A Multimodal and Sentiment-Based Trading System for Financial Portfolio Optimisation

Zhenglong Li, Vincent Tam, Kwan L. Yeung
Department of Electrical and Electronic Engineering, The University of Hong Kong
E-mails: lzlong@connect.hku.hk, {vtam, kyeung}@eee.hku.hk

***Abstract*—The keys to accurately predicting future price trends and correlations between assets in a portfolio are not only learning from historical price data, but also considering the current market sentiment among investors. Yet there is a lot of news published by different media everyday, some of which may describe irrelevant or even contradictory information, providing noise signals to estimate market sentiment. The existing portfolio management frameworks usually treat all news in the same way to extract sentiment intention, which unavoidably results in biased trading decisions. Accordingly, a Multimodal and Sentiment-based Adaptive trading framework, namely the MuSA, is proposed in this paper where the sentiment analyser collects relevant news from different media and evaluates the sentiment behind the news by using large language models. Afterwards, an entropy-based confidence learning mechanism measures the confidence of the news by checking the reliability of information sources to reduce the effects of irrelevant messages. Meanwhile, the price analyser learns the future trends of assets from historical price data. Finally, combining sentiment information and price estimation generates a new portfolio to adapt to the ever-changing financial markets. The empirical results in two real-world datasets clearly reveal the proposed framework can help extract useful sentiment information for achieving higher returns and lower risks.**

***Index Terms*—Portfolio Optimisation, Financial Sentiment Analysis, Multimodal data, Information fusion**

I. INTRODUCTION

In contrast to traditional finance that mainly relies on subjective personal judgement, quantitative finance [1] is a promising research topic utilizing intelligent approaches to help users handle various practical problems in the financial area. The previous studies [2] have proven the success of quantitative finance in many applications such as credit risk assessment, fraud detection, and order optimisation. However, portfolio optimisation (PO), as one of the fundamental financial problems to allocate capital to different assets for maximising the total returns of portfolios while diversifying investment risks, is still a challenge for investors to balance the opposite objectives. Originally inspired by the modern portfolio theory [3], many researchers explored financial-based algorithmic methods [4] where some valuable trend patterns of price are concluded to predict the future movements of assets. Among these approaches, directional change (DC) [5] brings a new perspective to record asset price movement instead of using the conventional time series with fixed time intervals, which has been successfully applied in foreign exchange markets. Recently, more efforts in PO [6], [7] have been made in exploring the benefits of applying deep reinforcement learning (DRL) approaches as virtual traders to make trading decisions by learning asset movements and correlations from historical prices. However, except for using historical prices to deduce the momentum of asset prices, price changes are significantly affected by the current market sentiment due to sudden events like regional conflicts and natural disasters.

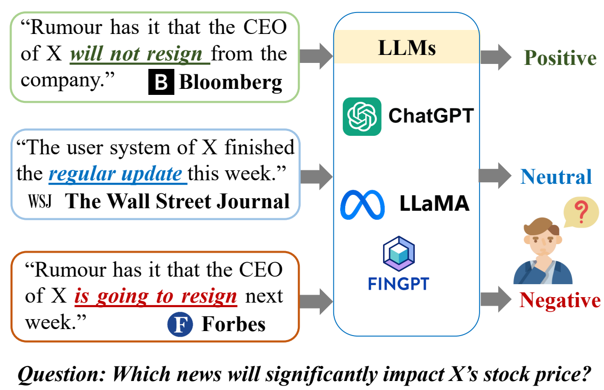

Fig. 1. The Challenge of Financial Sentiment Analysis from Multiple Information Sources

Generally speaking, investors' decisions are influenced by the current market sentiment, immediately resulting in large fluctuations of asset prices. Effectively analyzing sentiment in financial markets is the key to reducing investment risk exposures. In the past, it was subjective to evaluate the impact of news by humans as they judged the news in terms of their preferences and past experiences. With the rapid development of intelligent techniques, there are lots of advanced large language models (LLMs) like FinBERT [8], GPT [9], and LLaMA [10] that demonstrate remarkable ability on quantitatively analyzing the sentiment in a sentence or an article. Furthermore, some researchers [11], [12] attempt to combine price data and sentiment information to generate trading signals in portfolio optimization. Nevertheless, none of them studies the influences of news from different information sources before utilizing them in decision-making, which may include inaccurate messages or even fake news from unreliable media. Therefore, as shown in Fig. 1, recognizing reliable information sources and the most relevant news to the price changes will drive trading systems to generate more rational and profitable trading strategies.

To address the above concerns, a **Mu**ltimodal and

Sentiment-based **A**daptive framework for portfolio optimisation namely the **MuSA** is proposed in this paper in which three intelligent modules are designed to cooperatively analyze and fuse the multimodal data such that the adjusted portfolio can balance the overall returns and potential sudden risks over the trading period. Firstly, the pre-trained LLM in the sentiment analyser respectively evaluates the implied sentiment of financial texts that are recently collected from different news media. Furthermore, to reduce the impacts of irrelevant and inaccurate information from unreliable sources, an entropy-based confidence learning mechanism is introduced in this work to measure the reliability of each involved source by learning the correlations between the price trends of assets and the sentiment provided by information sources. An information source where the sentiment tendency of the provided financial texts to an asset is inconsistent with the price trends of the asset is regarded as an unreliable source. Subsequently, the impact of sentiment from those recognised unreliable sources is reduced in decision-making. In addition to the sentiment data, the price analyser, implemented by algorithmic approaches or DRL-based methods, tries to extract useful underlying patterns from price data to generate profitable trading signals (i.e., portfolio weights). Lastly, the portfolio aggregator of the proposed framework fuses the sentiment information of different information sources based on the confidence matrix, followed by the integration with the trading signals to produce a newly generated portfolio for adapting to the highly volatile financial markets. The main contributions of the proposed framework are summarised as follows.

1) The MuSA framework can recognize unreliable information sources and extract influential news for effectively estimating financial market sentiment, avoiding generating biased trading decisions due to irrelevant news.
2) As a plug-and-play tool, the MuSA framework can be conveniently integrated with other portfolio optimization methods to help analyze sentiment information.
3) The empirical results reveal a remarkable enhancement of the MuSA framework when considering the confidence of news from different media.

II. THE PRELIMINARIES

A. Portfolio Optimisation in Quantitative Finance

Even after being studied for a few decades, portfolio optimisation is a very popular yet challenging task in the financial area. As classical trend-tracking strategies, [13]–[16] attempt to capture the momentum or mean-reversion patterns of assets for obtaining arbitrage opportunities. Recently, there are many interesting studies [7], [17], [18] on applying deep reinforcement learning to optimise portfolios where the DRL-based agents simulate trading and rectify decisions by the interaction with the environment to approximate optimal trading policies for maximizing overall returns. They attempt to utilize different types of neural network architectures like convolution-based, recurrent-based, or attention-based to extract temporal as well as spatial information from historical prices. However, those models rarely consider current market sentiment in the trading. To fill the gap, [11], [12] tries to encode news as embedding vectors and fuse together with price embedding vectors to eventually produce portfolio weights. Yet their sentiment analysis modules trained by the limited data from a single information source may be restricted to the collection of more useful financial texts. Among the advanced DRL approaches, a model-free and actor-critic based approach namely the twin delayed deep deterministic policy gradient (TD3) [17], [19] has demonstrated a remarkable ability to enhance the policy of trading actions for maximising the total profits at the whole trading period. Yet the current prices of stock are highly sensitive to the news of related companies, most DRL approaches following the distribution learnt by the historical data may fail to immediately absorb the impactful information for quickly adapting to the currently changed market. This may increase the potential risks and result in great losses within a short period of time due to the possibly sudden crises. To further discuss portfolio optimisation, there are some essential financial concepts to be introduced as below.

Definition II.1. (Annualized Return) Given a portfolio involving N assets, the weight $w_{i,t}$ of i^{th} asset and the close price $p_{i,t}$ of i^{th} asset at time t, the portfolio value at time t can be denoted as $C_t = \sum_{i=1}^{N} w_{i,t}p_{i,t}$. Then, the annualized return (AR) is defined as $AR = ((\frac{C_T}{C_1})^{\frac{T_y}{T}} - 1) \times 100\%$, where T_y is the number of trading days per year and T is the total number of trading days at the whole trading period. Furthermore, to fulfil the regulations of trading, for any $w_{i,t} \in \mathbf{w}_t$, all portfolios are required to follow the long-only constraint (i.e., $w_{i,t} \geq 0$) and capital budget constraint (i.e., $\sum_{i=1}^{N} w_{i,t} = 1$). $\mathbf{w}_t \in \mathbb{R}^N$ is the weight vector of a portfolio at time t.

Definition II.2. (Maximum Drawdown) The maximum drawdown (MDD) is used to measure the risk of the worst case during the trading period. Assume that $t1 < t2 \leq t$, the MDD is defined as $MDD_t = \text{Max}(\frac{C_{t1}-C_{t2}}{C_{t1}})$.

Definition II.3. (Sharpe Ratio) The Sharpe Ratio (SR) as one of the most common indicators to evaluate the comprehensive performance of a portfolio is denoted as SR $= \frac{AR-r_f}{Vol}$. where AR is the annualized returns, r_f is the risk-free rate to measure the rates of returns on zero-risk assets. The yield of a 10-year treasury bond is usually regarded as the zero-risk asset. The annualized volatility Vol of the trading strategy is $Vol = \sqrt{\frac{T_{yr}}{T-1}\sum_{t=1}^{T}\left(\mathbf{r}_t^\top \mathbf{a}_{t-1} - \overline{\mathbf{r}_t^\top \mathbf{a}_{t-1}}\right)^2}$, where $\overline{\mathbf{r}_t^\top \mathbf{a}_{t-1}} = \frac{1}{T}\sum_{t=1}^{T}\mathbf{r}_t^\top \mathbf{a}_{t-1}$ is the average daily returns of a trading strategy.

B. Directional Change

Instead of exploring time series-based patterns in the most of algorithmic approaches, directional change (DC) [5] is an event-based algorithmic approach that captures the significant movement of asset prices while filtering the minor fluctuations at a specific period. According to the definition of DC, there are two basic events namely the downtrend event and uptrend event act as features to indicate the current state of an asset. For instance, as the pseudo-code shown in Algorithm 1, an asset is confirmed to fall into a downward run when $p_t \leq p^h_{t-1} \times (1 - \Delta p_{dc})$, whereas it comes back to an upward run when $p_t \geq p^l_{t-1} \times (1 + \Delta p_{dc})$. p^h_{t-1} and p^l_{t-1} are the highest and lowest close price in the current run until time $t-1$. Δp_{dc} is the DC threshold to adjust what extent the changes of prices will be captured. There are many DC-based approaches [20] integrating with meta-heuristic algorithms that have been successfully applied in the foreign exchange markets. Yet the most algorithmic approaches do not take into account the multimodal data like news or company reports when deducing trading signals.

Algorithm 1 The Pseudo-code for Generating DC Events

Input: Δp_{dc} as the DC threshold, $p_{i,t}$ as the close price of i^{th} asset at time t, $p^h_{i,t-1}$ and $p^l_{i,t-1}$ as the highest and lowest close price in the current run until time $t-1$, and $event_{i,t-1}$ as the DC event at time $t-1$.
Output: The current DC event $event_{i,t}$, the highest close price $p^h_{i,t}$ in the current run, and the lowest close price $p^l_{i,t}$ in the current run.
if $event_{i,t-1}$ is $UptrendEvent$ **then**
  if $p_{i,t} \leq p^h_{i,t-1} \times (1 - \Delta p_{dc})$ **then**
    $event_{i,t} \leftarrow DowntrendEvent$
    $p^l_{i,t} \leftarrow p_{i,t}$
    $p^h_{i,t} \leftarrow p^h_{i,t-1}$
  else
    $event_{i,t} \leftarrow UptrendEvent$
    $p^l_{i,t} \leftarrow p^l_{i,t-1}$
    if $p^h_{i,t-1} < p_{i,t}$ **then**
      $p^h_{i,t} \leftarrow p_{i,t}$
    else
      $p^h_{i,t} \leftarrow p^h_{i,t-1}$
    end if
  end if
else
  if $p_{i,t} \geq p^l_{i,t-1} \times (1 + \Delta p_{dc})$ **then**
    $event_{i,t} \leftarrow UptrendEvent$
    $p^l_{i,t} \leftarrow p^l_{i,t-1}$
    $p^h_{i,t} \leftarrow p_{i,t}$
  else
    $event_{i,t} \leftarrow DowntrendEvent$
    $p^h_{i,t} \leftarrow p^h_{i,t-1}$
    if $p^l_{i,t-1} > p_{i,t}$ **then**
      $p^l_{i,t} \leftarrow p_{i,t}$
    else
      $p^l_{i,t} \leftarrow p^l_{i,t-1}$
    end if
  end if
end if
return $event_{i,t}$, $p^h_{i,t}$ and $p^l_{i,t}$.

III. METHODOLOGY

A. The Overview of the Framework

To fill up the pitfall of the existing PO approaches that neglect the market sentiment reflected by financial texts, a **Mu**ltimodal and **S**entiment-based **A**daptive framework for portfolio optimisation namely the **MuSA** is proposed in which three key modules, including the sentiment analyser, price analyser, and portfolio aggregator, cooperatively handle the different types of financial data and generate rational trading decisions for dynamically balancing the long-term profits and potential sudden risks under the volatile financial markets.

Algorithm 2 The Inference Procedure of the MuSA

Input: T as the number of trading days, the pre-trained large-language models, the trained price analyser, and the settings of trading.
Output: The annualized return (AR), maximum drawdown (MDD), and Sharpe ratio (SR).
for $t = 1$ to T **do**
  Collect and pre-process the latest financial text data from multiple information sources.
  Collect the price data $\mathbf{P}_t$ of assets in a portfolio from stock exchanges.
  Invoke large-language models to deduce the sentiment $\mathbf{S}_t$ of the latest text data.
  Update the confidence matrix $\mathbf{R}_t$.
  Invoke the price analyser to generate the weight vector $\mathbf{w}_{A,t}$ of a portfolio based on the price data $\mathbf{P}_t$.
  Invoke the portfolio aggregator to generate the final weight vector $\mathbf{w}_t$ by integrating the suggested weight vector $\mathbf{w}_{A,t}$ and the sentiment scores $\mathbf{S}_t$ based on the confidence matrix $\mathbf{R}_t$.
  Execute the portfolio order with $\mathbf{w}_t$.
end for
Calculate the AR, MDD, and SR.
return the performance of AR, MDD, and SR.

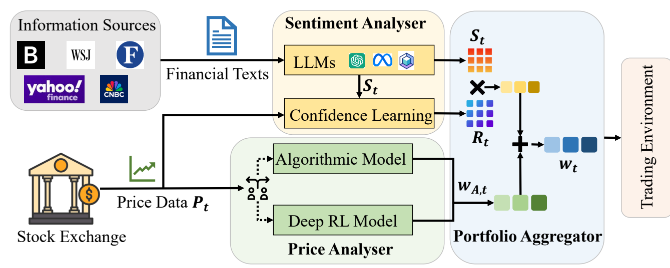

Fig. 2. The System Architecture of the Proposed MuSA Framework

Fig. 2 demonstrates the overall architecture of the proposed MuSA framework where the sentiment analyser evaluates the sentiment of financial texts in real time by LLMs and continuously assesses the confidence of multiple information sources while the price analyser is based on an algorithmic-based or a DRL-based approach to capture the complex patterns from historical prices. Specifically, as clearly depicted in Algorithm 2 where the pseudo-code of the inference procedure of the MuSA framework is given, all financial texts from different information sources are firstly fed into the LLMs to collect the matrix of current sentiment scores $\mathbf{S}_t \in \mathbb{R}^{N \times J}$, where N is the portfolio scale and J is the number of information sources. Subsequently, the confidence matrix $\mathbf{R}_t \in \mathbb{R}^{N \times J}$ is updated by the historical sentiment scores and price changes in previous days. Besides, the price-based trading signal $\mathbf{w}_{A,t}$ is generated by the price analyser in terms of the price momentum and the correlation between assets. Based on the confidence matrix $\mathbf{R}_t$, the portfolio aggregator calculates the final sentiment score at time t, followed by the integration with the trading signal $\mathbf{w}_{A,t}$ to produce the final trading decision $\mathbf{w}_t \in \mathbb{R}^N$, such that the MuSA can manage potential investment risks ahead of the large fluctuation of asset prices.

B. Sentiment Analyser

To evaluate the sentiment of texts with financial domain knowledge, FinBert [8], as the example of pre-trained LLMs, is applied to deduce the probability of different sentiment categories including positive, neutral, and negative. As different information sources may provide various news that possibly imply opposite feelings, some of the texts can be the noise to adversely affect the final trading decisions. To address the concern, an entropy-based learning mechanism is proposed to evaluate the reliability of information sources through measuring the correlations between the sentiment tendency of news from the same media and the price change of relevant assets in the past few days. When the sentiment score of j^{th}information source for i^{th} asset is positively correlated with the price changes of i^{th} asset, a higher confidence score is given for the news distributed by j^{th} media. For $\forall r_{i,j,t} \in \mathbf{R}_t$, the confidence coefficient $r_{i,j,t}$ is calculated as follows.

$$r_{i,j,t} = 1 - \frac{1}{M} \sum_{m=t-1}^{t-M} \sum_{k=1}^{K} (-\hat{p}_{i,m,k} \log q_{i,j,m,k}) \tag{1}$$

where $\hat{p}_{i,m,k}$ is the one-hot encoding of i^{th} asset price change at time m, $q_{i,j,m,k}$ is the probability of k^{th} sentiment class of i^{th} asset in j^{th} information source at m, M is the observation period, and K is the number of sentiment categories.

C. Price Analyser

Capturing the useful change patterns of asset prices and the correlations among assets in a portfolio is the key to optimising capital allocation. To investigate the improvement when integrating MuSA with different approaches, two representative methods of algorithmic approaches and DRL-based approaches are implemented as the example of price analyser in this work, including the revised DC method and the popular DRL algorithm TD3 [19]. As the previous DC studies focus on a single asset, this work extends it to the moving average DC for generating the weight vector of N assets in a portfolio. Assume that $\mathbf{x}_m^{dc} \in \mathbb{R}^N$ is the DC vector of N assets in which the DC signal can be set to 1 for uptrend events or 0 for downtrend events at time m, T_{dc} is the observation period, and $\text{Norm}(\cdot)$ is the max-min normalisation to fulfil the capital budget constraint. The weight vector $\mathbf{w}_{A,t} \in \mathbb{R}^N$ is updated by $\mathbf{w}_{A,t} = \text{Norm}(\frac{1}{T_{dc}} \sum_{m=t}^{t-T_{dc}} \mathbf{x}_m^{dc})$.

On the other hand, as a model-free and off-policy DRL approach, TD3 [17] has been widely used to capture temporal and spatial information from the interaction with environments. In this work, the TD3-based method is applied to mine the data patterns from a time-series perspective whereas the DC method focuses on an event-based perspective. Given the historical price $\mathbf{P}_t \in \mathbb{R}^{N \times N_f}$ of all assets in a portfolio, the weight vector $\mathbf{w}_{A,t}$ obtained by the TD3-based function $f(\cdot)$ with the optimized parameters θ is represented as $\mathbf{w}_{A,t} = \text{Softmax}(f(\mathbf{P}_t; \theta))$. N_f is the number of features of each asset.

It is worth noting that the price analyser of the proposed framework can be conveniently replaced by other algorithmic approaches and DRL-based approaches for further extension.

D. Portfolio Aggregator

As discussed in the above sections, straightforwardly summarising all sentiment scores may introduce noise to the final trading decision due to the reliability of different information sources. To carefully extract the useful message from financial texts to help enhance the overall investment performance, a simple yet effective weighted-sum method is given to generate the new portfolio with the use of additional sentiment information. According to the confidence matrix, the sentiment information coming from the source with higher confidence is more likely to be related to price changes, whereas the information may be irrelevant to price changes when it comes from the information source with lower confidence. Thus, the weighted sentiment score can better reflect the actual sentiment in the current environment. Given the weight $w_{A,i,t} \in \mathbf{w}_{A,t}$ suggested by the price analyser, the confidence vector $\mathbf{r}_{i,t} \in \mathbf{R}_t$ of all involved information sources to i^{th} asset, and the vector of sentiment scores $\mathbf{s}_{i,t} \in \mathbf{S}_t$ of i^{th} asset from all sources, the updated investment weight $w_{i,t} \in \mathbf{W}_t$ of i^{th} asset in a portfolio at time t is calculated by $w_{i,t} = w_{A,i,t} + \mathbf{r}_{i,t}^\top \mathbf{s}_{i,t}$.

IV. AN EMPIRICAL EVALUATION

Data sets and Comparative Methods: To demonstrate the effectiveness of the proposed MuSA framework in the real-world financial market, two challenging data sets [1] of Dow Jones Industrial Average (DJIA) and S&P 500 indexes from February 2010 to May 2020 are utilised to evaluate the performance of the MuSA and comparative methods. The datasets include the news headlines collected from two information sources. Among the trading period, the models are trained by the first five-year data and validated by the subsequent data set of two years. Lastly, the validated models of each approach are tested on the last three-year data. In terms of the company capital, the top 10 stocks of each market index are selected to construct the portfolio in the experiments. Moreover, to investigate the performance of the MuSA in different scales of portfolios, this work will conduct a scalability test on larger portfolios consisting of up to 100 stocks as well. Besides, all of the experiments are run on a machine installed with the AMD 12-core processor and two Nvidia RTX 3090 CPU cards. There are eight representative algorithmic-based or DRL-based methods selected to compare against the MuSA framework. The algorithmic methods are popular financial approaches that have been verified in different financial market all around the world, including the Best Constant Rebalanced Portfolio (BCRP) [21], Online Moving Average Reversion (OLMAR) [13], Passive Aggressive Mean Reversion (PAMR) [14], Correlation-driven Nonparametric

[1] https://www.kaggle.com/datasets/miguelaenlle/massive-stock-news-analysis-db-for-nlpbacktests

TABLE I
THE PERFORMANCE COMPARISON ON TWO CHALLENGING DATA SETS

Market	DJIA						S&P 500					
Metrics	AR		MDD		SR		AR		MDD		SR	
Models	Avg.(%) ↑	S.D.	Avg.(%) ↓	S.D.	Avg.↑	S.D.	Avg.(%) ↑	S.D.	Avg.(%) ↓	S.D.	Avg.↑	S.D.
BCRP	5.00	± 3.80E-03	47.53	± 4.34E-03	0.11	± 1.29E-02	9.96	± 4.28E-03	32.56	± 2.47E-03	0.30	± 1.58E-02
OLMAR	6.91	± 3.02E-03	34.78	± 2.94E-03	0.15	± 8.28E-03	12.25	± 5.20E-03	45.92	± 2.28E-03	0.35	± 1.73E-02
PAMR	3.72	± 3.52E-03	33.50	± 1.97E-03	0.06	± 1.00E-02	2.79	± 4.55E-03	47.45	± 6.14E-03	0.04	± 1.55E-02
CORN	-2.62	± 2.01E-03	35.61	± 5.04E-03	-0.15	± 7.02E-03	5.88	± 4.10E-03	36.96	± 7.90E-04	0.16	± 1.58E-02
RMR	10.46	± 6.99E-03	32.88	± 1.40E-03	0.25	± 1.96E-02	8.91	± 7.62E-03	54.37	± 3.87E-03	0.24	± 2.50E-02
DPM	16.77	± 6.92E-03	30.80	± 7.22E-03	0.67	± 6.55E-03	17.98	± 2.73E-03	31.60	± 3.09E-03	0.75	± 4.81E-03
PPN	16.42	± 4.08E-04	**30.42**	± 7.80E-04	0.67	± 6.26E-04	17.61	± 4.28E-04	**30.98**	± 8.00E-04	0.75	± 4.24E-04
RAT	16.49	± 5.77E-03	30.54	± 1.08E-02	0.67	± 5.85E-04	18.15	± 5.38E-03	31.87	± 8.91E-03	0.75	± 2.55E-03
MuSA-DC	17.68	± 1.42E-04	32.99	± 1.39E-04	0.66	± 1.05E-03	**20.15**	± 6.38E-05	33.66	± 2.88E-04	**0.79**	± 4.36E-04
MuSA-TD3	**18.00**	± 1.31E-04	33.12	± 1.62E-04	**0.67**	± 6.00E-04	19.54	± 5.41E-04	33.72	± 8.30E-05	0.77	± 2.52E-03

TABLE II
THE ABLATION STUDY OF THE MUSA FRAMEWORK ON THE DJIA INDEX AND S&P 500 INDEX

Market	DJIA						S&P 500					
Metrics	AR		MDD		SR		AR		MDD		SR	
Models	Avg.(%) ↑	S.D.	Avg.(%) ↓	S.D.	Avg.↑	S.D.	Avg.(%) ↑	S.D.	Avg.(%) ↓	S.D.	Avg.↑	S.D.
DC	16.79	± 4.15E-04	**32.37**	± 2.74E-04	0.64	± 1.57E-03	19.05	± 1.80E-04	34.47	± 2.20E-04	0.75	± 1.26E-03
DC+Src1	17.43	± 9.87E-05	33.02	± 1.25E-04	0.65	± 5.57E-04	20.00	± 1.07E-04	**33.51**	± 2.79E-04	0.79	± 1.08E-03
DC+Src2	17.51	± 1.76E-04	33.09	± 1.89E-04	0.66	± 9.96E-04	19.88	± 1.47E-04	33.72	± 1.82E-04	0.78	± 6.93E-04
DC+MultiSrc+EW	17.28	± 1.70E-04	33.18	± 1.87E-04	0.65	± 9.53E-04	19.81	± 8.75E-05	33.68	± 1.70E-04	0.78	± 2.89E-04
MuSA-DC	**17.68**	± 1.42E-04	32.99	± 1.39E-04	**0.66**	± 1.05E-03	**20.15**	± 6.38E-05	33.66	± 2.88E-04	**0.79**	± 4.36E-04
TD3	17.86	± 1.13E-04	**33.08**	± 7.58E-05	0.66	± 8.79E-04	19.23	± 4.31E-04	33.72	± 1.76E-04	0.75	± 1.88E-03
TD3+Src1	17.54	± 4.19E-04	33.17	± 3.24E-04	0.65	± 1.70E-03	19.29	± 7.86E-04	**33.56**	± 2.13E-04	0.75	± 2.91E-03
TD3+Src2	17.74	± 4.68E-04	33.23	± 2.89E-04	0.66	± 2.02E-03	19.11	± 1.94E-04	33.76	± 3.09E-04	0.75	± 1.84E-03
TD3+MultiSrc+EW	17.30	± 1.16E-04	33.39	± 2.91E-04	0.64	± 5.38E-04	19.01	± 2.03E-04	33.72	± 3.49E-04	0.74	± 1.43E-03
MuSA-TD3	**18.00**	± 1.31E-04	33.12	± 1.62E-04	**0.67**	± 6.00E-04	**19.54**	± 5.41E-04	33.72	± 8.30E-05	**0.77**	± 2.52E-03

Learning Strategy (CORN) [15], and Robust Median Reversion (RMR) [16]. Also, several latest DRL-based portfolio optimisation frameworks are compared as baselines in this work. They are Deep Portfolio Management (DPM) [22], Portfolio Policy Network (PPN) [23], and Relation-Aware Transformer (RAT) [6]. Besides, three widely used performance metrics including annualised return (AR), maximum drawdown (MDD), and Sharpe ratio (SR) of each approach are compared, and the average of results over 5 runs are reported.

Performance Analysis: As shown in Table I, some well-known approaches are compared with the proposed MuSA framework using different price analysers including DC and TD3. From the reported results on the DJIA index, the best MuSA framework attains the highest AR at 18% and the highest SR at 0.67 while maintaining the MDD at a relatively low level when compared against that of other algorithmic-based approaches and the DRL-based methods as well. Similarly, in the S&P 500 index, the MuSA significantly outperforms other comparative methods in both AR and SR, with at least 8% increase at AR than the best algorithmic methods and 1.5% AR higher than the best DRL-based approaches. This significantly shows the outstanding performance of the MuSA against the algorithmic approaches as well as the popular DRL-based methods.

Ablation Study: Table II reviews the results of the ablation study of the proposed MuSA framework in the DJIA and S&P 500 markets. MultiSrc+EW means equally dealing with the sentiment scores from all sources while the Src1 and Src2 stand for only using the first or second information source, respectively. The MuSA demonstrates a remarkable enhancement in achieving higher returns when compared with the original DC and TD3 without considering sentiment information in both DJIA and S&P 500 markets, which proves that utilising sentiment information in a proper way can further enhance the performance of the original models in a turbulent environment. On the other hand, sometimes using a single information source may bring biased or irrelevant messages to adversely affect the trading. For instance, the TD3 integrated with the second source suffers a loss at AR by 0.1% against that of the original TD3. Additionally, straightforwardly summarising sentiment scores from different sources may confuse the useful and noise messages such that the worse performance is obtained by DC/TD3+MultiSrc+EW when compared with that of using a single source. Based on the above discussion, the outstanding AR and SR of the MuSA clearly reveal the ability to achieve higher profits while managing the risks at a lower level after carefully analyzing sentiment information. More importantly, the proposed framework can work as a plug-and-play tool to integrate with the existing algorithmic-based or DRL-based algorithms for improving trading performance by utilising sentiment information in the financial area.

Scalability Analysis: The DJIA index has a total of 30 stocks, among which 8 stocks are filtered due to missing data. Therefore, there are 22 stocks that are available on the DJIA market during the trading period. Fig. 3 investigates the AR performance of the MuSA framework using the DC-based method as the price analyser against that of the original DC and its variants in managing the portfolios of different scales from 5 to 22 stocks in the DJIA index and from 10 to 100 stocks in the S&P 500 index. At a similar risk level, the two variants of the DC-based MuSA framework achieve around 1% to 2% excess returns in AR than the original DC in all experiments of the portfolios of different scales. Besides, the AR performance of DC variants using a single information source or the equal weight method is slightly higher than that of the original DC but is much lower than that of the MuSA.

This clearly reveals that the proposed MuSA can effectively utilise the multimodal data to optimise portfolios of different scales and sheds light on managing large-scale portfolios and funds in the highly turbulent financial market.

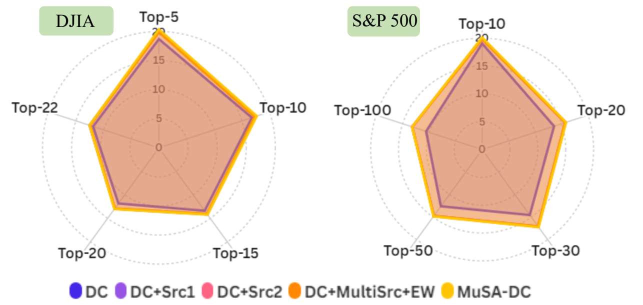

Fig. 3. The Annual Return (%) Comparison of the Original DC and the MuSA Framework for the Portfolio of Different Scales on Two Market Indexes

Hyper-parameter Analysis: Table III reviews the impacts of two key hyper-parameters introduced by the proposed MuSA framework in which the DC-based price analyser is utilised as the example in the experiments. From the investigation results of the observation period M of calculating confidence coefficients on the DJIA and S&P 500 markets, the proposed MuSA framework obtains the highest AR and SR on two markets when $M = 10$, which implies that the financial text data such as news and company reports in the involved data set will significantly affect the price movements within 10 days. On the other hand, the different observation periods T_{dc} of the moving average DC in the DC-based MuSA framework are studied in this section in which the variant of the MuSA framework achieves the highest AR and SR when $T_{dc} = 1$ on the DJIA market and $T_{dc} = 20$ on the S&P 500 market due to the different markets having various patterns of price movements, thus the hyper-parameter T_{dc} in the moving average DC can be further optimised by observing the historical price data to capture more potential patterns that may significantly affect stock prices in different markets.

TABLE III
THE HYPER-PARAMETER ANALYSIS OF THE MUSA FRAMEWORK ON THE DJIA INDEX AND S&P 500 INDEX

Market		DJIA			S&P 500		
Metrics		**AR**(%) ↑	**MDD**(%) ↓	**SR**↑	**AR**(%) ↑	**MDD**(%) ↓	**SR**↑
$M =$	5	17.59	**32.89**	0.66	19.86	33.83	0.78
	10	**17.68**	32.99	**0.66**	**20.15**	**33.66**	**0.79**
	30	17.54	32.97	0.66	19.96	33.67	0.79
	60	17.59	32.97	0.66	19.95	33.66	0.79
	90	17.58	32.98	0.66	19.96	33.66	0.79
	180	17.58	32.97	0.66	19.96	33.66	0.79
$T_{dc} =$	1	**19.45**	**32.23**	**0.74**	20.00	**32.46**	0.80
	5	18.28	32.72	0.69	19.64	33.31	0.78
	10	17.68	32.99	0.66	20.15	33.66	0.79
	20	17.16	33.15	0.64	**20.48**	33.11	**0.81**
	30	17.37	33.24	0.65	20.43	33.09	0.80

V. CONCLUSION

Recently, many portfolio optimization approaches tried to assess market sentiment from news for making trading decisions. Yet the existing frameworks rarely concern the reliability of different information sources in which the media may provide irrelevant or inaccurate information. Accordingly, a multimodal and sentiment-based adaptive framework namely the MuSA is proposed where an entropy-based confidence learning mechanism calculates the confidence of news from different sources, filtering irrelevant news from unreliable media and reducing potential investment risks. The empirical results demonstrate an impressive improvement of the MuSA in achieving excess returns by effectively utilizing sentiment information. More importantly, the proposed MuSA sheds light on many fields. For instance, the MuSA can be extended to evaluate the reliability of large-language models to judge financial texts from more information sources. Besides, more data with different modalities can be considered in the MuSA to deduce the price movements for higher profits under the volatile market.

REFERENCES

[1] F. Rundo *et al.*, "Machine learning for quantitative finance applications: A survey," *Applied Sciences*, vol. 9, no. 24, p. 5574, 2019.
[2] A. M. Ozbayoglu *et al.*, "Deep learning for financial applications: A survey," *Applied Soft Computing*, vol. 93, p. 106384, 2020.
[3] H. M. Markowitz and G. P. Todd, *Mean-variance Analysis in Portfolio Choice and Capital Markets*. John Wiley & Sons, 2000.
[4] B. Li and S. C. Hoi, "Online portfolio selection: A survey," *ACM Computing Surveys (CSUR)*, vol. 46, no. 3, pp. 1–36, 2014.
[5] J. Gypteau *et al.*, "Generating directional change based trading strategies with genetic programming," in *Applications of Evolutionary Computation: 18th European Conference*. Springer, 2015, pp. 267–278.
[6] K. Xu, Y. Zhang, D. Ye, P. Zhao, and M. Tan, "Relation-aware transformer for portfolio policy learning," in *IJCAI*, 2021.
[7] Y.-H. Lien *et al.*, "Contrastive learning and reward smoothing for deep portfolio management," in *IJCAI*, 2023, pp. 3966–3974.
[8] D. Araci, "Finbert: Financial sentiment analysis with pre-trained language models," *arXiv preprint arXiv:1908.10063*, 2019.
[9] OpenAI, "Gpt-4 technical report," 2023.
[10] H. Touvron *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
[11] Q. Liang, M. Zhu, X. Zheng, and Y. Wang, "An adaptive news-driven method for cvar-sensitive online portfolio selection in non-stationary financial markets." in *IJCAI*, 2021.
[12] Y. Ye *et al.*, "Reinforcement-learning based portfolio management with augmented asset movement prediction states," in *AAAI*, 2020.
[13] B. Li and S. C. Hoi, "On-line portfolio selection with moving average reversion," in *ICML*, 2012.
[14] B. Li *et al.*, "Pamr: Passive aggressive mean reversion strategy for portfolio selection," *Machine learning*, 2012.
[15] B. Li, S. C. Hoi, and V. Gopalkrishnan, "Corn: Correlation-driven nonparametric learning approach for portfolio selection," *ACM Transactions on Intelligent Systems and Technology*, 2011.
[16] D. Huang *et al.*, "Robust median reversion strategy for online portfolio selection," *IEEE TKDE*, vol. 28, no. 9, pp. 2480–2493, 2016.
[17] N. Baard *et al.*, "Twin-delayed deep deterministic policy gradient algorithm for portfolio selection," in *IEEE Symposium on CIFEr*, 2022.
[18] Z. Li, V. Tam, and K. L. Yeung, "Developing a multi-agent and self-adaptive framework with deep reinforcement learning for dynamic portfolio risk management," *AAMAS*, 2024.
[19] S. Fujimoto *et al.*, "Addressing function approximation error in actor-critic methods," in *ICML*, 2018, pp. 1587–1596.
[20] R. J. Almeida *et al.*, "Portfolio re-balancing and optimization using directional changes and genetic algorithms," in *IEEE CEC*, 2023.
[21] T. M. Cover, "Universal portfolios," *Mathematical finance*, vol. 1, no. 1, pp. 1–29, 1991.
[22] Z. Jiang *et al.*, "A deep reinforcement learning framework for the financial portfolio management problem," *arXiv preprint arXiv:1706.10059*, 2017.
[23] Y. Zhang *et al.*, "Cost-sensitive portfolio selection via deep reinforcement learning," *IEEE TKDE*, 2020.