

Instructor: A.A. Ahmadi

TAs: Aydin, Budway, Hua, Tziampazis

Due on Thursday, October 24, 2024, at 1:30pm EST, on Gradescope

For all problems that involve coding, please include your code.

Problem 1: Support Vector Machines (SVMs)

Recall our Support Vector Machines application of convex optimization from lecture. We have m feature vectors $x_1, \dots, x_m \in \mathbb{R}^n$ with each x_i having a label $y_i \in \{-1, 1\}$. The goal is to find a linear classifier, that is a hyperplane $a^T x - b$, where $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$, by solving the optimization problem

$$\begin{aligned} \min_{a,b} \quad & \|a\| \\ \text{s.t.} \quad & y_i(a^T x_i - b) \geq 1, \quad \forall i = 1, \dots, m. \end{aligned} \tag{1}$$

We will then use this classifier to classify new data points.

1. Uniqueness of the optimal solution.
 - (a) Is the objective function $\|a\|$ convex? Strictly convex?
 - (b) What about $\|a\|^2$? Is it convex? Strictly convex?
 - (c) Prove that the solution to (1) is unique.
2. We would like to show that the optimization problem (1) is equivalent to

$$\begin{aligned} \max_{a,b,t} \quad & t \\ \text{s.t.} \quad & y_i(a^T x_i - b) \geq t, \quad \forall i = 1, \dots, m \\ & \|a\| \leq 1, \end{aligned} \tag{2}$$

which is easier to interpret in terms of finding a classifier with maximum margin.

Show that if (1) is feasible (with a positive optimal value), then (2) is feasible (and has a positive optimal value). Conversely, show that if (2) is feasible (with a positive optimal value), then (1) is feasible (and has a positive optimal value). You can assume that there is at least one data point with $y_i = 1$ and one with $y_i = -1$ as otherwise there is nothing to classify.

3. Assume the optimal value of (2) is positive. Show that an optimal solution of (2) always satisfies $\|a\| = 1$.

Problem 2: SVMs with linearly separable data

Open the Matlab file `HWSVM.mat`. To do this, download the file into your working directory and open it by calling `load HWSVM` in Matlab. This will load 6 vectors into MATLAB. You will need three of these vectors ("`x1part2`", "`x2part2`" and "`ypart2`") for this part of the problem. These three vectors correspond to $m = 53$ points in \mathbb{R}^2 whose components $(x_1, x_2)_{i,i=1,\dots,m}$ are given in the first two vectors and whose labels y_i are given in the vector `ypart2`. For Python users, you can use the following code to load the data file.

```
1 import scipy
2 mat = scipy.io.loadmat('HWSVM.mat')
3 x1 = mat['x1part2']
4 x2 = mat['x2part2']
5 y = mat['ypart2']
```

1. Plot all the 53 points on a graph. We need to be able to tell the difference between points that are labelled 1 and points that are labelled -1 .
2. Solve optimization problem (1) and plot on the same graph the optimal linear classifier (hyperplane) and the two shifted hyperplanes corresponding to the boundaries of the margin. Give the equations of these three lines.
3. Which points are the support vectors? Give their coordinates.

Problem 3: SVMs with data that is not linearly separable

You will now need the data vectors "`x1part3`", "`x2part3`" and "`ypart3`" from "`HWSVM.mat`". These three vectors correspond to $m = 100$ points $(x_1, x_2)_{i,i=1,\dots,m}$ in \mathbb{R}^2 and an associated vector y which has the label of each point.

1. Let S be a set consisting of s points z_1, \dots, z_s in \mathbb{R}^k . The convex hull of S is defined as

$$\text{conv}(S) = \left\{ \sum_{i=1}^s \lambda_i z_i \mid z_i \in S, \lambda_i \geq 0, \text{ and } \sum_{i=1}^s \lambda_i = 1 \right\}.$$

In words, this is the set of points that can be written as a convex combination of the points in S . A geometric interpretation of this definition is given in Figure 1.

Define

$$A = \{(x_1, x_2)_{i,i=1,\dots,m} \mid y_i = 1\}$$




Figure 1: Convex hull of a set

and

$$B = \{(x_1, x_2)_{i=1, \dots, m} | y_i = -1\}.$$

We say that the sets A and B are linearly separable if there exists a hyperplane $a^T x - b$ that takes value ≥ 1 on A and ≤ -1 on B . Prove that if A and B are linearly separable, then their convex hulls do not intersect.

- For the numerical data given, find a point that is both in $\text{conv}(A)$ and $\text{conv}(B)$ using CVX/CVXPY. Plot this point on the graph and give its coordinates.

Hint: Write the problem as a convex optimization problem.

- Recall the following convex optimization problem from lecture that attempts to simultaneously minimize the number of misclassified points and maximize the length of the margin:

$$\begin{aligned} \min_{a, b, \eta} \quad & \|a\| + \gamma \|\eta\|_1 \\ \text{s.t.} \quad & y_i(a^T x_i - b) \geq 1 - \eta_i, \quad \forall i = 1, \dots, m \\ & \eta_i \geq 0, \quad \forall i = 1, \dots, m. \end{aligned} \tag{3}$$

Solve this problem for $\gamma = 1, 2, \dots, 10$ and generate two plots: The first one will give the length of the margin (counting both sides) as a function of γ ; the second one will give the number of misclassified points as a function of γ . Discuss the overall trends of the two plots; are they what you were expecting?

Problem 4: Hillary or Bernie?

You would like to use the knowledge you've acquired in optimization over the past few weeks to see if you could have predicted the outcome of each Hillary-Bernie race in the Democratic primaries. To make things easier, you consider only the counties in the tri-state area and New England, i.e., those that belong to the states of New York, New Jersey, Maine, New Hampshire, Pennsylvania, Vermont, Massachusetts, Connecticut, or Rhode Island.

Your goal is to find a linear classifier that, for each county, labels it either as a Bernie win or as a Hillary win. To do this, you have access to a feature vector comprising the following features: mean income, percentage of hispanics, percentage of whites, percentage of residents with a Bachelor's degree or higher, and population density.

1. Load the data file `Hillary_vs_Bernie` in MATLAB/Python. This file includes 4 parts: `features_train`, `features_test`, `labels_train` and `labels_test`.

In `features_train`, we have given you the feature vectors for 175 counties and in `label_train`, their corresponding labels (-1 is a Bernie win and 1 is a Hillary win). As there was a wide disparity in the orders of magnitude of the original data (average income is around 10^4 whereas the percentages are between 0 and 1), each feature vector has already been normalized by its standard deviation. The original data can be found at <https://www.kaggle.com/benhamner/2016-us-election> (as fact checking is popular at the moment :)). Solve problem (3) to build a linear classifier for this training set for $\gamma = 0.1, 1, 10$. For each value of γ , specify the optimal a^*, b^* obtained.

2. Test the performance of your classifier using the feature vectors from 21 other counties (given in `features_test`) by comparing the labels obtained to the ones given in `label_test`. Which γ gives you the highest success rate in terms of prediction? Take a look at the entries of a^* in this case – what does this suggest about the people who vote for Hillary compared to those who vote for Bernie?