

Data C100, Final Exam

Summer 2023

Name: _____

Email: _____@berkeley.edu

Student ID: _____

Examination room: _____

Name of the student to your left: _____

Name of the student to your right: _____

Instructions:

Do not open the examination until instructed to do so.

This exam consists of **80 points** spread out over **7 questions** and must be completed in the **110 minute** time period on August 10, 2023, from 5:10 PM to 7:00 PM unless you have pre-approved accommodations otherwise.

For multiple-choice questions with circular bubble options, **select one choice**. For multiple-choice questions with box options, **select all choices that apply**. In both cases, please **shade in** the box/circle to mark your answer. Do not use checkmarks or “×”s.

Make sure to write your student ID on each sheet (in the provided blanks) to ensure that your exam is graded.

Honor Code [1 pt]:

As a member of the UC Berkeley community, I act with honesty, integrity, and respect for others. I am the person whose name is on the exam, and I completed this exam in accordance with the Honor Code.

Signature: _____

1 Weepings of the Realm [10 Pts]

This question involves SQL databases. All code for this question must be written as SQLite queries. In each blank, you may write as much code as is necessary, provided it fits the given skeleton code. Throughout this question, you may assume that any numeric data is stored as floats.

Link, a heroic knight, has turned to data science to help him optimize his adventures.

He creates the SQL table `quests` to record the quests he has been offered on his journeys. The first few rows of `quests` are shown to the right. The table contains three columns:

`name`: the name of the person offering the quest

`location`: the location of the quest

`reward`: the monetary reward for completing the quest

name	location	reward
Hestu	Woodlands	10
Impa	Village	50
Sidon	Lake	20
Tulin	Mountain	5
Paya	Village	45
NULL	Underground	100

- (a) [4 Pts] Link, being an honorable hero, only wants to accept quests from people he knows. Write a SQL query to return the name, location, and reward associated with the **one** quest with the **highest reward** offered by someone with a **non-NULL name**. The desired output is displayed below to the right.

```
SELECT *
FROM quests
_____ A _____
_____ B _____
_____ C _____;
```

name	location	reward
Impa	Village	50

- (i) Fill in Blank A:

- (ii) Fill in Blank B:

- (iii) Fill in Blank C:

Link is an honorable *and* practical hero.

He creates a second table, `travel`, to store the travel cost associated with journeying to each location. The first few rows of `travel` are shown to the right.

location	travel_cost
Village	3
Lake	25
Woodlands	10
Mountain	50

(b) [6 Pts] Fill in the blanks below to write a SQL query that returns a table with **two columns**:

`location`: a location with one or more quests

`avg_reward`: the average reward of all quests in that location

Only include locations **where the total reward** of all quests in the location is **greater** than the travel cost of going to that location. If a location is not present in one or both of Link's tables, this location should be excluded.

```
SELECT q.location, _____A_____ AS avg_reward
FROM quests AS q _____B_____ travel AS t
ON q.location = t.location
_____C_____
_____D_____;
```

(i) Fill in Blank A:

(ii) Fill in Blank B:

(iii) Fill in Blank C:

(iv) Fill in Blank D:

2 Yodel Model [12 Pts]

It's time for the 100th annual Berkeley Yodeling Competition! Competitors will perform their best yodels (a type of singing) to win fame and glory. Too eager to wait for the final results, you decide to apply your modeling skills to predict the length of each competitor's longest yodel.

- (a) [4 Pts] You collect the following training dataset of yodel lengths (y), time spent practicing (x_1), and lung capacity (x_2). You decide to fit a **constant model** $\hat{y} = \theta_0$ using **absolute (L1) loss** to compute the optimal parameter $\hat{\theta}_0$.

x_1	x_2	y
1	3	1
3/2	5	1
4	1	4

What is the **root mean squared error (RMSE)** of the constant model on the training data above when using the optimal parameter $\hat{\theta}_0$? Simplify your answer as much as possible.

RMSE = _____

- (b) [2 Pts] Your friend wants to create their own model and fit it using **ordinary least squares (OLS)**. For which of the following loss (empirical risk) functions can OLS be used to compute the optimal value of the parameters θ_i ? Assume that x_{i1} and x_{i2} represent the first and second feature, respectively, of the i th observation.

- A. $R(\theta) = \frac{1}{n} \sum_{i=1}^n (y_i - \theta_0)^2$
- B. $R(\theta) = \frac{1}{n} \sum_{i=1}^n |y_i - \theta_0 - \theta_1 x_{i1}|$
- C. $R(\theta) = \frac{1}{n} \sum_{i=1}^n (y_i - \theta_1 x_{i1} - \theta_2 x_{i2})^2$
- D. $R(\theta) = \frac{1}{n} \sum_{i=1}^n (y_i - \theta_1 |x_{i1}| - \theta_2 x_{i2})^2$

- (c) [2 Pts] Your friend decides to take a different approach and instead fit a model using the following training data. Note that this dataset is identical to the original data on the previous page, with the addition of two new features x_3 and x_4 .

x_1	x_2	x_3	x_4	y
1	3	2	1/2	1
3/2	5	2	7	1
4	1	8	3	4

Which of the following is/are true when fitting a model to this training data?

- A. A linear regression model fitted using OLS would achieve a training mean squared error (MSE) of 0
 - B. A linear regression model fitted using OLS would achieve a training mean absolute error (MAE) of 0
 - C. A ridge regression model with $\lambda > 0$ would achieve a training MSE of 0
 - D. A ridge regression model with $\lambda > 0$ would achieve a training MAE of 0
 - E. None of the above
- (d) [4 Pts] You would rather use **gradient descent (GD)** to fit a model. You define the following loss function with one parameter θ_0 . You decide to first sample a single datapoint from your dataset. You then run gradient descent considering **only this sampled datapoint** at **every** GD update step. This single sampled datapoint is displayed below.

$$L(\theta_0) = (y - \theta_0)^4$$

x_1	x_2	x_3	x_4	y
1	3	2	1/2	1

Which of the following is/are true when using gradient descent to compute the optimal model parameter $\hat{\theta}_0$ with a learning rate of $\alpha = \frac{1}{4}$ for your new choice of loss function and data? Assume a starting guess of $\theta_0^{(0)} = 0$.

- A. The loss function is non-convex, so gradient descent won't converge to a solution
- B. The procedure described above is stochastic gradient descent
- C. The GD estimate for $\hat{\theta}_0$ will not update after the first iteration of the algorithm
- D. The solution for $\hat{\theta}_0$ produced by this gradient descent procedure achieves a training MSE of 0 on this single datapoint

3 Son of a Pitch(fork) [18 Pts]

Pitchfork is a music blog that reviews and rates music albums. For each music album, Pitchfork provides a rating (on a scale of 0 to 10), a review article, and a whether or not it deems the record “Best New Music”. Adrienne obtains a DataFrame `reviews` that contains all reviews Pitchfork has published from 1997 to 2021. `reviews` contains 26,036 rows and 5 columns. The columns are:

"URL": Simplified URLs of the review blog post, stored as `str`

"Score": Pitchfork’s rating of the record, on a scale of 0 to 10, stored as `float`

"Best New Music?": Pitchfork’s decision on whether or not the record is “the finest music of the current moment.” Only possible values are integers 1 or 0, with 1 representing “Yes” and 0 representing “No”

"Review body": The review article, stored as `str`

"Genre": the genre of the album, stored as `str`

The first 7 rows of `reviews` are shown below.

	URL	Score	Best New Music?	Review body	Genre
0	/reviews/albums/Eminem/928/Revival	5.0	0	Few texts in hip-hop...	Rap
1	/reviews/albums/Radiohead/2910/In_Rainbows	9.3	1	Like many music love...	Rock
2	/reviews/albums/Nicki_Minaj/6688/The_Pinkprint	7.5	0	Nicki Minaj is fed u...	Rap
3	/reviews/albums/Sufjan_Stevens/8020/Illinois	9.2	1	The best travel writ...	Folk/Country
4	/reviews/albums/Big_Thief/11808/U.F.O.F.	9.2	1	Big Thief are their ...	Rock
5	/reviews/albums/Ed_Sheeran/12622/+	2.8	0	Ed Sheeran needs you...	Pop/R&B
6	/reviews/albums/100_gecs/18330/1000_gecs	7.4	0	If you're the type t...	Experimental

- (a) [4 Pts] Adrienne realizes that there is no column for the title or the artist of each album. She notices that information about the title and artist is stored in the "URL" column of `reviews`. In particular, each entry in the "URL" column is in the format:

```
/reviews/albums/[artist]/[review id]/[album title]
```

You may assume that:

- The `review id` is a unique identifier for the review blogs, from 0 to 26035
- Spaces in the artists’ names and record titles are replaced by underscores (" _ ")
- There are no backslashes (" / ") in the artists’ names and titles.

For example, the first entry in the previewed DataFrame above corresponds to an album named “Revival” from the artist “Eminem”.

Complete the skeleton code to create a DataFrame named `artist_title` with two columns, one containing all the artists and one containing all the album titles.

```
pat = r"/reviews/albums/____A____"
artist_title = reviews["URL"].__B__
artist_title
```

The desired output for the first 7 rows of `reviews` is shown on the right.

	0	1
0	Eminem	Revival
1	Radiohead	In_Rainbows
2	Nicki_Minaj	The_Pinkprint
3	Sufjan_Stevens	Illinois
4	Big_Thief	U.F.O.F.
5	Ed_Sheeran	÷
6	100_gecs	1000_gecs

Note: you **do not** need to consider any edge cases beyond the example URLs shown in the first seven rows of `reviews`. As long as your code works for the "URL" strings shown in the DataFrame, and is not trivializing the problem, you are eligible for full credit.

(i) Fill in blank A:

(ii) Fill in blank B:

(b) [3 Pts] Define the **length** of a review article as the number of characters in that article.

Create a DataFrame named `genre_lengths` that is **indexed** by "Genre" and contains **two columns**: "Length" and "Review body". The "Review body" column should contain the review article that has the **longest length** out of all reviews for the corresponding genre. "Length" should contain the length of this longest review for the genre.

You may use as many or as few of the provided blanks below as you believe are necessary.

```
reviews["Length"] = reviews["Review body"]._____
```

```
genre_lengths = _____
```

Buck is also interested in the dataset and asks Adrienne to share it. Adrienne draws 100 records from `reviews` and gives the sample to Buck. For **all remaining parts** of this question, assume the records were drawn **uniformly at random with replacement** and **independently** of one another.

- (c) [2 Pts] Let X_i denote the value of the "Best New Music?" column for entry i in Buck's sample, for $i \in \{1, 2, 3, \dots, 100\}$. For example, if row 6 from the DataFrame on Page 3 is entry j in the sample, then $X_j = 0$.

"Best New Music?"	
population mean	0.1
population SD	0.2

Adrienne also gives Buck some population parameters for this column in `reviews`, displayed above.

What is the distribution of $S = \sum_{i=1}^{10} X_i$? State the distribution's **name** and **parameters**.

- (d) [3 Pts] Adrienne also gives Buck some **population parameters** for the "Score" column in `reviews`, displayed on the right. Let Y_i denote the "Score" value of entry i in Buck's sample, for $i \in \{1, 2, 3, \dots, 100\}$.

"Score"	
population mean	7.1
population SD	1.2

Assume the underlying data generating process for the "Score" column is a constant model: $Y = \theta + \epsilon$, where $\theta = 7.1$ is a fixed population parameter and ϵ is a random noise with expectation 0 and variance 1.44.

Buck uses the following model to fit his sample of Y_i 's: $\hat{Y} = \hat{\theta}$.

Buck estimates θ using the mean of the first 10 "Score" values: $\hat{\theta} = \frac{1}{10} \sum_{i=1}^{10} Y_i$. What is his **model bias**? Show all work and simplify your answer as much as possible.

model bias = _____

(e) [3 Pts] What is his **model variance**? Simplify your answer as much as possible.

model variance = _____

(f) [3 Pts] What is his **model risk**? Simplify your answer as much as possible.

Hint: You might find the results from the previous two parts helpful.

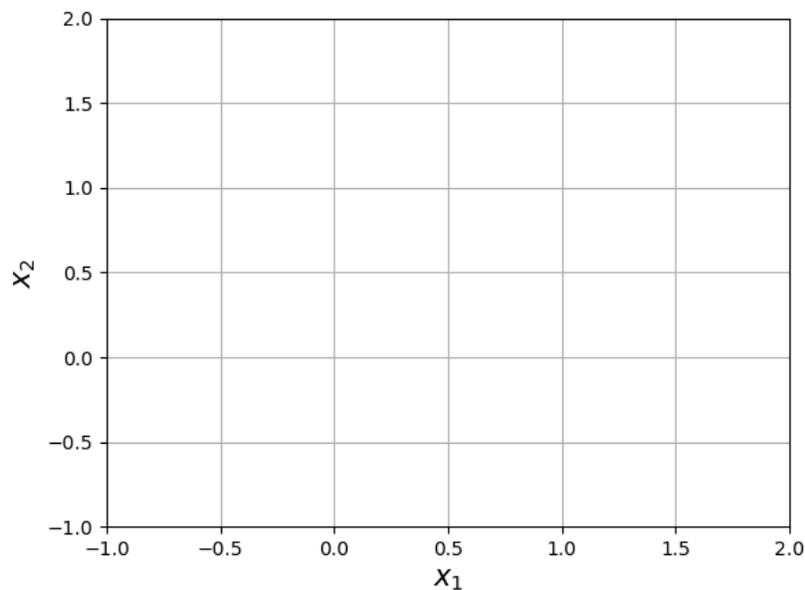
model risk = _____

4 Classifying XOR [12 Pts]

- (a) [2 Pts] We want to use **logistic regression** to train a binary classifier on four data points with two features: x_1 and x_2 .

x_1	x_2	y
0	0	0
0	1	1
1	0	1
1	1	0

- (i) Mark the four data points in the empty plot below. Write down the corresponding class (0 or 1) next to each point.



- (ii) Based on the plot from the previous part, is the dataset linearly separable?

Yes No

- (b) [2 Pts] We consider generating new features by transforming the original features x_1 and x_2 from Part (a). Which set of features lead to a **linearly separable** dataset?

- A. $\{x_1 - x_2, x_2\}$
- B. $\{x_1, x_1 + x_2\}$
- C. $\{x_1 + x_2, x_1x_2\}$
- D. $\{x_1, x_2^2\}$

This page has been intentionally left blank.

- (c) [4 Pts] We transform the original data to generate the new design matrix displayed below. We create a logistic regression model using a bias term, x_3 , and x_4 .

bias	x_3	x_4	y
1	1.0	-1.5	0
1	1.5	-2.0	0
1	0.5	3.0	1
1	0.0	1.0	0

We find that the optimal parameter vector for this model is as follows:

$$\hat{\theta} = [\log(1/2), \log(4/3), \log(1/3)]^\top$$

Now we want to classify a new data point: $x_* = [1, 1, -1]$. You may assume this new data point x_* is expressed in terms of the transformed features: [bias, x_3 , x_4].

- (i) Find the **predicted probability that point x_* is in class 1** $P(y_* = 1|x_*)$ and simplify your result as much as possible. Show all work for full credit. Recall that \log denotes the natural logarithm with base e . Hint: $\log(a) + \log(b) = \log(ab)$ and $-\log(a) = \log(\frac{1}{a})$

$P(y_* = 1|x_*) = \underline{\hspace{2cm}}$

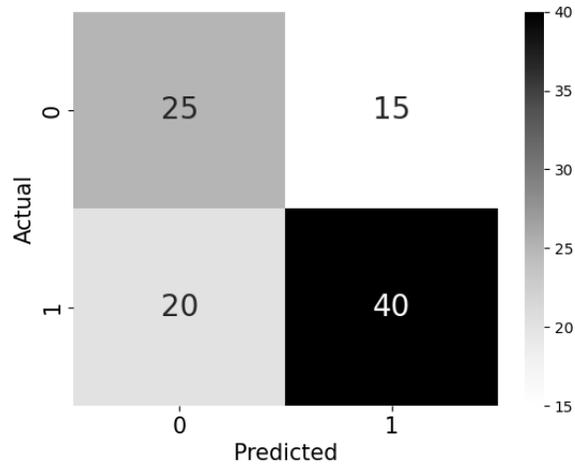
- (ii) For which threshold values T would the predicted class for x_* be $\hat{y}_* = 1$?

- A. 0 B. 0.3 C. 0.8 D. 1

- (d) [2 Pts] Now that our classifier has been trained, we wish to assess its performance via classification metrics. Select all of the **correct statements** from the list below.

- A. The area under a receiver operating characteristic (ROC) curve is always greater than or equal to 0.5
- B. The ROC curve of the ideal classifier passes through the (0, 1) point
- C. Decreasing the threshold T can never decrease the number of true negative predictions
- D. Increasing the threshold T can never decrease the number of false negative predictions

- (e) [2 Pts] We use the fitted logistic regression classifier from Part (c) to predict the labels of 100 new data points. The **confusion matrix** corresponding to the predictions of this classifier on the new data points is presented below.



Compute the **recall** of this classifier. Show all work for full credit.

Recall of logistic regression classifier: _____

5 The Shape of U [11 Pts]

Suppose we have collected a design matrix of data X . We wish to use principal component analysis (PCA) to reduce the dimensionality of our data.

- (a) [4 Pts] We use the **singular value decomposition (SVD)** to decompose the matrix X as follows:

$$X = U\Sigma V^T = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 4 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} & 0 & 1/\sqrt{2} \\ 0 & 1 & 0 \\ 1/\sqrt{2} & 0 & -1/\sqrt{2} \end{bmatrix}$$

Find the **second principal component (PC)** of X . Show all work for full credit.

Second PC: _____

- (b) [2 Pts] Which statements are **true** about principal component analysis performed on any matrix A with singular value decomposition $A = U\Sigma V^T$? Consider the general case where U , Σ and V^T may not have the specific values given in the previous part.

- A. All columns of the matrix A should have the same mean value to perform PCA
- B. Each principal component is a linear combination of the columns of A
- C. Each principal component is a linear combination of the columns of U
- D. An $n \times d$ matrix A has d non-zero singular values

- (c) [3 Pts] Assume we store the matrix X from Part (a) as the NumPy array X . **Fill in the blanks** to create a **scree plot** for the variance ratio captured by each each principal component of X . Assume `numpy` and `matplotlib.pyplot` are imported as `np` and `plt`, respectively.

```
U, S, Vt = np.linalg.svd(X, full_matrices=False)
numbers = np.arange(len(S))
y = _____A_____ / _____B_____
plt.plot(x = numbers, y = y)
```

- (i) Fill in blank A:

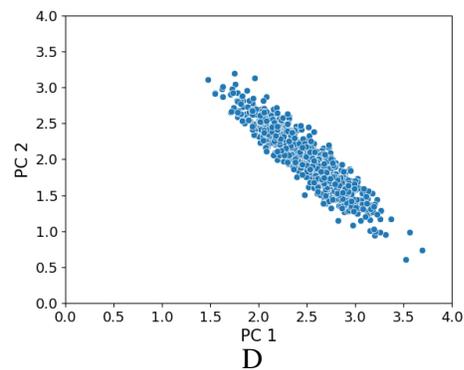
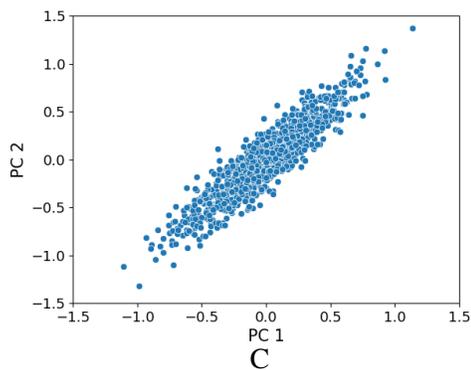
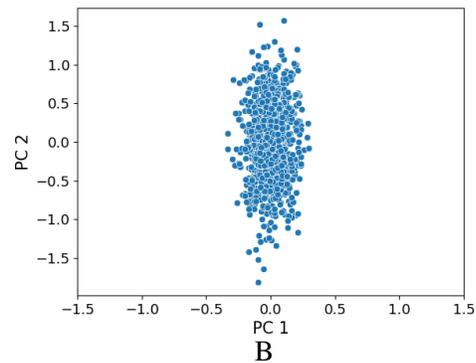
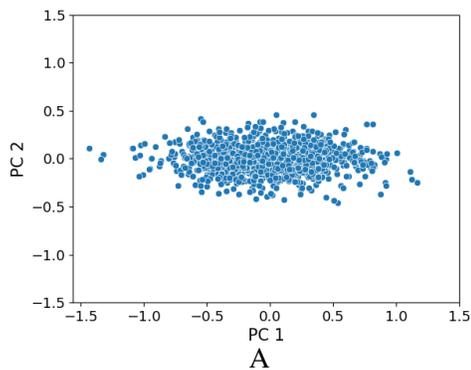
- (ii) Fill in blank B:

- (d) [2 Pts] Suppose we project a new matrix A onto the **directions of its first two principal components**. Which of the following plots could **possibly** display the projected data with the first PC plotted along the **horizontal axis** and the second PC plotted along the **vertical axis**? Select all that apply.

 A

 B

 C

 D


6 Yes or No [3 Pts]

The following statements are about decision trees or random forests. Determine if each of them is true or false. Each subpart is worth 0.5 points.

For Parts (a) and (b), consider a decision tree fitted to the dataset of two classes (0 and 1) shown below.

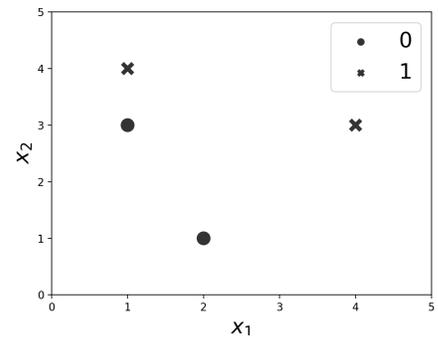
- (a) A decision tree will generate a linear decision boundary on this dataset.

True False

- (b) The top (root) node of a decision tree for this dataset has an entropy of 1.

True False

x_1	x_2	y
1	3	0
1	4	1
2	1	0
4	3	1



For Parts (c) to (f), consider the general behavior of decision trees and random forests on *any* data.

- (c) A decision tree with no restrictions on maximum depth always achieves 100% training accuracy on a linearly separable dataset.

True False

- (d) When training a decision tree, the maximum depth of the tree is a hyperparameter.

True False

- (e) In any random forest, each tree uses a random subset of some number m features to train the entire tree.

True False

- (f) A random forest will always achieve higher accuracy than a decision tree on the training set.

True False

7 Cross(words)-Validation [13 Pts]

Ella, an avid crossword enthusiast, wants to build a model to help her analyze crossword puzzles.

- (a) [2 Pts] Ella decides to build the following model with p numerical features and an intercept:

$$\hat{y} = \hat{\theta}_0 + \sum_{j=1}^p \hat{\theta}_j x_j$$

She uses **mean squared error with L1 regularization**, using a regularization hyperparameter of λ . Which of the following are correct objective function(s) for her model? Select all that apply.

- A. $\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 + \lambda \sum_{j=1}^p \theta_j^2$
 - B. $\frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| + \lambda \sum_{j=1}^p |\theta_j|$
 - C. $\frac{1}{n} \|\mathbb{Y} - \hat{\mathbb{Y}}\|_2^2 + \lambda \sum_{j=1}^p |\theta_j|$
 - D. $\frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| + \lambda \sum_{j=1}^p \theta_j^2$
- (b) [2 Pts] After performing feature engineering, Ella finds that the model's training error increases. Which of the following could **not** have been the reason? Select all that apply.
- A. Increasing the regularization penalty hyperparameter λ
 - B. Centering all features in the design matrix other than the bias column
 - C. Adding polynomial features
 - D. Removing the intercept/bias column from the design matrix
- (c) [2 Pts] Ella is worried that the variance of her model may be too large. Which of the following will **not** increase Ella's model variance? Select all that apply.
- A. Decreasing the regularization penalty hyperparameter λ to 0
 - B. Increasing the number of data points used for training
 - C. Using stochastic gradient descent to optimize the objective function instead of batch gradient descent
 - D. Introducing a new categorical variable by adding one-hot encoded columns
- (d) [2 Pts] Ella is not sure if she should use L1 or L2 regularization in her model. In which of the following scenarios is L1 regularization a **better** choice than L2? Select only one option.
- A. When the columns of the design matrix are collinear
 - B. When we have many features but don't know which are most helpful for modeling
 - C. When we want a closed-form analytical solution for optimal model parameters
 - D. When we have more features than data points

Ella wants to use gradient descent to determine the optimal parameter values for her regularized model. To do so, she needs to select a value for α , the gradient descent learning rate.

She performs **4-fold cross-validation** using a dataset of 60 observations to help her choose from three possible values of α . The validation errors computed in each fold of her cross-validation (CV) procedure are shown below. Assume that she is not considering a test set, so all 60 observations are used in the cross-validation procedure.

	Fold #1	Fold #2	Fold #3	Fold #4
$\alpha = 0.1$	2	10	3	5
$\alpha = 1$	6	1	2	3
$\alpha = 10$	4	9	1	2

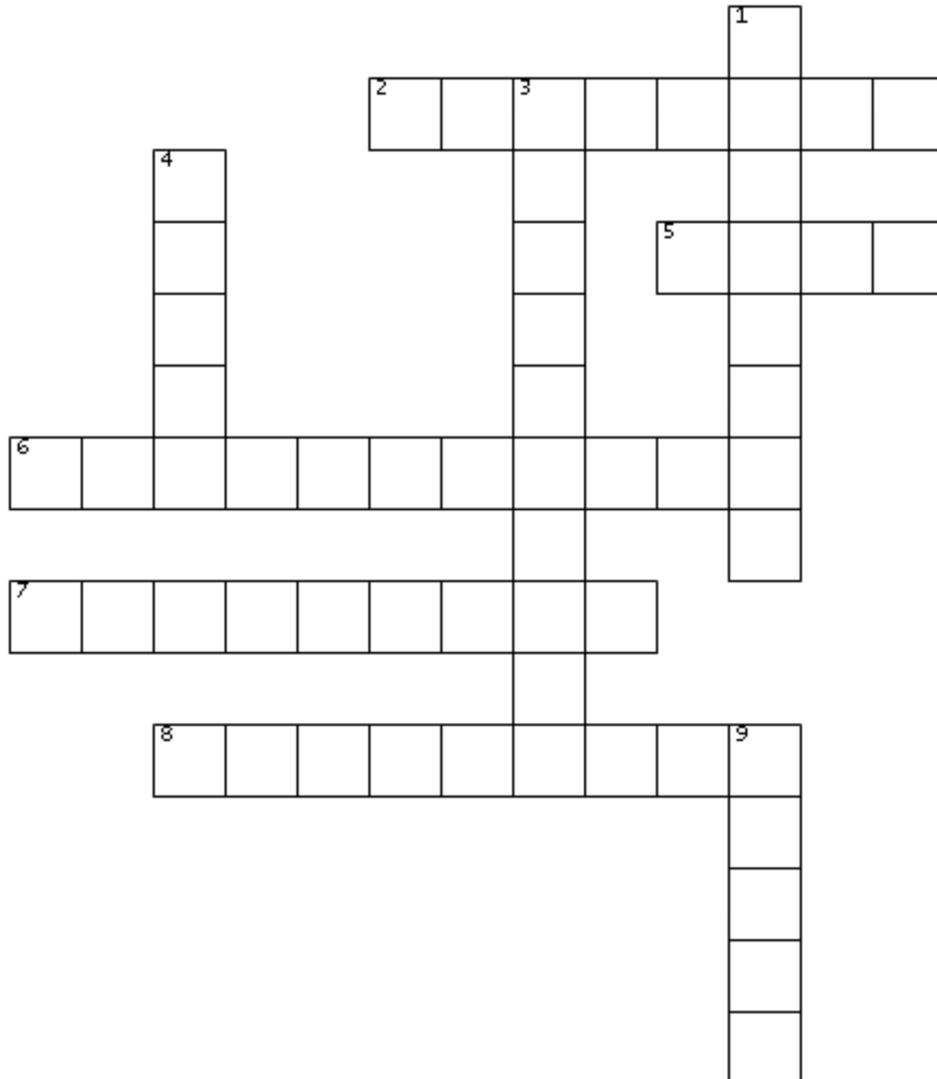
(e) [2 Pts] How many observations are in each validation fold of Ella's CV procedure?

Number of observations: _____

(f) [3 Pts] Given the validation errors displayed above, which value of α should Ella choose?

- A. $\alpha = 0.1$
- B. $\alpha = 1$
- C. $\alpha = 10$
- D. α should be determined using ordinary least squares, not cross-validation

- (g) [0 Pts] To celebrate completing her model, Ella creates the following crossword puzzle. Help her fill it out! Words can go across or down. Letters are shared when the words intersect. **This question is optional and worth 0 points.** You should only fill it out if you have completed all other questions in the exam.



Across

2. A metric for how much an estimator tends to vary from its mean value
5. A metric for how far off an estimator is from the true parameter
6. A tax system where the percentage tax rate is higher for individuals with higher net income
7. A numerical function of a random sample
8. A numerical function of a population

Down

1. When an estimator gets the population parameter right, on average
3. A tax system where the percentage tax rate is higher for individuals with lower net income
4. A regularization method that penalizes coefficients using L1 norm
9. A regularization method that penalizes coefficients using L2 norm

8 Congratulations [0 Pts]

Congratulations! You have completed the Final.

- **Make sure that you have written your student ID number on *every other page* of the exam.** You may lose points on pages where you have not done so.
- Also ensure that you have **signed the Honor Code** on the cover page of the exam for 1 point.
- If more than 10 minutes remain in the exam period, you may hand in your paper and leave. If ≤ 10 minutes remain, please **sit quietly** until the exam concludes.

Congrats on finishing the class! We're so happy to have spent the past few months with you.

[Optional, 0 pts] What did Data 100 “look” like this summer? Draw a picture!

A large empty rectangular box with a thin black border, intended for the student to draw a picture of what Data 100 "looked" like this summer.

Summer 2023 Data C100 Final Reference Sheet

Pandas

Suppose `df` is a `DataFrame`; `s` is a `Series`. `import pandas as pd`

Function	Description
<code>pd.read_csv(filepath, delimiter)</code>	Loads the data file at <code>filepath</code> into a <code>DataFrame</code> with column fields separated by <code>delimiter</code> ('', '\t')
<code>df[col]</code>	Returns the column labeled <code>col</code> from <code>df</code> as a <code>Series</code> .
<code>df[[col1, col2]]</code>	Returns a <code>DataFrame</code> containing the columns labeled <code>col1</code> and <code>col2</code> .
<code>s.loc[rows] / df.loc[rows, cols]</code>	Returns a <code>Series/DataFrame</code> with rows (and columns) selected by their index values.
<code>s.iloc[rows] / df.iloc[rows, cols]</code>	Returns a <code>Series/DataFrame</code> with rows (and columns) selected by their positions.
<code>s.isnull() / df.isnull()</code>	Returns boolean <code>Series/DataFrame</code> identifying missing values
<code>s.fillna(value) / df.fillna(value)</code>	Returns a <code>Series/DataFrame</code> where missing values are replaced by <code>value</code>
<code>s.isin(values) / df.isin(values)</code>	Returns a <code>Series/DataFrame</code> of booleans indicating if each element is in <code>values</code> .
<code>df.drop(labels, axis)</code>	Returns a <code>DataFrame</code> without the rows or columns named <code>labels</code> along <code>axis</code> (either 0 or 1)
<code>df.rename(index=None, columns=None)</code>	Returns a <code>DataFrame</code> with renamed columns from a dictionary <code>index</code> and/or <code>columns</code>
<code>df.sort_values(by, ascending=True)</code>	Returns a <code>DataFrame</code> where rows are sorted by the values in columns by
<code>s.sort_values(ascending=True)</code>	Returns a sorted <code>Series</code> .
<code>s.unique()</code>	Returns a <code>NumPy</code> array of the unique values
<code>s.value_counts()</code>	Returns the number of times each unique value appears in a <code>Series</code> , sorted in descending order of count.
<code>pd.merge(left, right, how='inner', on='a')</code>	Returns a <code>DataFrame</code> joining <code>left</code> and <code>right</code> on the column labeled <code>a</code> ; the join is of type <code>inner</code>
<code>left.merge(right, left_on=col1, right_on=col2)</code>	Returns a <code>DataFrame</code> joining <code>left</code> and <code>right</code> on columns labeled <code>col1</code> and <code>col2</code> .
<code>df.pivot_table(index, columns, values=None, aggfunc='mean')</code>	Returns a <code>DataFrame</code> pivot table where columns are unique values from <code>columns</code> (column name or list), and rows are unique values from <code>index</code> (column name or list); cells are collected <code>values</code> using <code>aggfunc</code> . If <code>values</code> is not provided, cells are collected for each remaining column with multi-level column indexing.
<code>df.set_index(col)</code>	Returns a <code>DataFrame</code> that uses the values in the column labeled <code>col</code> as the row index.
<code>df.reset_index()</code>	Returns a <code>DataFrame</code> that has row index 0, 1, etc., and adds the current index as a column.
<code>df.sample(n=1, replace=False)</code>	Returns <code>n</code> randomly sampled rows from <code>df</code> . By default, sampling is performed without replacement.

Let `grouped = df.groupby(by)` where `by` can be a column label or a list of labels.

Function	Description
<code>grouped.count()</code>	Return a <code>Series</code> containing the size of each group, excluding missing values
<code>grouped.size()</code>	Return a <code>Series</code> containing size of each group, including missing values
<code>grouped.mean().min().max()</code>	Return a <code>Series/DataFrame</code> containing mean/min/max of each group for each column, excluding missing values
<code>grouped.filter(f)</code> <code>grouped.agg(f)</code>	Filters or aggregates using the given function <code>f</code>

Text Wrangling and Regular Expressions

Pandas str methods

Function	Description
<code>s.str.len()</code>	Returns a Series containing length of each string
<code>s.str[a:b]</code>	Returns a Series where each element is a slice of the corresponding string indexed from <code>a</code> (inclusive, optional) to <code>b</code> (non-inclusive, optional)
<code>s.str.lower()/s.str.upper()</code>	Returns a Series of lowercase/uppercase versions of each string
<code>s.str.replace(pat, repl)</code>	Returns a Series that replaces occurrences of substrings matching the regex <code>pat</code> with string <code>repl</code>
<code>s.str.contains(pat)</code>	Returns a boolean Series indicating if a substring matching the regex <code>pat</code> is contained in each string
<code>s.str.extract(pat)</code>	Returns a Series of the first subsequence of each string that matches the regex <code>pat</code> . If <code>pat</code> contains capturing group(s), outputs a DataFrame with one integer-named column for each group.
<code>s.str.split(pat)</code>	Splits the strings in <code>s</code> at the delimiter <code>pat</code> . Returns a Series of lists, where each list contains strings of the characters before and after the split.

Regex patterns

Operator	Description	Operator	Description
<code>.</code>	Matches any character except <code>\n</code>	<code>*</code>	Matches preceding character/group zero or more times
<code>\</code>	Escapes metacharacters	<code>+</code>	Matches preceding character/group one or more times
<code> </code>	Matches expression on either side of expression; has lowest priority of any operator	<code>^</code>	Matches the beginning of the string
<code>\d, \w, \s</code>	Predefined character group of digits (0-9), alphanumerics (a-z, A-Z, 0-9, and underscore), or whitespace, respectively	<code>\$</code>	Matches the end of the string
<code>\D, \W, \S</code>	Inverse sets of <code>\d, \w, \s</code> , respectively	<code>()</code>	Capturing group or sub-expression
<code>{m}</code>	Matches preceding character/group exactly <code>m</code> times	<code>[]</code>	Character class used to match any of the specified characters or range (e.g. <code>[abcde]</code> is equivalent to <code>[a-e]</code>)
<code>{m, n}</code>	Matches preceding character/group at least <code>m</code> times and at most <code>n</code> times. If either <code>m</code> or <code>n</code> are omitted, set lower/upper bounds to 0 and ∞ , respectively	<code>[^]</code>	Invert character class; e.g. <code>[^a-c]</code> matches all characters except <code>a, b, c</code>

Python re methods

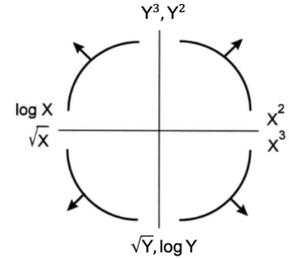
Function	Description
<code>re.match(pattern, string)</code>	Returns all matching characters if zero or more characters at beginning of <code>string</code> matches <code>pattern</code> , else <code>None</code>
<code>re.search(pattern, string)</code>	Returns all matching characters if zero or more characters anywhere in <code>string</code> matches <code>pattern</code> , else <code>None</code>
<code>re.findall(pattern, string)</code>	Returns a list of all non-overlapping matches of <code>pattern</code> in <code>string</code> (if none, returns empty list). If <code>pattern</code> includes capturing groups, only return captured characters.
<code>re.sub(pattern, repl, string)</code>	Returns <code>string</code> after replacing all occurrences of <code>pattern</code> with <code>repl</code>

Visualization

Matplotlib: x and y are sequences of values. `import matplotlib.pyplot as plt`

Function	Description
<code>plt.plot(x, y)</code>	Creates a line plot of x against y
<code>plt.scatter(x, y)</code>	Creates a scatter plot of x against y
<code>plt.hist(x, bins=None)</code>	Creates a histogram of x; bins can be an integer or a sequence
<code>plt.bar(x, height)</code>	Creates a bar plot of categories x and corresponding heights height

Tukey-Mosteller Bulge Diagram.



Seaborn: x and y are keyword arguments assigned to string column names in a DataFrame data. `import seaborn as sns`

Function	Description
<code>sns.countplot(data=None, x=None)</code>	Create a barplot of value counts of variable x from data
<code>sns.histplot(data=None, x=None, stat='count', kde=False)</code>	Creates a histogram of x from data, where bin statistics stat is one of 'count', 'frequency', 'probability', 'percent', and 'density'; optionally overlay a kernel density estimator.
<code>sns.displot(data=None, x=None, stat='count', rug=False, kde=True)</code>	
<code>sns.boxplot(data=None, x=None, y=None)</code> <code>sns.violinplot(data=None, x=None, y=None)</code>	Create a boxplot of a numeric feature (e.g., y), optionally factoring by a category (e.g., x), from data. <code>violinplot</code> is similar but also draws a kernel density estimator of the numeric feature
<code>sns.scatterplot(data=None, x=None, y=None)</code>	Create a scatterplot of x versus y from data
<code>sns.lmplot(data=None, x=None, y=None, fit_reg=True)</code>	Create a scatterplot of x versus y from data, and by default overlay a least-squares regression line
<code>sns.jointplot(data=None, x=None, y=None, kind)</code>	Combine a bivariate scatterplot of x versus y from data, with univariate density plots of each variable overlaid on the axes; kind determines the visualization type for the distribution plot, can be scatter, kde or hist
<code>sns.kdeplot(data=None, x=None)</code>	Create a kernel density estimate (KDE) of the distribution of x from data

Modeling

Concept	Formula	Concept	Formula
Variance, σ_x^2	$\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2$	Correlation r	$r = \frac{1}{n} \sum_{i=1}^n \frac{x_i - \bar{x}}{\sigma_x} \frac{y_i - \bar{y}}{\sigma_y}$
L_1 loss	$L_1(y, \hat{y}) = y - \hat{y} $	Linear regression estimate of y	$\hat{y} = \theta_0 + \theta_1 x$
L_2 loss	$L_2(y, \hat{y}) = (y - \hat{y})^2$	Least squares linear regression	$\hat{\theta}_0 = \bar{y} - \hat{\theta}_1 \bar{x} \quad \hat{\theta}_1 = r \frac{\sigma_y}{\sigma_x}$

Empirical risk with loss L

$$R(\theta) = \frac{1}{n} \sum_{i=1}^n L(y_i, \hat{y}_i)$$

Ordinary Least Squares

Multiple Linear Regression Model: $\hat{Y} = X\theta$ with design matrix X , response vector Y , and predicted vector \hat{Y} . If there are p features plus a bias/intercept, then the vector of parameters $\theta = [\theta_0, \theta_1, \dots, \theta_p]^T \in \mathbb{R}^{p+1}$. The vector of estimates $\hat{\theta}$ is obtained from fitting the model to the sample (X, Y) .

Concept	Formula	Concept	Formula
Mean squared error	$R(\theta) = \frac{1}{n} \ Y - X\theta\ _2^2$	Normal equation	$X^T X \hat{\theta} = X^T Y$
Least squares estimate, if X is full rank	$\hat{\theta} = (X^T X)^{-1} X^T Y$	Residual vector, e	$e = Y - \hat{Y}$
		Multiple R^2 (coefficient of determination)	$R^2 = \frac{\text{variance of fitted values}}{\text{variance of } y}$

Regularization

Concept	Formula	Concept	Formula
Ridge Regression L2 Regularization	$\frac{1}{n} \ Y - X\theta\ _2^2 + \lambda \ \theta\ _2^2$	Squared L2 Norm of $\theta \in \mathbb{R}^p$	$\ \theta\ _2^2 = \sum_{j=1}^p \theta_j^2$
Ridge regression estimate (closed form)	$\hat{\theta}_{\text{ridge}} = (X^T X + n\lambda I)^{-1} X^T Y$		
LASSO Regression L1 Regularization	$\frac{1}{n} \ Y - X\theta\ _2^2 + \lambda \ \theta\ _1$	L1 Norm of $\theta \in \mathbb{R}^p$	$\ \theta\ _1 = \sum_{j=1}^p \theta_j $

Gradient Descent

Let L be an objective function to minimize with respect to θ ; assume that some optimal parameter vector $\hat{\theta}$ exists. Suppose $\theta^{(0)}$ is some starting estimate at $t = 0$, and $\theta^{(t)}$ is the estimate at step t . Then for a learning rate α , the gradient update step to compute $\theta^{(t+1)}$ is:

$$\theta^{(t+1)} = \theta^{(t)} - \alpha \nabla_{\theta} L$$

where $\nabla_{\theta} L$ is the partial derivative/gradient of L with respect to θ , evaluated at $\theta^{(t)}$.

Classification and Logistic Regression

Confusion Matrix

Columns are the predicted values \hat{y} and rows are the actual classes y .

	$\hat{y} = 0$	$\hat{y} = 1$
$y = 0$	True negative (TN)	False Positive (FP)
$y = 1$	False negative (FN)	True Positive (TP)

Classification Performance

Suppose you predict n datapoints.

Metric	Formula	Other Names
Accuracy	$\frac{TP+TN}{n}$	
Precision	$\frac{TP}{TP+FP}$	
Recall/TPR	$\frac{TP}{TP+FN}$	True Positive Rate, Sensitivity
FPR	$\frac{FP}{FP+TN}$	False Positive Rate, Specificity

An ROC curve visualizes TPR vs. FPR for different thresholds T .

Logistic Regression Model: For input feature vector x , $\hat{P}_{\theta}(Y = 1|x) = \sigma(x^T \theta)$, where $\sigma(z) = 1/(1 + e^{-z})$. The estimate $\hat{\theta}$ is the parameter θ that minimizes the average cross-entropy loss on training data. For a single datapoint, define cross-entropy loss as $-[y \log(p) + (1 - y) \log(1 - p)]$, where p is the probability that the response is 1.

Logistic Regression Classifier: For a given input x and trained logistic regression model with parameter θ , compute $p = \hat{P}(Y = 1|x) = \sigma(x^T \theta)$. predict response \hat{y} with classification threshold T as follows:

$$\hat{y} = \text{classify}(x) = \begin{cases} 1 & p \geq T \\ 0 & \text{otherwise} \end{cases}$$

Scikit-Learn

Package: sklearn.linear_model

Linear Regression	Logistic Regression	Class/Function(s)	Description
✓	-	LinearRegression(fit_intercept=True)	Returns an ordinary least squares Linear Regression model.
-	✓	LogisticRegression(fit_intercept=True, penalty='l2', C=1.0)	Returns an ordinary least squares Linear Regression model. Hyperparameter C is inverse of regularization parameter, C = 1/λ.
✓	-	Lasso(), Ridge()	Returns a Lasso (L1 Regularization) or Ridge (L2 regularization) linear model, respectively.
✓	✓	model.fit(X, y)	Fits the scikit-learn model to the provided X and y.
✓	✓	model.predict(X)	Returns predictions for the X passed in according to the fitted model.
-	✓	model.predict_proba(X)	Returns predicted probabilities for X passed in according to the fitted model. If binary classes, returns probabilities for both classes 0 and 1.
✓	✓	model.coef_	Estimated coefficients for the linear model, not including the intercept.
✓	✓	model.intercept_	Bias/intercept term of the linear model. Set to 0.0 if fit_intercept=False.

Package: sklearn.model_selection

Function	Description
train_test_split(*arrays, test_size=0.2)	Returns two random subsets of each array passed in, with 0.8 of the array in the first subset and 0.2 in the second subset.

Probability

Let X have a discrete probability distribution $P(X = x)$. X has expectation $\mathbb{E}[X] = \sum_x xP(X = x)$ over all possible values x , variance $\text{Var}(X) = \mathbb{E}[(X - \mathbb{E}[X])^2]$, and standard deviation $\text{SD}(X) = \sqrt{\text{Var}(X)}$.

The covariance of two random variables X and Y is $\mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])]$. If X and Y are independent, then $\text{Cov}(X, Y) = 0$.

Notes	Property of Expectation	Property of Variance
X is a random variable.		$\text{Var}(X) = E[X^2] - (E[X])^2$
X is a random variable, $a, b \in \mathbb{R}$ are scalars.	$\mathbb{E}[aX + b] = a\mathbb{E}[X] + b$	$\text{Var}(aX + b) = a^2\text{Var}(X)$
X, Y are random variables.	$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$	$\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y) + 2\text{Cov}(X, Y)$
X is a Bernoulli random variable that takes on value 1 with probability p and 0 otherwise.	$\mathbb{E}[X] = p$	$\text{Var}(X) = p(1 - p)$

Central Limit Theorem

Let (X_1, \dots, X_n) be a sample of independent and identically distributed random variables drawn from a population with mean μ and standard deviation σ . The sample mean $\bar{X}_n = \sum_{i=1}^n X_i$ is normally distributed, where $\mathbb{E}[\bar{X}_n] = \mu$ and $\text{SD}(\bar{X}_n) = \sigma/\sqrt{n}$.

Model Risk

Suppose for each individual with fixed input x , we observe a random response $Y = g(x) + \epsilon$, where g is the true relationship and ϵ is random noise with zero mean and variance σ^2 .

For a new individual with fixed input x , define our random prediction $\hat{Y}(x)$ based on a model fit to our observed sample (\mathbb{X}, \mathbb{Y}) . The model risk is the mean squared prediction error between Y and $\hat{Y}(x)$: $\mathbb{E}[(Y - \hat{Y}(x))^2] = \sigma^2 + \left(\mathbb{E}[\hat{Y}(x)] - g(x)\right)^2 + \text{Var}(\hat{Y}(x))$.

SQL

```

SELECT [DISTINCT]
    { * | expr [[AS] c_alias]
    {, expr [[AS] c_alias] ... } }
FROM tableref {, tableref}
[[INNER | LEFT ] JOIN table_name ON qualification_list]
[WHERE search_condition]
[GROUP BY colname {, colname...}]
[HAVING search_condition]
[ORDER BY column_list]
[LIMIT number]
[OFFSET number of rows];
    
```

Syntax	Description
SELECT column_expression_list	List is comma-separated. Column expressions may include aggregation functions (MAX, SUM, COUNT, AVG, etc). AS renames columns. DISTINCT selects only unique rows.
WHERE a IN cons_list WHERE a IS NOT val WHERE a LIKE 'p'	<ul style="list-style-type: none"> Select rows for which the value in column a is among the values in a cons_list. Selects rows for which the value in column a is not equal to val (of any data type). Matches each entry in the column a to the text pattern p. The wildcard % matches at least zero characters. _ matches exactly one character.
ORDER BY RANDOM() LIMIT n	Draw a simple random sample of n rows.
ORDER BY a, b DESC	Order by column a (ascending by default) , then b (descending).
CASE WHEN pred THEN cons ELSE alt END	Evaluates to cons if pred is true and alt otherwise. Multiple WHEN/THEN pairs can be included, and ELSE is optional.
LIMIT number	Keep only the first number rows in the return result.
OFFSET number	Skip the first number rows in the return result.

Principal Component Analysis (PCA)

The i -th Principal Component of the matrix X is defined as the i -th column of $U\Sigma$ defined by Singular Value Decomposition (SVD).

$X = U\Sigma V^T$ is the SVD of X if U and V^T are matrices with orthonormal columns and Σ is a diagonal matrix. The diagonal entries of Σ , $[s_1, \dots, s_r, 0, \dots, 0]$, are known as singular values of X , where $s_i > s_j$ for $i < j$ and $r = \text{rank}(X)$.

Define the design matrix $X \in \mathbb{R}^{n \times p}$. Define the total variance of X as the sum of individual variances of the p features. The amount of variance captured by the i -th principal component is equivalent to s_i^2/n , where n is the number of datapoints.

Syntax	Description
np.linalg.svd(X, full_matrices = True)	SVD of X with shape (M, N) that returns u, s, vt, where s is a 1D array of X's singular values. If full_matrices=True, u and vt have shapes (M, M) and (N, N) respectively; otherwise shapes are (M, K) and (K, N), respectively, where $K = \min(M, N)$.

Decision Trees

Suppose you have a **decision tree** classifier for k classes. For each node, define the probability for class $C \in \{1, \dots, k\}$ as $p_C = d_C/d$, where d_C is the number of datapoints in class C (of the total d in the node). The entropy of the node (in bits) is defined as $S = -\sum_C p_C \log_2 p_C$, and the weighted entropy of the split is the average entropy of child nodes weighted by the number of datapoints in each.

Decision tree generation algorithm: all of the data starts in the root node. Repeat until every node is either pure or unsplitable.

- Pick the best feature x and best split value β to maximize the change in weighted entropy.
- Split data into two nodes, one where $x < \beta$, and one where $x \geq \beta$

A node that only has samples from one class is called a "pure" node. A node that has overlapping datapoints from different classes and thus cannot be split is called "unsplitable."

A **random forest** is a collection of many decision trees fit to variations of the same training data (e.g. bootstrapped samples, also called bagging). It is an ensemble method.