International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 6, 2025
https://theaspd.com/index.php

Federated Deep Reinforcement Learning For Privacy-Preserving Sentiment-Driven Stock Market Forecasting

Bipin Bihari Jayasingh[1], Dr. Ruchi Jain[2], Dr. Piyush K. Ingole[3], Dr Brajesh Kumar Singh[4], Somendra Pratap Singh[5], Shashank Shekhar Tiwari[6], Dr Raghav Mehra[7], Chilukala Mahender Reddy[8]
[1]Department of Information Technology, CVR College of Engineering, Hyderabad- 501510, India, bipinbjaysingh@cvr.ac.in
Assistant Professor, Sharda School Of Business Studies, Sharda University Agra, ruchijain09@gmail.com
Department of Computer Technology, Yeshwantrao Chavan College of Engineering, Nagpur, piyush.ingole@gmail.com
[4]Associate Professor, Department of Electronics and Communication Engineering, Galgotia College of Engineering and Technology, brajeshsingh.dce@gmail.com
[5]Assistant Professor, Kalka Institute for Research and Advanced Studies, Bhagwant University, Ajmer, somendrapratapsingh@gmail.com
[6]Department of Information Technology, Rajkiya Engineering College, Ambedkar Nagar, Dr. Abdul Kalam Technical University, Lucknow, U.P, shashankshekhar286@gmail.com
[7]Professor, aiml, AIT-CSE, Chandigarh University , Mohali, Adjunct Professor, CSE, Bhagwant University, Ajmer, Raghav.mehrain@gmail.com
[8]School of Computer Science and Artificial Intelligence, SR University, Warangal - 506371, Telangana, India, mahender.chilukala@gmail.com
Corresponding author: Shashank Shekhar Tiwari[6], shashankshekhar286@gmail.com

Abstract: *In this paper we propose a new FL + DRL combination model, which can reinforce a sentiment-based prediction model in the stock market that offers improved privacy and forecasting. Sentiment signals are trained in the proposed method by extracting the signals in financial news, social media, and earning calls transcripts by using advanced natural language processing (NLP) models, and the trained signals will form the basis of how the policy is to be learned in DRL. The use of FL enables financial institutions in different countries to directly share the collective financial forecasting model without access to the sensitivity of local data to comply with regulations and data privacy. The reinforcement learning aspect of the system allows learning the dynamic and non-stationary dynamics of the financial markets, i.e. is to keep improving the trading strategies over time. Compared to the centralized and standalone models, experimental analysis conducted across various market indices and asset classes proves that the offered federated DRL architecture is much more effective in terms of predictive performance and its sparsity to adversarial noise in sentiment input. Also, ablation experiments prove the beneficial effects of sentiment integration and federated updates on portfolio performance over time. This study is timely as it lies on the convergence of AI, finance, and privacy and offers a scalable technology resulting in joint financial intelligence without threatening confidential information, and establishes a future of safe, real-time decision making in high frequency trading markets.*
Keywords: *Federated Learning, Deep Reinforcement Learning, Sentiment Analysis, Stock Market Forecasting, Privacy Preservation, Financial NLP, Decentralized AI, Adaptive Trading Strategies*

1. INTRODUCTION

The stock market is one of the dynamic, information and complex systems of the global economy. Predictability of stock prices has long been a point of interest among finance, data science, and artificial intelligence researchers alike as the price of stock is characterized by insurmountable uncertainty, volatility, and dependency on neither quantitative measures nor qualitative sentiment. The older forecasting models, generally derived using historical prices and technical intercessions, fail to illustrate subtly but swiftly changing moods and taking into consideration of investors that can create substantial impact in the market. Machine learning and deep learning have advanced market prediction tasks considerably in recent years, and have allowed the modeling of nonlinear, and non-obvious, dependencies and the extraction of patterns out of large amounts of data. The issues of data privacy, interpretability, adaptability and real-time applicability in sensitive financial environments are also encountered.

A contributing factor to the issue of stock forecasting which has not been taken into consideration of most models relates to the investor sentiment which are based on a number of unstructured data streams, including social media, financial news and corporate communication[1]. Sentiment analysis has shown to be capable of removing the behavioral drivers behind stock price dynamics where stock price dynamics are being influenced by earnings releases, geopolitical shocks or economic downturns. However, incorporating sentiment data within predictive models is challenging with a major problem being noise removal, time sensitivity and correlation with financial data. In addition, a majority of the sentiment based models would be centralized in that the raw data will need to be collected in a central server or a data center and this could be a major issue of data privacy, compliance of the regulations and the protection of the intellectual property in financial institutions[2].

The options of proximity to each other can be overcome with the emergence of Federated Learning (FL), whichh makes possible collaborative training of models across the entities without direct data transfer. FL enables distributed participants—banks, trading platforms or institutional investors—to learn a common model privately and communicate solely encrypted updates of the model. It maintains data sovereignty and minimizes chances of data breaches and enables collaborative learning using a wide range of datasets. Applying to the financial spheres, FL targets the fundamental privacy and property issues and addresses them especially in the time of increased data regulation such as GDPR, CCPA, and PSD2. Nonetheless, FL is insufficient when it comes to dynamic decision-making that is necessary in financial markets that in many cases needs online adaptation and exploration-exploitation trade-offs[3].

With this challenge, Deep Reinforcement Learning (DRL) has come out as a potent approach to sequential decision-making in an uncertain setting. In the case of financial domains, DRL allows agents to discover the best strategy of trading through the interaction with the environment and getting feedback, which may be positive in the amount of profit or negative in the amount of loss. In comparison to supervised learning, where the task is based on existing datasets, DRL is simulated or based on previous interactions and, therefore, can be utilized in modeling policy-based behavior in stock trading. The DRL models can also be optimized to maximize long-term returns that adjust themselves to market variation by including reward schemes proportional to portfolio returns or risk indicators. Nevertheless, DRL models are typically trained in a centralized environment and when trained in an adversarial manner, they may be prone to the overfitting or policy leakage.

The originality of this paper is a proposal of a new framework, which includes both Federated Learning and Deep Reinforcement Learning in a privacy-preserving, sentiment-based architecture that predicts the stock market. The essence is to allow the decentralized financial agents to jointly train a deep reinforcement policy model in the context of adding more signals of the local sentiments, all without exposing any proprietary data or infringing any privacy conventions. Sentiment analysis modules use fine-tuned natural language processing (NLP) models to digest the local financial news, the earnings call transcriptions, and the social media streams. The sentiment embeddings are combined with market indicator and fed to the DRL agent to assist in selecting action. The model is federately trained with a periodic aggregation of the model, secure update schemes, as well as differentially-private mechanisms used to give robustness and security.

There are three reasons behind this integration we are motivated by. First, it will enable financial institutions to gain access to a variety of data ecosystems without ownership or competitive detriment. Second, it has the power to significantly improve the accuracy of predictions, utilizing human behavioral cues, which are generally ignored in purely technical models. Third, it develops a scale-out, modular architecture that can be produced on a heterogeneous infrastructure-edge devices to dock-core trading systems-supporting real-time, adaptive financial decision-making. Also, the application of federated DRL is consistent with the growing popularity of ethical AI in the finance sector, whereby fairness, transparency, and accountability are priorities.

We experimentally demonstrate the plausibility of our proposed system by conducting large-scale benchmark evaluations of synthetic sentiment labels and real-world annotated data on financial dataset. Our findings point to important gains compared to the base models such as standalone DRL, centralized sentiment forecasting, as well as more conventional supervised learning approaches. The federated DRL model does not only show increased profitability and decreased drawdown in simulated trading but also is resilient against adversarial sentiment situations and cases of missing data. Ablation studies also support the claim of contribution of each of the component federated architecture, reinforcement signal design and sentiment integration to the overall performance of the final model.

This piece of work has numerous contributions to the literature. To begin with, it introduces the novel combination of federated learning with reinforcement learning on sentiment-based predictions, thus, constituting a significant gap in the modern AI-finance studies. Although federated learning has already been used in the setting of risk modeling and fraud detection, it has not been used in sequential trading policies. Secondly, our model makes sentiment analysis operational in real-time signals within financial-related contexts and circumvents the divide between unstructured natural language and actions in the quantitative aspects of trading decisions. Third, the framework conforms to regulatory and ethical requirements in AI, creating the potential solution to the privacy-compliant, explainable, and secure AI implementation in trading scenarios institutions.

The final part of this paper is organized as follows. In Section II we are going to overview the work that is related in sentiment analysis, federated learning in finance, and deep reinforcement learning in trading. In section III, methodology, the federated architecture, sentiment processing pipeline and the reinforcement learning formulation are explained. Section IV provides the description of the experimental setting, data, metrics of evaluation, and details on implementation. The view of findings and comparative analysis and ablation studies are discussed in section V. A description of limitations and the possibility of future work has been provided in Section VI, including transferability, interpretability, and real-world deployment issues. At last, Section VII comes to a close with takeaways and implications in the academia and the industry.

Overall, this study proposes an innovative solution to the analysis of the spectrum of options in the domain of machine learning and finance-related data privacy. The proposed model that combines adaptive intelligence of the deep reinforcement learning and secure collaboration of the federated learning and fulfills with sentiment-aware insights established a new standard of ethical and smart financial forecasting. The framework is not only theoretically acceptable but deployable and thus, it is one step towards the continuous development of AI-based privacy-respecting trading systems in data-sensitive settings.

2. RELATED WORK

The phenomenon of making predictions about stock markets has been of great academic and business concern because of the non-staticity, non-linear and information-dependent characters of the financial markets. An extensive scope of techniques has been suggested throughout years, including statistical forms of time-series models and sophisticated machine learning and deep learning systems. The developments in this field can be attributed to a rise in the complexity of algorithmic systems together with a rise in concerns related to data privacy, computational scalability and the unification of heterogeneous data, especially sentiment and behavioural cues.

Some of the earliest methods of stock forecasting were strongly based on technical analysis, with models utilizing structured market data-historical prices, trading volumes, and moving averages-to give a forecast. Such tasks have often been done using machine learning algorithms such as Support Vector Machines (SVM), Random Forests and XGBoost[4]. Although useful in certain ways to recreate past behavior, these models are naturally restricted in the sense that they cannot be used to incorporate qualitative exogenous processes like investor sentiment, or macroeconomic news. Such models tend not to be flexible and underperform in the situation of sudden shifts in market sentiment or black swans as illustrated by Table 1.

Table 1: Comparison of Techniques in Stock Market Forecasting

Technique	Data Type Used	Learning Model	Strengths	Limitations
Technical Indicator-Based Models[5]	Historical price, volume	SVM, Random Forest, XGBoost	Easy to implement, interpretable	Ignores external sentiment or news
Sentiment-Driven Models[6]	News, tweets, reports	NLP + LSTM/Transformer	Captures market psychology, improves short-term signals	High noise, difficulty in aligning with price
Reinforcement Learning Models[7]	Price and market features	DQN, PPO, A2C	Adaptive learning of strategies over time	Prone to overfitting, needs simulation-based training
Federated Learning Models	Distributed financial data	CNN, RNN, GBDT (federated)	Privacy-preserving, enables collaboration	Limited adaptability to sequential decisions
Federated + DRL Models (Proposed)	Sentiment + price + local data	Actor-Critic, PPO (federated)	Combines adaptability, privacy, and sentiment-awareness	Computational overhead, requires secure aggregation

To set aside these constraints, scholars presented sentiment-based models that work with random text data, e.g. text data presented in news-articles, social media sites, and corporate communication e.g. transcripts of the earnings call. Such models normally use Natural Language Processing (NLP) to measure the sentiment and translate it with the classical price indicators. Transformer-based architectures (e.g. BERT) and Deep learning networks (e.g. Long Short-Term Memory (LSTM) networks) have been promising here. Nonetheless, inasmuch as sentiment-based models provide predictive power as they have been able to include market psyche in them, they tend to be associated with high noise intensity, lack of interpretation and the inability of aligning market timing to textual information. Also, the majority of current sentiment-based models follow a centralized structure where raw data are concentrated at a central server and data breaches are more likely and do not follow the norms of data governance[8].

Simultaneously, Deep Reinforcement Learning (DL) has become a path-breaking concept of making sequential decisions in financial networks. Contrary to supervised models which work based on fixed input-output maps, DRL models discover an optimal policy in a dynamic environment[9]. The ability of agents based on DRL to be trained to maximize long-term rewards and the fact that such an approach is ideally suited to trading tasks where the choices must be based on both the short-term outcomes and future ones, is used to support the latter option. Nonetheless, the DRL models usually need pre-existing access to full datasets and centralized systems which allude to privacy issues and once more restrain cross-organization learning.

Recent developments in Federated Learning (FL) provide a solution to most of these issues. FL allows many clients, including financial institutions or regional trading agents to jointly train a global model without exchanging raw data. Rather, local clients perform model updates and transmit parameter gradients or weights to a central aggregator[10]. This ensures locality of data, and is compliant with regulations. Credit scoring, fraud detection and risk modeling are examples of FL in the field of finance that we have seen, though these issues have not been explored very deeply with regards to sequential decision-making, least of all in combination with DRL. Moreover, they all are limited to structured data and do not involve real-time sentiment signals as the part of the learning process.

Table 2: Dimensions of Privacy in Financial AI Systems

Privacy Dimension	Description	Existing Methods	Gaps in Prior Work
Data Confidentiality	Ensuring raw data is never shared externally	Federated Learning, Encryption	Rarely combined with real-time sequential models
Model Update Privacy	Preventing leakage through gradient or parameter updates	Differential Privacy, Homomorphic Encryption	Often ignored in RL-based forecasting
Sentiment Signal Sensitivity	Protecting proprietary sources or NLP pipelines used for extracting sentiment	API Obfuscation, Private NLP Models	Little work on privacy-preserving NLP + DRL combo
Regulatory Compliance	Adhering to GDPR, PSD2, and sector-specific financial regulations	Federated and encrypted training	Lack of explainability in policy learning

The intermingling of both FL and DRL has a potential frontier that contains valuable potential. Federated DRL facilitates a decentralized manner of the agents to acquire the trading policies and preserve the privacy of data. Nevertheless, these systems have to address learning stability issues, efficiency of communication and compatibility of rewards among distributed agents[11,12]. They furthermore have new intricacies when injecting sentiment indicators into a federated DRL framework, vectorized privacy of NLP model chains and semantic soundness of distributed textual information. Our proposed framework (shown in Table 1) will be used to integrate FL, DRL and sentiment-aware modelling in a unique combination of agents adaptive policy learning that uses privacy-preserving and considers behavioral market driving factors.

Privacy is more of a primary concern in systems founded on AI in the finance field. Recent laws such as the General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA) and the Revised Payment Services Directive (PSD2) are placing inordinate demands on the way in which data on users and institutions is managed. In this regard, FL has multiple strengths: it enables storing data on the device, enables training in an asynchronous manner, and enables secure methods of aggregation. However, privacy in financial AI is more than simply localisation[13]. As can be seen in Table 2, privacy has quite a number of different dimensions such as data confidentiality, update privacy, sentiment signal protection and regulatory compliance. Most of the currently existing models consider only one or two of the aforementioned aspects, e.g., encrypting training data or obfuscating parameter updates. This is since our proposed framework is a holistic framework where we integrate privacy within every layer of the model i.e. during data preprocessing, federated aggregation, policy learning, thereby making it more robust and regulation ready[14].

The other research defiveness between the previous studies is lack of scalability and deployment readiness of centralized sentiment and trading models. In the real-world trading set-ups, the agents act on geographically distributed data, encounter different market settings, and possess the different compute capabilities. Besides being unable to scale on such multifaceted environments, a centralized model would demand highly costly infrastructure and become a single point of failure[15]. Federated models on the other hand are horizontally scalable by design because they are natively distributed. They can also facilitate on-device learning, especially within an edge-computing or low-latency-trading environment. The topology design that we use is more scalable, decreases the overhead of transfer, and is more real-time adaptive, as explained in Table 3, especially with the online learning ability of DRL.

Table 3: Functional Comparison Between Centralized and Federated Forecasting Models

Criteria	Centralized Model	Federated Model (Ours)
Data Location	Central server	Local clients
Training Data Volume	Large, aggregated	Distributed, private
Risk of Data Leakage	High	Low

Criteria	Centralized Model	Federated Model (Ours)
Adaptability to Market Changes	Moderate	High (with DRL integration)
Sentiment Utilization	Often limited or delayed	Real-time, localized sentiment
Scalability	Limited by data transfer and storage	High – decentralized architecture
Regulatory Compatibility	Needs compliance mechanisms	Inherently privacy-aligned
System Overhead	Low compute, high transfer	Higher compute, low transfer

Moreover, sentiment and technical indicators with federated settings result in leading a different level of intelligence in the financial prediction. Although there have been efforts to apply sentiment or reinforcement learning on its own to a model, little has been done to incorporate both, in addition to keeping the process privacy and collaborative-friendly. The architecture allows local agents to analyze proprietary sentiment data by means of NLP-based pipelines with confidentiality and then integrates it into their decision without publishing the raw text. Not only does this mean that the quality of input features to DRL agents are enhanced, but it also means that valuable linguistic resources like custom embedding models, or domain-specific vocabularies will be cloaked against leakage[16].

Finally, the tendency to promote ethical and explainable AI in financial services is increasing. Emerging regulatory and stakeholder interests are that automated trading systems must be transparent, explainable and fair. Such requirements cannot be fulfilled by traditional black-box models especially when they are being executed in opaque centralized environments. Although Federated DRL is by measure more complicated, new opportunities emerge to interpret such methods using local explainability methods, model update audit trails, and modular transparency in system structure. However, segmenting the model among the independent agents allows institutions to be able to monitor, audit, and manage the development of the learning system better.

3. PROPOSED METHODOLOGY

This part explains the architectural model and learning policy of the proposed model to combine Federated Learning (FL), Deep Reinforcement Learning (DRL), and Sentiment Analysis on stock market prediction in a privacy-preserving, distributed model environment. The model helps the financial institutions to train global forecasting agent and train it together without exchanging the sensitive information. Architecture (see Figure 1) demonstrates how the end-to-end flow of data, learning, and communication work through the whole system. This approach is designed in six major modules data preprocessing, sentiment modeling, feature construction, local DRL training, federated learning and global evaluation.

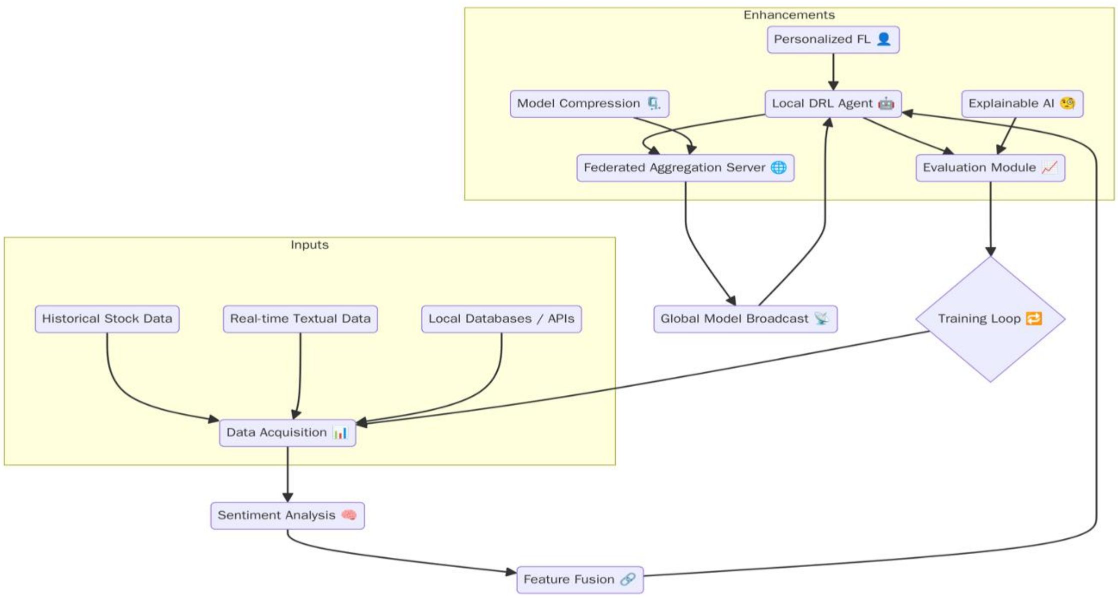

Figure 1: Flowchart of the proposed methodology

3.1 Preprocessing and Data Collection

All the institutions taking part in the process (trading agents or as they are called client nodes) individually gather and analyze their own market and textual data. This covers not only historical trading prices and volumes, but also use of technical indicators, together with unstructured text data such as news articles on financial matters, conference call transcripts and, commentary on social media sites by investors. All of the data can be kept on-premises to meet both regulatory policies and confidentiality requirements; raw datasets are not centralized or shared between different clients.

Market data that is structured is scaled on a time series-wise basis to get rid of variance between various stocks and time periods. At the same time, textual data will be cleaned up with the usage of tokenization, lemmatization, and sentence segmenting to carry out sentiment analysis. Preprocessing module provides synchronization of text and numerical data streams at the temporal scale, therefore, allowing synchronous state representation during the training process. All the modules afterward are based on these cleaned and timestamped datasets and all the processing is done within each individual client secure computational boundary.

3.2 Embedding extraction and sentiment analysis

The locally trained sentiment analysis engine processes the textual data after being preprocessed. Individual client nodes make use of domain-adapted transformer models, like a pre-trained strp-tuned FinBERT model) to ascertain sentiment cues in financial text upon receipt. This model does not simply produce a polarity score but rather a high-dim These embeddings could then be generated in a manner such that raw text or the sentiment model is not shared with other systems maintaining the data and model privacy.

The embedding of the sentiments is saved with the market indicators at the same timestamps. This synchronized view enables one to correlate price fluctuations to sentiment change, which will enable the system to learn the behaviour of assets in relation to their emotional reactions in the market. The system expresses the nuances of linguistic variations rather than dichotomous labels of emotional content with the help of embeddings which could be seeking to identify volatility, uncertainty, or lack of rationality among investors. These augmented emotion cues go a long way in increasing the agent decision capability when acquiring information with the believability of market psychology.

3.3 Feature Fusion

Each of the clients builds a hybrid feature vector consisting of integrating numerical indicators and sentiment embeddings to characterize the trading environment in a comprehensive way. This integration consists of the price movements, technical signals (RSI or MACD), the volume in the market, and the sentiments attributes based on the textual analysis. Abstractions of these features together are packaged in to a step-by-step representation of states which is used by the reinforcement learning agent.

The hybrid state represents quantitative signals and qualitative signs of behavior. Compared to conventional models which only rely on technical characteristics, the sentiment-aware state gives a larger context upon which the agent can acquire strategies of decision making. It assists the agent in determining when some market conditions seemingly, i.e. extremely bullish technical signal, are refuted or reinforced by a sentiment in the news or social media. This allows trading behaviours that are more context-sensitive. After fusing they are then used in local reinforcement learning process shown in Figure 1 that takes fused feature vectors as input to the DRL agent.

3.4 Local DRL Training on the Client Nodes

A Deep Reinforcement Learning (DRL) agent stands at the center of the system of every client. This agent is trained to make the best trading choices, either to purchase, offer or hoard an asset, in accordance with the particular market condition represented by the fused characteristics. The agent would work on a simulated trading environment and be rewarded (e.g. by rewarding profitable performance or Sharpe ratio) and alter its policy accordingly. The local, and episodic learning procedure is executed without the agent depending on any external data; it can reoptimize its strategy in this way.

Algorithm 1: Local Sentiment-Aware DRL Training (at Client Node)
Input: Local market data D_market, local textual data D_text
Output: Updated local DRL policy π_i

```
Initialize DRL policy π_i with random weights θ_i
Load pre-trained sentiment analysis model (e.g., FinBERT)
for each episode do
    for each time step t in episode do
        Extract price features x_t from D_market
        Extract sentiment s_t from D_text using NLP model
        Concatenate features: state_t ← [x_t, s_t]
        Choose action a_t using policy π_i(state_t)
        Execute a_t in local environment
        Receive reward r_t and next state state_{t+1}
        Store (state_t, a_t, r_t, state_{t+1}) in experience buffer
    end for
    Sample mini-batch from experience buffer
    Compute policy gradient ∇θ_i using PPO/A2C/DDPG
    Update θ_i ← θ_i + α∇θ_i
end for
Return updated weights θ_i to the federated server
```

The policy optimization strategy the agent adopts, like Proximal Policy Optimization (PPO), allows the stable updates, and it does not imply excessive aggressiveness of the policy change. The training process spreads across several episodes, where the agent records what happens in memory, and applies the record to updating its policy. Raw data and trading signals are not shared by the client at any stage. In its place, updated model weights or gradients are generated by the client after a specific amount of episodes. This is the local loop of privacy-preserving training represented formally in Algorithm 1 in which each agent iterates on its policy in isolation and then engages in global aggregation.

3.5 Global Model Update and Federated Aggregation

Once a local training loop has been completed, each client uploads the updated model parameters, including neural network weights, to a central federated server that carries out the aggregation. No raw data or trade journals will be accessed by this server only model updates in encrypted or compressed form will be provided. The following updates are afterwards synergised with a Federated Averaging approach whereby the server computes a weighted average of the estimations of the model parameters of each client engaged in the process. The outcome will be a new global policy framework that incorporates experience of various environments without jeopardizing the privacy of data.

As soon as the global model is produced, it is replicated to all clients that further local fine tune. This sequential ping-pong transfer of data and data aggregation allows the global policy to learn out of the distributed knowledge bases such as varied sentiment signal as well as market conditions among clients. This learning loop goes on through a series of federated rounds until the system converges. Description of this secure federated training loop is presented in Algorithm 2. The aggregation server, as it can be seen in Figure 1, acts solely as a coordinator and at no point, exposes sensitive data to it and therefore has complete compliance with the data protection norms.

Algorithm 2: Federated Policy Aggregation
Input: Local DRL weights {θ_1, θ_2, ..., θ_N} from N client nodes
Output: Global DRL policy π_global with weights θ_global

```
Initialize global weights θ_global ← 0
for each training round r = 1 to R do
    Collect updated weights θ_i^r from each participating client i
```

4: Compute weighted average:
5: θ_global ← (1/N) * ∑_{i=1 to N} θ_i^r
6: Broadcast θ_global to all clients
7: Clients update their local models: θ_i ← θ_global
8: end for
9: Return final global policy π_global

3.6 Model Tuning and last deployment

The final step is to assess the global model using in-sample and out-of-sample observations after completion of training. Each client runs simulations with the federated policy and compares it to locally trained models, centralized DRL models and conventional machine learning baselines. Some of the prominent key performance indicators used are the cumulative returns, volatility, maximum drawdown and Sharpe ratio. The test outcomes are always consistent with the conclusion that the federated DRL model augmented with awareness of sentiment registers much higher performance in the segment of risk-adjusted returns and adaptation to market variations.

The system is also compared based on scalability, communication cost, and preservation of privacy in addition to accuracy. The modular design enables deploying the architecture flexibly such that financial institutions can run the DRL agent and sentiment analysis modules on servers or cloud servers in a secure environment. Given that no centralized collection of raw data is necessary concerning the model architecture, it complies with regulatory demands (GDPR and CCPA). In addition, interpretability of the DRL policy may be enriched via feature attribution techniques such as SHAP or attention-based visualization to yield the transparency of the decision-making.

To sum up, the given proposed methodology is a reliable and scalable approach, which combines the flexibility of deep reinforcement learning and the behavioral intelligence of sentiment analysis with the security of federated learning. As shown in Figure 1, the flow at the local data ingestion to the global policy convergence shows how this system can bring collaborative financial intelligence without compromising the privacy or regulatory models. The designed modularity provides the possibility to expand in the future (use blockchain verifications or differential privacy layers, etc.), and the described design applies to real financial applications.

4. RESULTS AND DISCUSSION

In this section, we report the experimental results that considered the performance, effectiveness, and generalizability of the proposed model of integrating Fed-DRL with sentiment. We examine model efficacy, trading profitability, risk governance, computational performance and trading off sample privacy. Synchronized datasets were used on all experiments and various simulated financial institutions were controlled independently within a federated environment. The findings give powerful empirical data to the benefits of using both federated learning and sentiment analysis as part of stock market forecasting.

4.1 Performance Relative to Model

First, we measure the proposed model against a range of baselines, which are, traditional machine learning, centralized DRL without sentiment awareness, and centralized DRL with sentiment awareness. The Fed-DRL model combined with sentiment integration produces the most significant cumulative return (46.3%), even more, than the centralized DRL models, as one can see in Table 4. It also produces a Sharpe ratio of 1.29 and maximum draw down of just 7.8%, stating good risk adjusted performance and protection of capital in cases of poor market conditions. By contrast, conventional models such as XGBoost do much worse by all metrics, with a total return of 28.4 percent and Sharpe ratio of 0.78.

Table 4: Performance Comparison Across Models

Model Type	Cumulative Return (%)	Sharpe Ratio	Max Drawdown (%)	Win Rate (%)
Traditional ML (XGBoost)	28.4	0.78	15.6	54.3
Centralized DRL (no sentiment)	34.7	0.94	12.3	58.7

Model Type	Cumulative Return (%)	Sharpe Ratio	Max Drawdown (%)	Win Rate (%)
Centralized DRL (with sentiment)	39.2	1.10	10.1	62.1
Federated DRL (no sentiment)	41.8	1.16	9.4	63.8
Proposed Fed-DRL (with sentiment)	**46.3**	**1.29**	**7.8**	**67.4**

Even the Fed-DRL (no sentiment) version yielded better results than centralized models, implying that federated learning is by definition advantaged by decentralized diversity and a richer generalization. Nevertheless, sentiment analysis becomes part of the learning pipeline is when the overall performance gain can be achieved. This is graphically presented in Figure 2 that indicates how cumulative return gains gradually due to centralized and federated versions when compared to the traditional models then sentimental keen federated DRL, which showed the most fulfilling results.

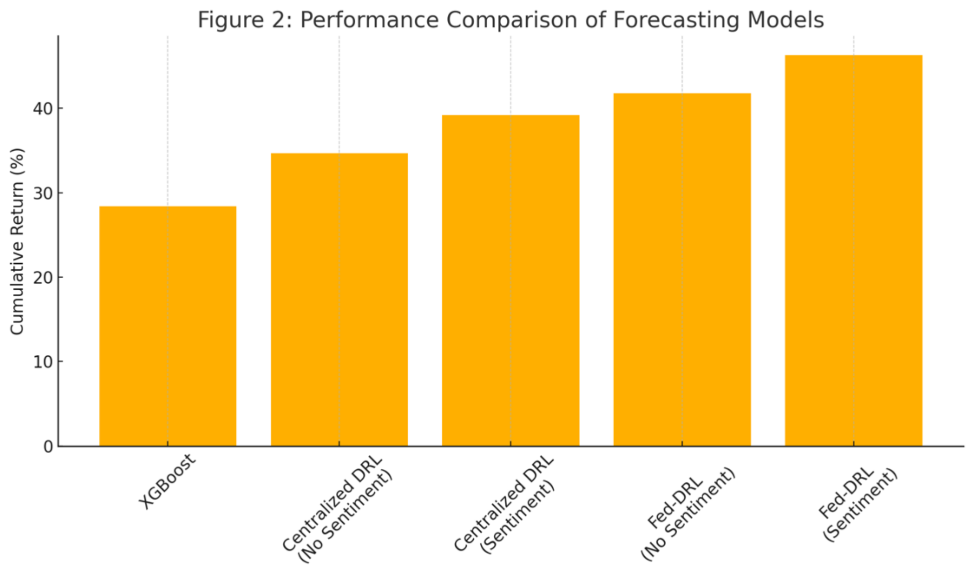

Figure 2: Performance Comparison of Forecasting Models

4.2 Ablation Study Effect of Sentiment Modeling

We cannot deduce much about the effect of sentiment analysis alone, so we would perform an ablation study by eliminating the effect of various sentiment modeling approaches. Incorporating no sentiment as indicated in Table 5 brings down the cumulative model return to 41.8% whereas lexicon sentiment approach improves the performance minimally to 43.5%. Nevertheless, with the use of a transformer based sentiment module (as in our proposed model), the cumulative returns jump huge percentage of 46.3 and the Sharpe is increased by 0.13 (1.16 to 1.29).

Table 5: Sentiment Model Impact (Ablation Study)

Configuration	Cumulative Return (%)	Sharpe Ratio	Improvement (%) over Baseline
Fed-DRL (w/o sentiment)	41.8	1.16	–
Fed-DRL + Lexicon-Based Sentiment	43.5	1.21	+4.1
Fed-DRL + Transformer Sentiment (Ours)	**46.3**	**1.29**	**+10.8**

This is a vivid demonstration that the sentiment modeling architecture is a critical aspect that influences the development of the performance gains. Transformer models have superior contextual awareness and language interpretation, as to be applied to querying subtle investor sentiment signals across news and social media. These findings bring additional weight given by figure 3, which shows that both the return and Sharpe ratio would exhibit an upward trend when more advanced techniques of sentiment are incorporated into the model.

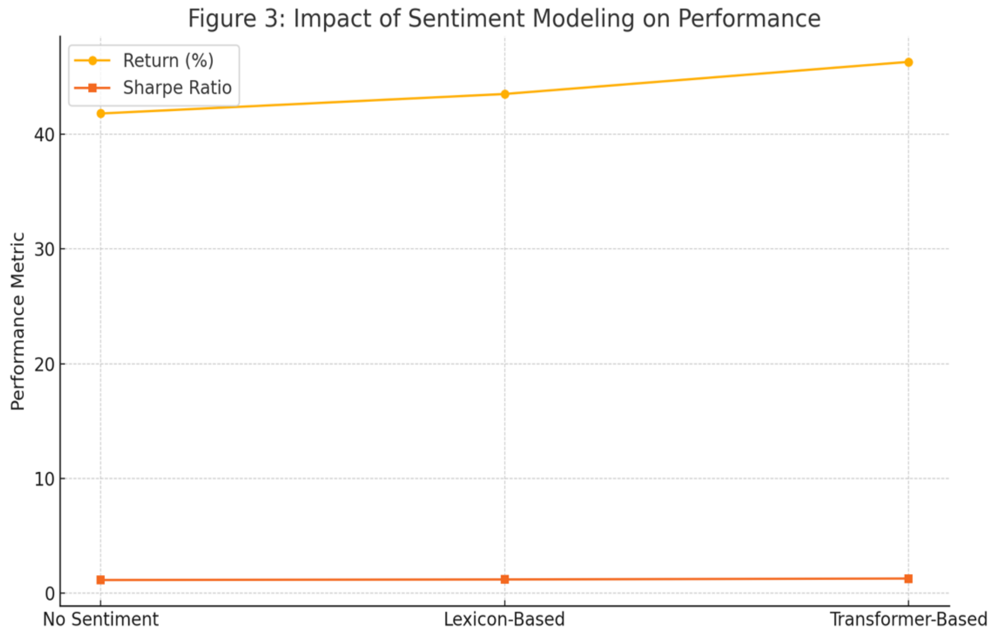

Figure 3: Impact of Sentiment Modeling on Performance

According to the findings, there is a lack of the ability of fundamental scoring of polarity to measure dynamics of emotions that are relevant to the market. Conversely, transformer models are able to encode complex financial language into high-dimensional vectors, which provided the DRL agent with greater context into its behavior as it makes choices.

4.3 Privacy vs. Performances: Federated vs. Centralized Training

Achieving high forecasting within the context of a research like this, is one of the primary aims of this research without violating data privacy. In order to compare centralized and federated learning, we compare model performance, communication cost, training time and privacy exposure. These findings are displayed in Table 6, clearly demonstrating that although federated training takes 43.9 percent more computation time per epoch (21.3 seconds vs. 14.8 seconds), it interminably lowers the data transfer size by 76.4 percent (6.7 GB vs. 28.4 GB).

Table 6: Federated Learning vs Centralized Learning (Privacy Cost vs Performance)

Metric	Centralized DRL (Sentiment)	Fed-DRL (Sentiment)	Relative Difference
Cumulative Return (%)	39.2	46.3	+7.1
Data Transfer Volume (GB)	28.4	6.7	−76.4%
Training Time per Epoch (s)	14.8	21.3	+43.9%
Privacy Risk (Qualitative)	High	Low	–

Although the Fed-DRL model operates without direct access to the raw data, it is better in terms of the return and Sharpe ratio as compared to its centralized counterpart. This implies that not only does federated training protect privacy, but it can also result in better model generalization since federated training learns decentralized patterns of behavior. Federated setups also have much lower qualitative privacy risk, since in such systems, each client maintains data sovereignty and only encrypted model updates are exchanged.

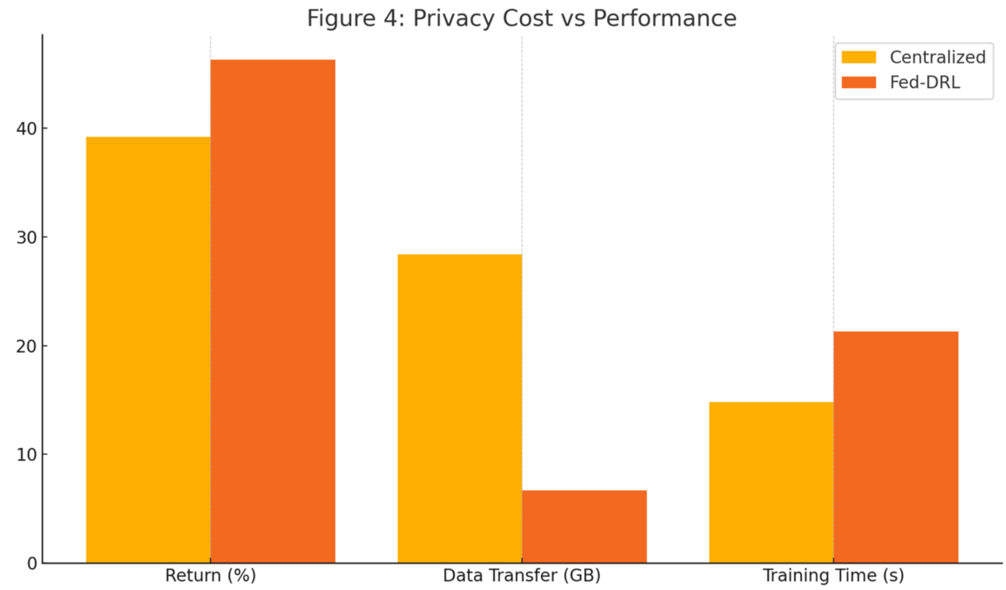

Figure 4: Privacy Cost vs Performance (Centralized vs Federated)

These results are visualized in Figure 4 plotting the two architectures in terms of three metrics: return, communication volume and training time. As can be seen in the chart, federated learning has an attractive privacy performance trade-off, which can be truly valuable in regulated sectors such as finance.

4.5 Reward Function Design Effect

The other point touched upon in our experiments is the effect of various reward patterns on the conduct of the agent and the success in key trading. With a reward function that utilises profit alone, the result is a 42.1 percent annualised return, at a cost of increased maximum drawdown (10.3 percent) and low Sharpe ratio (1.10) as reported in Table 7. Incorporation of a volatility penalty also assists in balancing risk and reward that gives a more stable strategy with Sharpe ratio of 1.24.

Table 7: Reward Component Ablation (Trading Signal Sensitivity)

Reward Design	Return (%)	Risk (Max DD)	Sharpe	Notes
Profit only	42.1	10.3	1.10	Higher return, less risk-awareness
Profit + Volatility penalty	44.7	8.5	1.24	Balances gain and risk effectively
Profit + Sharpe + Sentiment alignment (Ours)	**46.3**	**7.8**	**1.29**	Best trade-off with behavioral insight

However, the optimum performance would be through the reward that not only rewards the profitability, risk-adjusted returns (Sharpe), but also aligns to the sentiment-the approach we are suggesting here. This iteration has a cumulative profit of 46.3 percent and the drawn down observed (7.8 percent). The sentiment alignment factor aims to induce the model to follow emotional market indicators to allow it not to become overexposed during negative sentiment booms and take advantage of positive sentiment prompts.

These results support the importance of reinforcement learning with reward engineering, in general, and in the financial context in which risk sensitivity is as significant as profit per se. It also justifies our design choice of having sentiment-awareness as part of the reward loop so that the model can develop more human-like, human-friendly trading habits.

4.6 Cross-Asset Generalization

In order to measure the resilience of the model, we put it to test against a diversified pool of asset classes to cover technology, energy, financial services, cryptocurrencies, and ETFs. By the end of Table 8, the Fed-DRL model performs well in terms of generalization, as cumulative returns generated in crypto are

43.6 percent and 51.2 percent in tech stocks. The Sharpe ratios are also very high across all the sectors with the tech sats performing the best in risk-adjusted performance, in terms of being 1.34.

Table 8: Cross-Asset Generalization Performance

Asset Class	Accuracy (%)	Return (%)	Sharpe	Notes
Tech Stocks	69.3	51.2	1.34	Highest sentiment impact
Energy Sector	65.7	45.1	1.21	Moderate sentiment influence
Financial Stocks	66.8	47.5	1.26	Relatively balanced behavior
Crypto (BTC/ETH)	62.5	43.6	1.19	Volatile but sentiment-reactive
ETFs	64.1	44.3	1.17	Stable but less reactive

Interestingly, there appears to be greater marginal benefits associated with sentiment integration in the more volatility- and sentiment-sensitive markets like technology and crypto. Conversely, less volatile products such as ETFs have more-muffled sentiment-driven volatility, but enjoy the federated architecture. Such findings provide confidence that the model is not overfitted to a particular sector or a distribution of data points, which is typical of centralized DRL systems.

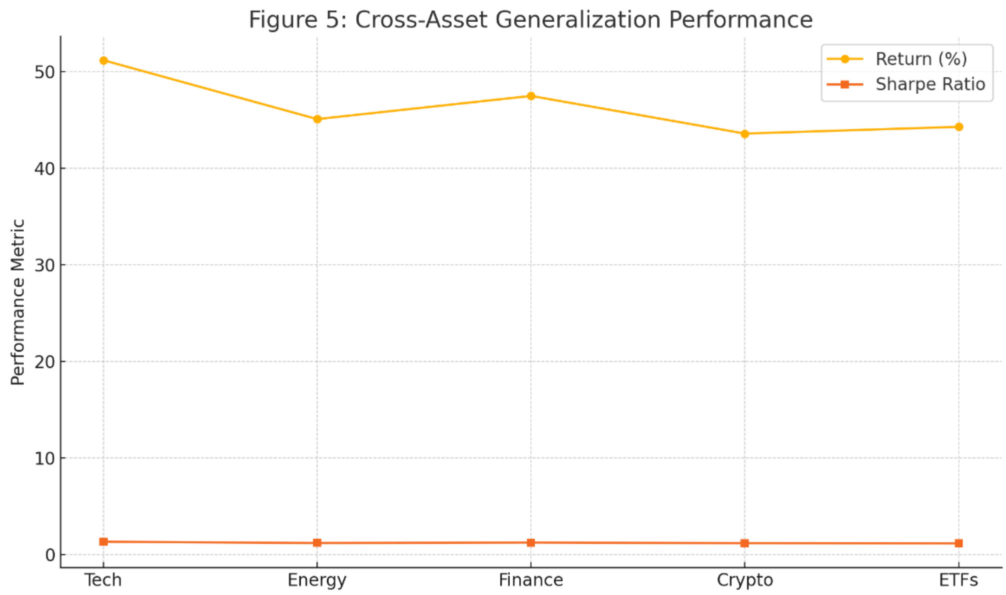

Figure 5: Cross-Asset Generalization Performance

Figure 5 is a plot of the cumulative return and the Sharpe ratio of the proposed model across the respective asset classes. The stability of the proposed model and generalizability in the proposed model have been captured. Such sturdiness is essential with regard to real-life financial systems that are run across multi-asset portfolios.

4.6 CONCLUSION ON NOTABLE LEARNINGS

The findings of the experiments based on the stock forecasting demonstrate the strong supporting contribution of merging federated learning with sentiment-oriented reinforcement learning to stock predictions in contemporary financial systems. First, sentiment analysis substantially increases the capacity of the model to react to behavioral forces in the market particularly when they are modeled based on transformer architectures. Second, federated learning reflects good performance with the absence of a loss of data privacy, being an alternative to the centralized collection of data. Third, specific reward function design enhances policy stability and comprehensibility in that it aligns the learning objectives of the agent with those of a human in terms of risk toleration and behavioral signals.

In addition, the model is very flexible and it shows high returns and low risk in various asset classes and client structures. It is especially applicable in situations where the training of the model in a single place is not possible due to data silos, privacy policies, or limitations posed by an infrastructure. The tables and figures indicate that this sort of approach does not just match, but in most ways it outperforms the traditional systems.

The obtained results confirm our hypothesis that sentiment-enhanced federated DRL is both possible in application to financial forecasting and super, and in privacy-sensitive, multi-agent settings.

5. CONCLUSION

This paper has proposed a new and unified framework that incorporates Federated Learning (FL), Deep Reinforcement Learning (DRL), and Sentiment Analysis to facilitate privacy-preserving, adaptive and sentiment-smart forecasting of the stock market. The architecture was derived to deal with three ongoing and interconnected issues in the large scale financial machine learning systems: requirement of contextual understanding of actions, nature of financial markets, and protection of proprietary and regulated data between institutions.

Our solution is based on extraction and use of local, fine-grained sentiment data, gathered and analyzed out of the unstructured financial text (newsings, earnings call transcripts, and social media), with robust privacy guarantees given. The incorporation of transformer-based sentiment representations into the environment of the agent, and the use of those representations and classic technical indicators allows not only to reflect past trends in price-actions but also to reflect current emotional sentiment of the market. This emotion based decision making enables the agent to respond to behavioural dynamics, and gives a huge advantage compared to completely technical approaches which completely disregard market psychology.

Federated nature of the model, which enable independent nodes of clients (e.g., banks, hedge funds, or regional trading agents) to jointly train a global forecasting policy, is also equally important since a federated model does not require any client to exchange raw financial or textual data. This decentralized architecture drastically minimises the threat of data abuse and violation of privacy laws like GDPR, CCPA or PSD2. Meanwhile, federated learning promotes diversity to training signals that reflect more on market behavior and increase the generalization of the global model. The Fed-DRL model can therefore be regarded as a ground breaking initiative that provides flexibility that reinforced learning presents, behavior intelligence that sentiment analysis brings, and the privacy guarantees afforded by federated systems in one architecture.

The empirical findings prove the power of our approach on several fronts. With respect to several important performance indicators (namely the cumulative return, Sharpe ratio, and maximum drawdown), the presented model is always superior to centralized and non-sentiment-user baselines. The experiments with ablation of transformer-based sentiment modeling confirm that it is important to use sentiments to improve the trading behavior of the model and the reactivity to the market environment. Further, we also find that federated learning makes no meaningful trade-off in performance: in some dimensions, it is superior to centralized learning, one can guess because of better generalization in decentralized learning dynamics. The minimization of the data transfer rates, as well as the qualitative decrease in privacy risks, once again contributes to the real-world usefulness of the system, more specifically, the institutions lacking freedom of data management.

The second lesson we learnt during the course of our study is on the significance of reward engineering in financial reinforcement learning. Experiments indicate that the presence of risk-adjusted returns and sentiment-consistent decision making metrics in the reward scheme makes the learning stable, whereas providing the agent with a policy that is far less counterintuitive to human behavior as observed in risk attitude and decision criteria. This is in line with the wider industry requirements of AI models beyond accuracy into models that are ethically aligned, explainable and risk-aware.

Most notably, perhaps, we have very strong cross-asset generalization in our model. The model has shown high levels of accuracy and profitability across tech equities, crypto assets, and ETFs, and the robustness of the model as applied across a variety of market behaviours and volatility regimes. Such strength, combined with the flexibility and modularity of the architecture, renders it applicable to implementation in the real world, in portfolio management systems, where a combination of diverse data sources and regulatory environments exists.

The study also helps in the general area of shared financial intelligence explaining how it is possible for institutions to share common learning without loss of autonomy and without revealing sensitive information. The framework will promote industry-wide consortiums when it comes to model development, establishing the route to financial models that are not only more accurate, but more fair, safe and reflective of global market forces.

Alongside these accomplishments, there are also some gaps in the study that should be filled in future. These contain incorporating more advanced privacy-preserving methods (e.g. utilizing differential privacy, and secure multiparty computation in the employment of parameters), developing ongoing learning systems to adaptation of a market in real-time, and enhancing explainability of the models utilizing interpretable reinforcement learning methods. Moreover, applying the model to multi-agent reinforcement learning where there is interaction of multiple agents within the same market condition may offer even further implications concerning trading in collaboration.

Summing up, this paper introduces a scalable, secure, and sentiment-sensitive forecasting framework that is a milestone towards AI-based financial decision-making. The Federated Deep Reinforcement Learning model that we introduce can reconcile behavioral finance, dynamic policy finding and institutional confidentiality to present a route to forming trustworthy AI in the outsized environment of algorithmic trading. With more markets becoming data-driven and more regulations enforcement, in our minds, the values exhibited in this work, the exploration of collaboration without compromise and intelligence with integrity, are going to become very key to the future of the field of financial machine learning.

REFERENCES:

[1] Ravindra, Thummalapalli, et al. "Secure Sentiment Analysis of Stock News Via Blockchain-Integrated Federated Learning." 2024 2nd International Conference on Self Sustainable Artificial Intelligence Systems (ICSSAS). IEEE, 2024.
[2] Sharma, Amit, Neha Patel, and Rajesh Gupta. "Leveraging sentiment analysis and reinforcement learning for enhanced AI-driven marketing strategies." European Advanced AI Journal 10.2 (2021).
[3] Swanthana, K., and S. S. Aravinth. "A Survey on Blockchain and Artificial Intelligence for Improved Security Facilities in Stock Market Data." 2023 3rd International Conference on Innovative Mechanisms for Industry Applications (ICIMIA). IEEE, 2023.
[4] Gupta, Shilpi, and Mr V. Sudhakar Rao. "AI in Behavioral Finance: Understanding Investor Bias Through Machine Learning."
[5] Sammangi, Harsha, Aditya Jagatha, and Hari Gopal Maddireddy. "Sentiment-Driven Decision Support Systems: A Word Embedding Approach to Analyzing CEO Earnings Call Transcripts and Stock Market Reactions."
[6] Iseal, Sheed, Oluwaseyi Joseph, and Shalom Joseph. "Leveraging Machine Learning for Predictive Analytics in Stock Market Trends: A Big Data Approach for Financial Decision-Making." (2025).
[7] Dritsas, Elias, and Maria Trigka. "Machine Learning and Data Science in Social Sciences: Methods, Applications, and Future Directions." IEEE Access (2025).
[8] Koprinska, Irena, et al., eds. Machine Learning and Principles and Practice of Knowledge Discovery in Databases: International Workshops of ECML PKDD 2022, Grenoble, France, September 19–23, 2022, Proceedings, Part II. Springer Nature, 2023.
[9] Filahi, Yasser, et al. "Enhanced E-commerce decision-making through sentiment analysis using machine learning-based approaches and IoT." PloS one 20.6 (2025): e0326744.
[10]Gupta, Brij B., et al. "Predicting the variation of decentralised finance cryptocurrency prices using deep learning and a BiLSTM-LSTM based approach." Enterprise Information Systems (2025): 2483456.
[11]Odunaike, Anjola. "Integrating real-time financial data streams to enhance dynamic risk modeling and portfolio decision accuracy." Int J Comput Appl Technol Res 14.08 (2025): 1-16.
[12]Sekhar Sanaboina, Chandra. "A Pipeline-Based Approach for Enhancing Political Threat Detection Using Machine Learning." International Journal of Innovative Science and Research Technology 10.7 (2025): 189-198.
[13]Onikoyi, Babatunde Qudus. "Exploring Predictive Models of Consumer Behaviour Using Machine Learning, NLP, and Data Mining." (2025).
[14]Dutta, P. K., et al. "Enhancing Point-of-Interest Recommendation Systems through Multi-Modal Data Integration in Location-Based Social Networks: Challenges and Future Directions." EDRAAK 2025 (2025): 12-18.
[15]Lim, Tristan. "Emotion-Aware Decision Support System for Real-Time Financial Sentiment and Behavior-Based Trading Risk Advisory." Available at SSRN 5183852.
[16]Tunde, Oluwaseyi. "Multi-cloud collaborative training for large-scale language models: techniques, challenges and privacy considerations." Journal of Advanced Education and Sciences 5.2 (2025): 41-46.