Proceedings of the 19th Conference on Computer Science and Intelligence Systems (FedCSIS) pp. 731–737 DOI: 10.15439/2024F7912 ISSN 2300-5963 ACSIS, Vol. 39

# Predicting Stock Trends Using Common Financial Indicators: A Summary of FedCSIS 2024 Data Science Challenge Held on KnowledgePit.ai Platform

Aleksandar M. Rakićević*, Pavle D. Milošević*, Ivana T. Dragović*, Ana M. Poledica*, Milica M. Zukanović*, Andrzej Janusz†‡, Dominik Ślęzak†§

* University of Belgrade - Faculty of Organizational Sciences, Jove Ilića 154, 11000 Belgrade, Serbia
Email: {aleksandar.rakicevic, pavle.milosevic, ivana.dragovic, ana.poledica, milica.zukanovic}@fon.bg.ac.rs
†Institute of Informatics, University of Warsaw, Banacha 2, 02-097 Warsaw, Poland
Email: {a.janusz, slezak}@mimuw.edu.pl
‡School of Information Systems, Queensland University of Technology, 2 George Street, 4000 Brisbane, Australia
Email: andrzej.janusz@qut.edu.au
§QED Software, Mazowiecka 11/49, 00-052 Warsaw, Poland
Email: dominik.slezak@qedsoftware.com

***Abstract*—Predictive analytics aims to empower finance professionals to make data-driven decisions, anticipate customer behavior, and navigate the complexities of the financial landscape. One of the tasks in this domain is the prediction of stock trend movements. The goal of the FedCSIS 2024 Data Science Challenge was to build such predictive models based on the financial fundamental data. Such models could have a vital role in algorithmic or manual trading, providing trading signals for making decisions about the time and direction of stock trades. We describe the prepared dataset and challenge task. We also summarize the challenge outcomes and provide insights about the most successful machine learning techniques applied.**

***Index Terms*—data science competitions; KnowledgePit.ai platform; stock market data; automatic trading**

## I. Introduction

THE STOCK market is a vital indicator of economic health, reflecting the dynamic interaction of investor sentiment, corporate performance, and macroeconomic trends. The global equity market has almost doubled its value for the last decade. According to statista[1], the total market capitalization of companies listed on stock exchanges worldwide increased from 65.04 trillion US dollars in 2013 to 111 trillion in 2023. As of December 2023, America's region has the largest equity market share, with NYSE and NASDAQ as the largest stock exchange operators, followed by Asia-Pacific and EMEA.

Stock market prediction typically implies forecasting price, trend, and direction of movement of stocks and stock market indexes. It is considered a rather challenging task, being volatile, stochastic, nonlinear, and influenced by a large number of factors. Traditionally, the stock market has been analyzed using technical and fundamental analysis. Further, as a sequence of historical data points (e.g., daily, monthly, quarterly) it is frequently modeled using traditional times series approaches such as statistical ARIMA, exponential smoothing (ES), Facebook's Prophet well-known as the industry standard, and its nonlinear extension NeuralProphet [1]. Since 2015, there has been an exponential growth of research papers investigating ML algorithms for stock market predictions [2]. A comprehensive literature review can be found in [3].

Understanding the dynamics of the stock market and forecasting stock markets have been receiving continuous research attention throughout the previous decades [3]. In most research papers, the US stock market indices the Standard & Poor's 500 (S&P 500), NYSE, NASDAQ and DJIA prevailed as the most common data sources [2]. Among the numerous indices that measure stock market performance, S&P 500 holds a prominent position. S&P 500 comprises 500 of the largest publicly traded companies in the United States and is regarded as a strong indicator of the US economy, but it is also a benchmark for global equity markets. Despite geopolitical tensions and an anticipated recession, the S&P 500 has shown a significant increase in value in the last 5 years. In 2023, the index recorded the highest value, closing at 4,769.83. The S&P 500 provides insights into market trends and risk management strategies. It also helps investors construct portfolios based on S&P 500 constituents. The stock movements impact not only individual investors but also influence institutional strategies, government policies, and international investment flows.

FedCSIS 2024 Data Science Challenge aimed at stock trend forecasting of S&P 500 companies. The scope is reduced to 300 companies that have been part of S&P 500 for the

[1]www.statista.com/statistics/274490/global-value-of-share-holdings-since-2000/

last 10 years. Due to the availability of financial statements and the long history of data, these publicly traded companies have been particularly appealing for machine learning (ML) research and performance analysis [4]. The prevailing markets and the most famous stock market indexes are surveyed in [5]. The task of stock trend prediction is valuable as it guides investment decisions and trading, risk management, and portfolio optimization. It provides useful insights into economic and market conditions. Even though there is a vast amount of publications about ML in stock market forecasting, this topic remains attractive for scientists and financial professionals.

The paper is organized as follows: Section II reviews the literature on ML algorithms applied in stock market forecasting. Section III summarizes the history of data science challenges held at KnowledgePit.ai. Section IV outlines the challenge's objective, gives details of the prediction problem that was solved, and describes our baseline solution. Section V discusses some insights from the post-competition analysis of submitted solutions. SectionVI concludes the paper.

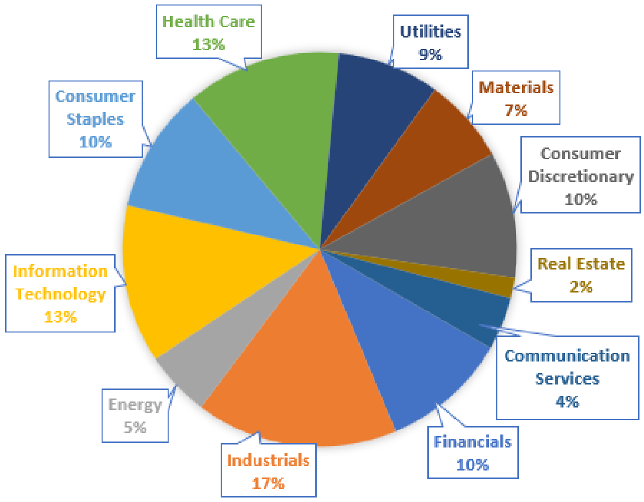


Fig. 1: The share of different industries in the dataset.

## II. Related Literature

In numerous ML studies the research focuses on several directions: classical ML, Ensemble Learning (EL), and Deep Learning (DL) [6], [7], [8]. In the last decade, Support Vector Machines (SVM) and Multi-layer Perceptrons (MLP) have been predominant ML algorithms (approximately 30 %) in stock market forecasting, followed by a group of regression algorithms (linear, logistic, and decision trees), Naïve Bayes (NB), k-nearest Neighbors (kNN), etc. [8].

EL techniques, especially Random Forest (RF), Light Gradient Boosting (LGB), and Extreme Gradient Boosting (XGBoost) have also shown promising forecasting outcomes [6]. A comprehensive evaluation of EL for stock market prediction is given in [9]. Due to their high performance in various data science challenges [10], where they consistently outperformed other algorithms, EL techniques also prevailed throughout our FedCSIS 2024 Data Science Challenge, including particularly the submitted solutions that exceeded the baseline.

In recent years, DL has deserved special attention in stock market prediction. This is because of the availability of financial data with a long history, and the fact that stock forecast is influenced by sentiment described by text [6]. The most common DL algorithms for stock market prediction include deep feed-forward, convolutional, and recurrent neural networks (DNN, CNN, RNN), as well as long short-term memory (LSTM), Gated Recurrent Unit (GRU), and bi-directional LSTM [5], [11]. Among all ML/DL algorithms, based on the recent analysis [12], LSTM is the most preferred model for predicting stock price movements. It is followed by classical ML models (SVM, MLP), next to a bigger family of DL models.

Apart from ML algorithms, it is critical to identify features that affect ML performance. According to a research study on feature selection and extraction for stock market prediction from 2011–2022 [13], the techniques most widely used for this purpose in stock market applications are based on correlations, RF, Principal Component Analysis (PCA), and autoencoders.

## III. History of KnowledgePit.ai challenges

FedCSIS data science challenges have been held on KnowledgePit.ai platform since 2014 [14], [15]. The topics included recognizing firefighters' activities based on sensor readings (2015) [16], [17], predicting seismic activity in coal mines (2016) [18], [19], video game data analytics (2017–2019) [20], [21], [22], predicting network device workloads (2020) [23], [24], predicting costs of freight forwarding contracts (2022) [25], [26], and detecting cyber-attacks on IoT devices (2023) [27], [28]. These challenges were highly successful, with more than 1,600 participating teams and thousands of solutions reflecting state-of-the-art methods in the fields such as feature extraction [19], [29], time series forecasting [24], [30], and EL-based prediction models [21], [31].

KnowledgePit.ai has evolved along with FedCSIS. Over the years, the platform's goals shifted from smaller projects to becoming a host for international data science challenges. The functionalities offered by the platform have also expanded to facilitate post-competition data analysis [10]. The most prestigious events in recent years were those hosted for industry clients such as Security on Demand (currently DeepSeas) [32], Information Builders [33], EMCA Software [23] or, as in the case of this particular challenge, Yettel.Bank.

## IV. Task and Results

The challenge focused on predicting trends in the US stock market based on fundamental financial indicators. A unique, hand-crafted dataset has been prepared for this challenge, encompassing quarterly data from financial statement announcements of 300 companies that constitute the S&P 500 index. Fig. 1 illustrates the sectoral distribution of the selected companies. The dataset covers 10 years, from 2014 to 2023.

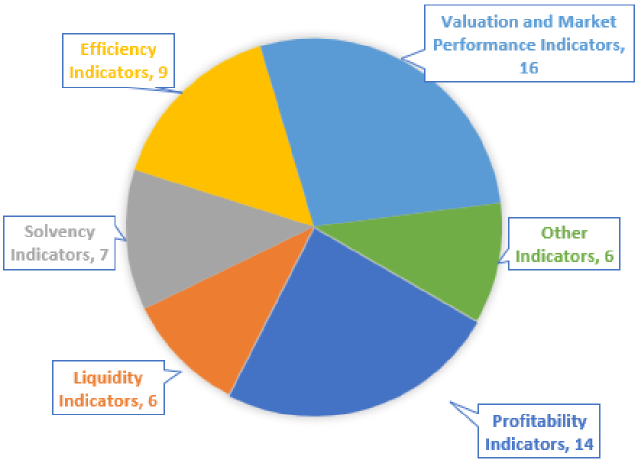


Fig. 2: Types of financial indicators used in the dataset.

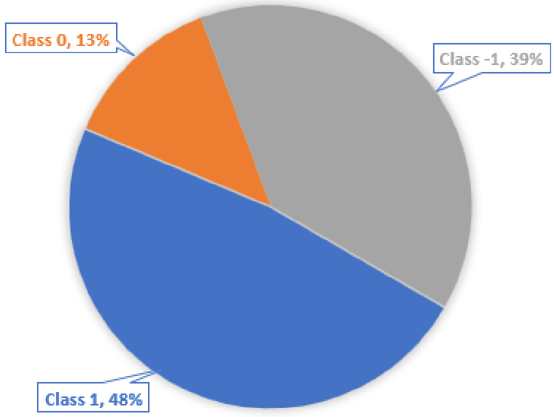


Fig. 3: The class distribution.

Companies were selected based on data availability; those without sufficient data for most of the ten years and those with a high percentage of missing values were omitted. The final dataset contains 10,000 instances. It has been made available to the community to facilitate post-competition research.

Each data instance contains information on the company's sector, values for 58 financial indicators, their 1-year (absolute) changes, target class information (the column 'Class'), and risk-return performance for a period after the announcement (the column 'Perform'). The indicators were chosen based on data availability, literature review, and advice from domain experts. Fig. 2 illustrates the number of indicators within each indicator category. All indicator values are annualized using the Trailing Twelve Months method to neutralize seasonal variation. To prevent participants from gaining an unfair advantage by "looking-ahead", the dataset was anonymized by removing the names of companies and timestamps.

### A. *Data preparation*

The dataset contains two distinct types of missing values with different semantics. NA values indicate that a certain financial indicator does not apply to a company. Empty cells represent conventional missing values (nulls) due to unavailable or missing information. The dataset contains approximately 2.75% NA values and 0.29% null values. The maximum NA percentage for a single attribute is 19.16%. The maximum null value percentage for a single attribute is 5.64%. This presents serious challenges when working with this dataset.

### B. *Evaluation procedure*

In this challenge, participants were asked to solve a three-class classification problem, where:

- **Class 1** means the stock should be bought, as its price will experience a significant uptrend after the announcement of financial statements.
- **Class 0** means no action should be taken, as its price will stay in a sideways trend or it will experience a small but risky (high volatility) uptrend.
- **Class -1** means the stock should be sold, as its price will experience a deteriorating performance.

These classes are obtained based on the Sharpe ratio, a commonly used measure of investing performance [11]:

$$SR_s = \frac{\overline{r}_s - r_f}{\sigma_s} \tag{1}$$

where $\overline{r}_s$ is a mean return of stock $s$, $r_f$ is a risk-free return, and $\sigma_s$ is a standard deviation of excess returns for $s$.

$SR_s$ represents the standardized excess return of an investment. Taking into calculation $\sigma_s$, it penalizes riskier investments. For the purpose of this challenge, the Sharpe ratio was calculated based on stock price movement for a period following the announcement of financial statements announcement until the end of that quarter (approximately two months). The class distribution is given in Fig. 3.

Although the temporal component is very important in stock predicting problems, this dataset does not comprise it explicitly. Namely, the information about a company name and the timestamp are omitted. Further, all instances are then shuffled randomly to mask information about a company and times. Still, the temporal component is included through the 1-year (absolute) change for each of these indicators.

To validate the obtained results, we performed a standard 80%/20% train-test split, ensuring that the class, NA, and null distributions, are maintained in both datasets. To avoid overfitting by submitting multiple solutions to the evaluation system, participants receive information about their preliminary score based on a small fixed subset of the test records after submitting a solution. The final evaluation is conducted after the challenge concludes, using the remaining test data.

The quality of submissions was evaluated using the average error cost measure with the error cost matrix given in Table I. The misclassifications where buying (class 1) is recommended instead of selling (class -1), and vice versa (class -1 instead

TABLE I: Evaluation cost matrix.

| | -1 | 0 | 1 |
|---|---|---|---|
| -1 | 0 | 1 | 2 |
| 0 | 1 | 0 | 1 |
| 1 | 2 | 1 | 0 |

of class 1), are penalized twice as much as misclassifications that resulted in taking no action (class 0). The rationale behind this was to penalize not just the actual loss because of the mistakenly predicted trend, but also the opportunity cost that comes from missed profit opportunity. It is also worth mentioning that using this cost matrix could be regarded as very similar to considering the challenge task as a regression problem with the MAE measure selected as the cost function.

Additionally, a special prize was awarded to the solution that achieved the highest cumulative risk-return performance:

$$CS_s = \sum PC_s \cdot SR_s \tag{2}$$

where $PC_s$ is a predicted class for stock $s$.

### C. The baseline solution

The baseline model was constructed to give participants a reference for the quality of their submissions. The model was trained using XGBoost [34]. The available training dataset was preprocessed to one-hot encode the categorical attribute indicating the company's sector. Two types of missing values were handled by setting them all to NA and adding new binary features to indicate their specific semantics. However, an investigation of the attribute importance for the final model did not reveal their substantial impact on the model's predictions. In the future, it could be worthwhile to extend this analysis using various methods to determine the significance of distinguishing between the considered two types of missing values [35].

The model's hyperparameters were not tuned extensively. The impact of several settings was checked, however, a notably large variance in evaluation results on small random subsets of training data was noticed. The final settings involved changing the default value of the learning rate $\eta$ to $0.001$, and the maximum depth of trees to 6. Moreover, strong regularization was enforced by subsampling features during the construction of trees with the factor $0.5$ and setting the $\lambda$ and $\alpha$ parameters to 10 and 100, respectively. Lastly, due to the observed high instability of predictions, the misclassification risk had to be taken into account. The predictions were adjusted by lowering the classification threshold for class 0, i.e., instead of simply selecting the class with the highest probability, the prediction was set to class 0 if its marginal probability was greater than $0.15$. The final result for such predictions achieved a score of $0.8548$, which gave it 21st position in the final ranking.

### D. Participation statistics

The challenge attracted 194 teams comprising 259 individuals, which makes it one of the most popular in the history of FedCSIS. 77 enrolled teams were deemed active, having submitted at least one solution. Analysis of the IP addresses of team leaders revealed that participants hailed from 28 different countries around the world, with the highest representation observed from Germany (58), Poland (50), Italy (41), Turkey (24), and Serbia (18). The collective efforts of the contestants culminated in nearly 3,000 submitted solutions. Table II presents the final rankings, scores, and submission counts for the top-performing teams, while Fig. 4 portrays the overall daily submission trends throughout the competition. The dynamics of daily submissions witnessed noticeable fluctuations. The most active phase occurred towards the end of May, with high participant engagement remaining consistent in the final week. On the final day, we noted the highest participant activity resulting in over 250 submissions.

TABLE II: Final results. Preliminary and Final score columns show the average prediction costs obtained by top-ranked teams (the lower the better). Cumulative risk-return column (the higher the better) presents the results of the additional evaluation metric computed after the challenge's completion.

| Rank | Team name | Preliminary | Final score | Cumulative risk-return | # subs |
|---|---|---|---|---|---|
| 1 | NxGTR | 0.6584 | 0.7720 | 69.61 | 52 |
| 2 | hieuvq | 0.7376 | 0.8003 | 58.09 | 220 |
| 3 | StockTrends | 0.7228 | 0.8020 | 54.44 | 52 |
| 4 | beamon | 0.6980 | 0.8059 | 46.76 | 90 |
| 5 | Pattern Pioneers | 0.7822 | 0.8076 | 53.51 | 51 |
| 6 | Team1 | 0.7525 | 0.8098 | 47.78 | 20 |
| 7 | Stokastik Heinz | 0.7970 | 0.8187 | 52.14 | 27 |
| 8 | Data_Bombers | 0.7624 | 0.8237 | 47.07 | 53 |
| 9 | The Singleton | 0.7921 | 0.8259 | 45.06 | 5 |
| 10 | No-Name | 0.7426 | 0.8270 | 41.95 | 39 |
| ... | ... | ... | ... | ... | ... |
| 15 | O.W.C.A. | 0.8020 | 0.8432 | 70.23 | 18 |
| ... | ... | ... | ... | ... | ... |

## V. Post-competition Analysis

20 teams exceeded the baseline score, thus their solutions were further analyzed. As in the previous FedCSIS challenges, most teams followed general steps in data science project methodologies [36]: data processing, data cleaning, feature engineering, feature selection and extraction, model construction, and evaluation. Unlike in previous challenges, the utilization of feature selection and extraction did not help participants enhance results substantially. As for data preprocessing and elementary feature engineering, a wide spread of techniques were exercised (normalization and simple forms of data aggregation such as count, sum, prod, std, etc.), but without significant performance score improvement since the dataset was already standardized. Furthermore, top teams employed feature extraction techniques such as PCA, correlation-based filtering, feature importance estimation based on RF, XGBoost, and wrapper-based feature selection algorithms.

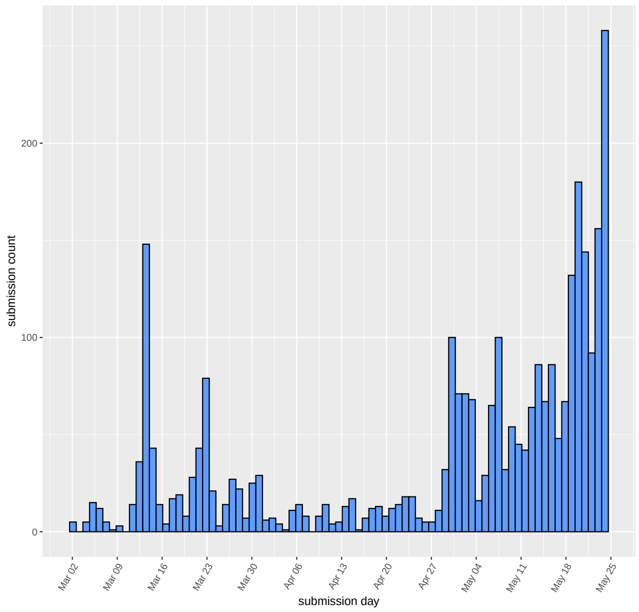


Fig. 4: Daily submissions over the course of the competition.

Most top teams distinguished null and NA missing values using one-hot encoding, setting different tags, or adding dummy columns. Missing values were regularly treated using standard techniques such as imputation with median, mean, zero, 2NN, or omitted if there were too many missing values in the column. Teams treated the challenge problem as classification, regression, or a combination of these two tasks. On the one hand, several teams were applying traditional ML models, e.g. MLP (in 3 solutions) or SVM (1 solution), while one team had tested a novel Kolmogorov-Arnold network. On the other, prevailing algorithms in the best solutions are more advanced, e.g., XGBoost (6 solutions), Gradient Boosting Machines (GBM, 4 solutions), LightGBM (4 solutions), and RF (3 solutions). Notably, several winning teams achieved solid performance scores using the AutoML approach. Despite the rising trend of DL usage in stock trend prediction in the research community, a low number of teams have submitted DL-based solutions. Still, there was an attempt to solve the problem with LSTM, as expected. However, as for this challenge, DL models underperformed compared to EL algorithms.

This year's challenge is unique compared to the previous ones [10], as it included two different measures of success (1, 2). The first one was used to determine the most successful predictors and to define the final rankings. The second was used to determine the most successful investor among participants. One could think that being the most successful predictor would inevitably lead to the best investment results. However, this is not true. The investment result depends not only on how accurately one predicts the trend but also on the magnitude of change of hits (correctly predicted trends) and misses (incorrectly predicted trends). Therefore, it is more significant to correctly predict the trend for stocks that generate returns of higher magnitude (both positive and negative).

Regarding the quality of the predictor, the best three teams were NxGTR, hieuvq and StockTrends. NxGTR had significantly better performance (lower cost function) than the rest of the top 10 teams. It is also worth mentioning that all teams from the top 10 had lower final scores than the preliminary ones, which can indicate an overfitting problem. Regarding the investment performance, the best team was O.W.C.A. as it achieved the highest cumulative risk-return performance. What is interesting regarding the O.W.C.A.'s result is that the team was ranked 15th in predicting stock price movements.

Table II shows a positive correlation between the quality of the predictor and its investment performance. The order of rankings for the best three predictors is the same as for their investment performance. However, in the final results, the best predictor did not achieve the best investment performance. The only plausible explanation of this result could be that the O.W.C.A.'s model was more successful in predicting stocks with more significant upward/downward movements.

## VI. Conclusions and Future Work

The focus of FedCSIS 2024 Data Science Challenge was on predicting trends (uptrend, sideways, downtrend) of stocks constituting the S&P 500 index based on fundamental financial indicators. In addition to providing an overview of the ML algorithms used in stock market prediction, this report paper includes a detailed description of the financial dataset, evaluation procedure, and the baseline model. Furthermore, we explored solutions exceeding the baseline score, including the one achieving the highest cumulative risk-return performance.

With 194 teams and nearly 3,000 solutions, this is one of the most successful challenges at KnowledgePit.ai. Participants employed diverse combinations of data preprocessing techniques. Detailed analysis revealed that the best solutions were mainly obtained using gradient boosting algorithms, such as XGBoost, GBM, and LightGBM. These algorithms outperformed both, classical ML algorithms and the examined DL models.

Due to the high level of attention our challenge received from the ML community and its field of application being consistently a trending topic in finance, the financial dataset will be publicly available for further improvements at us.fon.bg.ac.rs/data/fedcsis2024 and KnowledgePit.ai.

## VII. Acknowledgements

This work is a part of SENSEI project co-financed by EU Smart Growth Operational Programme 2014–2020 under GameINN project POIR.01.02.00-00-0184/17-00.

Also, this study is supported by University of Belgrade – Faculty of Organizational Sciences and Ministry of Science, Technological Development, and Innovation of the Republic of Serbia, institutional funding, Grant no. 200151.

## References

[1] O. Triebe, H. Hewamalage, P. Pilyugina, N. Laptev, C. Bergmeir, and R. Rajagopal, "NeuralProphet: Explainable Forecasting at Scale," *arXiv preprint*, p. arXiv:2111.15397, 2021.

[2] M. M. Kumbure, C. Lohrmann, P. Luukka, and J. Porras, "Machine Learning Techniques and Data for Stock Market Forecasting: A Literature Review," *Expert Systems with Applications*, vol. 197, p. 116659, 2022.

[3] T. Kehinde, F. T. S. Chan, and S. H. Chung, "Scientometric Review and Analysis of Recent Approaches to Stock Market Forecasting: Two Decades Survey," *Expert Systems with Applications*, vol. 213, p. 119299, 2023.

[4] M. Aché, A. Janusz, K. Zbikowski, D. Ślęzak, M. Kryszkiewicz, H. Rybinski, and P. Gawrysiak, "ISMIS 2017 Data Mining Competition: Trading Based on Recommendations," in *Foundations of Intelligent Systems – 23rd International Symposium, ISMIS 2017, Warsaw, Poland, June 26-29, 2017, Proceedings*, ser. Lecture Notes in Computer Science, vol. 10352, 2017, pp. 697–707. [Online]. Available: https://doi.org/10.1007/978-3-319-60438-1_68

[5] W. Jiang, "Applications of Deep Learning in Stock Market Prediction: Recent Progress," *Expert Systems with Applications*, vol. 184, p. 115537, 2021.

[6] G. Sonkavde, D. S. Dharrao, A. M. Bongale, S. T. Deokate, D. Doreswamy, and S. K. Bhat, "Forecasting Stock Market Prices Using Machine Learning and Deep Learning Models: A Systematic Review, Performance Analysis and Discussion of Implications," *International Journal of Financial Studies*, vol. 11, no. 3, p. 94, 2023.

[7] D. Kumar, P. K. Sarangi, and R. Verma, "A Systematic Review of Stock Market Prediction Using Machine Learning and Statistical Techniques," *Materials Today: Proceedings*, vol. 49, pp. 3187–3191, 2022.

[8] N. Rouf, M. B. Malik, T. Arif, S. Sharma, S. Singh, S. Aich, and H.-C. Kim, "Stock Market Prediction Using Machine Learning Techniques: A Decade Survey on Methodologies, Recent Developments, and Future Directions," *Electronics*, vol. 10, no. 21, p. 2717, 2022.

[9] I. K. Nti, A. F. Adekoya, and B. A. Weyori, "A Comprehensive Evaluation of Ensemble Learning for Stock-Market Prediction," *Journal of Big Data*, vol. 7, no. 1, p. 20, 2020.

[10] A. Janusz and D. Ślęzak, "KnowledgePit Meets BrightBox: A Step Toward Insightful Investigation of the Results of Data Science Competitions," in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, vol. 30, 2022, pp. 393–398. [Online]. Available: https://doi.org/10.15439/2022F309

[11] K. Olorunnimbe and H. Viktor, "Deep Learning in the Stock Market – A Systematic Survey of Practice, Backtesting, and Applications," *Artificial Intelligence Review*, vol. 56, no. 3, pp. 2057–2109, 2023. [Online]. Available: https://link.springer.com/article/10.1007/s10462-022-10226-0

[12] P. Balasubramanian, C. P, S. Badarudeen, and H. Sriraman, "A Systematic Literature Survey on Recent Trends in Stock Market Prediction," *PeerJ Computer Science*, vol. 10, p. e1700, 2024.

[13] H. H. Htun, M. Biehl, and N. Petkov, "Survey of Feature Selection and Extraction Techniques for Stock Market Prediction," *Financial Innovation*, vol. 9, no. 1, p. 26, 2023.

[14] A. Janusz, A. Krasuski, S. Stawicki, M. Rosiak, D. Ślęzak, and H. S. Nguyen, "Key Risk Factors for Polish State Fire Service: A Data Mining Competition at Knowledge Pit," in *Proceedings of the 2014 Federated Conference on Computer Science and Information Systems, Warsaw, Poland, September 7-10, 2014*, ser. Annals of Computer Science and Information Systems, vol. 2, 2014, pp. 345–354. [Online]. Available: https://doi.org/10.15439/2014F507

[15] E. Zdravevski, P. Lameski, A. Kulakov, and D. Gjorgjevikj, "Feature Selection and Allocation to Diverse Subsets for Multi-label Learning Problems with Large Datasets," in *Proceedings of the 2014 Federated Conference on Computer Science and Information Systems, Warsaw, Poland, September 7-10, 2014*, ser. Annals of Computer Science and Information Systems, vol. 2, 2014, pp. 387–394. [Online]. Available: https://doi.org/10.15439/2014F500

[16] J. Lasek and M. Gagolewski, "The Winning Solution to the AAIA'15 Data Mining Competition: Tagging Firefighter Activities at a Fire Scene," in *2015 Federated Conference on Computer Science and Information Systems, FedCSIS 2015, Łódź, Poland, September 13-16, 2015*, ser. Annals of Computer Science and Information Systems, vol. 5, 2015, pp. 375–380. [Online]. Available: https://doi.org/10.15439/2015F418

[17] M. Grzegorowski and S. Stawicki, "Window-based Feature Extraction Framework for Multi-sensor Data: A Posture Recognition Case Study," in *2015 Federated Conference on Computer Science and Information Systems, FedCSIS 2015, Łódź, Poland, September 13-16, 2015*, ser. Annals of Computer Science and Information Systems, vol. 5, 2015, pp. 397–405. [Online]. Available: https://doi.org/10.15439/2015F425

[18] A. Janusz, D. Ślęzak, M. Sikora, and Ł. Wróbel, "Predicting Dangerous Seismic Events: AAIA'16 Data Mining Challenge," in *Proceedings of the 2016 Federated Conference on Computer Science and Information Systems, FedCSIS 2016, Gdańsk, Poland, September 11-14, 2016*, ser. Annals of Computer Science and Information Systems, vol. 8, 2016, pp. 205–211. [Online]. Available: https://doi.org/10.15439/2016F560

[19] M. Grzegorowski, "Massively Parallel Feature Extraction Framework Application in Predicting Dangerous Seismic Events," in *Proceedings of the 2016 Federated Conference on Computer Science and Information Systems, FedCSIS 2016, Gdańsk, Poland, September 11-14, 2016*, ser. Annals of Computer Science and Information Systems, vol. 8, 2016, pp. 225–229. [Online]. Available: https://doi.org/10.15439/2016F90

[20] Ł. Grad, "Helping AI to Play Hearthstone Using Neural Networks," in *Proceedings of the 2017 Federated Conference on Computer Science and Information Systems, FedCSIS 2017, Prague, Czech Republic, September 3-6, 2017*, ser. Annals of Computer Science and Information Systems, vol. 11, 2017, pp. 131–134. [Online]. Available: https://doi.org/10.15439/2017F561

[21] Q. H. Vu, D. Ruta, A. Ruta, and L. Cen, "Predicting Win-rates of Hearthstone Decks: Models and Features that Won AAIA'2018 Data Mining Challenge," in *Proceedings of the 2018 Federated Conference on Computer Science and Information Systems, FedCSIS 2018, Poznań, Poland, September 9-12, 2018*, ser. Annals of Computer Science and Information Systems, vol. 15, 2018, pp. 197–200. [Online]. Available: https://doi.org/10.15439/2018F363

[22] A. Janusz, Ł. Grad, and M. Grzegorowski, "Clash Royale Challenge: How to Select Training Decks for Win-rate Prediction," in *Proceedings of the 2019 Federated Conference on Computer Science and Information Systems, FedCSIS 2019, Leipzig, Germany, September 1-4, 2019*, ser. Annals of Computer Science and Information Systems, vol. 18, 2019, pp. 3–6. [Online]. Available: https://doi.org/10.15439/2019F365

[23] A. Janusz, M. Przyborowski, P. Biczyk, and D. Ślęzak, "Network Device Workload Prediction: A Data Mining Challenge at Knowledge Pit," in *Proceedings of the 2020 Federated Conference on Computer Science and Information Systems, FedCSIS 2020, Sofia, Bulgaria, September 6-9, 2020*, ser. Annals of Computer Science and Information Systems, vol. 21, 2020, pp. 77–80. [Online]. Available: https://doi.org/10.15439/2020F159

[24] D. Ruta, L. Cen, and Q. H. Vu, "Deep bi-directional lstm networks for device workload forecasting," in *Proceedings of the 2020 Federated Conference on Computer Science and Information Systems, FedCSIS 2020, Sofia, Bulgaria, September 6-9, 2020*, ser. Annals of Computer Science and Information Systems, vol. 21, 2020, pp. 115–118. [Online]. Available: https://doi.org/10.15439/2020F213

[25] A. Janusz, A. Jamiołkowski, and M. Okulewicz, "Predicting the Costs of Forwarding Contracts: Analysis of Data Mining Competition Results," in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, vol. 30, 2022, pp. 399–402. [Online]. Available: https://doi.org/10.15439/2022F303

[26] E. Kannout, M. Grodzki, and M. Grzegorowski, "Considering Various Aspects of Models' Quality in the ML Pipeline – Application in the Logistics Sector," in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, vol. 30, 2022, pp. 403–412. [Online]. Available: https://doi.org/10.15439/2022F296

[27] M. Czerwiński, M. Michalak, P. Biczyk, B. Adamczyk, D. Iwanicki, I. Kostorz, M. Brzeczek, A. Janusz, M. Hermansa, Ł. Wawrowski, and A. Kozłowski, "Cybersecurity Threat Detection in the Behavior of IoT Devices: Analysis of Data Mining Competition Results," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems, FedCSIS 2023, Warsaw, Poland, September 17-20, 2023*, ser. Annals of Computer Science and Information Systems, vol. 35, 2023, pp. 1289–1293. [Online]. Available: https://doi.org/10.15439/2023F3089

[28] C. Lin, "Tackling Variable-length Sequences with High-cardinality Features in Cyber-attack Detection," in *Proceedings of the 18th Conference on Computer Science and Intelligence Systems, FedCSIS 2023, Warsaw, Poland, September 17-20, 2023*, ser. Annals of

Computer Science and Information Systems, vol. 35, 2023, pp. 1295–1299. [Online]. Available: https://doi.org/10.15439/2023F2385

[29] M. Boullé, "Predicting Dangerous Seismic Events in Coal Mines under Distribution Drift," in *Proceedings of the 2016 Federated Conference on Computer Science and Information Systems, FedCSIS 2016, Gdańsk, Poland, September 11-14, 2016*, ser. Annals of Computer Science and Information Systems, vol. 8, 2016, pp. 221–224. [Online]. Available: https://doi.org/10.15439/2016F21

[30] J. K. Milczek, R. Bogucki, J. Lasek, and M. Tadeusiak, "Early Warning System for Seismic Events in Coal Mines Using Machine Learning," in *Proceedings of the 2016 Federated Conference on Computer Science and Information Systems, FedCSIS 2016, Gdańsk, Poland, September 11-14, 2016*, ser. Annals of Computer Science and Information Systems, vol. 8, 2016, pp. 213–220. [Online]. Available: https://doi.org/10.15439/2016F420

[31] M. Trajanoska, P. Gjorgovski, and E. Zdravevski, "Application of Diversified Ensemble Learning in Real-life Business Problems: The Case of Predicting Costs of Forwarding Contracts," in *Proceedings of the 17th Conference on Computer Science and Intelligence Systems, FedCSIS 2022, Sofia, Bulgaria, September 4-7, 2022*, ser. Annals of Computer Science and Information Systems, vol. 30, 2022, pp. 437–446. [Online]. Available: https://doi.org/10.15439/2022F297

[32] A. Janusz, D. Kałuża, A. Chądzyńska-Krasowska, B. Konarski, J. Holland, and D. Ślęzak, "IEEE BigData 2019 Cup: Suspicious Network Event Recognition," in *2019 IEEE International Conference on Big Data (IEEE BigData), Los Angeles, CA, USA, December 9-12, 2019*, 2019, pp. 5881–5887. [Online]. Available: https://doi.org/10.1109/BigData47090.2019.9005668

[33] A. Janusz, G. Hao, D. Kałuża, T. Li, R. Wojciechowski, and D. Ślęzak, "Predicting Escalations in Customer Support: Analysis of Data Mining Challenge Results," in *2020 IEEE International Conference on Big Data (IEEE BigData 2020), Atlanta, GA, USA, December 10-13, 2020*, 2020, pp. 5519–5526. [Online]. Available: https://doi.org/10.1109/BigData50022.2020.9378024

[34] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2016*, 2016, pp. 785–794. [Online]. Available: https://doi.org/10.1145/2939672.2939785

[35] A. Janusz, D. Ślęzak, S. Stawicki, and K. Stencel, "A Practical Study of Methods for Deriving Insightful Attribute Importance Rankings Using Decision Bireducts," *Information Sciences*, vol. 645, p. 119354, 2023. [Online]. Available: https://doi.org/10.1016/j.ins.2023.119354

[36] C. Schröer, F. Kruse, and J. M. Gómez, "A Systematic Literature Review on Applying CRISP-DM Process Model," in *CENTERIS 2020 – International Conference on ENTERprise Information Systems / ProjMAN 2020 – International Conference on Project MANagement / HCist 2020 – International Conference on Health and Social Care Information Systems and Technologies 2020*, ser. Procedia Computer Science, vol. 181, 2021, pp. 526–534. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877050921002416