Homework Assignment 3b

Due: Friday, March 15, 2024, 11:59 p.m.
Total marks: 55

Question 1. [55 MARKS]

In this question, you will implement multivariate linear regression, and polynomial regression, and learn their parameters using mini-batch stochastic gradient descent (SGD). You will implement 3 different stepsize rules—ConstantLR, HeuristicLR, and AdaGrad—and compare their performance. You will implement linear regression first, with all the stepsize rules. Then you will use this implementation to do polynomial regression, by creating polynomial features and then calling the linear regression procedure. You will compare linear regression and polynomial regression with a constant learning rate. Then you will compare the three different stepsize rules for polynomial regression.

Initial code has been given to you in a notebook, called `A3.jl`, to run the regression algorithms on a dataset. Detailed information for each question is given in the notebook. Baseline algorithms (random and mean) are sanity checks; we should be able to outperform them.

You will be running the algorithms on the GraduateAdmissions_v1.0 data set on Kaggle, which has $n = 500$ samples and $d = 7$ features. The features are composed of GPA, TOEFL grades and a few other criteria. You are asked to train some models to predict the admission probability based on these features. The features are augmented to have a column of ones (to create the bias term), in `A3.jl` (not in the data file itself).

(a) [10 MARKS] Implement the `get_features` function.

(b) [5 MARKS] Implement the `epoch!` function.

(c) [5 MARKS] Implement the `loss` function.

(d) [5 MARKS] Implement the `gardient` function.

(e) [5 MARKS] Implement the `update!` function when the `ConstantLR` optimizer is used.

(f) [5 MARKS] Implement the `update!` function when the `HeuristicLR` optimizer is used.

(g) [5 MARKS] Implement the `update!` function when the `AdaGrad` optimizer is used.

(h) [5 MARKS] **WRITTEN:** Compare the `LinearModel` vs. `Polynomial2Model`. Report their average error and standard error on the admissions dataset. Provide 1-2 sentences commenting on the results.

(i) [10 MARKS] **WRITTEN:** Run the experiment for the `Polynomial2Model` with the three different stepsize rules on the normalized admissions dataset. You do not need to report the result here, it is simply a baseline result for you to compare to.

Now, run the experiment for the `Polynomial2Model` with the three different stepsize rules on the unnormalized admissions dataset. Report the average error and standard error for the three different stepsize rules. Provide 1-2 sentences commenting on the results.

Then, increase the number of epochs to 50 and again report the average errors and standard errors for the three different stepsize rules. Provide 1-2 sentences commenting on the results.

Homework policies:

Your assignment should be submitted as two pdf documents and a .jl notebook, on eClass. **Do not** submit a zip file with all three. One pdf is for the written work, the other pdf is generated from the .jl notebook. The first pdf should contain your answers for questions starting with "**WRITTEN:**". Your answers must be written legibly and scanned or must be typed (e.g., Latex). This .pdf should be named Firstname_LastName_Sol.pdf, For your code, we want you to submit it both as .pdf and .jl. To generate the .pdf format of a Pluto notebook, you can easily click on the circle-triangle icon on the right top corner of the screen, called Export, and then generate the .pdf file of your notebook. The .pdf of your Pluto notebook as Firstname_LastName_Code.pdf while the .jl of your Pluto notebook as Firstname_LastName.jl. All code should be turned in when you submit your assignment.

Because assignments are more for learning, and less for evaluation, grading will be based on coarse bins. **The grading is atypical**. For grades between (1) 81-100, we round-up to 100; (2) 61-80, we round-up to 80; (3) 41-60, we round-up to 60; and (4) **0-40, we round down to 0**. The last bin is to discourage quickly throwing together some answers to get some marks. The goal for the assignments is to help you learn the material, and completing less than 50% of the assignment is ineffective for learning.

We will not accept late assignments. There is no late penalty policy. The assignments must be submitted electronically via eClass on time, by 11:59 pm Mountain time on the due date. There is a grace period of 48 hours when assignments will be accepted. No submissions will be accepted after 48 hours after the deadline, and the assignment will be considered as incomplete if not submitted.

All assignments are individual. All the sources used for the problem solution must be acknowledged, e.g. web sites, books, research papers, personal communication with people, etc. Academic honesty is taken seriously; for detailed information see the University of Alberta Code of Student Behaviour.

Good luck!