CARNEGIE MELLON UNIVERSITY 07-280

HOMEWORK 7

DUE: Thursday, Oct. 29, 2026

https://www.cs.cmu.edu/~07280

INSTRUCTIONS

- **Format:** Use the provided LaTeX template to write your answers in the appropriate locations within the *.tex files and then compile a pdf for submission. We try to mark these areas with STUDENT SOLUTION HERE comments. Make sure that you don't change the size or location of any of the answer boxes and that your answers are within the dedicated regions for each question/part. If you do not follow this format, we may deduct points.

  You may also digitally annotate the pdf. Illegible handwriting will lead to lost points. However, we suggest that try to do at least some of your work directly in LaTeX.

- **How to submit written component:** Submit to Gradescope a pdf with your answers. Again, make sure your answer boxes are aligned with the original pdf template.

- **How to submit programming component:** See section Programming Submission for details on how to submit to the Gradescope autograder.

- **Policy:** See the course website for homework policies, including late policy, and academic integrity policies.

- **Proofs and Derivations:** For full credit, you must clearly show all nontrivial steps. Explicit reasons for each step are not required but can be helpful, especially if the step is not obvious.

Name	
Andrew ID	
Hours to complete all components (nearest hour)	

1 [18 pts] MLE

Consider the following distribution with parameters k and α:

$$p(x \mid k, \alpha) \;=\; \mathbb{I}(x \geq k)\frac{\alpha k^{\alpha}}{x^{\alpha+1}} \;=\; \begin{cases} \frac{\alpha k^{\alpha}}{x^{\alpha+1}} & x \geq k \\ 0 & \text{otherwise} \end{cases}$$

We also have that $k \in (0, \infty)$ and $\alpha \in (0, \infty)$.

This distribution is often used for modeling the distribution of wealth in society, fitting the trend that a large portion of wealth is held by a small fraction of the population. This is due to its nature as a skewed, heavy-tailed distribution (notice that the probability decays polynomially in value of the random variable x, unlike Gaussian, exponential, Laplace or Poisson distributions where the probability decays exponentially in the value of the random variable).

Suppose you have a dataset $\mathcal{D}$ which contains N i.i.d samples $x^{(1)}$,$x^{(2)}$,...,$x^{(N)}$ drawn from the above distribution.

Note: For convenience in this optimization problem, let us define $\log 0$ to be $-\infty$; this won't affect the answers in the end; it essentially just acts as a notational convenience.

1. **[4 pts]** Write the likelihood $\mathcal{L}(k, \alpha; \mathcal{D})$ and log-likelihood $\ell(k, \alpha; \mathcal{D})$. You may keep the indicator function in your answers (rather than trying to write cases or otherwise avoid the indicator function).

Your Answer

2. **[2 pts]** What is the numerical value of the likelihood (not the log-likelihood) if we have $\alpha = 2$, $k = 5$, and $N = 4$ data points $\mathcal{D} = \{2, 3, 7, 4\}$?

3. **[8 pts]** Give the MLE for the parameter α, assuming the parameter k is fixed and $k \leq x^{(i)}\ \forall i$. *Hint*: There is a closed-form solution.

$\hat{\alpha}_{MLE}$

4. **[4 pts]** Next, give the MLE for the parameter k. Here α may be any fixed value or set to its MLE.

   Hint: You may be tempted to conclude that MLE of k is infinity, but when $k = \infty$, what happens to $p(x \mid k, \alpha)$?

$\hat{k}_{MLE}$

2 [20 pts] MLE for Categorical Variables

The categorical distribution is a discrete distribution with K possible values (often corresponding to K discrete events). We'll represent this distribution as a one-hot vector of K random binary random variables, $Y = [Y_1, Y_2, \ldots, Y_K]^\top$, where Y_k takes on the value 1 if the outcome is in class k and 0 otherwise. It is described by a parameter vector $\boldsymbol{\phi} = (\phi_1, \phi_2, \ldots, \phi_k)$ where each ϕ_k determines the probability that the random variable is in category k. Also note that $\phi_k \geq 0$ and $\sum_{k=1}^{K} \phi_k = 1$. Note that the Bernoulli distribution is a special case of the categorical distribution where $K = 2$.

For a dataset $\mathcal{D}$ of N i.i.d. samples drawn from the categorical distribution, the likelihood function can be written as:

$$\mathcal{L}(\boldsymbol{\phi}; \mathcal{D}) = \prod_{i=1}^{N} p\left(\mathbf{y}^{(i)} \mid \boldsymbol{\phi}\right) = \prod_{i=1}^{N} \prod_{k=1}^{K} \phi_k^{\mathbb{I}\left(y_k^{(i)}=1\right)}$$

To help us to remove the indicator function notation, we can define N_k as the number of points in $\mathcal{D}$ that belong to the k-th category.

1. **[4 pts]** Write the likelihood and log-likelihood functions $p(\mathcal{D} \mid \boldsymbol{\phi})$ and $\log p(\mathcal{D} \mid \boldsymbol{\phi})$ in terms of N_k.

$\mathcal{L}(\boldsymbol{\phi}; \mathcal{D})$

$\ell(\boldsymbol{\phi}; \mathcal{D})$

If we just took the derivative of the log-likelihood and set it equal to zero, we would get an estimate saying that each ϕ_k should be equal to infinity, regardless of the value of N_k. So, something is a little broken here...

What we're missing is to account for the constraint that $\sum_{k=1}^{K} \phi_k = 1$. Here is the constrained optimization problem:

$$\begin{aligned} &\underset{\boldsymbol{\phi}}{\text{argmin}} - \ell(\boldsymbol{\phi}; \mathcal{D}) \\ &\text{s.t. } \sum_{k=1}^{K} \phi_k = 1 \end{aligned}$$

To solve this constrained optimization problem, we're going to use a really cool trick called Lagrange multipliers. We'll give you just enough information on Lagrange multipliers in this write-up, but if you are curious and want more information, you could start with Lagrange multiplier tutorials in Khan Academy or Paul's Notes.

- **Super brief instructions for using Lagrange multipliers to solve a generic constrained optimization problem**

  Goal: Solve the following optimization:

  $$\begin{aligned} &\underset{\boldsymbol{\phi}}{\text{argmin}}\ f(\boldsymbol{\phi}) \\ &\text{s.t. } g(\boldsymbol{\phi}) = 0 \end{aligned}$$

  Step 1: Construct the "Lagrangian", $L(\boldsymbol{\phi}, \lambda)$, with the added scalar variable λ. Note that we're going to use the letter L for the Lagrangian rather than the typical $\mathcal{L}$ to avoid confusion with the notation for the likelihood.

  $$L(\boldsymbol{\phi}, \lambda) = f(\boldsymbol{\phi}) + \lambda g(\boldsymbol{\phi})$$

  Step 2: Find the "saddle point" of the Lagrangian, specifically the $\boldsymbol{\phi}$ and λ where:

  $$\nabla L(\boldsymbol{\phi}, \lambda) = \mathbf{0}$$

For our categorical MLE problem, we can formulate the Lagrangian (Step 1) with $f(\boldsymbol{\phi})$ as the negative log-likelihood and $g(x) = (\sum_{k=1}^{K} \phi_k) - 1$.

2. **[4 pts]** (Lagrange multipliers Step 2a) Write the partial derivatives of the Lagrangian with respect to parameter ϕ_k and with respect to λ:

$\partial L / \partial \phi_k$

$\partial L / \partial \lambda$

3. **[8 pts]** (Lagrange multipliers Step 2b) Set your expressions for $\partial L/\partial \phi_k$ and $\partial L/\partial \lambda$ both equal to zero and use them as a system of equations to solve for the MLE estimate for ϕ_k.

ϕ_k MLE

4. **[4 pts]** Finally, let's plug in some numbers to see if our results above match our intuition. For a dataset with $N = 100$ points $K = 4$ categories and $N_1 = 10, N_2 = 20, N_3 = 30$, and $N_4 = 40$, what is the MLE of the vector $\boldsymbol{\phi}$. Write your answer as a vector of numerical values.

$\boldsymbol{\phi}_{\text{MLE}}$

3 [48 pts] Programming Experiments

Complete this written section after finishing the programming notebook `hw7_student.ipynb` (Questions 3, 4, 5b, 6b). Use your own generated outputs and printed statistics from those sections.

1. **[18 pts]** Training AlexNet on Small ImageNet (Programming Q3)

   You have just trained your AlexNet from scratch on the small `ImageNet10_10` dataset with 100 epochs, and produced 2 plots, 1 on loss and 1 on accuracy. Put your plots here.

   AlexNet Scratch on ImageNet10 10

   Look at both plots, and consider these questions (you do not have to write down anything):

   - Does the training accuracy increase quickly?
   - Does the validation curve lag behind?
   - Do the plots show the kind of overfitting you would expect from such a small dataset?

   Then, you also trained your AlexNet on the larger `ImageNet10_100` (also 100 epochs) and produced loss and accuracy plots. Put your plots here. They should look different than the plots you produced on `ImageNet10_10`.

AlexNet Scratch on ImageNet10 100

Do these training curves look better or worse than the ones for `ImageNet10_10`? Explain why you think that is (what changed about our training process, and how does that affect the training?).

◯ Better ◯ Worse

Explain

If you wanted to further improve the performance of this model, name one thing you could try doing, and why you think that might help. Credit will be given for any answer that is well-supported by your explanation.

Improving Performance

2. **[6 pts]** Comparing AlexNet and MobileNet (Programming Q4)

   How many parameters are in AlexNet Scratch? Torchvision AlexNet? Torchvision MobileNet? How many Conv2d and Linear layers are in each of these? Fill in the table below.

Parameters and Layers

Model	**# Parameters**	**# Conv2d Layers**	**# Linear Layers**
AlexNet (Scratch)			
TorchVision AlexNet			
TorchVision MobileNet			

   Why does our AlexNet from scratch have the same number of layers but less parameters than TorchVision AlexNet?

Answer

3. **[18 pts]** Fine-tuning Experiments on a Larger 10-Class ImageNet Subset (Programming Q5b)

   You ran 4 tests: AlexNet Unfrozen, AlexNet Frozen, MobileNet Unfrozen, and MobileNet Frozen. Put the plot you produced for each test in the corresponding boxes below:

AlexNet Unfrozen

AlexNet Frozen

MobileNet Unfrozen

MobileNet Frozen

If you were to pick one of the four fine-tuned models above, which one would you pick, and why? Consider the following questions to help guide your decision: What metrics are you using to decide the best model? Are there any red flags that are more prominent in some models more than others?

◯ AlexNet Unfrozen ◯ AlexNet Frozen ◯ MobileNet Unfrozen ◯ MobileNet Frozen

Explain

4. **[6 pts]** Transfer Learning Experiments (Programming Q6b)

   Put your plots from Q6b here:

AlexNet pretrained + frozen on ASL

MobileNet pretrained + frozen on ASL

4 Collaboration Policy

After you have completed all other components of this assignment, report your answers to the following collaboration questions.

1. Did you receive any help whatsoever from anyone in solving this assignment? If so, include full details including names of people who helped you and the exact nature of help you received.

Your Answer

2. Did you give any help whatsoever to anyone in solving this assignment? If so, include full details including names of people you helped and the exact nature of help you offered.

Your Answer

3. Did you find or come across code that implements any part of this assignment? If so, include full details including the source of the code and how you used it in the assignment.

Your Answer