

# Final Exam 2023

Wine

Quality

Study

1. Download the wine Quality Dataset (dataset #2) from the link provided on the 568 website.
2. Implement a fully-connected NN for regression on wine quality scores. I suggest 2 hidden layers with 30 nodes each, but feel free to experiment. Choose activation functions for the hidden layers. Clearly describe and motivate your choices.
3. Split the data into 70% training & 30% test sets. Train the NN (choose an optimizer and describe/explain your choice). If you have the resources, run the training process multiple times and choose the best run: starting from random weights/biases on the training/test sets use RMSE as a measure of performance; use early stopping if necessary.

Alternatively, play with various constraint terms to avoid overfitting.

4. For your best run, plot RMSE (training) and RMSE (test) as a function of the number of iterations. Write down the best value of RMSE (test) you've been able to achieve.

good luck!