A Vision on Sentiment Analysis and other AI Applications on Investments Portfolio Optimization

Anonymous ACL submission

Abstract

Many investors still rely on their emotions as their primary guide for asset allocation, blindly following one or two news sources that may or may not be an actual synthesis of the market, without any mathematical formulation to support them, even if slightly. This work aims to conduct a Systematic Review using Kitchenham's protocol to understand the state of the art regarding the combination of Portfolio Optimization techniques and Artificial Intelligence in the context of trading assets, with a special focus on Sentiment Analysis techniques. This was achieved through the definition of search keywords used in 4 different major research databases and selection through inclusion and exclusion criteria. 384 articles published between 2019–2025 were identified, among which 27 articles were selected that fit all the selection criteria. The research questions address the specific techniques employed for optimization, with the most proliferated being Deep Reinforcement Learning.

1 Introduction

Even though feelings are inherent to the investor's will, they are part of decision-making. Over the years, investors have started to use portfolio selection and optimization tools to make decision-making less emotionally biased. However, others remained confident that news sources can be enlightening for a more lucid decision when it comes to investments. There is also a select group of researchers who study the combination of those ways of investing, contemplating the benefits of both approaches, reducing the investor's emotional bias and quantifying the market's perception in the form of sentiment analysis of news and social media posts.

The main objective of this work is to conduct a Systematic Literature Review (SLR) that answers research questions about Sentiment Analysis and Artificial Intelligence applied to Portfolio Optimization, defines which article databases the study will be conducted and their respective search keywords, and sets inclusion and exclusion criteria to research, select, and extract information from papers. This will verify the state of the art regarding the interaction of sentiment analysis and other artificial intelligence techniques on portfolio optimization context.

Section 2 intends to inform the process adopted for the selection of works, which includes: The definition of the databases used as well as their respective search keywords, search questions, inclusion, exclusion and quality criteria; definition of data extracted from the selected articles. The following section (3) concisely presents the context of this work in relation to portfolio optimization and sentiment analysis. Section 4 will present the observations made after the systematic review is completed, answers to the research questions and more information about the works that fit the quality criterion. Section 5 will be reserved for conclusion, final comments and Section 6 the limitations of this work.

2 Systematic Literature Review

The Systematic Review, according to (Kitchenham, 2004), which this work will use as a model, is a means of identifying, evaluating, and interpreting all available scientific work that is relevant to a research question, topic, or phenomenon of interest. The topics used for the planning protocol of this systematic review were:

Selection of research questions, the most important activity during the protocol according to (Kitchenham, 2004), which determines the questions that the review intends to answer.

Databases are organized collections of peer-reviewed scientific papers, which can be physical or digital, making it possible to conduct detailed searches with keywords pertinent to the subject.

Selection criteria for studies, which determine rules for including or excluding papers from the systematic review. The subsequent execution of the protocol should describe how these criteria will be used (Kitchenham, 2004).

The definition of quality criteria serves as a guide of recommendations for future works, investigates whether the difference concerning these criteria provides an explanation for the differences in the results of the studies, and measures the importance of each study when its results are analyzed (Kitchenham, 2004).

Data extraction stage aims to develop forms to accurately record the information obtained from the analyzed studies (Kitchenham, 2004).

2.1 Literature Review Planning Protocol

The research questions that this study aims to answer are:

- RQ1 – From the articles, which Artificial Intelligence techniques are used in the context of Portfolio Optimization and for which purpose?
- RQ2 – What types of data are used to conduct the experiments in the paper?
- RQ3 – If it makes use of it, which language(s) is (are) the natural language data written?
- RQ4 – If any Sentiment Analysis technique is used, which way does it help with portfolio optimization?

This study was conducted in 4 well-known databases in the scientific scope, namely: "ACM Digital Library"[1], "Elsevier"[2], "IEEE Digital Library"[3] e "Science@Direct" [4].

The exclusion criteria that were considered relevant to the scope of the study and available for international verification were:

- E_1 – Articles not available online;
- E_2 – In a language other than English;
- E_3 – Studies that are not primary;
- E_4 – Does not have as its main scope investment in: stock market, assets, cryptocurrencies or similar;
- E_5 – Does not present Portfolio Optimization techniques in the scope;
- E_6 – Does not present Artificial Intelligence techniques in the scope.

The greatest interest of this study is to select the quality criterion to guide future works and measure the importance of certain studies in relation to others that fit the criteria.

- $CQ1$ – The work applies Sentiment Analysis in conjunction with some Portfolio Optimization technique.

For data extraction, those that detail the selected works or expand on the research questions were selected, they are:

- $D1$ – Artificial Intelligence technique used;
- $D2$ – Data used for Sentiment Analysis and its predicted output, such as type and period used;
- $D3$ – Language of the data used for Sentiment Analysis.

2.2 Execution

The search strings were determined to collect articles that would potentially be suitable for this study, with terms that are common in papers that involve Portfolio Optimization and also Artificial Intelligence. It was not possible to use the same string in all databases, it was necessary to adapt in order to perform a search in the "Science@Direct" database, since the site limits the number of logical operators per search to 8. The strings used are described in Table 1.

Figure 1 reveals the number of articles found and compares it with the number of selected articles. After searching with the strings and the selection of the year of publication from 2019-2025 [5] in the

[1] https://dl.acm.org/
[2] https://www.elsevier.com/
[3] https://ieeexplore.ieee.org/
[4] https://www.sciencedirect.com/

[5] Whenever possible, since some databases only allow selecting year span of 2019-2024.

Base(s) de Dados	String de Busca
ACM, IEEE Explore & Elsevier	("Sentiment Analysis" OR "Natural Language Processing" OR "Large Language Models" OR ("Artificial Intelligence")) AND ("Time Series" OR "Technical Analysis" OR "Fundamental Analysis") AND ("Portfolio Optimization" OR "Markowitz" OR "Monte Carlo Simulation" OR "Liability-Driven Investment" OR "Risk Budgeting" OR "Hierarchical Risk Parity" OR "Robust Optimization" OR "Factor Models" OR "Capital Asset Pricing" OR "Portfolio Management" OR "Modern portfolio theory" OR "Black-Litterman" OR "Mean-Variance")
Science@Direct	("Natural Language Processing" OR "Large Language Models") AND ("Time Series" OR "Technical Analysis" OR "Fundamental Analysis") AND ("Portfolio Optimization" OR "Markowitz" OR "Portfolio Management" OR "Modern portfolio theory")

Table 1: Bases de dados utilizadas na pesquisa e suas respectivas strings de busca.

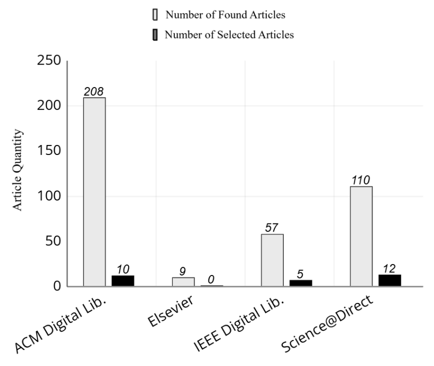

Figure 1: Number of collected and accepted papers in its respective database utilizing the rejection criteria.

respective sites of each database, 384 articles were found, with 1 being duplicated among databases and 356 rejected by criteria E_1-E_6, in that order, resulting in 27 relevant articles for this research, this process is better illustrated in Appendix A.1.

3 Theoretical Background

In the classical approaches to portfolio optimization, what defines an optimized portfolio is built based on the allocation, diversification (Viceira and Wang, 2018) and occasional rebalancing (Carroll et al., 2017) of assets. To achieve the optimum, the techniques implemented often face several difficulties rooted in their often rigorous structures, with their main difficulties being estimating the expected return and risk (Bartram et al., 2021). In the context of investments in assets, some authors devised possible solutions to these challenges with help of Artificial Intelligence techniques (Sutiene et al., 2024), among them, Sentiment Analysis is used under the pretext that the prices observed in the financial market reflect all the information available to the traded assets (Malkiel, 2003), learning patterns through news and feedback from social media communities.

Over the years, developments in Natural Language Processing (NLP), deep-learning, and transfer learning have significantly improved the extraction of sentiments from financial news and other sources (Agaian and Kolm, 2017; Man et al., 2019; Du et al., 2019; Colasanto et al., 2022), being a great potential tool that promises to obtain an optimized portfolio that better suits most financial scenarios of stock markets.

3.1 Portfolio Optimization (PortOp)

Portfolio optimization has as one of the most relevant works the Modern Portfolio Theory (MPT), proposed in 1952 by Harry Markowitz, explaining how to optimize the returns of an investment portfolio for a certain level of risk (Markowitz, 1952). MPT has become popular among stock market investors as a viable strategy that can structure an optimized portfolio, that is, a set of assets and their respective weights chosen so that the risk of the portfolio is minimized and its return, maximized. According to MPT, the expected return of a portfolio is the average value of its assets over a certain period, and the risk associated with a portfolio is calculated by the standard deviation of all its historical prices. For each level of risk, an optimized allocation of assets is formulated to produce the best relationship between expected return and risk. Given a static set of assets, the graphical representation of all portfolios with the best relationship between expected return and risk is called the "Efficient Frontier". The graph formed by it serves as a balance that helps investors decide whether the expected return is worth the risk that the portfolio provides.

As a response to the publication and consequent popularization of Markowitz's work, other authors wrote their own papers criticizing some aspects and assumptions made in his theory, among them are:

- The assumption that markets are efficient, and each stock has its value correctly stipulated by only its price, not taking into account insider trading or other types of market manipulation;
- According to Markowitz's theory, the correlations between two assets remain the same over time;
- Another assumption is that investors are rational and are always inclined towards investing on the lowest risk portfolio.

3.2 Sentiment Analysis

The findings of (Mishev et al., 2020) suggest that models using Sentiment Analysis can be applied to: predict stock market trends and corporate profits, make decisions in bond trading, portfolio management, brand reputation management, as well as fraud detection and regulation (Li et al., 2014; Smailović et al., 2013; Rao and Srivastava, 2012).

(Zhang and Skiena, 2010) mined data from blogs and news to extract specific sentiments from companies, which were then used to rank the best and worst companies for a long-short strategy.

For sentiments in the context of social media, (Bollen et al., 2011) used a fuzzy neural network and demonstrated that the collective public temperament, extracted from Twitter, is correlated to the value of the Dow Jones Index.

(Pagolu et al., 2016) and (Sul et al., 2017) also found that public sentiments in tweets about specific companies were strongly related to stock market movements of those companies in the following days. More recently, it has also been shown that strongly negative tweets from just one individual account with global reach can cause a short-term reduction in the market value of the company mentioned in the post (Brans and Scholtens, 2020). However, in the mentioned approaches, the studies focused mainly on news and specific company forecasts and there was no comparison of prediction performance with any other benchmark model.

Other researchers claim that stock prices have a strong correlation with information about an event that has already occurred, and others have experimented with using both event information and numerical time series data as input data for Sentiment Analysis networks (Yoo et al., 2005).

Sentiment Analysis has proven to be a viable option when acting in predictions of expected return and risk for PortOp, possibly replacing mean and standard deviation, which are metrics used in MPT (Markowitz, 1952) and do not focus sufficient information from a time series to make a prediction. With Sentiment Analysis, it would be possible to more easily identify insider trading, changes in correlations, and reduce human bias in asset trading.

4 Systematic Literature Review Results

4.1 Publications Over the Years

Figure 2 shows the increasing number of publications found on the subjects, which mostly involve exclusively Artificial Intelligence or PortOp, compared to the number of selected articles, which use Artificial Intelligence techniques in the context of PortOp, demonstrating that the application of Artificial Intelligence in conjunction with PortOp techniques is still a recent theme with publication potential.

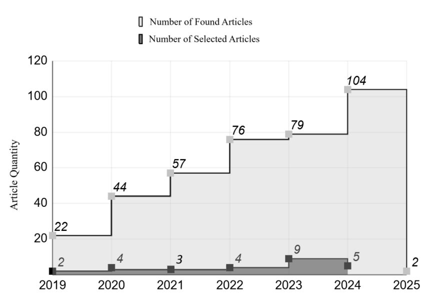

Figure 2: Number of articles found and selected between 2019-2025.

4.2 Distribution of publications among journals and conferences

Out of the 27 selected articles, 15 were published in journals and 12 in conferences. Of the 15 published in journals, 9 of them are unique. Among the journals, the ones with the most publications are: Expert Systems with Applications with 4, followed by Knowledge-Based Systems with 3 and IEEE Access with 2 publications (Figure 3).

The journal *Expert Systems with Applications* [6] has as its scope the technology and application of intelligent systems in the fields of, but not limited to: finance, information retrieval, telecommunications, stock trading, intelligent database management systems, etc.

IEEE Access [7] does not have a specific focus and considers itself multidisciplinary, covering all fields of interest, with the articles highlighted on its platform being from the Artificial Intelligence and telecommunications fields.

The journal *Knowledge-Based Systems* [8] is interdisciplinary in the realm of Artificial Intelligence, as long as the article aims to: support human prediction and decision-making through data science and computing techniques; provide a balanced coverage of theory and practical study in the study field; and to encourage new developments and implementations of models, methods, systems and software tools of knowledge-based intelligence, with applications in business, government, education, engineering, and health.

[6] https://www.sciencedirect.com/journal/expert-systems-with-applications

[7] https://ieeeaccess.ieee.org/

[8] https://www.sciencedirect.com/journal/knowledge-based-systems

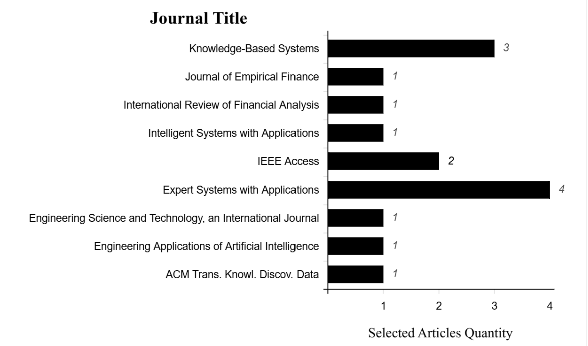

Figure 3: Number of paper per journal.

Out of the selected articles, 10 were published on unique conferences, with ASONAM'18 and ICAIF'23 conferences bearing the highest number of articles, 2 each (Figure 4). ICAIF [9] is an important conference reviewed by specialists that aims to bring together researchers from academia and industry to share challenges and advances in the impact of Artificial Intelligence and Machine Learning in the field of finance. While ASONAM [10]focuses on articles from researchers from various fields, as long as they are involved with Social Network Analysis and Data Mining, and that the study aims to raise discussions with a specific focus on emerging trends, industry needs, and involve real life applications.

4.3 Citation Analysis

Bibliometric analysis, such as the number of citations, must have a purpose, which for this work is: to have a way of comparing scientific articles with their peers (LEWISON, 2020). Table 2 compares the number of citations among the 15 most cited selected articles, and an average of 15.069 citations.

The citation analysis reveals that the most cited article (Koratamaddi et al., 2021), proposes an Adaptive Sentiment-Aware Deep Deterministic Policy Gradients (ASA DDPG) network that learns historical asset price trends as well as market sentiment, and that in terms of annual return on investment and Sharpe index (Sharpe, 1994) the proposed model is compared with techniques such as mean-variance (Markowitz, 1952), minimum variance and the classic DDPG.

Among the most cited selected articles, we can cite (Koratamaddi et al., 2021),(Leow et al., 2021), (Day et al., 2020b),(Xing et al., 2019) and (Muthivhi and van Zyl, 2022) as examples of works that applied Sentiment Analysis.

4.4 Research methods analysis

As for the limit of this paper content being 8 pages, the oldest selected articles published in 2019 (Figure 2) (Xing et al., 2019) and (Day and Lin, 2019) were chosen to be omitted from this section.

4.4.1 Recurrent Neural Networks (RNN)

Recurrent Neural Networks (RNNs) are a class of deep learning models that possess internal memory, enabling them to capture sequential dependencies (Shiri et al., 2023). Unlike traditional neural networks that treat inputs as independent entities, RNNs consider the temporal order of inputs, making them suitable for tasks involving sequential information (Shiri et al., 2023; Abbaspour et al., 2020) due to their recursive structure (Hoffmann et al., 2017).

In (Day et al., 2020a) proposes a modular system for robo-advisors that integrates big data analysis, deep learning methods, and the Black-Litterman model to generate asset weight allocations. The Long Short-Term Memory (LSTM) module performs the function of training the model to capture the long-term tendencies and pattern for making forecasts. Aiming to provide an optimized and personalized portfolio for investors, based on their preferences. Was used daily adjusted closing price data from 01/03/2007 to 12/30/2016.

In (Yun et al., 2020) was proposed a two-stage deep learning architecture, called the Grouped-ETFs Model (GEM) which focus on using a two-stage architecture, which includes a joint cost function for the training of LSTM used in GEM, in order to optimize asset predictions and portfolio diversification. The study uses daily closing prices from MSCI ACWI Index from 05/20/2002 to 06/08/2017.

In (Ong and Herremans, 2023), a multi-task deep learning architecture is used to improve the performance of a time-series momentum strategy and uses LSTM layers as shared layers to encode the input data, to better capture the temporal dependence of financial data. It was used daily adjusted prices from 01/1990 to 12/2020 of 78 US assets.

(Yin et al., 2021) presents the Wealth Flow Model (WFM), which uses recursive neural networks, specifically LSTM to capture the temporal dynamics of wealth flows matrices, alongside with attention mechanisms, improving the accuracy of

[9] https://ai-finance.org/
[10] https://asonam.cpsc.ucalgary.ca/

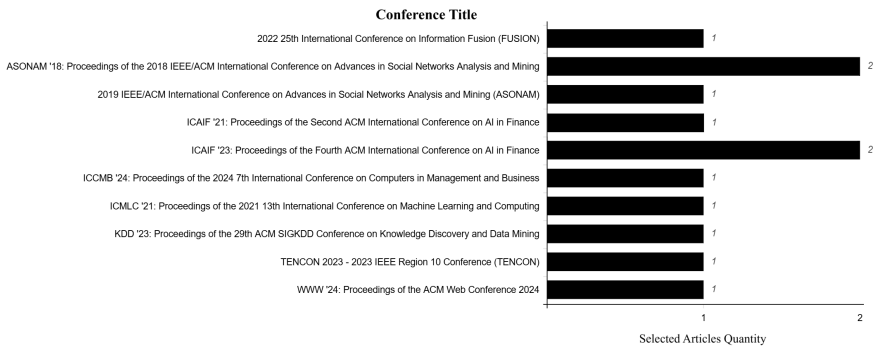

Figure 4: Number of paper per conference.

Título	Ano de Publicação	N° de Citações
Market sentiment-aware deep reinforcement learning approach for stock portfolio allocation (Koratamaddi et al., 2021)	2021	75
AI robo-advisor with big data analytics for financial services (Day et al., 2020a)	2020	54
Portfolio management via two-stage deep learning with a joint cost (Yun et al., 2020)	2020	54
Robo-advisor using genetic algorithm and BERT sentiments from tweets for hybrid portfolio optimisation (Leow et al., 2021)	2021	51
Artificial intelligence for conversational robo-advisor (Day et al., 2020b)	2020	44
Artificial Intelligence for ETF Market Prediction and Portfolio Optimization (Day and Lin, 2019)	2019	27
Harnessing a Hybrid CNN-LSTM Model for Portfolio Performance: A Case Study on Stock Selection and Optimization (Singh et al., 2023)	2023	21
Deep risk model: a deep learning solution for mining latent risk factors to improve covariance matrix estimation (Lin et al., 2021)	2022	19
Growing semantic vines for robust asset allocation (Xing et al., 2019)	2019	13
Efficient Continuous Space Policy Optimization for High-frequency Trading (Han et al., 2023)	2023	11
Generative Machine Learning for Multivariate Equity Returns (Tepelyan and Gopal, 2023)	2023	7
A General Framework on Enhancing Portfolio Management with Reinforcement Learning (Li et al., 2024)	2024	7
Automated stock picking using random forests (Breitung, 2023)	2023	7
Fusion of Sentiment and Asset Price Predictions for Portfolio Optimization (Muthivhi and van Zyl, 2022)	2022	6
Wealth Flow Model: Online Portfolio Selection Based on Learning Wealth Flow Matrices (Yin et al., 2021)	2021	6

Table 2: 15 most cited papers (last accessed 12/18/2024).

predictions. This paper also emphasizes the use of a DRL algorithm to train the WFM model. Using 5 different datasets and as indicators: annualized return (AR), Sharpe ratio (SR), maximum drawdown (MDD), and turnover. From dates that spans from 07/03/1962 to 06/30/2010.

(Lin et al., 2021) proposes the Deep Risk Model (DRM), a deep learning model that uses Gated Recurrent Units (GRU) to model temporal transformations in risk data and represent the evolution of risk over time. The model also uses Graph Attention Networks (GAT) to model cross-sectional transformations. Using daily indicators from 04/30/2009 to 02/10/2020.

In (Muthivhi and van Zyl, 2022), a Semantic Attention Model is used to process textual data to identify investor sentiment, and the results are used as part of the prediction process for a LSTM prediction model to capture temporal dependencies in stock price data, aiming to improve the accuracy of asset price prediction and generate the best capital allocations. Uses weekly sentiment and daily close adjusted stock prices from 01/01/2001 to 12/31/2018.

(Day et al., 2020b) addresses the development of a conversational robo-advisor that uses deep learning models such as LSTM to predict the increase in investment assets and, in combination with Black-Litterman, to find the best portfolio. Additionally, the system uses generative models to produce responses that simulate human conversation to better understand the investor risk profile, using a Seq2Seq architecture. Use adjusted close price and over 400,000 sentences and movie of dialogues as features from 01/03/2007 to 12/30/2016.

4.4.2 Deep Reinforcement Learning (DRL)

DRL, which is an in-depth combination of ANN and RL, has achieved great success in various kinds of complex tasks (Li, 2023) and enables RL techniques to scale to decision-making problems that were previously intractable (Arulkumaran et al., 2017), learning the policy and/or value function with or without the plant model (Spielberg et al., 2017).

In (Wekwete et al., 2023), DRL is used to create an autonomous decision-making agent in Asset Liability Management (ALM) that learns ALM objectives and actions through trial and error, with continuous feedback from the environment. The data used are generated through Monte Carlo simulations and include liability duration and asset portfolio maturity terms. The study focuses on zero-coupon bonds for simplicity, but results are applicable to coupon bonds as well. Used 10,000 simulated datapoints.

(Foo et al., 2023) goal is to develop TQ2CPT, a DRL-based trading agent that incorporates human risk behavior through Cumulative Prospect Theory (CPT) and Truncated Quantile Critics (TQC), improving portfolio allocation and risk management. The article uses a realistic trading environment, implemented using OpenAI Gym and FinRL, and employs market data for training and evaluation. Used daily features from 01/2016 to July/2020.

(Gao et al., 2021) objective is to build an intelligent portfolio management model that can obtain an optimal investment portfolio dynamically, using DDPG for online decision-making and CNN to process complex state data, correcting correlations between data. The article uses opening and closing price data for 5 stocks from the Shanghai Stock Exchange 50 Index and data from a risk-free asset (currency). The data is preprocessed using a time-window standardization method. The data used was SSE 50 stock price from 2017 to 2019. opening, closing, highest and lowest price, trading volume and other features.

(Choi and Kim, 2024) uses a combination of imitation learning and RL, with algorithms such as Soft Actor-Critic (SAC) and Deep Deterministic Policy Gradient (DDPG), as well as neural networks. Rule-based "tutor" models are created, and "student" models are trained via RL to outperform them. The objective is to maximize risk-adjusted investment returns in the long-term while maintaining model "explainability". This work uses historical monthly price data for diverse assets such as stocks, bonds, and commodities, collected from sources like Yahoo Finance and the Federal Reserve Bank of St. Louis (FRED), covering a period of 02/1984 to 05/2017.

(Han et al., 2023) focus is on directly inferring equity weights in the action space to maximize accumulated returns through a deep reinforcement learning-based policy optimization (DRPO) method for high-frequency trading using recurrent neural networks (LSTM) as part of the policy network. The data used was mainly prices including stock indices (Dow Jones 30, SSE 50) and cryptocurrencies (Bitcoin, Ethereum, and Cardano), with different levels of frequency (daily, seconds, minutes).

In (Koratamaddi et al., 2021), DRL is used for portfolio allocation decision-making and sentiment analysis is done using Valence Aware Dictionary and sEntiment Reasoner (VADER) to calculate the polarities of news and tweets, with the aim of reflecting market sentiment more accurately aiding PortOp. The article uses confidence scores (for sentiment analysis) and closing stock prices data from the 30 companies of the Dow Jones Industrial Average (DJIA) index and sentiment data collected from Google News and Twitter in English, from 01/01/2001 to 10/02/2018.

4.4.3 Reinforcement Learning (RL)

RL is the problem faced by an agent that must learn behavior through trial-and-error interactions with a dynamic environment, and it promises a way of programming agents by reward and punishment without needing to specify how the task is to be achieved (Kaelbling et al., 1996).

(Zhang et al., 2024) presents EarnMore, an approach with masking mechanism, self-supervised learning and a re-weighting mechanism that are used to handle changes in customizable stock pools. This method allows RL agents to operate with customizable stock pools without the need for retraining. The paper uses daily data from 10,273 US stocks from Yahoo Finance, with 95 derived technical indicators, and focuses on two financial markets: SP500 and DJ30, in different periods, from 09/26/2007 to 06/26/2022.

The RL algorithms employed by (Li et al., 2024) include Policy Gradient with Actor-Critic (PGAC), Proximal Policy Optimization (PPO), and Evolution Strategies (ES). The objective is to optimize the reallocation of assets continuously, considering transaction costs and short selling restrictions. The data used in the experiments consists of daily prices of European bond futures and commodities (gold, copper, oil), along with US Treasury bond prices and the SP500 index, from 2000 to 2020.

(Lu et al., 2024) employs text-based networks (TBN) to determine the target covariance matrix, which is built based on the similarity between the products/services of different companies and the PPO method is used to define an optimal correlation matrix shrinkage policy that incorporates future hedging needs increasing the Sharpe ratio (Sharpe, 1994) compared to traditional shrinkage techniques. The data includes stock returns from 430 companies listed on US exchanges and English textual data from 10-K database, covering the period from 1995 to 2022.

By integrating RL algorithms such as Deep Q-Network (DQN), Dual Deep Q-Network (DDQN), and Dueling DQN, together with sentiment analysis, (Bhakar et al., 2023) platform helps portfolio managers to make optimal investment decisions. RL helps the agent to learn from market interactions and maximize its reward function and sentiment analysis helped users gauge market mood and make more informed decisions about their assets. Yahoo Finance dataset for SENSEX National Stock Exchange of India (NSE) (open, close, high and low) was used along with news from newsapi.org (sentiment scores, stocks affected by news and analysis), from 2010 to 2023 (It wasn't mentioned the language of textual data, but the example shown was in English).

4.4.4 Other Techniques

(Breitung, 2023) focus on stock selection using Random Forest technique to identify stocks with potential for future returns, based on characteristics like momentum, trend, volatility, volume and daily high, low, open and close daily data, covering a period from 1990 to 2018.

(Begušić and Kostanjčar, 2020) employs clustering to identify latent factors in financial time series. It aims to find risk factors that affect specific assets within clusters, whether by country, asset class, or sector. The study uses weekly asset return data from global financial markets, more specifically 797 NASDAC assets (2005 - 2020) and 621 ETFs (2010 - 2020).

In (Tepelyan and Gopal, 2023) conditional importance-weighted autoencoders (CIWAE) and Normalizing Flows are used for learning the joint distribution of stock returns with applications in synthetic data generation, volatility and risk estimation, and PortOp. The data includes daily stock returns and market factors from 1996 to 2022.

(Mita and Takahashi, 2023) uses a state space model to simulate the impact of different agents on the stock market. It is not a traditional PortOp model, but aims for proactive risk management. The data used was comprised of future prices, trading volume, interest rate, price-to-earnings from the US and Japanese stock markets from 01/01/2000 to 08/30/2022.

In (Kouloumpris et al., 2024) it is proposed Stochastic-Aware Bootstrap Ensemble Ranking (SABER), a learning-to-rank method with an ensemble for stock selection that considers the stochasticity of rankings. The goal is to improve portfolio performance through more precise stock selection and a dynamic bootstrap selector for the number of stocks. The data includes market data, fundamental data (balance sheets, income statements), and macroeconomic indicators. The data 4925 US stocks was organized in quarterly reports with daily values from 2005-Q3 to 2022-Q1, using open, high, low, close and volume of stocks.

In (Du and Tanaka-Ishii, 2022) makes use of stock representations (stock embeddings) based on prices and texts, acquired via neural computation for PortOp named NEws-STock space with Event Distribution (NESTED). The goal is to balance tail risk with variance risk, using texts as a source of information about extreme events. Data includes news articles and historical stock prices from the US, UK, and Chinese markets, and stock representations are generated using natural language processing techniques. The data periodicity for the price data and news articles varied by dataset, with trading days typically spanning daily intervals. The news corpora were collected over different periods, such as from 2000/01 to 2018/12 for some datasets, while others ranged from 2005/01 to 2015/12 or 2006/10 to 2013/11, depending on the source.

(Leow et al., 2021) combines genetic algorithms with sentiment analysis of tweets using the BERT model for PortOp to improve portfolio performance by capturing market conditions via Twitter sentiments (Sentimental All-Weather (SAW) and Sentimental MPT (SMPT) models). The data used includes tweets in English and US stock data both with daily periodicity, from 08/2018 to 12/2019.

5 Conclusão

This paper explored a systematic literature review, covering the main papers of PortOp using Artificial Intelligence techniques, and answering the research questions described in the literature review planning protocol. As a result, it was possible to identify that each proposed approach addresses unique fusion of different techniques for different purposes, so it becomes more difficult to compare each work. During this review, it was noted that AI techniques are gradually being applied for designing new PortOp applications. In some applications, the integration of PortOp and AI leads to positive results with risk reduction and profit increase. However, it can be seen that the articles which integrates PortOp techniques with Sentiment Analysis are among the most cited (Table 2). Additionally, it should be remarked that AI techniques, such as, Sentiment Analysis, DRL, RL, RNN and BERT, have been successfully applied to design PortOp solutions.

6 Limitations

The main limitation of this work was space, some interesting articles and ?? were left out of this work, but they can be inserted again after possible evaluation and 9 pages leniency. Other limitation was that future works were left out our conclusion, so it could give out more space to other important information. Some articles, published on 2018 conferences were also updated to 2020 and this work used their updated dates, rather than the original. Partly because the search engine informed the latter date rather then the original publication date.

References

Saedeh Abbaspour, Faranak Fotouhi, Ali Sedaghatbaf, Hossein Fotouhi, Maryam Vahabi, and Maria Linden. 2020. A comparative analysis of hybrid deep learning models for human activity recognition. *Sensors*, 20(19):5707.

Sarkis Agaian and Petter Kolm. 2017. Financial sentiment analysis using machine learning techniques. *International Journal of Investment Management and Financial Innovations*, 3(1):1–9.

Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. 2017. A brief survey of deep reinforcement learning. *arXiv preprint arXiv:1708.05866*.

Söhnke M Bartram, Jürgen Branke, Giuliano De Rossi, and Mehrshad Motahari. 2021. Machine learning for active portfolio management. *Journal of Financial Data Science*, 3(3):9–30.

Stjepan Begušić and Zvonko Kostanjčar. 2020. Cluster-specific latent factor estimation in high-dimensional financial time series. *IEEE access*, 8:164365–164379.

Anirudh Bhakar, Priyanshu Senabaya Deori, Yash Vardhan Gautam, and KG Srinivasa. 2023. Maximizing returns with reinforcement learning: A paradigm shift in stock market portfolio management. In *TENCON 2023-2023 IEEE Region 10 Conference (TENCON)*, pages 393–398. IEEE.

Johan Bollen, Huina Mao, and Xiaojun Zeng. 2011. Twitter mood predicts the stock market. *Journal of computational science*, 2(1):1–8.

Heleen Brans and Bert Scholtens. 2020. Under his thumb the effect of president donald trump's twitter messages on the us stock market. *PloS one*, 15(3):e0229931.

Christian Breitung. 2023. Automated stock picking using random forests. *Journal of Empirical Finance*, 72:532–556.

Rachael Carroll, Thomas Conlon, John Cotter, and Enrique Salvador. 2017. Asset allocation with correlation: A composite trade-off. *European Journal of Operational Research*, 262(3):1164–1180.

Chanwoo Choi and Juri Kim. 2024. Outperforming the tutor: Expert-infused deep reinforcement learning for dynamic portfolio selection of diverse assets. *Knowledge-Based Systems*, 294:111739.

Francesco Colasanto, Luca Grilli, Domenico Santoro, and Giovanni Villani. 2022. Bert's sentiment score for portfolio optimization: a fine-tuned view in black and litterman model. *Neural Computing and Applications*, 34(20):17507–17521.

Min-Yuh Day, Tun-Kung Cheng, and Jheng-Gang Li. 2020a. Ai robo-advisor with big data analytics for financial services. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 1027–1031. IEEE.

Min-Yuh Day and Jian-Ting Lin. 2019. Artificial intelligence for etf market prediction and portfolio optimization. In *Proceedings of the 2019 IEEE/ACM international conference on advances in social networks analysis and mining*, pages 1026–1033.

Min-Yuh Day, Jian-Ting Lin, and Yuan-Chih Chen. 2020b. Artificial intelligence for conversational robo-advisor. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 1057–1064. IEEE.

Chi-Han Du, Ming-Feng Tsai, and Chuan-Ju Wang. 2019. Beyond word-level to sentence-level sentiment analysis for financial reports. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1562–1566. IEEE.

Xin Du and Kumiko Tanaka-Ishii. 2022. Stock portfolio selection balancing variance and tail risk via stock vector representation acquired from price data and texts. *Knowledge-Based Systems*, 249:108917.

Marcus Jun Rong Foo, Nixie S Lesmana, and Chi Seng Pun. 2023. Drl trading with cpt actor and truncated quantile critics. In *Proceedings of the Fourth ACM International Conference on AI in Finance*, pages 574–582.

Ni Gao, Yiyue He, Yuming Jiao, and Zhuo Chang. 2021. Online optimal investment portfolio model based on deep reinforcement learning. In *Proceedings of the 2021 13th International Conference on Machine Learning and Computing*, pages 14–20.

Li Han, Nan Ding, Guoxuan Wang, Dawei Cheng, and Yuqi Liang. 2023. Efficient continuous space policy optimization for high-frequency trading. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4112–4122.

Javier Hoffmann, Osvaldo Navarro, Florian Kastner, Benedikt Janßen, and M Hubner. 2017. A survey on cnn and rnn implementations. In *PESARO 2017: The Seventh International Conference on Performance, Safety and Robustness in Complex Systems and Applications*, volume 3.

Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285.

Barbara Kitchenham. 2004. Procedures for performing systematic reviews. *Keele, UK, Keele University*, 33(2004):1–26.

Prahlad Koratamaddi, Karan Wadhwani, Mridul Gupta, and Sriram G Sanjeevi. 2021. Market sentiment-aware deep reinforcement learning approach for stock portfolio allocation. *Engineering Science and Technology, an International Journal*, 24(4):848–859.

Eleftherios Kouloumpris, Konstantinos Moutsianas, and Ioannis Vlahavas. 2024. Saber: Stochastic-aware bootstrap ensemble ranking for portfolio management. *Expert Systems with Applications*, 249:123637.

Edmund Kwong Wei Leow, Binh P Nguyen, and Matthew Chin Heng Chua. 2021. Robo-advisor using genetic algorithm and bert sentiments from tweets for hybrid portfolio optimisation. *Expert Systems with Applications*, 179:115060.

GRANT LEWISON. 2020. Preparation of bibliometrics papers.

Shengbo Eben Li. 2023. *Deep Reinforcement Learning*, pages 365–402. Springer Nature Singapore, Singapore.

Xiaodong Li, Haoran Xie, Li Chen, Jianping Wang, and Xiaotie Deng. 2014. News impact on stock price return via sentiment analysis. *Knowledge-Based Systems*, 69:14–23.

Yinheng Li, Junhao Wang, and Yijie Cao. 2024. A general framework on enhancing portfolio management with reinforcement learning. In *Proceedings of the 2024 7th International Conference on Computers in Management and Business*, pages 119–125.

Hengxu Lin, Dong Zhou, Weiqing Liu, and Jiang Bian. 2021. Deep risk model: a deep learning solution for mining latent risk factors to improve covariance matrix estimation. In *Proceedings of the Second ACM International Conference on AI in Finance*, pages 1–8.

Cheng Lu, Papa Momar Ndiaye, and Majeed Simaan. 2024. Improved estimation of the correlation matrix using reinforcement learning and text-based networks. *International Review of Financial Analysis*, 96:103572.

Burton G Malkiel. 2003. The efficient market hypothesis and its critics. *Journal of economic perspectives*, 17(1):59–82.

Xiliu Man, Tong Luo, and Jianwu Lin. 2019. Financial sentiment analysis (fsa): A survey. In *2019 IEEE international conference on industrial cyber physical systems (ICPS)*, pages 617–622. IEEE.

Harry Markowitz. 1952. Modern portfolio theory. *Journal of Finance*, 7(11):77–91.

Kostadin Mishev, Ana Gjorgjevikj, Irena Vodenska, Lubomir T Chitkushev, and Dimitar Trajanov. 2020. Evaluation of sentiment analysis in finance: from lexicons to transformers. *IEEE access*, 8:131662–131682.

Daiya Mita and Akihiko Takahashi. 2023. Multi-agent model based proactive risk management for equity investment. *Engineering Applications of Artificial Intelligence*, 125:106701.

Mufhumudzi Muthivhi and Terence L van Zyl. 2022. Fusion of sentiment and asset price predictions for portfolio optimization. In *2022 25th International Conference on Information Fusion (FUSION)*, pages 1–8. IEEE.

Joel Ong and Dorien Herremans. 2023. Constructing time-series momentum portfolios with deep multi-task learning. *Expert Systems with Applications*, 230:120587.

Venkata Sasank Pagolu, Kamal Nayan Reddy, Ganapati Panda, and Babita Majhi. 2016. Sentiment analysis of twitter data for predicting stock market movements. In *2016 international conference on signal processing, communication, power and embedded system (SCOPES)*, pages 1345–1350. IEEE.

Tushar Rao and Saket Srivastava. 2012. Analyzing stock market movements using twitter sentiment analysis.

William F Sharpe. 1994. The sharpe ratio. *Journal of portfolio management*, 21(1):49–58.

Farhad Mortezapour Shiri, Thinagaran Perumal, Norwati Mustapha, and Raihani Mohamed. 2023. A comprehensive overview and comparative analysis on deep learning models: Cnn, rnn, lstm, gru. *arXiv preprint arXiv:2305.17473*.

Priya Singh, Manoj Jha, Mohamed Sharaf, Mohammed A El-Meligy, and Thippa Reddy Gadekallu. 2023. Harnessing a hybrid cnn-lstm model for portfolio performance: A case study on stock selection and optimization. *Ieee Access*, 11:104000–104015.

Jasmina Smailović, Miha Grčar, Nada Lavrač, and Martin Žnidaršič. 2013. Predictive sentiment analysis of tweets: A stock market application. In *Human-Computer Interaction and Knowledge Discovery in Complex, Unstructured, Big Data: Third International Workshop, HCI-KDD 2013, Held at SouthCHI*

2013, Maribor, Slovenia, July 1-3, 2013. Proceedings, pages 77–88. Springer.

SPK Spielberg, RB Gopaluni, and PD Loewen. 2017. Deep reinforcement learning approaches for process control. In *2017 6th international symposium on advanced control of industrial processes (AdCONIP)*, pages 201–206. IEEE.

Hong Kee Sul, Alan R Dennis, and Lingyao Yuan. 2017. Trading on twitter: Using social media sentiment to predict stock returns. *Decision Sciences*, 48(3):454–488.

Kristina Sutiene, Peter Schwendner, Ciprian Sipos, Luis Lorenzo, Miroslav Mirchev, Petre Lameski, Audrius Kabasinskas, Chemseddine Tidjani, Belma Ozturkkal, and Jurgita Cerneviciene. 2024. Enhancing portfolio management using artificial intelligence: literature review. *Frontiers in Artificial Intelligence*, 7:1371502.

Ruslan Tepelyan and Achintya Gopal. 2023. Generative machine learning for multivariate equity returns. In *Proceedings of the Fourth ACM International Conference on AI in Finance*, pages 159–166.

Luis M Viceira and Zixuan Kevin Wang. 2018. Global portfolio diversification for long-horizon investors. Technical report, National Bureau of Economic Research.

Takura Asael Wekwete, Rodwell Kufakunesu, and Gusti van Zyl. 2023. Application of deep reinforcement learning in asset liability management. *Intelligent Systems with Applications*, 20:200286.

Frank Z Xing, Erik Cambria, and Roy E Welsch. 2019. Growing semantic vines for robust asset allocation. *Knowledge-Based Systems*, 165:297–305.

Jianfei Yin, Ruili Wang, Yeqing Guo, Yizhe Bai, Shunda Ju, Weili Liu, and Joshua Zhexue Huang. 2021. Wealth flow model: Online portfolio selection based on learning wealth flow matrices. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 16(2):1–27.

Paul D Yoo, Maria H Kim, and Tony Jan. 2005. Machine learning techniques and use of event information for stock market prediction: A survey and evaluation. In *International conference on computational intelligence for modelling, control and automation and international conference on intelligent agents, web technologies and internet commerce (CIMCA-IAWTIC'06)*, volume 2, pages 835–841. IEEE.

Hyungbin Yun, Minhyeok Lee, Yeong Seon Kang, and Junhee Seok. 2020. Portfolio management via two-stage deep learning with a joint cost. *Expert Systems with Applications*, 143:113041.

Wenbin Zhang and Steven Skiena. 2010. Trading strategies to exploit blog and news sentiment. In *Proceedings of the international AAAI conference on web and social media*, volume 4, pages 375–378.

Wentao Zhang, Yilei Zhao, Shuo Sun, Jie Ying, Yonggang Xie, Zitao Song, Xinrun Wang, and Bo An. 2024. Reinforcement learning with maskable stock representation for portfolio management in customizable stock pools. In *Proceedings of the ACM on Web Conference 2024*, pages 187–198.

A Appendix

A.1 Article filtering method detailing the order of exclusion criteria by chronological phases