A Synthetic Limit Order Book Dataset for Benchmarking Forecasting Algorithms under Distributional Shift

Anonymous Author(s)
Affiliation
Address
email

Abstract

In electronic trading markets, limit order books (LOBs) provide information about pending buy/sell orders at various price levels for a given security. Recently, there has been a growing interest in using LOB data for resolving downstream machine learning tasks (e.g., forecasting). However, dealing with out-of-distribution (OOD) LOB data is challenging since distributional shifts are unlabeled in current publicly available LOB datasets. Therefore, it is critical to build a synthetic LOB dataset with labeled OOD samples serving as a testbed for developing models that generalize well to unseen scenarios. In this work, we utilize a multi-agent market simulator to build a synthetic LOB dataset with and without market stress scenarios, which allows for the design of controlled distributional shift benchmarking. Using the proposed synthetic dataset, we provide a holistic analysis on the forecasting performance of three different state-of-the-art forecasting methods. Our results reflect the need for increased researcher efforts to develop algorithms with robustness to distributional shifts in high-frequency time series data.

1 Introduction

Increasingly large market volumes are traded today electronically across multiple asset classes. Electronic trading is typically facilitated by limit order books (LOBs) - which present the list of orders that is maintained by a trading venue to indicate "buy" and "sell" interest of market participants for a given security. Specifically, LOBs dynamically record volume and price information about the buy and sell orders that are being placed in the market at different times [3]. LOBs present highly complex and noisy environments which enable multiple market participants to trade. In addition, trading is usually being performed using automated trading algorithms - part of which assumes ability to forecast LOB prices and volumes [1, 28].

Distributional shift refers to the fundamental issue that the underlying distributions of training and testing datasets are different from each other, which often causes machine learning systems to fail in handling out-of-distribution (OOD) inputs [34, 13]. For example, financial downstream tasks involving time series data can suffer from OOD inputs over time as a result of exogenous factors (e.g., macro shocks, earnings announcements, global pandemic, etc.). Moreover, such sudden distributional shifts can affect the LOB trading algorithms, specifically their price and volume prediction components.

Different from the computer vision (CV) and natural language processing (NLP) downstream tasks, where one can easily confirm whether the data suffers from distributional shift [21, 15, 26, 19], correlation structure of multivariate time series data is inherently different than that of images and

Submitted to the Workshop on Distribution Shifts, 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

text. Therefore, previous distributional shift algorithms applied to the CV and NLP domains are not necessarily suitable for the time series domain. Furthermore, publicly available datasets typically used for benchmarking forecasting algorithms are difficult to be used directly to verify whether a particular model has the ability to handle OOD. This are two major reasons: (1) free access to publicly available LOB data is limited, and (2) distributional shifts (e.g., market shocks) are not labeled in real data, making the assessment of a model's ability to account for distributional shift difficult.

Typically, test performance on OOD inputs is worse than that of test inputs following the same distribution, i.e., independent and identically distribution (IID), as the training data. The over-reliance on IID inputs makes machine learning systems challenging to deploy in real-world settings, where distributional shifts are common [36, 24]. However, to the best of our knowledge, distributional shifts on financial market data have not been thoroughly explored, despite the fact that they are common in real-world financial market. In this work, we aim to consider a domain adaption task on LOBs, where training and testing data are generated from related but different domains via a multi-agent market simulator. To this end, we propose a synthetic dataset which can be used to benchmark forecasting algorithms on variety of different downstream tasks, e.g., mid-price trend prediction under distributional shifts. Our proposed dataset can test the OOD adaption capabilities of state-of-the-art forecasting models, which should be excepted to generalize to unseen samples in spite of the occurrent of a distributional shift. Specifically, we utilize the multi-agent limit order book market simulator called ABIDES [4] to build a synthetic LOB dataset, where randomly introduced shock are utilized to construct distributional market shifts (i.e., OOD inputs)[33]. Since the LOB dataset is generated in a controlled manner, each snapshot of LOB is labeled, allowing for straightforward benchmarking on IID vs. OOD inputs. Furthermore, our proposed configuration of the simulator allows for the parametric characterization of multiple types of shocks (e.g., shocks of different magnitude), which can be used to understand the robustness of each forecasting algorithm as a function of shock parameters.

We summarize the main **contributions** as follows:

- To facilitate research in both machine learning and finance communities, we propose a dataset that allows to model distributional shifts due to market shocks in LOBs data using a bottom up multi-agent approach. This approach allows for easy adjustment of the parameters of the market agents to model wide spectrum of counterfactual shock scenarios.
- To apply a rigorous comparison on the proposed synthetic dataset, we choose three categories of time series forecasting algorithms including: (1) AdaRNN-based method focusing on distributional shift problem of time series data; (2) transformer-based method dealing with traditional time series prediction task with long term dependencies; and (3) DeepLOB specializing in high frequency LOB data.
- Evaluation results on both IID setting and OOD setting support two primary conclusion: (1) domain shifts can cause algorithms without considering OOD generalization fail to work; (2) there is significant room for improvement in machine learning solutions focusing on time series with OOD samples.

2 Limit Order Book Data

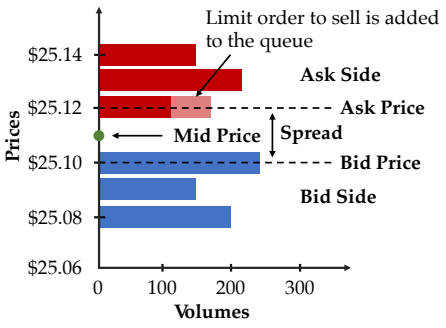

Figure 1: A example of LOB data

As shown in Figure. 1, a LOB record represents a snapshot of the supply and demand for a given security at a given time instance. It serves as a record of all the outstanding buy (ask)/sell (bid) orders organized by price levels. Additionally, LOBs provide information about the order size (volume) at each price level. Mathematically speaking, one can regard a snapshot of a LOB as matrix, where each row vector corresponds the associated price and order size of the traded asset at a particular level. In LOBs, order types can either be limit orders or market orders. A limit order specifies a price level at which a trader is willing to buy or sell the asset of interest. In other words, limit orders are passive orders in the LOB on

the side of the book of the market participant (buy or sell side). In contrast, when a trader places a market order, it indicates that they are willing to buy or sell the asset at the best available price. In this work, we use 100 most recent LOB records with 40 features as the input $X = [x_1, x_2, \ldots, x_{100}]^T$, where $x_t = [p_a^i(t), v_a^i(t), p_b^i(t), v_b^i(t)]_{i=1}^{n=10}$. In addition, for the given LOB data at time t, p^i denotes the price and v_i denotes the volume size at i-th level on both ask side (p_a^i) and bid side (p_b^i). Therefore, the mid-price $p_t = \frac{p_a^1(t)+p_b^1(t)}{2}$ can be used to create the label $y_t = p_t$ for price forecasting.

2.1 Generation of LOB Data Using Multi-Agent Market Simulation

The synthetic dataset used for benchmarking purposes was generated using a multi-agent limit order book market simulator called ABIDES [4]. ABIDES is an event-based simulation environment that is composed of simulation kernel, a single exchange, and various market participants. The simulation kernel manages the flow of time and handles all inter-agent communication. For example, all requests/orders completed by background agents (e.g., placing limit and market orders) is managed via the simulation kernel. The exchange is a NASDAQ-like exchange agent that lists a security for trade against a LOB with FIFO matching rules. The market participants are so-called background agents that represent market agents with different types of trading strategies. Please see the Appendix C for explicit details on ABIDES and different types of background agents.

2.1.1 ABIDES with Distributional Shifts

To introduce a distributional shift to LOB data, we modify the trading behavior of a type of background agent called a value agent. Value agents are strategic agents that base their trading actions on an internal estimate of the fundamental value of the asset of interest. Value agents derive these estimates from noisy observations of the true fundamental (e.g., a mean-reverting process), which is managed by ABIDES. If the estimate of the fundamental value implies that the price of the asset will go up, then the agent will place a buy order. Similarly, if the estimate of the fundamental value implies that the price of the asset will go down, then the agent will place a sell order. We make two important changes to the design of value agents in order to introduce distributional shift into our market data:

1. We introduce a Gaussian shock S to the observed fundamental that occurs at random time T_s in a random direction $d_s \in \{-1, 1\}$, i.e.,
$$S \sim \mathcal{N}(d_s \mu_s, \sigma_s^2), \tag{1}$$
where μ_s and σ_s^2 denote the mean and variance of S, respectively, and $P(d_s = -1) = P(d_s = 1) = 0.5$. Once the value agents observe the shocked fundamental, their belief about the value of the asset will change, leading to a strong shift in their trading actions.
2. To emulate agent behavior under realistic shocks, we modify the counting process driving the arrival of value agents. Specifically, rather than using a homogeneous Poisson process to model value agents arrivals, we consider a non-homogeneous Poisson process with arrival rate function $\lambda_{\text{value}}(t)$. In the presence of a shock, agents arrive at a higher rate, where in this work, we consider the following arrival rate function:
$$\lambda_{\text{value}}(t) = \begin{cases} \bar{\lambda}_{\text{value}}, & t < T_s \\ \bar{\lambda}_{\text{value}}(1 + A_s \exp(-\theta_s(t - T_s)), & t \geq T_s \end{cases}, \tag{2}$$
where A_s is a hyperparameter controlling the scaling factor of the arrival rate and θ_s controls the reversion of the arrival rate to the mean value $\bar{\lambda}_{\text{value}}$. Intuitively, agent arrivals will arrive with constant arrival rate $\bar{\lambda}_{\text{value}}$ before the shock occurs. Once the shock occurs, the arrival rate spikes to a value of $(1 + A_s)\bar{\lambda}_{\text{value}}$ and decays at exponential rate θ_s to $\bar{\lambda}_{\text{value}}$ at $t \to \infty$.

We provide a more detailed background on this configuration in the Appendix C.2.1, as well as the parameter settings utilized to generate the synthetic LOB dataset used in the experiments of this work.

3 Deep Learning Benchmarks and Evaluation

In this work, we can define the distributional shift on LOB data according to the definition of temporal covariate shift [8]: Given a LOB dataset D with n labeled segments according to the type of shocks,

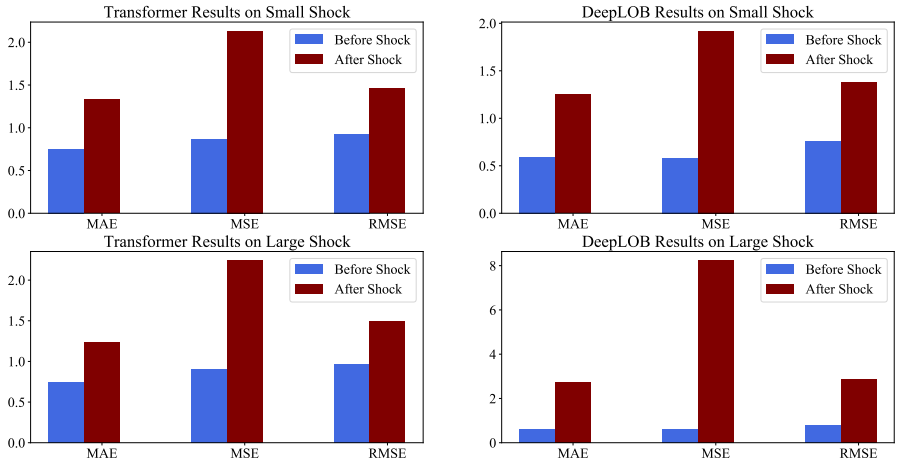

Figure 2: Results on Distributional Shifts according to the shock type.

i.e., $D = \{D_i, \ldots, D_n\}$, the distributional shift is referred to the case that all the segments under same type of shock's influence follow the same data distribution $P_{D_i}(x, y)$, while for different types of shocks where $1 \leq i \neq j \leq n$, marginal probability distributions are different, and the conditional distributions are the same, i.e., $P_{D_i}(x) \neq P_{D_j}(x)$ and $P_{D_i}(y|x) = P_{D_j}(y|x)$.

We plot the time series mid-price data derived from the generated LOB dataset with distributional shift in Figure. 4 and 5 as well as the IID data in Figure. 3 (Refer to Appendix). After generating this dataset, we adapt three baseline algorithms to the mid-price prediction under distributional shift task. Specifically, these baseline algorithms includes (1) AdaRNN [8], which is a deep learning model fits for time series distribution-shift problem; (2) Transformer [39], which is proved the effectiveness of processing time series and (3) DeepLOB [41], which dedicates to handle LOB data.

	IID	Small Shock	Large Shock
AdaRNN	1.02 ± 2.24e-4	1.00 ± 8.07e-5	0.99 ± 3.58e-5
Transformer	0.87 ± 1.98e-3	1.02 ± 6.29e-3	1.08 ± 0.01
DeepLOB	0.66 ± 0.11	1.08 ± 0.07	2.25 ± 0.15

Table 1: RMSE results on synthetic LOB dataset.

We first use independent and identically distributed (IID) data with no shock comes up, training the baseline models, then we build two setting including IID set using the test data before the shock appears and OOD set using the data after the shock appears with distributional shift. Experimental results are shown on Table 1, where we can find that models without any special treatment on OOD problem (Transformer and DeepLOB) suffer from performance degradation against with models dedicated to dealing with distributional shift (AdaRNN). In addition, DeepLOB, which is specifically designed to analyze LOB data, achieves the best performance dealing with IID high-speed trading data. We further analyze how shocks of different magnitude impact the model's performance. We plot the performance metrics for Transformer and DeepLOB applied to small and large shocks before and after the shock on Figure. 2. We find that for a drastic distributional shift (large shock), the performance of the model relying only on the training data degrades more than for the case of a smaller shock. These experimental results suggest that our synthetic dataset can be used to guide the building of machine learning models that are more robust to distributional shifts.

4 Conclusion

This work builds a large-scale synthetic financial LOB dataset which has distributional shift after some random shocks happen and can be used to verify machine learning benchmarks for financial trend prediction tasks. Next, we evaluate this dataset on three different types of forecasting models and show that considering both distributional shift and inherent dependencies is still challenging for current machine learning models. In the future, we plan to refine our agent-based model for shock events to cover more granular stress scenarios.

References

[1] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 1999.

[2] Tucker Hybinette Balch, Mahmoud Mahfouz, Joshua Lockhart, Maria Hybinette, and David Byrd. How to evaluate trading strategies: Single agent market replay or multiple agent interactive simulation? *ArXiv*, abs/1906.12010, 2019.

[3] J. P. Bouchaud, M. Mezard, and M. Potters. Statistical properties of stock order books: Empirical results and models. *Quantitative Finance*, pages 251–256, 2002.

[4] David Byrd, Maria Hybinette, and Tucker Hybinette Balch. Abides: Towards high-fidelity market simulation for ai research. *arXiv preprint arXiv:1904.12066*, 2019.

[5] Diane K Denis, John J McConnell, Alexei V Ovtchinnikov, and Yun Yu. S&p 500 index additions and earnings expectations. *the Journal of Finance*, 58(5):1821–1840, 2003.

[6] Mucong Ding, Kezhi Kong, Jiuhai Chen, John Kirchenbauer, Micah Goldblum, David Wipf, Furong Huang, and Tom Goldstein. A closer look at distribution shifts and out-of-distribution generalization on graphs. In *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021.

[7] Michael E Drew, Tony Naughton, and Madhu Veeraraghavan. Firm size, book-to-market equity and security returns: Evidence from the shanghai stock exchange. *Australian Journal of Management*, 28(2):119, 2003.

[8] Yuntao Du, Jindong Wang, Wenjie Feng, Sinno Pan, Tao Qin, Renjun Xu, and Chongjun Wang. Adarnn: Adaptive learning and forecasting of time series. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 402–411, 2021.

[9] William B Elliott and Richard S Warr. Price pressure on the nyse and nasdaq: Evidence from s&p 500 index changes. *Financial Management*, pages 85–99, 2003.

[10] Mor Geva, Yoav Goldberg, and Jonathan Berant. Are we modeling the task or the annotator? an investigation of annotator bias in natural language understanding datasets. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1161–1166, Hong Kong, China, November 2019. Association for Computational Linguistics.

[11] Martin D. Gould, M. A. Porter, Stacy Williams, Mark McDonald, Daniel J. Fenn, and Sam D. Howison. Limit order books. *Quantitative Finance*, 13:1709 – 1742, 2010.

[12] Martin D Gould, Mason A Porter, Stacy Williams, Mark McDonald, Daniel J Fenn, and Sam D Howison. Limit order books. *Quantitative Finance*, 13(11):1709–1742, 2013.

[13] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization, 2020.

[14] Lin Lawrence Guo, Stephen R Pfohl, Jason Fries, Alistair EW Johnson, Jose Posada, Catherine Aftandilian, Nigam Shah, and Lillian Sung. Evaluation of domain generalization and adaptation on improving model robustness to temporal dataset shift in clinical medicine. *Scientific reports*, 12(1):1–10, 2022.

[15] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *International Conference on Learning Representations*, 2019.

[16] Hailin Hu, MingJian Tang, and Chengcheng Bai. Datsing: Data augmented time series forecasting with adversarial domain adaptation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2061–2064, 2020.

[17] Yeping Hu, Xiaogang Jia, Masayoshi Tomizuka, and Wei Zhan. Causal-based time series domain generalization for vehicle intention prediction. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 7806–7813. IEEE, 2022.

[18] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard L. Phillips, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. Wilds: A benchmark of in-the-wild distribution shifts. In *ICML*, 2021.

[19] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems*, 2020.

[20] Andrey Malinin, Neil Band, German Chesnokov, Yarin Gal, Mark John Francis Gales, Alexey Noskov, Andrey Ploskonosov, Liudmila Prokhorenkova, Ivan Provilkov, Vatsal Raina, Vyas Raina, Mariya Shmatova, Panos Tigas, and Boris Yangel. Shifts: A dataset of real distributional shift across multiple large-scale tasks. *ArXiv*, abs/2107.07455, 2021.

[21] Nathan Ng, Kyunghyun Cho, and Marzyeh Ghassemi. Ssmba: Self-supervised manifold based data augmentation for improving out-of-domain robustness. *arXiv preprint arXiv:2009.10195*, 2020.

[22] Adamantios Ntakaris, Martin Magris, Juho Kanniainen, Moncef Gabbouj, and Alexandros Iosifidis. Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods. *Journal of Forecasting*, 37(8):852–866, 2018.

[23] Adamantios Ntakaris, Martin Magris, Juho Kanniainen, Moncef Gabbouj, and Alexandros Iosifidis. Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods. *Journal of Forecasting*, 2018.

[24] Alan Ramponi and Barbara Plank. Neural unsupervised domain adaptation in NLP—A survey. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6838–6855, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics.

[25] J. Ren, Peter J. Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark A. DePristo, Joshua V. Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. *ArXiv*, abs/1906.02845, 2019.

[26] Jie Ren, Peter J. Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark Depristo, Joshua Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[27] Lukas Ruff, Jacob R. Kauffmann, Robert A. Vandermeulen, Grégoire Montavon, Wojciech Samek, Marius Kloft, Thomas G. Dietterich, and Klaus-Robert Müller. A unifying review of deep and shallow anomaly detection. *CoRR*, abs/2009.11732, 2020.

[28] J. Sirignano. Deep learning for limit order books. *arXiv:1601.01987*, 2016.

[29] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, D. Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015.

[30] Svitlana Vyetrenko, David Byrd, Nick Petosa, Mahmoud Mahfouz, Danial Dervovic, Manuela Veloso, and Tucker Balch. Get real: Realism metrics for robust limit order book market simulations. In *Proceedings of the First ACM International Conference on AI in Finance*, pages 1–8, 2020.

[31] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.

[32] Ziyu Wang, Bin Dai, David Wipf, and Jun Zhu. Further analysis of outlier detection with deep generative models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20, Red Hook, NY, USA, 2020. Curran Associates Inc.

[33] Olivia Wiles, Sven Gowal, Florian Stimberg, Sylvestre Alvise-Rebuffi, Ira Ktena, Taylan Cemgil, et al. A fine-grained analysis on distribution shift. *arXiv preprint arXiv:2110.11328*, 2021.

[34] Olivia Wiles, Sven Gowal, Florian Stimberg, Sylvestre Alvise-Rebuffi, Ira Ktena, Krishnamurthy Dvijotham, and Taylan Cemgil. A fine-grained analysis on distribution shift, 2021.

[35] Garrett Wilson, Janardhan Rao Doppa, and Diane J Cook. Multi-source deep domain adaptation with weak supervision for time-series sensor data. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1768–1778, 2020.

[36] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021.

[37] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *CoRR*, abs/2110.11334, 2021.

[38] Nanyang Ye, Kaican Li, Haoyue Bai, Runpeng Yu, Lanqing Hong, Fengwei Zhou, Zhenguo Li, and Jun Zhu. Ood-bench: Quantifying and understanding two dimensions of out-of-distribution generalization. 2021.

[39] George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff. A transformer-based framework for multivariate time series representation learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2114–2124, 2021.

[40] Haoran Zhang, Natalie Dullerud, Laleh Seyyed-Kalantari, Quaid Morris, Shalmali Joshi, and Marzyeh Ghassemi. An empirical framework for domain generalization in clinical settings. *Proceedings of the Conference on Health, Inference, and Learning*, 2021.

[41] Zihao Zhang, Stefan Zohren, and Stephen Roberts. Deeplob: Deep convolutional neural networks for limit order books. *IEEE Transactions on Signal Processing*, 67(11):3001–3012, 2019.

[42] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.

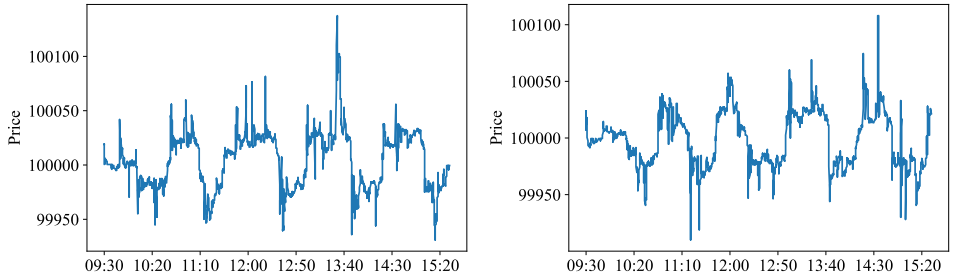

(a) Day 5 without shock coming (b) Day 8 without shock coming

Figure 3: Examples of training data's mid-price.

A Related Work on Financial Datasets.

There are no unifying benchmark datasets for financial tasks. For instance, S&P 500 [5]: Standard & Poor's 500 Index is a market-capitalization-weighted index of the 500 largest U.S. publicly traded companies. NASDAQ index [9]: The NASDAQ is the world's fast electronic stock exchange which operates through computer and telephones, as opposite to traditional method. NASDAQ lists only technology-based companies. Shanghai stock exchange [7]: The Shanghai Stock Exchange (SSE) represents the largest stock exchange in China. It is a non-profit organization run by the China Securities Regulatory Commission (CSRC). Stocks, funds, bonds, and derivatives are all traded on the exchange. Without the definition of a unique dataset and appropriate performance indicators, researchers cannot make a complete comparison between the proposed studies in order to select a suitable solution for a specific problem. Also, the majority of investigated primary studies provide different evaluation metrics for time-series forecasting. However, most of they are extracted individually by authors using different splits and date ranges, which may introduce annotators bias [10]. Thus, there is an urgent need for a comprehensive suite of real-world tasks that combine a diverse set of datasets of various sizes coming from financial institutions. Data splits as well as evaluation metrics are important so that progress can be measured in a consistent and reproducible way.

LOB for a given asset are dispersed across several exchanges, creating a fragmentation of liquidity, which poses a problem for empirical studies. As [11] points out, differences between matching rules and transaction costs of different trading platforms complicate comparisons between different limit order books for the same asset. However, these issues related to fragmentation are not present in the data obtained from the less fragmented Nasdaq Nordic market. In addition, the Helsinki Exchange is a pure limit order market in which market makers have a limited role. For research purposes, the FI-2020 dataset [22] collects high-frequency limit order data of five stocks from the Nasdaq Nordic stock market in 10 consecutive days. [2] show that Interactive Agent-based Simulation (IABS) market environment can be adapted for using as a backtester. LOB-ITCH [23] includes 5 assets over 10 trading days from Nasdanq Helsinki SE, 4,000,000 observations. The proposed task of LOB-ITCH is classification of mid-price movements for 1, 2, 3, 5, and 10 predicted horizons. [23] uses ridge regression and MLP-like network regression methods to verify the quality of LOB-ITCH.

B Related work on Distributional Shift Downstream Tasks

As a way to capture this failure of machine learning models on distributional shifts, many works have been done to find and normalize datasets with out-of-distribution samples [18, 38]. According to the availability information from target domain, we can split the distributional shifts tasks into domain adaptation [31, 35, 16] and domain generalization [6, 42, 17], where the test domains can be visible during the training process of the former, while they are not available during the training of the latter. Previous work use deep auto-regressive generative models [25] or GANs [32] to deal with OOD samples. Specifically, they propose scoring metrics, such as likelihood estimation, to obtain good OOD detectors. Those models have been shown to be effective in evaluating the likelihood of input data and estimating data distributions. For time series domain, anomaly detection (AD) setting can

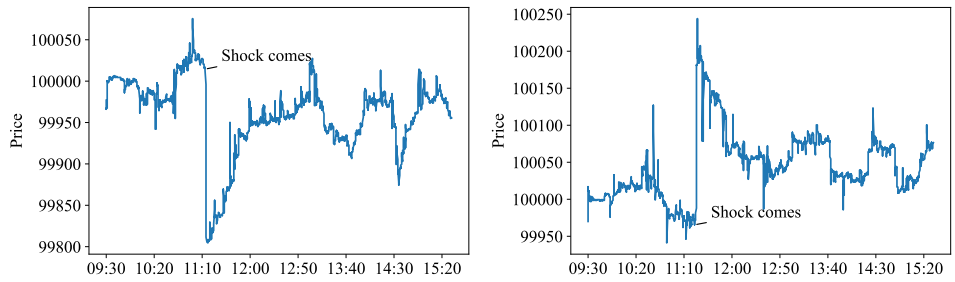

(a) Day 10 with small shock coming (b) Day 120 with small shock coming

Figure 4: Examples of distributional shift (small shock) on synthetic data's mid-price.

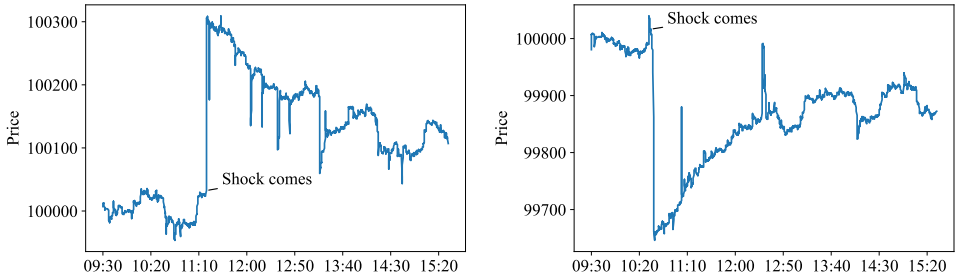

(a) Day 2 with large shock coming (b) Day 19 with large shock coming

Figure 5: Examples of distributional shift (large shock) on synthetic data's mid-price.

be used to solve OOD problems [27, 37]. However, AD is different from OOD in the following 2 aspects: first, OOD samples cannot be used as labeled examples during the training process in AD as distribution of OOD space is ambiguous; second, AD assumes that the observations of normal samples are homogeneous which will fail to detect OOD samples. For the time series dataset, [40] used in-hospital death records and lung x-rays to build distributional shift in the clinical setting, and [14] used patient health records from the ICU in grouping according to year of data collection. Besides, [20] investigate temporal shifts in a large number of weather data. General framework for working with distributional shift problems was developed in [13]. Several classes of approaches to solving distributional shift problems were tested in [33].

C LOB data

C.1 Limit Order Books (LOBs)

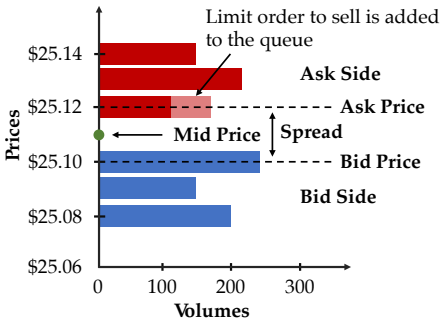

Figure 6: A example of LOB data

A limit order book (LOB) represents a snapshot of the supply and demand for an asset at a given time in the financial markets. It is an electronic record of all outstanding buy and sell limit orders, organized by price level. A buy (sell) order is an order to buy (sell) an asset at a specified price or below (above) a specified price. Limit orders are lined up in the LOB in the corresponding position for resting orders. A limit order above or below a certain price level is sometimes referred to as an under-offer. A market order indicates that the trader is willing to accept the best available price immediately. Figure 6 provides a

diagram illustrating the structure of the LOB. Please refer to [12] for more information on LOBs.

C.2 Background Agents Descriptions for ABIDES

In this work, we mainly looked at the following types of background agents:

- **Noise agents** are non-strategic agents that do not base their trading actions based on intelligent strategies. In particular, noise agents place limit orders of random size (volume) and of random direction (buy or sell) with interarrival times independently sampled from discrete uniform distribution form 1 to 100 nanoseconds.
- **Momentum agents** are agents that place market orders based on mid-price trends. More precisely, given lookback periods $T_{\min}$ and $T_{\max}$ where $T_{\max} > T_{\min}$, momentum agents use a moving average filter to compute the average mid-price over each of the lookback periods. If the mid-price based on the shorter lookback period is larger than that of the longer lookback period, intuitively, the momentum agent believes the price is increasing and will place a buy order of random size (or a sell order in the opposite case). Unlike value agents, momentum agents are configured to have deterministic interarrival times, i.e., they arrive to the market every T_{MOM} seconds.
- **Market maker agents** are agents that supply liquidity to the market by placing orders on both sides of the LOB at various price levels ever T_{MM} seconds. For more information about the market maker used in this configuration of agents, please see [30].
- **Value agents** are strategic agents that base their trading actions based on an internal estimate of the fundamental value of the asset being traded obtained from some noisy observation (see Fig. XX). If the estimate of the fundamental value implies that the price of the asset will go up, then the agent will place a buy order. On the other hand, if the noisy observation of the fundamental implies that the price of the asset will go down, then the agent will place a sell order. Let x_t denote the fundamental value at time t and y_t denote its corresponding noisy observation. In ABIDES, the fundamental is modelled via an Ornstein-Uhlenbeck (OU) process. An OU process is a mean-reverting process and the probability density function fundamental value at time t' given the fundamental value at time $t < t'$ is Gaussian, i.e.,

$$p(x_{t'}|x_t) = \mathcal{N}\left(x_{t'}|\mu + (x_t - \mu)e^{-\theta\Delta_t}, \frac{\sigma_x^2}{2\theta}(1 - e^{-2\theta\Delta_t})\right), \tag{3}$$

  where $\Delta_t = (t' - t)$, μ is mean of the process, σ_x is the volatility of the process, and θ is the mean-reversion parameter. The value agent's observation model of the fundamental is also Gaussian, i.e. at time t' the value agent believes that the observed fundamental $y_{t'}$ is simply a Gaussian perturbation of the true fundamental:

$$p(y_{t'}|x_{t'}) = \mathcal{N}(y_{t'}|x_{t'}, \sigma_y^2), \tag{4}$$

  where σ_y^2 is the observation noise variance. Given the parameters of the OU process and the observation model, value agents obtain an estimate of the fundamental using a simple Bayesian estimation procedure to determine an estimate of the true fundamental that is used to drive their trading decision. For more details on this, please see [XX]. Like the noise agents, the arrivals of a value agent are modelled as a Poisson process in a standard configuration with mean arrival rate λ_{value}.

Given a configuration which specifies the number of each type of agent, as well as their corresponding parameters, ABIDES can be used to simulate LOB data. Next, we will describe the details of the configuration as well as how distributional shift was modellled into the LOB dataset.

C.2.1 Configuration Parameters for Synthetic LOB Dataset

Using ABIDES, we simulate $N_{\mathrm{days}} = 365$ trading days worth of data of which 50% are under ordinary market conditions (no market shocks) and the other 50% are days which experience a market shock of random magnitude/direction. The configuration includes background agents, of which $N_{\mathrm{noise}} = 50$ are noise agents, $N_{\mathrm{value}} = 100$ are value agents with mean arrival rate $\bar{\lambda}_{\mathrm{value}} = 0.005$ seconds, $N_{\mathrm{momentum}} = 10$ are momentum agents with lookback parameters $T_{\min} = 20$ and $T_{\max} = 50$, and

$N_{\mathrm{MM}} = 1$ are market makers with wake-up period $T_{\mathrm{MM}} = 5$ seconds. For all trading days, the parameters of the fundamental are given by: $\mu = 100000$ and $\sigma^2 = 1 \times 10^{-12}$. In total, our synthetic dataset simulates three different scenarios corresponding to both IID and OOD settings:

- **Ordinary market conditions:** This comprises 50% of the simulated data, where no shock is introduced to the fundamental. Under this setting, the mean-reversion parameter of the OU process in this configuration is $\theta = 1 \times 10^{-12}$. The data produced from this setting are used as the training data for the forecasting algorithms (IID).
- **Small shock:** This comprises 25% of the simulated data and assumes a reversion rate of $\theta = 1 \times 10^{-12}$ on the fundamental. For the shock , we model the shock time according to a uniform distribution $T_s \sim \mathcal{U}(1 \text{ hours}, 2 \text{ hours})$ after market open with shock mean and variance as $\mu_s = 200$ and $\sigma_s^2 = 400$, respectively. Furthermore, after the shock occurs, value agent arrivals are governed by he amplification rate $A_s = 2$ and arrival rate reversion parameter $\theta_s = 1 \times 10^{-12}$. The data produced from this scenario are considered OOD inputs.
- **Large shock**: This comprises 25% of the simulated data and assumes a reversion rate of $\theta = 5 \times 10^{-13}$ on the fundamental. For the shock , we model the shock time according to a uniform distribution $T_s \sim \mathcal{U}(1 \text{ hours}, 2 \text{ hours})$ after market open with shock mean and variance as $\mu_s = 400$ and $\sigma_s^2 = 1600$, respectively. Furthermore, after the shock occurs, value agent arrivals are governed by he amplification rate $A_s = 3$ and arrival rate reversion parameter $\theta_s = 1 \times 5^{-13}$. This scenario corresponds to OOD setting, where the distributional shift is larger than that of the small shock scenario. Since the reversion rate of the fundamental and arrival rate amplification are smaller in this setting, the simulated market shock also lasts longer.

D Benchmarks

AdaRNN is the specific model dealing with distribution-shift problem where statistical properties of time series can change over time. The first module of AdaRNN, called Temporal Distribution Characterization (TDC), aims to better characterize the distribution information in the time series. The second module is Temporal Distribution Matching (TDM), which uses a boosting-based procedure to learn the hidden representation and to reduce distribution mismatch in time series. By combining TDC and TDM, AdaRNN can utilize the common knowledge by matching the distribution to learn the efficiency representation and then finish the prediction task with that representation.

Transformer is a encoder and decoder structure to process sequence data, which has achieved excellent performance in several time series tasks with the ability of handling with long term dependence. Each encoder/decoder block consists of a multi-head self-attention module and a position embedding neural network. In addition, each decoder module inserts a cross-attention module between the multi-head self-attention module and the position embedding neural network. Different from LSTM or RNN, transformer uses positional encoding to embed positional information for the input instead of any iterative or convolutional operations. In the task of mid-price prediction under distributional shifts, we use a simple encoder model with attention layer for learning representations and the FFN layer is used to predict the final price.

DeepLOB is a deep neural network architecture containing convolutional layers as well as long short-term memory (LSTM) units for predicting future stock price movements in large-scale high-frequency LOB data. DeepLOB contains with three modules, where CNN module with convolutional and pooling layers extract features automatically to avoid the limitations of hand-crafted features, then the Inception module [29] helps to infer local interactions over different time scales. After that, the resulting feature maps are passed to an LSTM unit that captures the dynamic temporal behavior. In the work, we will use our synthetic data to verify the stability of DeepLOB in the face of OOD problems.

E Discussion on future work

We will describe potential directions of future work in this section. First, we discuss the major challenges of financial system and then list some of major tasks according to the challenges, especially

under OOD setting. After that, we summarize what role can our future benchmark play in machine learning algorithms of financial system.

E.1 Major Challenges in Finance System

The first challenge is that how to create and sharing synthetic data. As we described above, real world finance dataset suffers many limitations which will hurt the academic research in this domain. Thus, we need to build high-quality synthetic data with the following characteristics:

- Capability for containing many different data types, such as numeric, binary and categorical.
- High-quality data with high fidelity.
- Selecting appropriate evaluation metrics for financial data.
- Data generation with privacy guarantees.

The section challenge is there is lack of deep learning benchmarks for relevant financial tasks:

1. **Mid-Price Forecasting under distributional shifts**: Time-series price forecasting under distributional shifts. There will be flash crashes and potential distributional shifts in agents trading behaviors. A good performance on non-shocked data does not imply good performance on shocked data. This is motivated by COVID, previous market crash, earning calls, and etc... This task measures the error in predicted prices of all assets. *Note:* Most recent benchmark work do not consider advances in transformer encoders and attention mechanisms.
2. **Portfolio Optimization**: Deep learning with decision-making. Higher prediction accuracy sometimes might not reflect a profitable model, hence risk and reward structure must also be taken into consideration. This task measures the expected returns of a portfolio over a period of time, factoring in uncertainty, risks, and costs. This can be represented in the form of decision-focused learning with an optimization-focused objective or deep RL and agent-based automated systems that can decide what to hold and when to buy.
3. **Trading Behavior Prediction**: This needs a bit more clarification, but the idea is to predict the behaviors of trading agents in the market, which could be {aggressive, passive, more noise, more momentum}. This is unique because we have the annotated labels from the ABIDES simulation and the task would be financially interesting to understand and predict trading behaviors of the agents. The input here would be different from price forecasting/trend prediction because we have individual orders placed by each agent in the LOB.
4. **Synthetic Data Generation**: Learning generative models for either tabular or time-series data. One branch could be aiming to learn to generate data as realistic as possible while still maintaining statistics of real data. This is interesting from technical standpoints but also poses the challenges of data generation under certain market constraints. Another branch could be learning to generate data from few samples/points in the cases of adverse and shocked events. This task measures how well generative models can learn to generate data without or without complex simulations.
5. **Anomaly Detection + Early Classification**: This last task a bit more general, but the goal here is to detect adverse events in any of the assets over any long period of trading days. The earlier we can detect this events, the better and could be an interesting financial task. Input would be similar to price forecasting/trend prediction but the output results would be different if we have annotated labels of when the adverse event happened.

E.2 Potential of LOB benchmark of Machine Learning in Financial System

According to the above tasks, we can list two future directions of our benchmark:

- Improving prediction, classification, anomaly detection algorithms with current state-of-the-art and data-hungry deep learning algorithms. Our datasets and benchmarks will provide a playground to further advance research in time-series ML.

- Going from predictions to decision-making, with realistic decisions being made on the predicted from ML algorithms, our datasets can provide opportunities for operation research via the the predict-then-optimize framework.