

WEEKLY PARTICIPATION 8: LAYER NORMALIZATION VS BATCH NORMALIZATION

Modify the code given in class to compare the impact of layer normalization, batch normalization, and no normalization on the performance of a three hidden-layer MLP for classification on the FashionMNIST data set.

Submit plots of the validation accuracies and losses per training epoch. Use the same architectures, learning rates, and other hyperparameters as in the code given in class.