

Exam Review

Name:

Probability & Sampling


1. A political scientist is interested in answering a question about a country composed of three states with exactly 10000, 20000, and 30000 voting adults. To answer this question, a political survey is administered by randomly sampling 25, 50, and 75 voting adults from each town in each state, respectively. Which sampling plan was used in the survey?
(a) cluster sampling (b) stratified sampling (c) quota sampling (d) census
2. Suppose Sam visits your store to buy some items. He buys toothpaste for \$2.00 with probability 0.5. He buys a toothbrush for \$1.00 with probability 0.1. Let the random variable X be the total amount Sam spends. Find $\mathbb{E}[X]$.
3. Suppose we have a coin that lands heads 80% of the time. Let the random variable X be the *proportion* of times the coin lands tails out of 100 flips. What is $\text{Var}[X]$?

EDA & Visualization

4. For each of the following scenarios, determine which plot type is *most* appropriate to reveal the distribution of and/or the relationships between the following variable(s). For each scenario, select only one plot type. Some plot types may be used multiple times.
A. histogram B. pie chart C. bar plot D. line plot
E. side-by-side boxplots F. scatter plot G. stacked bar plot H. overlaid line plots
 - (a) Sale price and number of bedrooms for houses sold in Berkeley in 2010.
 - (b) Sale price and date of sale for houses sold in Berkeley between 1995 and 2015.
 - (c) Infant birth weight (grams) for babies born at Alta Bates hospital in 2016.
 - (d) Mother's education-level (highest degree held) for students admitted to UC Berkeley in 2016.
 - (e) SAT score and HS GPA of students admitted to UC Berkeley in 2016.
 - (f) The percentage of female student admitted to UC Berkeley each year from 1950 to 2000.
 - (g) SAT score for males and females of students admitted to UCB from 1950 to 2000

Prediction

5. When developing a model for a donkey's weight, we consider the following box plots of weight by age category.



This plot suggests:

- (a) Age is not needed in the model
- (b) Some of the age categories can be combined
- (c) Age could be treated as a numeric variable
- (d) None of the above

Optimization

6. Fix the following buggy Python implementation of gradient descent:

```

1 def grad_descent(X, Y, theta0, grad_function, max_iter = 1000):
2     """X: A 2D array, the feature matrix.
3     Y: A 1D array, the response vector.
4     theta0: A 1D array, the initial parameter vector.
5     grad_function: Maps a parameter vector, a feature matrix, and a
6         response vector to the gradient of some loss function at the
7         given parameter value. The return value is a 1D array."""
8     theta = theta0
9     for t in range(1, max_iter+1):
10        grad = grad_function(theta, X, Y)
11        theta = theta0 + t * grad
12    return grad

```

7. Suppose you are given a dataset $\{(x_i, y_i)\}_{i=1}^n$ where $x_i \in \mathbb{R}$ is a one dimensional feature and $y_i \in \mathbb{R}$ is a real-valued response. You use f_θ to model the data where θ is the model parameter. You choose to use the following regularized loss:

$$L(\theta) = \frac{1}{n} \sum_{i=1}^n (y_i - f_\theta(x_i))^2 + \lambda \theta^2$$

- (a) This regularized loss is best described as:
- Average absolute loss with L^2 regularization.
 - Average squared loss with L^1 regularization.
 - Average squared loss with L^2 regularization.
 - Average Huber loss with λ regularization.
- (b) Suppose you choose the model $f_\theta(x_i) = \theta x_i^3$. Using the above objective derive the loss minimizing estimate for θ .

Inference

8. **True or False.** Determine whether the following statements are true or false.
- (a) Suppose we have 100 samples drawn independently from a population. If we construct a 95% confidence interval for each sample, we expect 95 of them to include the **sample** mean.
 - (b) We often prefer a pseudo-random number generator because our simulations results can be exactly reproduced by controlling the seed.
9. Suppose we have a Pandas Series called **thePop** which contains a census of **25000 subjects**. We also have a simple random sample of **400 individuals** saved in the Series **theSample**. We are interested in studying the behavior of the bootstrap procedure on the simple random sample. Fill in the blanks in the code below to construct **1000 bootstrapped estimates** for the **median**.

```
boot_stats = [  
    _____  
    .sample(n = _____, replace = _____)  
    ._____()  
  
    for j in range(_____)  
]
```