 0, \\[2mm]
\underline{\hspace{3cm}}, & 0 \le x < 1, \\[2mm]
\underline{\hspace{3cm}}, & 1 \le x.
\end{cases}
$$

(b) Randy wonders whether it might be possible to fit the following two-layer neural network with ReLU activations to perfectly predict his altitude.

With our notation, e.g., w_2^2 denotes the second weight of the second layer, and multiplies the first layer activation from the bottom ReLU, a_2^1. Here w_2^1, w_4^1 and w_0^2 are constant offsets.

Is it possible to find weights to fit the dataset? If so, specify the weights as they appear in the diagram above. If not, explain why it is not possible using the provided neural network architecture.

○ Possible   ○ Impossible

$w_1^1 =$ $w_2^1 =$

$w_3^1 =$ $w_4^1 =$

$w_0^2 =$ $w_1^2 =$ $w_2^2 =$

(c) Randy wants to see whether PyTorch can find a good fit to the data. Can you help him remember how to express this neural network architecture in PyTorch? Select one of the following.

○ `[Linear(in_features=1, out_features=2, bias=True), ReLU(),`
`  Linear(in_features=2, out_features=2, bias=True),`
`  Softmax(dim=-1)]`

○ `[Linear(in_features=1, out_features=1, bias=True),`
`  Linear(in_features=1, out_features=1, bias=True), ReLU(), ReLU(),`
`  Linear(in_features=2, out_features=1, bias=True), ReLU()]`

○ `[Linear(in_features=1, out_features=2, bias=False), ReLU(),`
`  Linear(in_features=2, out_features=1, bias=False), ReLU()]`

○ `[Linear(in_features=1, out_features=2, bias=True), ReLU(),`
`  Linear(in_features=2, out_features=1, bias=True), ReLU()]`

○ `[Linear(in_features=1, out_features=1, bias=True), ReLU(),`
`  Linear(in_features=2, out_features=1, bias=True), ReLU()]`

(d) Randy wants to understand what PyTorch is doing behind the scenes, so he decides to manually derive gradient descent updates for each of the weight parameters. He chooses the squared-error loss function, $\mathcal{L}(y, \hat{y}) = (y - \hat{y})^2$.

He initializes the weights of the neural network as:

$$w_1^1 = 1, \ w_2^1 = 0, \ w_3^1 = 0, \ w_4^1 = 1, \ w_0^2 = 1, \ w_1^2 = 1, \ w_2^2 = 1.$$

These initial weights are shown on the figure below, in italics.
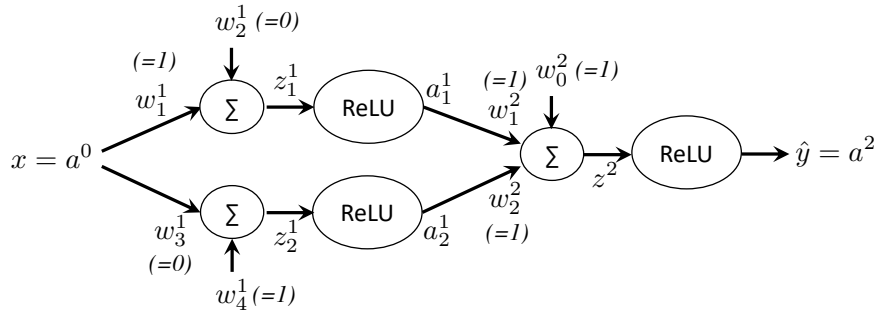
With these initialized weights, what does the network predict for each of the following data points? What is the value of the loss function evaluated at each data point?

$(x, y) = (0, 1):$ $\hat{y} =$ loss=
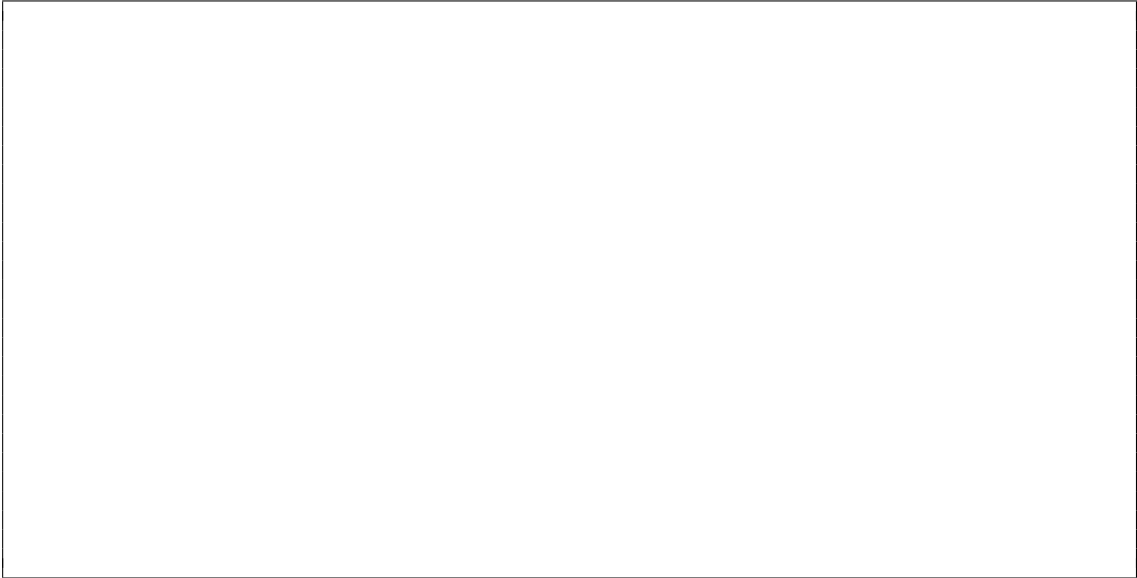
$(x, y) = (1, 0):$ $\hat{y} =$ loss=

(e) What is the gradient of the loss function with respect to the weight w_1^1? Derive the gradient in symbolic form (i.e., as a function of x, z's, etc.) using backpropagation. Then, compute the value of the gradient at the datapoint $(x, y) = (1, 0)$. Show all of your work.

$\frac{\partial \mathcal{L}(y, \hat{y})}{\partial w_1^1} =$

Evaluated at data point $(x, y) = (1, 0)$:

(f) With a step size $\eta = 1$, what is the new value of w_1^1 after one iteration of stochastic gradient descent with respect to the data point $(x, y) = (1, 0)$?

Grey Matter

2. (10 points) Grey Matter, Inc. wants to build a neural network to convert 2×2 RGB color images into greyscale. The color images have three channels: channel 0 is red, channel 1 is green, and channel 2 is blue. For each pixel, Grey Matter use the following `RGB2Grey` conversion:

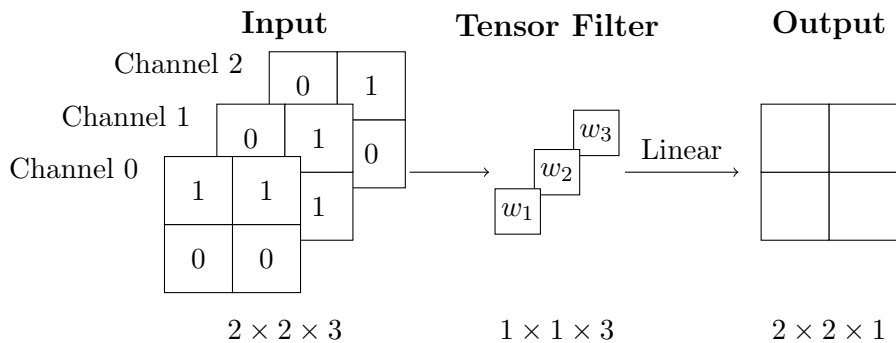$$\text{grey} = 0.30 \text{ red} + 0.59 \text{ green} + 0.11 \text{ blue}.$$

Each pixel value (in both greyscale images and color image channels) is a real number in the range $[0, 1]$.

(a) Consider the following pixel values for a 2×2 color image:

Channel 0 (red)   Channel 1 (green)   Channel 2 (blue)

1	1
0	0

0	1
0	1

0	1
1	0

What are the corresponding pixel values in the greyscale image?

(b) Blake is charged with building a convolutional neural network (CNN) to implement the `RGB2Grey` conversion. They heard a rumor that 1×1 convolutional tensor filters are useful in CNNs. They try the following architecture, where the input RGB color images have shape $2 \times 2 \times 3$, the tensor filter has shape $1 \times 1 \times 3$ with bias term $w_0 = 0$ and stride 1, a Linear activation function is used, and the output greyscale image has shape $2 \times 2 \times 1$:

Input   **Tensor Filter**   **Output**

Channel 2
Channel 1
Channel 0

0	1

0	1
	0

1	1
	1
0	0

w_3
w_2
w_1

$\xrightarrow{\text{Linear}}$

$2 \times 2 \times 3$ $1 \times 1 \times 3$ $2 \times 2 \times 1$

With hundreds of different training samples, can this CNN learn the greyscale conversion? If so, provide a set of values for the weight parameters w_1, w_2, and w_3 that implements the converter. If not, explain why the converter cannot be learned with this CNN.

(c) Cade believes that the architecture above needs more capacity to deal with a large variety of images. He replaces the 1×1 convolutional filter with a $3 \times 3 \times 3$ tensor filter with offset term $w_0 = 0$, and replaces the Linear activation with a ReLU:

Input   **Tensor Filter**   **Output**

Channel 2

| 0 | 1 |

Channel 1

| 0 | 1 |
| | 0 |

Channel 0

| 1 | 1 |
| | 1 |

| 0 | 0 |

ReLU $\longrightarrow$

$2 \times 2 \times 3$ $3 \times 3 \times 3$ $2 \times 2 \times 1$

Assume the filter is applied with stride 1. In order to obtain a $2 \times 2 \times 1$ output, is padding of the input necessary? If so, describe where padding values are needed, and what the padding values should be.

(d) With hundreds of different training samples, can Cade's CNN from part (c) learn the greyscale conversion, assuming that zero-padding has been applied if or as necessary? If so, fill in a set of weights for the $3 \times 3 \times 3$ tensor filter above. If not, explain why the grey-scale converter cannot be learned with this CNN.

Yes, or No with explanation:

If yes, fill in the tensor filter weights:

Channel 0   Channel 1   Channel 2

Rainbow Matter

3. (10 points) Grey Matter, Inc., is hoping to learn beyond just filters to spin-off a new company, now called Rainbow Matter, Ltd.!

(a) They task Cade to consult with Dana (a photographer) to conduct some market research. Cade asks Dana, "What are some of your 'go-to' moves when you edit photos? We'd like to try machine learning a model to achieve those moves!" Dana replies, "Hmm. I often need to rotate my photos."

Cade believes that a convolutional filter can be learned to rotate photos, i.e., to move the pixel values by one-quarter turn (90 degrees) clockwise in the image matrix:

Input   Desired output (after convolution)

1	2
4	3

4	1
3	2

As a "proof-of-concept," Cade tries to hand-design a two-dimensional convolutional filter to achieve the goal. To his delight, Cade finds that the following filter (with offset 0) succeeds in rotating the image above, using zero-padding of size 1 on the input, and an (unusual!) **stride of 2**.

a	b
d	c

$=$

$\frac{2}{3}$	$\frac{3}{4}$
$\frac{1}{2}$	c

What is the missing filter weight c?

[answer box]

(b) Cade is super excited; he believes he has found this magical rotational filter. Dana has serious doubts though. She says that Cade's filter is a one-time wonder, believing that it can only rotate that particular original image. She challenges Cade to demonstrate his success again, by rotating the previous output another 90 degrees using his filter.

Keeping the zero-padding and stride scheme the same as before, when the filter Cade found in part (a) is applied to the new input image, will it produce the desired output below?

New Input   Desired output (after convolution)

4	1
3	2

3	4
2	1

○ Yes   ○ No

(c) Dana's influencer friend Emily heard about the collaboration. She suggests: "the more business-savvy choice might be to learn a detector that can predict whether a filtered photo will get any 'Like's or not!" Dana was like, yeah, cool idea! She collected 1000 of her recent photos posted on social media, and extracted the Like counts.

The goal is now to predict whether a post will receive any Likes or not, (i.e., more than zero), and we need to help Cade design a complete CNN architecture.

    i. What output encoding could Cade use on top of Dana's Like counts as the output? Specifically, describe how Cade should transform Dana's Like counts to obtain the labels used for training, the number of output units, and the loss function that should be used.

    ii. Suppose that Cade has access to these six layer and activation types:

- Convolutional Layer
- Fully Connected Layer
- Max Pooling Layer

- ReLU Activation
- Sigmoid Activation
- tanh Activation

Fill in the blanks below with the preferable layer or activation to help Cade complete the architecture for his model (preferable in the computational-cost sense):

(Input Image) → Convolution layer → _____ → _____

→ Convolutional → _____ → _____

→ Fully connected → _____ → (Output)

(d) Cade grew unsatisfied with just predicting if a photo will get any Likes. He thought, wouldn't it be more fun to predict $[p_1, p_{10}, p_{100}]$, i.e., the probability that the image would receive at least one like, at least 10 likes, and at least 100 likes? He believes a simple change of his network architecture to use a softmax activation at the output layer would achieve this new goal.

Is he correct? If not, please explain what activation function would be more appropriate.

(e) Cade's friend Fish raises another question to Cade: regardless of whether the softmax is a valid idea, your CNN couldn't possibly have accurate predictions out in the real world. Specifically, it isn't possible to accurately to predict whether an image will have at least 100 likes by only looking at the image. You need more features to make the prediction!

Do you agree with Fish? What could be such a real-world feature that's missing as input, in addition to the photo itself?

(This is an open-ended question. We're only expecting you to identify a feature; you do not need to specify how to encode this feature or how to construct a model taking in this feature.)

Decisions, Decisions

4. (8 points) You have 16 data points in your training set, plotted below. These data points in the plot just below are located exactly where gridlines intersect; for instance, there is a point at $(x_1, x_2) = (1, 1)$. There are two classes: circles and dashes. You are interested in building a decision tree for classification.

(a) What is the entropy for the entire data set? Your final answer should be a single numeric value (not an equation).

(b) First, we think about making a decision tree with a single split. What is the weighted average entropy of the split in dimension 1 at location 0 for this data set? Your final answer should be a single numeric value (not an equation).

Name: _____

(c) What is the weighted average entropy of the split in dimension 1 at location 1.5 for this data set? Your final answer should be a single numeric value (not an equation).

(d) Now consider any possible split, on dimension j at location s. Which such split of this data set has the smallest weighted average entropy? If there are ties, describe all the splits with smallest weighted average entropy.

(e) More generally, suppose we take a greedy approach to growing a tree of the form depicted just below on any data set. That is, we first choose an optimal split for the root node (according to weighted average entropy). Second, we choose optimal splits for the left and right children of the root node. Is this procedure guaranteed to find the tree of the form drawn below that has the highest accuracy on the training data? Be sure to justify your answer.

(f) Now suppose a little noise is added to our data; see the plot just below. So no two data points share the exact same feature value anymore. Will the value of the smallest weighted average entropy across possible splits change from part d? Will the split with smallest weighted average entropy change? Why or why not? (Note: you do **not** need to report any numeric values for weighted average of the new, noisy data in your answer.)

(g) Now let's return to the original data. Draw the decision boundary and predicted classes for 1-nearest neighbors on the plot in the box below. (The plot is duplicated from the plot at the start.)

Game, Set, Match

5. (14 points) Rowan has been trying out various different classifiers on her data, but she forgot which is which. She hopes you can help her remember. Rowan's data has two classes: circles and dashes. Each picture below represents a binary classifier, where the decision boundary is in black. The region shaded gray predicts the dash class; the region shaded white predicts the circle class.

In the questions below, Rowan is curious which of the pictured classifiers, if any, might belong to a stated hypothesis class. You can choose one picture, multiple pictures, or no pictures. Assume the features x_1 and x_2 are used directly; there are no feature transformations prior to applying the predictor.

Unless the question refers specifically to the pictured data, you can assume Rowan is asking about *any* member of the hypothesis class, *not* just the optimal hypothesis with respect to the data.


(A)

(B)

(C)

(D)
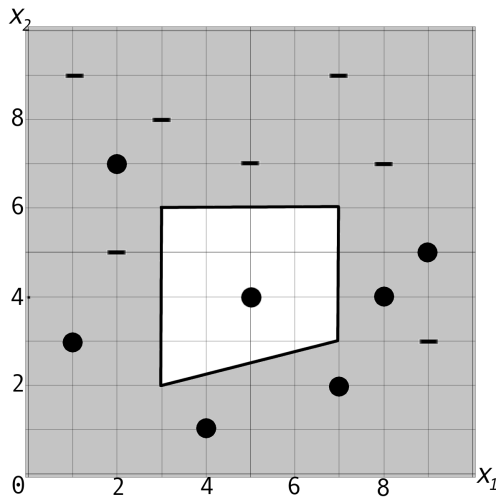
(a) Rowan thinks she might have tried a linear classifier. Identify all of the plots above which could represent such a predictor; mark with an X all that apply.

○ (A)   ○ (B)   ○ (C)   ○ (D)   ○ None

(b) For every linear classifier (A) through (D) that you identified in the previous part, provide the decision boundary (i.e., the separator) in the form $0 = \theta^\top x + \theta_0$. Identify θ and θ_0. Additionally, write the normal vector for the classifier that points to the half-space classified as positive $(+1)$, where we think of the circle class as corresponding to a positive $(+1)$ prediction.

(A) $\theta =$ $\theta_0 =$ normal vector $=$

(B) $\theta =$ $\theta_0 =$ normal vector $=$

(C) $\theta =$ $\theta_0 =$ normal vector $=$

(D) $\theta =$ $\theta_0 =$ normal vector $=$

(c) Rowan thinks she might have tried a decision tree with two splits of the form $x_j \geq s$. Identify all of the plots above which could represent such a predictor; mark with an X all that apply.

○ (A)   ○ (B)   ○ (C)   ○ (D)   ○ None

(d) For every classifier you identified in the previous part, fill in the corresponding tree. In particular, we've set up the tree structure for you below. You will need to fill in the splits of the form $x_j \geq s$, fill in the leaf values, and choose "yes" or "no" as appropriate. As a reminder, you only need to specify a tree that produces the decision boundary (if any), without worrying about whether any particular learning algorithm would find it.

(A)

(B)

(C)

(D)
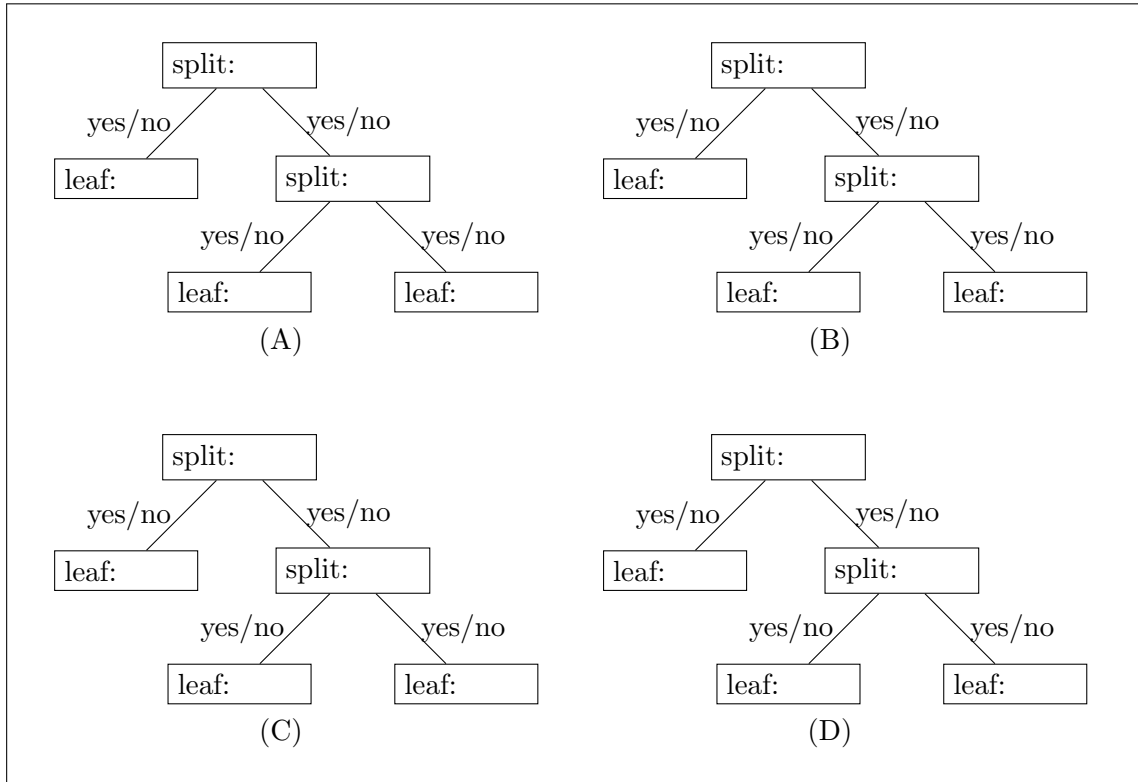
(e) Rowan thinks she might have tried a decision tree with an unconstrained number of splits of the form $x_j \geq s$. Identify all of the plots above which could represent such a predictor; mark with an X all that apply.

○ (A)   ○ (B)   ○ (C)   ○ (D)   ○ None

Name: _____

(f) Rowan thinks she might have tried a 2-layer neural network with two hidden neurons and one output neuron, where each neuron uses the ReLU activation function. That is, the first layer takes the two features (x_1 and x_2) as inputs and has two outputs; the second layer takes the two outputs from the first layer as inputs and has one output (the final predicted class). Identify all of the plots above which could represent such a predictor; mark with an X all that apply.

○ (A)   ○ (B)   ○ (C)   ○ (D)   ○ None

(g) Rowan thinks she might have tried a 2-layer neural network with an unconstrained number of hidden neurons and one output neuron, where each neuron uses the ReLU activation function. Identify all of the plots above which could represent such a predictor; mark with an X all that apply.

○ (A)   ○ (B)   ○ (C)   ○ (D)   ○ None

(h) Rowan thinks she might have tried a 1-nearest neighbors classifier using Euclidean distance on the full set of pictured data. (Unlike the previous questions, this question really does depend on Rowan's full data set, shown in the picture.) Identify all of the plots above which could represent such a predictor; mark with an X all that apply.

○ (A)   ○ (B)   ○ (C)   ○ (D)   ○ None

Page 18

Word

6. (18 points) "Word" is a word game where one tries to guess a three-letter word in four guesses, where each guess must be a three-letter word. In our game, we will use five common letters in the English language $\{A, E, R, S, T\}$.

   (a) We will consider all valid three-letter English words as character sequences $c = [c_0, c_1, c_2]$. For the letters $\{A, E, R, S, T\}$, we will use the following encoding function ϕ:

$$\phi(A) = [1, 0, 0, 0, 0],$$
$$\phi(E) = [0, 1, 0, 0, 0],$$
$$\phi(R) = [0, 0, 1, 0, 0],$$
$$\phi(S) = [0, 0, 0, 1, 0],$$
$$\phi(T) = [0, 0, 0, 0, 1].$$

   Consider the character sequence $[E, A, R]$. Provide the encoding of this sequence, i.e., $[\phi(E), \phi(A), \phi(R)]$.

   $[\phi(E), \phi(A), \phi(R)] =$

   (b) We want to build a model to help us play the game better. We want the model to take a game state (which we will encode later in the problem) as input, and to predict the probability of all letters at all positions in the word. How should we encode the output and what is the dimension of the output?

(c) Which (potentially nonlinear) output function should we use to build this output, and which sets of dimensions should this output be applied to?

(d) We indicate the probability for letter i in position j as $a_j^{(i)}$, where $j \in \{1, 2, 3\}$. Give an example output for the word ERA with the probability for the first position $a_1^{(E)} = 1$; for the second position $a_2^{(R)} = 0.9$ and $a_2^{(T)} = 0.1$; and for the third position $a_3^{(A)} = 0.9$ and $a_3^{(T)} = 0.1$.

(e) Suppose we know that the first letter is an E, the last two letters are not an E, S, or T, and we want the output to indicate 80% probability that the word is EAR and 20% probability that the word is ERA (and zero probability that it is any other word, like ERR). Can we represent that output with the representation above? If so, show the output. If not, explain why not.

(f) Rather than having our model making predictions for all the letters at once, we instead decide to use a recurrent neural network (RNN) to output its prediction as a sequence of letters. We will use the all-zeros encoding $\phi('-') = [0,0,0,0,0]$ for the token $'-'$ (dash), a special character that indicates no letter. Our RNN has the following design made of three identical blocks:

1. The RNN first takes as input the game state s and $x_0 = [\phi('-'), \phi('-')]$ and outputs the probability that each letter appears in the first position, a_1. We then sample the first letter with the likelihood given by the probabilities a_1 to produce our guess of the first letter, g_1.

2. To the second block, we pass in the same game state s from before and $x_1 = [\phi(g_1), \phi('-')]$, with the guess g_1 from the previous block and a blank second letter. The second block outputs the probability that each letter appears in the second position, a_2. We then sample the second letter with the likelihood given by the probabilities a_2 to produce our guess g_2.

3. Finally we pass the same game state s and $x_2 = [\phi(g_1), \phi(g_2)]$ to the final block. The third block predicts g_3, the third letter, with the likelihood given by the probabilities a_3.

Our final prediction is the sequence of letters $[g_1, g_2, g_3]$. Note that while in the diagram we are using the letters as input, we actually would use the encoding ϕ from part (a).
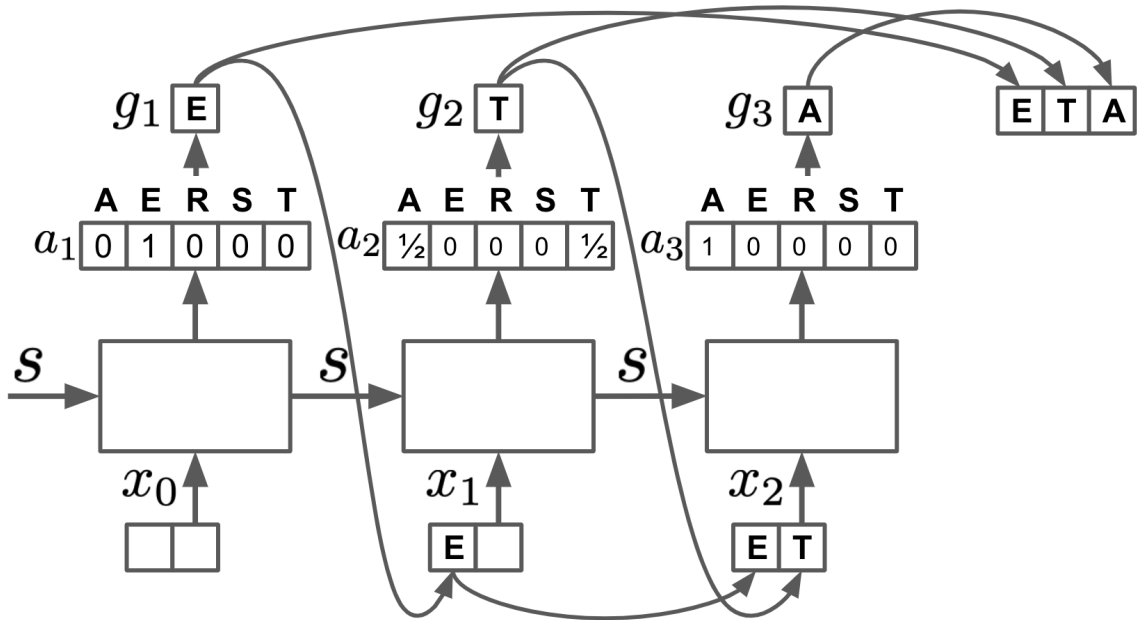
i. For every guess, the game will return the game state s, which reflects what we know about the word based on our previous guesses. For each position $j \in \{1, 2, 3\}$ in the word, we use the following encoding for game state at that position, s_j:

  1. $[0, 0]$ for when the letter is not in this position in the word (including not in the word at all),

  2. $[0, 1]$ for when we don't know if the letter is in this position or not,

  3. $[1, 1]$ for when the letter is in the word in this position.

  What is the name for this type of encoding?

ii. The game has given us the following information from our previous guesses:

  1. A is not the second letter (not in the second position in the word);

  2. E is the first letter (the first position in the word);

  3. R is not the third letter (not in the third position in the word);

  4. S is not in the word;

  5. and our guesses have not generated additional information about T.

  The encoding of the game state for the first position (first letter) in the word is

  $$s_1 = [[0, 0], [1, 1], [0, 0], [0, 0], [0, 0]]$$

  where each element in the list is the encoding corresponding to knowledge about letter A, E, R, S, and T, in sequence. The game state for the second letter in the word is:

  $$s_2 = [[0, 0], [0, 1], [0, 1], [0, 0], [0, 1]] .$$

  Using the encoding above, give the encoding of the game state for the third position in the word:

  $s_3 =$

iii. Recall the output activation function from part (c). If we want the predicted probability of a given letter to be (very close to) zero, what value should the pre-activated input to the output activation function be?

iv. We want to guarantee that the model will not predict letters that we know from the game state s cannot be at the given position. To do this, we will apply the following neural network to the pre-activations z_j of the output a_j, where j indicates the j^{th} position. The pre-activations take into consideration x_{j-1}, and this neural network accounts for s.

```
# s_j : game state at position j shape [5, 2] -> [letter, encoding]
# z_j : pre-activation of a_j, shape [5]
# w : shape [1], b : shape [1]
s_j_summed = s_j.sum(axis=1)   # shape [5]
condition, if_true, if_false = s_j_summed < 1, s_j_summed, 1
s_where = np.where(condition, if_true, if_false)   # shape [5]
c_j = f_1(w * s_where + b)
a_j = f_2(z_j - c_j)
```

Assume `w` and `b` are forced to be values between -10 and 10, and `f_2` is the output activation function. Specify what the function `f_1` should be, and values for `w` and `b`, such that `c_j` is nonzero when `s_j[i]` `==` `[0, 0]` (the state for the i^{th} letter in the j^{th} position) for all i but zero when `s_j[i]` is `[0, 1]`, or `[1, 1]` for any i. Note that the value `c_j` is subtracted from `z_j` and keep in mind your answer to part (iii).

`f_1` =

`w` =

`b` =

Markovian Spins

7. (18 points) Pauli is in the lab measuring the direction of the spin of a qubit. We're going to help Pauli model their observations. The spin of the qubit can be in the state up ↑ or down ↓ along specific directions x or z, denoted as ↑ x, ↓ x, ↑ z, and ↓ z.

Pauli can apply one of two measurements, M_x or M_z, to observe the spin along the directions x or z.

Pauli knows that qubits obey the following physics. If Pauli applies measurement...

1. M_x on ↑ x or ↓ x, they get back the same state, ↑ x or ↓ x, respectively.

2. M_z on ↑ z or ↓ z, they get back the same state, ↑ z or ↓ z, respectively.

3. M_x on state ↑ z or ↓ z, they observe a transition to ↑ x with 50% likelihood or to ↓ x with 50% likelihood.

4. M_z on state ↑ x or ↓ x, they observe a transition to ↑ z with 50% likelihood or ↓ z with 50% likelihood.

We would like to model Pauli's measurements of the qubit as a Markov Decision Process. To do this, we must define our states, actions, and transitions between states for each action.

(a) What are the state space $\mathcal{S}$ and action space $\mathcal{A}$?

$$\mathcal{S} =$$

$$\mathcal{A} =$$

(b) Draw a diagram showing the transitions between states via actions M_x. Be sure to write the probability of transitioning between states.

M_x (↑ x)   (↓ x)

(↑ z)   (↓ z)

(c) Draw a diagram showing the transitions between states via actions M_z. Be sure to write the probability of transitioning between states.

M_z (↑ x) (↓ x)

(↑ z) (↓ z)

(d) We want to calculate the value of the following policy π_0:

- $\pi_0(\uparrow x) = \pi_0(\downarrow x) = \pi_0(\uparrow z) = M_z$
- $\pi_0(\downarrow z) = M_x$

Suppose we get a reward of 100 for measuring the state $\uparrow z$ with M_z, $R(\uparrow z, M_z) = 100$, and the reward of all other state-action pairs is 0. Calculate the value of the policy π_0 for horizon of 0 and horizon of 1 starting in $\uparrow z$ or $\downarrow z$. Where appropriate, first express your answer in terms reward function $R(s, a)$ for the appropriate state s and action a determined by the policy π_0, before reducing the answer to the numerical values.

$V_{\pi_0}^0(\uparrow z) =$

$V_{\pi_0}^1(\uparrow z) =$

$V_{\pi_0}^0(\downarrow z) =$

$V_{\pi_0}^1(\downarrow z) =$

(e) Calculate the value of the policy for horizon of 2, starting from $\uparrow x$. Assume a discount factor of γ, meaning leave your answers in terms of γ. Where appropriate, first express your answer in terms reward function $R(s, a)$ and transition model $T(s, a, s')$ for the appropriate state s and s' and action a determined by the policy π_0, before reducing the answer to the numerical values.

$V_{\pi_0}^2(\uparrow x) =$

(f) We turn on a magnetic field pointing along the $\uparrow x$ direction, and this changes what transitions occur when we make a measurement, in the following ways.

1. If we are in state $\uparrow z$ or $\downarrow z$ and measure with M_x, we are 80% likely to transition to $\uparrow x$ and 20% likely to transition to $\downarrow x$.
2. If we are in state $\downarrow x$ and measure with M_x, we are 30% likely to stay $\downarrow x$ and 70% likely to transition to $\uparrow x$.
3. All other transitions stay the same.

Draw a diagram showing the new transitions between states via actions M_x. Be sure to write the probability of transitioning between states.

M_x $\quad$ $(\uparrow x)$ $\qquad$ $(\downarrow x)$

$\qquad\quad$ $(\uparrow z)$ $\qquad$ $(\downarrow z)$

(g) We continue with the magnetic field turned on. Suppose we get a reward of 100 for measuring the state $\uparrow z$ with M_z and now we also get a reward of 10 for measuring the state $\uparrow x$ with M_x, i.e., $R(\uparrow z, M_z) = 100$ and $R(\uparrow x, M_x) = 10$, respectively. Does the value of the policy π_0 for horizon of 2 starting in $\uparrow x$ change from part (e)? Why or why not? If it does change, please calculate it. Assume a discount factor of γ, meaning leave your answer in terms of γ. Where appropriate, first express your answer in terms reward function $R(s, a)$ for the appropriate state s and action a determined by the policy π_0, before reducing the answer to the numerical values.

Change or not (and explanation)?

$V_{\pi_0}^2(\uparrow x) =$

(h) We want to calculate the value of a new policy π_1:

- $\pi_1(\uparrow x) = \pi_1(\downarrow x) = \pi_1(\downarrow z) = M_x$
- $\pi_1(\uparrow z) = M_z$

We continue with the rewards and transition probabilities from part (g) above.

If we start in $\uparrow x$, what is value of new policy π_1 with a horizon of 2? Again, assume a discount factor of γ, meaning leave your answer in terms of γ. Where appropriate, first express your answer in terms reward function $R(s, a)$ for the appropriate state s and action a determined by the policy π_1, before reducing the answer to the numerical values.

$V_{\pi_1}^2(\uparrow x) =$

(i) To maximize expected reward, for what range of discount factors γ should Pauli choose policy π_1 instead of π_0, for horizon 2 measurement of state $\uparrow x$?

Griddy-Greedy World

8. (10 points) We're riding a bike in a 5-by-5 grid-world. We always start at the square with our bicycle, (0,1). The shaded squares (1,2), (1,3), and (3,2) are pillars. Four possible actions, North, South, East and West, allow us to navigate this world (the transitions will be detailed in the parts below).

The state's rewards are written on the squares, regardless of the action taken. For instance, we receive a reward of −10 for any action taken in state (0,0). Those squares without numbers written on them all incur 0 reward for any actions taken in them.

Squares (2,2) and (4,2) are special. When we reach either square, we get to make one final action, collect a 1 reward if we were in (2,2), or a 10 reward if we were in (4,2), and the game terminates. Our goal is to maximize the total reward we collect along the path we choose.

(a) In this part, we assume a discount factor of 1.0 and deterministic transitions. If our action would take us out of this world or into a pillar, we stay where we are. For instance, in state $(0,1)$, a North action takes us to state $(0,2)$; whereas in state $(1,1)$, a North action leaves us in state $(1,1)$.

Suppose we run the value iteration algorithm for 1000 iterations. What would be the optimal value of our initial state $(0,1)$? Is the corresponding best path unique (unique up to but not including the last action chosen)? Draw out the actions as arrows taken at the states along your best (or one of the best) path(s).

Value of best path:

Unique?

Draw best path:

(b) In this part, we keep all of the setup from the previous part, except that we change the discount factor to 0.1.

Suppose we run the value iteration algorithm for 1000 iterations. What would be the optimal value of our initial state (0,1)? Is the best path unique (unique up to but not including the last action chosen)? Draw out the actions as arrows taken at the states along your best (or one of the best) path(s).

Value of best path:

Unique?

Draw best path:

(c) In this part, we now assume a discount factor of 0.99 (close to 1). As before, if our action would take us out of this world or into a pillar, we deterministically stay where we are.

However, now all other transitions are no longer deterministic; instead, they all have a "noisy transition" probability. Specifically, with probability 0.5, the next state is the "intended" target state (i.e, the state in the direction indicated by the action). However, with probability 0.5, the next state will instead be one of the horizontal or vertical reachable neighbors of the target state (with equal probability).

As a concrete example, in state $(0, 1)$, an East action takes us to state $(1, 1)$ with probability 0.5, or to one of the other three reachable immediate neighbors of $(1, 1)$ (which are $(0, 1)$, $(1, 0)$, and $(2, 1)$ but not $(1, 2)$ because that is a pillar), each with probability $0.5 \cdot \frac{1}{3}$.

Would we still be able to run the value iteration algorithm? If yes, suppose we run the value iteration algorithm for 1000 iterations. What would be the best path (or one of the best paths) with respect to the expected value? Draw out the path.

If we cannot run the value iteration algorithm, explain what variable(s) might be missing that prevents us from running value iteration in this part.

Is it possible to run the value iteration algorithm? Circle one: Yes   No
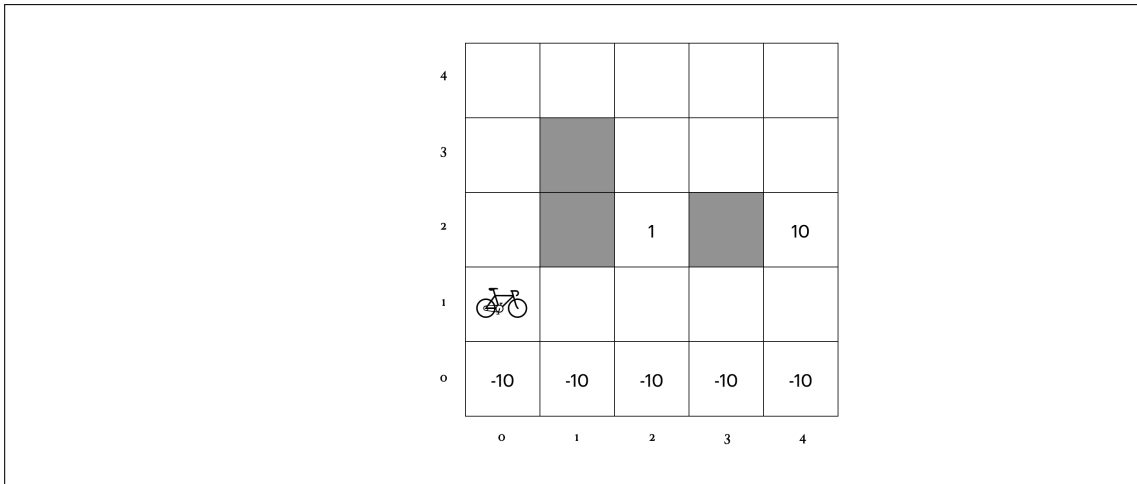
If yes, draw the best path:

If no, explain your reasoning:

(d) In this part, we keep all of the setup from part (c), including noisy transitions. We will try reinforcement learning using Q-learning. We use a learning rate of 0.2 and "greedy" action selection (that is, we always pick action with best Q value). In case of action selection ties, we randomly select from the tied actions.

We run the Q-learning algorithm for 10000 episodes. Each episode starts in the square (0,1), and we take actions until the game in that episode terminates.

Draw out the best path learned by our Q-learning algorithm.

y \ x	0	1	2	3	4
4					
3		(gray)			
2		(gray)	1	(gray)	10
1	🚲				
0	-10	-10	-10	-10	-10

Work space

Work space