

HOMEWORK 2 - V1

STA 414/2104 WINTER 2021

University of Toronto

VERSION HISTORY: V0 → V1: ADD ϵI TO COVARIANCE IN Q2.B

- **Deadline:** Mon, Feb. 22, at 13:59.
- **Submission:** You need to submit your solutions through Crowdmark, including all your derivations, plots, and your code. You can produce the file however you like (e.g. L^AT_EX, Microsoft Word, etc), as long as it is readable. Points will be deducted if we have a hard time reading your solutions or understanding the structure of your code. For this assignment, you should submit the filled out starter code right after your answers.

MNIST dataset. In this assignment, you will fit both generative and discriminative models using the MNIST dataset of handwritten numbers.

Each datapoint in the MNIST <http://yann.lecun.com/exdb/mnist/> dataset is a 28x28 black-and-white image of a handwritten digit in $\{0 \dots 9\}$, and a label indicating which digit.

MNIST is the 'fruit fly' of machine learning - a simple standard problem useful for comparing the properties of different algorithms. A starter Python code that loads and plots the MNIST dataset is attached. For this assignment, we will *binarize* the data, converting grey-scale pixels to either black or white (0 or 1) with > 0.5 being the cutoff (already done in the starter code).

The starter code `hw2-train.py` is to be used in both questions below. You will need to write the missing parts of the functions and return it back for evaluation. Note that each missing part should be typically a few lines of code, so make sure your code is compact. When comparing models, you will need a training and test set. Build a dataset of only 2000 training samples (controlled by `N_data`) to use when coding or debugging, to make loading and training faster. Inspect the starter code carefully, before you start coding.



Fig 1: Samples from MNIST data set.

1. Multi-class Logistic Regression Classifier - 50 pts. In this question, you will fit a discriminative model using gradient descent. Our model will be multi-class logistic regression:

$$p(t_k = 1 | \mathbf{x}, \mathbf{w}) = \frac{\exp(\mathbf{w}_k^T \mathbf{x})}{\sum_{i=0}^9 \exp(\mathbf{w}_i^T \mathbf{x})}$$

Omit bias (intercept) parameters for this question.

- (a) (5pts) How many parameters does this model have?
- (b) (10pts) Write down the log-likelihood and convert it into a minimization problem over the cross-entropy loss E . Derive the gradient of E with respect to each \mathbf{w}_k , i.e., $\nabla_{\mathbf{w}_k} E(\mathbf{w})$.
- (c) (30pts) Code up a gradient descent optimizer using the starter code provided to you, and minimize the cross-entropy loss. Report the final training and the test accuracy achieved. The training must be done over the full training dataset, unless there are computational issues, in which case you can reduce the number of training samples depending on the memory available. Report the number of samples used to obtain the final result. Hint: For `log_softmax` function, use `scipy.special.logsumexp` (already imported in the starter code) or its equivalent to make your code more numerically stable. Avoid nested `for` loops, and instead use matrix operations to keep your code fast. Each missing chunk should be a few lines of code!
- (d) (5pts) Plot the final weights obtained as 10 images.

What to submit?

- a) *Number of parameters.*
- b) *Log-likelihood, resulting cross-entropy minimization, and the gradient.*
- c) *Final training and test errors as well as the number of samples used in training.*
- d) *Figure containing each weight \mathbf{w}_k as an image.*
- e) *Your entire code should be attached to the end of your answers.*

2. Gaussian Discriminant Analysis - 50 pts. In this part, we train a generative model using the MNIST dataset. Assuming that the data generating distribution is Gaussian, i.e.

$$(2.1) \quad p(\mathbf{x} | \mathcal{C}_k) = \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}).$$

We know that the posterior $p(\mathcal{C}_k | \mathbf{x})$ can be written in terms of the softmax function

$$(2.2) \quad p(\mathcal{C}_k | \mathbf{x}) = \frac{\exp\{a_k\}}{\sum_j \exp\{a_j\}} \quad \text{where} \quad a_k = \mathbf{w}_k^T \mathbf{x} + w_{k0}.$$

Here, we also know that

$$(2.3) \quad \mathbf{w}_k = \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_k \quad \text{and} \quad w_{k0} = -\frac{1}{2} \boldsymbol{\mu}_k^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_k + \log(p(\mathcal{C}_k)).$$

- (a) (5pts) Write down the log-likelihood implied by this model and find the maximum likelihood estimator (MLE) for the priors $p(\mathcal{C}_k) = \pi_k$ and the class means $\boldsymbol{\mu}_k$, for $k = 1, \dots, K$. Note that you do not need to derive the MLE for the covariance matrix.

- (b) (20pts) Compute the MLEs obtained in the previous part together with the following estimator for the covariance matrix

$$(2.4) \quad \widehat{\Sigma} = \sum_{k=1}^K \frac{N_k}{N} \widehat{\Sigma}_k \quad \text{where} \quad \widehat{\Sigma}_k = \frac{1}{N_k} \sum_{n \in \mathcal{C}_k} (\mathbf{x}_n - \boldsymbol{\mu}_k)(\mathbf{x}_n - \boldsymbol{\mu}_k)^T,$$

where N_k is the number of images that belong to class k , and N is the total number of images. In order to make $\widehat{\Sigma}$ invertible, add ϵI for ϵ small e.g. $\epsilon = 1/N$. Plot the means of each class as an image. Hint: In this part, if you use the entire training dataset to train your model, your computer's memory will likely run out. Start with a small number `N_data = 2000` and slowly increase it. In your final model, use as many samples as permitted by the computer memory. Report this number below. Try to avoid for loops as much as possible. Many of these operations can be written as matrix-matrix products. Take advantage of 1-of-K encoding.

- (c) (15pts) Using the MLE estimators obtained in previous part as well as the posterior (2.2), make predictions on both training and test sets and report the obtained accuracy in each dataset. Also, report the number of training images used to compute the MLE estimators.
- (d) (5pts) Briefly compare the performance of this model to that of logistic regression.
- (e) (5pts) Using the generative model you trained, generate 10 images from digit 0 and 10 images from digit 3.

What to submit?

- a) *Log-likelihood, MLE for class means and the priors, your derivations.*
- b) *Figure containing each class mean $\boldsymbol{\mu}_k$ as an image.*
- c) *Final training and test errors as well as the number of samples used in training.*
- d) *Brief comparison of final accuracies.*
- e) *20 images you generated.*
- f) *Your entire code should be attached to the end of your answers.*