

Section 5

Based on the work of many CS109 staffs

Overview of Section Materials

The warmup questions provided will help students practice concepts introduced in lectures. The section problems are meant to apply these concepts in more complex scenarios similar to what you will see in problem sets and quizzes.

1 Warmups


1.1 Conditional Expectation

Let $X \sim Geo(p)$. Use the Law of Total Expectation to prove that $E[X] = 1/p$, by conditioning on whether the first flip is heads or tails.

1.2 General Inference

Suppose X_1, \dots, X_4 are discrete random variables. We will abuse notation and write $p(x_1, x_2, x_3, x_4)$ to represent $P(X_1 = x_1, X_2 = x_2, X_3 = x_3, X_4 = x_4)$. In your answers, feel free to do the same. For example, $p(x_1, x_3) = P(X_1 = x_1, X_3 = x_3)$. For the following cases, decompose into four terms, with each being as simple as possible.

1. If there is no assumption of independence, what is $p(x_1, x_2, x_3, x_4)$?
2. If all variables are assumed independent, what is $p(x_1, x_2, x_3, x_4)$?
3. Assuming the variables follow the Bayesian network structure below, what is $p(x_1, x_2, x_3, x_4)$?



2 Problems

2.1 Random Number of Random Variables

Let N be a non-negative integer-valued random variable; that is, takes values in $\{0, 1, 2, \dots\}$. Let X_1, X_2, X_3, \dots be an infinite sequence of iid random variables (independent of N), each with mean μ , and $X = \sum_{i=1}^N X_i$ be the sum of the first N of them.

Before doing any work, what do you think $E[X]$ will turn out to be? Show it mathematically.

2.2 Fairness in AI.

In their 2018 paper “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification,” Joy Buolamwini and Timnit Gebru showed that commercial gender classifiers performed significantly worse on face images of darker-skinned females (error rates up to 34.7%) than lighter-skinned males (maximum error rate of 0.8%.) This disparity may result from training the classifiers on unbalanced datasets. To evaluate the classifiers, the authors designed their own dataset by collecting photos of national parliamentarians from three African countries and three European ones.

The probability table below shows the joint distribution of the dataset between two random variables: the demographic (D) of the photo subject and their country (C).

Demographic	South Africa	Senegal	Rwanda	Sweden	Finland	Iceland
Darker Female	0.12	0.05	0.04	0.01	0	0
Darker Male	0.15	0.07	0.02	0.01	0	0
Lighter Female	0.02	0	0	0.12	0.06	0.02
Lighter Male	0.05	0	0	0.14	0.09	0.03

- What is the marginal probability distribution for demographic D ? Provide your result as a mapping from values that D can take to probabilities.
- What is the conditional probability of country given that the subject is a lighter female, $P(C|D = \text{Lighter Female})$? Provide your result as a mapping from values that C can take to probabilities. Is this mapping a probability distribution?
- What is the conditional probability that the subject is from Senegal given their demographic, $P(C = \text{Senegal}|D)$? Provide your answer as a mapping from values that D can take to probabilities. Is this mapping a probability distribution?
- What are the pitfalls in using this dataset for a purpose beyond what the authors intended?

2.3 Binary Trees

Consider the following function for constructing binary trees:

```
struct Node {
    Node *left;
    Node *right;
};

Node *randomTree(float p) {
    if (randomBool(p)) { // returns true with probability p
        Node *newNode = new Node;
        newNode->left = randomTree(p);
        newNode->right = randomTree(p);
        return newNode;
    } else {
        return nullptr;
    }
}
```

The `if` branch is taken with probability p (and the `else` branch with probability $1 - p$). A tree with no nodes is represented by `nullptr`; so a tree node with no left child has `nullptr` for the left field (and the same for the right child).

Let X be the number of nodes in a tree returned by `randomTree`. You can assume $0 < p < 0.5$. What is $E[X]$, in terms of p ?