

Dynamic Portfolio Management Driven by GNNs and Reinforcement Learning

Leveraging LLMs for Alternative Data Extraction from News and SEC Filings

Graduate



Sri Ram Kalyan Venkat Kovela

Introduction: Active portfolio management constitutes one of the most important and challenging decision-making processes in the financial industry. It involves allocating dynamically a certain amount of capital into a portfolio of assets, with the objective of maximizing the profits constrained to a given risk level, thereby outperforming a passively managed portfolio. In this master's thesis I explored the use of Reinforcement Learning (RL) to address the complex challenge of active stock portfolio management. The model advances beyond conventional approaches based on market data by integrating alternative data sources, including unstructured textual data from stock news. Large Language Models (LLMs) are utilized to extract meaningful insights from this unstructured data, transforming vast amounts of unstructured information into valuable inputs that enhance the allocation model's decision-making process. To better capture the relationships between assets, the model replaces the traditional Pearson correlation with quantile correlation, which measures tail dependence. This approach is expected to reduce portfolio risk more effectively by accounting for extreme market movements.

Approach: Several types of financial data were considered in this work, including text data. In order to get sentiment scores from the SEC filings, FIN GLM 6b was fine-tuned on SEC filings. As the pricing data for the securities is daily, later it was decided to include the daily news data for the stocks and use the text as their embeddings instead of a sentiment score. The stock price dataset for ~200 securities along with their fundamentals information from 2016-01-01 to 2024-04-01 (yyyy-mm-dd) was considered and trained on two different models i.e. THGNN and GPM. The embeddings were obtained using an API from JINA. Dimensionality reduction on the embeddings was performed using PCA. Back Testing was done on several time periods and market scenarios and thus the performance of the models was evaluated.

Conclusion: The performance metrics of our trained FIN CHAT GLM 6-B model was within 5% of the performance of the FINGPT 13-B V3.1 model when tested against four different standard financial text benchmarks. This proves that smaller language models trained on robust datasets achieve performance matching that of larger models even with more than double the parameters. Though the THGNN didn't yield good results, GPM generated some interesting ones. We were able to achieve Maximum Draw Down of less than 10% when tested on data before the covid market crash. We have displayed that with the combination of pricing, fundamentals and news, we can achieve good predictive performance. Though the quantile correlation which could have supported better in the extreme market events like covid, processing that

knowledge as a part of compressed node embeddings instead of directly through the model training could have had their effect reduced, which can be looked into for future work.

Fig 1- GPM, Model Architecture

<https://doi.org/10.1016/j.neucom.2022.04.105>




Fig 2- Performance of the GPM model.

Source: Own Representation





Fig 3- Performance of the trained FIN CHAT GLM -6B model

Source: Own Representation



Advisor

Dr. Shao Jü Woo

Co-Examiner

Mag. Dr. Peter Reiter

Subject Area

Data Science,
Computer Science

