Generative Adversarial Network for Stock Market Price Prediction

Ricardo Carrillo

racr@stanford.edu

INTRODUCTION

MOTIVATION
- Despite our extensive exploration with GANs, we found that the relative performance of GAN models with respect to traditional deep learning models such as LSTM has not been assessed.

INPUTS/OUTPUTS
- Input: 20-days sequence of trade data, including open price, close price, highest price, lowest price and volume.
- Output: A prediction of the movement of the stocks close price on the 21th day.

APPROACH
- Baseline Model: ARIMA model, Shallow LSTM and Deep LSTM.
- Experimental Model: The GAN architecture used has a three-layer dense network as a generator and a three-layer CNN as the discriminator. The discriminator was implemented with three convolutional layers, followed by a flattened layer and one dense layer with sigmoid activation.

RESULTS
- The experimental model doesn’t outperform traditional deep learning models such as LSTM.
- GAN results were more consistent in regard to up and down predictions than LSTM.

Data

SOURCE
- We used **Alpha Vantage** and US bundle provided by **Quandl**.

STRUCTURE
- 20 daily values of the following daily price data: open price, high price, low price, close price, and volume for 500 US Standard & Poor's Companies.

SPLIT
- **373K** training samples, **18K** validation, **18K** test

Methods

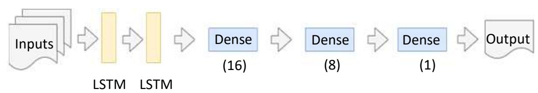

Figure 1: Deep LSTM, two LSTM layers with 64 and 32 hidden units concatenated with three fully connected layers. The first three dense layers used tanh activation and the final dense unit a linear activation.

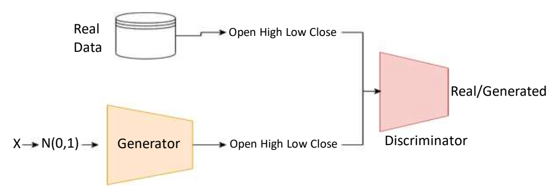

Figure 2:The GAN architecture used has a three-layer dense network as a generator and a three-layer CNN as the discriminator. The discriminator was implemented with three convolutional layers, followed by a flattened layer and one dense layer with sigmoid activation.

Results

Model	Train Set Up Movement Accuracy	Test Set Up Movement Accuracy
ARIMA	N/A	59.16%
Shallow LSTM	59.95%	74.16%
Deep LSTM	79.26%	62.85%
GAN	73.04%	72.68%

Results Cont.

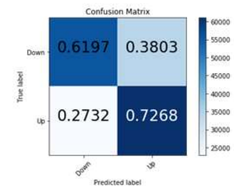

Figure 3: GAN Test Set Confusion Matrix, 50K Epochs, Adam optimizer and learning rate of 1e-4. Training Set data from 500 companies that are components of S&P 500.

Conclusion

Summary:
- GAN network architectures can be used to make representations of time series and specifically representations of the distribution of the stock price of capital markets.
- There are no significant differences with GAN and models traditionally used such as LSTM. Another finding is the fact that more effort was required for the tuning of models with many features.

Future Work:
- GAN Architecture. Explore deep learning models used in time series (LSTM) as GAN Generator.
- GAN loss. Explore loss functions different from traditional ones with GANs, such as WGAN, which uses Wasserstein distance.
- More GAN Exploration. Test time series with Metropolis Hastings GAN (MHGAN), where after training the generator, instead of discarding, the discriminator uses it to select the distribution closest to the actual distribution created by the generator.

References

S.Siami-Naminiand A.S.Namin, “Forecasting Economics and Financial Time Series: ARIMA vs. LSTM,” pp. 1–19, 2018.

K. Zhang, G. Zhong, J. Dong, S. Wang, and Y. Wang, “Stock Market Prediction Based on Generative Adversarial Network,” Procedia Computer Science, vol. 147, pp. 400–406, 2019.