

Homework #0

CSEP 546: Machine Learning

Professor Jamie Morgenstern

Due: ~~Wednesday April 8, 2026 11:59pm~~ Updated: Thursday April 9, 2026 6:25pm (just before class)

38 points

Please review all homework guidance posted on the website before submitting to Gradescope. Reminders:

- All code must be written in Python, and all written work must be typeset (e.g. \LaTeX).
- Make sure to read the “What to Submit” section following each question and include all items.
- Please provide succinct answers and supporting reasoning for each question. Similarly, when discussing experimental results, concisely create tables and/or figures when appropriate to organize the experimental results. All explanations, tables, and figures for any particular part of a question must be grouped together.
- For every problem involving generating plots, please include the plots as part of your PDF submission.
- When submitting to Gradescope, please link each question from the homework in Gradescope to the location of its answer in your homework PDF. Failure to do so may result in deductions of up to 10% of the value of each question not properly linked. For instructions, see https://www.gradescope.com/get_started#student-submission.
- **If you used AI tools on any part of this assignment, you must turn in a complete transcript of all conversations with AI systems as part of your submission. Failure to do so is a violation of the collaboration policy.**
- **Make sure to submit your code by running `inv submit` from within the `csep546 uv` environment and submitting the zip folder created. Submissions that don't follow this procedure will fail the autograder on Gradescope. After the due date, we will not accept resubmissions. We suggest you wait for the autograder on Gradescope to complete and display your score before regarding the coding portion of assignments complete.**

Not adhering to these reminders may result in point deductions.

Important: By turning in this assignment (and all that follow), you acknowledge that you have read and understood the collaboration policy with humans and AI assistants alike: <https://courses.cs.washington.edu/courses/csep546/26sp/assignments/>. Any questions about the policy should be raised at least 24 hours before the assignment is due. There are no warnings or second chances. If we suspect you have violated the collaboration policy, we will report it to the college of engineering who will complete an investigation.

Quick note about `macro.tex`

The `macro.tex` file provided on the course website provides the definitions for different macros that are referenced in the CSEP 546 homework \LaTeX files. For example, it includes the new command that makes the point values for problems pink and italicized on the homework documents, and also includes commands for things like set notation and writing matrices.

While not required to use the provided homework \LaTeX files, if you would like to compile them, place the homework file in a directory and change the path to the macro file on line 3 (e.g. if the macro file is in the same directory as the homework file, the path should be `\subimport*{.}{macro}`). You can also choose to directly copy the contents of the macro file into your \LaTeX code.

Probability and Statistics

A1. [2 points] (From Murphy Exercise 2.4.) After your yearly checkup, the doctor has bad news and good news. The bad news is that you tested positive for a serious disease, and that the test is 99% accurate (i.e., the probability of testing positive given that you have the disease is 0.99, as is the probability of testing negative given that you don't have the disease). The good news is that this is a rare disease, striking only one in 10,000 people. What are the chances that you actually have the disease?

What to Submit:

- Final Answer
- Corresponding Calculations

A2. Let $\mathbf{X} = [X_1 \ \cdots \ X_n]^\top$ be a random vector with mean vector $\boldsymbol{\mu} \in \mathbb{R}^n$ and covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{n \times n}$. That is, for each $i \in [n]$, $\mathbb{E}[X_i] = \boldsymbol{\mu}_i$, and for each $i, j \in [n]$, $\text{Cov}(X_i, X_j) = \boldsymbol{\Sigma}_{i,j}$.

- [2 points] Let $i \in [n]$. Find an expression for $\mathbb{E}[X_i^2]$ in terms of only $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.
- [2 points] Recall that the trace of a square matrix is the sum of its diagonal entries. Show that $\mathbb{E}[\mathbf{X}^\top \mathbf{X}] = \text{tr}(\boldsymbol{\Sigma}) + \boldsymbol{\mu}^\top \boldsymbol{\mu}$.

What to Submit:

- **Part a:** Expression for $\mathbb{E}[X_i^2]$ and corresponding calculations
- **Part b:** Proof

A3. Let $X_1, X_2, \dots, X_n \sim \mathcal{N}(\mu, \sigma^2)$ be i.i.d. random variables. Compute the following:

- [1 point] $a \in \mathbb{R}, b \in \mathbb{R}$ such that $aX_1 + b \sim \mathcal{N}(0, 1)$.
- [2 points] Using a and b from the previous part, $\mathbb{E}[|aX_1 + b|]$.
Hint: You may need to use u-substitution.
- [1 point] $\mathbb{E}[X_1 + 2X_2], \text{Var}[X_1 + 2X_2]$.
- [2 points] Setting $\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n X_i$, the mean and variance of $\sqrt{n}(\hat{\mu}_n - \mu)$.

What to Submit:

- **Part a:** a, b
- **Part b:** $\mathbb{E}[|aX_1 + b|]$
- **Part c:** $\mathbb{E}[X_1 + 2X_2], \text{Var}[X_1 + 2X_2]$
- **Part d:** $\mathbb{E}[\sqrt{n}(\hat{\mu}_n - \mu)], \text{Var}[\sqrt{n}(\hat{\mu}_n - \mu)]$
- **Parts a-d** Corresponding calculations

Linear Algebra and Vector Calculus

A4. Let $\mathbf{A} = \begin{bmatrix} 0 & 1 & -1 & -2 \\ 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & 2 \end{bmatrix}$ and $\mathbf{B} = \begin{bmatrix} -1 & 2 & 0 \\ -2 & 1 & 2 \\ 2 & 0 & -3 \\ -1 & -1 & 2 \end{bmatrix}$.

- [1 point] Compute \mathbf{AB} .
- [2 points] What is the rank of \mathbf{AB} ? Use this to determine the rank of both \mathbf{A} and \mathbf{B} .
- [1 point] Find a matrix \mathbf{C} such that $\mathbf{ABC} = \mathbf{I}$, or explain why such a matrix does not exist.
- [1 point] Let $b = [1 \ 1 \ 1]^\top$. What is the solution to the linear system $\mathbf{AB}x = b$?

What to Submit:

- **Parts a-d:** Solution and corresponding calculations

A5. For possibly non-symmetric $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ and $c \in \mathbb{R}$, let $f(x, y) = x^\top \mathbf{A}x + y^\top \mathbf{B}y + c$. Define

$$\nabla_z f(x, y) = \left[\frac{\partial f}{\partial z_1}(x, y) \quad \frac{\partial f}{\partial z_2}(x, y) \quad \dots \quad \frac{\partial f}{\partial z_n}(x, y) \right]^\top \in \mathbb{R}^n.$$

Note: If you are unfamiliar with gradients, you may find the resources available on the course website useful. Section 4 of Zico Kolter and Chuong Do's Linear Algebra Review and Reference may be particularly helpful.

- [2 points] Explicitly write out the function $f(x, y)$ in terms of the components $A_{i,j}$ and $B_{i,j}$ using appropriate summations over the indices.
- [2 points] What is $\nabla_x f(x, y)$ and $\nabla_y f(x, y)$ in terms of summations over indices?
- [2 points] What is $\nabla_x f(x, y)$ and $\nabla_y f(x, y)$ in vector notation?

What to Submit:

- **Part a:** Explicit formula for $f(x, y)$
- **Parts b:** Summation form and corresponding calculations. Summation form includes writing out what each component of the resultant vector is, where each component is expressed as a summation. Intermediate components may be indicated by ellipses, like in the equation given in the problem description.
- **Part c:** Vector form and corresponding calculations. Vector form includes writing the final answer only in terms of products, sums (or differences), and/or transposes of the input matrices and vectors.

A6. Show the following:

- [2 points] Let $g: \mathbb{R} \setminus \{0\} \rightarrow \mathbb{R}$ and $v, w \in \mathbb{R}^n$ such that $v_i \neq 0$ and $g(v_i) = w_i$ for $i \in [n]$. Find an expression for g such that $\text{diag}(v)^{-1} = \text{diag}(w)$.
- [2 points] Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be orthonormal and $x \in \mathbb{R}^n$. An orthonormal matrix is a square matrix whose columns and rows are orthonormal vectors, such that $\mathbf{AA}^\top = \mathbf{A}^\top \mathbf{A} = \mathbf{I}$ where \mathbf{I} is the identity matrix. Show that $\|\mathbf{A}x\|_2^2 = \|x\|_2^2$.
- [2 points] Let $\mathbf{B} \in \mathbb{R}^{n \times n}$ be invertible and symmetric. A symmetric matrix is a square matrix satisfying $\mathbf{B} = \mathbf{B}^\top$. Show that \mathbf{B}^{-1} is also symmetric.
- [2 points] Let $\mathbf{C} \in \mathbb{R}^{n \times n}$ be positive semi-definite (PSD). A positive semi-definite matrix is a symmetric matrix satisfying $x^\top \mathbf{C}x \geq 0$ for any vector $x \in \mathbb{R}^n$. Show that its eigenvalues are non-negative.

What to Submit:

- **Part a:** Explicit formula for g
- **Parts a-d:** Proof

Programming

These problems are available in a .zip file, with some starter code. All coding questions in this class will have starter code. Before attempting these problems, you will need to set up a uv environment that you will use for every assignment in the course. Unzip the HW0-A.zip file and read the instructions in the README file to get started.

A7. For $\nabla_x f(x, y)$ as solved for in Problem ??:

- [1 point]** Using native Python, implement `vanilla_solution` using your `vanilla_matmul` and `vanilla_transpose` functions.
- [1 point]** Now implement `numpy_version` using NumPy functions.
- [1 point]** Report the difference in wall-clock time for parts a-b, and discuss reasons for the observed difference.

What to Submit:

- **Part c:** Plot that shows the difference in wall-clock time for parts a-b
- **Part c:** Explanation for the difference (1-2 sentences)
- **Code** on Gradescope through coding submission

A8. Two random variables X and Y have equal distributions if their CDFs, F_X and F_Y , respectively, are equal, i.e. for all x , $|F_X(x) - F_Y(x)| = 0$. The central limit theorem says that the sum of k independent, zero-mean, variance $1/k$ random variables converges to a (standard) Normal distribution as k tends to infinity. We will study this phenomenon empirically (you will use the Python packages NumPy and Matplotlib). Each of the following subproblems includes a description of how the plots were generated; these have been coded for you. The code is available in the .zip file. In this problem, you will add to our implementation to explore `matplotlib` library, and how the solution depends on n and k .


- [2 points]** For $i = 1, \dots, n$ let $Z_i \sim \mathcal{N}(0, 1)$. Let $x \mapsto F(x)$ denote the true CDF from which each Z_i is drawn (i.e., Gaussian). Define $\hat{F}_n(x) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_i \leq x\}$ for $x \in \mathbb{R}$ and we will choose n large enough such that, for all $x \in \mathbb{R}$,

$$\sqrt{\mathbb{E}\left[\left(\hat{F}_n(x) - F(x)\right)^2\right]} \leq 0.0025 .$$

Plot $x \mapsto \hat{F}_n(x)$ for x ranging from -3 to 3 .

- [2 points]** Define $Y^{(k)} = \frac{1}{\sqrt{k}} \sum_{i=1}^k B_i$ where each B_i is equal to -1 and 1 with equal probability and the B_i 's are independent. We know that each $\frac{1}{\sqrt{k}} B_i$ is zero-mean and has variance $1/k$. For each $k \in \{1, 8, 64, 512\}$ we will generate n (same as in part a) independent copies $Y^{(k)}$ and plot their empirical CDF on the same plot as part ??.

Be sure to always label your axes. Your plot should look something like the following (up to styling) (Tip: checkout `seaborn` for instantly better looking plots.)



What to Submit:

- **Part ??:** Value for n (You can simply print the value determined by the code provided for you). **Part ??:** In 1-2 sentences: How does the empirical CDF change with k ?
- **Parts ?? and ??:** Plot of $x \mapsto \hat{F}_n(x)$ for $x \in [-3, 3]$
- **Code** on Gradescope through coding submission