(Un)certain asset price direction prediction using Bayesian Neural Networks

Authors

Aliaksandra Kupreyeva
aliaksandra.kupreyeva.4325@student.uu.se

Gustav Lind
gustav.lind.3933@student.uu.se

Andreas Lindgren
Andreas.lindgren.8632@student.uu.se

Supervisor

Robin Eriksson
robin.eriksson@it.uu.se

Introduction

The aim of this project is to implement a Bayesian neural network to predict the direction of a stock price and to analyze the uncertainties in the predictions.

Bayesian Neural Networks

Bayesian neural networks (BNNs) differ from standard artificial neural networks (ANNs) in that each weight is represented by a probability distribution. Thus, BNNs offer a way to quantify the uncertainties of the predictions by sampling from the distributions to obtain several different models and studying how those models predict on the same sample.

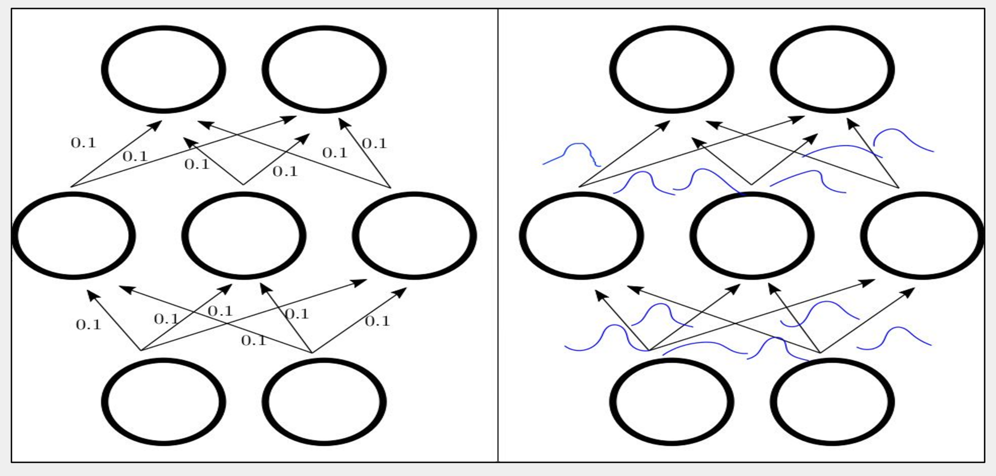

Fig 1. In the standard ANN (left) the weights are fixed values, and in the BNN (right) each weight is represented by a distribution of values.

Data

As data, we use historical prices from initial public offering until 1st of January 2021 for five large American technology companies. These are Microsoft, Facebook, Alphabet, Amazon, and Apple.

Method

We compute six different technical indicators to serve as the feature space, and assign binary labels, indicating whether the stock price returns are positive or negative.

When evaluating data points, we compute Monte Carlo sample predictions from the BNN. Each sample estimates the probability of the stock value increasing or decreasing and the distribution of the ensemble of predictions result in a normal distribution, as shown in Fig 2.

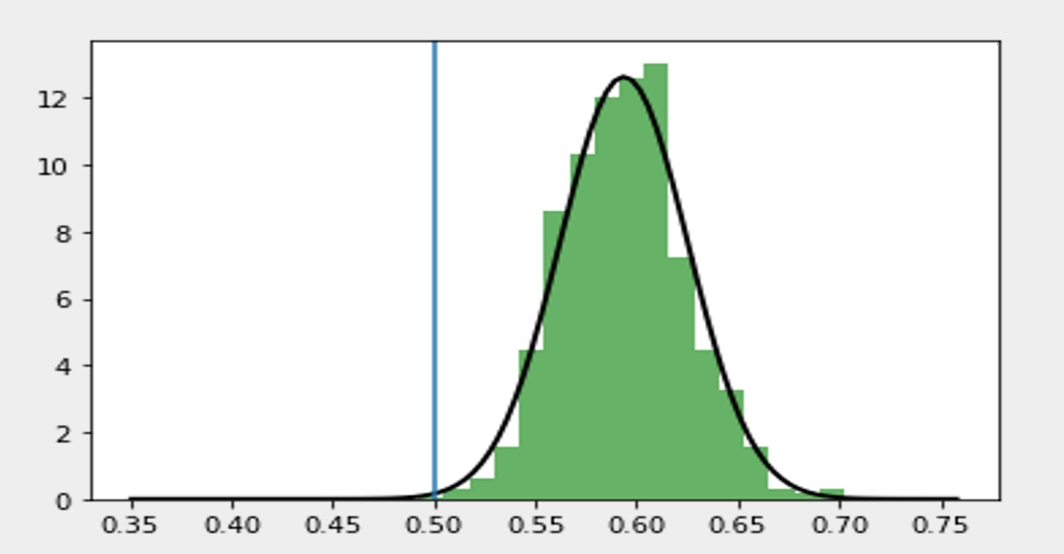

Fig 2. Empirical distribution of the predicted probability of positive 30-day forward return. Almost all of the predictions are above 0.5 (blue line) which indicates that the network is quite confident that the output is positive.

We use balanced accuracy(BA) as a metric to evaluate the performance of the model. BA is defined as the mean of recall and specificity. BA is especially useful when the classes are imbalanced in the data, i.e. more positive than negative returns.

Results

Accuracy

Tab 1 shows how the smoothing parameter α affected the BA of the model.

Tab 1. Balanced accuracy for different smoothing parameters

α	0.04	0.2	1.0
BA [%]	71.4	55.5	50.8

We successively reduced the test set to increase focus on the distributions indicating lower uncertainty. This resulted in an increased accuracy up to roughly 1.8-2.3 standard deviations, as seen in Fig 3.

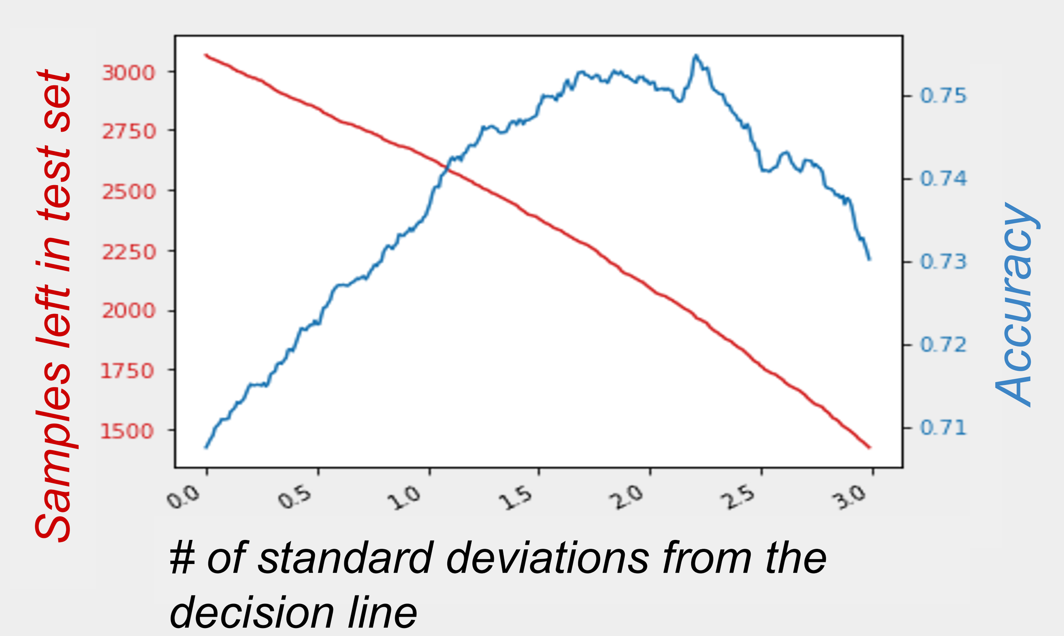

Fig 3. The effect on the accuracy as the test set is reduced depending on the uncertainty of each sample.

Extreme events

We also examine the network's performance on extreme events. Fig 4 shows the standard deviation of the model as it predicts 30 days forward in time for the Alphabet stock. The network successfully identifies that the situation is something it had limited experience of from training.

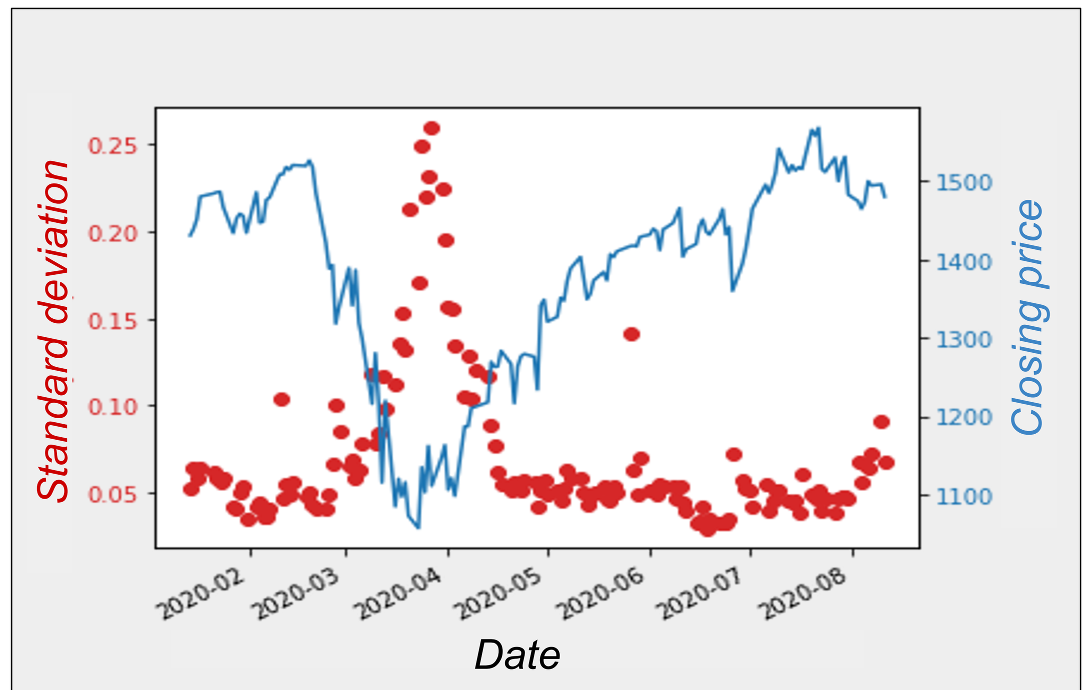

Fig 4. The standard deviation of predictions in relation to the price of the Alphabet stock.

Conclusions

- The parameter α, governing the smoothing of the data, proved to have a large impact on the classification performance.
- The random nature of stock movements makes it difficult to accurately predict them. However, we achieved 50.8 - 71.4% BA on the test set, depending on the α.
- The model successfully detected uncertain predictions for previously unseen scenarios, e.g. the 2020 covid19 crash.
- The prediction uncertainties did provide valuable information, and could be further studied.