

CS 446 / ECE 449 — Homework 1

your NetID here

Version 1.3

Instructions.

- Homework is due **Tuesday, February 16, at noon CST**; no late homework accepted.
- Everyone must submit individually at gradescope under **hw1** and **hw1code**.
- The “written” submission at **hw1** **must be typed**, and submitted in any format gradescope accepts (to be safe, submit a PDF). You may use L^AT_EX, markdown, google docs, MS word, whatever you like; but it must be typed!
- When submitting at **hw1**, gradescope will ask you to mark out boxes around each of your answers; please do this precisely!
- Please make sure your NetID is clear and large on the first page of the homework.
- Your solution **must** be written in your own words. Please see the course webpage for full academic integrity information. Briefly, you may have high-level discussions with at most 3 classmates, whose NetIDs you should place on the first page of your solutions, and you should cite any external reference you use; despite all this, your solution must be written in your own words.
- We reserve the right to reduce the auto-graded score for **hw1code** if we detect funny business (e.g., your solution lacks any algorithm and hard-codes answers you obtained from someone else, or simply via trial-and-error with the autograder).
- When submitting to **hw1code**, only upload **hw1.py** and **hw1_utils.py**. Additional files will be ignored.

Version History.

1. Initial version.
- 1.1 Fixed typo in hint for 1c.
- 1.2 Clarified wording of 1a, 1b.
- 1.3 Fixed bolding of vectors in 2c.

1. Linear Regression/SVD.

Throughout this problem let \mathbf{X} be the $n \times d$ matrix with the feature vectors $(\mathbf{x}_i)_{i=1}^n$ as its rows. Suppose we have the singular value decomposition $\mathbf{X} = \sum_{i=1}^r s_i \mathbf{u}_i \mathbf{v}_i^\top$.

- (a) Let the training examples $(\mathbf{x}_i)_{i=1}^n$ be the standard basis vectors \mathbf{e}_i of \mathbb{R}^d with each \mathbf{e}_i repeated $n_i > 0$ times having labels $(y_{i_j})_{j=1}^{n_i}$. That is, our training set is:

$$\bigcup_{i=1}^d \left\{ (\mathbf{e}_i, y_{i_j}) \right\}_{j=1}^{n_i},$$

where $\sum_{i=1}^d n_i = n$. Show that for a vector \mathbf{w} that minimizes the empirical risk, the components w_i of \mathbf{w} are the averages of the labels $(y_{i_j})_{j=1}^{n_i}$: $w_i = \frac{1}{n_i} \sum_{j=1}^{n_i} y_{i_j}$.

Hint: Write out the expression for the empirical risk with the squared loss and set the gradient equal to zero.

Remark: This gives some intuition as to why “regression” originally meant “regression towards the mean.”

- (b) Returning to a general matrix \mathbf{X} , show that if the label vector \mathbf{y} is a linear combination of the $\{\mathbf{u}_i\}_{i=1}^r$ then there exists a \mathbf{w} for which the empirical risk is zero (meaning $\mathbf{X}\mathbf{w} = \mathbf{y}$).

Hint: Either consider the range of \mathbf{X} and use the SVD, or compute the empirical risk explicitly with $\mathbf{y} = \sum_{i=1}^r a_i \mathbf{u}_i$ for some constants a_i and $\hat{\mathbf{w}}_{\text{ols}} = \mathbf{X}^+ \mathbf{y}$.

Remark: It’s also not hard to show that if \mathbf{y} is not a linear combination of the $\{\mathbf{u}_i\}_{i=1}^r$, then the empirical risk must be nonzero.

- (c) Show that $\mathbf{X}^\top \mathbf{X}$ is invertible if and only if $(\mathbf{x}_i)_{i=1}^n$ spans \mathbb{R}^d .

Hint: Recall that the squares of the singular values of \mathbf{X} are eigenvalues of $\mathbf{X}^\top \mathbf{X}$.

Remark: This characterizes when linear regression has a unique solution due to the normal equation (note that we always have at least one solution obtained by the pseudoinverse). We would not have had a unique solution for part (a) if we had an $n_i = 0$.

- (d) Provide a matrix \mathbf{X} such that $\mathbf{X}^\top \mathbf{X}$ is invertible and $\mathbf{X}\mathbf{X}^\top$ is not. Include a formal verification of this for full points.

Hint: Use part (c). It may be helpful to think about conditions under which a matrix is not invertible.

Solution.

2. Loss functions/MLE.

Recall that maximum likelihood estimation is a technique used to maximize the probability of observed data given parameters w of a probability model: $P(\text{data}|w)$. In machine learning, we want to find the most likely set of parameters w given the observed data. This can be expressed in terms of the likelihood function $\mathcal{L}(w) = P(\text{data}|w)$ as follows:

$$\arg \max_w P(w|\text{data}) = \arg \max_w \frac{P(\text{data}|w)P(w)}{P(\text{data})} = \arg \max_w P(\text{data}|w)P(w). \quad (1)$$

Assuming all parameters w are equally likely, (1) reduces to $\arg \max_w P(\text{data}|w)$ (we operate under this assumption for parts (a) and (b) of this problem). We represent the probability our model with parameters w assigns to data $(x_i, y_i)_{i=1}^n$ by the likelihood function $\mathcal{L}(w) = \prod_{i=1}^n p_w(y_i|x_i)$ where $p_w(y_i|x_i)$ is our model's estimated probability of label y_i given x_i .

- (a) Suppose we are implementing a multiclass classifier with k target classes. Call the target class of the i 'th training example $t_i \in \{1, \dots, k\}$. We can encode the label y_i of the i 'th training example as a one-hot vector with a 1 at index t_i and 0s elsewhere. Consider the likelihood function

$$\mathcal{L}(w) = \prod_{i=1}^n \prod_{j=1}^k p_w(t_j|x_i)^{y_{ij}},$$

where y_{ij} is the j 'th element of the vector y_i . Show that maximizing this likelihood function is equivalent to minimizing the cross entropy loss (from Lecture 4) over the dataset

$$-\sum_{i=1}^n \sum_{j=1}^k y_{ij} \log(p_w(t_j|x_i)).$$

Remark: Why the double product? The inner one, $\prod_{j=1}^k p_w(t_j|x_i)^{y_{ij}}$, is just a clever way of writing the probability of the i 'th training example: all of the exponents y_{ij} will be zero except for y_{t_i} , which is 1. Therefore, the product reduces to the probability of the true label given the input.

- (b) Recall that for logistic regression we minimize the logistic loss $\ell_{\text{logistic}}(y w^\top x)$ when we model the probability of a positive label by $p_w(1|x) = \frac{1}{1+e^{-w^\top x}}$ and the labels $y \in \{-1, 1\}$. Another common loss function for logistic regression is binary cross entropy, defined by

$$-y \log(p_w(1|x)) - (1-y) \log(1-p_w(1|x)),$$

where the labels $y \in \{0, 1\}$ (the cross entropy derived in part (a) is a generalization of this). Show that $\ell_{\text{logistic}}(y w^\top x)$ with $y = -1$ is equal to binary cross entropy with $y = 0$ and that $\ell_{\text{logistic}}(y w^\top x)$ with $y = 1$ is equal to binary cross entropy with $y = 1$.

Remark: This shows that the loss functions are equivalent up to a relabeling.

- (c) Now we will drop the assumption that all parameters w are equally likely. Specifically, assume that the w_i are i.i.d. and follow a normal (Gaussian) distribution with mean 0 and variance $1/\lambda$: $w_i \sim \mathcal{N}(0, 1/\lambda)$. Independence of the w_i implies $P(w) = \prod_{i=1}^d P(w_i)$. Further, assume that $y_i|x_i, w$ follows a normal distribution with mean $w^\top x_i$ and variance 1: $y_i|x_i, w \sim \mathcal{N}(w^\top x_i, 1)$. Show that

$$\arg \max_w \prod_{i=1}^n p_w(y_i|x_i)P(w) = \arg \min_w \left\{ \frac{1}{2} \sum_{i=1}^n (w^\top x_i - y_i)^2 + \frac{\lambda}{2} \sum_{j=1}^d w_j^2 \right\}.$$

Remark: This is an interpretation of the regularization term for ridge regression. The probability of very positive or very negative w_i is low if w_i follows a normal distribution, which is reflected by a larger penalty $\sum_i w_i^2$. (Note that we also dropped the $1/n$ factor of the empirical risk by maximum likelihood convention).

Solution.

3. Linear Regression.

Recall that the empirical risk in the linear regression method is defined as $\hat{\mathcal{R}}(\mathbf{w}) := \frac{1}{2n} \sum_{i=1}^n (\mathbf{w}^\top \mathbf{x}_i - y_i)^2$, where $\mathbf{x}_i \in \mathbb{R}^d$ is a data point and y_i is an associated label.

- (a) Implement linear regression using gradient descent in the `linear_gd(X, Y, lrate, num_iter)` function of `hw1.py`. You are given as input a training set \mathbf{X} as an $n \times d$ tensor, training labels \mathbf{Y} as an $n \times 1$ tensor, a learning rate `lrate`, and the number of iterations of gradient descent to run `num_iter`. Using gradient descent, find parameters \mathbf{w} that minimize the empirical risk $\hat{\mathcal{R}}(\mathbf{w})$. Use $\mathbf{w} = 0$ as your initial parameters, and return your final w as output. Prepend a column of ones to \mathbf{X} in order to accommodate a bias term in \mathbf{w} .

Library routines: `torch.matmul (@)`, `torch.tensor.shape`, `torch.tensor.t`, `torch.cat`, `torch.ones`, `torch.zeros`, `torch.reshape`.

- (b) Implement linear regression by using the pseudoinverse to solve for w in the `linear_normal(X,Y)` function of `hw1.py`. You are given a training set \mathbf{X} as an $n \times d$ tensor and training labels \mathbf{Y} as an $n \times 1$ tensor. Return your parameters w as output. As before, make sure to accommodate a bias term by prepending ones to the training examples \mathbf{X} .

Library routines: `torch.matmul (@)`, `torch.cat`, `torch.ones`, `torch.pinv`.

- (c) Implement the `plot_linear()` function in `hw1.py`. Use the provided function `hw1_utils.load_reg_data()` to generate a training set \mathbf{X} and training labels \mathbf{Y} . Plot the curve generated by `linear_normal()` along with the points from the data set. Return the plot as output. Include the plot in your written submission.

Library routines: `torch.matmul (@)`, `torch.cat`, `torch.ones`, `plt.plot`, `plt.scatter`, `plt.show`, `plt.gcf` where `plt` refers to the `matplotlib.pyplot` library.

Solution.

4. Polynomial Regression.

In Problem 3 you constructed a linear model $\mathbf{w}^\top \mathbf{x} = \sum_{i=1}^d x_i w_i$. In this problem you will use the same setup as in the previous problem, but enhance your linear model by doing a quadratic expansion of the features. Namely, you will construct a new linear model $f_{\mathbf{w}}$ with parameters

$$(w_0, w_{01}, \dots, w_{0d}, w_{11}, w_{12}, \dots, w_{1d}, w_{22}, w_{23}, \dots, w_{2d}, \dots, w_{dd}),$$

defined by

$$f_{\mathbf{w}}(x) = \mathbf{w}^\top \phi(\mathbf{x}) = w_0 + \sum_{i=1}^d w_{0i} x_i + \sum_{i \leq j}^d w_{ij} x_i x_j.$$

- (a) Given a 3-dimensional feature vector $\mathbf{x} = (x_1, x_2, x_3)$ completely write out the quadratic expanded feature vector $\phi(\mathbf{x})$.
- (b) Implement the `poly_gd()` function in `hw1.py`. The input is in the same format as it was in Problem 3. Implement gradient descent on this training set with \mathbf{w} initialized to 0. Return \mathbf{w} as the output with terms in this exact order: bias, linear, then quadratic. For example, if $d = 3$ then you would return $(w_0, w_{01}, w_{02}, w_{03}, w_{11}, w_{12}, w_{13}, w_{22}, w_{23}, w_{33})$.

Library routines: `torch.cat`, `torch.ones`, `torch.zeros`, `torch.stack`.

Hint: You will want to prepend a column of ones to \mathbf{X} , and append to \mathbf{X} the squared features in the specified order. You can generate the squared features in the correct order using a nested for loop.

- (c) Implement the `poly_normal` function in `hw1.py`. You are given the same data set as from part (b), but this time determine \mathbf{w} by using the pseudoinverse. Return \mathbf{w} in the same order as in part (b).

Library routines: `torch.pinv`.

Hint: You will still need to transform the matrix \mathbf{X} in the same way as in part (b).

- (d) Implement the `plot_poly()` function in `hw1.py`. Use the provided function `hw1_utils.load_reg_data()` to generate a training set \mathbf{X} and training labels \mathbf{Y} . Plot the curve generated by `poly_normal()` along with the points from the data set. Return the plot as output and include it in your written submission. Compare and contrast this plot with the plot from Problem 3. Which model appears to approximate the data best? Justify your answer.

Library routines: `plt.plot`, `plt.scatter`, `plt.show`, `plt.gcf`.

- (e) The Minsky-Papert XOR problem is a classification problem with data set:

$$X = \{(-1, +1), (+1, -1), (-1, -1), (+1, +1)\}$$

where the label for a given point (x_1, x_2) is given by its product $x_1 x_2$. For example, the point $(-1, +1)$ would be given label $y = (-1)(1) = -1$. Implement the `poly_xor()` function in `hw1.py`. In this function you will load the XOR data set by calling the `hw1_utils.load_xor_data()` function, and then apply the `linear_normal()` and `poly_normal()` functions to generate labels for the XOR points. Include a plot of contour lines that show how each model classifies points in your written submission. You may use `contour_plot()` in `hw1_utils.py`. Return the labels for both the linear model and the polynomial model in that order. Do both models correctly classify all points? (Note that red corresponds to larger values and blue to smaller values when using `contour_plot`).

Solution.

5. Logistic Regression.

Recall the empirical risk $\hat{\mathcal{R}}$ for logistic regression (as presented in lecture 3):

$$\hat{\mathcal{R}}_{\log}(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \ln(1 + \exp(-y_i \mathbf{w}^\top \mathbf{x}_i)).$$

Here you will minimize this risk using gradient descent.

- (a) In your written submission, derive the gradient descent update rule for this empirical risk by taking the gradient. Write your answer in terms of the learning rate η , previous parameters \mathbf{w} , new parameters \mathbf{w}' , number of examples n , and training examples \mathbf{x}_i . Show all of your steps.
- (b) Implement the `logistic()` function in `hw1.py`. You are given as input a training set `X`, training labels `Y`, a learning rate `lrate`, and number of gradient updates `num_iter`. Implement gradient descent to find parameters \mathbf{w} that minimize the empirical risk $\hat{\mathcal{R}}_{\log}(\mathbf{w})$. Perform gradient descent for `num_iter` updates with a learning rate of `lrate`, initializing $\mathbf{w} = 0$ and returning \mathbf{w} as output. Don't forget to prepend `X` with a column of ones.

Library routines: `torch.matmul` (`@`), `torch.tensor.t`, `torch.exp`.

- (c) Implement the `logistic_vs_ols()` function in `hw1.py`. Use `hw1_utils.load_logistic_data()` to generate a training set `X` and training labels `Y`. Run `logistic(X,Y)` from part (b) taking `X` and `Y` as input to obtain parameters \mathbf{w} (use the defaults for `num_iter` and `lrate`). Also run `linear_gd(X,Y)` from Problem 3 to obtain parameters \mathbf{w} . Plot the decision boundaries for your logistic regression and least squares models along with the data `X`. Which model appears to classify the data better? Explain why you believe your choice is the better classifier for this problem.

Library routines: `torch.linspace`, `plt.scatter`, `plt.plot`, `plt.show`, `plt.gcf`.

Solution.