

Will Monroe
CS 109

August 12, 2017

Practice Final Examination

Problems by Mehran Sahami and Chris Piech

Final exam: Saturday, August 19, 12:15pm-3:15pm in NVIDIA Auditorium

This is a closed book, closed calculator, closed computer exam. You are, however, allowed to use two 8½" × 11" page (front and back; or four pages, one side only) of notes in the exam. The last page of the exam is a Standard Normal Table, in case you need it.

You have 3 hours (180 minutes) to take the exam. The exam is 180 points, meant to roughly correspond to one point per minute of the exam. You may want to use the point allocation for each problem as an indicator for pacing yourself on the exam.

In the event of an incorrect answer, any explanation you provide of how you obtained your answer can potentially allow us to give you partial credit for a problem. For example, describe the distributions and parameter values you used, where appropriate. It is fine for your answers to include summations (but not integrals), products, factorials, exponentials, and combinations, unless the question specifically asks for a numeric quantity or closed form. Where numeric answers are required, the use of fractions is fine.

THE STANFORD UNIVERSITY HONOR CODE

- A. The Honor Code is an undertaking of the students, individually and collectively:
 - (1) that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;
 - (2) that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.
- B. The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid as far as practicable, academic procedures that create temptations to violate the Honor Code.
- C. While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to create optimal conditions for honorable academic work.

I acknowledge and accept the letter and spirit of the Honor Code:

Signature: _____

Name (print): _____

SUNetID: _____

Email (preferably as used on Gradescope): _____

1. [15 points] Say we have a standard 52-card deck of cards, where the cards have the usual ordering of ranks: 2, 3, 4, 5, 6, 7, 8, 9, 10, Jack, Queen, King, and Ace. Assume each card in the deck is equally likely to be drawn. Two cards are then drawn sequentially (without replacement) from the deck.

What is the probability that the second card drawn has a rank *greater than* the rank of the first card drawn? Show the derivation you used to compute the probability and also provide your final answer in closed form (if a closed form exists).

2. [20 points] You decide to buy one share of stock in “Happy Computing Inc.” (abbreviated: HCI) for \$10. Each day, the price of HCI either increases \$1 (with probability p) or decreases \$1 (with probability $1 - p$). You decide that you will only sell your stock at some point if *either* it gains \$2 (i.e., reaches a price of \$12) or loses \$2 (i.e., reaches a price of \$8).
- a. [7 points] What is the probability that you will sell your share of HCI (for either a gain or a loss of \$2) *exactly* 4 days after you buy it?
- b. [13 points] Assuming you just bought a share of HCI, so you don't know how long it will be until you sell it, what is the probability (in terms of p) that when you eventually sell it, you will sell it for a gain (i.e., sell it for \$12)? Give a closed form answer.

3. [20 points] Say you are monitoring energy usage for two different types of computers (type W and type X) in an office. There are 10 machines of type W and 10 machines of type X in the office, and the energy usage of all 20 machines are independent.

Let W_i = the number of kilowatt-hours used by machine i of type W in a day. Let X_j = the number of kilowatt-hours used by machine j of type X in a day.

Say that $W_i \sim \text{Poi}(4)$ for $1 \leq i \leq 10$ and $X_j \sim \text{N}(5, 3)$ for $1 \leq j \leq 10$.

Now say that you have a separate machine that (probabilistically) monitors the energy usage of the 10 type W and 10 type X machines (call these 20 machines the machine pool). Each day, a new “watch list” is generated, where each machine in the machine pool independently has a 0.2 probability of being placed on the “watch list”. Let Y = the total energy usage (in kilowatt-hours) of all the machines on the “watch list”.

- a. [4 points] What is $E[Y]$?

- b. [8 points] Say that three type W (and no type X) machines are on the “watch list” on a particular day. What is $P(Y \geq 20)$ that day?

- c. [8 points] Now say that three type X (and no type W) machines are on the “watch list” on a particular day. What is $P(Y \geq 20)$ that day? Give a numeric answer, if possible.

4. [20 points] We have a random number generator that generates random *integers* independently uniformly distributed from 1 to 5, inclusive. This random number generator is used to generate a sequence of n independent random integers.

Let X = number of 1's generated in the sequence of n integers.

Let Y = number of 5's generated in the sequence of n integers.

(Hint: it might be useful to define indicator variables denoting whether the i -th integer generated is a 1, and likewise for whether the i -th integer generated is a 5.)

- a. [13 points] What is $\text{Cov}(X, Y)$?

- b. [7 points] What is the correlation $\rho(X, Y)$?

5. [25 points] Say that our favorite manufacturer of computers produces machines whose lifetimes (independently) each have a mean of 100 weeks with a variance of 25 weeks². We buy and use a computer built by this manufacturer until it dies. When the machine dies, we immediately buy a new machine to use by the same manufacturer, and repeat this process every time the machine we are using dies. Let X = the total number of weeks that we use a machine built by our favorite manufacturer.

What is the minimal number of machines we would have to buy so that the total number of weeks we use a machine built by our favorite manufacturer is greater than 2000 with 95% probability? Justify your answer mathematically.

(Hint: derive an expression for the number of machines that will be necessary to purchase in order to have $P(X > 2000) \geq 0.95$, and use that to determine the answer.)

6. [25 points] Consider the following recursive functions:

```
int near() {
    int b = randomInteger(1, 4); // equally likely to be 1, 2, 3 or 4
    if (b == 1) return 2;
    else if (b == 2) return 4;
    else if (b == 3) return (6 + near());
    else return (8 + near());
}

int far() {
    int a = randomInteger(1, 3); // equally likely to be 1, 2 or 3
    if (a == 1) return 2;
    else if (a == 2) return (2 + near());
    else return (4 + far());
}
```

Let Y = the value returned by `far()`.

a. [10 points] What is $E[Y]$? Give a numeric answer.

b. [15 points] What is $\text{Var}(Y)$? Give a numeric answer.

7. [20 points] Consider a hash table with n buckets. Assume that all buckets of the hash table are equally likely to get hashed to. Let X = the number of strings that are hashed up to and including the first time a string is hashed to a non-empty bucket (i.e., a “collision” occurs).

a. [6 points] Say the number of buckets in the hash table is $n = 4$. Calculate $P(X = k)$ for $1 \leq k \leq 5$. In other words, calculate the probability that *exactly* k strings are hashed at the time that the first string is hashed to a non-empty bucket.

$$P(X = 1) =$$

$$P(X = 2) =$$

$$P(X = 3) =$$

$$P(X = 4) =$$

$$P(X = 5) =$$

b. [7 points] Still considering the case where $n = 4$, calculate $E[X]$.

c. [7 points] Now considering the general case of a hash table with n buckets, give a well-defined mathematical expression to compute $E[X]$ in terms of n (for arbitrary positive integer values of n). It’s fine for your expression to include summations, products, etc., as long as it is fully defined.

8. [15 points] Consider a sample of independently and identically distributed (I.I.D.) random variables X_1, X_2, \dots, X_n , that each have geometric distributions. In other words, $X_i \sim \text{Geo}(p)$ for all $1 \leq i \leq n$.
- a. [11 points] Derive the maximum likelihood estimate for the parameter p of the geometric distribution. Give a numeric answer. Explicitly show all the steps in your derivation. (For this problem, the derivation is worth many more points than the final answer, so please be clear how you derived your final answer.)
- b. [4 points] Say that we have a sample of five such I.I.D. geometric variables with the following values: $X_1 = 4, X_2 = 3, X_3 = 4, X_4 = 2, X_5 = 7$. What value of p in the geometric distribution would maximize the likelihood of these observations? Give a numeric answer.

9. [20 points] We are able to observe three real-valued input variables X_1 , X_2 , and X_3 and want to predict a single binary output variable Y (which can have values 0 or 1). We know the functional forms for the input variables are all uniform distributions, namely: $X_1 \sim \text{Uni}(a_1, b_1)$, $X_2 \sim \text{Uni}(a_2, b_2)$, and $X_3 \sim \text{Uni}(a_3, b_3)$, but we are not given the values of the parameters a_1, b_1, a_2, b_2, a_3 , or b_3 . We are, however, given the following dataset of 8 training instances:

X_1	X_2	X_3	Y
0.1	0.8	0.4	0
0.7	0.6	0.1	0
0.3	0.7	0.2	0
0.4	0.4	0.6	0
0.8	0.2	0.5	1
0.5	0.7	0.8	1
0.9	0.4	0.7	1
0.6	0.6	0.4	1

- a. [10 points] Use maximum likelihood estimation to estimate the parameters a_1, b_1, a_2, b_2, a_3 , and b_3 in the case where $Y = 0$ as well as the case $Y = 1$. (That is, estimate the distribution $P(X_i | Y)$ for $i = 1, 2$, and 3.) Note that the parameter values for a_1, b_1, a_2, b_2, a_3 , and b_3 may be different when $Y = 0$ versus when $Y = 1$.

- b. [10 points] You are given the following 3 testing instances, numbered 1, 2 and 3. (Note that the testing instances do not have output variable Y specified).

	X_1	X_2	X_3
test instance 1	0.5	0.6	0.4
test instance 2	0.7	0.7	0.7
test instance 3	0.5	0.4	0.3

Using the Naïve Bayes assumption and your probability estimates from part (a), predict the output variable Y for each instance (you should have 3 predictions). Show how you derived the prediction by showing the computations you made for each test instance.

