DNN-TRND

A Comparison of Modern Deep Learning Methods for Time-series Trend Prediction

Summary

In this study TCNs and Bi-LSTMS were compared to MLPs, RNNs, CNNs and LSTM networks on their ability to predict trends in time series data using only a sequence of historical trends. Three datasets were used for the study: daily Cape Town air temperatures, daily S&P500 closing prices, and a household voltage dataset. Two experiments were run: the first was to use the traditional approach of training a model to predict both components of a trend simultaneously, and the second was to train independent models to predict trend slope and trend duration separately. The best performance was measured when using a TCN with a single prediction approach - though all DNNs performed well on the problem (each DNN was able to predict S&P 500 trends with above 84% directional accuracy). Across all DNNs, there was a slight performance improvement when using the proposed single prediction approach.

Objectives

The aim of the project was to compare TCNs and Bi-LSTMs with more commonly used DNNs (RNN, MLP, CNN, LSTM) in the problem of trend prediction. At the time of writing, both of these DNNs were relatively unexplored in trend prediction. Additionally, an alternative approach to trend prediction of using two separate DNNs to predict each trend component independently was tested and compared to the traditional approach of using one DNN to predict both components.

Data Preprocessing

The following preprocessing methods were applied to the time series:

1. **Missing Data Imputation**:
   Missing data was filled in using the most recent known value preceding it.
2. **Min-max Normalization:**
   Inputs were scaled to be between 0 and 1 for consistency.
3. **Median Filtering:**
   Each value was replaced with the median of itself and it's neighbors to reduce signal noise and help to isolate the underlying pattern associated with the data.
   See the alongside figure:

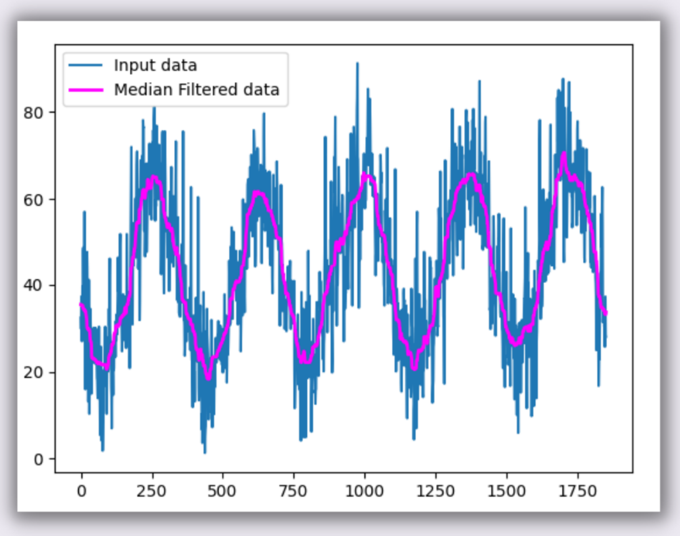

4. **Piecewise Linear Approximation:**
   The final and perhaps most important step of pre-processing was to use piecewise linear approximation to divide the data sets into trend line approximations of the data. This effectively converts the time series into a series of trend lines which are then used as the input for the neural network.

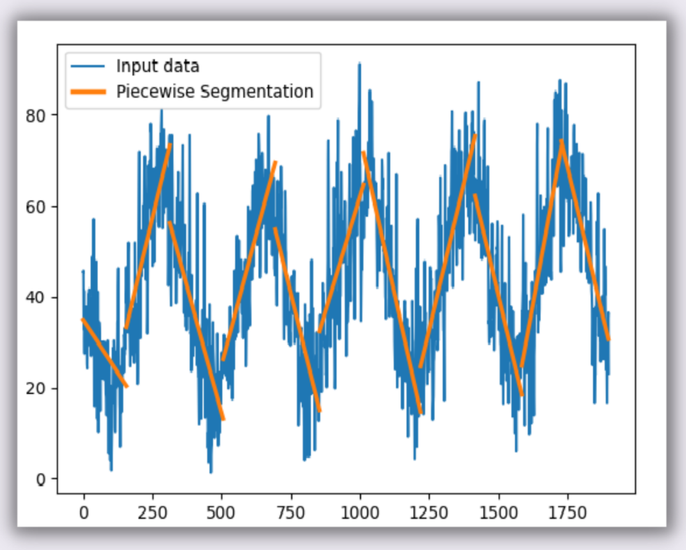

Performance Metrics

1. **Root Mean Square Error**
2. **Accuracy** (how often direction was predicted correctly)
3. **Sensitivity** (how often uptrends were predicted correctly)
4. **Specificity** (how often downtrends were predicted correctly)
5. **F-Score** (A composite score of classification performance)

Results

The RMSE performance hierarchy of each approach is shown below:

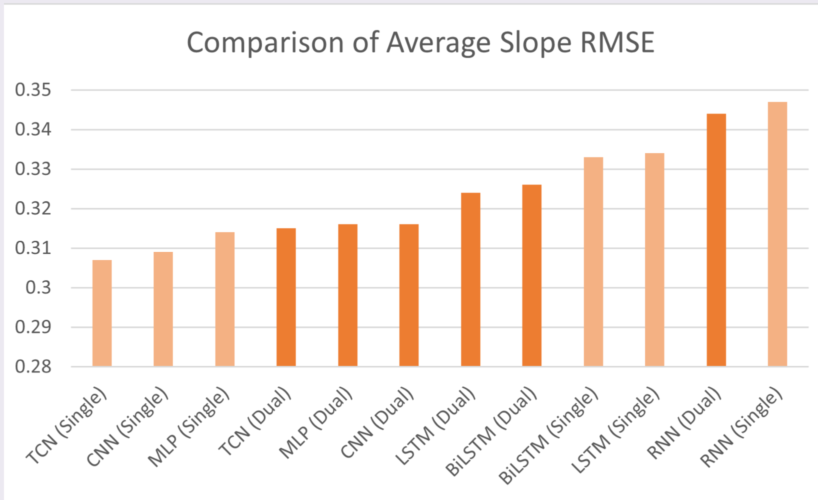

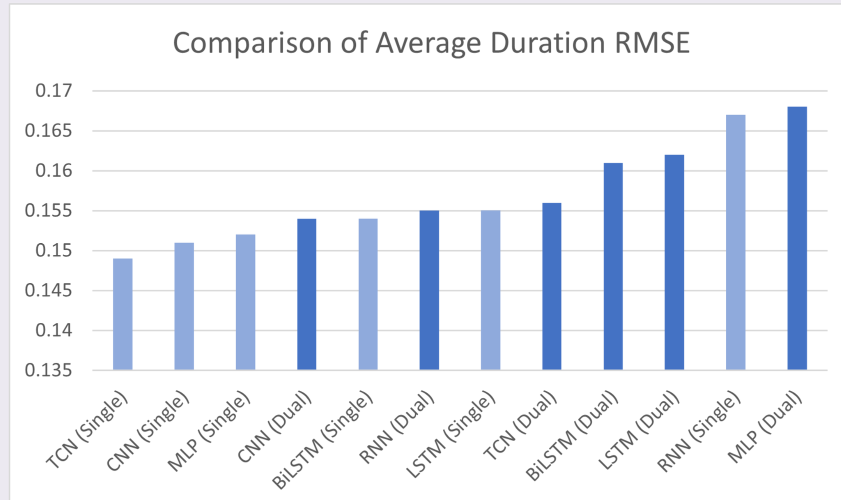

The results indicate that using a TCN with a single prediction approach yields the best results for both components. This is followed by the single output approach CNN and MLP. The recurrent neural networks had slightly worse performance, which may be attributed to the small size of the datasets after segmentation. To better understanding the RMSE metric, a variety of classification metrics were also calculated based on the models' ability to predict the direction of the trend. Below is a summary of the directional accuracy recorded by each model:

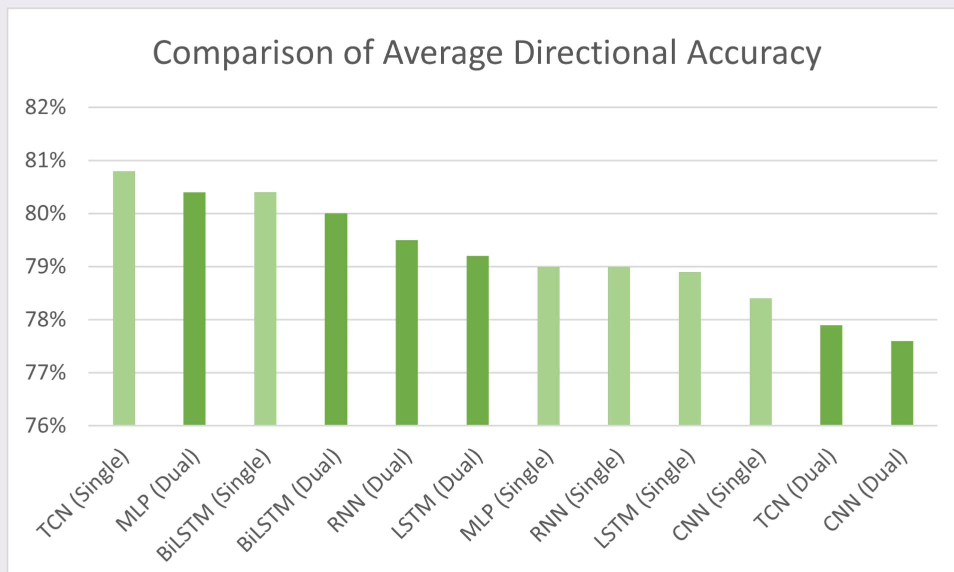

The single output TCN predicts the trend directional most accurately out of all models, followed by the dual output MLP and single output BiLSTM.

Conclusions

The single output approach was shown to yield better performance than the dual output approach for both slope and duration predictions. Additionally, the TCN was found to be the best performing model overall, while the single output standard RNN and dual output MLP were the two worst models for slope and duration, respectively. Although a clear performance hierarchy exists, it must be noted that the improvements are only slight. When analysing the classification metrics, the single output approach TCN is again the best performing model, followed by the dual output MLP model. Another consideration that should be made, however is that TCNs were found to have much slower train times than all other models which may impact model choice.

Authors

Morgan Blem
Chesney Cristaudo

Deshen Moodley (Supervisor)
Patrick Marais (2nd reader)

Contact

blmmor002@myuct.ac.za
crspie005@myuct.ac.za

deshen.moodley@uct.ac.za
patrick@cs.uct.ac.za

Institution

Department of Computer Science
University of Cape Town
Private Bag X3
Rondebosch, 7701
South Africa
Tel: +27 21 650 2663