

---

# BDLOB: Bayesian Deep Convolutional Neural Networks for Limit Order Books

---

Zihao Zhang, Stefan Zohren, Stephen Roberts

Department of Engineering Science, Oxford-Man Institute of Quantitative Finance, University of Oxford  
{zihao, zohren, sjrob}@robots.ox.ac.uk

## Abstract

We showcase how dropout variational inference can be applied to a large-scale deep learning model that predicts price movements from limit order books (LOBs), the canonical data source representing trading and pricing movements. We demonstrate that uncertainty information derived from posterior predictive distributions can be utilised for position sizing, avoiding unnecessary trades and improving profits. Further, we test our models by using millions of observations across several instruments and markets from the London Stock Exchange. Our results suggest that those Bayesian techniques not only deliver uncertainty information that can be used for trading but also improve predictive performance as stochastic regularisers. To the best of our knowledge, we are the first to apply Bayesian networks to LOBs.

## 1 Introduction

Electronic Limit Order Books (LOBs) [7] are commonly used in modern financial markets. We can consider LOBs as the most granular financial data that records all limit orders<sup>1</sup> of an exchange. Limit order data, like any other financial time-series data, is generally non-stationary and has a notoriously low signal-to-noise ratio.

Many models used historically for LOB data, such as the vector autoregressive model (VAR) [15], operate on handcrafted features extracted of the data to improve predictability. It is far from trivial to select informative features by hand. Recent research [11, 10, 14] has shown that deep learning models can deliver good predictive performance using only raw LOBs data with the process of feature extraction being automated by convolutional layers. However, such deep learning techniques are (typically) non-probabilistic and thus provide only heuristic uncertainty measures associated with output variables. For any decision making process, especially for high-frequency trading, trust and risk are primary concerns.

Uncertainty may be quantified through Bayesian inference. Given the complexity of network models, such *Bayesian Neural Networks* [1] are often achieved by approximation such as variational inference [12]. The work in [3] proposed *dropout variational inference*, also known as dropout sampling, as an approximation to BNNs. While variational dropout has been well established and deployed to many application domains, we have not found any literature applying it to LOBs.

*Our contributions:* We apply variational dropout to a large-scale deep learning model designed in [14] to predict short term movements in cash equity prices from LOBs. We showcase that dropout techniques, as stochastic regularizers, achieve better predictive performance and we design trading strategies that can deliver improved returns using derived uncertainty information. Further, we show how uncertainty relates to position sizing, avoiding bad trades and amplifying profits.

---

<sup>1</sup>Limit orders are resting order to either buy or sell an instrument, i.e. those that do not match immediately. All resting limit orders form the LOB which represents the supply and demand of a given financial instrument at any moment in time.

## 2 Methods

We here briefly introduce the dropout variational inference [2] and the network architecture, known as DeepLOB, in [14]. In addition, we discuss different trading strategies and demonstrate how uncertainty information derived from Bayesian networks can be used to decide entry and exit points and further to upsize or downsize our positions.

### 2.1 Dropout Variational Inference and Network Architecture

For Bayesian networks that have more than one hidden layer, the true posterior distributions of model parameters can not usually be evaluated analytically. Instead, we can define a variational distribution, which has a simple structure to evaluate, to approximate the true posterior distribution. Following this variational interpretation, the work of [2] showed that dropout can be considered as an approximating distribution to the posterior in a Bayesian network. In other words, the weights obtained from the optimisation of a network using dropout are the same as the variational parameters of a network that has the same structure. We can then model uncertainty in Bayesian networks by a technique called *Monte-Carlo (MC) dropout*, which is essentially running a number of stochastic forward passes through our network and averaging the results,

The network architecture in [14] is composed of three main building blocks: convolutional layers, an Inception Module and LSTMs. It is often difficult to extract representative features in financial applications as there are enormous observations and financial data is notoriously noisy. Hand-engineered features, such as technical indicators or PCA, often rely on tacit assumptions that may not well describe financial data and the process is non-trivial. The usage of CNNs and Inception Modules is then to automate the process of feature extraction, removing the above constraints. Feature maps obtained from convolutional layers are then passed into a LSTM layer instead of fully connected layers since these features still contain time dependencies. We can use LSTMs to capture temporal relationship among the resulting feature maps.

Dropout was not used in the above work but we apply it to obtain uncertainty estimates. However, in order to get well-calibrated uncertainty, we need to adapt the dropout probability and this process is often time consuming for the large model and dataset. As a result, we apply a dropout layer with rate 0.2 after the Inception Module, decided using grid-search methods. Note that the feature maps are still time-series, so we keep the dropout mask the same for all timesteps. Doing this would not break time dependencies but regularize resulting features. Further, we applied Concrete Dropout [4] to convolutional layers in the Inception Module and found that all dropout rates are pushed towards zero due to the large number of observations we have available. Further, we note that the convolutional layers have considerably less parameters compared to the dense layers because of the parameter sharing structure, so they are less likely to be overfitted [13].

### 2.2 Different Trading Strategies

The work of [14] classified future price movements into three classes (up=+1, neutral=0, down=-1) and used a softmax layer to give a probability vector at each time stamp  $\mathbf{p}_t = [p_{+1,t}, p_{0,t}, p_{-1,t}]^T$ . The basic trading strategy designed in [14] to test the quality of the prediction was based on discrete outputs and did not use the softmax probability vector. We briefly describe this strategy below and refer to it as **Normal trading strategy**. We then show how returns can be amplified by incorporating information from the softmax layer (**Softmax trading strategy**) and finally show how positions can be sized by using uncertainty information from the Bayesian networks (**Bayesian trading strategy**).

**Normal trading strategy** The model delivers a softmax probability vector,  $\mathbf{p}_t$ , at each time. We enter the long position if  $\max(\mathbf{p}_t) = p_{+1,t}$  or the short position if  $\max(\mathbf{p}_t) = p_{-1,t}$ . Only one position is allowed and we hold until the opposite movement is predicted (i.e.  $\max(\mathbf{p}_{t+k}) = p_{-1,t+k}$  in case of a long position where  $k$  is the holding period). We go long or short (buy or sell) the same number of shares ( $\mu$ ) at each time and set  $\mu$  to be 30% of first ask or bid level depending on the position we enter.

**Softmax trading strategy** Same as above, but we use softmax probability vector to constrain our timing of entering positions. Instead of entering the trades when  $\max(\mathbf{p}_t)$  is equal to  $p_{+1,t}$  or  $p_{-1,t}$ ,

we place a simple rule to only enter the trade if  $\max(\mathbf{p}_t)$  is greater than a threshold  $\alpha$ . We place no constraint to exit the trade.

**Bayesian trading strategy** While we use softmax outputs in the previous strategy, the network outputs could be biased and it is often argued [2] that softmax probability vectors can not be interpreted as true model confidences. Instead, three methods, variation ratios, predictive entropy and mutual information are suggested in [2] to summarise uncertainty in classification when using variational dropout. We tested on all three quantities and discover that predictive entropy,  $\mathbb{H}$ , gives most consistent uncertainty information and we use this information to size our positions. For a given input  $\mathbf{x}$ , we run  $M(= 100)$  stochastic forward passes through our network and average the probabilities of each class to obtain the averaged probability vector  $\bar{\mathbf{p}}_t$ . As before, we still enter trades if  $\max(\bar{\mathbf{p}}_t) > \alpha$ , but we upsize our positions to  $1.5 \times \mu$  if  $\mathbb{H} < \beta_1$  or keep original size  $\mu$  if  $\beta_1 < \mathbb{H} < \beta_2$  but downsize our positions to  $0.5 \times \mu$  if  $\mathbb{H} > \beta_2$ . We exit the position if  $\mathbb{H} < \beta_2$ . In next section, we show how these trading parameters affect our profits and risk level.

### 3 Results

The data we use are limit order book data for five stocks: Lloyds Bank, Barclays, Tesco, BT and Vodafone listed on London Stock Exchange for the entire 2017 year, totalling more than 134 million market quotes. We take the first 6 months as training data, the next 3 months as validation data and the last 3 months as test data. Our test data contains millions of observations that can thoroughly verify our model performance. We follow [14] to prepare our input data.

We denote our model as BDLOB and compare it to two baseline models: the CNN model presented in [11] and DeepLOB5 in [14]. We show each model’s precision, recall, F1 score and average AUC score in Table 1 and we present confusion matrices and boxplots for daily accuracy in Figure 1. BDLOB achieves the best results since dropout, as stochastic regularizers, improves generalization ability. Also, from Figure 1, we can see that BDLOB has better ability to classify stationary class and lead to consistent better performance with narrow interquartile range (IQR) and less outliers.

Table 1: Evaluation metrics for different baseline models

Model	Precision	Recall	F1 Score	AUC
CNN	0.52	0.52	0.48	0.672
DeepLOB5	0.60	0.60	0.58	0.803
BDLOB	0.60	0.61	0.60	0.811

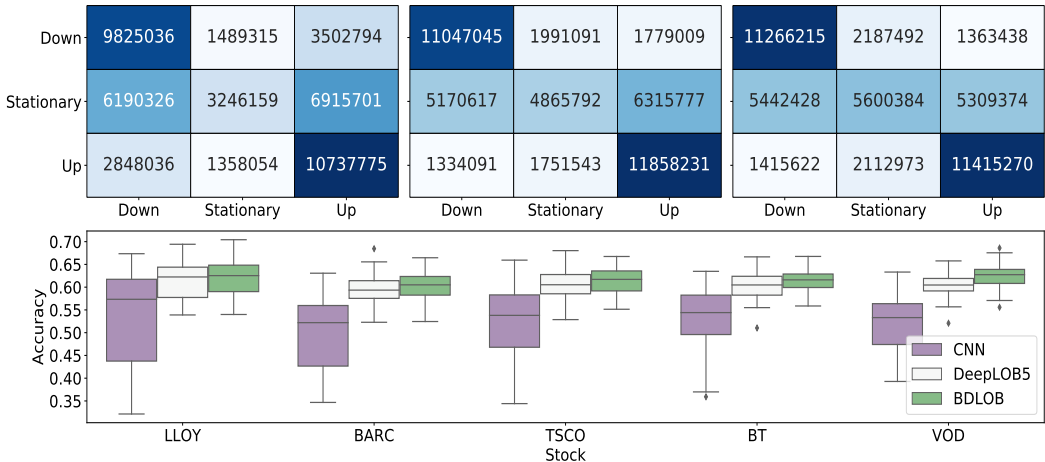


Figure 1: **Top:** confusion matrices for CNN, DeepLOB5 and BDLOB; **Bottom:** boxplots of daily accuracy for the different models across 5 stocks.

We follow the same assumptions, mid-price simulation and no transaction costs, as in [14] and present how different trading strategies affect profitability. Our focus is not to design a stand-alone trading strategy but to showcase how uncertainty information can be beneficial to decide enter and exit points and also position sizing. However, due to the short holding periods, most high-frequency trading firms are engaged in market making. Being able to trade passively on one side of the trade effectively corresponds to a mid-mid trade, being able to trade both sides passively reduces the fee further. The above, simplified trading simulation still approximates the relative value of the here presented model when added to an existing market making strategy.

We hold no stocks overnight and close all positions by the end of the day. As each strategy leads to different trading volumes, we normalise each day’s return by total volumes executed on that day. In the literature, the Sharpe Ratio [9] is often used to measure a strategy’s risk. However, the Sharpe Ratio is based on variance of returns that treats both positive and negative returns to be “risky”. In practice, few traders would consider large positive returns as risky so we use the Downward Deviation Ratio (DDR) [6] as a measure of risk. The DDR penalizes negative returns and rewards large positive returns. The DDR is defined as,

$$\text{DDR} = \frac{E(R_t)}{\text{DD}_T} \quad \text{where} \quad \text{DD}_T = \left( \frac{1}{T} \sum_{t=1}^T \min(R_t, 0)^2 \right)^{\frac{1}{2}} \quad (1)$$

Figure 2 presents the boxplots for daily profits normalized by total volumes on that day. Any strategy using uncertainty information leads to substantial improvements in returns and BDLOB delivers the most profits and maintain a similar risk level. Figure 3 shows the cumulative profits across testing periods and we observe consistent and better performance for BDLOB. We set  $\alpha = 0.7$  for all strategies that use softmax probability vectors and set  $\beta_1 = 0.1, \beta_2 = 0.9$  for Bayesian trading strategy. The value of these hyperparameters are selected to balance profits and risk. We discuss these choices below.

Figure 4 shows how profit and risk change according to these hyperparameters. We can see that a higher  $\alpha$  generally improves profits but  $\alpha = 0.8$  or  $\alpha = 0.9$  can lead to larger risk or lower profits – this is due to few trades taking place so any wrong prediction has a stronger impact on performance. As a result, we select  $\alpha = 0.7$  as the best value and apply it to all Bayesian trading strategies in Figure 4. We set  $\beta_1 = 0.1$  and find profit or risk is not hugely affected by this value because our model is very confident in its predictions when a high  $\alpha = 0.7$  is used. The effect of  $\beta_2$  is obvious as it not only penalizes unnecessary trades but also controls the exit time point. We can see that high profits can be obtained with a smaller  $\beta_2$  but with a much higher risk level. A strict exit rule allows us to take even fewer trades amplifying profits, but the impact of wrong predictions is also more obvious and we can hold a bad trade for too long without leaving it, leading to large variance.

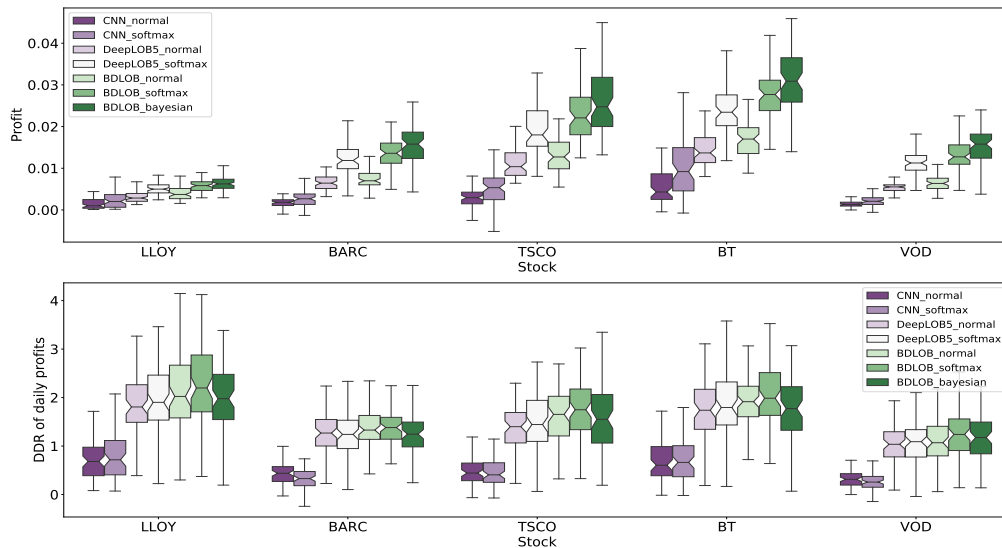


Figure 2: Boxplots and DDR for normalized daily profits. Profits are in GBX(=GBP/100).

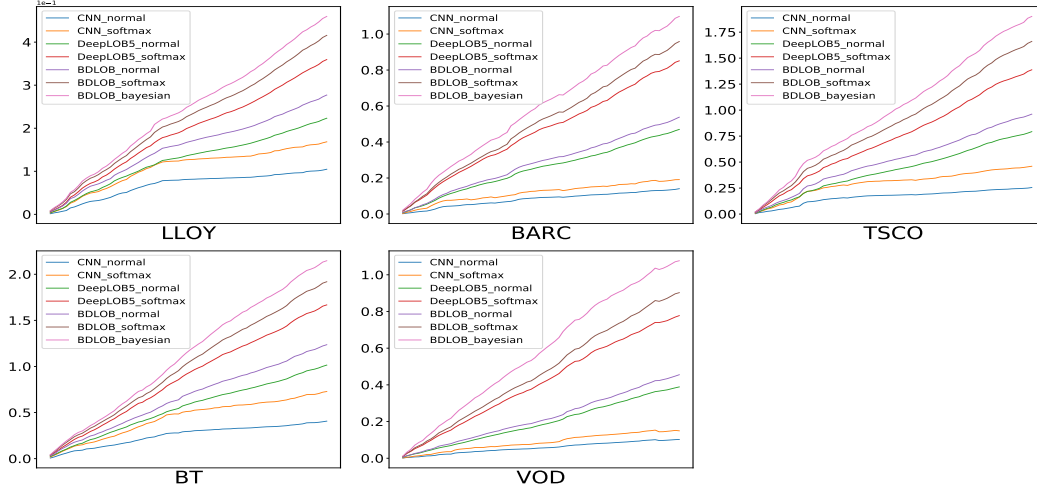


Figure 3: Cumulative profits for test periods for different models. Profits are in GBX(=GBP/100).

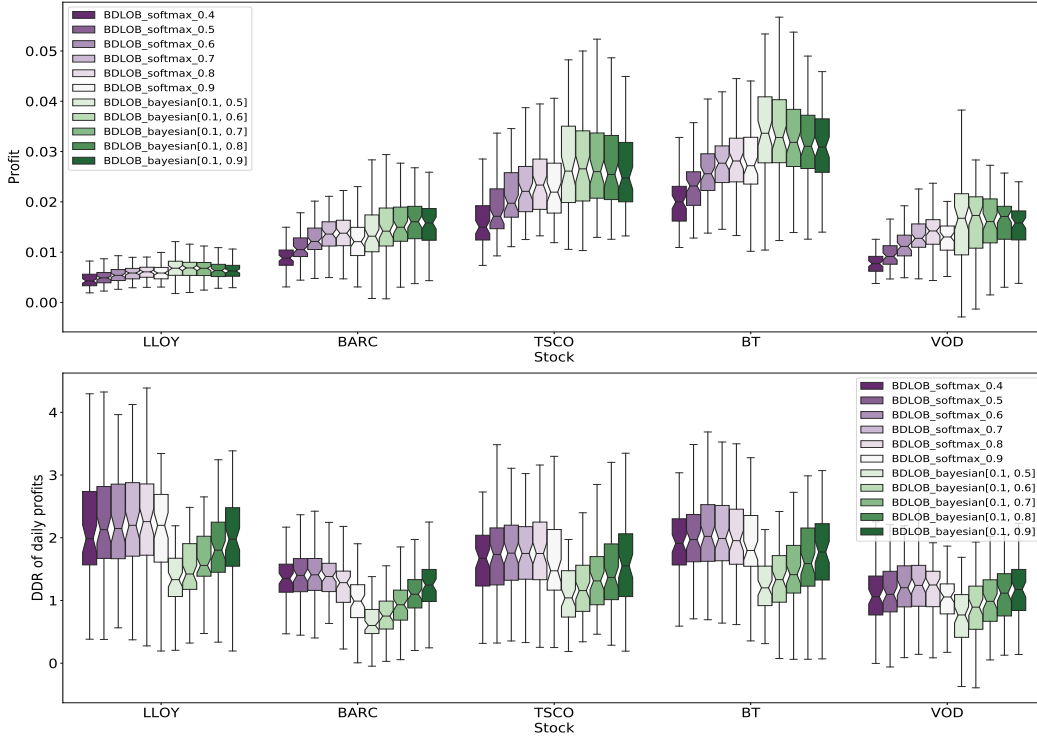


Figure 4: Boxplots for normalized profits and DDR for different trading parameters. We denote a softmax strategy with rate  $\alpha$  as  $\text{BDLOB\_softmax}_\alpha$  and Bayesian strategy as  $\text{BDLOB\_bayesian}[\beta_1, \beta_2]$ . We set  $\alpha = 0.7$  for all Bayesian trading strategies. Profits are in GBX(=GBP/100).

## 4 Conclusion

We show that Dropout inference improves performance for predicting stock price movements and demonstrate that uncertainty information can be utilised to design better trading strategies. A promising direction for future work is to explore recent Bayesian techniques such as [5] and [8] and consider trading strategies that monitor uncertainty information during holding periods.

## References

- [1] Christopher M Bishop. *Neural networks for pattern recognition*. Oxford University Press, 1995.
- [2] Yarin Gal. *Uncertainty in deep learning*. PhD thesis, University of Cambridge, 2016.
- [3] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059, 2016.
- [4] Yarin Gal, Jiri Hron, and Alex Kendall. Concrete dropout. In *Advances in Neural Information Processing Systems*, pages 3581–3590, 2017.
- [5] Mohammad Emtiyaz Khan, Didrik Nielsen, Voot Tangkaratt, Wu Lin, Yarin Gal, and Akash Srivastava. Fast and Scalable Bayesian Deep Learning by Weight-Perturbation in Adam. *arXiv preprint: 1806.04854*, 2018.
- [6] John Moody and Matthew Saffell. Learning to trade via direct reinforcement. *IEEE Transactions on Neural Networks*, 12(4):875–889, 2001.
- [7] Christine A Parlour and Duane J Seppi. Limit order markets: A survey. *Handbook of financial intermediation and banking*, 5:63–95, 2008.
- [8] Arnold Salas, Stefan Zohren, and Stephen Roberts. Practical Bayesian Learning of Neural Networks via Adaptive Subgradient Methods. *arXiv preprint:1811.03679*, 2018.
- [9] William F Sharpe. The Sharpe Ratio. *Journal of Portfolio Management*, 21(1):49–58, 1994.
- [10] Justin Sirignano and Rama Cont. Universal features of price formation in financial markets: perspectives from Deep Learning. *arXiv preprint: 1803.06917*, 2018.
- [11] Avraam Tsantekidis, Nikolaos Passalis, Anastasios Tefas, Juho Kannianen, Moncef Gabbouj, and Alexandros Iosifidis. Forecasting stock prices from the limit order book using convolutional neural networks. In *2017 IEEE 19th Conference on Business Informatics (CBI)*, volume 1, pages 7–12. IEEE, 2017.
- [12] Dimitris G. Tzikas, Aristidis C. Likas, and Nikolaos P. Galatsanos. The variational approximation for Bayesian inference. *IEEE Signal Processing Magazine*, 25(31), 2008.
- [13] Haibing Wu and Xiaodong Gu. Towards dropout training for convolutional neural networks. *Neural Networks*, 71:1–10, 2015.
- [14] Zihao Zhang, Stefan Zohren, and Stephen Roberts. DeepLOB: Deep convolutional neural networks for limit order books. *arXiv preprint:1808.03668*, 2018.
- [15] Eric Zivot and Jiahui Wang. Vector autoregressive models for multivariate time series. *Modeling Financial Time Series with S-PLUS*, pages 385–429, 2006.