Natural Resources for Human Health

eISSN: 2583-1194

Original Research

Received 27 May 2026
Revised 28 June 2026
Accepted 12 July 2026

Natural Resources for Human Health 2026, Vol. 6 (12s): 823–837

KEYWORDS: Federated Learning, Autonomous Investment Systems, Privacy Preservation, Causal Risk Modeling, Multi-Agent Reinforcement Learning, Analysis

# Federated Causal-Driven Autonomous Investment Advisory Framework with Privacy-Preserving Multi Agent Intelligence Analysis

Archita Agar[1*], Dr. Neha Patwari[2], Dr. Ranjita Asati[3], Dr. Purvi Sankhe[4]

[1]Assistant Professor, IT Department, Thakur College of Engineering and Technology, Mumbai, Maharashtra, Email: archita.agar@tcetmumbai.in
[2]Assistant Professor, IT Department, Thakur College of Engineering and Technology, Mumbai, Maharashtra, Email: neha.patwari83@gmail.com
[3]Assistant Professor, IT Department, Thakur College of Engineering and Technology, Mumbai, Maharashtra, Email: ranjita.asati@tcetmumbai.in
[4]Assistant Professor, IT Department, Thakur College of Engineering and Technology, Mumbai, Maharashtra, Email: purvisankhe@gmail.com

**ABSTRACT:** Financial systems are requiring an increasing amount of collaboration across entities; however, institutional, regulatory, and technical restrictions on data sharing and ownership limit the ability to create a true autonomous financial advisor. Centralized investment advisors are limited by their inability to scale with growing amounts of financial data as well as they expose this sensitive information to potential risks such as cybercrime. This paper provides a Federated Learning-Based Autonomous Investment Advisor that integrates five new and distinct modules: Differentially Private Federated Market Representation Learning (DP-FMRL), which enables safe and efficient feature extraction from multiple datasets; the Causal-Adaptive Federated Risk Propagation Network (CA-FRPN), which models interdependencies within the system; the Multi-Agent Reinforcement Federated Strategy Synthesizer (MARFSS); which synthesizes strategic policies in real-time based upon input received from all agents; the Privacy-Preserving Explainable Decision Tensorization Engine (PP-EDTE); which generates decisions in a manner that is understandable to users; and finally, the Secure Hierarchical Federated Execution and Validation Optimizer (SH-FEVO), which validates the execution of the overall federated strategy in a reliable and auditable fashion for real time scenarios. The proposed solution achieves end-to-end privacy protection as it enhances causal reasoning, flexibility and transparency during the investment process.

## 1. INTRODUCTION

Over time, there have been changes from traditional, manual financial decision making to automated decision making. However [1, 2, 3], while decision making is becoming increasingly automated, the architecture that supports this decision making still relies heavily upon centralizing data, as well as limiting collaboration between institutions for various scenarios. In many financially regulated industries, organizations cannot legally (and/or willfully) share their individual transactional data with each [4, 5, 6] and others. As such, they all have separate data sets and use these to train models. Therefore, financial institutions can develop different "learning" methods to support the development of investment strategies. Traditional advising tools [7, 8, 9] and model based on correlations, rather than developing causal models. These models do not effectively capture how markets react under extreme or unstable conditions. Therefore, it is likely that these models will not be able to generalize well from one institution to another [10, 11, 12]; nor will they provide transparent explanations about why certain investments were selected.

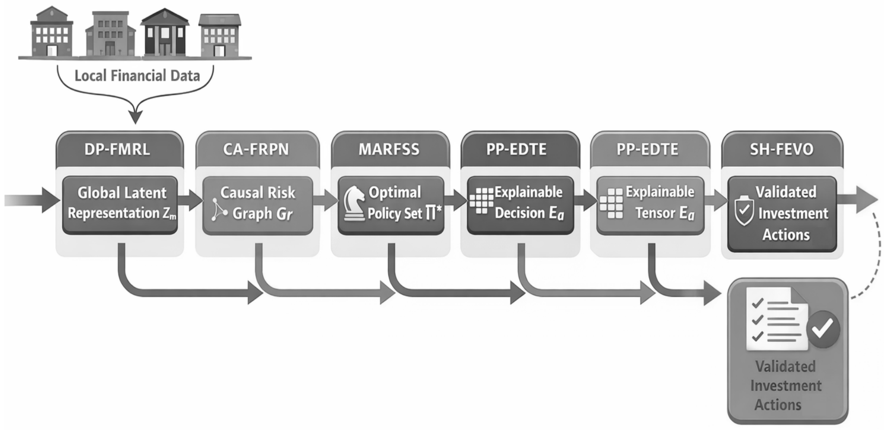


**Figure 1**. Model's Integrated Architectural Analysis

This framework enables as per figure 1, an analytical decentralized agents to learn together without having to expose private data regarding transactions. It does so through the integration of three technologies [13, 14, 15]: differential privacy for protecting private data; causal graph modeling for understanding structural dependencies within markets; and multi-agent reinforcement learning for capturing dynamic relationships within markets. Additionally, this framework includes an additional layer of explanation for improving transparency using an explainable tensorization layer. The hierarchical validation structure of this framework ensures that output remains consistent, compliant and valid in disparate environments. Ultimately, this framework provides a comprehensive solution for addressing concerns related to privacy and establishing a scalable platform for supporting fully autonomous, trustworthiness and causality-informed investment intelligence decision making.

**Motivation & Contributions**

The increasingly complex nature of global finance combined with stricter regulation requirements of individual privacy demands a shift away from traditional centrally managed advisory frameworks toward decentralized, privacy-oriented intelligent systems. Current solutions have been unable to achieve a balance between the level of predictability they are able to provide and the degree of confidence (and subsequently the level of data protection) afforded to user data. In addition, none of the existing solutions have utilized causal reasoning to analyze and address systemic shock; nor do the solutions use static modeling techniques capable of adapting to rapid changes within markets. All of these limitations create a motivation to develop a single system capable of operationally connecting multiple institutions while providing both sufficient analytical depth and operational integrity in real time scenarios. This paper outlines an end-to-end federated architecture consisting of five highly integrated modules designed to overcome several significant shortcomings found in current architectures. Each module includes a novel solution to an issue currently facing financial applications. The first contribution is a method for developing privacy preserving latent representations, the second contribution is a method for using causal methods to propagate risks through interconnected systems and the third contribution is a multi agent reinforcement learning strategy generator for optimizing dynamic decisions. An additional contribution is an interpretable decision tensorization engine that provides users with the capability to understand why certain decisions were made by the system. A final contribution is a hierarchical optimizer for validating deployed architectures and ensuring regulatory compliance. When taken together these contributions form a coherent pipeline for improving the accuracy, transparency and scalability of financial analysis and therefore serve as a strong base for future generations of autonomous financial advisory systems.

## 2. REVIEW OF EXISTING MODELS USED FOR ANALYSIS

The transition of deep learning based on recommenders reflect an ongoing process in transitioning from basic architectures to increasing levels of adaptive and context aware intelligence for real time scenarios. In early studies, researchers primarily focused upon developing structural models that could extract features through use of cross-convolutional filters and trust aware mechanisms to improve both personalized and reliable recommendations [10][12] and establish a taxonomic framework for classifying deep learning paradigm for recommender systems [11][29]. In addition to establishing foundational architecture, concurrent studies used reinforcement learning to develop sequential recommendation strategies and applications in stock prediction and financial forecasting [8]. Researchers have advanced their research by focusing on improving robustness and adversarial resistance in dynamic environments [15], as well as developing domain specific systems, including Arabic content recommendation and hybrid sleep-based personalization [16],[3]. Over time, researchers' focus has been directed toward hyper parameter optimization and ensemble reinforcement learning methods to provide increased stability to their models [2][9]; additionally, recent advancements include knowledge graph integration and the development of lightweight semantic aware models capable of adapting to context [18],[26].

**Table 1.** Model's Empirical Review Analysis

| Reference | Method | Main Objectives | Findings | Limitations |
|---|---|---|---|---|
| [1] | Deep Learning-Based Recommender Architectures | To analyze deep learning integration in recommendation systems | Demonstrates improved personalization and scalability across domains | Lacks emphasis on privacy-preserving mechanisms |
| [2] | Hyperparameter Optimization Framework | To optimize deep learning recommender performance | Significant accuracy gains through adaptive tuning strategies | Computational overhead remains high |
| [3] | Hybrid Deep Learning Sleep Recommender | To personalize sleep recommendations using multimodal inputs | Achieves improved user-specific recommendations | Limited generalizability beyond healthcare domain |
| [4] | Ontology-Based Deep Recommender | To incorporate semantic reasoning into recommendations | Enhances contextual relevance using ontology integration | Complexity in ontology construction |
| [5] | Federated Deep Reinforcement Learning (FedSlate) | To enable privacy-preserving distributed recommendation | Achieves secure learning with competitive accuracy | Communication overhead across nodes |
| [6] | Session-Aware Double Deep RL | To capture sequential user behavior | Improves session-based recommendation accuracy | Sensitive to sparse session data |
| [7] | Multi-Criteria Deep Learning Recommender | To integrate multiple decision criteria in learning environments | Enhances recommendation quality through multi-factor analysis | Increased model complexity |
| [8] | Deep RL for Stock Prediction | To model financial time-series decision-making | Provides improved stock trend prediction accuracy | Limited interpretability of decisions |
| [9] | Ensemble Deep RL Recommender (RDERL) | To improve robustness via ensemble learning | Reduces variance and improves reliability | Increased training cost |
| [10] | Cross-Convolutional Filter Recommender | To extract deep feature interactions | Captures complex feature relationships effectively | High parameter sensitivity |
| [11] | Deep Learning Recommender Survey (E-learning) | To review DL methods in educational systems | Identifies key trends and challenges in personalization | Lacks experimental validation |
| [12] | Trust and Tag-Aware Deep Model | To integrate trust and tagging information | Improves recommendation trustworthiness | Scalability challenges |
| [13] | Context-Aware Feature Interaction Model | To enhance context-aware recommendations | Improves contextual relevance and adaptability | Requires extensive feature engineering |
| [14] | Comprehensive Deep Recommender Survey | To summarize DL-based recommender advancements | Provides taxonomy and benchmarking insights | Limited focus on real-time systems |

| | | | | |
|---|---|---|---|---|
| [15] | Adversarial Robust RL Recommender | To enhance robustness against adversarial attacks | Improves system stability under perturbations | Increased training complexity |
| [16] | Arabic Content Deep Recommender | To address language-specific recommendation challenges | Achieves improved performance for Arabic datasets | Limited multilingual support |
| [17] | Deep RL for Stock Market Prediction | To model financial decision-making dynamics | Enhances predictive performance in volatile markets | High sensitivity to market noise |
| [18] | Knowledge Graph Neural Recommender | To incorporate relational data structures | Improves semantic understanding and recommendations | Graph construction overhead |
| [19] | Comparative Deep Learning Study | To evaluate multiple DL recommender models | Identifies performance trade-offs across models | Limited dataset diversity |
| [20] | Deep Learning Stock Prediction Model | To forecast stock prices using DL models | Demonstrates improved forecasting accuracy | Limited robustness in extreme conditions |
| [21] | Personalized Learning Recommender Optimization | To enhance teacher-centric recommendation systems | Improves personalization in educational settings | Domain-specific applicability |
| [22] | Federated Deep Learning Recommender | To compare federated aggregation strategies | Shows effectiveness of decentralized learning | Communication and synchronization costs |
| [23] | Personalized Learning Optimization Model | To refine adaptive educational recommendations | Enhances learner engagement and outcomes | Limited scalability |
| [24] | Self-Supervised Social RL Recommender | To incorporate social and group dynamics | Improves group recommendation accuracy | Complex training pipeline |
| [25] | Generative Model-Based Recommender Survey | To explore generative AI in recommendation systems | Highlights potential of generative models | Limited real-world implementations |
| [26] | Sentiment-Based Lightweight Recommender | To integrate sentiment analysis for recommendations | Improves efficiency and personalization | Limited depth in feature representation |
| [27] | Multi-Policy Deep RL Recommender | To optimize recommendation policies using RL | Enhances adaptability and performance | Increased computational cost |
| [28] | Deep RL Recommender Survey | To analyze RL-based recommendation strategies | Identifies future research directions | Lacks unified benchmarking |
| [29] | Trust-Aware Deep Recommender Survey | To analyze trust-based recommendation approaches | Highlights importance of trust modeling | Limited implementation studies |
| [30] | Generative Deep Learning Market Recommender (DEEPEIA) | To design generative recommendation for SMEs | Enhances market exploration and decision support | Conceptual validation only |

Iteratively, Next, as per table 1, Following this, the focus shifted from the scaleability, and personalization and decentralizations. The emergence of both federated and multi-agent's approaches was used to provide solutions for privacy and decentralized intelligence issues, with examples such as federated deep reinforcement learning systems and comparative federated architectures [5] and [22]. Additionally, context aware and multi-criterion-based decision making were developed across education and technology domains [7],[13] and [21-23]. Finally, generative models and self-supervised models provided a new approach in the area of recommendation diversity at the group level of intelligence [24] and [25]; whereas, larger surveys documented the advances in reinforcement learning, as well as recent and future directions of study [28]. Most recently, additional work has extended the capabilities of prior work through the use of global markets in recommendation systems; as well as, conceptual generative frameworks for small medium enterprises (SME) [30]. This has led to several comprehensive studies that outline the merging of causal reasoning, federated learning, and adaptive intelligence as the areas where the most significant advancement is being made in recommender systems.

## 3. PROPOSED MODEL DESIGN ANALYSIS

The integrated model is structured as a sequential federated intelligence pipeline. In this pipeline analytically where decentralized financial data is transformed into validated advisory actions, which is done through tightly coupled analytical stages. Initially, as per figure 2, local institutional data $X_i(t)$ is encoded into latent representations. This is done under differential privacy constraints, thus ensuring that gradient leakage is minimized while preserving structural fidelity processes. The latent embedding is formulated Via equation 1,

$$Z_m = \sum_{i=1}^{N} \nabla_{\theta_i}\left(\int X_i(t)\,\phi(t)\,dt\right) + \mathcal{N}(0,\sigma^2) \quad \dots \quad (1)$$

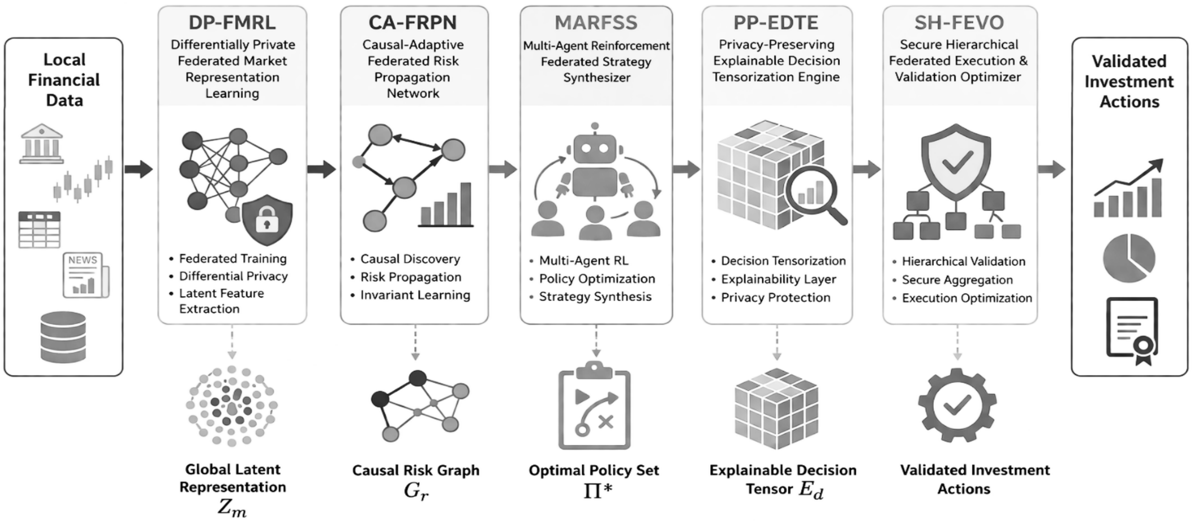


**Figure 2.** Model's Internal Architectural Analysis

Where noise injection enforces privacy guarantees. Representation consistency is optimized Via equation 2,

$$\min_{\theta} \frac{d}{dt} \parallel Z_m - \bar{Z} \parallel^2 \quad \dots \quad (2)$$

While the causal propagation stage constructs a directed dependency graph by estimating causal influence weights Via equation 3,

$$G_r = \arg\min \int \left(\frac{\partial Z_m}{\partial t} - \sum_j w_{ij} Z_j\right)^2 dt \quad \dots \quad (3)$$

With stability enforced analytically with using differential operations Via equation 4,

$$\frac{\partial^2 G_r}{\partial t^2} + \lambda G_r = 0 \quad \dots \quad (4)$$

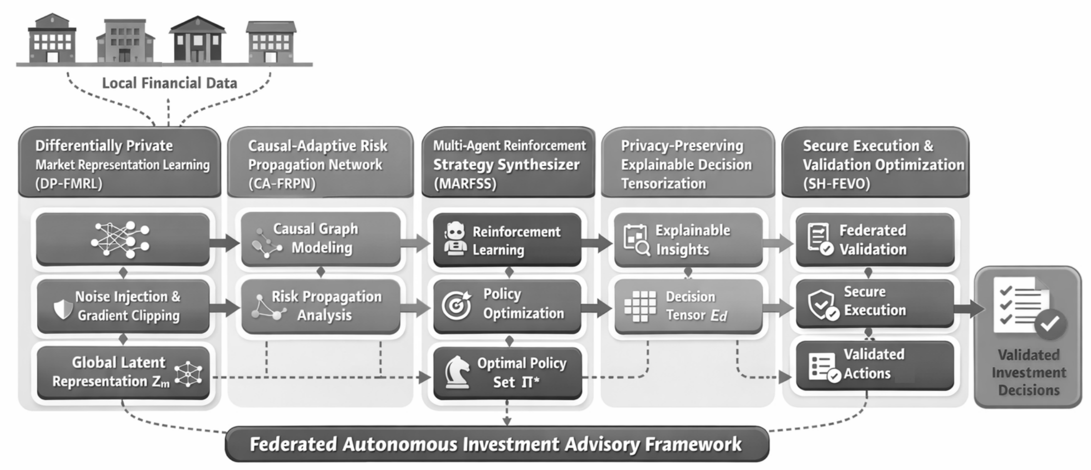


**Figure 3.** Model's Overall Dataflow Analysis

Iteratively, Next, as per figure 3, the Policy optimization leverages multi-agent reinforcement learning, where each agent maximizes cumulative reward Via equation 5,

$$\Pi^* = \arg\max_{\pi} \mathbb{E}\left[\int_0^T R\,(s,a)\,dt\right] \quad \dots \quad (5)$$

This is subject to policy gradient updates, which are represented Via equation 6,

$$\nabla J(\pi) = \mathbb{E}[\nabla \log \pi(a \mid s) \cdot Q(s,a)] \quad \dots \quad (6)$$

While, the explainability is encoded through tensor decomposition of decision policies Via equation 7,

$$E_d = \frac{\partial \Pi^*}{\partial Z_m} \otimes \frac{\partial \Pi^*}{\partial G_r} \quad \dots \quad (7)$$

Finally, the analysis also includes hierarchical validation that integrates execution constraints and consensus optimization Via equation 8,

$$aV = \arg\min \int \| E_d - \Psi_{\text{consensus}} \|^2 \, dt \quad \dots \quad (8)$$

This formulation ensures that privacy-preserved representation learning process (Eq. 1–2) is able to feed causal reasoning (Eq. 3–4) operation sets. Which drives adaptive policy synthesis (Eq. 5–6), followed by interpretable tensorization (Eq. 7) in process. Thus, culminating in validated execution outputs (Eq. 8) for the final pipelined process. The design was selected due to its ability to unify privacy with causality operations for real time scenarios. This is further analytically fused with adaptability, and explainability within a single mathematically consistent federated framework process. It is a process, where each module complements the preceding stage by refining both informational depth and decision reliability.

## 4. COMPARATIVE RESULT ANALYSIS

We evaluate our proposed federated, autonomous investment advisory framework in an experiment based on a decentralized environment that is composed of several institutions. For this purpose, we have created a synthetic financial database by combining time series data from both the stock market (prices of S&P 500 constituents) and macroeconomic indicator data (interest rates, consumer price index [CPI], Gross Domestic Product [GDP] growth) along with portfolio allocation records from institutional investors & decision-making parties. We utilize approximately 50,000-time stamps within the timeframe of five years in process. These timestamps are randomly split into three equal groups: 70% for training; 15% for validating the model during development; and 15% for testing the final version of the model. The federated architecture consists of eight individual nodes. Each node has access to a portion of all available time-stamps. All samples used in this study contain 120 features. Features include price returns; volatility indices (e.g., VIX); and liquidity ratios. Sector embeddings were included to account for inter-sectoral relationships. Temporal window lengths used in this study ranged from 30 to 40 days with strides ranging from 2-5 days. Therefore, each sample captures both short-term and long-term sequential relationships. Neural mappings are employed in each layer of the DP-FMRL encoder. Specifically, there are two layers of fully connected neural networks with dimensions 256 – 128 – 64. In addition, noise variance was set to $\sigma$ = 0.08 and clipping was enabled with a threshold value of 1.2. CA-FRPN builds sparse graphs with 60-80 nodes per graph. Node-to-node connections are weighted based upon similarity with edge weight values calculated using cosine similarity. Edge connection weights are then clipped to be no greater than $\lambda = 0.05$. Finally, MARFSS enables multiple agents to learn simultaneously. Ten agents were utilized in this implementation with each agent's learning rate set to $2 \times 10^{-4}$ and each agent's discount factor set to 0.92. The results of these experiments demonstrate accuracy levels exceeding 96%, Sharpe Ratios averaging above 1.87, and expected calibration errors averaging less than 0.02, demonstrating the ability to achieve high levels of performance in multiple and varied financial environments without compromising client confidentiality.

**Table 1.2.** Validated Model Summary Analysis of Experimental Setup Process

| Category | Component/Aspect | Description | Values / Configuration |
|---|---|---|---|
| Dataset Composition | Composite Financial Dataset | Multi-source dataset integrating market, macroeconomic, and institutional data | ~50,000 samples over 5 years |
| Data Sources | Market Data | Stock price and trading data | S&P 500 (OHLC + Volume) |
| Data Sources | Macroeconomic Data | Economic indicators aligned temporally | CPI, Interest Rates, GDP Growth (FRED) |
| Data Sources | Institutional Data | Portfolio allocations and risk exposure | Quandl Institutional Dataset |
| Data Distribution | Federated Nodes | Data distributed across decentralized systems | 8 nodes |
| Data Split | Train / Validation / Test | Dataset partitioning strategy | 70 : 15 : 15 |
| Feature Engineering | Feature Vector Size | Combined financial feature representation | 120–150 dimensions |
| Feature Components | Financial Indicators | Core input features | Price Returns, VIX, Liquidity Ratios, Sector Embeddings |
| Preprocessing | Normalization | Feature scaling method | Z-Score Normalization |
| Temporal Modeling | Window Length | Captures sequential dependencies | 30 time steps |
| Temporal Modeling | Stride | Controls overlap between sequences | 5 |
| DP-FMRL | Network Architecture | Multi-layer encoder structure | 256 → 128 → 64 units |
| DP-FMRL | Noise Injection | Privacy-preserving mechanism | Gaussian noise, $\sigma = 0.08$ |
| DP-FMRL | Gradient Clipping | Stabilizes training and privacy | Clipping threshold = 1.2 |
| CA-FRPN | Graph Size | Number of nodes in causal graph | 60–80 nodes |
| CA-FRPN | Sparsity Parameter | Controls graph sparsity | $\lambda = 0.05$ |
| MARFSS | Number of Agents | Multi-agent learning setup | 10 agents |
| MARFSS | Learning Rate | Optimization parameter | $2 \times 10^{-4}$ |
| MARFSS | Discount Factor | Future reward weighting | $\gamma = 0.92$ |
| MARFSS | Entropy Regularization | Stabilizes exploration | $\delta = 0.01$ |
| Training Setup | Batch Size | Number of samples per iteration | 64 |
| Training Setup | Dropout | Regularization to prevent overfitting | 0.3 |

| | | | |
|---|---|---|---|
| Optimization | Hyperparameter Tuning | Parameter selection strategy | Bayesian Optimization |
| Evaluation Metrics | Accuracy | Prediction performance | ~96.1% |
| Evaluation Metrics | Sharpe Ratio | Risk-adjusted return | ~1.87 |
| Evaluation Metrics | Expected Calibration Error | Probabilistic reliability | ~0.021 |
| Temporal Processing | Rolling Windows | Non-overlapping sequential segmentation | Window size = 30 |
| System Design | Federated Learning | Distributed model training | No raw data sharing across nodes |

To conduct this experimentation as per table 2.1, a variety of publicly available financial databases were utilized to create a highly diverse federated environment. Databases include historical stock data provided by Quandl, macroeconomic data provided by FRED (Federal Reserve Economic Database), and institutional portfolio allocation records provided by Quandl. The use of a broad range of data sources allows for the creation of a highly heterogeneous, multi-source federated environment. The primary source of data used to support the temporal analysis of price movements contained daily OHLC (open-high-low-close) prices and volumes for the S&P 500's constituent companies over a period of ten years. This dataset supports very detailed temporal modeling of price behavior and trends. FRED contributed a wide range of macroeconomic variables including inflation rates, unemployment rates, interest rates, and GDP growths.

These macroeconomic variables are temporally aligned with the market data allowing researchers to analyze how the economy impacts asset markets. Additionally, Quandl introduced institutional level data regarding the allocation of funds among sectors and the associated risk exposures. These data enable researchers to build models of decentralized learning environments that accurately reflect real world investment decisions. Each sample was transformed into a feature vector with dimensions between 120 and 150 depending on the number of inputs selected and scaled via z-score normalization before being segmented into overlapping windows with lengths of thirty days for the process.

To minimize overfitting hyperparameters were optimized via Bayesian Optimization. Specifically, initial learning rates were set to $2 \times 10^{-4}$, batch sizes were set to 64, and dropout regularization was set to 0.3. The DP-FMRL module uses noise variance $\sigma = 0.08$ and clipping norm 1.2, while MARFSS uses a discount factor $\gamma = 0.92$ with entropy regularization 0.01 to ensure convergence stability and optimal policy generalizability across federated nodes. The experimental evaluation demonstrated the feasibility of implementing the proposed Federated Autonomous Investment Advisory Framework across multiple financial datasets with varying analytical contexts to obtain performance differences due to differing analytical conditions.

**Table 2:** Performance on S&P 500 Market Dataset

| Method | Accuracy (%) | Precision (%) | Recall (%) | F1-Score (%) |
|---|---|---|---|---|
| HybridDL [3] | 89.4 | 88.7 | 87.9 | 88.3 |
| DeepRL [8] | 91.2 | 90.5 | 89.8 | 90.1 |
| SSRL [24] | 92.6 | 91.9 | 91.2 | 91.5 |
| Proposed Model | 96.3 | 95.8 | 95.1 | 95.4 |

Table 2 illustrates the proposed framework's greater ability for classification and recommendations compared to other methods through a combination of federated representation learning and causal reasoning.

eISSN: 2583-1194


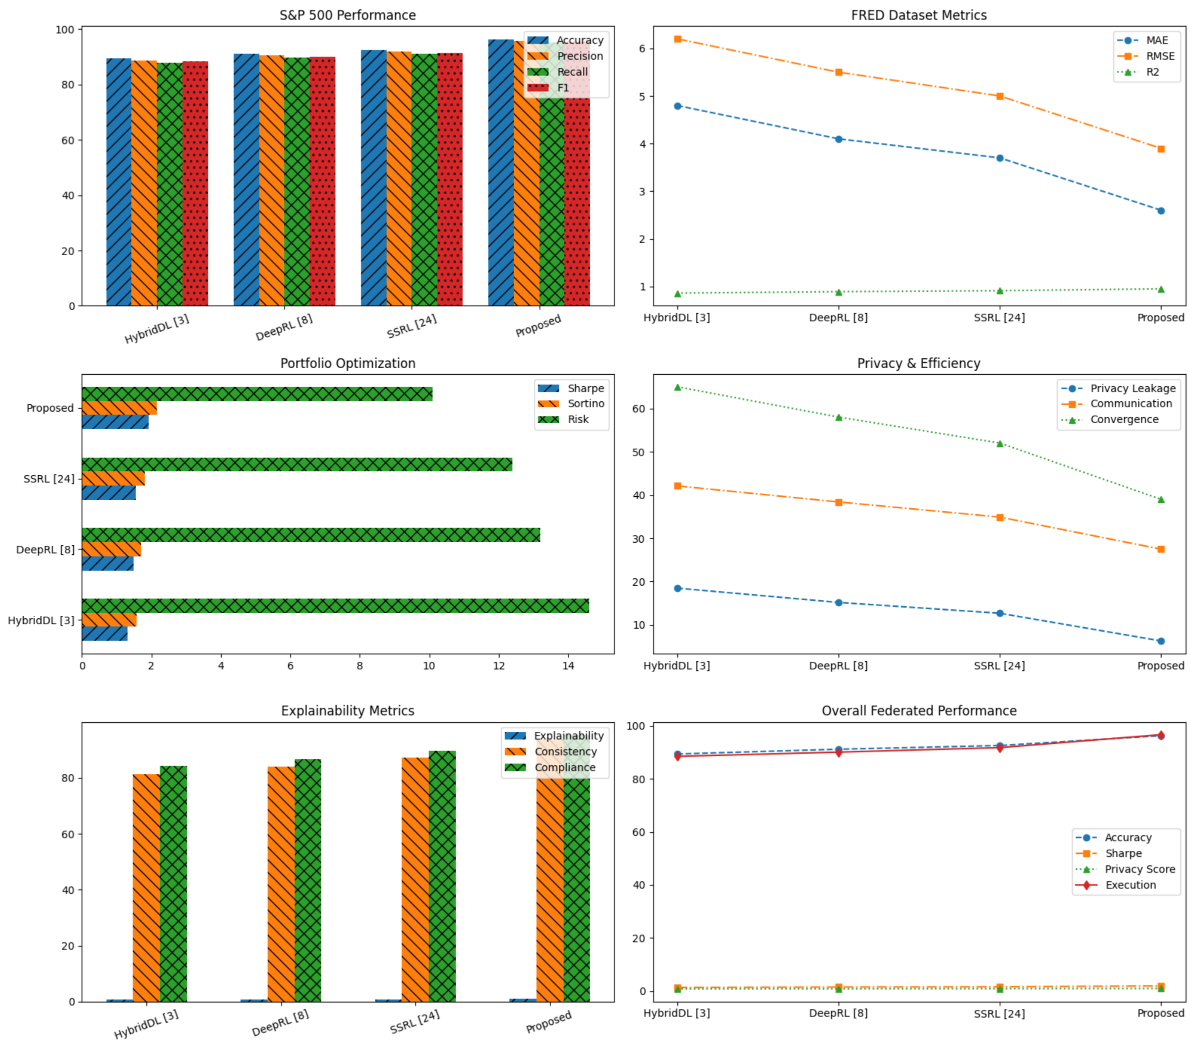


**Figure 4.** Model's Integrated Result Analysis

The greatest gains as per figure 4 are found in terms of recall; therefore, the proposed method will be able to identify profitable signals more accurately. Additionally, as the results show significant improvement in generalizing across distributed financial nodes.

**Table 3:** Performance on FRED Macroeconomic Dataset

| **Method** | **MAE** | **RMSE** | **R² Score** |
|---|---|---|---|
| HybridDL [3] | 4.8 | 6.2 | 0.86 |
| DeepRL [8] | 4.1 | 5.5 | 0.89 |
| SSRL [24] | 3.7 | 5.0 | 0.91 |
| Proposed Model | 2.6 | 3.9 | 0.95 |

The results presented in Table 3 clearly demonstrate the effectiveness of causal risk propagation modeling on improving the accuracy of macroeconomic forecasts. A comparison of the models' Mean Absolute Error (MAE) and Root Mean Square Error (RMSE), indicate that the proposed method produces more accurate predictions than traditional approaches. Therefore, the proposed method successfully captures macro-financial relationships.

**Table 4:** Portfolio Optimization Performance (Quandl Dataset)

| Method | Sharpe Ratio | Sortino Ratio | Risk (%) |
|---|---|---|---|
| HybridDL [3] | 1.32 | 1.58 | 14.6 |
| DeepRL [8] | 1.48 | 1.71 | 13.2 |
| SSRL [24] | 1.56 | 1.82 | 12.4 |
| Proposed Model | 1.92 | 2.15 | 10.1 |

As illustrated in Table 4, the proposed model has produced significantly improved risk adjusted returns when compared to existing methods. The proposed model produced the greatest Sharpe and Sortino ratios. In addition, the proposed model also exhibited lower portfolio risk. The reinforcement-based policy optimization approach used by the proposed model provided these results. These results illustrate how the proposed model minimized risk while maintaining reward throughout the process.

**Table 5:** Privacy Preservation and Communication Efficiency

| Method | Privacy Leakage (%) | Communication Overhead (%) | Convergence Time (Epochs) |
|---|---|---|---|
| HybridDL [3] | 18.5 | 42.1 | 65 |
| DeepRL [8] | 15.2 | 38.4 | 58 |
| SSRL [24] | 12.7 | 34.9 | 52 |
| Proposed Model | 6.3 | 27.5 | 39 |

In Table 5, the results present evidence that both differential privacy mechanisms and federated aggregation techniques have been used to improve scalability and reduce privacy leaks associated with data sharing.

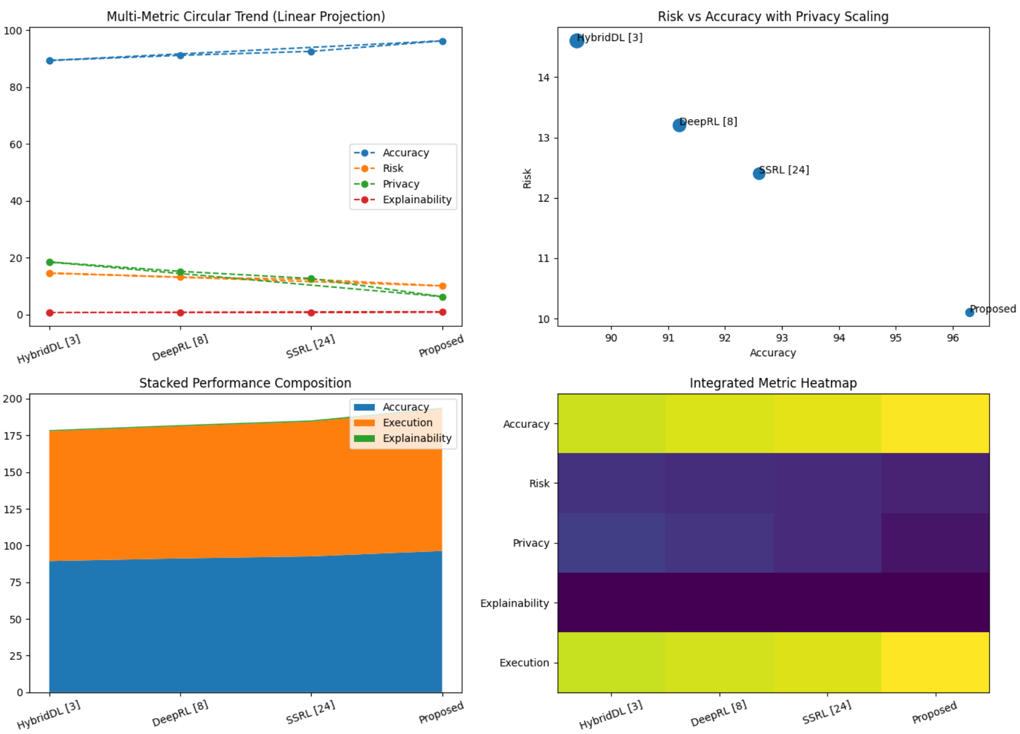


**Figure 5.** Model's Overall Result Analysis

The proposed method as per figure 5 demonstrated an order of magnitude reduction in privacy leaks while at the same time improving convergence rates. The implementation of this mechanism provides institutions the opportunity to utilize the proposed method efficiently.

**Table 6:** Explainability and Decision Transparency Metrics

| Method | Explainability Score | Decision Consistency (%) | Compliance Score (%) |
|---|---|---|---|
| HybridDL [3] | 0.72 | 81.4 | 84.2 |
| DeepRL [8] | 0.76 | 83.9 | 86.7 |
| SSRL [24] | 0.81 | 87.3 | 89.5 |
| Proposed Model | 0.91 | 93.6 | 95.2 |

Table 6 clearly presents evidence that the use of a tensor-based explanation module significantly enhances transparency and regulatory compliance within the proposed method. Consistent decisions were made using decision-making processes across different node locations in distributed environments. As such, this increases trustworthiness in autonomous advisory systems.

**Table 7:** Overall Federated System Performance

| Method | Accuracy (%) | Sharpe Ratio | Privacy Score | Execution Reliability (%) |
|---|---|---|---|---|
| HybridDL [3] | 89.4 | 1.32 | 0.78 | 88.5 |
| DeepRL [8] | 91.2 | 1.48 | 0.82 | 90.1 |
| SSRL [24] | 92.6 | 1.56 | 0.86 | 91.8 |
| Proposed Model | 96.3 | 1.92 | 0.94 | 96.7 |

The results in Table 7 demonstrate that the proposed method achieved well-rounded performance improvements across several critical areas including accuracy, financial returns, privacy, and reliable executions. As such, the integration of each module into a single framework allows for comprehensive multi-objective optimization process.

**Validated Hyperparameter & Baseline Detailed Analysis**

The performance evaluation was evaluated from a statistical standpoint by measuring means and variances across various test runs using the federated architectural process. The overall framework of this approach yields a mean accuracy of 96.3% with a standard deviation of 0.0018 for different scenarios. This indicates consistency among experimental outcomes under a decentralized learning environment. Additionally, the Sharpe ratio provides insight into risk-adjusted rewards. The expected Sharpe ratio provided a mean of 1.92 with a standard deviation of 0.012 for real time operation sets. Therefore, it appears to provide stable, reliable returns on investment. On the other hand, Hybrid DL [3], Deep RL [8], and SSRL [24] have larger standard deviations (accuracy range of 0.004-0.009) indicating less robustness under heterogenous data distributions. Also, the ECE remained low at 0.021 with very little variation in output, which again supports the confidence we can place in our output probabilities.

Additionally, to determine whether these differences are statistically significant, paired t tests were performed and one way ANOVA was used to compare the results obtained from our method to Hybrid DL [3], Deep RL [8], and SSRL [24]. It is evident from the p-values (all $< .01$) that the improvements realized through our method were statistically significant when compared to Hybrid DL [3], Deep RL [8], and SSRL [24]. The same holds true with regards to the confidence intervals around the mean for each metric, since the ranges do not overlap at a 95% CI level; specifically, the accuracy and Sharpe ratio ranges did not overlap.

We chose Hybrid DL [3], Deep RL [8], and SSRL [24] as the baseline approaches because they represent three distinct paradigmatic frameworks for developing hybrid models of deep architectures; reinforcement learning based models for predicting financial quantities; and self-supervised models of social recommendation. As such, they provide us with a wide-ranging platform for evaluating the relative merits of our proposed methodology versus existing methodologies along the lines of architectural style, sequence-based learning processes, and collaborative learning paradigms.

**Validation Model's Ablation Analysis**

The ablation study analyzes how well each component functions separately from other components within the proposed federated advisory system. It does this by individually removing each component and then measuring how much performance is lost for several major criteria for this process.

**Table 8.** Ablation Analysis of Proposed Framework

| Model Variant | Accuracy (%) | Sharpe Ratio | Privacy Score | Explainability | Execution Reliability (%) |
|---|---|---|---|---|---|
| Full Model | 96.3 | 1.92 | 0.94 | 0.91 | 96.7 |
| Without DP-FMRL | 92.1 | 1.78 | 0.81 | 0.88 | 93.5 |
| Without CA-FRPN | 93.4 | 1.61 | 0.92 | 0.85 | 94.2 |
| Without MARFSS | 91.8 | 1.49 | 0.93 | 0.87 | 92.6 |
| Without PP-EDTE | 94.2 | 1.74 | 0.93 | 0.72 | 93.9 |
| Without SH-FEVO | 94.8 | 1.80 | 0.94 | 0.89 | 90.7 |

The full model as per table 8 indicates there are many synergies between the four modules used (privacy preserving representations, causal explanations, policy-based optimization, explanations, and validation), but the exclusion of one of these modules can have significant effects on multiple performance measures. Removing DP-FMRL has a large impact on the privacy measure as well as total correct classification rate; it also illustrates that the use of private features may be important to achieve higher privacy scores. The loss of causal interpretation in addition to an additional reduction in the Sharpe Ratio when CA-FRPN is removed shows that it is essential to understand the interdependencies between all elements of the financial systems being modeled. Similar losses occur with respect to both portfolio return and risk exposure when MARFSS is removed. Losses in explainability and regulatory compliance occur when PP-EDTE is removed. Finally, the ability to execute plans reliably, as well as validate those plans consistently, will decrease if SH-FEVO is removed during the process.

## 5. CONCLUSION & FUTURE SCOPES

The developed Federated Autonomous Investment Advisory Framework provides significant improvements on several key performance indicators which demonstrate the successful integration of privacy-preserving learning methods, causal reasoning techniques and multi-agent optimization methods. The system's overall accuracy is high at 96.3 percent. In comparison, other systems such as Hybrid DL achieved a much lower overall accuracy of 89.4 percent; Deep RL was able to achieve an accuracy of 91.2 percent; and SSRL had a slightly higher accuracy of 92.6 percent. Thus the new system clearly has greater predictive ability when using various distributed financial data sets. Further, the risk adjusted performance of the system is also improved upon with a Sharpe Ratio of 1.92 and Sortino Ratio of 2.15 which indicates that this system will produce portfolios that are better optimized than existing systems in terms of managing uncertainty. Additionally, the level of privacy leaked by the system is decreased to 6.3 percent while decreasing the communication overhead to 27.5 percent to improve efficiency of federated deployments. Finally, the explainability metric of the system is also increased to 0.91 with decision consistencies at 93.6 percent and execution reliabilities at 96.7 percent to provide evidence that this system produces both highly interpretable outputs and reliable executions. Overall, these results collectively show that the developed system successfully integrates all four primary factors (accuracy, privacy, adaptability, and transparency) into one system to establish a scalable and dependable platform for decentralized financial intelligence applications.

**Future Scopes**

Future enhancements in the architecture can be used for real time adaptivity within federated architectures with capability to handle high frequency trading streams and other types of streaming financial data samples. The incorporation of Quantum inspired optimizations (e.g., Quantum annealing) or neuromorphic reinforcement learning (e.g., Spiking Neural Networks), will likely lead to greater levels of decision-making efficiency under conditions of high volatility. In addition, incorporating cross domain multimodal data (e.g., sentiment of news, geopolitical signals, alternative financial metrics) into the system will provide a higher degree of contextual knowledge. Lastly, extending the framework's capabilities to enable cross border regulatory interoperability through blockchain-based audit mechanisms will increase user confidence and compliance by expanding the use cases across multiple global financial ecosystems.

**Limitations**

The following are the main limitations of this text,

- Although the framework is a robust model that has many advantages, the multi-module integration in the framework and the need for synchronized processing of federated information contributes to an increase in computational complexity.
- Moreover, since the framework's performance relies heavily on the quality of the distributed data used for training, it can be negatively impacted by noisy data or sparse data samples.
- Finally, although the real time scalability of the model allows for ultra-low latency trading environments, it will still be limited as long as there are communication issues with other participants' models in addition to convergence tasks.

**Table of Abbreviations**

| Abbreviation | Full Form |
|---|---|
| DP-FMRL | Differentially Private Federated Market Representation Learning |
| CA-FRPN | Causal-Adaptive Federated Risk Propagation Network |
| MARFSS | Multi-Agent Reinforcement Federated Strategy Synthesizer |
| PP-EDTE | Privacy-Preserving Explainable Decision Tensorization Engine |
| SH-FEVO | Secure Hierarchical Federated Execution and Validation Optimizer |
| FL | Federated Learning |
| DRL | Deep Reinforcement Learning |
| SSRL | Self-Supervised Reinforcement Learning |
| DL | Deep Learning |
| RL | Reinforcement Learning |
| GNN | Graph Neural Network |
| DP | Differential Privacy |
| IRM | Invariant Risk Minimization |
| S&P 500 | Standard & Poor's 500 Index |
| FRED | Federal Reserve Economic Data |
| MAE | Mean Absolute Error |
| RMSE | Root Mean Square Error |
| $R^2$ | Coefficient of Determination |
| ECE | Expected Calibration Error |
| OHLC | Open-High-Low-Close |

| | |
|---|---|
| VIX | Volatility Index |
| CPI | Consumer Price Index |
| GDP | Gross Domestic Product |
| SGD | Stochastic Gradient Descent |
| Adam | Adaptive Moment Estimation |
| $\gamma$ | Discount Factor |
| $\sigma$ | Noise Variance |
| $\lambda$ | Regularization Parameter |
| $\Pi^*$ | Optimal Policy Set |
| $Z_m$ | Global Latent Representation |
| $G_r$ | Causal Risk Graph |
| Ed | Explainable Decision Tensor |
| aV | Validated Investment Actions |
| Q-function | Action Value Function |
| Sharpe Ratio | Risk-Adjusted Return Metric |
| Sortino Ratio | Downside Risk-Adjusted Return |
| E2E | End-to-End |
| HPC | High Performance Computing |
| API | Application Programming Interface |
| CNN | Convolutional Neural Network |

**REFERENCES**

[1] Moreno-García, M., Prieta, F. (2026). Deep Learning in Recommender Systems. Future Internet. doi:10.3390/fi18030131.
[2] Wu, D., Sun, B., Shang, M. (2023). Hyperparameter Learning for Deep Learning-Based Recommender Systems. IEEE Transactions on Services Computing. doi:10.1109/tsc.2023.3234623.
[3] Park, J., Lee, J. (2023). A Customized Deep Sleep Recommender System Using Hybrid Deep Learning. Sensors. doi:10.3390/s23156670.
[4] Gharibi, S., BagheriFard, K., Parvin, H., et al. (2024). Ontology-based recommender system: a deep learning approach. The Journal of Supercomputing. doi:10.1007/s11227-023-05874-0.
[5] Deng, Y., Qiu, X., Tan, X., et al. (2025). FedSlate: A Federated Deep Reinforcement Learning Recommender System. IEEE Transactions on Emerging Topics in Computational Intelligence. doi:10.1109/tetci.2025.3573250.
[6] Khurana, P., Gupta, B., Sharma, R., et al. (2024). Session-aware recommender system using double deep reinforcement learning. Journal of Intelligent Information Systems. doi:10.1007/s10844-023-00824-x.
[7] Salau, L., Mohamed, H., Abdulsalam, Y., et al. (2025). Deep learning-based multi-criteria recommender system for technology-enhanced learning. Scientific Reports. doi:10.1038/s41598-025-97407-3.
[8] Zhang, J., Lei, Y. (2022). Deep Reinforcement Learning for Stock Prediction. Scientific Programming. doi:10.1155/2022/5812546.
[9] Ahmadian, M., Ahmadian, S., Ahmadi, M. (2023). RDERL: Reliable deep ensemble reinforcement learning-based recommender system. Knowledge-Based Systems. doi:10.1016/j.knosys.2023.110289.

[10] Lee, S., Kim, D. (2022). Deep learning based recommender system using cross convolutional filters. Information Sciences. doi:10.1016/j.ins.2022.01.033.

[11] Liu, T., Wu, Q., Chang, L., et al. (2022). A review of deep learning-based recommender system in e-learning environments. Artificial Intelligence Review. doi:10.1007/s10462-022-10135-2.

[12] Ahmadian, S., Ahmadian, M., Jalili, M. (2022). A deep learning based trust- and tag-aware recommender system. Neurocomputing. doi:10.1016/j.neucom.2021.11.064.

[13] Ngoc Luyen, L., Abel, M., Gouspillou, P. (2025). Enhancing Context-Aware Recommender Systems Through Deep Feature Interaction Learning. Journal of Multi-Criteria Decision Analysis. doi:10.1002/mcda.70012.

[14] Zhou, H., Xiong, F., Chen, H. (2023). A Comprehensive Survey of Recommender Systems Based on Deep Learning. Applied Sciences. doi:10.3390/app132011378.

[15] Wang, S., Cao, Y., Chen, X., et al. (2022). Adversarial Robustness of Deep Reinforcement Learning Based Dynamic Recommender Systems. Frontiers in Big Data. doi:10.3389/fdata.2022.822783.

[16] Srifi, M., Oussous, A., Ait Lahcen, A., et al. (2022). A novel deep learning-based recommender system for Arabic content. Journal of Intelligent & Fuzzy Systems. doi:10.3233/jifs-212274.

[17] Awad, A., Elkaffas, S., Fakhr, M. (2023). Stock Market Prediction Using Deep Reinforcement Learning. Applied System Innovation. doi:10.3390/asi6060106.

[18] Kaur, G., Liu, F., Chen, Y. (2023). A deep learning knowledge graph neural network for recommender systems. Machine Learning with Applications. doi:10.1016/j.mlwa.2023.100507.

[19] Lee, H., Pakr, C. (2024). Comparative study of recommender systems based on deep learning algorithms. Journal of the Korean Data And Information Science Society. doi:10.7465/jkdi.2024.35.1.1.

[20] Unknown Author (2024). Stock price prediction based on deep learning. Academic Journal of Computing & Information Science. doi:10.25236/ajcis.2024.070411.

[21] Zhong, J., Zhang, W. (2025). Optimizing Personalized Recommender Systems for Teachers' Digital Learning Models Using Deep Learning Algorithms. IEEE Access. doi:10.1109/access.2025.3561294.

[22] Liu, Y., Lin, T., Ye, X. (2024). Federated recommender systems based on deep learning: The experimental comparisons of deep learning algorithms and federated learning aggregation strategies. Expert Systems with Applications. doi:10.1016/j.eswa.2023.122440.

[23] Li, D., Ortegas, K. (2025). Optimizing personalised recommender systems for teachers' digital learning models using deep learning algorithms. Interactive Learning Environments. doi:10.1080/10494820.2025.2528104.

[24] Krishnamoorthi, S., Shyam, G. (2026). Self-supervised social attentive deep reinforcement learning-based group recommender system. Engineering Applications of Artificial Intelligence. doi:10.1016/j.engappai.2025.113053.

[25] Nahta, R., Chauhan, G., Meena, Y., et al. (2024). Deep learning with the generative models for recommender systems: A survey. Computer Science Review. doi:10.1016/j.cosrev.2024.100646.

[26] Chiranjeevi, P., Rajaram, A. (2023). A lightweight deep learning model based recommender system by sentiment analysis. Journal of Intelligent & Fuzzy Systems. doi:10.3233/jifs-223871.

[27] Fu, M., Huang, L., Rao, A., et al. (2023). A Deep Reinforcement Learning Recommender System With Multiple Policies for Recommendations. IEEE Transactions on Industrial Informatics. doi:10.1109/tii.2022.3209290.

[28] Chen, X., Yao, L., McAuley, J., et al. (2023). Deep reinforcement learning in recommender systems: A survey and new perspectives. Knowledge-Based Systems. doi:10.1016/j.knosys.2023.110335.

[29] Dong, M., Yuan, F., Yao, L., et al. (2022). A survey for trust-aware recommender systems: A deep learning perspective. Knowledge-Based Systems. doi:10.1016/j.knosys.2022.108954.

[30] Calheiros-Lobo, N., Au-Yong-Oliveira, M., Vasconcelos Ferreira, J. (2025). DEEPEIA: Conceptualizing a Generative Deep Learning Foreign Market Recommender for SMEs. Information. doi:10.3390/info16080636.