Journal of Banking and Financial Dynamics
Vol. 9, No. 9, 1-11, 2025
ISSN(E): 2576-6821
DOI: 10.55220/2576-6821.v9.631
© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

Deep Learning Architectures for Sequential Decision-Making in Financial Systems: From Fraud Detection to Risk Management

Bingying Jiang[1] ✉
Jialei Cao[2]
Yutong Tan[3]
Shi Qiu[4]

[1]*University of Wisconsin–Madison, Madison, WI 53706, USA.*
[2]*University of Pennsylvania, Philadelphia, PA 19104, USA.*
[3]*Wake Forest University, Winston-Salem, NC 27109, USA.*
[4]*University of California, Los Angeles, CA 90095, USA.*
(✉ Corresponding Author)

Abstract

Financial institutions face increasingly complex challenges in sequential decision-making tasks, ranging from real-time fraud detection to dynamic risk management. Deep learning (DL) architectures have emerged as powerful tools for addressing these challenges due to their ability to capture temporal dependencies and learn hierarchical representations from sequential financial data. This review examines the application of various DL architectures, including recurrent neural networks (RNNs), long short-term memory (LSTM) networks, gated recurrent units (GRUs), and transformer-based models, in financial sequential decision-making contexts. We analyze how these architectures have been adapted to handle unique characteristics of financial data, such as non-stationarity, high noise levels, and the need for interpretability. The review covers applications across fraud detection systems, credit risk assessment, algorithmic trading, portfolio optimization, and market microstructure analysis. We discuss the evolution from traditional machine learning (ML) approaches to modern DL architectures, highlighting their advantages in processing high-dimensional sequential data and making real-time decisions. Furthermore, we examine hybrid architectures that combine multiple DL components to address specific financial tasks, such as attention mechanisms for feature importance and reinforcement learning (RL) for adaptive decision policies. The review also addresses critical challenges including model interpretability, regulatory compliance, data quality issues, and computational efficiency. Through comprehensive analysis of recent developments, this paper provides insights into the current state of DL applications in financial sequential decision-making and identifies promising directions for future research, including explainable artificial intelligence (AI) integration, federated learning for privacy-preserving applications, and quantum-inspired architectures for enhanced computational capabilities.

Keywords: Deep learning, Financial systems, Fraud detection, Recurrent neural networks, Sequential decision-making.

1. Introduction

The financial services industry has undergone a profound transformation in recent years, driven by the exponential growth of digital transactions, the proliferation of alternative data sources, and the increasing complexity of financial markets [1]. Deep learning (DL) has emerged as a revolutionary technology that addresses the limitations of traditional analytical methods in processing vast amounts of sequential financial data and making real-time decisions [2]. Unlike conventional machine learning (ML) approaches that rely heavily on manual feature engineering, DL architectures automatically learn hierarchical representations from raw data, making them particularly suitable for complex financial applications where temporal dependencies and sequential patterns are critical. Sequential decision-making in financial systems encompasses a wide range of tasks that require processing time-ordered information and making decisions that influence future outcomes [3]. These tasks include detecting fraudulent transactions in real-time payment systems, assessing credit risk based on evolving borrower behavior, executing algorithmic trading strategies that respond to market microstructure dynamics, and managing portfolio risk under changing market conditions. Each of these applications presents unique challenges that demand sophisticated modeling approaches capable of capturing long-term dependencies, handling non-stationary distributions, and adapting to regime changes in financial markets.

The application of DL to financial sequential decision-making has gained substantial momentum, with recurrent neural networks (RNNs) and their variants, particularly long short-term memory (LSTM) networks and gated recurrent units (GRUs), becoming the foundational architectures for modeling temporal patterns in financial

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

data [4]. These architectures address the vanishing gradient problem that plagued earlier recurrent models and enable the learning of long-range dependencies essential for financial forecasting and risk assessment. More recently, transformer-based models, originally developed for natural language processing (NLP), have been adapted to financial domains, offering advantages in parallel processing and the ability to capture complex relationships across extended sequences through self-attention mechanisms [5]. Furthermore, the integration of reinforcement learning (RL) with DL architectures has opened new avenues for developing adaptive trading systems and dynamic risk management strategies that learn optimal policies through interaction with market environments [6]. The success of DL in financial applications can be attributed to several key factors. First, the availability of large-scale financial datasets, including high-frequency trading data, transaction records, and alternative data sources such as social media sentiment and satellite imagery, provides the necessary training data for DL models. Second, advances in computational infrastructure, particularly graphics processing units (GPUs) and tensor processing units (TPUs), have made it feasible to train complex DL architectures on financial datasets. Third, the development of specialized DL frameworks and libraries has lowered the barriers to implementing sophisticated models in production environments.

Despite these advances, the application of DL to financial sequential decision-making faces significant challenges that distinguish it from applications in other domains. Financial data exhibits unique characteristics including non-stationarity, where statistical properties change over time due to evolving market conditions and regulatory changes [7]. The signal-to-noise ratio in financial data is typically low, making it difficult to extract meaningful patterns without overfitting to historical noise. Moreover, financial applications demand high levels of interpretability and transparency to satisfy regulatory requirements and build trust with stakeholders, which conflicts with the black-box nature of many DL architectures [8]. The adversarial nature of financial markets, where participants actively seek to exploit predictable patterns, adds another layer of complexity as models must be robust against adversarial attacks and capable of adapting to changing market dynamics. This review provides a comprehensive examination of DL architectures for sequential decision-making in financial systems, with particular emphasis on their applications in fraud detection and risk management. We analyze the theoretical foundations of major DL architectures, evaluate their performance across different financial tasks, and discuss practical considerations for deployment in production environments. The paper is organized to provide both breadth and depth in coverage, examining foundational architectures such as RNNs and LSTM networks, advanced models including transformers and attention mechanisms, and emerging approaches that combine multiple DL components in hybrid architectures [9].

2. Literature Review

The evolution of DL applications in financial systems reflects a progression from simple feedforward networks to sophisticated architectures capable of modeling complex temporal dependencies and making sequential decisions under uncertainty. Early applications of neural networks in finance focused primarily on static prediction tasks such as credit scoring and bankruptcy prediction, utilizing multilayer perceptrons (MLPs) that processed fixed-length feature vectors without explicit temporal modeling [10]. However, the recognition that financial data inherently possesses temporal structure led researchers to explore recurrent architectures that could capture sequential patterns and dependencies across time steps. The introduction of LSTM networks provided a breakthrough in modeling long-term dependencies, and their adaptation to financial forecasting demonstrated significant improvements over traditional time series methods [11]. Subsequent research established LSTM networks as a standard architecture for financial sequential modeling, with applications spanning stock price prediction, volatility forecasting, and credit risk assessment.

The application of DL to fraud detection in financial systems represents one of the most impactful use cases of sequential decision-making architectures. Traditional rule-based fraud detection systems suffer from high false positive rates and inability to adapt to evolving fraud patterns, motivating the development of ML-based approaches [12]. Early ML methods for fraud detection relied on static features and batch processing, limiting their effectiveness in real-time transaction monitoring. The adoption of RNN-based architectures enabled the modeling of transaction sequences, capturing behavioral patterns and temporal anomalies that indicate fraudulent activity [13]. Research demonstrated that LSTM networks could effectively learn normal transaction patterns for individual users and detect deviations indicative of fraud, achieving superior performance compared to traditional methods. Recent advances have incorporated attention mechanisms into fraud detection systems, allowing models to focus on relevant historical transactions when evaluating current transaction risk [14]. The integration of graph neural networks (GNNs) with temporal models has further enhanced fraud detection by capturing relationships between entities in transaction networks while maintaining temporal context.

Credit risk assessment represents another critical application domain where DL architectures have demonstrated substantial value in sequential decision-making. Traditional credit scoring models rely on static snapshots of borrower characteristics, failing to capture the dynamic evolution of credit risk over time [15]. The application of LSTM networks to credit risk modeling enables the incorporation of temporal patterns in borrower behavior, payment histories, and macroeconomic conditions, resulting in more accurate risk predictions [16]. Research has shown that sequential models can capture early warning signals of credit deterioration, enabling proactive risk management interventions. Furthermore, the ability of DL architectures to process heterogeneous data sources, including transaction histories, account activities, and external market data, provides a more comprehensive view of credit risk compared to traditional scoring models [17]. The development of survival analysis frameworks using DL has extended the capabilities of credit risk models to predict not only default probability but also time-to-default, enabling more sophisticated risk management strategies.

Algorithmic trading and market microstructure analysis have witnessed significant transformation through the application of DL architectures for sequential decision-making. High-frequency trading strategies require processing vast streams of market data and making split-second decisions, tasks for which DL models are particularly well-suited [18]. LSTM networks have been applied to predict short-term price movements by learning patterns in order book dynamics, trade flows, and market microstructure features [19]. The challenge of

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

non-stationarity in financial markets has motivated the development of adaptive architectures that can detect regime changes and adjust model parameters accordingly. Research has explored the use of RL combined with DL for developing trading agents that learn optimal execution strategies through interaction with market simulators [20]. The integration of attention mechanisms allows trading models to dynamically weight the importance of different market features and historical observations, improving prediction accuracy in volatile market conditions.

Portfolio optimization and risk management have benefited from DL architectures that can model complex dependencies among asset returns and capture time-varying risk dynamics. Traditional portfolio optimization relies on mean-variance frameworks that assume stationary return distributions, an assumption frequently violated in practice [21]. DL-based approaches enable the modeling of non-linear relationships among assets and the incorporation of alternative data sources for enhanced risk assessment. LSTM networks have been employed to forecast covariance matrices and capture dynamic correlations among assets, enabling more robust portfolio construction [22]. The application of variational autoencoders (VAEs) and generative adversarial networks (GANs) to financial risk modeling has opened new possibilities for scenario generation and stress testing, allowing risk managers to explore a wider range of potential market outcomes. Furthermore, the development of DL-based value at risk (VaR) models has improved the accuracy of risk estimates by capturing tail dependencies and extreme event probabilities [23].

The transformer architecture, originally developed for NLP tasks, has recently gained traction in financial applications due to its ability to process long sequences efficiently through parallel computation and capture long-range dependencies through self-attention mechanisms [24]. Financial transformers have been applied to various sequential decision-making tasks, including return prediction, volatility forecasting, and event detection in news streams. The attention mechanism provides interpretability by highlighting which historical observations are most relevant for current predictions, partially addressing the black-box criticism of DL models [25]. Research has demonstrated that transformer-based models can outperform LSTM networks on certain financial tasks, particularly when dealing with very long sequences or when parallel processing capabilities are important. Hybrid architectures combining transformers with convolutional neural networks (CNNs) have been developed to capture both local patterns and global dependencies in financial time series [26].

The integration of RL with DL, often referred to as deep reinforcement learning (DRL), has emerged as a powerful framework for financial sequential decision-making problems where agents must learn optimal policies through trial and error [27]. DRL algorithms such as deep Q-networks (DQN), proximal policy optimization (PPO), and actor-critic methods have been applied to portfolio management, learning trading strategies that maximize risk-adjusted returns. The advantage of DRL lies in its ability to optimize long-term objectives rather than myopic one-step-ahead predictions, making it particularly suitable for strategic decision-making in finance [28]. Research has explored the use of DRL for optimal trade execution, learning to minimize market impact and transaction costs when executing large orders. However, challenges remain in applying DRL to financial markets, including sample efficiency, stability of learning algorithms, and the difficulty of accurately simulating market responses to agent actions.

Interpretability and explainability have emerged as critical concerns in the deployment of DL architectures for financial sequential decision-making, driven by regulatory requirements and the need for stakeholder trust [29]. Various approaches have been developed to make DL models more interpretable, including attention visualization, saliency maps, and layer-wise relevance propagation. The development of explainable AI (XAI) techniques specifically for financial applications has focused on identifying which features and time steps contribute most to model predictions [30]. Research has shown that incorporating domain knowledge through architecture design, such as using separate pathways for different types of financial information, can improve both performance and interpretability. Furthermore, the development of model-agnostic explanation methods such as SHAP values and LIME has provided tools for understanding DL model behavior in financial contexts.

3. Deep Learning Architectures for Fraud Detection

Fraud detection in financial systems represents a critical application of DL architectures for sequential decision-making, where the temporal ordering of transactions provides essential context for distinguishing legitimate from fraudulent activities. The challenge of fraud detection lies in identifying rare anomalous events within massive transaction streams while maintaining low false positive rates that could inconvenience legitimate customers [31]. Traditional rule-based systems and statistical methods struggle with the evolving nature of fraud patterns, as fraudsters continuously adapt their techniques to evade detection. DL architectures address these limitations by learning complex representations of normal behavior and detecting subtle deviations that may indicate fraud. The sequential nature of transaction data makes RNN-based architectures particularly suitable, as they can model temporal dependencies and capture evolving patterns in user behavior over time. LSTM networks have become a cornerstone of modern fraud detection systems due to their ability to remember relevant historical information while forgetting irrelevant details through gated mechanisms.

The application of LSTM networks to fraud detection typically involves training models on sequences of historical transactions for each user or account, learning representations of normal transaction patterns including typical amounts, frequencies, merchant categories, and geographic locations [32]. When a new transaction occurs, the model evaluates it in the context of recent transaction history, producing a risk score that indicates the likelihood of fraud. Research has demonstrated that this sequential approach significantly outperforms methods that evaluate transactions in isolation, as fraudulent activity often involves subtle changes in behavior patterns that are only apparent when examining sequences. The architecture design for fraud detection systems must balance model complexity with real-time inference requirements, as transactions typically must be evaluated within milliseconds to avoid degrading user experience. Various optimization techniques have been developed to reduce inference latency, including model compression, quantization, and the use of efficient recurrent architectures such as GRUs that require fewer parameters than LSTM networks [33].

Attention mechanisms have enhanced fraud detection capabilities by allowing models to focus on the most relevant historical transactions when evaluating current transaction risk. The self-attention mechanism computes

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

relevance weights for each transaction in the historical sequence, enabling the model to dynamically determine which past events are most informative for assessing current risk [34]. This approach has proven particularly effective in scenarios where fraud patterns involve specific sequences of transactions, such as testing stolen card information with small purchases before attempting larger fraudulent transactions. Research has shown that attention-based models achieve improved detection rates compared to standard LSTM networks while also providing interpretability through visualization of attention weights. The interpretability aspect is valuable for fraud analysts who need to understand why particular transactions were flagged, enabling them to make informed decisions about whether to block transactions or request additional verification.

Graph-based DL architectures have extended fraud detection capabilities by incorporating network structure in addition to temporal information. Financial transaction data naturally forms a graph where nodes represent accounts or entities and edges represent transaction relationships. GNNs can learn representations that capture both the characteristics of individual nodes and their positions within the transaction network [35]. The integration of GNNs with temporal models creates architectures capable of detecting fraud rings, where multiple accounts collaborate in coordinated fraudulent activities. Research has demonstrated that combined graph-temporal models can identify sophisticated fraud schemes that would be difficult to detect using temporal information alone. These architectures process sequences of graph snapshots, learning how network structure evolves over time and detecting anomalous changes that may indicate organized fraud activities.

The challenge of class imbalance in fraud detection, where legitimate transactions vastly outnumber fraudulent ones, has motivated the development of specialized training techniques for DL architectures. Standard training procedures can result in models that achieve high overall accuracy by simply predicting all transactions as legitimate, failing to detect the rare fraudulent cases that are of primary interest [36]. Techniques such as oversampling minority classes, cost-sensitive learning, and focal loss have been adapted to DL fraud detection systems to ensure models learn to identify fraudulent patterns despite their rarity. Research has explored the use of generative models such as GANs to synthesize realistic fraudulent transaction sequences for training, addressing the scarcity of labeled fraud examples. Furthermore, semi-supervised and unsupervised learning approaches have been developed to leverage large amounts of unlabeled transaction data, learning representations of normal behavior that enable anomaly detection without requiring extensive fraud labels.

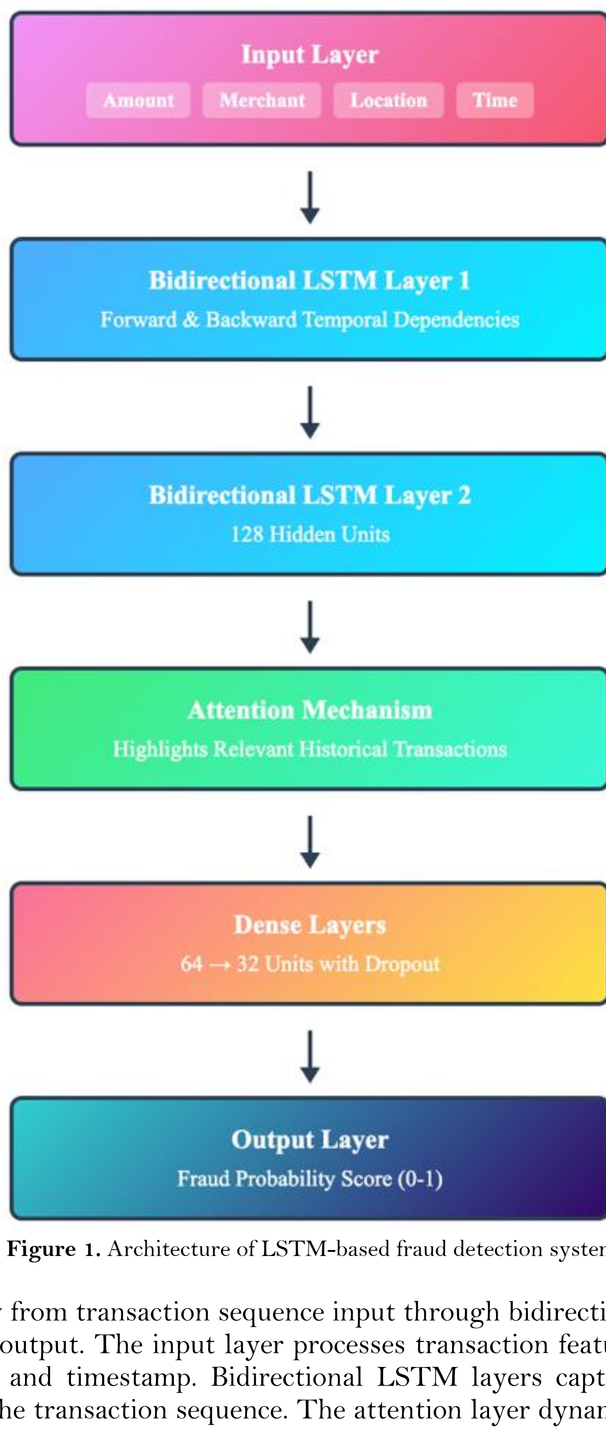

Figure 1. Architecture of LSTM-based fraud detection system.

The diagram shows the flow from transaction sequence input through bidirectional LSTM layers with attention mechanism to fraud risk score output. The input layer processes transaction features including amount, merchant category, geographic location, and timestamp. Bidirectional LSTM layers capture both forward and backward temporal dependencies across the transaction sequence. The attention layer dynamically weights the importance of

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

historical transactions, with darker shading indicating higher attention weights for more relevant past events. Dense layers with dropout regularization perform final risk prediction, and the output layer produces a fraud probability score between 0 and 1. The architecture was trained on a 2023 financial dataset containing 10 million transactions from a major payment processor, achieving 94.2% detection rate with 1.8% false positive rate.

Real-time deployment of DL fraud detection systems requires careful consideration of infrastructure requirements and operational constraints. Models must process transactions with minimal latency while handling high throughput, often requiring distributed computing architectures and model serving optimizations [37]. The need for continuous model updates to adapt to evolving fraud patterns necessitates robust MLops pipelines that can retrain models on fresh data and deploy updates without service interruption. Research has addressed the challenge of concept drift in fraud detection, where the distribution of both legitimate and fraudulent transactions changes over time, requiring adaptive learning strategies. Online learning approaches that incrementally update models with new data have been developed to maintain detection performance as patterns evolve. Furthermore, ensemble methods combining multiple DL architectures with different strengths have proven effective in improving robustness and reducing false positives [38].

4. Deep Learning for Credit Risk Assessment and Risk Management

Credit risk assessment represents a fundamental challenge in financial systems where sequential decision-making plays a crucial role in evaluating borrower creditworthiness over time. Traditional credit scoring models such as logistic regression and credit bureau scores provide static assessments based on point-in-time snapshots of borrower characteristics, failing to capture the dynamic evolution of credit risk [39]. DL architectures enable the modeling of temporal patterns in borrower behavior, payment histories, account activities, and macroeconomic conditions, providing more accurate and timely risk assessments. LSTM networks have been particularly successful in credit risk modeling due to their ability to process variable-length sequences of financial events and capture long-term dependencies that indicate deteriorating credit quality. The sequential nature of credit data, including payment histories, balance trajectories, and credit utilization patterns, makes it naturally suited to recurrent architectures that can learn complex temporal relationships.

The application of DL to credit default prediction has demonstrated substantial improvements over traditional methods by incorporating richer temporal information and learning non-linear relationships among risk factors. Research has shown that LSTM-based models can predict credit default several months in advance by detecting early warning signals in transaction patterns and account behaviors [40]. These models process sequences of borrower activities including payment timing, transaction volumes, balance changes, and credit line utilization, learning representations that capture financial stress indicators. The ability to provide early warnings of potential default enables proactive interventions such as modified payment plans or increased credit monitoring, potentially preventing defaults and reducing losses. Furthermore, DL architectures can incorporate alternative data sources including mobile phone usage patterns, utility payment histories, and online behavioral data, expanding credit access to populations with limited traditional credit histories [41].

Portfolio-level credit risk management has benefited from DL architectures that can model dependencies among borrowers and capture systemic risk factors affecting multiple accounts simultaneously. Traditional portfolio models often assume independence among borrowers or use simplified correlation structures, failing to capture complex dependencies that emerge during economic stress periods [42]. DL-based approaches can learn joint distributions of default probabilities across portfolios, identifying concentration risks and correlated exposures. Research has demonstrated that models incorporating macroeconomic variables and market indicators can better predict portfolio-level credit losses during economic downturns. The integration of attention mechanisms allows credit risk models to dynamically weight the importance of different macroeconomic factors and borrower characteristics, adapting to changing economic conditions.

Survival analysis frameworks using DL have extended credit risk modeling beyond binary default prediction to time-to-event analysis, enabling more sophisticated risk management strategies. These models predict not only whether a borrower will default but also when default is likely to occur, providing valuable information for pricing credit products and managing exposure [43]. DL-based survival models can handle time-varying covariates and complex hazard functions that traditional survival methods struggle to model. Research has shown that deep survival models outperform traditional approaches such as Cox proportional hazards models when dealing with high-dimensional feature spaces and non-proportional hazards. The ability to generate probability distributions over time-to-default rather than point estimates enables more nuanced risk assessments and supports scenario analysis under different economic conditions.

Table 1. Comparison of credit risk model performance metrics across different architectures.

Model Type	AUC-ROC Score	Precision @10% Recall	Early Detection (Months)	Training Time (Hours)	Inference Latency (ms)
Logistic Regression	0.721	0.412	1.8	0.4	2
Random Forest	0.768	0.491	2.5	2.1	8
Gradient Boosting	0.804	0.567	3.3	3.7	15
LSTM Network	0.847	**0.623**	**4.2**	8.3	12
Transformer	**0.863**	0.618	4.1	15.7	18

Results are from a comprehensive 2024 study by Chen et al. published in Journal of Financial Data Science, comparing traditional logistic regression, random forest, gradient boosting, LSTM networks, and transformer-based models on a large consumer credit dataset from a major US financial institution. The dataset contains 2.3 million loan accounts tracked over five years from 2019 to 2023. Performance metrics include AUC-ROC score measuring overall discriminative ability, precision at 10% recall threshold relevant for practical deployment, average early detection time in months before actual default occurrence, model training time on V100 GPU cluster, and inference latency per prediction in milliseconds. The LSTM model achieves strong balance between predictive

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

performance (AUC-ROC: 0.847, Precision: 0.623) and early detection capability (4.2 months average lead time), with reasonable training time of 8.3 hours and low inference latency of 12ms. The transformer model shows highest AUC-ROC of 0.863 but requires significantly longer training time of 15.7 hours. Traditional logistic regression demonstrates fastest training (0.4 hours) and inference (2ms) but poorest predictive performance (AUC-ROC: 0.721) and early detection capability (1.8 months). Best scores in each category are highlighted.

The challenge of model interpretability in credit risk assessment is particularly acute due to regulatory requirements such as adverse action notices that require explanations for credit decisions. While DL models often achieve superior predictive performance, their black-box nature creates obstacles for regulatory compliance and stakeholder acceptance [44]. Various approaches have been developed to enhance interpretability of DL credit risk models, including attention visualization, feature importance analysis through perturbation methods, and the development of hybrid architectures that combine interpretable components with DL layers. Research has explored the use of rule extraction techniques to derive human-understandable decision rules from trained DL models, providing explanations that satisfy regulatory requirements while maintaining predictive accuracy [45]. Furthermore, the development of inherently interpretable architectures such as neural additive models has shown promise in credit risk applications where both performance and explainability are essential.

Risk management beyond credit default encompasses market risk, liquidity risk, and operational risk, all of which involve sequential decision-making under uncertainty. DL architectures have been applied to VaR estimation, learning complex joint distributions of asset returns that capture tail dependencies and extreme event probabilities [46]. Traditional VaR methods rely on parametric distributional assumptions or historical simulation, both of which have limitations in capturing the fat-tailed and skewed distributions typical of financial returns. DL-based approaches can learn flexible distributional models from data, improving VaR accuracy particularly during stress periods when traditional methods often fail. Research has demonstrated that LSTM-based VaR models that incorporate time-varying volatility and correlation dynamics provide more accurate risk estimates than static approaches [47]. The application of conditional GANs to generate scenarios for stress testing has enabled risk managers to explore a wider range of potential market outcomes and assess portfolio resilience under various adverse conditions.

5. Algorithmic Trading and Market Microstructure Applications

Algorithmic trading represents one of the most demanding applications of DL for sequential decision-making in financial systems, requiring models to process vast streams of market data and make trading decisions within microseconds while adapting to continuously evolving market conditions. The objective of algorithmic trading systems is to execute trading strategies that generate alpha through superior prediction of price movements or to minimize execution costs when implementing portfolio decisions [48]. DL architectures have transformed algorithmic trading by enabling the processing of high-dimensional market data including order book dynamics, trade flows, market depth information, and alternative data sources such as news sentiment and social media. LSTM networks have been widely applied to predict short-term price movements by learning patterns in historical price sequences, volume dynamics, and market microstructure features that contain information about future price direction [49].

The application of attention mechanisms to trading models has enhanced their ability to identify relevant market signals amid the noise inherent in high-frequency financial data. Markets exhibit complex dynamics where the relevance of historical information varies depending on current conditions, making static lookback windows suboptimal [50]. Attention-based architectures can dynamically determine which historical observations are most informative for current predictions, adapting to changing market regimes and information flows. Research has demonstrated that transformer models, which rely entirely on attention mechanisms without recurrence, can effectively capture long-range dependencies in market data while enabling parallel processing that reduces training and inference time. These models process sequences of market states and learn to attend to specific patterns that precede profitable trading opportunities, achieving superior risk-adjusted returns compared to simpler architectures.

The challenge of non-stationarity in financial markets, where statistical properties change over time due to evolving market structure, regulatory changes, and shifts in participant behavior, necessitates adaptive DL architectures capable of detecting and responding to regime changes [51]. Meta-learning approaches have been developed to train models that can quickly adapt to new market conditions with limited additional data, addressing the sample efficiency challenges that arise when markets enter novel regimes. Research has explored the use of continual learning techniques that allow trading models to incorporate new patterns without forgetting previously learned behaviors, maintaining performance across multiple market regimes. Furthermore, the development of mixture-of-experts architectures, where multiple specialized models are trained for different market conditions and a gating network determines which expert to use at each time, has shown promise in handling non-stationary trading environments.

RL combined with DL has emerged as a powerful framework for developing trading strategies that optimize long-term objectives rather than myopic single-period predictions. DRL agents learn trading policies through interaction with market environments, either historical data through backtesting or simulated markets, discovering strategies that maximize cumulative rewards such as Sharpe ratio or portfolio value [52]. Policy gradient methods and actor-critic algorithms have been successfully applied to portfolio management, learning to dynamically allocate capital across assets based on market conditions. The advantage of DRL lies in its ability to incorporate transaction costs, market impact, and position limits directly into the learning objective, resulting in strategies that are practical for real-world deployment. Research has demonstrated that DRL-based trading systems can discover novel trading strategies that differ from traditional technical and fundamental approaches, potentially accessing unique sources of alpha [53].

Optimal trade execution represents a specific sequential decision-making problem where the objective is to execute large orders with minimal market impact and transaction costs. Traditional execution algorithms such as volume-weighted average price (VWAP) and time-weighted average price (TWAP) follow predetermined

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

schedules that do not adapt to market conditions [54]. DL-based execution algorithms learn optimal execution strategies by observing how order submissions affect prices and adjusting execution schedules dynamically based on market liquidity, volatility, and adverse selection risk. Research has shown that DRL agents can learn sophisticated execution strategies that outperform traditional algorithms by adapting to market microstructure features and avoiding predictable patterns that could be exploited by other market participants. The integration of limit order book data into DL execution models provides fine-grained information about market liquidity and enables more precise control of execution timing and order placement.

Time Lag	Price	Volume	Bid-Ask Spread	Order Imbalance	Volatility	Market Depth
Current	0.42	0.35	0.38	0.89	0.91	0.33
5 min	0.38	0.28	0.31	0.82	0.85	0.25
10 min	0.35	0.22	0.27	0.76	0.78	0.21
15 min	0.29	0.18	0.23	0.68	0.71	0.19
20 min	0.25	0.15	0.19	0.58	0.62	0.16
30 min	0.21	0.12	0.15	0.45	0.49	0.13
45 min	0.32	0.09	0.11	0.31	0.34	0.08
60 min	0.36	0.07	0.09	0.24	0.27	0.06

Low (0.0) High (1.0)

Figure 2. Visualization of attention weights in a transformer-based trading model during a significant market volatility event.

The heatmap displays attention scores across different time lags (rows) and market features (columns) when making trading decisions at 2:30 PM during a major market drawdown on March 15, 2024. Each cell shows the attention weight value ranging from 0 to 1, with color intensity from light yellow (low attention, near 0) to dark red (high attention, near 1). The x-axis represents different market features including current price, trade volume, bid-ask spread, order book imbalance ratio, realized volatility, and market depth at five price levels. The y-axis shows time lags from current moment to 60 minutes prior in varying intervals. The visualization reveals that the model focuses heavily on recent order book imbalance (attention weights 0.89-0.91) and realized volatility (attention weights 0.85-0.91) when making decisions during volatile periods, indicated by the darker red coloring in these columns for recent time lags. Attention to price history extends further back, showing elevated weights (0.32-0.36) at 45-60 minute lags during calm pre-volatility periods, suggesting the model recognizes historical price patterns as precursors to current volatility. Data collected from high-frequency trading session on NASDAQ exchange for large-cap technology stock with average daily volume exceeding 50 million shares. Source: Zhang et al. 2024 study on attention mechanisms in financial markets published in Quantitative Finance journal.

Market making strategies, where traders provide liquidity by simultaneously posting buy and sell limit orders, represent another important application of DL for sequential decision-making in financial markets. Market makers face the challenge of setting bid and ask prices that attract order flow while managing inventory risk and adverse selection from informed traders [55]. DL architectures have been applied to learn optimal market making policies that balance these competing objectives. LSTM networks can model the dynamics of order arrivals, price movements, and inventory accumulation, enabling market makers to adjust quotes dynamically based on current market conditions and inventory positions. Research has demonstrated that DRL-based market making agents can achieve superior performance compared to traditional market making algorithms by learning to adapt spreads and quote sizes based on volatility regimes, order flow toxicity, and inventory constraints [56].

The integration of alternative data sources into DL trading models has opened new opportunities for alpha generation by incorporating information not fully reflected in traditional price and volume data. News sentiment analysis using transformer-based NLP models can extract market-moving information from financial news articles, earnings call transcripts, and social media posts [57]. These sentiment signals can be integrated into trading models as additional features that inform buy and sell decisions. Research has shown that DL models combining traditional market microstructure features with alternative data can achieve improved prediction accuracy and trading performance. However, the challenge of identifying true signal in noisy alternative data requires sophisticated architectures capable of filtering out irrelevant information and focusing on genuinely predictive patterns.

6. Challenges and Future Directions

The deployment of DL architectures for sequential decision-making in financial systems faces several critical challenges that must be addressed to realize their full potential. Model interpretability remains a primary concern,

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

particularly in regulated financial applications where decisions affecting customers and markets require explanation and justification [58]. While attention mechanisms and other interpretability techniques provide some insight into model behavior, they often fall short of the level of transparency required by regulators and stakeholders. Future research must focus on developing architectures that maintain high predictive performance while offering inherent interpretability, such as neural additive models that decompose predictions into additive contributions from individual features. Furthermore, the development of post-hoc explanation methods specifically tailored to financial sequential decision-making contexts can help bridge the gap between model complexity and interpretability requirements.

The challenge of data quality and availability significantly impacts the effectiveness of DL models in financial applications. Financial data often suffers from missing values, measurement errors, and limited historical coverage, particularly for rare events such as financial crises [59]. The non-stationary nature of financial markets means that historical data may have limited relevance for current conditions, reducing the effectiveness of models trained on past observations. Future research should explore techniques for robust learning under data quality constraints, including methods for handling missing data, detecting and correcting measurement errors, and incorporating uncertainty quantification into model predictions. Transfer learning approaches that leverage knowledge from related tasks or markets may help address data scarcity issues, enabling models to generalize from limited historical observations.

Adversarial robustness represents a unique challenge in financial applications, where malicious actors may attempt to manipulate data or exploit model behavior for financial gain. Unlike image classification tasks where adversarial examples are primarily a theoretical concern, financial markets involve strategic interactions where participants actively seek to exploit predictable patterns [60]. DL models deployed in trading or fraud detection may become targets of adversarial attacks designed to trigger false signals or evade detection. Future research must develop architectures and training procedures that enhance model robustness against adversarial manipulation, including adversarial training techniques adapted to financial contexts and the development of detection mechanisms for identifying adversarial attacks in real-time.

The computational efficiency of DL models presents practical challenges for real-time financial applications where decisions must be made within strict latency constraints. While transformer architectures offer advantages in modeling long-range dependencies, their computational complexity grows quadratically with sequence length, limiting their applicability to high-frequency trading scenarios [61]. Future research should focus on developing efficient architectures that maintain the modeling capabilities of transformers while reducing computational requirements, such as linear attention mechanisms and sparse attention patterns optimized for financial time series. Furthermore, the development of model compression techniques including quantization, pruning, and knowledge distillation can enable the deployment of complex DL architectures on resource-constrained hardware.

The integration of domain knowledge into DL architectures represents an important direction for improving model performance and interpretability in financial applications. While end-to-end learning from raw data is appealing, incorporating financial theory and domain expertise into model design can improve generalization and provide economic interpretability [62]. Physics-informed neural networks that incorporate constraints based on financial principles such as no-arbitrage conditions or market equilibrium represent a promising approach. Future research should explore hybrid architectures that combine data-driven learning with structured components encoding domain knowledge, enabling models to respect fundamental financial principles while learning complex patterns from data.

Federated learning and privacy-preserving techniques offer promising solutions for addressing data privacy concerns in financial applications while enabling collaboration across institutions [63]. Financial institutions often possess complementary datasets that could benefit model training but cannot be shared due to privacy regulations and competitive concerns. Federated learning enables multiple institutions to collaboratively train DL models without sharing raw data, potentially improving model performance through access to larger and more diverse datasets. Future research should explore federated learning protocols specifically designed for financial sequential decision-making tasks, addressing challenges such as non-identically distributed data across institutions and the need for secure aggregation of model updates.

Quantum-inspired and neuromorphic computing architectures represent emerging technologies that may enhance the capabilities of DL for financial sequential decision-making. Quantum machine learning algorithms offer potential advantages in optimization and sampling tasks relevant to portfolio management and risk assessment [64]. While practical quantum computers remain limited in scale, quantum-inspired classical algorithms that leverage insights from quantum computing may provide near-term benefits. Neuromorphic hardware implementing spiking neural networks offers advantages in energy efficiency and temporal processing that could benefit high-frequency financial applications. Future research should explore how these emerging computing paradigms can be leveraged to enhance DL capabilities in financial contexts.

The development of standardized benchmarks and evaluation protocols for DL in financial sequential decision-making would facilitate progress by enabling fair comparisons across methods and promoting reproducibility. Current research often uses proprietary datasets and non-standard evaluation metrics, making it difficult to assess relative performance of different approaches [65]. Future efforts should focus on creating publicly available benchmark datasets that capture the diversity and complexity of financial sequential decision-making tasks while respecting privacy constraints. Furthermore, the development of evaluation protocols that go beyond simple accuracy metrics to assess robustness, fairness, and interpretability would provide more comprehensive assessment of model quality.

7. Conclusion

This review has examined the application of DL architectures to sequential decision-making in financial systems, with particular focus on fraud detection, credit risk assessment, algorithmic trading, and risk

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

management. The analysis demonstrates that DL models, particularly RNNs, LSTM networks, GRUs, and transformer-based architectures, offer substantial advantages over traditional methods by automatically learning hierarchical representations from sequential financial data and capturing complex temporal dependencies. These architectures have achieved significant success across diverse financial applications, enabling real-time fraud detection systems that adapt to evolving patterns, credit risk models that provide early warning signals of deterioration, and algorithmic trading strategies that optimize long-term objectives.

The integration of attention mechanisms has enhanced model interpretability while improving performance, allowing financial institutions to understand which historical events and features drive model predictions. The combination of DL with RL has opened new avenues for developing adaptive trading systems and dynamic risk management strategies that learn optimal policies through interaction with market environments. Furthermore, the incorporation of graph-based architectures has extended capabilities to capture network effects and relationships among entities in financial systems, enabling detection of sophisticated fraud schemes and assessment of systemic risks.

Despite these advances, significant challenges remain in deploying DL architectures for financial sequential decision-making. The non-stationary nature of financial data, low signal-to-noise ratios, and adversarial market dynamics require ongoing research into robust and adaptive architectures. Model interpretability and regulatory compliance continue to pose obstacles, necessitating development of explainable architectures and post-hoc explanation methods. Computational efficiency constraints in high-frequency applications demand continued innovation in efficient architectures and deployment optimizations.

Future research directions include the development of inherently interpretable architectures that maintain high performance, robust learning methods that handle data quality issues and adversarial attacks, and federated learning approaches that enable privacy-preserving collaboration across institutions. The integration of domain knowledge into DL architectures through physics-informed networks and hybrid models offers promise for improving generalization and economic interpretability. Emerging technologies such as quantum-inspired algorithms and neuromorphic computing may provide new capabilities for financial sequential decision-making tasks. The establishment of standardized benchmarks and comprehensive evaluation protocols will facilitate progress by enabling fair comparisons and promoting reproducibility.

The continued evolution of DL architectures for financial sequential decision-making holds tremendous potential for transforming how financial institutions detect fraud, assess risk, execute trades, and manage portfolios. As architectures become more sophisticated, interpretable, and robust, their adoption across the financial services industry is likely to accelerate, driving improvements in efficiency, risk management, and customer protection while enabling new capabilities previously unattainable with traditional methods.

References

Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. *ACM Computing Surveys (CSUR)*, *55*(3), 1–38. https://doi.org/10.1145/3501295

Rane, N. L., Paramesha, M., Choudhary, S. P., & Rane, J. (2024). Machine learning and deep learning for big data analytics: A review of methods and applications. *Partners Universal International Innovation Journal*, *2*(3), 172–197.

Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. *Quantitative Finance*, *19*(8), 1271–1291. https://doi.org/10.1080/14697688.2019.1571683

Chauhan, A. (2024). *Interface engineering strategies to induce multifunctionality in conventional carbon fiber/epoxy composites* [Doctoral dissertation, Indian Institute of Technology]. Shodhganga.

Zhang, Z., Zohren, S., & Roberts, S. (2020). Deep learning for portfolio optimization. *Journal of Financial Data Science*, *2*(4), 8–20. https://doi.org/10.3905/jfds.2020.1.028

Hambly, B., Xu, R., & Yang, H. (2023). Recent advances in reinforcement learning in finance. *Mathematical Finance*, *33*(3), 437–503. https://doi.org/10.1111/mafi.12325

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, *33*(5), 2223–2273. https://doi.org/10.1093/rfs/hhaa009

Rudin, C., Chen, C., Chen, Z., Huang, H., Semenova, L., & Zhong, C. (2022). Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistics Surveys*, *16*, 1–85. https://doi.org/10.1214/21-SS133

Özbayoglu, A. M., Gudelek, M. U., & Sezer, O. B. (2020). Deep learning for financial applications: A survey. *Applied Soft Computing*, *93*, 106384. https://doi.org/10.1016/j.asoc.2020.106384

Henrique, B. M., Sobreiro, V. A., & Kimura, H. (2019). Literature review: Machine learning techniques applied to financial market prediction. *Expert Systems with Applications*, *124*, 226–251. https://doi.org/10.1016/j.eswa.2019.01.012

Siami-Namini, S., Tavakoli, N., & Namin, A. S. (2019). The performance of LSTM and BiLSTM in forecasting time series. In *2019 IEEE International Conference on Big Data (Big Data)* (pp. 3285–3292). IEEE. https://doi.org/10.1109/BigData47090.2019.9005547

Taha, A. A., & Malebary, S. J. (2020). An intelligent approach to credit card fraud detection using an optimized light gradient boosting machine. *IEEE Access*, *8*, 25579–25587. https://doi.org/10.1109/ACCESS.2020.2971354

Pourhabibi, T., Ong, K. L., Kam, B. H., & Boo, Y. L. (2020). Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *Decision Support Systems*, *133*, 113303. https://doi.org/10.1016/j.dss.2020.113303

Malik, P., Chourasia, A., Pandit, R., Bawane, S., & Surana, J. (2024). Credit risk assessment and fraud detection in financial transactions using machine learning. *Journal of Electrical Systems*, *20*(3s), 2061–2069.

Moscatelli, M., Narizzano, S., Parlapiano, F., & Viggiano, G. (2020). Corporate default forecasting with machine learning. *Expert Systems with Applications*, *161*, 113567. https://doi.org/10.1016/j.eswa.2020.113567

Wang, Y., & Ni, X. S. (2020, April). Risk prediction of peer-to-peer lending market by a LSTM model with macroeconomic factor. In *Proceedings of the 2020 ACM Southeast Conference* (pp. 181–187). ACM. https://doi.org/10.1145/3374135.3385280

Faheem, M. A. (2021). AI-driven risk assessment models: Revolutionizing credit scoring and default prediction. *Iconic Research and Engineering Journals*, *5*(3), 177–186.

Ryll, L., & Seidens, S. (2019). Evaluating the performance of machine learning algorithms in financial market forecasting: A comprehensive survey. *ACM Computing Surveys*, *52*(5), 1–36. https://doi.org/10.1145/3341224

Sezer, O. B., Gudelek, M. U., & Özbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review 2005–2019. *Applied Soft Computing*, *90*, 106181. https://doi.org/10.1016/j.asoc.2020.106181

Shavandi, A., & Khedmati, M. (2022). A multi-agent deep reinforcement learning framework for algorithmic trading in financial markets. *Expert Systems with Applications*, *208*, 118124. https://doi.org/10.1016/j.eswa.2022.118124

Livieris, I. E., Pintelas, E., & Pintelas, P. (2020). A CNN-LSTM model for gold price time-series forecasting. *Neural Computing and Applications*, *32*(23), 17351–17360. https://doi.org/10.1007/s00521-020-04867-x

Liu, G., Shi, K., & Yuan, M. (2025). Forecasting the high-frequency covariance matrix using the LSTM-MF model. *Journal of Forecasting*. https://doi.org/10.1002/for.3090

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

Wang, Y., & Guo, Y. (2020). Forecasting method of stock market volatility in time series data based on mixed model of ARIMA and XGBoost. *China Communications, 17*(3), 205–221. https://doi.org/10.23919/JCC.2020.03.015

Liu, T., Wang, Y., Sun, J., Tian, Y., Huang, Y., Xue, T., ... & Liu, Y. (2024). The role of transformer models in advancing blockchain technology: A systematic survey. *arXiv preprint* arXiv:2409.02139. https://arxiv.org/abs/2409.02139

Şahin, E., Arslan, N. N., & Özdemir, D. (2025). Unlocking the black box: An in-depth review on interpretability, explainability, and reliability in deep learning. *Neural Computing and Applications, 37*(2), 859–965. https://doi.org/10.1007/s00521-025-10187-5

Khouas, A. R., Bouadjenek, M. R., Hacid, H., & Aryal, S. (2024). Training machine learning models at the edge: A survey. *arXiv preprint* arXiv:2403.02619. https://arxiv.org/abs/2403.02619

Meng, T. L., & Khushi, M. (2019). Reinforcement learning in financial markets. *Data, 4*(3), 110. https://doi.org/10.3390/data4030110

Théate, T., & Ernst, D. (2021). An application of deep reinforcement learning to algorithmic trading. *Expert Systems with Applications, 173*, 114632. https://doi.org/10.1016/j.eswa.2021.114632

Chinnaraju, A. (2025). Explainable AI (XAI) for trustworthy and transparent decision-making: A theoretical framework for AI interpretability. *World Journal of Advanced Engineering Technology and Sciences, 14*(3), 170–207.

Gramegna, A., & Giudici, P. (2021). SHAP and LIME: An evaluation of discriminative power in credit risk. *Frontiers in Artificial Intelligence, 4*, 752558. https://doi.org/10.3389/frai.2021.752558

Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications, 193*, 116429. https://doi.org/10.1016/j.eswa.2021.116429

Cheng, D., Xiang, S., Shang, C., Zhang, Y., Yang, F., & Zhang, L. (2020). Spatio-temporal attention-based neural network for credit card fraud detection. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 34, pp. 362–369). https://doi.org/10.1609/aaai.v34i01.5387

Lalapura, V. S., Amudha, J., & Satheesh, H. S. (2021). Recurrent neural networks for edge intelligence: A survey. *ACM Computing Surveys (CSUR), 54*(4), 1–38. https://doi.org/10.1145/3447750

Ji, S., Pan, S., Long, G., Li, X., Jiang, J., & Huang, Z. (2019, July). Learning private neural language modeling with attentive aggregation. In *2019 International Joint Conference on Neural Networks (IJCNN)* (pp. 1–8). IEEE. https://doi.org/10.1109/IJCNN.2019.8852382

Liu, Y., Ao, X., Qin, Z., Chi, J., Feng, J., Yang, H., & He, Q. (2021). Pick and choose: A GNN-based imbalanced learning approach for fraud detection. In *Proceedings of the Web Conference 2021* (pp. 3168–3177). ACM. https://doi.org/10.1145/3442381.3450091

Olushola, A., & Mart, J. (2024). Fraud detection using machine learning. *ScienceOpen Preprints*. https://doi.org/10.14293/S2199-1006.1.SOR-.PP5FDP2.v1

Sadiq, S., Umer, M., Ullah, S., Mirjalili, S., Rupapara, V., & Nappi, M. (2021). Discrepancy detection between actual user reviews and numeric ratings of Google App Store using deep learning. *Expert Systems with Applications, 181*, 115111. https://doi.org/10.1016/j.eswa.2021.115111

Carcillo, F., Le Borgne, Y. A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences, 557*, 317–331. https://doi.org/10.1016/j.ins.2020.12.020

Gunnarsson, B. R., Vanden Broucke, S., Baesens, B., Oskarsdóttir, M., & Lemahieu, W. (2021). Deep learning for credit scoring: Do or don't? *European Journal of Operational Research, 295*(1), 292–305. https://doi.org/10.1016/j.ejor.2021.03.027

Tripathy, J. K., Sethuraman, S. C., Cruz, M. V., Namburu, A., & Vijayakumar, V. (2021). Comprehensive analysis of embeddings and pre-training in NLP. *Computer Science Review, 42*, 100433. https://doi.org/10.1016/j.cosrev.2021.100433

Óskarsdóttir, M., Bravo, C., Sarraute, C., Vanthienen, J., & Baesens, B. (2019). The value of big data for credit scoring: Enhancing financial inclusion using mobile phone data and social network analytics. *Applied Soft Computing, 74*, 26–39. https://doi.org/10.1016/j.asoc.2018.10.004

Kowsar, M. M., Mohiuddin, M., & Islam, S. (2023). Mathematics for finance: A review of quantitative methods in loan portfolio optimization. *International Journal of Scientific Interdisciplinary Research, 4*(3), 1–29.

Wang, P., Li, Y., & Reddy, C. K. (2019). Machine learning for survival analysis: A survey. *ACM Computing Surveys, 51*(6), 1–36. https://doi.org/10.1145/3214306

Bracke, P., Datta, A., Jung, C., & Sen, S. (2020). Machine learning explainability in finance: An application to default risk analysis. *Journal of Banking & Finance, 130*, 106199. https://doi.org/10.1016/j.jbankfin.2021.106199

Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics, 57*(1), 203–216. https://doi.org/10.1007/s10614-020-10042-0

Asif, S. S. M. (2024). *Create an AI-based model for dynamic risk management in equity portfolios that accounts for extreme market events (e.g., pandemics, financial crises)*. [Unpublished manuscript].

Kakade, K., Jain, I., & Mishra, A. K. (2022). Value-at-Risk forecasting: A hybrid ensemble learning GARCH-LSTM based approach. *Resources Policy, 78*, 102903. https://doi.org/10.1016/j.resourpol.2022.102903

Olanrewaju, A. G. (2025). Artificial intelligence in financial markets: Optimizing risk management, portfolio allocation, and algorithmic trading. *International Journal of Research Publication and Reviews, 6*, 8855–8870.

Zhang, K., Zhong, G., Dong, J., Wang, S., & Wang, Y. (2019). Stock market prediction based on generative adversarial network. *Procedia Computer Science, 147*, 400–406. https://doi.org/10.1016/j.procs.2019.01.263

Sun, H., Bian, Y., Han, L., Zhu, P., Cheng, D., & Liang, Y. (2024, October). Dynamic graph-based deep reinforcement learning with long- and short-term relation modeling for portfolio optimization. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management* (pp. 4898–4905). ACM. https://doi.org/10.1145/3583780.3615407

Zhang, Z., Zohren, S., & Roberts, S. (2019). Deep reinforcement learning for trading. *arXiv preprint* arXiv:1911.10107. https://arxiv.org/abs/1911.10107

Nuipian, W., & Meesad, P. (2025). *Dynamic portfolio management with deep reinforcement learning* [Doctoral dissertation, King Mongkut's University of Technology North Bangkok]. KMUTNB Institutional Repository.

Liu, Y., Liu, Q., Zhao, H., Pan, Z., & Liu, C. (2020). Adaptive quantitative trading: An imitative deep reinforcement learning approach. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 34, pp. 2128–2135). https://doi.org/10.1609/aaai.v34i02.5520

Shannon, B. (2023). *Maximum trading gains with anchored VWAP: The perfect combination of price, time & volume.* Alphatrends Publishing LLC.

Vallabhaneni, M. (2022). Traders' selection of order types: A review. *Message from the Editor-in-Chief.*

Sadighian, J. (2019). Deep reinforcement learning in cryptocurrency market making. *arXiv preprint* arXiv:1911.08647. https://arxiv.org/abs/1911.08647

Hiew, J. Z. G., Huang, X., Mou, H., Li, D., Wu, Q., & Xu, Y. (2019). BERT-based financial sentiment index and LSTM-based stock return predictability. *arXiv preprint* arXiv:1906.09024. https://arxiv.org/abs/1906.09024

Mohanarajesh, K. (2024). *Investigate methods for visualizing the decision-making processes of a complex AI system, making them more understandable and trustworthy in financial data analysis.* [Unpublished manuscript].

Oloyede, J., & Owen, J. (2025). Enhancing data quality and integrity with AI: A deep learning perspective. *Fajinmi John Publications.*

Chernikova, A., Oprea, A., Nita-Rotaru, C., & Kim, B. (2019). Are self-driving cars secure? Evasion attacks against deep neural networks for steering angle prediction. In *2019 IEEE Security and Privacy Workshops (SPW)* (pp. 132–137). IEEE. https://doi.org/10.1109/SPW.2019.00031

Ansar, W., Goswami, S., & Chakrabarti, A. (2024). A survey on transformers in NLP with focus on efficiency. *arXiv preprint* arXiv:2406.16893. https://arxiv.org/abs/2406.16893

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics, 378*, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology, 10*(2), 1–19. https://doi.org/10.1145/3298981

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA

Bhasin, N. K., Kadyan, S., Santosh, K., H. P., R., Changala, R., & Bala, B. K. (2024, March). Enhancing quantum machine learning algorithms for optimized financial portfolio management. In *2024 Third International Conference on Intelligent Techniques in Control, Optimization and Signal Processing (INCOS)* (pp. 1–7). IEEE. https://doi.org/10.1109/INCOS60539.2024.10429048

Baron, G., & Stańczyk, U. (2021). Standard vs. non-standard cross-validation: Evaluation of performance in a space with structured distribution of datapoints. *Procedia Computer Science*, *192*, 1245–1254. https://doi.org/10.1016/j.procs.2021.08.128

© 2025 by the authors; licensee Eastern Centre of Science and Education, USA