

ECE 3803, Fall 2021

Homework #5

Due Thursday, October 28, at 11:59pm

1. Prepare a one paragraph summary of what we talked about in class since the last assignment. I do not want just a bulleted list of topics, I want you to use complete sentences and establish context (Why is what we have learned relevant? How does it connect with other classes?). The more insight you give, the better.

2. Consider the function

$$f(x) = \log(1 + e^{ax}),$$

where $x \in \mathbb{R}$ and a is a constant.

- (a) Show that f is *not* strongly convex. [Hint: Recall the second derivative computed in homework 1.]
- (b) Determine a value of M for which f is M -smooth.
- (c) Now consider

$$\tilde{f}(x) = \log(1 + e^{ax}) + \delta x^2,$$

where $\delta > 0$ is a constant. Now show that \tilde{f} is both strongly convex and M -smooth (and give the associated constants m and M).

3. **1D optimization.** Consider the function

$$f(x) = \begin{cases} 25x^2 & x < 1 \\ x^2 + 48x - 24 & 1 \leq x \leq 2 \\ 25x^2 - 48x + 72 & x > 2 \end{cases}$$

- (a) Prove that f is strongly convex with parameter $m = 2$ and is M -smooth with $M = 50$.
 - (b) What is the global minimizer of f ? Justify your answer.
 - (c) Run gradient descent with $\alpha_k = 1/50$, the heavy ball method with $\alpha_k = 1/18$ and $\beta_k = 4/9$, Nesterov's method with $\alpha_k = 1/50$ and $\beta_k = (k-1)/(k+2)$, and Newton's method with $\alpha_k = 1$. In all cases initialize at $x_0 = 5$. Report how many iterations each method takes to converge.
4. **Smoothed data fitting.** Suppose that we observe a sequence of N observations given by $y_n = x_n + e_n$, where e_n represents noise, and consider the case where x_n represents samples of some underlying smooth function, so that we do not expect x_n to be too different from x_{n-1} and x_{n+1} . One approach to "denoising" the observations y_n to obtain a smoothed estimate of x_n is to solve the optimization problem

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad \frac{1}{2} \sum_{n=1}^N (x_n - y_n)^2 + \frac{\lambda}{2} \sum_{n=1}^{N-1} (x_{n+1} - x_n)^2.$$

- (c) The key question in fitting a logistic regression model is deciding on how to set \mathbf{a} and b based on the training data. To see how this is done, note that if $y_n = 1$, we would like to have $g(\mathbf{a}^T \mathbf{x}_n + b) \approx 1$, in which case

$$\log(g(\mathbf{a}^T \mathbf{x}_n + b)) \approx 0.$$

Similarly, if $y_n = 0$, then we would like $g(\mathbf{a}^T \mathbf{x}_n + b) \approx 0$, or said differently, $1 - g(\mathbf{a}^T \mathbf{x}_n + b) \approx 1$, in which case

$$\log(1 - g(\mathbf{a}^T \mathbf{x}_n + b)) \approx 0.$$

Note that in both cases, the terms on the left-hand side are always negative, so that to make them ≈ 0 corresponds to making these terms as *large* as possible. Combining these desired criteria together, we can express the problem of fitting \mathbf{a} and b as

$$\underset{\mathbf{a}, b}{\text{maximize}} \quad \sum_{n: y_n=1} \log(g(\mathbf{a}^T \mathbf{x}_n + b)) + \sum_{n: y_n=0} \log(1 - g(\mathbf{a}^T \mathbf{x}_n + b)).$$

It is somewhat more convenient to re-express this as the equivalent problem:

$$\underset{\mathbf{a}, b}{\text{maximize}} \quad \sum_{n=1}^N y_n \log(g(\mathbf{a}^T \mathbf{x}_n + b)) + (1 - y_n) \log(1 - g(\mathbf{a}^T \mathbf{x}_n + b)).$$

Show that, by plugging in the formula for g , this problem reduces to

$$\underset{\mathbf{a}, b}{\text{maximize}} \quad \sum_{n=1}^N y_n (\mathbf{a}^T \mathbf{x}_n + b) - \log(1 + e^{\mathbf{a}^T \mathbf{x}_n + b}).$$

- (d) We would now like to use an iterative algorithm like gradient descent to solve this optimization problem. To do this, we will need to calculate the gradient. Note that if we use the notation

$$\tilde{\mathbf{x}}_n = \begin{bmatrix} \mathbf{x}_n \\ 1 \end{bmatrix} \quad \boldsymbol{\theta} = \begin{bmatrix} \mathbf{a} \\ b \end{bmatrix},$$

then we can re-write our objective function as

$$f(\boldsymbol{\theta}) = \sum_{n=1}^N y_n \boldsymbol{\theta}^T \tilde{\mathbf{x}}_n - \log(1 + e^{\boldsymbol{\theta}^T \tilde{\mathbf{x}}_n}).$$

With this notation, compute a formula for the gradient $\nabla f(\boldsymbol{\theta})$ and the Hessian matrix $\nabla^2 f(\boldsymbol{\theta})$. [Hint: You may use the answers from homework 2 to avoid any actual computation here, but be careful as the notation in this problem is a little different...]

- (e) Implement a solver for estimating $\boldsymbol{\theta}$ using gradient descent. Test your implementation on the dataset produced by the following code:

```
import numpy as np
from sklearn import datasets
```

```
np.random.seed(2020) # Set random seed so results are repeatable
x,y = datasets.make_blobs(n_samples=100,n_features=2,centers=2,cluster_std=6.0)
```

Use an initial guess of $\boldsymbol{\theta}_0 = \mathbf{0}$. To begin, consider a fixed step size of $\alpha \approx 0.001$ (although you should feel free to play around with α a bit). Report how many iterations your algorithm takes to converge. Remember that you are trying to *minimize* $-f(\boldsymbol{\theta})$. If your algorithm is not converging, a sign error is a likely culprit.

- (f) Now implement the backtracking line search strategy described in the notes, and create a modified version of gradient descent in which you update α at each iteration using a line search. Experiment with the c and ρ . Report the c and ρ that work best, and how many iterations gradient descent requires when using backtracking.
- (g) Next implement a solver using the “heavy ball method”. First try setting $\alpha \approx 0.001$ and $\beta \approx 0.95$ and run the code with these fixed step sizes. Play around with α and β to see if you can get any improvement. Report the values of α and β that work best, and how many iterations the method requires. Then try using backtracking (using the same parameters as before) to adjust α at each iteration (you can leave β fixed here). Now how many iterations are required?
- (h) Next implement a solver using “Nesterov’s optimal method”. Try both using a fixed stepsize of $\alpha \approx 0.001$ as well as backtracking. Use the rule of thumb $\beta = (k - 1)/(k + 2)$. Report how many iterations are required for both versions.
- (i) Next, implement a solver using Newton’s method. Try both using a fixed stepsize of $\alpha = 1$ as well as using backtracking to choose α (be sure to start with $\bar{\alpha} = 1$). Report how many iterations are required for both versions.
- (j) Finally, implement a solver using the BFGS method, again setting α using backtracking with $\bar{\alpha} = 1$. Report the number of quasi-Newton steps required.