

Final Exam

600.464/664 Artificial Intelligence
Spring 2025

Name:

Instructions

- Please be sure to write your name in the space above!
- Please be sure to read through the entire exam before you start, and be mindful of the time. If one question is taking too long, it may be worth moving on and coming back to the problem question(s) after the rest of the exam is complete.
- Remember that you are only allowed one sheet (both sides) of notes, everything else besides that and the test itself must be off of your workspace.
- Please show ALL relevant work for your answers and provide explanations when prompted. Failure to do either will result in loss of credit.

1. Probabilistic Reasoning

10 points

A group of 10 crewmates is playing a game of Among Us. Out of the 10 players:

- 2 are imposters
- 8 are crewmates

Each round, a random player is selected to fix a wiring task. Crewmates always complete the task successfully. There is a 50% chance that impostors complete the task successfully. After the task, the system announces whether it was "successfully completed."

Suppose a player named Red is selected for the wiring task, and the system announces the task was "successfully completed."

Let I be the event that Red is an imposter and S be the event that the task was successfully completed.

1. (2 points) What is the prior probability that Red is an imposter?

$$P(I) = \frac{2}{2+8} = .2$$

2. (6 points) What is the probability that Red is an imposter given that the task was successfully completed? Show your work.

$$P(I | S) = \frac{P(S|I) \cdot P(I)}{P(S)}$$

$$P(S) = P(S | I) \cdot P(I) + P(S | \neg I) \cdot P(\neg I)$$

$$P(S) = 0.5 \cdot 0.2 + 1 \cdot 0.8 = 0.1 + 0.8 = 0.9$$

$$P(I | S) = \frac{0.5 \cdot 0.2}{0.9} = \frac{0.1}{0.9} \approx 0.111$$

3. (2 points) Are the events I (Red is an imposter) and S (task successfully completed) independent? Justify your answer mathematically.

$$P(I) = 0.2, \quad P(I | S) \approx 0.111, \text{ so they are not independent}$$

2. Bayesian Network

15 points

A smart kitchen robot named ChefBot prepares meals. Whether the robot successfully makes a tasty dish depends on multiple factors. The robot has sensors and a Bayesian network that tracks the following variables:

Variables:


- F = Fresh ingredients (True/False)
- P = Proper preheating of the oven (True/False)
- T = Tasty dish (True/False)
- H = Human approval (True/False)

Structure of the Bayesian Network:

- Fresh ingredients and proper preheating influence whether the dish is tasty: $F \rightarrow T \leftarrow P$
- Human approval depends only on whether the dish is tasty: $T \rightarrow H$

Answer the following questions

1. (6 points) Draw the structure of this Bayesian network (nodes and arrows).



2. (6 points) Identify two pairs of variables that are conditionally independent, and specify the conditioning variable(s).

F and H are conditionally independent given T

P and H are conditionally independent given T

3. (3 points) Suppose you observe that the dish is not tasty. Does learning that the oven was properly preheated make it more or less likely that the ingredients were fresh? Explain this intuitively — no formal proof is required.

If we observe $T = \text{False}$, we suspect at least one of the causes (F, P) might have failed. If we then learn $P = \text{True}$, that increases the suspicion that $F = \text{False}$ (i.e., ingredients were not fresh).

3. Deep Learning

25 points

Consider a neural network that has a 2-dimensional input vector x , a 2-dimensional hidden vector h , and a 1-dimensional output vector y . The layers are fully connected by weight matrices W_1 and W_2 , no bias terms, and a ReLU activation function.

The output vector is used for a classification task. Samples fall either into the first class $y = (1)$ or the second class $y = (0)$. We use the L2 norm as loss.

Recall:
$$\text{ReLU}(x) = \begin{cases} x & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \quad \text{L2} = \frac{1}{2}(t - y)^2$$

1. (5 points) Draw the computation graph for this network (including the loss function).

$$x + W_1 \rightarrow \text{PROD} \rightarrow \text{ReLU} + W_2 \rightarrow \text{PROD} \rightarrow \text{ReLU} \rightarrow \text{L2}$$

2. (5 points) Given the weight matrices $W_1 = \begin{bmatrix} -1 & 0 \\ 1 & 1 \end{bmatrix}$, $W_2 = \begin{bmatrix} -1 & 1 \end{bmatrix}$ process the sample $x = (1 \ 2)^T$.

$$h = \text{ReLU}(-1 \ 3) = (0 \ 3), y = \text{ReLU}(3) = (3)$$

3. (5 points) Compute the derivatives for all functions in your computation graph.

$$\frac{\partial \text{L2}}{\partial y} = y - t$$

$$\frac{\partial \text{ReLU}}{\partial s} = \begin{cases} 1 & \text{if } s \geq 0 \\ 0 & \text{if } s < 0 \end{cases}$$

$$\frac{\partial \text{PROD}}{\partial w_i} = h_i, \frac{\partial \text{PROD}}{\partial h_i} = w_i$$

4. (5 points) If the target value for this sample is $t = (1)$, what are the gradient values back-propagated to W_2 (no need to go further than that)?

$$\text{gradient for L2 node: } y - t = 3 - 1 = 2,$$

$$\text{gradient for ReLU node: } 1 \times 2 = 2,$$

$$\text{gradient for PROD node (towards } W_2): (0 \ 3) \times 2 = (0 \ 6)$$


5. (5 points) We want to expand this task to three classes and predict a 1-hot vector, i.e., $(1 \ 0 \ 0)$ for Class 1, $(0 \ 1 \ 0)$ for Class 2, and $(0 \ 0 \ 1)$ for Class 3. How would you change activation functions and the loss function?

The final activation function should be softmax. The loss function could be cross-entropy loss, i.e., log of the score assign to the correct class element in the vector.

4. Reinforcement Learning

20 points

Consider the *deterministic* reinforcement environment drawn below. Let $\gamma = 1$. *Immediate* rewards are indicated at the nodes (the number after R). The current state of the Q table is indicated on the arcs. Once the agent reaches the **end** state the current episode ends.



Recall the Q-learning update rule:

$$\Delta Q(s, a) = \alpha \left(R(s') + \gamma \cdot \max_{a'} Q(s', a') - Q(s, a) \right)$$

- (4 points) Assuming our RL agent exploits its policy, what is the path it will take from start to goal?
start, y, w, goal
- (12 points) Following a random path selection, the agent takes the path (start, x, z, goal). Assume a learning rate of $\alpha = 1$. How are the Q values updated?
 - start $\rightarrow$ x
Update: $2 + 2 - 0 = 4$, new value: $0 + 4 = 4$
 - x $\rightarrow$ z
Update: $-5 + 4 - (-3) = 2$, new value: $-3 + 2 = -1$
 - z $\rightarrow$ goal
Update: $3 + 0 - 4 = -1$, new value: $4 + (-1) = 3$
- (4 points) After the updates in part (2), what action sequence will the agent choose from the start state if it exploits its updated Q-values?

start, x, y, w, goal

5. Generative Artificial Intelligence

20 points

1. (10 points) In the Vision Model lecture we discussed autoencoders that compress an image into a single "latent" vector. Design an autoencoder that compresses natural language sentences into a single vector using recurrent neural networks. Describe the network architecture and how would you train it.

The recurrent neural network first walks through a sentence, and then a copy of the sentence. The hidden state between the two copies of the sentence acts as latent vector that encodes the whole sentence. It would be trained on predicting the copy of the sentence from the initial sentence.

2. (10 points) Imagine you are building a Generative Adversarial Network (GAN) to generate images of hand-written digits that look like those from the MNIST dataset. The generator G takes in random noise and outputs an image; the discriminator D takes in an image and outputs a score indicating whether it thinks the image is real or fake.

Describe in your own words how you would train the generator and the discriminator. Include in your answer:

- What kind of data each network sees
- What the objective function for each network is
- How they influence each other's training

The discriminator is trained to distinguish between real MNIST digits and fake digits produced by the generator. It gets a mix of real and generated images and learns to output high scores for real images and low scores for fake ones. The objective function is the prediction error of the label false/true.

The generator is trained to fool the discriminator. It generates images from random noise and learns to make them more realistic so that the discriminator gives them higher scores. Its objective function is a mix of the original learning objective (e.g., reconstruction loss) and the opposite of the discriminator's objective (fooling the discriminator).

Over time, the generator gets better at producing realistic images, and the discriminator gets better at telling them apart — until ideally, the generated images are so good that the discriminator can't tell the difference.