

CS 4320: Machine Learning

Assignment: Neural Network Classification

Continue using the [heart_failure_clinical_records_dataset](#) data set at Kaggle. Update your attempts to achieve a good model fit to this data by adding an artificial neural network model.

Your report will use the data exploration and analysis from the previous assignment, with any updated observations. It will also include the final results of your best decision tree and support vector classification models.

Additionally, the report will include details about the neural network model configurations attempted and the configuration of the model that achieved the best validation score.

It is expected that you will use the already modified hyper parameter search program, but add options for dealing with neural networks, similar to the changes made in class to the Titanic model fitter. It is expected that you will use the validation techniques provided by Tensorflow when fitting your models.

Create a report that includes the data exploration plots and analysis. The report will also include for each type of model previously fit (decision tree and svc), which hyper parameters produced the best $F1$ cross-validation score, and the $F1$ scores obtained on the training data and testing data.

The report will also include which neural network hyper parameters were searched, and the ranges or categories of values used in your search. Report the best validation score obtained. Note that you are probably using the cross entropy loss, which isn't the same as $F1$. Also include the $F1$ scores for the training and testing data.

Include a comparison of the full training $F1$ scores between the three models, and which model you would select, based only on those scores. Finally, discuss whether the $F1$ scores on the testing data support your decision or not.

Required Steps

- Download your data.
- Explore and analyze your data.
- Split the data 80%/20%, for training/testing.
- Write (or modify) a Python program to do the fitting and scoring.
- *AFTER* finding your best fit models, measure each model's $F1$ score on your test data.
- Create a report with the contents mentioned above.
- Commit and push your code in the git repository.
- Submit the report (as PDF) to Canvas.