

# Predicting Stock Movements Using LSTM

Matt Roche

## Overview

- Stocks and finance are an area saturated with data
- Using past sequential price data to govern future predictions

## Problem Statement

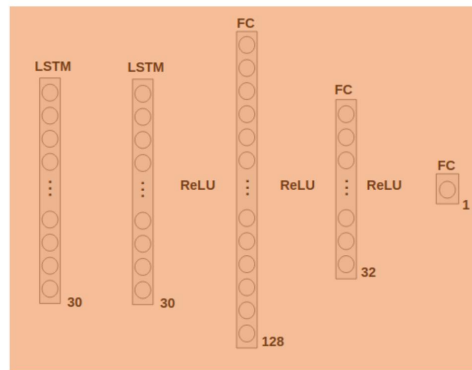
- Input: Vector of length 30
- Represents 30 days of closing stock prices
- Output: positive scalar, representing the predicted standardized price of the next day's closing

## Model Details

- 1<sup>st</sup> layer: LSTM with 30 units
- 2<sup>nd</sup>: LSTM with 30 units
- 3<sup>rd</sup>: FC with 64 units
- 4<sup>th</sup>: FC with 32 units
- 5<sup>th</sup>: output scalar

## Network Inputs

- Raw dataset contains daily closing prices from 1960s - present
- Inputs are segmented into 30-day sequences
- Sequences are individually standardized (divided by mean)



## Results

- Correct prediction of movement direction did not exceed 63%
- The model did not remarkably out-perform chance
- However, predictions in the correct direction were fairly accurate

## Analysis

- Stock price prediction is a difficult task
- It is often said that there is too much randomness in stock prices to make accurate predictions
- Other attempts have faced similar problems
- Optimistically, it is still possible to profit in expectation by following the model