Predicting stock price dynamics using stacked GRU's and LSTM's

Raul Overdijk Girbal
rgirbal@stanford.edu

Glib Stronov
gstronov@stanford.edu

OVERVIEW

Motivation: Financial markets have evolved to provide the modern investors with an abundance of data that can be used to make better investment decisions. However, it is not clear how to sift through that data to find important parameters that should influence our trading strategies. The goal of this project is to develop a model that can take past data about a stock or an index and predict the price change over some period of time into the future.

Models: We employed linear regression as the baseline, and built a multi-stack LSTM-GRU model that can predict prices over different time periods.

Results: Our best vanilla model has a RMSE test error of 0.022 for 1 day prediction, and our best model for 50-day predictions is the 5 block model with error of 0.53..

DATA & FEATURES

Our dataset spans all trading days on the NYSE from January 2000 to September 2018 . It includes 8 features: Opening Price, High Price, Low Price, Closing Price, Trading Volume, Moving Average Convergence Divergence, Moving Average Convergence Divergence Signal, and Relative Strength Index for S&P 500. We scrapped the price and volume data directly from Yahoo Finance and calculated the other indicators based on this price data. All of the data is then normalized with respect to the starting period.

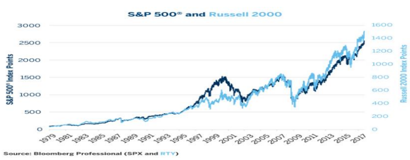

LOSS FUNCTION

We use Root Mean Square Error to evaluate the performance of our model. We chose this function to reflect the nature of trading strategies – big mistakes can lead to a prolonged drawdown period or even the collapse of the fund, while small losses don't have as much of a profound negative effect.

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n}(y_{pred,i} - y_i)^2}{n}}$$

OUR STANDARD BLOCK MODEL

We will use an LSTM-GRU stacking method along with a 20% dropout for our model. Our 2 stack model is shown as example on the right. Additionally we have some vanilla models detailed below with which to compare our performance with.

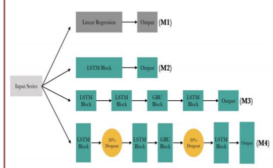

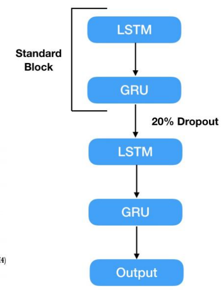

VANILLA RESULTS

The first is a trend we expected: the fact that at first linear regression performs far worse than our other models and then converges as we attempt to predict more and more days ahead. This reflects the fact that as we attempt to look further into the future, the information in the data relevant to price prediction decreases and more stochastic events can occur that dislodge the price from the trends that our model expected to hold. Additionally, dropout tends to help longer-term prediction performance since it helps create a more robust model.

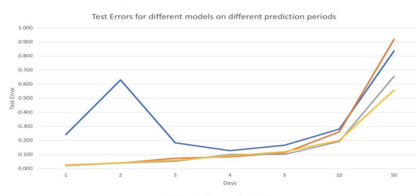

References:
Helmbold, David P., and Philip M. Long. "Surprising Properties of Dropout in Deep Networks." Proceedings of Machine Learning Research
Aungiers, Jakob. LSTM Neural Network for Time Series Prediction.
Schuermann, Lucas. Notes on LSTMs for Time Series Prediction in Finance
Shen, Guizhu, et al. Deep Learning with Gated Recurrent Unit Networks for Financial Sequence Predictions.

STACKING RESULTS

We note that the 10 block model performs worse than all other models, probably because the additional blocks are only causing more over-fitting despite the inclusion of 20% dropouts between each block. Secondly, we note that the ideal blocking number seems to be somewhere in between 3 and 5 blocks, especially given the large divergence in performance in the 10 and 50 predictions.

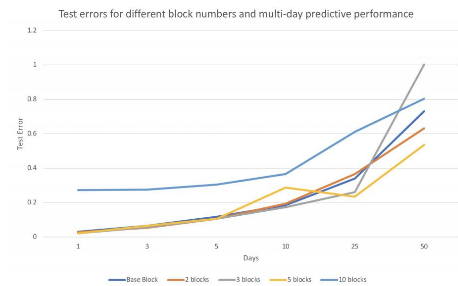

DISCUSSION AND FUTURE

Possible improvements

Currently, our model relies only on price and volume indicators as well as some macroeconomic factors. Given that there is only so much information that can be extracted from these indicators, we suspect that adding features that measure other market tendencies can result in a better performance.

A) One feature that we want to add is investor sentiment. Changes in sentiment can greatly affect the market and tracking those changes can help us predict the direction of this change and its effect on the price.

B) It would also be interesting to experiment with fundamental data, such as P/E ratios. Even though this data is available on a quarterly basis instead of daily, it can still be a valuable source of information.

Future work

In addition to adding the features discussed above, we would like to generalize our model to predict not only S&P 500, but also individual stock prices. We would then want to connect our model to a trading strategy based on our predictions and measure its performance in the market.