

MIDTERM EXAM REVIEW

10-701: INTRODUCTION TO MACHINE LEARNING

Spring 2026

1 Decision Trees

1. Suppose you have a training dataset consisting of 3 boolean features X_1, X_2 , and X_3 where $X_i \in \{0, 1\}$. Furthermore, assume the label is defined as $Y = X_1 \vee X_2$ i.e., $Y = 1$ if $X_1 = 1$ or $X_2 = 1$ and $Y = 0$ otherwise. Suppose that your training dataset contains all of the 8 possible feature vectors:

X_1	X_2	X_3	Y
0	0	0	0
1	0	0	1
0	1	0	1
1	1	0	1
0	0	1	0
1	0	1	1
0	1	1	1
1	1	1	1

- (a) What is the training error rate of a depth-0 decision tree on this dataset?

- (b) If your splitting criterion is training error rate, which feature or features would you split on first? Briefly justify your answer.

- (c) If your splitting criterion is information gain, which feature or features would you split on first? Briefly justify your answer.

2. You are given a dataset for binary classification with two features x_1 and x_2 . x_1 can have two possible values 0,1 and x_2 can have three possible values 0,1,2. Figure 1 provides a depiction of this dataset.




Figure 1: Binary classification dataset.

- (a) What is the lowest possible training error rate any decision tree could achieve on this dataset? The tree need not be a binary tree but each node should split on at most one feature.

- (b) Draw a decision tree that achieves this training error rate.

3. We would like to learn a decision tree. We have n samples for training. You can assume that n is large and we are using continuous features. In the following questions, suppose we only use the first feature, x_1 .

(a) We would like to split according to x_1 at the root with 3 branches: data points are split at the root to three different subtrees by findings values a and b such that the three subtrees are $x_1 \leq a$, $a < x_1 < b$ and $x_1 \geq b$. What is the runtime for finding the values of x_1 that should be used by the root for such a split?

$O(\log n)$

$O(\sqrt{n})$

$O(n)$

$O(n^2)$

$O(n^3)$

(b) Following our first split, we would like to *split again* on x_1 in each of the three sub-trees resulting from the split in the previous question. Again, for each subtree we would split three ways as we did in the root. What is the total runtime for determining the optimal splits for ALL three subtrees?

$O(n)$

$O(n^2)$

$O(n^3)$

$O(n^4)$

$O(n^8)$

(c) For the same dataset we would like to learn the *optimal* tree that:

- Only uses x_1
- Has two levels, a root and a level below it where each splits to three branches (see the figure below).



What is the runtime for computing the *optimal tree*, in terms of minimizing the training error rate?

- $O(n)$
- $O(n^2)$
- $O(n^4)$
- $O(n^8)$
- $O(n^9)$

2 kNN

1. **True or False:** Regardless of the training dataset size, the true error rate of a 1-NN model will never exceed 2 * the true error rate of the optimal classifier on a binary classification problem with stochastic labels.

True

False


2. Suppose you are deciding between the following k NN models for a binary classification problem where $y \in \{-1, +1\}$ (for both models, assume we use the Euclidean distance):

- Model 1: Probabilistic 3NN. This model predicts the label of a query data point with probabilities corresponding to the relative proportion of labels in the 3 nearest neighbors, e.g., if the 3 nearest neighbors have labels $\{-1, -1, +1\}$ then this model predicts -1 with probability $2/3$ and $+1$ with probability $1/3$.
- Model 2: Distance weighted 2NN. This model predicts

$$\hat{y} = \text{sign} \left(\sum_{(\vec{x}^{(i)}, y^{(i)}) \in N_2(\vec{x})} \frac{y^{(i)}}{\|\vec{x}^{(i)} - \vec{x}\|_2^2 + 1} \right)$$

where $N_2(\vec{x})$ are the 2 nearest neighbors in the training data to the query data point \vec{x} (the plus one in the denominator is to avoid division by 0). For the purposes of this model, let $\text{sign}(0) = +1$.

You gather a training dataset of 6 points: 3 red $-$'s, which correspond to label $y = -1$, and 3 black $+$'s, which correspond to label $y = +1$. You also gather a validation dataset consisting of points A and B, both of which have label $+1$.



- (a) What is the *expected training* error rate for Model 1? You may express your answer as a fraction if necessary.

- (b) What is the *training* error rate for Model 2? You may express your answer as a fraction if necessary.

- (c) What is the *expected validation* error rate for Model 1? You may express your answer as a fraction if necessary.

- (d) What is the *validation* error rate for Model 2? You may express your answer as a fraction if necessary.

3. For the points shown in Figure 2 suppose x_1, x_4, x_5, x_7 , and x_9 have label $+1$, and the other points have label -1 . For a 3NN classifier using the Euclidean distance, what is the LOOCV error?




Figure 2

3 Linear Regression & Regularization

1. Using monthly stock averages, x_1, x_2, \dots, x_n , you run linear regression without regularization to estimate some future stock value y . However, your test error rate turns out to be very high. Your friend suggests that maybe the price *differences* between consecutive months may offer better features, so you add additional features $(x_1 - x_2), (x_2 - x_3), (x_3 - x_4), \dots, (x_{n-1} - x_n)$. Do you expect an improvement? Briefly justify your answer.

2. Consider the following linear regression model:

$$\mathbf{y} = \mathbf{X}\mathbf{w} + \boldsymbol{\epsilon}$$

where we assume the residuals are normal and i.i.d.: $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. Recall from lecture that the MLE of \mathbf{w} is $\hat{\mathbf{w}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$. Using the fact that $\hat{\mathbf{w}}$ is an unbiased estimator of \mathbf{w} , derive an expression for the variance of $\text{var}(\hat{\mathbf{w}})$:

4 MLE/MAP

1. **True or False:** Recall that the MLE of a parameter θ given some dataset \mathcal{D} is

$$\hat{\theta}_{MLE} = \arg \max_{\theta} p(\mathcal{D}|\theta)$$

while the MAP estimate of θ is

$$\hat{\theta}_{MAP} = \arg \max_{\theta} p(\mathcal{D}|\theta)p(\theta).$$

For every dataset, there will always exist some prior distribution $p(\theta)$ for which the MAP estimate is equivalent to the MLE.

- True
 False

2. The probability density function of the Pareto distribution is

$$p(x|k, \alpha) = \begin{cases} \frac{\alpha k^\alpha}{x^{\alpha+1}} & x \in [k, \infty) \\ 0 & \text{otherwise} \end{cases}$$

$k, \alpha \in (0, \infty)$

Given n independent samples x_1, x_2, \dots, x_n drawn from a Pareto distribution, what is the MLE for the parameters k and α ? **Hint: first determine the MLE of k , then use that result to in your expression for the MLE of α .**

$$\hat{k}_{MLE} = \underline{\hspace{2cm}}$$

$$\hat{\alpha}_{MLE} = \underline{\hspace{2cm}}$$

5 Logistic Regression


1. Consider a dataset with only binary features, $\mathbf{x} \in \{0, 1\}^d$, where feature x_1 is rare and only takes on value 1 in the training dataset if the corresponding the label is 1. If you fit a logistic regression model to this dataset using gradient descent, what value will the coefficient on this feature, \hat{w}_1 , approach? Briefly justify your answer.

2. A generalization of logistic regression to a multiclass settings involves expressing the per-class probabilities $P(y = c|x)$ as the softmax function $\frac{\exp(w_c^T x)}{\sum_{d \in C} \exp(w_d^T x)}$, where c is some class from the set of all classes C and w_c is a class-specific weight vector.

Consider a binary classification problem with labels encoded as $y \in \{0, 1\}$. Rewrite the expression above for this situation to end up with expressions for $P(Y = 1|x)$ and $P(Y = 0|x)$ that we derived in class for binary logistic regression.

6 Regularization

- The plot below shows data from two classes (filled and unfilled circles) separated by a linear decision boundary (dashed) defined by a two-dimensional weight vector and a bias term.



Suppose we apply L_2 regularization *to the bias term only* i.e., the two-dimensional weight vector remains the same for the regularized and unregularized settings. Which of the following best describes the new decision boundary that would result from this change. Assume that the amount of regularization is mild.

- farther from the origin; parallel to the previous line
 - farther from the origin; not parallel to the previous line
 - closer to the origin; parallel to the previous line
 - closer to the origin; not parallel to the previous line
- Complete the following paragraph about L_1 and L_2 regularization by circling the best of the provided options for each of blanks:

When applied to linear regression, L_1 regularization tends to learn

denser / sparser weight vectors than L_2 regularization because the

constraint surface is non-convex / square-shaped.

The optimal weight vector for linear regression with L_1 regularization can / cannot

be solved for in closed form.

7 Neural Networks & Backpropagation

1. Your friend wants to build a fully-connected one-hidden-layer Neural Network. Her inputs have 9 features. She wants 7 hidden neurons in the first layer and 2 neurons in the output layer generated using a softmax function. Each layer is a linear layer fully-connected to the previous one and includes a bias term. Each scalar counts as one parameter.

- (a) What is the number of parameters she will need to create this neural network? Each scalar counts as one parameter.

- (b) Your friend shares with you that the output she is trying to predict is binary. Using this information, describe an architecture that can learn the same decision boundary while reducing the number of parameters.

2. Your friend trains a Neural Network on some training dataset and achieves an almost perfect training accuracy. However they perform very poorly when it comes to test time. In the boxes below, write either **{Increase, Decrease or N/A}** to indicate what they should do with their model to reduce the amount of overfitting (N/A here means changing the specified quantity won't effect overfitting). *You should assume the training and test data are representative of the true underlying distribution of the data.*

the number of training data

the number of test data

the regularization constant

the number of hidden layers

the number of neurons in each hidden layer

the number of nodes in the output layer

3. In backpropagation, why is it necessary that we loop through the layers in reverse-order, moving from the output layer backwards towards the input layer?

4. When training a neural network, which of the following techniques is/are *primarily* used to address the nonconvexity of the loss function?

- Momentum
- Weight decay
- Random restarts
- Early termination
- None of the above

8 CNNs

1. What is the *primary* purpose of a convolutional layer in a CNN?
 - To reduce the dimensions of the input image.
 - To detect features such as edges and textures in the input image.
 - To classify the input image into various categories.
 - To flatten the input image across channels.
2. Suppose you are building a CNN that takes input images with 3 channels, each of size 20×20 pixels. Your first convolutional layer takes the 3 input channels and produces 5 output channels; it uses 4×4 filters. It also uses a padding of 2 (along all sides of the image) and a stride of 2 (in both dimensions).

- (a) How many parameters are there in this convolutional layer, including the bias terms?

- (b) What is the dimensionality of the output of this convolutional layer? Include the channel dimension in your answer and make sure you clearly indicate which dimension is the channel dimension.