IBOVESPA index volatility forecasting using neural networks

Author: Gabriel Ogawa Cruz
Supervisor: Prof. Dr. Roberto Hirata Jr.
University of Sao Paulo - Institute of Mathematics and Statistics

Introduction

Volatility is a financial instrument to measure the price variation tendency over time. Specially for risk management, the ability to accurately forecast an asset's volatility over arbitrary time horizons is highly desirable and many methods and models have been developed towards that end.

In this work we use a Long Short-Term Memory **Recurrent Neural Network (RNN)** to forecast the IBOVESPA[a]index volatility (Fig. 2). Using a multitude of prices series correlated to the index to train the RNN, we show that complex related patterns can be learned to produce forecasts more accurate than more traditional models such as GARCH.

Recurrent Neural Network

A **RNN** is a neural network specialized for processing sequences of variable length allowing for dependencies through time to be learned, where each position in the input sequences is seen as a time step. In particular, a Long Short-Term Memory (LSTM) network is a recurrent network further specialized for handling long term dependencies by using gates and dynamically created paths to control the influence of past values over state updates, as can be seen in Figure 1.

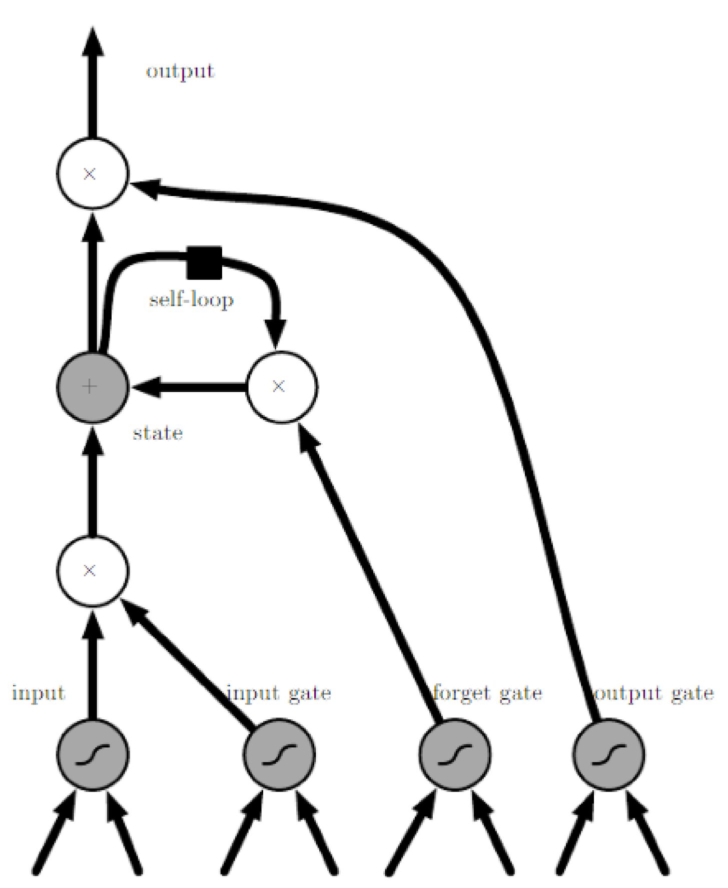

Figure 1: LSTM cell block diagram. [1]

Implementation

The neural network implementation was done in PyTorch. The chosen models were two RNNs with differing architectures but composed of a dual-layer LSTM unit, and a dual-stage attention-based RNN (DA-RNN) that added an attention mechanism between two LSTM layers.

A simple ARIMA was also fit to the data using *auto.arima* and a GARCH(1, 1) model was optimized with *garchFit* to generate forecasts for comparison.

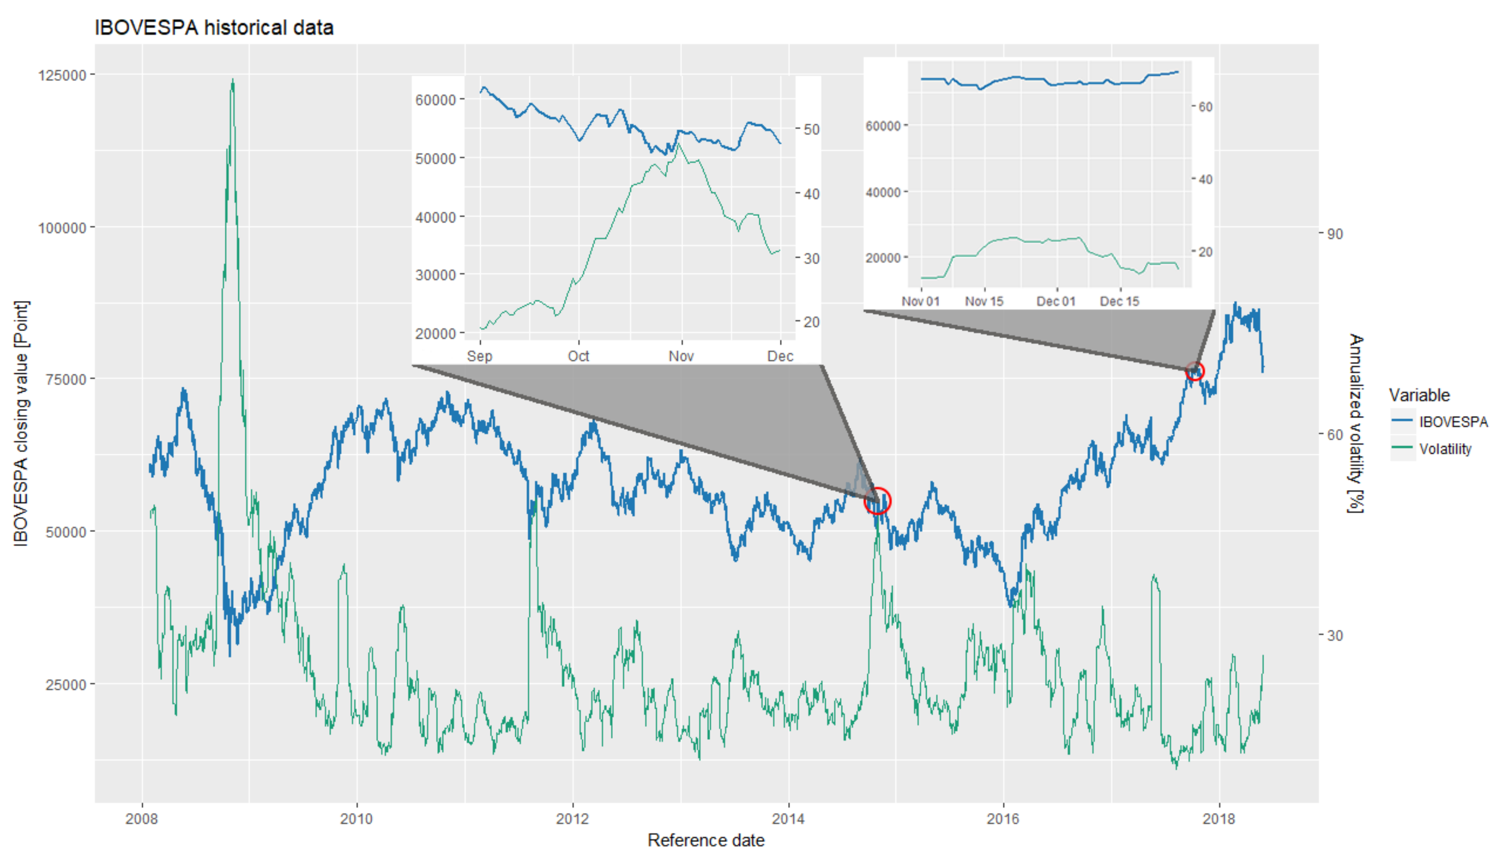

Figure 2: IBOVESPA historical data and its volatility series. A period of high volatility and another of low volatility are shown in the zoomed cut ins.

Dataset

From a dataset with 10 years of daily trading and economic indicators data the log-returns for the following series were chosen, by analyzing conditional correlation to the index, as inputs for training the neural network:

- IBOVESPA close points (IBOV)
- Dollar exchange rate
- Crude oil barrel spot price (WTI)
- 5 years interest rates
- Petrobras preferred stock spot price (PETR4)
- Vale common stock spot price (VALE3)
- Bradesco Bank preferred stock spot price (BBDC4)

The original (vanilla) and the outlier-filtered (cleaned) series were used separately for training and forecasting. The volatility measurement used was the standard deviation of log-returns over 20 days periods.

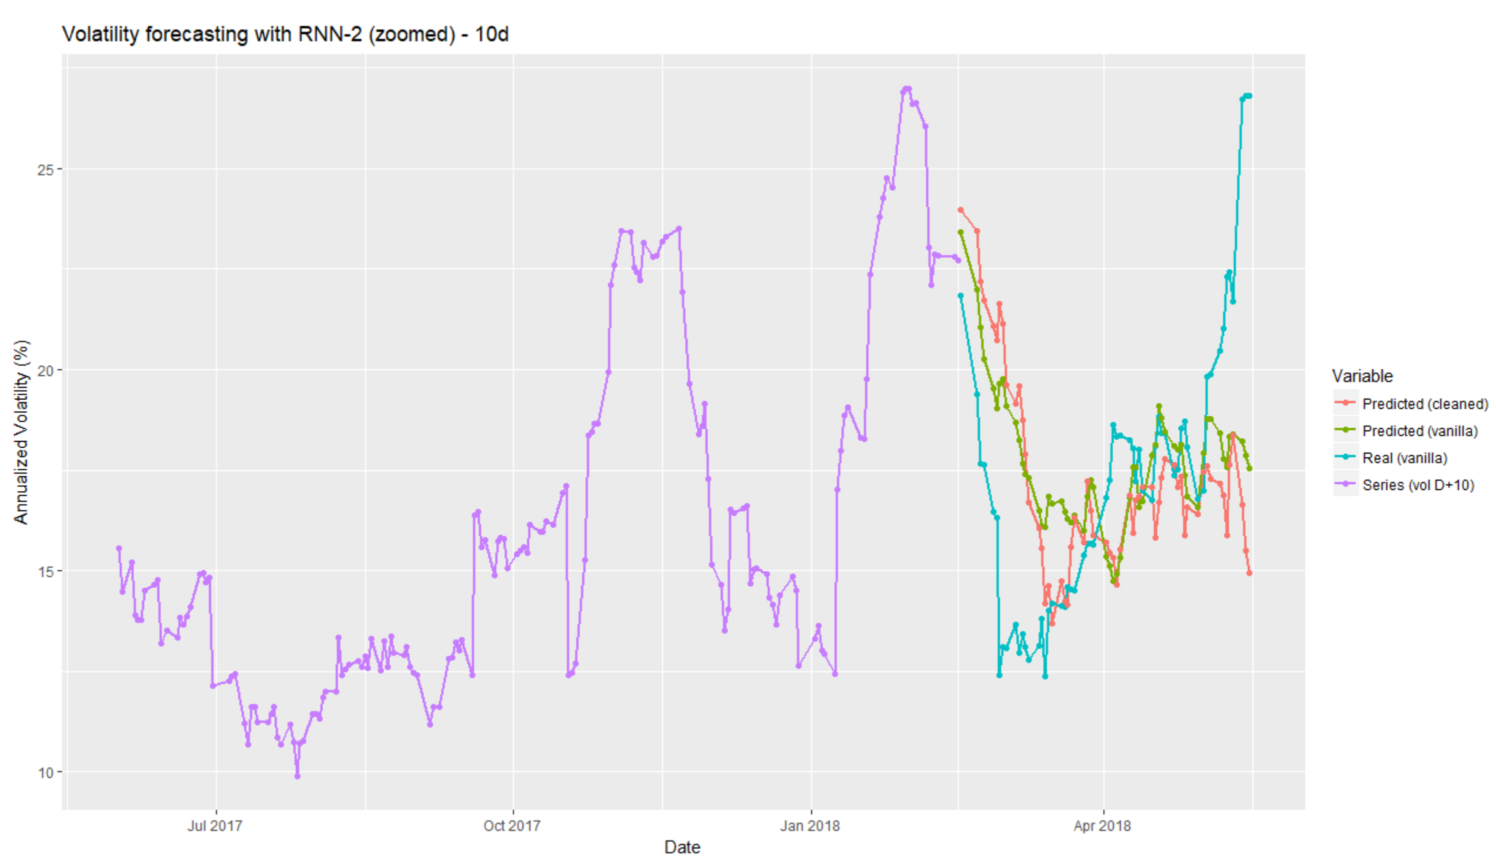

Figure 3: Forecasting results for the best performing recurrent neural network.

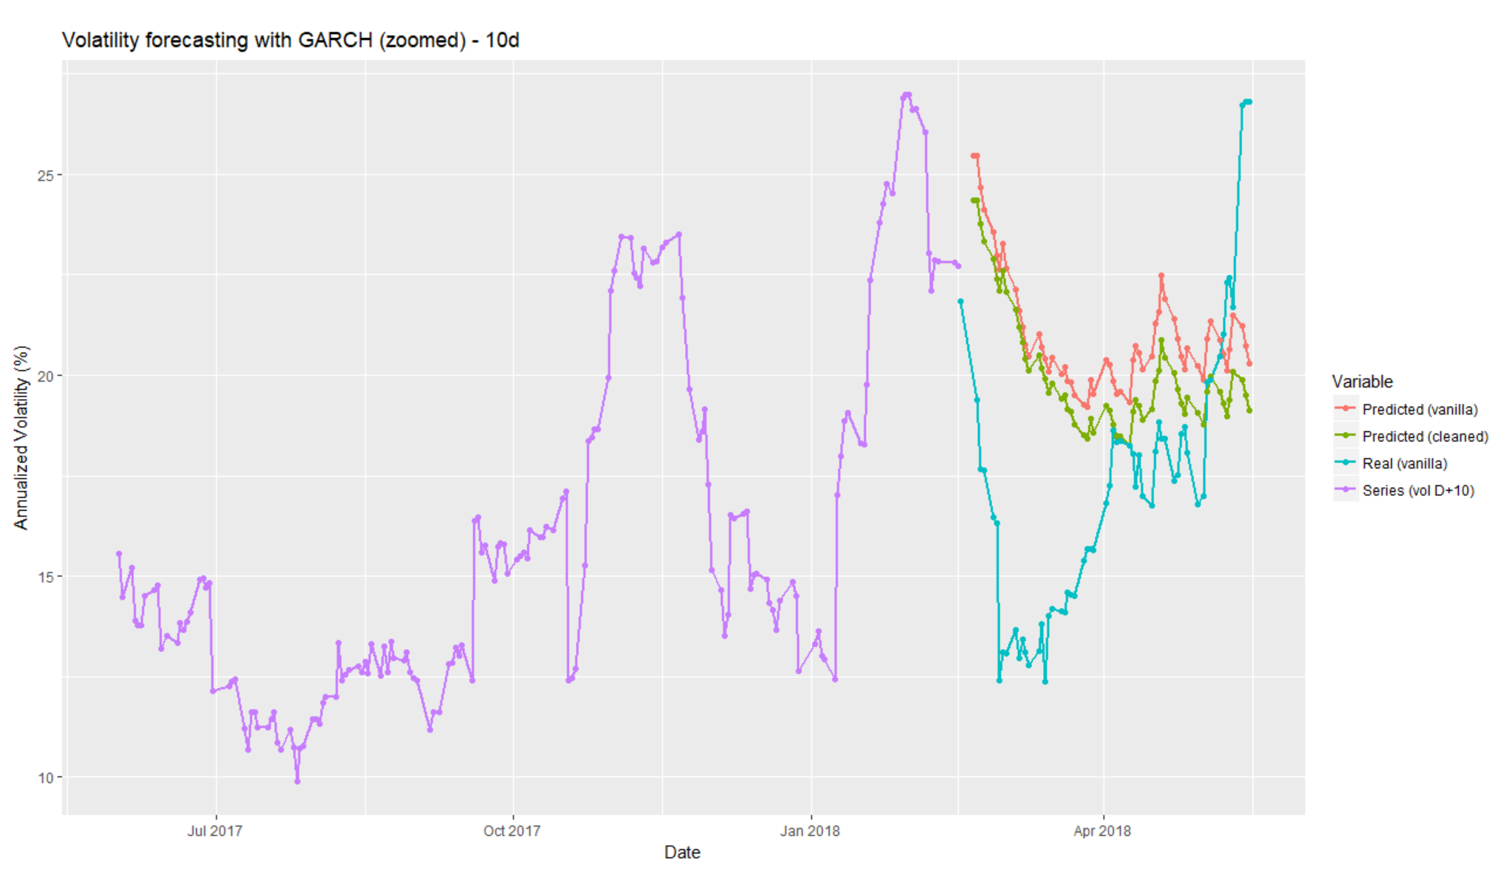

Figure 4: Forecasting results with a GARCH(1, 1) model.

Experiments and results

Every implemented model was fit over the vanilla (Van.) and cleaned (Clean) datasets in order to generate forecasts for time horizons of 10 and 21 days for the last 60 days of sample. The predictions were then compared to the real values by mean squares (MSE) and mean percent (MPE) errors. A comparison of the best performing model results to real values in the 10 days forecast case can be seen in Figure 3, the same graph for GARCH model results in Figure 4 and a comparison of the errors between models for the same horizon forecast is presented in Table 1.

Model	MSE	MPE
ARIMA Van.	24.13	23.11%
ARIMA Clean	293.75	80.30%
GARCH Van.	27.07	29.00%
GARCH Clean	22.53	24.85%
RNN-1 Van.	19.92	22.89%
RNN-1 Clean.	19.65	20.25%
RNN-2 Van.	11.96	16.06%
RNN-2 Clean	16.80	17.42%
DA-RNN Van.	21.52	21.97%
DA-RNN Clean	22.67	21.20%

Table 1: Forecast accuracy by model - 10 days prediction

Conclusion

Results suggest recurrent neural networks are capable of producing significantly more accurate forecasts when compared to other methods. The difference in results between networks suggests further architecture and training method tweaking would likely improve results even further. It is also possible that adding other relevant series to the network's input could lead forecasts even closer to real values.

References

[1] Ian Goodfellow, Yoshua Bengio, and Aaron Courville.
Deep Learning.
MIT Press, 2016.
http://www.deeplearningbook.org.

[2] BMFBOVESPA.
Metodologia do indice ibovespa.

[3] Source code, dataset and thesis.
https://github.com/ogaw4/MAC0499-Ibovespa-Volatility.

[a] IBOVESPA is Brazil's main stock market index, composed by stocks that represent 85% of total trading in the country's sole stock exchange, B3 [2].