

10-601 Machine Learning
Summer 2025
Practice Problems
Updated: May 22, 2025
Time Limit: N/A

Name:
Andrew ID:
Room:
Seat:
Exam Number:

Instructions:

- Verify your name and Andrew ID above.
- This exam contains 38 pages (including this cover page).
The total number of points is 0.
- Clearly mark your answers in the allocated space. If you have made a mistake, cross out the invalid parts of your solution, and circle the ones which should be graded.
- Look over the exam first to make sure that none of the 38 pages are missing.
- No electronic devices may be used during the exam.
- Please write all answers in pen or *darkly* in pencil.
- You have N/A to complete the exam. Good luck!

Question	Points
1. Decision Trees	0
2. K Nearest Neighbors	0
3. Model Selection and Errors	0
4. Perceptron	0
5. MLE/MAP	0
6. Logistic Regression and Regularization	0
7. Linear Regression	0
8. Optimization	0
9. Feature Engineering and Regularization	0
10. Neural Networks	0
Total:	0

Instructions for Specific Problem Types

For “Select One” questions, please fill in the appropriate bubble completely:

Select One: Who taught this course?

- ☒ Matt Gormley
- ☐ Marie Curie
- ☐ Noam Chomsky

If you need to change your answer, you may cross out the previous answer and bubble in the new answer:

Select One: Who taught this course?

- ☒ Henry Chai
- ☐ Marie Curie
- ☒ Noam Chomsky

For “Select all that apply” questions, please fill in all appropriate squares completely:

Select all that apply: Which are instructors for this course?

- ☒ Matt Gormley
- ☒ Henry Chai
- ☐ I don't know

Again, if you need to change your answer, you may cross out the previous answer(s) and bubble in the new answer(s):

Select all that apply: Which are the instructors for this course?

- ☒ Matt Gormley
- ☒ Henry Chai
- ☒ I don't know

For questions where you must fill in a blank, please make sure your final answer is fully included in the given space. You may cross out answers or parts of answers, but the final answer must still be within the given space.

Fill in the blank: What is the course number?

10-601

10-6301

1 Decision Trees (0 points)

- 1.1. To exploit the desirable properties of decision tree classifiers and perceptrons, Adam came up with a new algorithm called the “perceptron tree” that combines features from both. Perceptron trees are similar to decision trees, but each leaf node contains a perceptron rather than a majority vote.

To create a perceptron tree, the first step is to follow a regular decision tree learning algorithm (such as ID3) and perform splitting on attributes until the specified maximum depth is reached. Once maximum depth has been reached, at each leaf node, a perceptron is trained on the remaining attributes which have not yet been used in that branch. Classification of a new example is done via a similar procedure. The example is first passed through the decision tree based on its attribute values. When it reaches a leaf node, the final prediction is made by running the corresponding perceptron at that node.

Assume that you have a dataset with 6 binary attributes $\{A, B, C, D, E, F\}$ and two output labels $\{-1, 1\}$. A perceptron tree of depth 2 on this dataset is given below. Weights of the perceptron are given in the leaf nodes. Assume bias $b = 1$ for each perceptron.




Figure 1: Perceptron Tree of depth 2


- (a) **Numerical answer:** What would the given perceptron tree predict as the output label for the sample $\mathbf{x} = [1, 1, 0, 1, 0, 1]$?

- (b) **True or False:** The decision boundary of a perceptron tree will *always* be linear.
- ☐ True
- ☐ False
- (c) **True or False:** For small values of max depth (e.g., 2 or 3), decision trees are *more* likely to underfit the data than perceptron trees.
- ☐ True
- ☐ False
- 1.2. **Select all that apply:** Given a training dataset, \mathcal{D} , suppose you train two decision trees, one using mutual information as the splitting criterion and the other using training error rate. If both trees are trained until they achieve zero training error, which of the following statements must be true?
- ☐ Both trees split on the same feature at the root node.
- ☐ Both trees have the same depth
- ☐ Both trees make the same predictions for each data point in \mathcal{D}
- ☐ Both trees use all of the features in \mathcal{D}
- ☐ None of the above.
- 1.3. **True or False:** One advantage of the ID3 algorithm for training decision trees is that it is not susceptible to overfitting because it uses recursion.
- ☐ True
- ☐ False
- 1.4. **Fill in the blank:** Complete the following paragraph about pruning decision trees by circling the best of the provided options for each of blanks:

When pruning a decision tree, each split is replaced with the most common label among the training / validation data points that end up at each node. Splits are compared based on their reduction in training / validation error rate and if two splits are tied, we prune the node closer to / farther from the root.


1.5. The ID3 algorithm for learning decision trees greedily picks the split with the best mutual information at every node. In this question, you'll see an example of when this strategy may not result in a globally optimal tree.

- (a) **Drawing:** On the training dataset below, draw the decision boundaries learned by a decision tree with a max depth of 2. The decision tree uses mutual information as its splitting criterion and breaks ties by choosing $+$. You may **not** use a feature twice in a single path from the root node to a leaf node in the tree.



- (b) **Numerical answer:** What is the training error of the tree drawn in part (a)?

- (c) **Drawing:** On the training dataset below, draw the decision boundaries of the globally optimal tree—that is, the tree with the lowest training error. The same criteria apply from above: max depth of 2, ties are broken by in favor of $+$, and a feature may only be used once in a path from root to leaf.



- (d) **Numerical answer:** What is the training error of the tree drawn in part (c)?

2 K Nearest Neighbors (0 points)

2.1. Consider the following training dataset for a regression task:


$$\mathcal{D} = \{(x^{(1)}, y^{(1)}), (x^{(2)}, y^{(2)}), \dots, (x^{(N)}, y^{(N)})\}$$

with $x^{(i)} \in \mathbb{R}$ and $y^{(i)} \in \mathbb{R}$.

For regression with k -nearest neighbors, we make predictions on unseen data points similar to the classification algorithm, but instead of a majority vote, we take the mean of the output values of the k nearest points to some new data point x . That is,

$$h(x) = \frac{1}{k} \sum_{i \in \mathcal{N}(x, \mathcal{D})} y^{(i)}$$

where \mathcal{N} is the neighborhood function i.e., $\mathcal{N}(x, \mathcal{D})$ is the set of indices of the k closest training points to x .



In the dataset shown above, the red \times 's denote training points and the black semi-circles A, B, C denote test points of unknown output values. For convenience, all training data points have integer input and output values. Assume ties are broken by selecting the point with the lower x value.

- (a) **Numerical answer:** When $k = 1$, what is the mean squared error on the training set?

- (b) **Numerical answer:** When $k = 2$, what is the predicted value at A?

- (c) **Numerical answer:** When $k = 2$, what is the predicted value at B?

- (d) **Numerical answer:** When $k = 3$, what is the predicted value at C?

- (e) **Numerical answer:** When $k = 8$, what is the predicted value at C?


- (f) **Math:** With $k = N$, for any dataset \mathcal{D} with the form specified in the beginning of this question, write down a mathematical expression for the predicted value $\hat{y} = h(x)$. Your response shouldn't include a reference to the neighborhood function $\mathcal{N}()$.

- 2.2. **Select one:** Imagine you are using a k -Nearest Neighbor classifier on a dataset with lots of noise. You want your classifier to be *less* sensitive to the noise. Which of the following changes is likely to help achieve this goal and what is the corresponding effect on the prediction time?

- ☐ Increase the value of $k \rightarrow$ Increase in prediction time
- ☐ Decrease the value of $k \rightarrow$ Increase in prediction time
- ☐ Increase the value of $k \rightarrow$ Decrease in prediction time
- ☐ Decrease the value of $k \rightarrow$ Decrease in prediction time

- 2.3. You have just enrolled into your favourite course at CMU - Introduction to Machine Learning 10-301/601 - but you have not yet decided if you want to take it for a grade or as pass/fail. You want to use your performance in HW1 and HW2 to make this decision. You follow a general rule that if you can get at least a B in the course, you will take it for a grade, and if not, you will take it as pass/fail.

You have just learned the new classification technique, k -NN, and wish to employ it to make your decision. You start by collecting data from 10 prior students: you measure their performance on HW1 and HW2, along with their final letter grades. You then create a binary label based on the final grades such that you assign a label of $+$ if the final grade is at least a B and $-$ otherwise. Next, you train a k -NN model on this data set using the **Euclidean distance metric**.



- (a) **Numerical answer:** What is the training error rate on this dataset for a k -NN model where $k = 1$?

- (b) **Numerical answer:** What is the training error rate on this dataset for a k -NN model where $k = 3$?

- (c) **True or False:** Using Euclidean distance as the distance measure, the decision boundary of a k -NN model with $k = 1$ is a piece-wise straight line, that is, it contains only straight line segments. **Justify your answer.**

☐ True

☐ False

- (d) **Drawing:** In the image above, draw a rough decision boundary for $k = 1$. Clearly label the + and - sides of the decision boundary.

- (e) You have scored 60 in HW1 and 40 in HW2, and you now want to predict if your final grade would be at least a B.

- i. **Numerical answer:** What would be the predicted class when $k = 1$?

- ii. **Numerical answer:** What would be the predicted class when $k = 3$?

- (f) **Short answer:** Looking at the training errors, you choose the model with $k = 1$ as it has the lowest training error. Do you think this is the right approach to select a model? Why or why not?

2.4. **Select all that apply:** Which of the following statements about k -NN models is/are *always* true?

- ☐ Decreasing k makes a k -NN model have simpler or less complex decision boundaries.
- ☐ Increasing k makes a k -NN model less sensitive to outliers.
- ☐ k -NNs can be applied to classification problems but not regression problems.
- ☐ k -NNs can be applied to datasets with real-valued features but not categorical features.
- ☐ None of the above

3 Model Selection and Errors (0 points)


3.1. **Train and test errors:** In this problem, we will see how you can debug a classifier by looking at its train and test errors. Consider a classifier trained until convergence on some training data $\mathcal{D}^{\text{train}}$, and tested on a separate test set $\mathcal{D}^{\text{test}}$. You look at the test error, and find that it is very high. You then compute the training error and find that it is close to 0.

(a) **Short Answer:** What is this scenario called?


(b) **Select all that apply:** Which of the following actions are likely to address the issue you identified in the previous question?

- ☐ Increasing the training data size.
- ☐ Decreasing the training data size.
- ☐ Increasing model complexity (For example, if your classifier is a decision tree, increase the depth).
- ☐ Decreasing model complexity.
- ☐ Training on a combination of $\mathcal{D}^{\text{train}}$ and $\mathcal{D}^{\text{test}}$ and test on $\mathcal{D}^{\text{test}}$
- ☐ None of the above

(c) **Select one:** Say you plot the train and test errors as a function of the model complexity. Which of the following two plots is your plot expected to look like?



(a)



(b)

- ☐ Plot A
- ☐ Plot B

3.2. What are the effects of the following on overfitting? Choose the best answer.

- (a) Increasing the max depth a decision tree is allowed to grow to.
 - ☐ Less likely to overfit
 - ☐ More likely to overfit
- (b) Increasing the minimum mutual information threshold required to add a split to a decision tree.
 - ☐ Less likely to overfit
 - ☐ More likely to overfit
- (c) Increasing the minimum number of data points required to add a split to a decision tree.
 - ☐ Less likely to overfit
 - ☐ More likely to overfit
- (d) Increasing k in k -nearest neighbor.
 - ☐ Less likely to overfit
 - ☐ More likely to overfit
- (e) Increasing the training data size for decision trees. Assume that training data points are drawn independently from the true data distribution.
 - ☐ Less likely to overfit
 - ☐ More likely to overfit
- (f) Increasing the training data size for 1-nearest neighbor. Assume that training data points are drawn independently from the true data distribution.
 - ☐ Less likely to overfit
 - ☐ More likely to overfit

- 3.3. Consider a learning algorithm that uses two hyperparameters, γ and ω , and it takes 1 hour to train *regardless* of the size of the training set.

We choose to do random subsampling cross-validation, where we do K runs of cross-validation and for each run, we randomly subsample a fixed fraction αN of the dataset for validation and use the remaining for training, where $\alpha \in (0, 1)$ and N is the number of data points.

- (a) **Numerical answer:** In combination with the cross-validation method above, we choose to do grid search on discrete values for the two hyperparameters.

Given $N = 1000$ data points, $K = 4$ runs, and $\alpha = 0.25$, if we have 100 hours to complete the entire cross-validation process, what is the maximum number of discrete values of γ that we can include in our search if we also want to include 8 values of ω ? Assume that any computations other than training are negligible.

- (b) **Short answer:** In 1-2 concise sentences, describe one advantage of increasing the value of α .

- 3.4. **Select one:** When implementing k-Fold Cross-Validation for a Linear Regression model with a 1-dimensional input feature, which of the following best describes the correct procedure?

- ☐ The dataset is divided into k contiguous blocks, where each block consists of adjacent data points based on the sorting of the 1-dimensional feature values. The model is trained k times, each time using a different block as the test set and the remaining blocks for training.
- ☐ The dataset is randomly partitioned into k equal-sized subsets. For each fold, the Linear Regression model is trained on $k-1$ subsets and validated on the remaining subset to assess the model's performance.
- ☐ The model is trained once using the entire dataset as the training set and then tested k times, each time with a different single data point as the test set.
- ☐ We fit a Linear Regression model to every possible subset of size k of the training data (resulting in $\binom{N}{k}$ models) and then average their errors on a held-out test dataset.

4 Perceptron (0 points)

4.1. Suppose you are given the following dataset:

Example Number	X_1	X_2	Y
1	-1	2	-1
2	-2	-2	+1
3	1	-1	+1
4	-3	1	-1

You wish to perform the Batch Perceptron algorithm on this data. Assume you start with initial weights $\theta^T = [0, 0]$ and bias $b = 0$, and that you pass through all of the examples in order of their example number.

- (a) **Numerical answer:** What would be the updated weight vector θ after we pass example 1 through the Perceptron algorithm?

- (b) **Numerical answer:** What would be the updated bias b after we pass example 1 through the Perceptron algorithm?

- (c) **Numerical answer:** What would be the updated weight vector θ after we pass example 2 through the Perceptron algorithm?

- (d) **Numerical answer:** What would be the updated bias b after we pass example 2 through the Perceptron algorithm?

- (e) **Numerical answer:** What would be the updated weight vector θ after we pass example 3 through the Perceptron algorithm?

- (f) **Numerical answer:** What would be the updated bias b be after we pass example 3 through the Perceptron algorithm?

(g) **True or False:** Your friend stops you here and tells you that you do not need to update the Perceptron weights or bias anymore; is this true or false?

- ☐ True
- ☐ False

4.2. **Select all that apply:** Let $S = \{(\mathbf{x}^{(1)}, y^{(1)}), \dots, (\mathbf{x}^{(n)}, y^{(n)})\}$ be n linearly separable points by a separator through the origin in \mathbb{R}^d . Let S' be generated from S as: $S' = \{(c\mathbf{x}^{(1)}, y^{(1)}), \dots, (c\mathbf{x}^{(n)}, y^{(n)})\}$, where $c > 1$ is a constant. Suppose that we would like to run the perceptron algorithm on both data sets separately, with the same initial $\boldsymbol{\theta}$ from class, and that the perceptron algorithm converges on S . Which of the following statements are true?

- ☐ The mistake bound of perceptron on S' is larger than the mistake bound on S
- ☐ The perceptron algorithm when run on S and S' returns the same classifier, modulo constant factors (i.e., if \mathbf{w}_S and \mathbf{w}'_S are outputs of the perceptron for S and S' , then $\mathbf{w}'_S = c_1 \mathbf{w}_S$ for some constant c_1).
- ☐ The perceptron algorithm converges on S' .
- ☐ None of the above.

4.3. **True or False:** Given a linearly separable dataset, the convergence time of the perceptron algorithm depends on the sample size n .

- ☐ True
- ☐ False

4.4. **Select all that apply:** Which of the following are inductive biases of the perceptron algorithm?

- ☐ Most of the cases in a small neighborhood in feature space belong to the same class.
- ☐ The true decision boundary is linear.
- ☐ We prefer to correct the most recent mistakes.
- ☐ We prefer the simplest hypothesis that explains the data.
- ☐ None of the above.

4.5. **True or False:** All *data points* (\mathbf{x}, y) that the Perceptron learning algorithm sees during training affect the weights equally.

- ☐ True
- ☐ False

5 MLE/MAP (0 points)

- 5.1. Magnetic Resonance Imaging (MRI) scans are commonly used to generate detailed images of patients' internal anatomy at hospitals. The scanner returns an image with N pixels. For each pixel, we extract the noise from that pixel to obtain a vector of noise terms $\mathbf{x} \in \mathbb{R}^N$ s.t. $\forall i \in \{1 \dots N\}$, $x_i \geq 0$ and x_i is independent and identically distributed and follows a Rayleigh distribution. The probability density function of a Rayleigh distribution is given by:

$$f(x \mid \sigma) = \frac{x}{\sigma^2} \exp\left(\frac{-x^2}{2\sigma^2}\right)$$

for scale parameter $\sigma \geq 0$ and $x \geq 0$.

- (a) **Math:** Write the log-likelihood $\ell(\sigma)$ of a noise vector \mathbf{x} obtained from one image. Report your answer in terms of the variables x_i, i, N, σ , the function $\exp(\cdot)$, and any constants you may need. For full credit you must push the log through to remove as many multiplications/divisions as possible.

- (b) **Math:** Report the maximum likelihood estimator of the scale parameter, σ , for a single image's noise vector \mathbf{x} .

- 5.2. Suppose a random variable k follows a **Poisson distribution** with unknown rate parameter λ :

$$p(k \mid \lambda) = \frac{\lambda^k e^{-\lambda}}{k!} \quad k = 1, 2, \dots$$

The Poisson distribution is used for modeling the number of times an event occurs within a fixed time interval given a mean occurrence rate assuming that the occurrences are independent.

A conjugate prior for the rate parameter in a Poisson likelihood is a **gamma distribution** with shape parameter α and rate parameter β :

$$f(\lambda \mid \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda} \quad \lambda > 0$$

where Γ is some normalizing constant that does not depend on λ .

- (a) **Math:** Suppose we receive a set of N observations of k : $\mathcal{D} = \{k_1, k_2, \dots, k_N\}$; assume the observations are independent and identically distributed. Using the Poisson distribution and gamma prior with parameters α and β , derive an expression for the unnormalized log posterior of λ i.e. the sum of the log prior and the log likelihood of \mathcal{D} .

Express your answer in terms of α , β , λ , k_1, k_2, \dots, k_N , N and Γ ; simplify your answer as much as possible.

- (b) **Math:** Take the partial derivative of the expression you derived in the previous part with respect to λ .

Express your answer in terms of α , β , λ , k_1, k_2, \dots, k_N , and N ; simplify your answer as much as possible.

- (c) **Math:** Finally, compute the MAP estimate of λ given \mathcal{D} by setting the partial derivative you computed in the previous part equal to 0 and solving for λ .

Express your answer in terms of α , β , k_1, k_2, \dots, k_N , and N ; simplify your answer as much as possible.

- 5.3. **True or False:** A random variable x follows a probability distribution with a single, real-valued parameter, θ . In the limit of infinitely many samples of x , the MAP estimate of θ will *always* approach the MLE estimate, regardless of your choice of prior on θ . Briefly justify your answer in 1-2 sentences.

6 Logistic Regression and Regularization (0 points)

- 6.1. **Math:** A generalization of logistic regression to a multiclass settings involves expressing the per-class probabilities $P(y = c \mid x)$ 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 .

Consider a 2-class problem with labels 0 or 1. Rewrite the above expression for this situation to derive the expressions for $P(Y = 1|x)$ and $P(Y = 0|x)$ that we have already seen in class for binary logistic regression.

- 6.2. **Short answer:** Given 3 data points $(1, 1), (1, 0), (0, 0)$ with labels 0, 1, 0 respectively, consider 2 models that define $p(y = 1 \mid \mathbf{x})$:

$$\text{Model 1: } \sigma(w_1 x_1 + w_2 x_2)$$

$$\text{Model 2: } \sigma(w_0 + w_1 x_1 + w_2 x_2)$$

As usual, $\sigma(z)$ is the sigmoid function $\frac{1}{1+e^{-z}}$. Using the given dataset, suppose we learn parameters for both models, $\hat{\mathbf{w}}_1$ and $\hat{\mathbf{w}}_2$, by maximizing the conditional log-likelihood.

If we switched $(0, 0)$ to label 1 instead of label 0, would the parameters we learn for Model 1 change? What about Model 2?

- 6.3. Given a training dataset $\mathcal{D} = \{(\mathbf{x}^{(1)}, y^{(1)}), \dots, (\mathbf{x}^{(N)}, y^{(N)})\}$ where $\mathbf{x}^{(i)} \in \mathbb{R}^d$ is a feature vector and $y_i \in \{0, 1\}$ is a binary label, we want to find the parameters $\hat{\mathbf{w}}$ that maximize the likelihood of the training dataset, assuming a parametric model of the form

$$p(y = 1 \mid \mathbf{x}; \mathbf{w}) = \frac{1}{1 + \exp(-\mathbf{w}^T \mathbf{x})}.$$

The conditional log likelihood of \mathcal{D} is

$$\ell(\mathbf{w}) = \sum_{i=1}^n y_i \log p(y_i \mid \mathbf{x}_i; \mathbf{w}) + (1 - y_i) \log(1 - p(y_i \mid \mathbf{x}_i; \mathbf{w})),$$

and the gradient is

$$\nabla \ell(\mathbf{w}) = \sum_{i=1}^N (y_i - p(y_i \mid \mathbf{x}_i; \mathbf{w})) \mathbf{x}_i.$$

- (a) **Short answer:** Is it possible to solve for the optimal parameters $\hat{\mathbf{w}}$ in closed form? If yes, explain how and if no, describe how you would compute $\hat{\mathbf{w}}$ in practice?

- (b) **Math:** For a binary logistic regression model, we predict $y = 1$ when $p(y = 1 \mid \mathbf{x}) \geq 0.5$. Show that this is a linear classifier in \mathbf{x} .

- (c) Consider the case where all features are binary, i.e. $\mathbf{x} \in \{0, 1\}^d$. Furthermore, suppose feature x_1 is rare and just happens to take on value 1 in the training set only when the label is also 1.

What is \hat{w}_1 i.e. the optimal weight on the first feature? Is the gradient ever zero for any finite value of w_1 ?

- 6.4. **Math:** Given the following dataset, \mathcal{D} , and a fixed parameter vector, \mathbf{w} , write an expression for the conditional likelihood of a binary logistic regression model on this dataset.

$$\mathcal{D} = \{(\mathbf{x}^{(1)}, y^{(1)} = 0), (\mathbf{x}^{(2)}, y^{(2)} = 0), (\mathbf{x}^{(3)}, y^{(3)} = 1), (\mathbf{x}^{(4)}, y^{(4)} = 1)\}$$

- Write your answer in terms of \mathbf{w} , $\mathbf{x}^{(1)}$, $\mathbf{x}^{(2)}$, $\mathbf{x}^{(3)}$, and $\mathbf{x}^{(4)}$.
- Do not include $y^{(1)}$, $y^{(2)}$, $y^{(3)}$, or $y^{(4)}$ in your answer.
- Do not try to simplify your expression.

6.5. Suppose we apply feature engineering to a two-dimensional input, $x_1 \in \mathbb{R}$ and

$x_2 \in \mathbb{R}$, mapping it to a new input vector: $\mathbf{x} = \begin{bmatrix} 1 \\ x_1^2 \\ x_2^2 \end{bmatrix}$

- (a) Write an expression for the decision boundary of binary logistic regression with this feature vector \mathbf{x} and the corresponding parameter vector $\boldsymbol{\theta} = [b, w_1, w_2]^T$. Assume that the decision boundary occurs when $P(Y = 1 \mid x, \boldsymbol{\theta}) = P(Y = 0 \mid x, \boldsymbol{\theta})$. Write your answer in terms of x_1 , x_2 , b , w_1 , and w_2 .

Decision boundary expression:

- (b) Assume that $w_1 > 0$, $w_2 > 0$, and $b < 0$. What is the geometric shape defined by this equation?

- (c) If we add an L2 regularization term when learning $[w_1, w_2]^T$, what happens to the **parameters** as we increase the λ that scales this regularization term?

- (d) If we add an L2 regularization term when learning $[w_1, w_2]^T$, what happens to the **decision boundary shape** as we increase the λ that scales this regularization term?

7 Linear Regression (0 points)

7.1. **Select all that apply:** For some function $f: \mathbb{R} \rightarrow \mathbb{R}$, suppose the unique argmax of f is x^* . Which of the following functions are guaranteed to have the same argmax as f ?

☐ $g(x) = f(x) + 2$

☐ $g(x) = 2f(x)$

☐ $g(x) = f(x)^2$

☐ $g(x) = 2^{f(x)}$

☐ None of the above.

7.2. Consider linear regression on 1-dimensional data i.e., $x, y \in \mathbb{R}$. We apply linear regression in both directions on this data: we first fit a model to predict y using x and get $y = m_1x + b_1$ as the fitted line that minimizes the mean squared error, then we fit a model to predict x using y and get $x = m_2y + b_2$.

True or False: It is always the case that $m_2 = \frac{1}{m_1}$.

☐ True

☐ False

7.3. **Select all that apply:** Which of the following statements about the gradient of a function is/are true?

☐ The gradient is the same dimensionality as the function's output.

☐ The gradient points in the direction of steepest descent.

☐ The gradient of a function always exists everywhere.

☐ The gradient is the limit of the slope between two points on the function as the distance between the points approaches ∞ .

☐ None of the above.

- 7.4. Instead of mean squared error, suppose instead that you fit a linear regression model by minimizing the mean cubed error:

$$J(\theta) = \frac{1}{N} \sum_{i=1}^N (\theta^T \mathbf{x}^{(i)} - y^{(i)})^3$$

- (a) **Math:** What is the gradient of the mean cubed error with respect to the parameter vector θ ?

- (b) **True or False:** This loss function is more sensitive to outliers than the mean squared error.

☐ True

☐ False

- (c) **Short answer:** In 1-2 concise sentences, briefly describe the primary issue with this objective function for regression?

- 7.5. **True or False:** Consider a linear regression model with only one parameter, the intercept, i.e., $y = b$. Then, given N data points $(x^{(i)}, y^{(i)})$ (where $x^{(i)}$ is the feature and $y^{(i)}$ is the target), minimizing the sum of squared errors results in b being the *median* of the $y^{(i)}$ values. Briefly justify your answer

☐ True

☐ False

7.6. You decide to optimize the function $f(x) = x^2$ with respect to x using gradient descent. You initialize $x^{(0)}$ to -1 and use a step size of $\eta^{(0)} = 1.5$.

- (a) **Fill in the blanks:** Fill in the table below with the results you get from running gradient descent in this setting for two iterations; most of the 0th iteration has been filled in on your behalf; you should fill in the missing element in this line along with all the other missing elements.

t	$x^{(t)}$	$f(x^{(t)})$	$\nabla_x f(x^{(t)})$
0	-1	1	_____
1	_____	_____	_____
2	_____	_____	_____

- (b) **Select all that apply:** Based on your findings from the previous part, which of the following changes could help you achieve better results?

- ☐ Decreasing the step size.
- ☐ Running gradient descent for more iterations.
- ☐ Moving the initial value $x^{(0)}$ further from the origin.
- ☐ Negating the function and running gradient *ascent*.
- ☐ None of the above.

7.7. **Math:** Your friend accidentally solved linear regression using the wrong objective function for mean squared error; specifically, they used the following objective function that contains two mistakes: 1) they forgot the $1/N$ and 2) they have one sign error.

$$J(\mathbf{w}, b) = \sum_{i=1}^N \left(y^{(i)} - \left(\sum_{j=1}^M w_j x_j^{(i)} - b \right) \right)^2$$

You realize that you can still use the parameters that they learned, \mathbf{w} and b , to compute the minimizer of the mean squared error. Write the equation that implements this corrected prediction function $h(\mathbf{x}, \mathbf{w}, b)$ using your friend's learned parameters, \mathbf{w} and b .

8 Optimization (0 points)

- 8.1. (a) **Select all that apply:** Determine if the following 1-D functions are convex. Assume that the domain of each function is \mathbb{R} . The definition of a convex function is as follows:

$$f(x) \text{ is convex} \iff f(\alpha x + (1-\alpha)z) \leq \alpha f(x) + (1-\alpha)f(z), \forall \alpha \in [0, 1] \text{ and } \forall x, z.$$

- ☐ $f(x) = x + b$ for any $b \in \mathbb{R}$
- ☐ $f(x) = c^2 x$ for any $c \in \mathbb{R}$
- ☐ $f(x) = ax^2 + b$ for any $a \in \mathbb{R}$ and any $b \in \mathbb{R}$
- ☐ $f(x) = 0$
- ☐ None of the above

- (b) **Select all that apply:** Suppose we are trying to minimize the convex function $f(z) = z^2$ using gradient descent. Let α be the learning rate and assume that we use an initial value of $z^{(0)} = 1$.

For which values of α will gradient descent converge to the optimal value, $x^* = 0$?

- ☐ $\alpha = 0$
- ☐ $\alpha = \frac{1}{2}$
- ☐ $\alpha = 1$
- ☐ $\alpha = 2$
- ☐ None of the above

- 8.2. **Fill in the Blanks:** Complete the following sentence by circling the best option in each square (options are separated by “/”s):

The mean squared error objective function for linear regression is

non-convex / convex / strictly convex

which means that gradient descent will con-

verge to a

global minimum / local minimum which might not be a global minimum

if it

converges; however, a global minimum

does not / does

always exist.

8.3. **Select all that apply:** Which of the following statements about gradient descent (GD) and stochastic gradient descent (SGD) are correct?

- ☐ Each update step in SGD pushes the parameter vector closer to the parameter vector that minimizes the objective function.
- ☐ The gradient computed in SGD is, in expectation, equal to the gradient computed in GD.
- ☐ The gradient computed in GD has a higher variance than that computed in SGD, which is why in practice SGD converges faster in time than GD.
- ☐ Both SGD and GD are guaranteed to converge if and only if the objective function is strictly convex.
- ☐ None of the above.

8.4. **True or False:** For a given dataset and objective function, one *epoch* of stochastic gradient descent (i.e., one pass through the entire dataset) will always have the same big-O computational cost as one iteration of gradient descent.

- ☐ True
- ☐ False

9 Feature Engineering and Regularization (0 points)

- 9.1. Suppose you have D -dimensional data points $\mathbf{x} = [x_1, x_2, \dots, x_D]^T$ and you want to apply a 2-dimensional polynomial feature expansion to them, ϕ_2 . This feature expansion contains all first-order terms as well as *all possible second-order terms*, including combinations of features i.e.

$$\{x_4, x_1^2, x_2x_3\} \subset \phi_2(\mathbf{x})$$

- (a) **Math:** What is the dimensionality of $\phi_2(\mathbf{x})$? Express your answer in terms of D , the dimensionality of \mathbf{x} .

- (b) **Math:** Given a dataset consisting of N data points, suppose a unique closed-form solution exists for linear regression using the ϕ_2 -transformation. What is the big-O computational cost of computing $(\Phi^T \Phi)^{-1}$, where Φ is the design matrix of the transformed data points i.e.,

$$\Phi = \begin{bmatrix} 1 & \phi_2(\mathbf{x}^{(1)}) \\ 1 & \phi_2(\mathbf{x}^{(2)}) \\ \vdots & \vdots \\ 1 & \phi_2(\mathbf{x}^{(N)}) \end{bmatrix}$$

Express your answer in terms N and D .

Hint: don't forget about the computational costs associated with the bias or intercept parameters.

- 9.2. **Model Complexity:** In this question we will consider the effect of increasing the model complexity, while keeping the size of the training dataset fixed. To be concrete, consider a classification task on the real line \mathbb{R} with some unknown distribution over data points D and unknown target function $c^* : \mathbb{R} \rightarrow \{\pm 1\}$.

Suppose we have a randomly sampled dataset S of size N , drawn i.i.d. from D . For each degree d , let ϕ_d be the feature map given by $\phi_d(x) = (1, x, x^2, \dots, x^d)$ that maps points on the real line to $(d + 1)$ -dimensional space.

Now consider the learning algorithm that first applies the feature map ϕ_d to all the training data points and then runs logistic regression. A new example is classified by first applying the feature map ϕ_d and then using the learned classifier.

- (a) For a given dataset S , is it possible for the training error to increase when we increase the degree d of the feature map? Briefly justify your answer in 1 to 2 sentences.

- (b) Suppose we plot the true error of our algorithm as a function of d : we observe that it initially decreases and then increases as we increase the degree d . Briefly explain this trend in 2 to 3 sentences.

- 9.3. **Short Answer:** Your friend is training a logistic regression model with ridge regularization, where λ is the regularization constant. They run cross-validation for $\lambda = [0.01, 0.1, 1, 10]$ and compare train, validation and test errors. They choose $\lambda = 0.01$ because that had the lowest *test* error.

However, you observe that the test error linearly increases from $\lambda = 0.01$ to 10 and thus, there exists a value of $\lambda < 0.01$ that gives a lower test error. You tell your friend that they should run the cross-validation for $\lambda = [0.0001, 0.001, 0.01]$ to get the optimal model.

Do you think you did the right thing by giving your friend this suggestion? Briefly justify your answer in 1 to 2 concise sentences.

10 Neural Networks (0 points)

10.1. **Matching:** Match the corresponding neural network component to its role in the neural network.

- (a) Cross-Entropy
- (b) Identity
- (c) Mean Absolute Error
- (d) Mean Squared Error
- (e) ReLU
- (f) Sigmoid
- (g) Softmax
- (h) Stochastic Gradient Descent
- (i) Tanh

Activation Function

Loss Function

Optimizer




Figure 3: Neural Network

10.2. Consider the neural network architecture shown above for a binary classification problem. The values for the weights are shown in the figure. We define:

$$a_1 = w_{11}x_1 + b_{11}$$

$$a_2 = w_{12}x_1 + b_{12}$$

$$a_3 = w_{21}z_1 + w_{22}z_2 + b_{21}$$

$$z_1 = \text{ReLU}(a_1)$$

$$z_2 = \text{ReLU}(a_2)$$

$$z_3 = \sigma(a_3), \sigma(x) = \frac{1}{1+e^{-x}}$$

where $\text{ReLU}(x) = \max(0, x)$.

(a) **Math:** For $x_1 = 0.3$, compute z_3 in terms of e .

(b) **Select one:** Which class does the network predict for the data point $x_1 = 0.3$, assuming that $\hat{y} = 1$ if $z_3 > \frac{1}{2}$ and otherwise, $\hat{y} = 0$.

☐ 1

☐ 0

- (c) **Math:** Perform backpropagation on the bias term b_{21} by deriving the expression for the gradient of the loss function $L(y, z_3)$ with respect to the bias term b_{21} , $\frac{\partial L}{\partial b_{21}}$.

Express your answer in terms of partial derivatives of the form $\frac{\partial \alpha}{\partial \beta}$, where α and β can be any of $L, z_i, a_i, b_{ij}, w_{ij}, x_1$ for all valid values of i, j . Your backpropagation algorithm should be as explicit as possible — that is, make sure each partial derivative $\frac{\partial \alpha}{\partial \beta}$ cannot be decomposed further into simpler partial derivatives. Do *not* evaluate the partial derivatives.

- (d) **Math:** Perform backpropagation on the bias term b_{12} by deriving the expression for the gradient of the loss function $L(y, z_3)$ with respect to the bias term b_{12} , $\frac{\partial L}{\partial b_{12}}$.

Express your answer in terms of partial derivatives of the form $\frac{\partial \alpha}{\partial \beta}$, where α and β can be any of $L, z_i, a_i, b_{ij}, w_{ij}, x_1$ for all valid values of i, j . Your backpropagation algorithm should be as explicit as possible — that is, make sure each partial derivative $\frac{\partial \alpha}{\partial \beta}$ cannot be decomposed further into simpler partial derivatives. Do *not* evaluate the partial derivatives.

- 10.3. In this problem we will use a neural network to distinguish the crosses (\times) from the circles (\circ) in the simple data set shown in Figure 4a. Even though the crosses and circles are not linearly separable, we can break the examples into three groups, S_1 , S_2 , and S_3 (shown in Figure 4a) so that S_1 is linearly separable from S_2 and S_2 is linearly separable from S_3 . We will exploit this fact to design weights for the neural network shown in Figure 4b in order to correctly classify this training dataset. For all nodes, we will use the threshold activation function

$$\phi(z) = \begin{cases} 1 & z > 0 \\ 0 & z \leq 0. \end{cases}$$

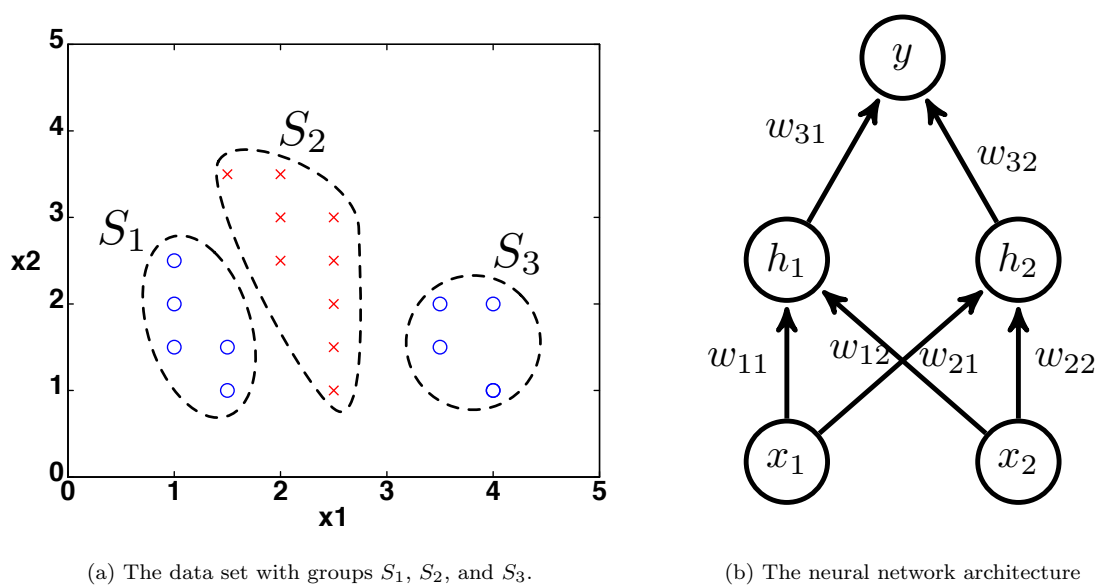


Figure 4





Figure 5: NN classification.

- (a) First we will set the parameters w_{11}, w_{12} and b_1 of the neuron labeled h_1 so that its output $h_1(x) = \phi(w_{11}x_1 + w_{12}x_2 + b_1)$ forms a linear separator between the sets S_2 and S_3 .

- i. On Fig 5a, draw a linear decision boundary that separates S_2 and S_3 .
- ii. Write down a possible setting of the weights w_{11}, w_{12} , and b_1 such that $h_1(x) = 0$ for all points in S_3 and $h_1(x) = 1$ for all points in S_2 .




- (b) Next we will set the parameters w_{21}, w_{22} and b_2 of the neuron labeled h_2 so that its output $h_2(x) = \phi(w_{21}x_1 + w_{22}x_2 + b_2)$ forms a linear separator between the sets S_1 and S_2 .

- i. On Fig 5b, draw a linear decision boundary that separates S_1 and S_2 .
- ii. Write down a possible setting of the weights w_{21}, w_{22} , and b_2 such that $h_2(x) = 0$ for all points in S_1 and $h_2(x) = 1$ for all points in S_2 .



- (c) Now we have two classifiers h_1 (to classify S_2 from S_3) and h_2 (to classify S_1 from S_2). We will set the weights of the final neuron of the neural network based on the results from h_1 and h_2 to classify the crosses from the circles. Let $h_3(x) = \phi(w_{31}h_1(x) + w_{32}h_2(x) + b_3)$.

- i. Write down a possible setting of the weights w_{31}, w_{32}, b_3 such that $h_3(x)$ correctly classifies the entire data set.



- ii. Draw your decision boundary in Fig 5c.

10.4. Consider the following neural network for a 2-D input, $x_1 \in \mathbb{R}$ and $x_2 \in \mathbb{R}$ where:




Figure 6: Neural Network

- All occurrences of the function g are the same arbitrary non-linear activation function with no parameters
- $\ell(y, \hat{y})$ is an arbitrary loss function with no parameters, and:

$$z_1 = w_A x_1 + w_B x_2, a_1 = g(z_1)$$

$$z_2 = w_C a_1, a_2 = g(z_2)$$

$$z_3 = w_D a_1, a_3 = g(z_3)$$

$$z_4 = w_E a_2 + w_F a_3, \hat{y} = g(z_4)$$

Note: There are no bias terms in this network.

(a) What is the chain of partial derivatives needed to calculate the derivative $\frac{\partial \ell}{\partial w_E}$?

Your answer should be in the form: $\frac{\partial \ell}{\partial w_E} = \frac{\partial ?}{\partial ?} \frac{\partial ?}{\partial ?} \dots$. Make sure each partial derivative $\frac{\partial ?}{\partial ?}$ in your answer cannot be decomposed further into simpler partial derivatives. **Do not evaluate the derivatives.** Be sure to specify the correct subscripts in your answer.

$$\frac{\partial \ell}{\partial w_E} =$$

- (b) The network diagram from above is repeated here for convenience: What is the




Figure 7: Neural Network

chain of partial derivatives needed to calculate the derivative $\frac{\partial \ell}{\partial w_C}$?

Your answer should be in the form:

$$\frac{\partial \ell}{\partial w_C} = \frac{\partial ?}{\partial ?} \frac{\partial ?}{\partial ?} \dots$$

Make sure each partial derivative $\frac{\partial ?}{\partial ?}$ in your answer cannot be decomposed further into simpler partial derivatives. **Do not evaluate the derivatives.** Be sure to specify the correct superscripts in your answer.

$$\frac{\partial \ell}{\partial w_C} =$$

- (c) We want to modify our neural network objective function to add an L2 regularization term on the weights. The new objective is:

$$\ell(y, \hat{y}) + \lambda \frac{1}{2} \|\mathbf{w}\|_2^2$$

where λ (lambda) is the regularization hyperparameter and \mathbf{w} is all of the weights in the neural network stacked into a single vector, $\mathbf{w} = [w_A, w_B, w_C, w_D, w_E, w_F]^T$.

Write the right-hand side of the new gradient descent update step for weight w_C given this new objective function. You may use $\frac{\partial \ell}{\partial w_C}$ in your answer.

Update: $w_C \leftarrow \dots$

10.5. Define a function $\text{floor} : \mathbb{R}_n \rightarrow \mathbb{R}_n$ such that

$$\text{floor}(\mathbf{z}) = [\lfloor z_i \rfloor \text{ for } 0 \leq i \leq D]^T$$

or essentially, a function that produces an output vector by applying $\lfloor \cdot \rfloor$ element-wise to the input vector.

Neural wants to use this function as an activation function to train his neural network. Is this possible? Explain why or why not.
