

Final Exam

600.464/664 Artificial Intelligence
Spring 2024

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.

Probabilistic Reasoning

20 points

When surveying adults in the United States, on average we find the following probabilities.

	Prefers Blue	Prefers Pink
Is Tall	.32	.08
Is not Tall	.38	.22

1. (5 points) Show mathematically if color preference and being tall are independent or not independent of each other.


2. (5 points) A more detailed survey adds a question about gender. It finds the following probabilities.

	Prefers Blue / Is Tall	Prefers Pink / Is Tall	Prefers Blue / Is not Tall	Prefers Pink / Is not Tall
Female	.05	.05	.20	.20
Male	.27	.03	.18	.02

Show mathematically if color preference and being tall are conditionally independent given gender?

3. (10 points) Draw a Bayesian network that models these probabilities. Start the network with gender and have the model reflect any conditional independence you have found.

The following data samples are drawn from two Gaussian distributions. The positive data points from a distribution centered in (1,1), the negative examples from a distribution centered in (-1,-1).



1. (5 points) How would a trained linear classifier divide up the space? Sketch an estimate of the decision boundary into the graph.
2. (5 points) How would a powerful classifier that is able to classify all training examples correctly divide up the space? Sketch an estimate of the decision boundary into the graph.
3. (10 point) Which model will likely do better on new unseen test samples? Explain how this example illustrates the concepts: *Overfitting* and *Ockham's razor*.

Deep Learning

20 points

Consider the deep learning model defined by the following equations.

Input values: x_1, x_2 $s_1 = x_1 \times w_1$

State variables: s_i $s_2 = x_2 \times w_2$

Trainable parameters: w_i $s_3 = s_1 + s_2$

Computed value: y $s_4 = s_3 \times s_3$

Correct output value: t $y = s_4 \times w_3$


Loss: $\frac{1}{2}(t - y)^2$

1. (5 points) Draw a computation graph for this model, using similar notation to the one we used in class.
2. (5 points) If the input is $(x_1, x_2) = (2, -1)$ and all initial weights are 1, and the correct output value is 10, what is the loss?
3. (5 points) Add formulas to the graph that allow backpropagation to update the weights w_i .
4. (5 points) With a learning rate of 0.01, what is the value of updated weight w_3 after carrying out gradient descent training with this example?

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 update formula for Q-Learning (simplified from the class lecture).

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

- (4 points) Assuming our RL agent exploits its policy, what is the path it will take from start to end?
- (12 points) Following a random path selection, the agent RL takes the path (start, a, c, end). Assume a learning rate of $\alpha = 1$. How are the Q values updated?
 - start \rightarrow a
 - a \rightarrow c
 - c \rightarrow end
- (4 points) With sufficient number of random paths, will the policy learned with Q-learning converge to optimal values?

Generative Artificial Intelligence

20 points

You want to apply the idea of diffusion to language models. For the purpose of this question, we treat the problem of generative language models to generate a sentence of English, e.g., *My cat likes to jump on the table*.

1. (5 points) Define the process that generates training data for this approach. What would this process do with the sentence *My cat likes to jump on the table*? Give an example training item generated from it.
2. (10 point) We want to use a decoder-only Transformer model that learns from the generated training data. Draw a diagram and briefly describe the architecture of this model (no need for formulas or detailed steps such as layer normalization but describe inputs, outputs and intermediate representations).
3. (5 points) What is the key difference between an autoencoder and a variational autoencoder?