


## CS109 Final Exam

---

This is a closed calculator/computer exam. You are, however, allowed to use notes in the exam. The last page of the exam is a Standard Normal Table, in case you need it. The penultimate page has some common random variables. 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, 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.



I acknowledge and accept the letter and spirit of the honor code. I pledge to write more neatly than I have in my entire life:

Signature: \_\_\_\_\_

Family Name (print): \_\_\_\_\_

Given Name (print): \_\_\_\_\_

Email (preferably your gradescope email): \_\_\_\_\_





c. (8 points) You estimate that  $P(M|A) = \frac{1}{5}$  and that  $P(M|A^C) = \frac{1}{10}$ . What is  $P(A|M)$ ? You may leave your answer in terms of  $P(A)$ . To make your future life easier, simplify your answer as much as possible.

d. (8 points) Using the same assumption as the previous part, calculate  $P(A)$ . You may use a function `root(a, b, c)` which returns the value of  $x$  such that  $ax^2 + bx + c = 0$ .



- e. (10 points) Prove that  $P(A) = P(A|M^C)$  if people who did not respond are just as likely to vote for Ayesha as those who did:  $P(A|M) = P(A|M^C)$ .



## 2 Occupy Mt Gox [24 points]

Modern nations have such high disparity in incomes that we need to use special fat-tailed random variables to represent country wide distributions. Enter the ExamPareto distribution, a simplified version of the Pareto income distribution.  $X \sim \text{ExamPareto}(\alpha)$  represents a countries distribution of incomes—in Bitcoins (B):

$$\begin{array}{ll} \text{Pdf:} & f_X(x) = \begin{cases} \alpha(1+x)^{-(\alpha+1)} & x \geq 0 \\ 0 & \text{else} \end{cases} & \text{Mean:} & E[X] = \frac{1}{\alpha-1} \\ \text{CDF:} & F_X(x) = \begin{cases} 1 - (1+x)^{-\alpha} & x \geq 0 \\ 0 & \text{else} \end{cases} & \text{Variance:} & \text{Var}(X) = \frac{\alpha}{(\alpha-1)^2(\alpha-2)} \end{array}$$

- a. (4 points) Let  $X \sim \text{ExamPareto}(\alpha = 2)$ . What is the probability that  $X$  is greater than 4 B?
- b. (6 points) Let  $X \sim \text{ExamPareto}(\alpha = 3)$ . If we sample 100 IID individual incomes from  $X$ , approximate the probability that the sum of their incomes is greater than 55 B?





c. (4 points) You have samples of  $n$  IID incomes from a country  $(x_1, x_2, \dots, x_n)$ . Write an equation for the log likelihood of the  $n$  samples.

d. (10 points) Write pseudocode that can estimate the parameter  $\alpha$  from the  $n$  samples. You may use any parameter estimation technique we talked about in class. Make sure to work out any necessary math first.







c. (7 points) First the hacker needs to find out the length of the password. What is the probability that the time taken to test a guess of correct length (server executes 6 lines) is longer than the time taken to test a guess of an incorrect length (server executes 4 lines)? Assume that the first letter of the guess does not match the first letter of the password. Hint  $P(A > B)$  is the same as  $P(A - B > 0)$ .

d. (5 points) Let  $p$  be your answer to part (c). The hacker first tries a guess of length 5, then a guess of length 6, and so on up until a guess of length 10. What is the probability that the guess of length 7, which is the correct length (server executes 6 lines), takes longer than all the other guesses of incorrect length (where the server executes 4 lines)?



- e. (8 points) Now that our hacker knows the length of the password, to get the actual string she is going to try and figure out each letter one at a time, starting with the first letter. The hacker tries the string “aaaaaa” and it takes 27s. Based on this timing, how much more probable is it that first character matched (server executes 8 lines) than the first character did not match (server executes 6 lines)? Assume that all letters in the alphabet are equally likely to be the first letter.

- f. (3 points) If it takes the hacker 5 guesses to find the length of the password, and 26 guesses per letter to crack the password string, how many attempts must she make to crack our password, “gobayes”? Yikes!









Method	Description
<code>V = sampleBeta(a, b)</code>	Returns a real number value in the range $[0, 1]$ with probability defined by a PDF of a Beta with parameters $a$ and $b$ .
<code>R = giveDrug(i)</code>	Gives drug $i$ to the next patient. Returns a True if the drug was successful in curing the patient or False if it was not. Throws an error if $i \notin \{1, 2\}$ .
<code>I = argmax(list)</code>	Returns the index of the largest value in the list.

- c. (10 points) Write pseudocode to administer either of the two drugs to 100 patients using Thompson's Sampling Algorithm. Use functions from the table above. Your code should execute `giveDrug` 100 times.



- d. (4 points) After running Thompsons' Algorithm 100 times, you end up with the following Beta distributions:

$$\theta_1 \sim \text{Beta}(11, 11),$$

$$\theta_2 \sim \text{Beta}(76, 6)$$

What is the expected probability of success for each drug?

- e. (8 points) Write pseudocode to compute the p-value for the difference in expected probability of success between drug2 and drug1 observed from the 100 samples in part (d). Ie, under the null hypothesis assumption that all 100 samples are identically distributed, calculate the probability of a difference in expected probabilities that is greater than or equal to the one observed. Describe, don't implement any helper methods you want to use.



## 5 Expected Maximum Uniform [12 points]

Let  $X \sim \text{Uni}(0, 1)$  and  $Y \sim \text{Uni}(0, 1)$  be independent random variables. What is  $E[\max(X, Y)]$ ? Gain most of the points by writing an expression that you could solve using Wolfram Alpha. Provide a numerical answer for full credit.





## 6 Gaussian Naïve Bayes [16 points]

The version of Naïve Bayes that we used in class worked great when the feature values were all binary. If instead they are continuous, we are going to have to rethink how we estimate of the probability of the  $i$ th feature given the label,  $P(X_i|Y)$ . The ubiquitous solution is to make the *Gaussian Input Assumption* that:

$$\text{If } Y = 0, \text{ then } X_i \sim N(\mu_{i,0}, \sigma_{i,0}^2)$$

$$\text{If } Y = 1, \text{ then } X_i \sim N(\mu_{i,1}, \sigma_{i,1}^2)$$

For each feature, there are 4 parameters (mean and variance for both class labels). There is a final parameter,  $p$ , which is the estimate of  $P(Y = 1)$ . Assume that you have trained on data with **two** input features and have **already estimated** all 9 parameter values:

Feature $i$	$\mu_{i,0}$	$\mu_{i,1}$	$\sigma_{i,0}^2$	$\sigma_{i,1}^2$
1	5	0	1	1
2	0	3	1	4

$$p = 0.6$$


Write an inequality to predict whether  $Y = 1$  for input  $[X_1 = 5, X_2 = 3]$ . Use the Naïve Bayes assumption and the Gaussian Input Assumption. Your expression must be in terms of the learned parameters (either using numbers or symbols is fine):







- d. (10 points) Saliency is much more interesting for deep learning networks. Consider a simple deep learning network with **two** hidden neurons that predicts  $P(Y = 1 | \mathbf{X} = \mathbf{x}) \approx \hat{y}$ :



$$h_1 = \sigma\left(\sum_{i=1}^n \theta_{i,1}^{(h)} x_i\right)$$

$$h_2 = \sigma\left(\sum_{i=1}^n \theta_{i,2}^{(h)} x_i\right)$$

$$\hat{y} = \sigma\left(\theta_1^{(y)} h_1 + \theta_2^{(y)} h_2\right)$$

Calculate the saliency of a single feature ( $x_i$ ) in a training example  $(\mathbf{x}, y)$ .

- e. (5 points) Weights are not invariant to units of measurement. Imagine your logistic regression classifier has weight values:  $[\theta_0 = 10, \theta_1 = -4, \theta_2 = 8]$ . What values do you expect for the **weights** in a newly trained model if we change the units of the last feature ( $x_2$ ) to be centimeters instead of meters? That means all the values  $x_2$  will be 100x times as large.

$$\theta_0 =$$

$$\theta_1 =$$

$$\theta_2 =$$

[End of exam]

*That's all folks! Missing data requires probability theory to understand. Income distributions give us a way to measure inequality. Timing attacks are a real thing... watch your code. Thompson Sampling is a hot new algorithm for N-Armed-Bandit questions, a key part of MCTS, the algorithm used in AlphaGoZero. Gaussian Naïve bayes is good times for continuous inputs and saliency is a new tool for understanding deep neural nets. Thank you all for the wonderful quarter.*

## Continuous Random Variable Reference

### Uniform

All intervals in the range (of the same length) are equally probable.

---

Notation	$X \sim \text{Uni}(\alpha, \beta)$
----------	------------------------------------

---

Range(X):	$x \in \mathbb{R}, a \leq x \leq b$
-----------	-------------------------------------

---

pdf:	$f(X = k) = \frac{1}{b-a}$
------	----------------------------

---

E[X]:	$\frac{1}{2}(a + b)$
-------	----------------------

---

Var(X):	$\frac{1}{12}(b - a)^2$
---------	-------------------------

---

### Normal

Often used to represent real-valued random variables whose distributions are not known.

---

Notation	$X \sim N(\mu, \sigma^2)$
----------	---------------------------

---

Range(X):	$x \in \mathbb{R}$
-----------	--------------------

---

pdf:	$f(X = k) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(k-\mu)^2}{2\sigma^2}}$
------	---

---

cdf:	$F(k) = \Phi\left(\frac{k-\mu}{\sigma}\right)$
------	--

---

E[X]:	$\mu$
-------	-------

---

Var(X):	$\sigma^2$
---------	------------

---

### Beta

Often used to represent a distribution of belief of a probability.

---

Notation	$X \sim \text{Beta}(a, b)$
----------	----------------------------

---

Range(X):	$x \in \mathbb{R}, 0 \leq x \leq 1$
-----------	-------------------------------------

---

pdf:	$f(X = k) = x^{(\alpha-1)}(1-x)^{(\beta-1)}$
------	--

---

E[X]:	$\frac{\alpha}{\alpha+\beta}$
-------	-------------------------------

---

Var(X):	$\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$
---------	--

---



### Standard Normal Table

An entry in the table is the area under the curve to the left of  $z$ ,  $P(Z \leq z) = \Phi(z)$ .



<b>Z</b>	<b>0.00</b>	<b>0.01</b>	<b>0.02</b>	<b>0.03</b>	<b>0.04</b>	<b>0.05</b>	<b>0.06</b>	<b>0.07</b>	<b>0.08</b>	<b>0.09</b>
<b>0.0</b>	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359
<b>0.1</b>	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753
<b>0.2</b>	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141
<b>0.3</b>	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517
<b>0.4</b>	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879
<b>0.5</b>	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224
<b>0.6</b>	0.7257	0.7291	0.7324	0.7357	0.7389	0.7422	0.7454	0.7486	0.7517	0.7549
<b>0.7</b>	0.7580	0.7611	0.7642	0.7673	0.7703	0.7734	0.7764	0.7793	0.7823	0.7852
<b>0.8</b>	0.7881	0.7910	0.7939	0.7967	0.7995	0.8023	0.8051	0.8078	0.8106	0.8133
<b>0.9</b>	0.8159	0.8186	0.8212	0.8238	0.8264	0.8289	0.8315	0.8340	0.8365	0.8389
<b>1.0</b>	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554	0.8577	0.8599	0.8621
<b>1.1</b>	0.8643	0.8665	0.8686	0.8708	0.8729	0.8749	0.8770	0.8790	0.8810	0.8830
<b>1.2</b>	0.8849	0.8869	0.8888	0.8906	0.8925	0.8943	0.8962	0.8980	0.8997	0.9015
<b>1.3</b>	0.9032	0.9049	0.9066	0.9082	0.9099	0.9115	0.9131	0.9147	0.9162	0.9177
<b>1.4</b>	0.9192	0.9207	0.9222	0.9236	0.9251	0.9265	0.9279	0.9292	0.9306	0.9319
<b>1.5</b>	0.9332	0.9345	0.9357	0.9370	0.9382	0.9394	0.9406	0.9418	0.9429	0.9441
<b>1.6</b>	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515	0.9525	0.9535	0.9545
<b>1.7</b>	0.9554	0.9564	0.9573	0.9582	0.9591	0.9599	0.9608	0.9616	0.9625	0.9633
<b>1.8</b>	0.9641	0.9649	0.9656	0.9664	0.9671	0.9678	0.9686	0.9693	0.9699	0.9706
<b>1.9</b>	0.9713	0.9719	0.9726	0.9732	0.9738	0.9744	0.9750	0.9756	0.9761	0.9767
<b>2.0</b>	0.9772	0.9778	0.9783	0.9788	0.9793	0.9798	0.9803	0.9808	0.9812	0.9817
<b>2.1</b>	0.9821	0.9826	0.9830	0.9834	0.9838	0.9842	0.9846	0.9850	0.9854	0.9857
<b>2.2</b>	0.9861	0.9864	0.9868	0.9871	0.9875	0.9878	0.9881	0.9884	0.9887	0.9890
<b>2.3</b>	0.9893	0.9896	0.9898	0.9901	0.9904	0.9906	0.9909	0.9911	0.9913	0.9916
<b>2.4</b>	0.9918	0.9920	0.9922	0.9925	0.9927	0.9929	0.9931	0.9932	0.9934	0.9936
<b>2.5</b>	0.9938	0.9940	0.9941	0.9943	0.9945	0.9946	0.9948	0.9949	0.9951	0.9952
<b>2.6</b>	0.9953	0.9955	0.9956	0.9957	0.9959	0.9960	0.9961	0.9962	0.9963	0.9964
<b>2.7</b>	0.9965	0.9966	0.9967	0.9968	0.9969	0.9970	0.9971	0.9972	0.9973	0.9974
<b>2.8</b>	0.9974	0.9975	0.9976	0.9977	0.9977	0.9978	0.9979	0.9979	0.9980	0.9981
<b>2.9</b>	0.9981	0.9982	0.9982	0.9983	0.9984	0.9984	0.9985	0.9985	0.9986	0.9986
<b>3.0</b>	0.9987	0.9987	0.9987	0.9988	0.9988	0.9989	0.9989	0.9989	0.9990	0.9990
<b>3.1</b>	0.9990	0.9991	0.9991	0.9991	0.9992	0.9992	0.9992	0.9992	0.9993	0.9993
<b>3.2</b>	0.9993	0.9993	0.9994	0.9994	0.9994	0.9994	0.9994	0.9995	0.9995	0.9995
<b>3.3</b>	0.9995	0.9995	0.9995	0.9996	0.9996	0.9996	0.9996	0.9996	0.9996	0.9997
<b>3.4</b>	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9998
<b>3.5</b>	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998

Filler