

CS 472/572, Spring 2024  
Homework 3 (Programming): Linear Models  
Due: May 20th, 2024 at 11:59 (via Canvas)

Started code and autograder at: <https://ix.cs.uoregon.edu/~swalton2/cis472/hw3>

## 1 Overview

In this assignment, you will implement the perceptron algorithm and logistic regression. The file format is the same as in the previous project: binary attributes and a binary class label represented as comma-separated values.

Please use the provided templates `perceptron.py` and `lr.py` for the perceptron and logistic regression algorithms, respectively. Fill in your code for the training and prediction functions, `train_perceptron`, `predict_perceptron`, `train_lr`, and `predict_lr`. You may add any additional functions you find helpful.

Your code must be your own. **Undergraduates may complete the assignment in teams of 2.** Graduates must complete the assignment alone.

To test the code, you can run the `.py` files from the command line with the following arguments:

```
python3 perceptron.py --maxiter <iters> --model <model> <train> <test>
python3 lr.py --eta <eta> --l2 <lambda> --maxiter <iters> --model <model> <train> <test>
```

Where `train` is the name of a file containing training data, `test` is the name of a file containing test data to be labeled and `model` is the filename where you will save the resulting linear model. For logistic regression, `eta` is the learning rate and `lambda` is the scale of an L2 regularizer.

## 2 Model File Format

For saving model files, use the following format:

```
0.10976
foo -1.2857
bar 0.4811
```

where the first line contains the bias, and each line that follows lists one of the attributes and its corresponding weight.

## 3 Perceptron

`train_perceptron` should run the perceptron algorithm until convergence. This means running until all training examples are classified correctly. Your implementation should also halt after 100 passes through the training data, if it has

not yet converged. You do not need to shuffle the example order for this question.

`predict_perceptron` should return the activation  $wx+b$  for a given example  $x$  with a given model  $(w, b)$ .

## 4 Logistic Regression

`train_lr` should run the logistic regression learning algorithm for 100 iterations or until convergence. Each iteration should perform one step of batch gradient descent with the given learning rate – do not use stochastic/mini-batching gradient descent for this assignment. You may assume the algorithm has converged when the gradient is small (e.g., its magnitude is less than 0.0001). On many datasets, logistic regression will use the full 100 iterations without converging.

Use the  $\frac{\lambda}{2} \sum_i w_i^2$  as the regularizer, where  $\lambda$  is the strength of the L2 regularizer. The gradient of the regularizer with respect to each  $w_j$  is  $\lambda w_j$ . When implementing gradient descent, just add the gradient of the regularizer to the gradient of the logistic loss.

`predict_lr` should return the probability that  $y = +1$  for a given example  $x$  according to a logistic regression model with the given parameters  $(w, b)$ .

## 5 What to Upload to Canvas

1. `perceptron.py`
2. `lr.py`