

[Inspiratron.org](https://inspiratron.org) - [Natural language processing, machine learning and cybersecurity](#)

Predicting stock prices using recurrent neural networks (LSTM)

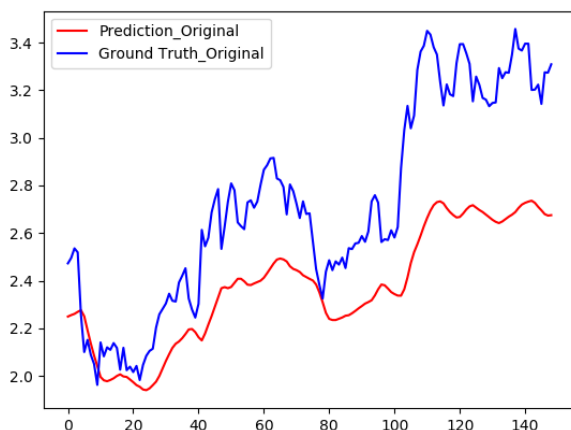
by Nikola Milošević? - Wednesday, January 10, 2018

<https://inspiratron.org/blog/2018/01/10/predicting-stock-prices-using-recurrent-neural-networks/>

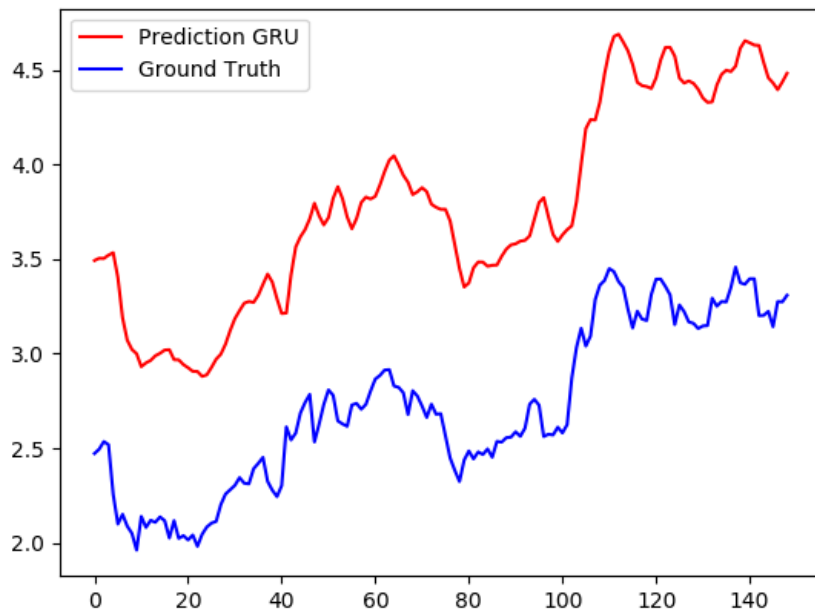
Just two days ago, I found an interesting project on GitHub. I was reminded about a paper I was reviewing for one journal some time ago, regarding stock price prediction using recurrent neural networks that proved to be quite good. I searched the web for recurrent neural networks for stock prediction and found the following project: <https://github.com/Kulbear/stock-prediction>

So I decided to try it out with Microsoft and Apple stocks and maybe try to do some slight modifications. The algorithm works with LSTM (Long Short Term Memory) neuron cells. LSTM neurons are a bit more complex cells than standard neurons, as they can be seen as a sequence of neurons, containing 3 gates. Memorization is achieved through the memory cell, which is implemented as identity activation function. There is an internal recursion loop, but also whole cell can have external recursion as well.


However, back to the project, I found. It stated that it can predict price movement quite accurately by taking into account the window of 20 prices. The prices I used was daily close prices of the stocks, namely Microsoft and Apple. The algorithm trains it on a network containing 2 relatively small layers of LSTM (I believe the size was the number of features, therefore 1), over 300 epochs and batch size of 748. The algorithm was scaling down stocks, therefore on the image prices are between 2-4. However, it can be scaled up. The image below is for Apple stock.



It generally can be said that it follows the movement of the stock prices, however, it is not totally accurate. Maybe over experimentation, I reduced the number of epochs. So I tried some modifications. One was to use GRU instead of LSTM. GRU stands for Gated Recurrent Unit, and it is also a type of recurrent neural network with gates having similar characteristics as the LSTMs for certain tasks. However, my modification reduced the number of epochs to 10-20, but increased the number of cells in the first layer to 1000, while in the second layer to 200. I won't include the code, because it can be easily done (in case you can't do it, it can be sold). The results were:




As I said it can be rescaled, so the



predictions would be

The algorithm predicts slightly higher prices but matches the direction of price movement almost perfectly. The original algorithm for Microsoft worked with the following plot:



Now the next experiment was to return to LSTM, but with an increased size of layers. They should theoretically give the better scores, since theoretically with huge enough size of the layer, a neural network can learn almost anything learnable.

I have a couple of more ideas what to try next in improving the performance of this algorithm. So far, it was trained only on one stock, however, I am wondering whether it can be really general for predicting stock price movement, so if I merge data from multiple stocks, I will have more data to train over. I spent last two days playing with sizes of layers, the best insight from it is that wider layers help. Maybe I can try to add some dense layer at the end but does not seem to help. Anyway, it seems that for short-term prediction of the stock prices this works quite satisfactory. It is worth testing it over longer periods of times, like a week, and month. I may create at some point some advisory platform that will use this and some other techniques to advise on investments.

All rights reserved and copyrighted by inspiratron.org and Nikola Milosevic