APPLIED NEURAL NETWORKS: STOCK PRICE PREDICTION

Ajay Aryan,
Deven Gill,
Dionel Jaime

THE ERDŐS INSTITUTE
Revolutionary Collaborations in Academia and Industry

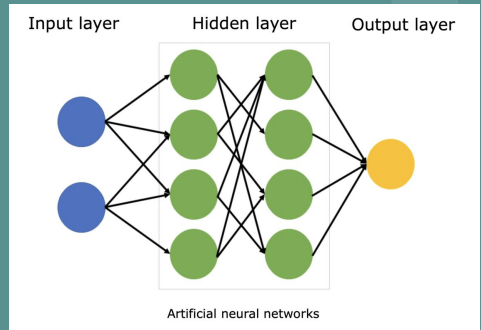

Table of Contents

What is the project about?

Predicting the stock prices is a complex task due to the volatile and dynamic nature of the financial markets.

Our Goal?

Build a model which can come close to accurately predicting the closing prices of a stock based on historical closing price data

Data Source:

From Kaggle for the SP 500 companies.

Data

- The data includes 5 years of stock data across 500 companies from 2013-02-08 to 2018-02-07.
- The data is clean without any missing or undefined entries.

date	# open	# high	# low	# close	# volume	Name
2013-02-08	15.07	15.12	14.63	14.75	8407500	AAL
2013-02-11	14.89	15.01	14.26	14.46	8882000	AAL
2013-02-12	14.45	14.51	14.1	14.27	8126000	AAL
2013-02-13	14.3	14.94	14.25	14.66	10259500	AAL
2013-02-14	14.94	14.96	13.16	13.99	31879900	AAL
2013-02-15	13.93	14.61	13.93	14.5	15628000	AAL
2013-02-19	14.33	14.56	14.08	14.26	11354400	AAL
2013-02-20	14.17	14.26	13.15	13.33	14725200	AAL
2013-02-21	13.62	13.95	12.9	13.37	11922100	AAL
2013-02-22	13.57	13.6	13.21	13.57	6071400	AAL

Advance Auto Parts (AAP) Data

- We test our model on the AAP stock whose first 200 observations are shown.
- The data has a trend but no seasonality.

2.) Our General Approach

- For each day we wish to forecast, our models are trained on all the days in our data which precede that day.
- The last 20% of stock data is used as testing data and will only be predicted by the model which has the best time series cross-validation error on the other 80% of data.
- Each of our methods involved a “sliding window” meaning that for each day we predict, we used a fixed number of preceding days to make our prediction.
- We used a neural network approach, a linear regression approach, and a naive method of predicting stocks.

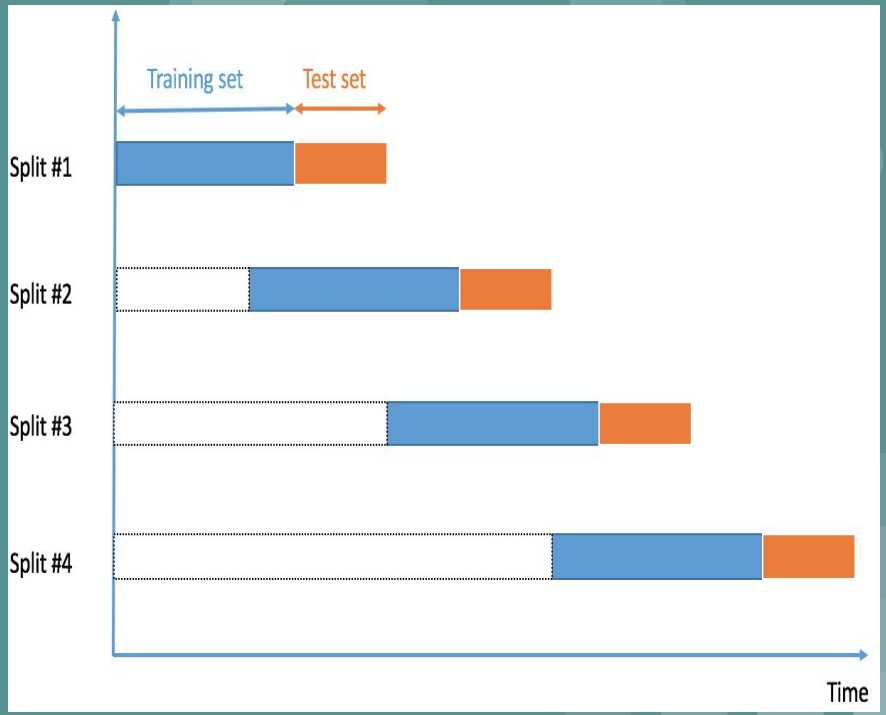

The Neural Network Approach: (LSTM)

Description: Long Short-Term Memory Models are neural networks which consider consecutive frames of data of some fixed window length and generate predictions based on this last group of frames.

Implementation:

- A window length of 7 days (3 business weeks) was used in training the model
- Architecture: Sequential model with dense layers: 1 input layer, 1 hidden layer, and 1 output layer.
- Average Validation MSE: $11.015

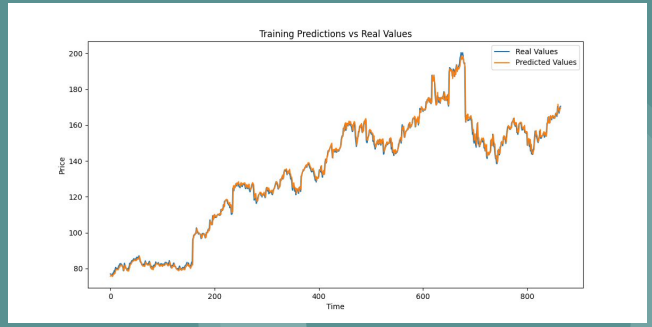

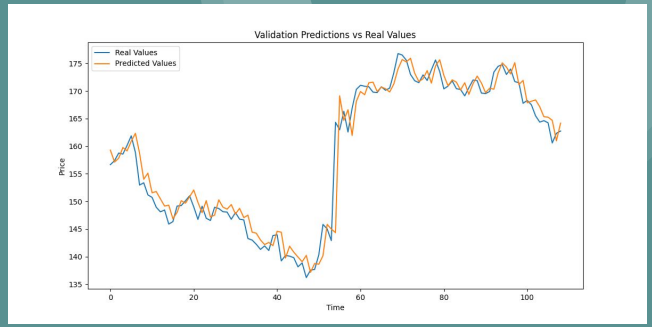

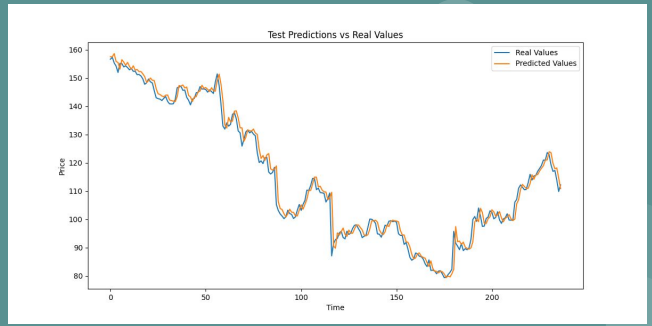

A Regression Approach:
Sliding Window Linear Regression

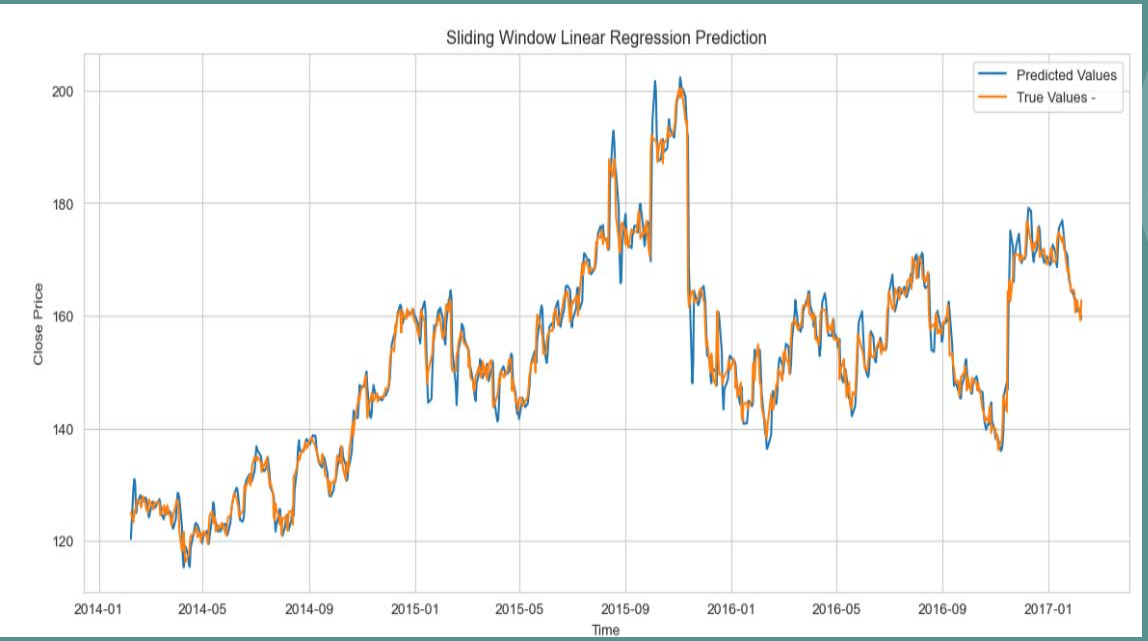

- A linear regression on 4 years of stock data evidently has its shortcomings in predicting stock data.
- Given a day which we wish to predict, *Sliding Windows Linear Regression* uses a relatively small number W, known as the *window size, and applies linear regression on the W preceding days to produce a prediction.*
- Since a stock's value on a given day is most influenced by its values from a few days preceding it, this should produce a low error.
- A window size of 5 gives the minimum time series cross-validation error of 10.58$.

Naive Prediction: Still the Best Model

- The Naive Prediction predicts the closing price to be the closing price of the previous day and thus uses a window size of only one day.
- The justification for this model is that stock closing prices, typically, do not change significantly from day to day.
- The Naive Prediction performed better than the other two models with a time series cross-validation error of 6.86$ and an MSE of 7.54$ on the test data.

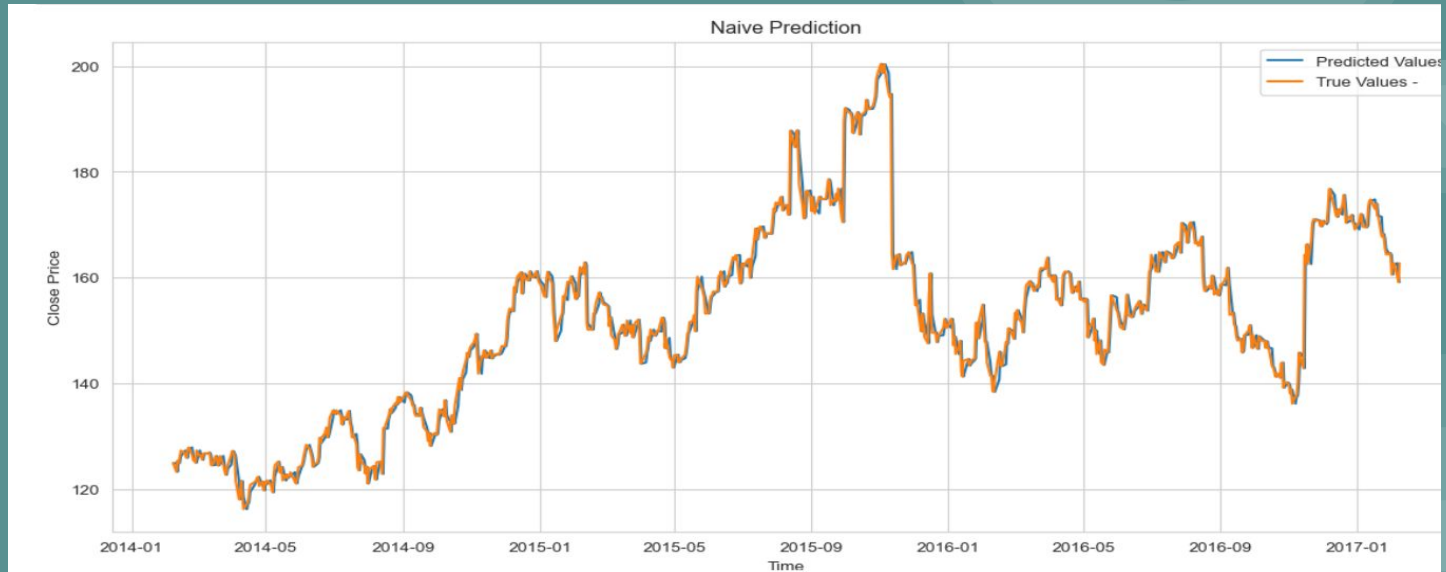

Future Work

- Currently, only the stock data for one company (AAP) is used for prediction. In future iterations, data for all companies will be incorporated to build a more comprehensive model.
- Incorporating additional drivers of stock prices in our models such as the political and economic environments, that affect the rates of change in prices.
- Develop the models and test them on stocks from countries that economically and significantly different from the US market.

Acknowledgements

- Steven Gubkin
- Alec Clott

THE ERDŐS INSTITUTE
Revolutionary Collaborations in Academia and Industry