

ECE 6254, Spring 2024

Homework # 6

Due Tuesday, April 23, at 11:59pm EST.

Problems:

1. **Neural networks outputs as probabilities.** Suppose we are applying a neural network to a classification problem with K classes. One approach to predicting a class label, as well getting some notion of “confidence” is to map an input to a discrete probability distribution over all K classes. One way to do this is to have the neural network will map an input data point $\mathbf{x} \in \mathbb{R}^D$ to a vector $\mathbf{z} \in \mathbb{R}^K$, then pass \mathbf{z} through the *softmax function* $\sigma : \mathbb{R}^K \rightarrow \mathbb{R}^K$. We will denote output vector of the softmax as $\mathbf{p} \in \mathbb{R}^K$, with i -th entry given as

$$p_i = \frac{e^{z_i}}{\sum_{k=1}^K e^{z_k}} \quad (1)$$

In this problem, we explore connections between the softmax function and projection onto the $K - 1$ dimensional simplex, defined as:

$$\Delta_{K-1} = \left\{ \mathbf{p} \in \mathbb{R}^K : \sum_{k=1}^K p_k = 1 \text{ and } p_k \geq 0 \forall k = 1, \dots, K \right\} \quad (2)$$

We will show that

$$\begin{aligned} \mathbf{p} = \sigma(\mathbf{z}) &= \arg \min_{\mathbf{x} \in \mathbb{R}^K} -\langle \mathbf{z}, \mathbf{x} \rangle - H(\mathbf{x}) \\ \text{s.t.} \quad &\sum_{k=1}^K x_k = 1 \\ &0 \leq x_k \leq 1 \quad \forall k = 1, \dots, K \end{aligned}$$

where H is the entropy of \mathbf{x} , defined as

$$H(\mathbf{x}) = -\sum_{k=1}^K x_k \log(x_k)$$

For this problem, you will need a slightly more general version of the Lagrangian/KKT conditions that the one in class in order to account for the equality constraint. This can be found here: https://en.wikipedia.org/wiki/KarushKuhnTucker_conditions.

- (a) What is the Lagrangian for the above projection?
- (b) Write down the KKT conditions for the above problem.
- (c) Using the KKT conditions, show that $p_i = \frac{e^{z_i}}{\sum_{k=1}^K e^{z_k}}$.

2. **Classification and regression with neural networks** In this problem, we'll investigate the performance of neural networks¹ as both classifiers and regressors. In particular, we will use a very basic type of neural network called the multi-layer perceptron (MLP) for both settings. To get started, download `neural_net.py`.
- (a) The provided code loads the same digits dataset used in the previous problem, but with all 10 classes. Train a MLP classifier with the provided training set. Set `max_iter=1000`. Feel free to adjust other parameters such as `hidden_layer_sizes` (both size of the layers and number of layers) or `activation`. Report the test accuracy and the configuration of parameters that achieved this accuracy.
 - (b) The provided code generates noisy samples of the function $\sin(9x) + x$. Train a MLP regressor with the provided training set. Set `max_iter=1000`. Feel free to adjust other parameters such as `hidden_layer_sizes` (both size of the layers and number of layers) or `activation`. Report the mean squared error on the test set and the configuration of parameters that achieved this error. Include a plot of the true function and the neural network output.

¹In practice, neural networks are not implemented with scikit-learn, but rather with dedicated frameworks such as PyTorch (<https://pytorch.org>). Such frameworks often have a bit of a learning curve, so we use the scikit-learn implementation for this problem.