MSE 448 Midterm Presentation

Aman Sawhney, Yang Fan, Chris Lazarus

April 27, 2021

Outline

Overview

- High frequency traders (HFTs) do not make directional predictions about stocks due to fundamental value.
- High frequency trading is a problem suited to reinforcement learning.
- Generative Adversarial Networks can produce stock market simulators that closely replicate real stock paths without enforcing strict stochastic modeling assumptions.

High Frequency Trading

HFTs do not make directional opinions on the movement of an asset due to fundamental value. Instead they concern themselves with micro movements as a result of another party trading, or changes in liquidity [WYLJ+15]. As a result, HFT strategies do not attempt to model the fundamental value of an asset, but rather trade strictly due to technical factors.

High Frequency Trading as an MDP

Since HFT strategies rely on taking and providing liquidity when it is appropriate, we make the modeling assumption that order book dynamics are a markov process. Hence we may formulate a HFT strategy as a Markov Descion Process in the following manner:

- We will assume discrete time intervals which will be determined by our time-scale, T.
- For each time step $0 \leq t \leq T$ we have
  - $s_t := (O_t, q_t)$ where O_t is the order book history at time t over a look back period steps and q_t is the amount of the asset the agent currently holds.
  - $a_t \in \{T, P, N\}$ where T is the act of taking liquidity, P is the act of providing liquidity, and N is the act of doing nothing.
  - $r_{t+1} = \frac{R_t - R_f}{\sigma_t}$ where R_t is the rolling return of the strategy over a look back period steps and σ_t is the rolling standard deviation over that same period.

This framework generalizes to a multidimensional asset space.

Our Data

- We use MayStreet as our primary data source.
- To avoid liquidity issues we opted to trade only assets within the S&P 500.
- We will aggregate our order book data over one second time-steps. This is because:
  - Research by Quantitative Brokers has demonstrated significant price impact over the 1s time scale for futures assets. Hence a HFT that can anticipate these trades and provide or take liquidity appropriately can generate significant alpha. We would expect similar price action to exist in equity markets [Sch20].

Baseline Model using RL

To investigate the efficacy of RL we opted to attempt to produce alpha using an agent trading on lower frequency data.

- We used the FinRL [LYC$^+$20] package to train a reinforcement learning agent on 5-minute aggregated data. We use 3 years of 5-minute AAPL aggregated data to train our reinforcement learning agent.
- This basic agent is only able to access the top of the book prices, and we assume the agent is able to eat the bid and the ask at any time. We use these top of book prices to calculate various technical features. Additionally we use Proximal Policy Optimization to train our RL agent.

Baseline Model using RL Cout.

We train on data from 2018-01-01 to 2020-01-01 and trade from 2020-01-01 to 2021-01-01.

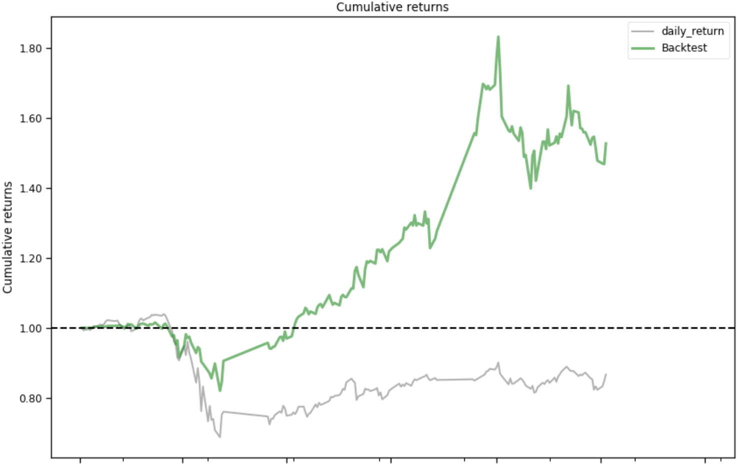

Baseline Model using RL Cout.

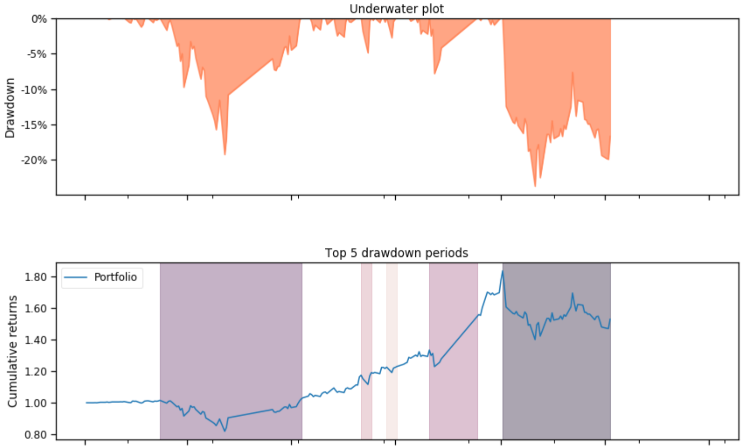

Baseline Model using RL Cout.

Our strategy obtains the following results:

- Annual return: 78.668%
- Annual volatility: 45.313%
- Sharpe ratio: 1.51
- Max drawdown: -23.697%
- Daily value at risk: -5.438%

Feature Engineering - DeepLOB [ZZR19]

Technical indicators such as MACD and the RSI and preprocessing mechanisms such as PCA often make implicit assumptions so it makes sense to use models without fixed parameters.

- 1-D backward convolutional filter on the time axis to automatically extract features
  - $y(n) = \sum_{k=0}^{M} b_k x(n-k)$ where output $y(n)$ is determined by past values of the input signal $x(n)$, and b_k's are coefficients to be learned.
- Inception modules to wrap several convolutions of different sizes together
- LSTM to replace feed forward network and Capture temporal relationships

Feature Engineering - Transformers [Wal20]

A transformer block consists of a specific combination of multi-head self-attention, residual connections, layer normalization and feedforward layers.

- The state-of-art model in the realm of natural language processing.
- Self-attention to replace Long Short-Term Memory (LSTM).
- Performs extremely well on the FI-2010 dataset (≈ 88 F1-score), which is which is a LOB of five instruments over a ten day period.
- We can pretrain the transformer on prediction tasks first and then connect heads to DQN.
- Easy to implement transformers in PyTorch.

Feature Engineering - Fourier Transform [RK20]

The paper compared performances of Gated Recurrent Unit (GRU), which is a variation on LSTM, when using traditional technical indicator vs Fourier transformed features.

- Discrete signals can be decomposed as $s(n) = \frac{1}{N}\sum_{k=0}^{N} c_k e^{\frac{i2\pi kn}{N}}$, where N is the total number of discrete points. $c_m \in \mathbb{C}$.
- Extract the frequency with the maximum magnitude from the decomposition of prices (LOB level, OHLC).
- Generate statistically significant improvement on the GRU model.

Data Augmentation - GANs

RL algorithms struggle with sample efficiency and are thus data hungry.

- Most RL algorithms learn by trial and error and require a simulator. Alternatively, massive amounts of data can be used for batch learning.
- There is only one history of financial data. We need more data to perform RL.
- GANs have been shown to successfully generate time series with close distributional properties as the training data. [WKKK20]

Plan for the project

- Set up testing framework
- Write our own implementation of the RL algorithm
- Implement Fourier Transform on LOB price data and feed it into Transformer layers before entering the DQN
- Data augmentation
- Analysis of results

References I

Xiao-Yang Liu, Hongyang Yang, Qian Chen, Runjia Zhang, Liuqing Yang, Bowen Xiao, and Christina Dan Wang, *Finrl: A deep reinforcement learning library for automated stock trading in quantitative finance*, Deep RL Workshop, NeurIPS 2020 (2020).

Dragana RadojiÄ iÄ and Simeon Kredatus, *The impact of stock market price fourier transform analysis on the gated recurrent unit classifier model*, Expert Systems with Applications **159** (2020), 113565.

I. Scheinfeld, *Causal forests for inferring market action price impacts*, Tech. report, Quantitative Brokers, September 2020.

J. Wallbridge, *Transformers for limit order books*, arXiv:2003.00130 (2020).

References II

Magnus Wiese, Robert Knobloch, Ralf Korn, and Peter Kretschmer, *Quant gans: deep generation of financial time series*, Quantitative Finance **20** (2020), no. 9, 1419â 1440.

K. Webster, M. Alvarez Y. Luo, J. Jussa, S. Wang, G. Rohal, A. Wang, D. Elledge, and G. Zhao, *A portfolio manager's guidebook to trade execution*, Tech. report, Deutsche Bank, July 2015.

Zihao Zhang, Stefan Zohren, and Stephen Roberts, *Deeplob: Deep convolutional neural networks for limit order books*, IEEE Transactions on Signal Processing **67** (2019), no. 11, 3001â 3012.