

Predicting Stock Prices with A Recurrent Neural Network

Introduction

Predicting stock prices is an fascinating pursuit for investors and researchers. A variety of models that forecast changes in stock market prices have been introduced. Keim et al. [1] showed that predicting stock returns with statistical significance was feasible if a model with several predetermined variables were used. French et al. [2] attempted to use the relationship between a stock's volatility and its return in the generalized autoregressive conditional heteroscedasticity (GARCH) model to predict stock prices. Fama and French [3] proposed two-factor models using size and book-to-market equity, both of which are related to a company's fundamental information, to predict stock prices. Jeantheau [4] suggested that the autoregressive conditional heteroskedastic model could be applied to predict stock prices under stationary conditions. Ariyo et al. [5] proposed a predictive model for short-term future stock prices using the autoregressive integrated moving average model.

With the significant improvement of computational devices for the past 20 years, machine learning and deep learning techniques have found their place in the financial institutions to predict time series data with high degrees of accuracy. A prominent technique involves the use of artificial neural network (ANN). The most common form of ANN used in stock market prediction is the feed forward network that uses the backward propagation of errors algorithm to update the network weights. This type of ANN is commonly referred to as the backpropagation network. Another form of ANN that is

considered to be more appropriate for stock prediction is the time recurrent neural network (RNN) or time delay neural network (TDNN).

In this study, we adopt a recurrent neural network (RNN) and Long Short-Term Memory (LSTM) approach to predict stock market indices of the Google Inc. Because of its internal memory, an RNN is able to remember important things about the input it receives, which makes predicting what is coming next more precise. That is to say, an RNN can form a much deeper understanding of a sequence and its context. This is the reason that it is a preferred algorithm for sequential data, such as time series financial data. An RNN however suffers from short-term memory. If a sequence is long enough, it will have a hard time carrying information from earlier time steps to later ones. Therefore the LSTM is usually incorporated in an RNN.

The core concept of the LSTM is the cell state and its various gates. The cell state act as a transport highway that transfers relative information all the way down the sequence chain. We can think of it as the “memory” of the network. The cell state, in theory, can carry relevant information throughout the processing of the sequence. So even information from the earlier time steps can make its way to later time steps, therefore reducing the effects of short-term memory. Combined with the LSTM, an RNN is then able to effectively associate memories and input remote in time; hence it is appropriate to grasp the structure of data dynamically over time with high prediction capacity, in that the LSTM uses introduces a unit of computation that replaces traditional artificial neurons in the hidden layer of the network.

Method

Data preparation

Historical stock data is collected from the Google stock price between January 2012 and December 2017, and this historical data is used for the prediction of future stock prices. After the dataset is transformed into a clean dataset, the dataset is divided into training and testing sets so as to evaluate. Creating a data structure with 60 time steps and 1 output. Only the features that are to be fed to the neural network are chosen. We choose such features as Date, open, high, low, close, and volume.

Training Neural Network

The data is imported to the neural network, assigned random biases and weights, and trained for prediction. Our LSTM model is composed of a sequential input layer followed by 3 LSTM layers and dense layer with activation and then finally a dense output layer with linear activation function.

To adjust the learning rate, Adam optimizer is chosen in this study. An optimizer can greatly affect how fast the algorithm converges. It is also important that there is some notion of randomness to avoid getting stuck in a local minimum and not reach the global minimum. In addition, a Tikhonov form of regularization to ensure no overfitting occurs, where the weights getting too large and focusing on one data point.

Figure 1 shows an RNN being unfolded into a full network. By unfolding, we simply mean that we write out the network for the complete sequence. For example, if the sequence we care about is a sentence of 5 words, the network would be unrolled into a 5-layer neural network, one layer for each word. The formulas that govern the computation happening in a RNN are as follows. x_t is the input at time step t . s_t is the hidden state at

time step t . It's the “memory” of the network. s_t is calculated based on the previous hidden state and the input at the current step: $s_t = f(Ux_t + Ws_{t-1})$.




Fig. 1. Recurrent Neural Network

The function f usually is a nonlinearity such as tanh or ReLU. Here in this study, we use tanh. s_{-1} , which is required to calculate the first hidden state, is typically initialized to all zeroes. o_t is the output at step t . For example, if we wanted to predict the next word in a sentence it would be a vector of probabilities across our vocabulary. $o_t = \text{softmax}(Vs_t)$.

Results and Discussion

Table 1 presents the result of the prediction for January of 2017. The performance of the RNN and LSTM method was measured by computing root mean square error (RMSE). Compared with the RMSE, the predicted values are closest to the high point with an RMSE of 11.43 and are off the lowest point the most with an MSE of 16.20.

Table 1. Predicted stock prices for Google of Jan. 2017

Date	Open	High	Low	Close	Predicted
1/3/2017	778.81	789.63	775.8	786.14	797.434021
1/4/2017	788.36	791.34	783.16	786.9	794.547302
1/5/2017	786.08	794.48	785.02	794.02	794.887695
1/6/2017	795.26	807.9	792.2	806.15	796.524963
1/9/2017	806.4	809.97	802.83	806.65	799.988647
1/10/2017	807.86	809.13	803.51	804.79	805.845764
1/11/2017	805	808.15	801.37	807.91	811.06488
1/12/2017	807.14	807.39	799.17	806.36	812.943359
1/13/2017	807.48	811.22	806.69	807.88	812.985779
1/17/2017	807.08	807.14	800.37	804.61	812.766174
1/18/2017	805.81	806.21	800.99	806.07	812.904358
1/19/2017	805.12	809.48	801.8	802.17	813.177246
1/20/2017	806.91	806.91	801.69	805.02	813.430115
1/23/2017	807.25	820.87	803.74	819.31	814.193665
1/24/2017	822.3	825.9	817.82	823.87	815.208557
1/25/2017	829.62	835.77	825.06	835.67	819.83197
1/26/2017	837.81	838	827.01	832.15	826.811829
1/27/2017	834.71	841.95	820.44	823.31	834.193115
1/30/2017	814.66	815.84	799.8	802.32	838.011536
1/31/2017	796.86	801.25	790.52	796.79	833.483154

Visually inspecting from figure 1, we can see that the trend of predicted and actual prices are not off much, suggesting the prediction is acceptable. However, looking at the graph, it is clearly that the prediction curve is a replication of actual stock curve with some delay. Whether it is due to the memory effect in the neuron cell needs more investigation.




Figure 1. Real stock price for Google from Jan. 2016 to Dec. 2017

The popularity of stock market trading is growing rapidly, which is encouraging researchers to find out new methods for the prediction using new techniques. The forecasting technique is not only helping the researchers but it also helps investors and any person dealing with the stock market. In order to help predict the stock indices, a forecasting model with good accuracy is required. In this preliminary study, we used the Recurrent Neural Network and Long Short-Term Memory unit to forecast the stock price of Google Inc. Although this research is still going on to make the models better, the result provides investors and researchers with knowledge of the future trend of the stock market.

Reference

1. Keim DB, Stambaugh RF. Predicting returns in the stock and bond markets. J financ econ. 1986;17(2):357–90.
2. French KR, Schwert GW, Stambaugh RF. Expected stock returns and volatility. J financ econ. 1987;19(1):3–29.
3. Fama E, French K. The cross-section of expected stock returns. J Finance [Internet]. 1992;47(2):427–65
4. Jeantheau T. A link between complete models with stochastic volatility and ARCH models. Vol. 8, Finance and Stochastics. 2004. p. 111–31.
5. Ariyo A a., Adewumi AO, Ayo CK. Stock Price Prediction Using the ARIMA Model. In: 2014 UKSim-AMSS 16th International Conference on Computer Modelling and Simulation Stock [Internet]. 2014. p. 106–12.