

Deep Learning for Market by Order Data

Zihao Zhang^a, Bryan Lim^a and Stefan Zohren ^a

^aOxford-Man Institute of Quantitative Finance, University of Oxford, Oxford, UK.

ARTICLE HISTORY

Compiled February 18, 2021

ABSTRACT

Market by order (MBO) data – a detailed feed of individual trade instructions for a given stock on an exchange – is arguably one of the most granular sources of microstructure information. While limit order books (LOBs) are implicitly derived from it, MBO data is largely neglected by current academic literature which focuses primarily on LOB modelling. In this paper, we demonstrate the utility of MBO data for forecasting high-frequency price movements, providing an orthogonal source of information to LOB snapshots. We provide the first predictive analysis on MBO data by carefully introducing the data structure and presenting a specific normalisation scheme to consider level information in order books and to allow model training with multiple instruments. Through forecasting experiments using deep neural networks, we show that while MBO-driven and LOB-driven models individually provide similar performance, ensembles of the two can lead to improvements in forecasting accuracy – indicating that MBO data is additive to LOB-based features.

KEYWORDS

Market by Order Data; Limit Order Books; Deep Learning; Long Short-Term Memory; Attention.

1. Introduction

High-frequency microstructure data has received growing attention both in academia and industry with the computerisation of financial exchanges and the increase capacity of data storage. The detailed records of order flow and price dynamics provide us with a granular description of short-term supply and demand, and we can take the dynamics of order books into account during the modelling process. Propelled by the publication of the benchmark dataset (Ntakaris et al. 2018) of high-frequency limit order book (LOB) data, there has been a growing interest in research studying LOB data. Recent works by Tsantekidis et al. (2017); Sirignano and Cont (2019); Zhang, Zohren, and Roberts (2019a); Briola, Turiel, and Aste (2020) demonstrate that strong predictive performance can be obtained from modelling high-frequency LOB data and with resulting predictions finding applications in market-making and trade execution which have short holding periods.

In this work, we introduce Market by Order (MBO) data for predictive modelling with deep learning algorithms. MBO data provides full resolution of the underlying market microstructure – with both LOB data and trade sequences being derived from

it.

A LOB is a record of all outstanding limit orders (passive orders) for an instrument at a given time point and it is sorted into different levels based on submitted prices. At each price level, a LOB only shows the total available quantity. However, any given price level actually consists of many individual orders with different sizes. MBO data is essentially a message-base data feed that allows us to infer the individual queue position for each individual order by reconstructing the order book step by step. A detailed description of MBO data and how it relates to LOB data is presented in Section 3.

Despite MBO data being the original raw data source, current literature on high-frequency predictive modelling focuses predominantly on LOBs and, to the best of our knowledge, MBO data has not been used for direct predicting modelling. In this work, we propose a first deep learning model based on MBO data. In particular, we adopt a classification framework to predict stock price movements by using MBO data. In doing so, we provide a complete analysis of MBO by carefully introducing the data structure and the components of the message-base data feed. A specific data normalisation scheme is introduced to model level information contained in LOBs and to allow model training with multiple instruments. Our dataset consists of MBO data over a period of one year for five highly liquid instruments from the London Stock Exchange. Our testing set contains millions of samples to verify the robustness and generalisation of the results.

In our proposed models, we apply deep learning architectures including LSTMs (Hochreiter and Schmidhuber 1997) and Attention mechanisms (Bahdanau, Cho, and Bengio 2014) to model the dynamics of MBO data for market predictions. Our experiments show consistent and robust results from MBO data that are comparable to models that utilise derived LOB data. We observe that predictive models based on MBO data are complementary to LOB models and we propose an ensemble approach which yields superior results. As such, we observe that MBO data adds diversification to the LOB model and improves prediction performance.

The remainder of the paper is organised as follows. After a short literature review in Section 2, we proceed in Section 3 by introducing MBO data, including data preprocessing, normalisation and labelling. Section 4 presents deep learning architectures. We next describe our experiments and present the results of predicting market movements from MBO data in Section 5. We conclude our findings and discuss promising future extensions in Section 6.

2. Literature

Research on the high-frequency microstructure data remains largely focused on modelling the limit order book (LOB), where the classical works are referred to O’Hara (1995); Harris (2003) and a review is presented in Gould et al. (2013). However, there is limited work on MBO data in the current literature. NASDAQ (OUCH 2020) and CME Group (CME 2020) provide a preliminary description on MBO data for introducing their exchange match engines, and the works of Byrd, Hybinette, and Balch (2019); Belcak, Calliess, and Zohren (2020) use MBO data for market simulation to model trading scenarios or to study latency effects. To the best of our knowledge, this paper is the first to use MBO data to predict market movements, filling in this literature gap by using deep learning models.

Deep Learning (Goodfellow et al. 2016) algorithms have been heavily used for

predicting high-frequency microstructure data (Tsantekidis et al. 2017; Sirignano and Cont 2019; Briola, Turiel, and Aste 2020; Wallbridge 2020). In particular, Zhang, Zohren, and Roberts (2018, 2019a,b) apply convolutional neural networks and LSTMs to model the dynamics of LOB and demonstrate accuracy improvements over linear models. Unlike traditional time-series models (Mills and Mills 1991; Hamilton 2020) or econometric models (Fama and French 1993) that assume a parametric process for the underlying time-series, deep learning methods are able to capture arbitrary nonlinear relationships without placing any specific assumptions on the input data. Our experiments also suggest that deep networks deliver better results than linear methods for modelling MBO data.

We investigate deep learning models, including LSTMs (Hochreiter and Schmidhuber 1997) and Attention (Bahdanau, Cho, and Bengio 2014), to model MBO data. Attention is used to solve the problem of diminishing performance with long input sequences by utilising information at each hidden state of a recurrent network (Bahdanau, Cho, and Bengio 2014; Dai et al. 2019), and it can be used for constructing multi-horizon forecasting models (Lim and Zohren 2020). Our experiment suggests that networks with a recurrent nature lead to good predictive results compared to the state-of-art networks trained with LOB data, suggesting the potential benefits of using MBO data as an additional data source.

3. Market by Order Data

3.1. Descriptions of Market by Order Data

MBO data is a message-base data feed that allows us to observe individual actions of market participants. Essentially, it is an order instruction that describes the action of a specific trader at a given time point. In what follows, we focus on the essential components of such messages ignoring certain auxiliary information. Table 1 shows an example of sequences of MBO data, where:

- **Time stamp** records the time point when an instruction is given;
- **ID** shows the unique ID for order identification which is anonymous to others;
- **Type** indicates the order type, here limit order (Type = 1) or market order (Type = 2);
- **Side** indicates whether an order is buy (1) or sell (2);
- **Action** represents the specific instruction where 0 means updating the price or size for the existing order, 1 means adding a new order and 2 means cancelling an existing order. If Action = 2, the entries of Side, Price and Size are N/A as the matching engine will be able to identify and cancel the existing order using the unique ID;
- **Price** shows the price level of the instruction;
- **Size** shows the size (i.e. number of stocks) of the instruction.

Table 1. An example of a sequence of market by order data.

Time stamp	ID	Type	Side	Action	Price	Size
2018-01-02 09:21:15.717500766	462805645163273214	1	N/A	2	N/A	N/A
2018-01-02 09:21:18.585446702	462805645163298476	1	1	1	68.54	8334.0
2018-01-02 09:21:20.680552032	462805645163297649	1	1	0	68.56	3227.0
2018-01-02 09:21:20.944574722	462805645163297649	1	N/A	2	N/A	N/A
2018-01-02 09:21:20.945483443	462805645163298567	1	2	1	68.59	5100.0

A LOB updates whenever there is a new message from the MBO data coming in and this process is illustrated in Figure 1, where we show how a MBO message affect a LOB. For example, if we look at the top of Figure 1, a new limit order (ID=46280) is added to the ask side of the order book with price at 70.04 and size of 7580. The order book updates its status and the new order is added to the right price level. In general, a LOB only shows the total available quantity at each price level but MBO data provides us with extra information by showing individual behaviour. Although, MBO data does not directly indicate which price level the order is added to, our normalisation scheme introduced in the next section allows us to consider this information and we not only obtain a smaller input space but also obtain relevant information comparable with LOB data.

In addition, the usage of MBO data increases transparency and improves the understanding of order book dynamics without disclosing customer identification. Although, we can access to unique order ID but this number is generally assigned sequentially by the exchange match engine (CME 2020) and a private link is provided to the customer, which keeps identification confidential. Further, unlike LOB data where we sometimes only view limited price levels, MBO data allows us to observe the entire order book with full-depth information. Such a granularity can improve traders' confidence in posting large order size as they can better evaluate the potential market impact by knowing individual queue positions.

3.2. *Data Preprocessing and Normalisation*

We focus on MBO data that represents limit orders because market orders only account for a tiny percentage of total order flow. Figure 2 illustrates the process of data preprocessing and normalisation. In particular, we process the MBO data for an unique ID as:

- **Side and Price:** Missing values correspond to updates and cancellations and we fill those with the corresponding values from the original order of that ID;
- **Size:** Missing values correspond to full cancellation and we fill those with 0 to indicate that no shares are outstanding after the action;
- **Action:** we change Action to have values -1, 0 and 1. -1 means cancelling an order, 0 means updating price or size for the existing order and 1 means adding a new order;
- **Change price and Change size:** We add these two new features to calculate the difference between entries for the price and size of a specific ID to reflect the intention of adding or decreasing positions for the given order.

Data preprocessing is applied to every unique order ID and we then normalise the data as:

- **Normalised price:** $(\text{price} - \text{mid-price}) / (\text{minimum tick size} \times 100)$.¹ This calculation transforms price to tick change, representing how many ticks the price is away from the mid-price. The deviation of minimum tick size is needed when we train models with multiple instruments as it maps price to a similar scale;
- **Normalised size:** $\text{size} / \text{mid-size}$. Mid-size is the mean between the current best ask and bid size, which is similar to mid-price.
- **Normalised change price:** $\text{change price} / \text{minimum tick size}$;

¹Mid-price is the mean between the best ask and bid price. Those references prices and their sizes which we use below can be obtained on the fly from the MBO data.

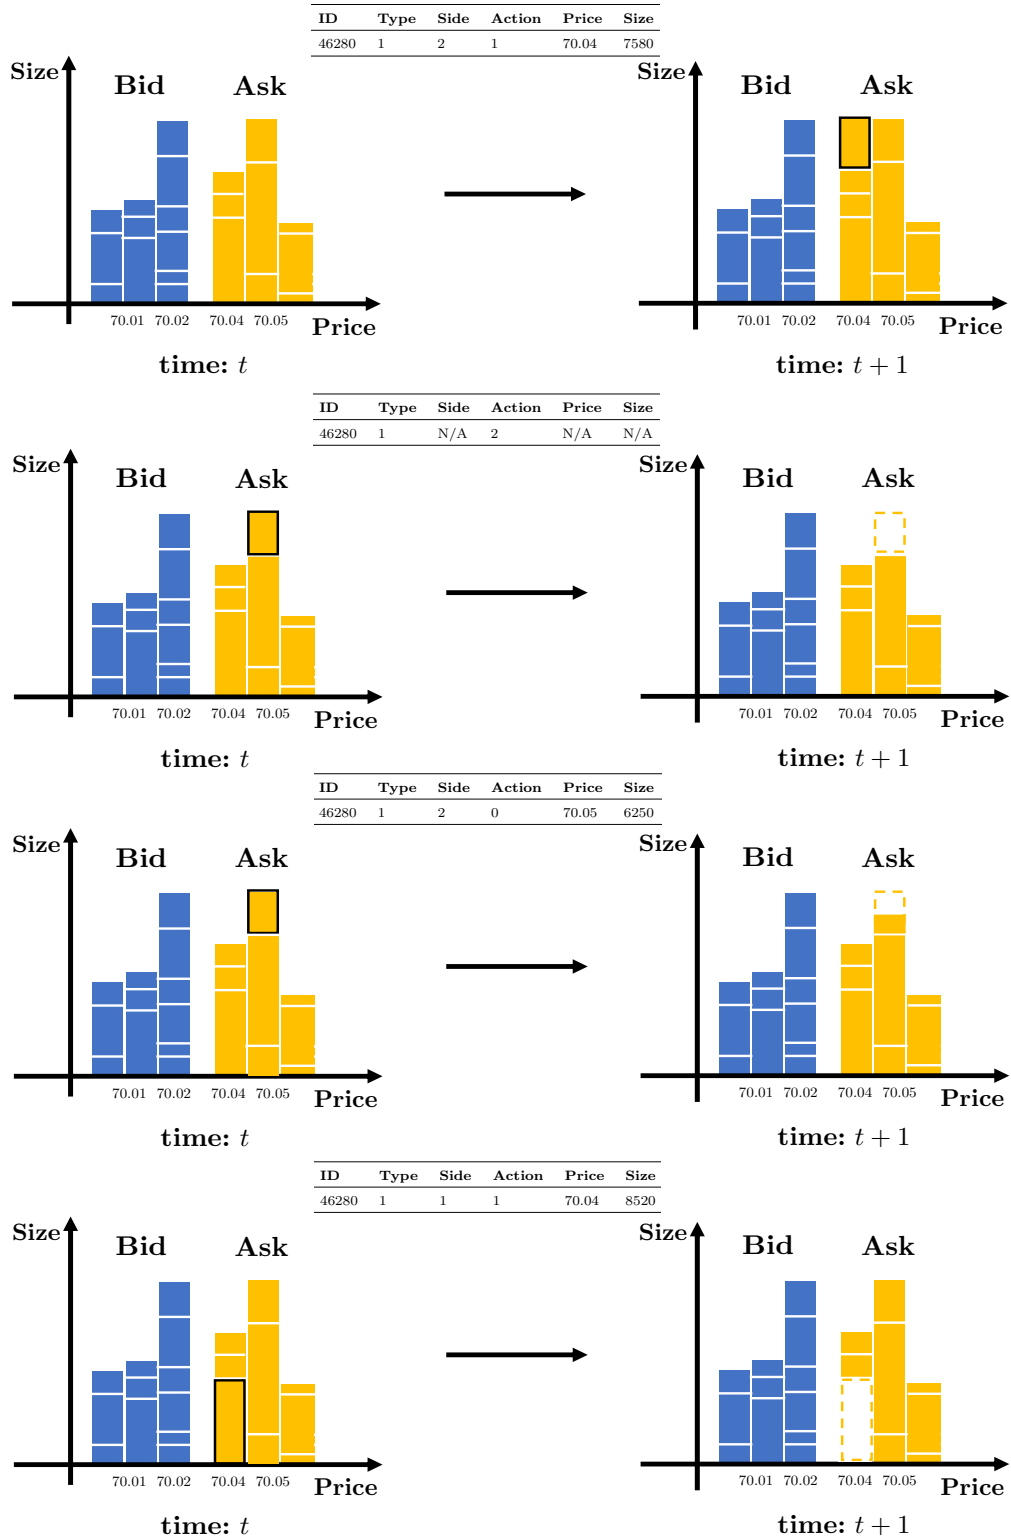


Figure 1. An illustration of how MBO data updates a LOB. **Top:** An addition of a new limit order; **Middle top:** A cancellation of an existing order; **Middle bottom:** An update for a partial cancellation; **Bottom:** A marketable buy limit order that crosses the spread.

Time stamp	ID	Side	Action	Price	Size			
2018-01-02 09:00:52.369737129	462805645163386861	2	1	68.47	2049.0			
2018-01-02 09:00:57.396438193	462805645163386861	2	0	68.46	2110.0			
2018-01-02 09:01:30.641547824	462805645163386861	2	0	68.46	2082.0			
2018-01-02 09:01:30.642333313	462805645163386861	2	0	68.45	2082.0			
2018-01-02 09:01:33.656776415	462805645163386861	N/A	2	N/A	N/A			

↓ *Data preprocessing*

Time stamp	Side	Action	Price	Size	Change price	Change size	Mid-price	Mid-size
2018-01-02 09:00:52.369737129	2	1	68.47	2049.0	0.00	2049.0	68.445	32900.0
2018-01-02 09:00:57.396438193	2	0	68.46	2110.0	-0.01	61.0	68.440	26750.0
2018-01-02 09:01:30.641547824	2	0	68.46	2082.0	0.00	-28.0	68.450	13250.0
2018-01-02 09:01:30.642333313	2	0	68.45	2082.0	-0.01	0.0	68.450	13250.0
2018-01-02 09:01:33.656776415	2	-1	68.45	0.0000	0.00	-2082.0	68.425	7750.00

↓ *Data normalisation*

Time stamp	Side	Action	Nor. price	Nor. size	Nor. change price	Nor. change size		
2018-01-02 09:00:52.369737129	2	1	2.5	0.0623	0.0	0.0623		
2018-01-02 09:00:57.396438193	2	0	2.0	0.0789	-1.0	0.0023		
2018-01-02 09:01:30.641547824	2	0	1.0	0.1571	0.0	-0.0021		
2018-01-02 09:01:30.642333313	2	0	0.0	0.1571	-1.0	0.0000		
2018-01-02 09:01:33.656776415	2	-1	2.5	0.0000	0.0	-0.2686		

Figure 2. An example of preprocessing and normalising the MBO data.

- **Normalised change size:** change size / mid-size.²
- **Side and Action:** remain unchanged.

At the end, we remove “Time stamp” and “ID”, leading to 6 features in our feature space. Note that the normalised price essentially represent the price level where the current order is in the order book, taking the level information into account.

3.3. Data Labelling

In this work, we study a classification framework where we want to predict the future market movements into three classes: the market going up, staying stationary or going down. We use mid-prices to create labels and adopt the labelling method in Zhang, Zohren, and Roberts (2019a) to classify movements. In particular, we define

$$\begin{aligned}
 l_t &= \frac{m_+(t) - m_-(t)}{m_-(t)}, \\
 m_-(t) &= \frac{1}{k} \sum_{i=0}^{k-1} p_{t-i}, \\
 m_+(t) &= \frac{1}{k} \sum_{i=1}^k p_{t+i},
 \end{aligned} \tag{1}$$

²By mid-size we denote the mean of best bid and ask size, in analogy to mid price. Those values again are obtained on the fly from the MBO data.

where k is the prediction horizon and p_t is the mid-price at time t . To decide on the label we compare l_t with a threshold (α), labelling it as up if $l_t > \alpha$, down if $l_t < -\alpha$ and stationary otherwise. The choice of α is related to the prediction horizon (k) and we set α for each instrument to obtain a balanced training set. Our choices of k and α are listed in Section 5 and we show that the dataset are balanced under our choice.

Note that Equation (1) introduces a smooth labelling that leads to consistent labels that are better for designing trading signals and the work of Zhang, Zohren, and Roberts (2019a) includes a more detailed discussion demonstrating the effects of different labelling methods. Interested readers are referred to their work for a detailed explanation.

4. Methodology

In this section, we introduce the different deep learning algorithms studied in our work. For a single input of any time-series, we write $\mathbf{x}_{1:T}$, where \mathbf{x}_t represents the features at time t and T is the length of the sequence which will later correspond to the length of the lookback of the input.

4.1. Multilayer perceptrons (MLPs)

MLPs are canonical neural network models where a typical network is organised into a series of layers in a chain structure, with each layer being a function of the layer that precedes it. We can define the hidden layer of a MLP as:

$$\mathbf{h}^{(l)} = g^{(l)}(\mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}), \quad (2)$$

where $\mathbf{h}^{(l)} \in \mathbb{R}^{N_l}$ represents the l -th hidden layer with weights $\mathbf{W}^{(l)} \in \mathbb{R}^{N_l \times N_{l-1}}$ and biases $\mathbf{b}^{(l)} \in \mathbb{R}^{N_l}$. Here $g^{(l)}(\cdot)$ is the activation function that allows networks to model nonlinearities. The final output is a function of the last hidden layer and we compute objective functions to minimise errors between target outputs and estimates.

However, for MLPs, we first need to flatten $\mathbf{x}_{1:T}$ and feed it to subsequent hidden layers. Doing this breaks the time dependences and treats features at different time stamps independently. We generally observe inferior results using MLPs and find that recurrent neural networks (RNNs) often deliver better performance. This is because a RNN acts as a memory buffer by summarising past information and recursively updating the hidden state with new observations at each time step of the input (Zhang, Zohren, and Roberts 2020a).

4.2. Long Short-Term Memory (LSTMs)

Standard RNNs suffer from vanishing or exploding gradient problems (Bengio, Simard, and Frasconi 1994) and Long Short-Term Memory networks (LSTMs) are proposed to solve this problem. This is done by operating a gating mechanism that efficiently controls the propagation of past information (Hochreiter and Schmidhuber 1997). A LSTM updates its hidden state recursively and has a cell state \mathbf{c}_t coupled with a series

of gates at each hidden state. In mathematical terms, we can write

$$\begin{aligned}
\text{Input gate: } \mathbf{i}_t &= \sigma(\mathbf{W}_{i,h}\mathbf{h}_{t-1} + \mathbf{W}_{i,x}\mathbf{x}_t + \mathbf{b}_i), \\
&\text{with } \mathbf{W}_{i,h} \in \mathbb{R}^{N_h \times N_h}, \mathbf{W}_{i,x} \in \mathbb{R}^{N_h \times N_x} \text{ and } \mathbf{b}_i \in \mathbb{R}^{N_h}, \\
\text{Output gate: } \mathbf{o}_t &= \sigma(\mathbf{W}_{o,h}\mathbf{h}_{t-1} + \mathbf{W}_{o,x}\mathbf{x}_t + \mathbf{b}_o), \\
&\text{with } \mathbf{W}_{o,h} \in \mathbb{R}^{N_h \times N_h}, \mathbf{W}_{o,x} \in \mathbb{R}^{N_h \times N_x} \text{ and } \mathbf{b}_o \in \mathbb{R}^{N_h}, \\
\text{Forget gate: } \mathbf{f}_t &= \sigma(\mathbf{W}_{f,h}\mathbf{h}_{t-1} + \mathbf{W}_{f,x}\mathbf{x}_t + \mathbf{b}_f), \\
&\text{with } \mathbf{W}_{f,h} \in \mathbb{R}^{N_h \times N_h}, \mathbf{W}_{f,x} \in \mathbb{R}^{N_h \times N_x} \text{ and } \mathbf{b}_f \in \mathbb{R}^{N_h},
\end{aligned} \tag{3}$$

where \mathbf{h}_{t-1} is the hidden state of a LSTM at time $t-1$ and $\sigma(\cdot)$ represents the sigmoid activation function. We use \mathbf{W} and \mathbf{b} to represent weights and biases at different gate operations. Subsequently, the current cell state and hidden state can be written as:

$$\begin{aligned}
\text{Cell state: } \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_{c,h}\mathbf{h}_{t-1} + \mathbf{W}_{c,x}\mathbf{x}_t + \mathbf{b}_c), \\
\text{Hidden state: } \mathbf{h}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t),
\end{aligned} \tag{4}$$

where $\mathbf{W}_{c,h} \in \mathbb{R}^{N_h \times N_h}$, $\mathbf{W}_{c,x} \in \mathbb{R}^{N_h \times N_x}$, $\mathbf{b}_c \in \mathbb{R}^{N_h}$, \odot is the element-wise product and $\tanh(\cdot)$ is the hyperbolic tangent activation function. The hidden state \mathbf{h}_t summarises the information from past states and current observations, and the gating mechanism efficiently addresses the vanishing gradient problem.

4.3. Attention Mechanism

The Attention Mechanism (Bahdanau, Cho, and Bengio 2014) is heavily used in machine translation and is proposed to solve the problem of diminishing performance for long input sequences. On the one hand, a LSTM calculates the final output as a function of only the last hidden state. An attention model, on the other hand, with an additional component called context vector, assigns trainable weights to all the hidden states of an input. We can write an attention mechanism for modelling many-to-one problem as:

$$\mathbf{h}_t = \mathbf{f}_t(\mathbf{h}_{t-1}, \mathbf{x}_t), \tag{5}$$

where \mathbf{h}_t can be the hidden state from a LSTM at time t for an input $\mathbf{x}_{1:T}$, and we define the context vector \mathbf{c}_T as:

$$\begin{aligned}
\text{Context vector: } \mathbf{c}_T &= \sum_{t=1}^T \alpha(t, T) \mathbf{h}_t, \\
\text{Attention weights: } \alpha(t, T) &= \frac{\exp(e(t, T))}{\sum_{t=1}^T \exp(e(t, T))}, \\
\text{Score: } e(t, T) &= \mathbf{v}^T \tanh(\mathbf{W}_h \mathbf{h}_t),
\end{aligned} \tag{6}$$

where $\mathbf{v} \in \mathbb{R}^{N_h}$ and $\mathbf{W}_h \in \mathbb{R}^{N_h \times N_h}$ are the trainable weights. We can then obtain the attention vector:

$$\mathbf{a}_T = \mathbf{f}(\mathbf{c}_T, \mathbf{h}_T) = \tanh(\mathbf{W}_c[\mathbf{c}_T; \mathbf{h}_T]), \tag{7}$$

Table 2. Label parameters (α) for different prediction horizons and instruments (units in 10^{-4}).

	LLOY	BARC	TSCO	BT	VOD
$k = 20$	0.25	0.35	0.10	0.40	0.22
$k = 50$	0.50	0.65	0.70	0.70	0.45
$k = 100$	0.75	0.95	1.20	1.00	0.70

where the final output is a function of \mathbf{c}_T , taking information at every hidden state into account.

5. Experiments

5.1. Descriptions of Datasets

Our datasets consist of MBO data for five highly liquid stocks, Lloyds (LLOY), Barclays (BARC), Tesco (TSCO), BT and Vodafone (VOD), for the entire year of 2018 from the London Stock Exchange. From the MBO data one can derive LOB data which we use for our benchmarks and for references prices. Our LOB dataset contains ask and bid information for an order book up to ten levels. For our modelling we remove messages outside ten levels from the MBO data to align the timestamps of two datasets allowing for fair comparisons in the performance analysis. Afterwards, we train two sets of models by separately using the MBO and LOB data with the same targets. A direct comparison can be then made to compare predictive performance using the MBO and LOB data respectively.

For each trading day, we take the data between 08:30:00 and 16:00:00, restricting ourselves to liquid continuous trading hours, excluding any auctions. Overall, we have more than 169 million samples in our dataset and we take the first 6 months as training data, the next 3 months as validation data and the last 3 months as testing data. In the context of high-frequency microstructure data, we have more than 46 million observations in our testing set, providing sufficient scope for verifying the robustness and generalisability of model performance.

We test our models at three prediction horizons ($k = 20, 50, 100$) and list the choices of label parameter (α) in Table 2. We choose α for each instrument to have a balanced training set and the proportion of different classes is presented in Figure B1 in Appendix B. Overall, the labels are roughly balanced for the testing set as well (noting that those were fixed on the training set). In terms of the lookback window (T) of the input, we take the 50 most recent updates of MBO data to form a single input and feed it to our model.

5.2. Training Procedure

For the MBO data, we study the deep learning models (MBO-MLP, MBO-LSTM and MBO-Attention) introduced in Section 4 along with a simple linear model (MBO-LM). We list the values of hyperparameters for different algorithms in Table 3, and the Gradient descent with the Adam optimiser (Kingma and Ba 2015) is used for training all models. The complete search space of hyperparameters is included in Appendix A and we use a grid-search method to select best hyperparameters.

For the LOB data, we include the 10 levels of a limit order book and past 50 observations as a single input. We follow the normalisation scheme in Zhang, Zohren,




Figure 3. Pearson correlation between different predictive signals for different prediction horizons (k). **Top:** $k = 20$; **Middle:** $k = 50$; **Bottom:** $k = 100$.

and Roberts (2019a) and both the MBO and LOB datasets share the same predictive targets, allowing a direct comparison between different models. We choose state-of-art network architectures as comparison models, including the LOB-LSTM (Sirignano and Cont 2019), LOB-CNN (Tsantekidis et al. 2017) and LOB-DeepLOB (Zhang, Zohren, and Roberts 2019a). The details of the network architecture and choices of hyperparameters can be found in their papers. We also include a linear model (LOB-LM) and a multilayer perceptron (LOB-MLP) as benchmark models.

We use categorical cross-entropy loss as our objective function and the learning is stopped when the validation loss does not decrease for more than 10 epochs. In general, it takes about 30 epochs to finish model training. TensorFlow and Keras (Girija 2016) are used to build all models and four NVIDIA GeForce RTX 2080 are used in our experiment.

Table 3. Choices of hyperparameters.

	<u># of layers</u>	<u># of units</u>	<u>Learning rate</u>	<u>Batch size</u>	<u># of parameters</u>
LM	-	-	0.0001	128	903
MLP	1	64	0.0001	128	19459
LSTM	2	64	0.0001	128	51907
Attention	2	64	0.0001	128	72067

5.3. Experimental Results

Table 4 summarises the results for all models studied (different rows) and one suitable for each different prediction horizons. We use four evaluation metrics (different columns) to make comparisons: Accuracy, Precision, Recall and F1-score. Kolmogorov-Smirnov (Massey Jr 1951) tests are used to check the statistical significance of results and all differences in evaluation metrics are significant.

We observe that the models trained with LOB data are comparable, but slightly outperform the ones using MBO data. While a priori, MBO data contains more information (contents of level and trades), it is harder to model the raw messages rather than LOB snapshots when can be seen as derived or handcrafted features from the MBO data. What is encouraging is that we are able to obtain comparable performance by modelling the raw messages directly. Furthermore, if we look at the Pearson correlation between predictive signals in Figure 3, we can see that predictive signals from the MBO data are less correlated with LOB’s signals. This means that we were indeed able to

Table 4. Experimental results for different prediction horizons (k).

Model	Accuracy %	Precision %	Recall %	F1 %
Prediction horizon k = 20				
MBO-LM	41.81	41.16	41.81	34.97
MBO-MLP	47.12	46.17	47.12	46.46
MBO-LSTM	61.94	61.60	61.94	61.75
MBO-Attention	61.19	62.83	61.19	61.73
LOB-LM	45.71	43.44	45.71	42.38
LOB-MLP	50.06	50.04	50.06	46.89
LOB-LSTM	66.09	67.53	66.09	66.68
LOB-CNN	63.39	67.31	63.39	64.64
LOB-DeepLOB	68.73	68.16	68.73	68.40
Ensemble-MBO	62.35	62.92	62.35	62.56
Ensemble-LOB	67.97	68.74	67.97	68.31
Ensemble-MBO-LOB	68.95	69.10	68.95	69.02
Prediction horizon k = 50				
MBO-LM	41.88	38.42	41.88	36.57
MBO-MLP	46.39	43.07	46.39	42.33
MBO-LSTM	58.84	59.65	58.84	59.18
MBO-Attention	59.31	56.10	59.31	56.88
LOB-LM	46.97	44.34	46.97	41.13
LOB-MLP	50.56	48.46	50.56	47.25
LOB-LSTM	64.49	64.88	64.49	64.65
LOB-CNN	64.77	62.55	64.77	63.26
LOB-DeepLOB	66.12	64.37	65.38	64.79
Ensemble-MBO	60.03	58.45	60.03	59.05
Ensemble-LOB	65.95	64.72	65.95	65.23
Ensemble-MBO-LOB	66.17	64.78	66.17	65.34
Prediction horizon k = 100				
MBO-LM	41.27	34.23	41.27	35.05
MBO-MLP	44.19	42.70	44.19	40.29
MBO-LSTM	57.96	54.10	57.96	54.79
MBO-Attention	56.36	53.66	56.36	53.75
LOB-LM	46.19	43.29	46.19	41.80
LOB-MLP	48.36	47.39	48.36	43.66
LOB-LSTM	61.27	58.47	62.82	57.97
LOB-CNN	61.78	56.91	61.78	55.40
LOB-DeepLOB	62.82	60.94	61.27	61.10
Ensemble-MBO	56.62	54.85	56.62	55.48
Ensemble-LOB	63.25	59.41	63.25	60.56
Ensemble-MBO-LOB	63.75	60.01	63.75	61.82

extract different information from the MBO data. It also suggests that a combination of two signals, from MBO and LOB data, can benefit from diversification that reduces signal variance given the low correlation.

To verify this statement, we include three ensemble models in our experiment, where Ensemble-MBO is obtained from MBO-LSTM and MBO-Attention; Ensemble-LOB is from LOB-LSTM, LOB-CNN, and LOB-DeepLOB; and Ensemble-MBO-LOB combines Ensemble-MBO and Ensemble-LOB. A equal weighting scheme is used to construct ensemble models and we can observe that ensemble approaches, in general, improve predictive performance. In particular, Ensemble-MBO-LOB delivers the best performance, indicating the potential benefits of combining the MBO and LOB data.

Since this work aims to study MBO data, we focus on analysing results from the models trained using the MBO data. We can see that the deep learning models outperform the simple linear model, suggesting the existence of nonlinear features in




Figure 4. Normalised confusion matrices for different prediction horizons (k). **Top:** $k = 20$; **Middle:** $k = 50$; **Bottom:** $k = 100$. From the left to right, we display MBO-LM, MBO-MLP, MBO-LSTM and MBO-Attention.

financial time-series, and networks are capable of extracting such features from the raw messages in MBO data. We observe that MBO-MLP delivers inferior results compared to other networks. This is most likely due to the structure of the MLP which has full connectivity between input and hidden units – leading MLPs to often underperform when compared to other networks in financial applications with low signal-to-noise ratio. MBO-LSTM and MBO-Attention all have a recurrent structure with parameter sharing that enables hidden states to summarise past information and update status with current observations. Such a process filters unnecessary input components and naturally models the propagation of order flow. This observation has also been reported by Lim, Zohren, and Roberts (2019); Zhang, Zohren, and Roberts (2020b,a) where they find that networks with a recurrent nature deliver better results than MLPs when modelling financial time-series.

Figure 4 shows the normalised confusion matrices which helps to understand how models perform at predicting each label class. We calculate the accuracy score for every instrument and for each testing day to understand the consistency of our results. This is summarised in the whisker plots in Figure 5. Each point in the whisker plot represents the accuracy score for one testing day, and we make the box represents the median and interquartile range from these scores. We can see that the MBO-LM and MBO-MLP have large interquartile ranges, suggesting high variances in results, while MBO-LSTM and MBO-Attention show consistent and robust results across the entire testing period. These whisker plots allow us to understand the model performance on a daily basis to ensure the generalisability of our methods. In particular, we see that performance is consistent across the entire testing period and not focused on a few days which could be due to noise.

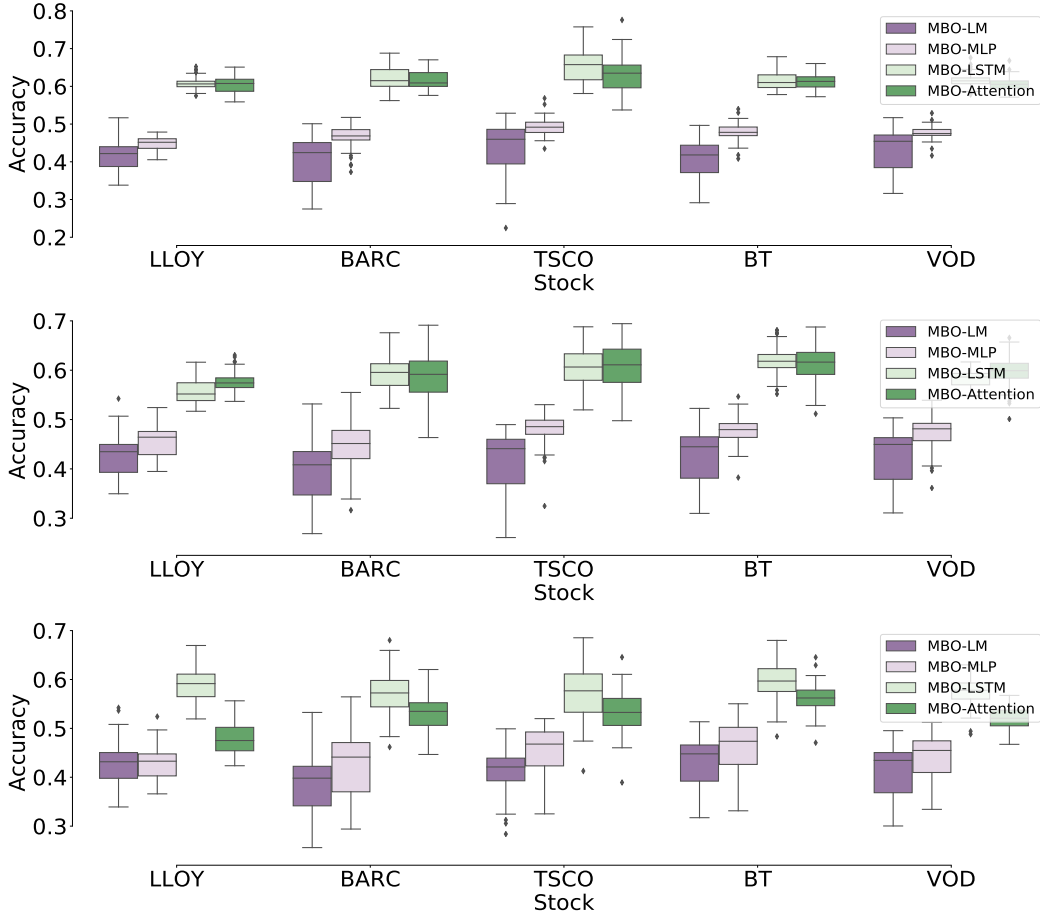


Figure 5. Whisker plots of daily accuracy for different prediction horizons (k). **Top:** $k = 20$; **Middle:** $k = 50$; **Bottom:** $k = 100$.

6. Conclusion

In this work we introduce deep learning models for Market by Order (MBO) data. To the best of our knowledge this is the first study of predictive modelling of MBO data using data-driven techniques in the academic literature. Current academic research in this direction is primarily focused on LOB data and we hope that this work helps to popularise the usage of MBO which we see as the next frontier in microstructure modelling in financial data science.

We carefully introduce the structure of MBO data and demonstrate a specific normalisation scheme that allows model training with multiple instruments using deep learning. We consider a wide range of deep learning architectures including MLP, LSTM and attention layers. Our dataset consists of millions of sample for highly liquid instruments from the London Stock Exchange, ensuring the consistency and generalisability of our methods.

We compare models trained using MBO and LOB data respectively. We show that we can obtain similar, but slightly inferior, performance by modelling raw MBO messages, when compared to modelling LOB data. While MBO data a priori contains more information, it is harder to model the raw messages rather than LOBs, which can be seen as derived features of the data. Importantly, we show that our models can extract

additional information from the MBO data which is not captured by models trained on LOB data. This means that they can add additional value as we demonstrate in an ensemble approach that combines signals from the MBO and LOB data and delivers the best performance.

In subsequent continuation of this work, we can apply MBO data to various financial applications including market-making or trade execution. Further, the work of Briola et al. (2021) applies Reinforcement Learning (RL) algorithms to high-frequency trading, and it would be interesting to test the effectiveness of using MBO data within a RL framework.

Acknowledgement(s)

The authors would like to thank members of Machine Learning Research Group at the University of Oxford for their useful comments. We are most grateful to the Oxford-Man Institute of Quantitative Finance for computing support and data access.

References

- Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio. 2014. “Neural machine translation by jointly learning to align and translate.” *arXiv:1409.0473* .
- Belcak, Peter, Jan-Peter Calliess, and Stefan Zohren. 2020. “Fast Agent-Based Simulation Framework of Limit Order Books with Applications to Pro-Rata Markets and the Study of Latency Effects.” *arXiv:2008.07871* .
- Bengio, Yoshua, Patrice Simard, and Paolo Frasconi. 1994. “Learning long-term dependencies with gradient descent is difficult.” *IEEE transactions on neural networks* 5 (2): 157–166.
- Briola, Antonio, Jeremy Turiel, and Tomaso Aste. 2020. “Deep Learning Modeling of the Limit Order Book: A Comparative Perspective.” *Available at SSRN 3714230* .
- Briola, Antonio, Jeremy Turiel, Riccardo Marcaccioli, and Tomaso Aste. 2021. “Deep Reinforcement Learning for Active High Frequency Trading.” *arXiv:2101.07107* .
- Byrd, David, Maria Hybinette, and Tucker Hybinette Balch. 2019. “Abides: Towards high-fidelity market simulation for AI research.” *arXiv:1904.12066* .
- CME. 2020. “Market by Order (MBO).” <https://www.cmegroup.com/education/market-by-order-mbo.html>.
- Dai, Zihang, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. 2019. “Transformer-xl: Attentive language models beyond a fixed-length context.” *arXiv:1901.02860* .
- Fama, Eugene F, and Kenneth R French. 1993. “Common risk factors in the returns on stocks and bonds.” *Journal of financial economics* 33 (1): 3–56.
- Girija, Sanjay Surendranath. 2016. “TensorFlow: Large-scale machine learning on heterogeneous distributed systems.” *Software available from tensorflow.org* 39 (9).
- Goodfellow, Ian, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. 2016. *Deep learning*. Vol. 1. MIT press Cambridge.
- Gould, Martin D, Mason A Porter, Stacy Williams, Mark McDonald, Daniel J Fenn, and Sam D Howison. 2013. “Limit order books.” *Quantitative Finance* 13 (11): 1709–1742.
- Hamilton, James Douglas. 2020. *Time series analysis*. Princeton university press.
- Harris, Larry. 2003. *Trading and exchanges: Market microstructure for practitioners*. OUP USA.
- Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. “Long short-term memory.” *Neural computation* 9 (8): 1735–1780.
- Kingma, Diederik P, and Jimmy Ba. 2015. “Adam: A method for stochastic optimization.” *Proceedings of the International Conference on Learning Representations* .

- Lim, Bryan, and Stefan Zohren. 2020. “Time series forecasting with deep learning: A survey.” *arXiv:2004.13408* .
- Lim, Bryan, Stefan Zohren, and Stephen Roberts. 2019. “Enhancing Time-Series Momentum Strategies Using Deep Neural Networks.” *The Journal of Financial Data Science* <https://jfds.pm-research.com/content/early/2019/09/09/jfds.2019.1.015>.
- Massey Jr, Frank J. 1951. “The Kolmogorov-Smirnov test for goodness of fit.” *Journal of the American statistical Association* 46 (253): 68–78.
- Mills, Terence C, and Terence C Mills. 1991. *Time series techniques for economists*. Cambridge University Press.
- Ntakaris, Adamantios, Martin Magris, Juho Kannianen, Moncef Gabbouj, and Alexandros Iosifidis. 2018. “Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods.” *Journal of Forecasting* 37 (8): 852–866.
- O’Hara, Maureen. 1995. *Market microstructure theory*. Vol. 108. Blackwell Publishers Cambridge, MA.
- OUCH. 2020. “NASDAQ, OUCH.” <http://www.nasdaqtrader.com/content/technicalsupport/specifications/TradingProducts/OUCH4.2.pdf>.
- Sirignano, Justin, and Rama Cont. 2019. “Universal features of price formation in financial markets: Perspectives from deep learning.” *Quantitative Finance* 19 (9): 1449–1459.
- Tsantekidis, Avraam, Nikolaos Passalis, Anastasios Tefas, Juho Kannianen, Moncef Gabbouj, and Alexandros Iosifidis. 2017. “Forecasting stock prices from the limit order book using convolutional neural networks.” In *2017 IEEE 19th Conference on Business Informatics (CBI)*, Vol. 1, 7–12. IEEE.
- Wallbridge, James. 2020. “Transformers for limit order books.” *arXiv:2003.00130* .
- Zhang, Zihao, Stefan Zohren, and Stephen Roberts. 2018. “BDLOB: Bayesian deep convolutional neural networks for limit order books.” *Third workshop on Bayesian Deep Learning (NeurIPS 2018), Montr al, Canada* .
- Zhang, Zihao, Stefan Zohren, and Stephen Roberts. 2019a. “DeepLOB: Deep convolutional neural networks for limit order books.” *IEEE Transactions on Signal Processing* 67 (11): 3001–3012.
- Zhang, Zihao, Stefan Zohren, and Stephen Roberts. 2019b. “Extending Deep Learning Models for Limit Order Books to Quantile Regression.” *Proceedings of Time Series Workshop of the 36 th International Conference on Machine Learning, Long Beach, California, PMLR 97, 2019* .
- Zhang, Zihao, Stefan Zohren, and Stephen Roberts. 2020a. “Deep learning for portfolio optimization.” *The Journal of Financial Data Science* 2 (4): 8–20.
- Zhang, Zihao, Stefan Zohren, and Stephen Roberts. 2020b. “Deep reinforcement learning for trading.” *The Journal of Financial Data Science* 2 (2): 25–40.

Appendix A. Complete Search Space for Hyperparameters

Linear Model (MBO-LM):

- Learning rate: [0.0001, 0.0005, 0.001]
- Minibatch Size: [64, 128, 256]

Multi-layer Perceptron (MBO-MLP):

- Number of hidden layer: [1, 2, 3]
- Number of neurons at each layer: [32, 64, 128]
- Learning rate: [0.0001, 0.0005, 0.001]
- Minibatch Size: [64, 128, 256]

Long Short-Term Memory (MBO-LSTM):

- Number of hidden layer: [1, 2, 3]
- Number of neurons at each layer: [32, 64, 128]
- Learning rate: [0.0001, 0.0005, 0.001]
- Minibatch Size: [64, 128, 256]

MBO-Attention:

- Number of hidden layer: [1, 2, 3]
- Number of neurons at each layer: [32, 64, 128]
- Learning rate: [0.0001, 0.0005, 0.001]
- Minibatch Size: [64, 128, 256]

Appendix B. Additional Results




Figure B1. Label class balancing for train, validation and test sets for different prediction horizons (k). **Top:** $k = 20$; **Middle:** $k = 50$; **Bottom:** $k = 100$.