Adapting Natural Language Processing Strategies for Stock Price Prediction

Frederic Voigt[1]

Abstract:

Due to the parallels between Natural Language Processing (NLP) and stock price prediction (SPP) as a time series problem, an attempt is made to interpret SPP as an NLP problem. As adaptable techniques word vector representations, pre-trained language models, advanced recurrent neural networks, unsupervised learning methods, and multimodal methods are introduced and it is outlined how they can be transferred into the stock prediction domain.

Keywords: Stock Price Prediction; Financial Analysis; Quantitative Analysis; Fundamental Analysis; NLP

1 Introduction

As in the field of economics various theories argue about the extent to which it is possible to predict future stock prices, stock price prediction (SPP) can be generally considered a difficult problem. The "Random Walk Theory" (RWT), proposed by Kendall and Hill [KH53], implies that, due to the random nature of stock prices, prediction is impossible. The "Efficient Market Hypothesis" (EMH), on the other hand, states that all information, available to the investors, is represented in the stock price [Fa70] (sometimes even set in a causal relationship with the RWT [Gi01]), which is seen by some researchers as a possibility to predict stock prices in accordance with the EMH [Li19].

While the market hypotheses and their consequences are debated in economics, many computer scientists, mathematicians, and engineers have taken the opportunity to build models for SPP based on a quantitative/technical point of view [De15]. The quantitative analysis aims to predict future stock prices by creating models based only on learning historical stock prices [Pé21]. In contrast, fundamental analysis aims to predict future stock prices by taking into account factors beyond mere price values [WHM15]. These factors encompass a wide range, from newspaper articles, tweets, social media data to business reports of individual companies.

[1] Hochschule für Angewandte Wissenschaften Hamburg, Stiftstraße 69, 20999 Hamburg, Germany, frederic.voigt@ventric.de

doi:10.18420/ki2023-dc-03

Stock prices can be understood as a sequence of individual prices, where the stock x_i of the company c_i has a value of x_i^t at a time t. The quantitative analysis aims to predict a future price x_i^{t+1} from the historical prices $\{x_i^j : j \leq t\}$.

Natural Language Processing (NLP) is not only one of the most intensively researched subdisciplines of machine learning (ML), but also one that has brought some of the most impressive results in the ML domain so far. NLP is intuitively suited to be transferred into another time series prediction problem. A stock price (of a single company) is nothing more than a sequence of individual stock prices, while a NLP text is nothing more than a sequence of tokens, that can be indexed. The definition of the goal of quantitative analysis, given at the beginning, can be translated into a "Causal Language Modeling" task, where the next token v^{t+1} has to be predicted for a given sentence beginning $\{v^j : j \leq t\}$ [Li22b]. If one would like to predict not only the $t+1$-th price of a stock i, but a whole price trend this could be formulated as the attempt to predict the probability for a certain price development

$$P(\ (x_i^{j_1}, x_i^{j_1+1}, ..., x_i^{j_2})\ )\ , \tag{1}$$

which can be expressed as

$$P(\ (x_i^{j_1}, x_i^{j_1+1}, ..., x_i^{j_2})\ ) = \prod_{l=j_1}^{j_2} p(x_i^l | x_i^{j_1}, ..., x_i^{l-1})\ . \tag{2}$$

This definition of a stock price development is the same as the definition for "Unconditional-Language Modeling", where the probability of a sentence $P(v)$ to occur is recursively computed following the chain rule of conditional probabilities [SBC19]. As shown by the examples of Next-Token Prediction and Unconditional-Language Modeling, the parallels between NLP and SPP are obvious. Nevertheless, to my knowledge, there is no approach to explore these parallels in a systematic way, using NLP techniques for SPP.

2 Motivation

Many of the techniques and models developed in NLP are transferred to other domains, where they are successfully used. One example is the application of (in an NLP context developed) transformer models (e.g. BERT [De19]) to the image-processing domain (e.g. ViT [Do20]). The transfer of ML models, between different ML subdisciplines, is one of the most important building blocks of research since it helps to promote the understanding of the models, enables the effective use of carefully developed research results, and can help to improve the performance of ML in certain domains by orders of magnitude.

In addition to the importance for ML research to adapt solution strategies to other problem domains, the technicality of the stock market is not to be underestimated. Next to the possibility of generating profits, risk management, and sensible profitability are also important aspects for many companies, to ensure economic stability.

3 Research Question

Considering the parallels highlighted in Sect. 1, the research gap that is intended to be addressed within the scope of my thesis becomes evident. In the existing research, in my opinion, there has been little to no exploration of the extent to which NLP techniques can be applied to other time series problems in general, and specifically to SPP.

Therefore the following questions can be derived:

- How helpful are NLP research approaches for time series problems?
- Can NLP techniques (not just model architectures) be transferred to other domains?
- Can SPP be conceptualized as an NLP problem, with the stock market being considered as a “language” that can be learned?

These questions highlight the unexplored territory where this thesis aims to make significant contributions.

4 Related Work

While some popular NLP models are being used for SPP, there is a lack of systematic adaptation. As a result, most approaches in SPP differ significantly from NLP, as they merely adapt the models without delving into the underlying training methods, training tasks, or data representation. To illustrate this further, the following provides an overview of the current state of SPP, with a particular focus on NLP-based models.

In [St20], four directions for the incorporation of ML in SPP were mentioned; neural networks (NN), support vector machines (SVM), genetic algorithms (GA), and other methods. This roughly aligns with [Pa21], where a certain relevance is also attributed to Naïve Bayes, linear regression, and clustering methods (along others).

Despite being slowly but surely outperformed by transformer models and at most appearing in combination with them, RNNs (as well as LSTMs and GRUs) [Hu22], are still very popular in the field of SPP and financial time series.

The use of RNNs in SPP can be broadly categorized in two directions: RNNs that analyze historical prices, in the sense of the quantitative approach, and those that include other modalities (mainly news and social media) like [Ak16]. If based on a quantitative approach the RNNs are usually used as modules of more complex architectures [Fe19, BYR17, Fe18, Qi17], but there have also been simple approaches where input features are fed in the RNN and the prediction is used directly [NPdO17, DPH16, CZD15]. A particular focus here is on the work of Zhang et al. [ZAQ17], which attempts to decompose different frequencies/patterns in historic stock data. This may seem a bit far from NLP at first, however, there is research like [Hi11] that manages to learn hierarchical structures in texts using models based on processing the input on different frequencies.

Due to the impressive success of various transformer-based models, especially BERT, transformers have gained popularity beyond the borders of the NLP subdiscipline. For purposes of SPP, Ding et al. [Di20] were arguably the first to use transformers for financial time series data back in 2020 while also following the quantitative approach. The model in [Di20] uses a sliding window over the data to predict future stock movements. This has been also done in slight variations by most other transformer-based models [Mu22, SS21, Zh22]. There are also transformer-based models in combinations with RNNs (LSTMs and GRUs) [LSZ18, Hu21, RPAGNV21, LQ23]. The model in [Qi17] is remarkable, not only because it is free of any multimodality, but also because exogenous time series are used, making a first step towards the interpretation of SPP as a two-dimensional (t *and* C) problem (which is similar to trying to predict x^{t+1}). More popular is the use of transformers in SPP in combination with textual inputs (from the news or social media), as a second modality, as done in [DWL20, DL21, Li22a, XC18, Li19, Hu18, Li17, Di15]. As one of the few that follows a generative approach, in [Mi22], a model was designed which has an encoder-decoder architecture in order to generate the stock prices of the next five days.

5 Approach

To explore the research gap, various NLP techniques will be selected and evaluated for their adaptability to the SPP domain. The following is a list of what I consider to be the most appropriate NLP techniques that could be adapted for SPP. Most of them can be combined with each other and do not have to be understood as separate projects.

Stock2Vec Model Texts in NLP are typically decomposed into individual word tokens $v_i \in X' \subseteq \mathbb{N}$. Each v_i is then assigned a (trainable) embedding vector $\hat{e}_i \in \mathbb{R}^H$ (with H as the embedding dimension), which facilitates the processing of the models. The dense vector representations $\hat{e}_i$ are often trained by using an auxiliary network to derive either surrounding context words (skip-gram) or a context is used to derive the current word (continuous bag of words) [Mi13]. In both cases, $\hat{e}$ is trained to represent the semantic expression of the word. This can for example result in the fact that different $\hat{e}$-s representing semantically similar words are close to each other or that pairs of $\hat{e}$-s with the same distances and angles to each other represent word pairs with same relationships to each other.

It would be desirable that stock prices can also be transformed into high-dimensional vectors, where vector pairs with similar angles and distances to each other are in a similar relation to each other. Similar research with simplified approaches already exists, for example, by Dolphin et al. [DSD22]. An example is illustrated in Fig. 1.

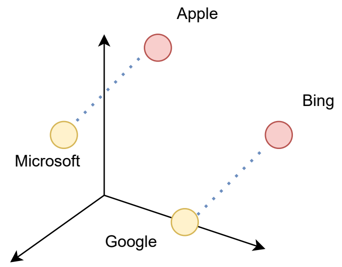

Fig. 1: Example of two vector pairs of price embeddings in a high dimensional (in this example for illustration purposes only three dimensional) vector space. The companies in the pairs are each other's largest competitors, as shown by the fact that the angles and distances of the two pairs of competitors are equal.

To define embeddings $e \in \mathbb{R}^H$ for stock prices, several options are available which can be tested in the course of the thesis. Vector representations for company prices e_i can be trained by using one auxiliary network to predict past/future prices of x_i and one auxiliary network to predict other company prices based on x_i^t. The idea is to assign a token to each company, which can be embedded in a multidimensional vector. Thus, the relationships of the companies can be encoded by the respective vector space positions. The current economic situation (e.g. current stock prices) could be indicated by shifts in the vector space.

Pre-trained Stock Price Prediction Models Few developments have influenced the NLP field in recent years as much as the use of pre-trained language models. These models, mostly based on transformers, are trained with different pre-training tasks to develop a generalized language understanding. With this basis, the models can then be used for any downstream tasks such as sentiment prediction, summary generation or similar. To the best of my knowledge, no approach in the SPP domain has yet dealt with the use of pre-trained models representing a generalized understanding of stock markets.

To change this, one could pre-train a transformer-encoder-based model on adapted pre-training tasks from NLP. As a first idea, "Mask-Language Modeling" (predicting a masked word in the context of the rest of the sentence) is adapted as "Masked-Price Prediction" (MPP). In MPP, x_i^t is to be predicted given a context x_i^{t-T},..., x_i^{t+T}. The "Next-Sequence Prediction" (NSP) task, where the model has to decide whether one sentence follows another, is also adaptable. In terms of SPP, a "Next-Price Matching" (NPM) task can be formulated, where the model has to decide whether one series of prices follows another. The actual use of the pre-trained model is, as in NLP, not limited to a specific task and it can be explored what the model can be used for, whether it is for SPP, stock movement prediction, prediction of single prices or whole market sectors.

Multimodal Transformer As discussed in Sect. 1, the fusion of quantitative stock analysis and NLP data from various sources is well-suited for multimodal analysis of the SPP problem. For NLP data analysis, utilizing pre-trained language processing models like BERT is practical. The insights derived from this analysis can then be seamlessly integrated in a multimodal fashion. Transformer models, particularly, have gained immense popularity for combining different modalities, evident in numerous sub-disciplines that connect NLP with other data sources. For instance, transformer models like ALBEF [Li21] have gained recognition for effectively linking visual and NLP data in the vision-language (V+L) reserach domain. The specific integration process for this thesis can be determined during the course of the research. Highlighting the significance of the multimodal aspect is crucial, as it enables the integration of the two main fields of interest in the thesis, namely SPP and NLP. The combination of multimodal approaches with SPP and NLP is currently a prominent area of research, as indicated by the publications mentioned in Sect. 4.

Unsupervised Neuron In the SPP literature, the problem of too low input dimensionality is often addressed [DL21]. Besides the trained contextual embeddings there is another NLP technique that is suitable for adaptation. Trying to enter a text letter(/byte)-wise into a model (often, even unsupervised) has already been shown by Radford et al. [RJS17] to achieve state-of-the-art performance in sentiment analysis. Here, a model was trained to predict the next letter in a text and meanwhile developed a single neuron suitable for sentiment prediction.

While other works have been investigating character-level processing [ZZL15] the idea to elaborate embedded domain knowledge from unsupervised (low dimensional) processing has also been followed [Ra19]. An approach that can be adapted for SPP by performing "Next-Price Prediction" on an RNN-based model to infer embedded domain knowledge.

Advanced Recurrent Neural Networks Another possible use of RNNs in NLP is trying to decompose the hierarchical structure of language, which has been done with models that pay a certain amount of attention to the frequency of the incoming data [Hi11]. As mentioned in Sect. 4, there are works in the SPP area like [ZAQ17] which deal with recognizing different frequencies in the stock market. Processes that occur at different frequencies in the stock market can also be regarded as hierarchical structuring. For example, the season can be a very coarse-grained (slow) frequency that affects stock performance, while fluctuations in productivity due to the day of the week appear in a high frequency. The most promising models are the Multiple-Timescale RNN (MTRNN) [Hi11] and Clockwork RNNs (CWRNN) [Ko14], which are not very popular at the moment but can be helpful for these tasks.

6 Evaluation

The evaluation of approaches at a technical level is a trivial task, as there is ample availability of historical data. These data can be compared with the predicted data using various metrics such as mean squared error, Matthews correlation coefficient, R2 score, RMSE, etc.

The contribution to ML research needs to be evaluated for each approach. For instance, the Stock2Vec model may successfully capture the relationships between different companies without necessarily improving the prediction of stock prices. However, this insight can still be valuable as it demonstrates the utility of NLP techniques in other domains. For pre-trained stock models, the performance of pre-training tasks can be measured. In the case of hierarchy-analysis-based RNN models, it is interesting to assess whether frequencies in stock price movements are detected.

7 Summary

To summarize the underlying thesis, the adaptability of techniques from the NLP subdiscipline to the SPP time series problem will be explored. Excerpts of models from the NLP field have already been adopted in previous research on SPP, but entire techniques, training approaches, or application strategies have not been adapted. As initial steps, the testing of word vector representations, pre-trained language models, advanced recurrent neural networks, unsupervised learning methods, and multimodal approaches for adaptation will be performed.

Bibliography

[Ak16] Akita, Ryo; Yoshihara, Akira; Matsubara, Takashi; Uehara, Kuniaki: Deep learning for stock prediction using numerical and textual information. In: 2016 IEEE/ACIS 15th International Conference on Computer and Information Science (ICIS). pp. 1–6, 2016.

[BYR17] Bao, Wei; Yue, Jun; Rao, Yulei: A deep learning framework for financial time series using stacked autoencoders and long-short term memory. PLOS ONE, 12(7):1–24, 07 2017.

[CZD15] Chen, Kai; Zhou, Yi; Dai, Fangyan: A LSTM-based method for stock returns prediction: A case study of China stock market. 2015 IEEE International Conference on Big Data (Big Data), pp. 2823–2824, 2015.

[De15] DeFusco, Richard; McLeavey, Dennis; Pinto, Jerald; Runkle, David: Quantitative Investment Analysis. 2015. . John Wiley & Sons. (Cited on pages 1 and 3).

[De19] Devlin, Jacob; Chang, Ming-Wei; Lee, Kenton; Toutanova, Kristina: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and

Short Papers). Association for Computational Linguistics, Minneapolis, Minnesota, pp. 4171–4186, jun 2019.

[Di15] Ding, Xiao; Zhang, Yue; Liu, Ting; Duan, Junwen: Deep Learning for Event-Driven Stock Prediction. In: Proceedings of the 24th International Conference on Artificial Intelligence. IJCAI'15. AAAI Press, p. 2327–2333, 2015.

[Di20] Ding, Qianggang; Wu, Sifan; Sun, Hao; Guo, Jiadong; Guo, Jian: Hierarchical Multi-Scale Gaussian Transformer for Stock Movement Prediction. In (Bessiere, Christian, ed.): Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20. International Joint Conferences on Artificial Intelligence Organization, pp. 4640–4646, 7 2020. Special Track on AI in FinTech.

[DL21] Daiya, Divyanshu; Lin, Che: Stock Movement Prediction and Portfolio Management via Multimodal Learning with Transformer. In: ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 3305–3309, 2021.

[Do20] Dosovitskiy, Alexey; Beyer, Lucas; Kolesnikov, Alexander; Weissenborn, Dirk; Zhai, Xiaohua; Unterthiner, Thomas; Dehghani, Mostafa; Minderer, Matthias; Heigold, Georg; Gelly, Sylvain; Uszkoreit, Jakob; Houlsby, Neil: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. CoRR, abs/2010.11929, 2020.

[DPH16] Di Persio, Luca; Honchar, Oleksandr: Artificial neural networks architectures for stock price prediction: Comparisons and applications. International journal of circuits, systems and signal processing, 10(2016):403–413, 2016.

[DSD22] Dolphin, Rian; Smyth, Barry; Dong, Ruihai: Stock Embeddings: Learning Distributed Representations for Financial Assets. ArXiv, abs/2202.08968, 2022.

[DWL20] Daiya, Divy; Wu, Min-Sheng; Lin, Che: Stock Movement Prediction That Integrates Heterogeneous Data Sources Using Dilated Causal Convolution Networks with Attention. pp. 8359–8363, 05 2020.

[Fa70] Fama, Eugene F.: Efficient Capital Markets: A Review of Theory and Empirical Work. The Journal of Finance, 25(2):383–417, 1970.

[Fe18] Feng, Fuli; Chen, Huimin; He, Xiangnan; Ding, Ji; Sun, Maosong; Chua, Tat-Seng: Enhancing Stock Movement Prediction with Adversarial Training, 2018.

[Fe19] Feng, Fuli; He, Xiangnan; Wang, Xiang; Luo, Cheng; Liu, Yiqun; Chua, Tat-Seng: Temporal Relational Ranking for Stock Prediction. ACM Trans. Inf. Syst., 37(2), mar 2019.

[Gi01] Gidófalvi, Györző: Using News Articles to Predict Stock Price Movements. 2001.

[Hi11] Hinoshita, Wataru; Arie, Hiroaki; Tani, Jun; Okuno, Hiroshi G.; Ogata, Tetsuya: Emergence of hierarchical structure mirroring linguistic composition in a recurrent neural network. Neural Networks, 24(4):311–320, 2011.

[Hu18] Hu, Ziniu; Liu, Weiqing; Bian, Jiang; Liu, Xuanzhe; Liu, Tie-Yan: Listening to Chaotic Whispers: A Deep Learning Framework for News-Oriented Stock Trend Prediction. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18, Association for Computing Machinery, New York, NY, USA, p. 261–269, 2018.

[Hu21] Hu, Xiaokang: Stock Price Prediction Based on Temporal Fusion Transformer. In: 2021 3rd International Conference on Machine Learning, Big Data and Business Intelligence (MLBDBI). pp. 60–66, 2021.

[Hu22] Hutchins, DeLesley; Schlag, Imanol; Wu, Yuhuai; Dyer, Ethan; Neyshabur, Behnam: Block-Recurrent Transformers, 2022.

[KH53] Kendall, M. G.; Hill, A. Bradford: The Analysis of Economic Time-Series-Part I: Prices. Journal of the Royal Statistical Society. Series A (General), 116(1):11–34, 1953.

[Ko14] Koutník, Jan; Greff, Klaus; Gomez, Faustino J.; Schmidhuber, Jürgen: A Clockwork RNN. CoRR, abs/1402.3511, 2014.

[Li17] Li, Qing; Wang, Jun; Wang, Feng; Li, Ping; Liu, Ling; Chen, Yuanzhu: The role of social sentiment in stock markets: a view from joint effects of multiple information sources. Multimedia Tools and Applications, 76(10):12315–12345, May 2017.

[Li19] Liu, Jintao; Lin, Hongfei; Liu, Xikai; Xu, Bo; Ren, Yuqi; Diao, Yufeng; Yang, Liang: Transformer-Based Capsule Network For Stock Movement Prediction. In: Proceedings of the First Workshop on Financial Technology and Natural Language Processing. Macao, China, pp. 66–73, August 2019.

[Li21] Li, Junnan; Selvaraju, Ramprasaath; Gotmare, Akhilesh; Joty, Shafiq; Xiong, Caiming; Hoi, Steven Chu Hong: Align before Fuse: Vision and Language Representation Learning with Momentum Distillation. In (Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P. S.; Vaughan, J. Wortman, eds): Advances in Neural Information Processing Systems. volume 34. Curran Associates, Inc., pp. 9694–9705, 2021.

[Li22a] Li, Yawei; Lv, Shuqi; Liu, Xinghua; Zhang, Qiuyue: Incorporating Transformers and Attention Networks for Stock Movement Prediction. Complexity, 2022:7739087, Feb 2022.

[Li22b] Liu, Hao; Geng, Xinyang; Lee, Lisa; Mordatch, Igor; Levine, Sergey; Narang, Sharan; Abbeel, Pieter: Towards Better Few-Shot and Finetuning Performance with Forgetful Causal Language Models, 2022.

[LQ23] Li, Chengyu; Qian, Guoqi: Stock Price Prediction Using a Frequency Decomposition Based GRU Transformer Neural Network. Applied Sciences, 13(1), 2023.

[LSZ18] Li, Hao; Shen, Yanyan; Zhu, Yanmin: Stock Price Prediction Using Attention-based Multi-Input LSTM. In (Zhu, Jun; Takeuchi, Ichiro, eds): Proceedings of The 10th Asian Conference on Machine Learning. volume 95 of Proceedings of Machine Learning Research. PMLR, pp. 454–469, 14–16 Nov 2018.

[Mi13] Mikolov, Tomás; Chen, Kai; Corrado, Greg; Dean, Jeffrey: Efficient Estimation of Word Representations in Vector Space. In: 1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings. 2013.

[Mi22] Mirjebreili, Seyed; Solouki, Ata; Soltanalizadeh, Hamidreza; Sabokrou, Mohammad: Multi-Task Transformer for Stock Market Trend Prediction. pp. 101–105, 11 2022.

[Mu22] Muhammad, Tashreef; Aftab, Anika Bintee; Ahsan, Md. Mainul; Muhu, Maishameem Meherin; Ibrahim, Muhammad; Khan, Shahidul Islam; Alam, Mohammad Shafiul: Transformer-Based Deep Learning Model for Stock Price Prediction: A Case Study on Bangladesh Stock Market, 2022.

[NPdO17] Nelson, David M. Q.; Pereira, Adriano C. M.; de Oliveira, Renato A.: Stock market's price movement prediction with LSTM neural networks. In: 2017 International Joint Conference on Neural Networks (IJCNN). pp. 1419–1426, 2017.

[Pa21] Patel, Ramkrishna; Choudhary, Vikas; Saxena, Deepika; Singh, Ashutosh Kumar: Review of stock prediction using machine learning techniques. In: 2021 5th International Conference on Trends in Electronics and Informatics (ICOEI). IEEE, pp. 840–846, 2021.

[Pé21] Pérez, Gerardo Alfonso: Forecasting and optimization of stock trades. 2021.

[Qi17] Qin, Yao; Song, Dongjin; Chen, Haifeng; Cheng, Wei; Jiang, Guofei; Cottrell, Garrison W.: A Dual-Stage Attention-Based Recurrent Neural Network for Time Series Prediction. CoRR, abs/1704.02971, 2017.

[Ra19] Radford, Alec; Wu, Jeff; Child, Rewon; Luan, David; Amodei, Dario; Sutskever, Ilya: Language Models are Unsupervised Multitask Learners. 2019.

[RJS17] Radford, Alec; Józefowicz, Rafal; Sutskever, Ilya: Learning to Generate Reviews and Discovering Sentiment. CoRR, abs/1704.01444, 2017.

[RPAGNV21] Ramos-Pérez, Eduardo; Alonso-González, Pablo J.; Núñez-Velázquez, José Javier: Multi-Transformer: A New Neural Network-Based Architecture for Forecasting S&P Volatility. Mathematics, 9(15), 2021.

[SBC19] Subramani, Nishant; Bowman, Samuel R.; Cho, Kyunghyun: Can Unconditional Language Models Recover Arbitrary Sentences? CoRR, abs/1907.04944, 2019.

[SS21] Sridhar, Sashank; Sanagavarapu, Sowmya: Multi-Head Self-Attention Transformer for Dogecoin Price Prediction. In: 2021 14th International Conference on Human System Interaction (HSI). pp. 1–6, 2021.

[St20] Strader, Troy J.; Rozycki, John J.; Root, Thomas H.; Huang, Yu-Hsiang: Machine Learning Stock Market Prediction Studies: Review and Research Directions. Journal of International Technology and Information Management, 2020.

[WHM15] Wafi, Ahmed. S.; Hassan, Hassan; Mabrouk, Adel: Fundamental Analysis Models in Financial Markets – Review Study. Procedia Economics and Finance, 30:939–947, 2015. IISES 3rd and 4th Economics and Finance Conference.

[XC18] Xu, Yumo; Cohen, Shay B.: Stock Movement Prediction from Tweets and Historical Prices. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Association for Computational Linguistics, Melbourne, Australia, pp. 1970–1979, July 2018.

[ZAQ17] Zhang, Liheng; Aggarwal, Charu; Qi, Guo-Jun: Stock Price Prediction via Discovering Multi-Frequency Trading Patterns. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '17, Association for Computing Machinery, New York, NY, USA, p. 2141–2149, 2017.

[Zh22] Zhang, Qiuyue; Qin, Chao; Zhang, Yunfeng; Bao, Fangxun; Zhang, Caiming; Liu, Peide: Transformer-based attention network for stock movement prediction. Expert Systems with Applications, 202:117239, 2022.

[ZZL15] Zhang, Xiang; Zhao, Junbo Jake; LeCun, Yann: Character-level Convolutional Networks for Text Classification. CoRR, abs/1509.01626, 2015.