EEE 6002: Selected Topics in Electrical and Electronic Engineering – Privacy Preserving Machine Learning

Final assignment instructions

1. Choose the MNIST dataset and consider digits 5 and 8 to be the two classes.
2. Split the dataset into train, validation and test sets (e.g., using 70%, 10% and 20% splits). You can use sklearn library for splitting the data into these partitions.
3. Write the Logistic Regression code by yourself (do not use any built-in classification function).
4. Modify the Logistic Regression code such that it satisfies differential privacy (you can use either Laplace mechanism or Gaussian mechanism).
5. Include the following plots for both non-private and private classifiers for \epsilon = 0.01 and \delta = 1e-5:
    a. Training and validation loss vs iterations
    b. Norm of the gradient vs iterations
    c. Training and validation loss vs number of training samples
    d. Training and validation loss vs regularization coefficient λ
    e. Required number of iterations vs step size α
    f. Precision and Recall vs threshold
    g. True Positive Rate vs False Positive Rate (ROC curve)
6. Plot the classification accuracy against various \epsilon values for fixed \delta and sample size.
7. Plot the classification accuracy against varying sample size for a fixed \epsilon and \delta.
8. The report must contain a brief description of the dataset and the complete code.