
DSC 140A - Homework 04

Due: Wednesday, May 1

Instructions: Write your solutions to the following problems either by typing them or handwriting them on another piece of paper. Show your work or provide justification unless otherwise noted. If you write code to solve a problem, include the code by copy/pasting or as a screenshot. You may use `numpy`, `matplotlib` (or another plotting library), and any standard library module, but no other third-party libraries unless specified. Submit homeworks via Gradescope by 11:59 PM.

Problem 1.

In lecture, we saw that the Soft-SVM optimization problem can be framed as the problem of finding the parameter vector \vec{w} that minimizes the *regularized* empirical risk with respect to the hinge loss:

$$R_{\text{svm}}(\vec{w}) = \|\vec{w}\|^2 + C \sum_{i=1}^n \max\left(0, 1 - y_i \vec{w} \cdot \text{Aug}(\vec{x}^{(i)})\right).$$

It can be shown that a subgradient of the empirical risk is given by

$$\text{subgrad } R_{\text{svm}}(\vec{w}) = 2\vec{w} + C \sum_{i=1}^n \begin{cases} -y_i \text{Aug}(\vec{x}^{(i)}) & \text{if } y_i \vec{w} \cdot \text{Aug}(\vec{x}^{(i)}) < 1, \\ 0 & \text{otherwise} \end{cases}$$

The file below contains data suitable for a binary classification problem.

https://f000.backblazeb2.com/file/jeldridge-data/003-two_clusters/data.csv

The file contains three columns: x_1 , x_2 , and y . The first two columns are the features, and the third column is the label.

Using subgradient descent, train a Soft-SVM model on this data using your choice of **either**:

- $C = 10$. This will earn full credit, as long as subgradient appears to converge to (close to) the optimal solution.
- $C = 1000$. This will earn one point of extra credit, as long as subgradient appears to converge to (close to) the optimal solution. While still convex, this optimization problem is much harder because the contours of the objective function are much more elongated. You'll have to fight with the learning rate and the stopping criterion to get it to converge.

You should turn in the following:

- The final parameter vector \vec{w} .
- A plot of the data, showing each class as a different color, as well as the learned decision boundary and the lines where $H = 1$ and $H = -1$, respectively.


Tips: you may find it difficult to get the subgradient descent algorithm to converge if you use a stringent stopping criterion. Instead, you may want to use a smaller threshold and/or a fixed number of iterations and increase it if the plotted decision boundary does not look correct. Remember, you can plot the decision boundary and the lines where $H = 1$ and $H = -1$ by either solving for x_2 in terms of x_1 or by using matplotlib's `contour` function.

Problem 2.

The data set linked below contains data for performing non-linear regression. The first column is x (the independent variable), and the second column is y (the dependent variable).

https://f000.backblazeb2.com/file/jeldridge-data/010-nonlinear_regression/data.csv

Plotting the data shows that there is a non-linear relationship between x and y :



a) Consider the function:

$$H(x) = w_0 + w_1\phi_1(x) + w_2\phi_2(x) + w_3\phi_3(x) + w_4\phi_4(x) + w_5\phi_5(x),$$

where each $\phi_i(x)$ is a Gaussian basis function:

$$\phi_i(x) = \exp\left(-\frac{(x - \mu_i)^2}{\sigma^2}\right).$$

Assume that the basis functions are equally spaced, with $\mu_1 = -10, \mu_2 = -5, \mu_3 = 0, \mu_4 = 5, \mu_5 = 10$. Also assume that all of the basis functions have the same width parameter, σ .

Write a Python function `plot_h(w_0, w_1, w_2, w_3, w_4, w_5, sigma)` that takes in w_0, \dots, w_5 and σ and plots the function $H(x)$ on top of the data.

Using this function, guess values for w_0, \dots, w_5 and σ that make the plot of $H(x)$ look like it fits the data well. Provide your code, the values you guessed, and the plot of $H(x)$ that results from your guesses.

Note: your function doesn't need to fit the data perfectly, but should be reasonably close. There will necessarily be places where H does not fit the data well.

b) Write a function `phi(x)` which takes in a scalar x in the input space and maps it to a vector in the feature space using the Gaussian basis functions above. The function should return a vector of length 5, where the i th feature is given by $\phi_i(x)$. Your function should use $\sigma = 2$ for the Gaussian width, but the same locations as in the previous part.

For this part, provide 1) your code, and 2) the output of `phi(3)`.

c) Fit a linear regression model in feature space by minimizing mean squared error, again using a Gaussian width parameter of $\sigma = 2$. What is the learned \vec{w} ? Show your work.

d) Using the learned \vec{w} , plot the prediction function $H(x)$ on top of the data.

Hint: Can you reuse the `plot_h` function from part (a)?