Deep Learning in Financial Time Series Analysis: A Comprehensive Review of Methods, Challenges, and Future Directions

Husam A.S Khalifa
Bani Waleed University- Libya
husam.khalifa@bwu.edu.ly

Submitted	Accepted	Published
25/02/2025	29/03/2025	31/03/2025

Abstract:
This comprehensive review examines the application of deep learning techniques in financial time series analysis, with a particular focus on asset price prediction and market volatility forecasting. Through a systematic analysis of over 80 studies and technical sources published between 2014 and 2025—including peer-reviewed journal articles, conference proceedings, preprints, and industry reports—we identify key architectural advances, methodological innovations, and emerging research directions. Our findings reveal that transformer-based models and hybrid architectures consistently outperform traditional econometric approaches, achieving up to 25% improvement in prediction accuracy. However, significant challenges remain in model interpretability, generalization across market regimes, and the integration of multimodal data sources. We identify critical research gaps in explainable AI (XAI) for financial decision-making, transfer learning for emerging markets, and online adaptive learning systems. This review provides a structured framework for understanding the current state of the field and highlights key opportunities for advancing deep learning-based financial forecasting.

Keywords: Deep Learning, Financial Time Series, Asset Price Prediction, Volatility Forecasting, Transformer Models, Hybrid Architectures, Explainable AI, Multimodal Learning.

Introduction

Financial time series forecasting represents one of the most challenging applications of machine learning, characterized by high noise levels, non-stationarity, and complex interdependencies. The ability to accurately predict asset prices and market volatility has profound implications for investment strategies, risk management, and economic policy formulation. Traditional econometric models, while theoretically grounded, often struggle to capture the nonlinear dynamics and regime changes inherent in financial markets. (Kanungo, 2025)

The emergence of deep learning has fundamentally transformed the landscape of financial forecasting. Unlike traditional statistical models that rely on linear assumptions and manual feature engineering, deep learning architectures can automatically discover complex patterns and nonlinear relationships within high-dimensional data. This capability is particularly valuable in financial markets where relationships between variables are often time-varying and subject to structural breaks. (Li, Schulwol, & Miikkulainen, 2025)

Recent advances in deep learning have introduced sophisticated architectures specifically designed for sequential data analysis. Long Short-Term Memory (LSTM) networks address the vanishing gradient problem in traditional recurrent neural networks, enabling the capture of long-term dependencies crucial for financial forecasting. Transformer models, originally developed for natural language processing, have demonstrated remarkable success in modeling temporal relationships through self-attention mechanisms. Hybrid architectures combining convolutional neural networks (CNNs) with recurrent layers have shown promise in capturing both local patterns and global temporal dependencies. (Li, Sun, Wu, & Tao, 2024)

The integration of multimodal data sources represents another significant advancement. Modern deep learning frameworks can simultaneously process traditional financial data alongside alternative information sources including news sentiment, social media data, and macroeconomic indicators. This capability addresses a fundamental limitation of traditional models that typically focus on single data modalities.(Kong, Chen, Liu, Ning, Zhang, Muhammad Marier, et al., 2025)

Despite these advances, several critical challenges persist. Model interpretability remains a significant concern, particularly in regulated financial environments where decision transparency is paramount. The stability of deep learning models across different market regimes and their ability to generalize to unseen market conditions continue to pose significant challenges. Additionally, the computational requirements and data dependencies of deep learning models raise practical concerns for real-world deployment. (Q. Chen, 2025; J. Li et al., 2024; Varadharajan, Smith, Kalla, Samaah, et al., 2024)

This review aims to provide a comprehensive analysis of deep learning applications in financial time series forecasting, examining architectural innovations, methodological advances, and empirical findings. We systematically evaluate the current state of the field, identify key research gaps, and propose future directions for advancing the intersection of deep learning and financial forecasting.

2. Traditional vs. Deep Learning Approaches

2.1 Traditional Econometric Models

Traditional financial time series analysis has been dominated by econometric models designed to capture specific stylized facts of financial data. **Autoregressive Integrated Moving Average (ARIMA)** models form the foundation of classical time series analysis, providing a framework for modeling linear dependencies and trend components. The **Box-Jenkins methodology** offers a systematic approach to model identification, estimation, and diagnostic checking, making ARIMA models highly interpretable and theoretically grounded.(Paras Varshney, n.d.)

Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models address volatility clustering, a fundamental characteristic of financial time series where periods of high volatility tend to cluster together. GARCH variants, including EGARCH and GJR-GARCH, incorporate asymmetric effects where negative shocks have different impacts on volatility compared to positive shocks of similar magnitude. These models have been extensively validated and remain widely used in risk management applications.(Mishra et al., 2024) (Kong, Chen, Liu, Ning, Zhang, Marier, et al., 2025)

Vector Autoregression (VAR) models extend univariate approaches to capture interdependencies among multiple financial variables. VAR models are particularly valuable for understanding spillover effects and Granger causality relationships among financial markets. However, their effectiveness diminishes rapidly as the number of variables increases due to the curse of dimensionality.(Moghar & Hamiche, 2020)

2.2 Deep Learning Architectures

Deep learning approaches fundamentally differ from traditional methods in their ability to automatically discover complex, nonlinear patterns without explicit specification of functional forms. **Recurrent Neural Networks (RNNs)** provide the basic framework for sequential data modeling, though they suffer from vanishing gradient problems that limit their ability to capture long-term dependencies.(Ozturk, 2020)

Long Short-Term Memory (LSTM) networks address these limitations through gating mechanisms that control information flow. The forget gate determines which information to discard, the input gate controls new information incorporation, and the output gate regulates the hidden state output. Recent empirical studies demonstrate that LSTM models achieve superior performance compared to traditional approaches, with improvements ranging from 15-30% in out-of-sample forecasting accuracy.(Z. Xu et al., 2024)

Transformer architectures represent a paradigm shift in sequential modeling through self-attention mechanisms that can model long-range dependencies without sequential processing constraints. The multi-head attention mechanism allows models to focus on different aspects of the input sequence simultaneously, proving particularly effective for capturing complex temporal patterns in financial data.(Zeng et al., 2023) (S. Li et al., 2025)

Convolutional Neural Networks (CNNs) have been adapted for financial time series through techniques that transform temporal data into image-like representations. This approach enables CNNs to identify local patterns and features that may be missed by purely sequential models.(Y. Li & Pan, 2022)

2.3 Comparative Performance Analysis

Empirical comparisons consistently demonstrate the superiority of deep learning approaches across multiple evaluation metrics.

Table 1 summarizes key performance improvements observed in recent studies.

Model Type	RMSE Improvement	MAE Improvement	Directional Accuracy	Sharpe Ratio Enhancement
LSTM vs ARIMA	18-25%	15-22%	8-12%	15-20%
Transformer vs GARCH	20-30%	18-28%	10-15%	18-25%

Model Type	RMSE Improvement	MAE Improvement	Directional Accuracy	Sharpe Ratio Enhancement
Hybrid CNN-LSTM vs VAR	22-35%	20-32%	12-18%	20-30%

These improvements are particularly pronounced during periods of market stress and volatility clustering, where traditional models' linear assumptions prove most limiting. Deep learning models demonstrate superior ability to adapt to regime changes and capture non-stationary behavior that characterizes financial markets during crisis periods.(Di-Giorgi et al., 2025)

2.4 Complexity and Interpretability Trade-offs

While deep learning models achieve superior predictive performance, they introduce significant challenges in terms of model complexity and interpretability. Traditional econometric models offer clear parameter interpretations and established hypothesis testing frameworks. In contrast, deep learning models operate as "black boxes" where decision processes are opaque and difficult to explain.(Tu, 2025)

This interpretability gap has significant implications for regulatory compliance and risk management applications where model transparency is crucial. Recent advances in explainable AI (XAI) attempt to address these concerns through techniques such as attention visualization, feature importance analysis, and counterfactual explanations. However, these approaches remain in early stages of development for financial applications.(Seseri, 2023a)

3. Architectural Advances in Deep Learning for Finance

3.1 Recurrent Neural Network Variants

The evolution of recurrent architectures has been driven by the need to address fundamental limitations in capturing long-term dependencies within financial time series. **Vanilla RNNs** suffer from vanishing gradients that prevent effective learning of relationships spanning extended time horizons, a critical limitation given the long memory effects observed in financial markets.(Feng et al., 2023)

LSTM networks represent a significant advancement through their sophisticated gating mechanisms. The architecture incorporates three specialized gates: the **forget gate** ($f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f)$) determines which information from the cell state should be discarded; the **input gate** ($i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$) controls which new information is stored; and the **output gate** ($o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o)$) regulates the output based on the cell state. This architecture enables LSTM networks to maintain information across hundreds of time steps, crucial for capturing seasonal patterns and long-term trends in financial data.(Shubham Ghelani, n.d.)

Gated Recurrent Units (GRUs) offer a simplified alternative that combines the forget and input gates into a single update gate, reducing computational complexity while maintaining comparable performance. Recent comparative studies in financial forecasting suggest that GRUs achieve 85-95% of LSTM performance with 30-40% fewer parameters, making them attractive for resource-constrained applications.(Z. Li & Huang, 2025; Varadharajan, Smith, Kalla, Kumar, et al., 2024)

Bidirectional RNNs process sequences in both forward and backward directions, enabling the model to incorporate future context when predicting intermediate time points. This capability proves particularly valuable for financial backtesting scenarios where future information is available for model training and validation.(Linné, 2024)

3.2 Transformer Architectures for Financial Time Series

The adaptation of transformer models to financial time series has yielded remarkable performance improvements across diverse forecasting tasks. **Multi-head self-attention** mechanisms enable models to identify relationships between distant time points without the sequential processing constraints that limit RNN-based approaches.(Zeng et al., 2023)

The **scaled dot-product attention** mechanism computes attention weights as: $\text{Attention}(Q,K,V) = \text{softmax}(QK^T/\sqrt{d_k})V$

where Q, K, and V represent query, key, and value matrices respectively. This formulation allows the model to assign different importance weights to various time steps when making predictions, providing insights into which historical periods are most relevant for current forecasting tasks.(Zeng et al., 2023)

Positional encoding addresses the lack of inherent sequential ordering in transformer architectures by injecting position information through sinusoidal functions. For financial applications, researchers have explored learnable positional embeddings that can adapt to market-specific temporal patterns and trading calendars.(Seseri, 2023b)

Recent studies demonstrate that transformer models achieve superior performance in multi-horizon forecasting scenarios, with **PatchTST** architectures showing particular promise for financial applications. These models segment time series into patches and apply attention mechanisms at the patch level, reducing computational complexity while maintaining modeling capacity.(Ferrari, 2022)

3.3 Convolutional Neural Networks for Financial Data

The application of CNNs to financial time series requires innovative approaches to transform temporal data into spatial representations amenable to convolutional processing. **Gramian Angular Fields (GAF)** transform time series into image-like matrices that preserve temporal dependencies while enabling CNN processing.(Tam, 2021)

1D Convolutional layers can directly process sequential financial data, with kernels acting as learnable filters that identify local patterns such as price reversals, momentum shifts, and volatility spikes. The hierarchical feature extraction capability of CNNs proves particularly effective for identifying multi-scale patterns that characterize financial markets across different time horizons.(Liladhar Rane et al., 2024)

Dilated convolutions extend the receptive field without increasing computational complexity, enabling the capture of long-range dependencies that are crucial for financial forecasting. This approach allows CNNs to model seasonal patterns and long-term trends while maintaining computational efficiency.(Bhambu et al., 2025)

3.4 Hybrid and Ensemble Architectures

The combination of different deep learning architectures has emerged as a powerful approach for leveraging the complementary strengths of various model types. **CNN-LSTM hybrids** utilize convolutional layers for local feature extraction followed by LSTM layers for temporal modeling. This architecture proves particularly effective for capturing both short-term price patterns and long-term trend dynamics. (W. Chen et al., 2023)

Ensemble methods combine predictions from multiple models to improve robustness and reduce overfitting risks. **Stacking ensembles** train a meta-learner to optimally combine base model predictions, while **voting ensembles** use simple averaging or majority voting schemes. Recent studies demonstrate that ensemble approaches can reduce prediction variance by 20-35% compared to individual models while maintaining comparable bias levels.(Awalee Consulting, 2023) (Y. Li & Pan, 2022)

Attention-based ensemble methods dynamically weight different models based on current market conditions, enabling adaptive model selection that responds to changing market regimes. These approaches show particular promise for handling the non-stationary nature of financial time series.(Corporate Finance Institute, n.d.)

3.5 Training Methodologies and Optimization

The training of deep learning models for financial applications requires specialized techniques to address the unique characteristics of financial data. **Walk-forward validation** ensures temporal consistency by using only past information for training and validation, preventing look-ahead bias that could artificially inflate performance metrics.(Mohammed & Kora, 2023)

Gradient clipping addresses the exploding gradient problem common in financial time series due to extreme price movements and volatility spikes. Learning rate scheduling approaches, including **cosine annealing** and **warm restarts**, help models navigate the complex loss landscapes characteristic of financial forecasting problems.(Milvus, n.d.)

Regularization techniques including dropout, batch normalization, and weight decay prove crucial for preventing overfitting in financial applications where noise levels are typically high. Early stopping based on validation set performance helps identify optimal model complexity and prevents overtraining.(Černevičienė & Kabašinskas, 2024)

4. Feature Engineering and Data Sources

4.1 Traditional Technical Indicators

The foundation of financial feature engineering rests on well-established technical indicators that quantify market dynamics and price behavior patterns. **Moving averages** capture trend direction and momentum, with **Simple Moving Averages (SMA)** and **Exponential Moving Averages (EMA)** providing different weighting schemes for historical prices. The crossover signals between short-term and long-term moving averages serve as fundamental trend reversal indicators.(Karadaş et al., 2025)

Relative Strength Index (RSI) measures momentum by comparing recent gains to recent losses over a specified period, typically 14 days. RSI values above 70 indicate overbought conditions while values below 30 suggest oversold conditions, providing insights into potential reversal points. (H. Wu, 2024)

Bollinger Bands combine moving averages with volatility measures by plotting bands at standard deviations above and below a central moving average. Price movements touching the upper or lower bands often signal potential reversal points or continuation patterns, depending on market context.(Wilson, 2025)

Volume-based indicators such as **On-Balance Volume (OBV)** and **Volume-Weighted Average Price (VWAP)** incorporate trading volume information to validate price movements and identify potential divergences. These indicators prove particularly valuable for institutional trading applications where volume patterns provide insights into market participation levels.(J. Liu, 2025)

4.2 Advanced Financial Features

Modern feature engineering extends beyond traditional technical analysis to incorporate sophisticated mathematical transformations and market microstructure variables. **Wavelet transforms** decompose price series into different frequency components, enabling the separate analysis of short-term noise and long-term trends. This decomposition proves particularly valuable for multi-horizon forecasting applications.(Abbasimehr & Paki, 2022; Gao & Kuruoğlu, 2024; J. Li et al., 2023; Shali et al., 2021; Wen & Li, 2023; Y. Xiao et al., 2021; X. Zhang et al., 2019; Zhou et al., 2020)

Fractal and complexity measures capture the self-similar properties of financial time series. The **Hurst exponent** quantifies long-range dependence and mean reversion tendencies, while **multifractal detrended fluctuation analysis (MFDFA)** reveals scaling properties across different time scales.(Mohsin & Nasim, 2025)

Option-implied features extract forward-looking information from derivatives markets. **Implied volatility** surfaces provide market expectations of future volatility across different strikes and maturities, while **volatility skew** measures capture asymmetric risk perceptions. **Put-call ratios** and **option flow indicators** offer insights into market sentiment and institutional positioning. (Che et al., 2024; L. Li et al., 2023; S. Li & Tang, 2024)

Market microstructure variables including bid-ask spreads, order book depth, and trade size distributions provide information about liquidity conditions and market efficiency. These features prove particularly valuable for high-frequency trading applications and short-term price prediction.(Kumar et al., 2023)

4.3 Alternative Data Sources

The integration of alternative data sources has emerged as a key differentiator in modern financial forecasting applications. **Social media sentiment** extracted from platforms such as Twitter, Reddit, and financial forums provides real-time insights into market psychology and retail investor behavior. **Natural Language Processing (NLP)** techniques including **sentiment analysis**, **named entity recognition**, and **topic modeling** transform unstructured text into quantitative features. (AbdElnapi et al., 2024; Abu Jamie et al., 2024)

News analytics processes financial news articles, earnings transcripts, and analyst reports to extract relevant information for price prediction. **Event detection algorithms** identify market-moving news events and quantify their potential impact on asset prices. **Earning surprise indicators** and **analyst revision patterns** provide insights into fundamental changes in company prospects.(J. Wang et al., 2023)

Satellite imagery and geospatial data offer unique insights for commodity and real estate markets. **Agricultural monitoring** through satellite imagery enables crop yield estimation, while **oil storage tracking** provides insights into supply-demand dynamics. **Economic activity indicators** derived from satellite data, including nighttime illumination and traffic patterns, offer real-time economic monitoring capabilities.(Bailey, n.d.-b; Bruederle & Hodler, 2018; Mateo-Sanchis et al., 2019; Yao, 2019)

Patent filings and research publications provide leading indicators of technological innovation and competitive positioning for technology companies. **Supply chain analytics** using shipping data and logistics information offers insights into operational efficiency and demand patterns. (Confraria et al., 2024; Érdi et al., 2013)

4.4 Multimodal Data Integration

The effective integration of diverse data sources requires sophisticated approaches that can handle different data types, frequencies, and quality levels. **Feature alignment techniques** address temporal mismatches between data sources, ensuring that features represent contemporaneous information.(S. Li & Tang, 2024; Mondal et al., 2025)

Dimensionality reduction methods including **Principal Component Analysis (PCA)**, **Independent Component Analysis (ICA)**, and **autoencoders** help manage the curse of dimensionality when integrating numerous alternative data sources. These techniques identify the most informative combinations of features while reducing computational complexity. (Che et al., 2024)

Attention mechanisms enable models to dynamically weight different data sources based on their relevance for current market conditions. This adaptive weighting proves particularly valuable when data quality or relevance varies over time or across different market regimes. (Alzahrani, 2024; X. Li et al., 2017)

Cross-modal learning approaches enable models to leverage relationships between different data modalities. For example, **image-text models** can connect satellite imagery with textual news reports about agricultural conditions, while **audio-visual models** can integrate earnings call transcripts with speaker sentiment analysis.(Tavakoli et al., 2024)

4.5 Preprocessing and Feature Selection

Effective preprocessing is crucial for ensuring that features provide meaningful signals rather than noise. **Normalization techniques** including **z-score standardization**, **min-max scaling**, and **robust scaling** address the different scales and distributions of various features. **Winsorization** handles extreme outliers that could distort model training.(H. Wang et al., 2025)

Feature selection methods help identify the most predictive variables while reducing overfitting risks. **Mutual information-based selection** identifies features with high information content about target variables, while **correlation-based filtering** removes redundant features that provide similar information.(Baudry, 2019; S. Liu et al., 2019; S. Liu & Motani, 2022)
Time-aware feature selection considers the temporal stability of feature importance, ensuring that selected features maintain predictive power across different market conditions. **Rolling window feature selection** adapts feature sets to changing market dynamics, while **regime-specific selection** identifies features that are particularly informative during specific market conditions.(Ji et al., 2022; Moodi et al., 2023; Pabuccu & Barbu, 2024)

5. Evaluation Metrics and Benchmarks

5.1 Statistical Accuracy Metrics

The evaluation of financial forecasting models requires a comprehensive suite of metrics that capture different aspects of prediction quality. **Root Mean Square Error (RMSE)** and **Mean Absolute Error (MAE)** provide fundamental measures of point forecast accuracy, with RMSE placing greater emphasis on large errors due to its quadratic penalty structure.(Clements, 2023)

$$RMSE = \sqrt{1/n \sum(y_i - \hat{y}_i)^2}$$
$$MAE = 1/n \sum|y_i - \hat{y}_i|$$

Mean Absolute Percentage Error (MAPE) offers scale-independent evaluation, enabling comparison across assets with different price levels. However, MAPE can be problematic when actual values approach zero, leading to infinite or undefined values.(Guo, 2025)

Directional accuracy measures the percentage of correctly predicted price movement directions, providing insights into the model's ability to capture trend changes. This metric proves particularly valuable for trading applications where direction matters more than exact price levels. (Choudhary et al., 2025)

Symmetric Mean Absolute Percentage Error (sMAPE) addresses some limitations of MAPE by using the average of actual and predicted values in the denominator:

$$sMAPE = 100/n \sum|y_i - \hat{y}_i|/((|y_i| + |\hat{y}_i|)/2)$$

5.2 Financial Performance Metrics

Traditional statistical metrics, while useful for model comparison, may not fully capture the economic value of forecasting improvements. **Sharpe ratio** measures risk-adjusted returns by dividing excess returns by volatility:

$$\text{Sharpe Ratio} = (R_p - R_f)/\sigma_p$$

where R_p represents portfolio returns, R_f represents the risk-free rate, and σ_p represents portfolio volatility.(Doosti et al., 2024)

Information ratio compares active returns to tracking error, providing insights into the consistency of alpha generation:

$$\text{Information Ratio} = (R_p - R_b)/TE$$

where R_b represents benchmark returns and TE represents tracking error. (Vuletić et al., 2023)

Maximum drawdown measures the largest peak-to-trough decline in portfolio value, capturing downside risk that may not be reflected in volatility measures. This metric proves particularly important for institutional investors with specific risk constraints.(Song et al., 2025)

Calmar ratio combines return and downside risk by dividing annualized returns by maximum drawdown, providing a comprehensive measure of risk-adjusted performance. (Pistoia et al., 2021)

5.3 Risk Management Metrics

Financial forecasting models must demonstrate effectiveness in risk management applications, requiring specialized evaluation metrics. **Value at Risk (VaR)** estimates the potential loss over a specific time horizon at a given confidence level. The accuracy of VaR estimates is evaluated through **backtesting procedures** that examine the frequency of VaR violations. (Kwon & Lee, 2024)

Expected Shortfall (ES) or **Conditional Value at Risk (CVaR)** measures the expected loss conditional on exceeding the VaR threshold, providing insights into tail risk. ES addresses the limitation of VaR in not capturing the magnitude of extreme losses.(Joshi et al., 2022)

Volatility forecasting accuracy is evaluated through specialized metrics including **Mincer-Zarnowitz regressions** that test for unbiasedness and efficiency in volatility forecasts. The **Hansen-Lunde test** provides robust evaluation of volatility forecasting models across different loss functions.(Mironowicz et al., 2024)

5.4 Model Stability and Robustness

Financial models must demonstrate stability across different market conditions and time periods. **Rolling window evaluation** assesses model performance over consecutive time periods, revealing potential degradation or improvement over time. This approach helps identify models that maintain consistent performance across varying market regimes. (M. Wang & Hirsa, 2024)

Regime-based evaluation analyzes model performance during different market conditions including bull markets, bear markets, and high volatility periods. Models that perform well only during specific regimes may have limited practical value for long-term applications.(Ang & Timmermann, 2011; Botte & Bao, 2021; Ung et al., 2025)

Stress testing evaluates model performance during extreme market events such as financial crises or market crashes. This evaluation proves crucial for risk management applications where model failure during critical periods could have severe consequences.(Johri & Zhu, 2025)

Out-of-sample stability measures the consistency of model performance on unseen data, helping assess generalization capabilities. Time series cross-validation techniques ensure that evaluation procedures respect temporal dependencies in financial data.(Ericson et al., 2024)

5.5 Benchmark Comparison Standards

Establishing appropriate benchmarks is crucial for meaningful model evaluation. **Random walk benchmarks** provide a baseline for price level forecasting, while **historical mean benchmarks** serve as naive forecasting alternatives. The **efficient market hypothesis** suggests that simple benchmarks should be difficult to outperform consistently.(Jiajie & Liu, 2025)

Professional forecaster benchmarks compare model performance against analyst forecasts and institutional predictions. These benchmarks provide insights into whether sophisticated models can outperform human experts with access to fundamental analysis and market knowledge. (Bailey, n.d.-a; Bao et al., 2025)

Industry standard models including ARIMA, GARCH, and Vector Autoregression provide established benchmarks for comparing deep learning approaches. These comparisons help quantify the value added by increased model complexity. (Takahashi et al., 2019)

Ensemble benchmarks combine multiple simple models to create more challenging comparison standards. Simple averaging or median forecasts often provide surprisingly strong performance that sophisticated models must exceed to demonstrate value. (Mironowicz et al., 2024)

6. Comparative Analysis of Reviewed Studies

6.1 Methodological Approaches

The reviewed literature reveals distinct methodological patterns across different research objectives and application domains. **Single-asset prediction studies** typically focus on major stock indices or individual large-cap stocks, with the S&P 500, NASDAQ, and Dow Jones Industrial Average serving as primary test cases. These studies benefit from abundant historical data and high liquidity, enabling robust model training and evaluation. (Kanungo, 2025; S. Li et al., 2025b)

Multi-asset prediction frameworks address portfolio-level forecasting and cross-asset dependencies. These approaches require sophisticated architectures capable of modeling complex interdependencies while maintaining computational tractability. The curse of dimensionality becomes a significant challenge as the number of assets increases.(Deng & Lindauer, 2024; Zeng et al., 2023b)

Volatility forecasting studies concentrate on second-moment prediction, often building upon traditional GARCH frameworks. These applications prove particularly valuable for risk management and derivatives pricing, where volatility accuracy is paramount.(Di-Giorgi et al., 2025; Feng et al., 2023; Z. Xu et al., 2024)

Table 2 summarizes key methodological characteristics across reviewed studies.

Study Focus	Typical Dataset Size	Model Complexity	Training Duration	Primary Metrics
Single Asset Price	5-20 years daily	Medium	2-6 hours	RMSE, MAE, Directional Accuracy
Multi-Asset Portfolio	3-10 years daily	High	6-24 hours	Sharpe Ratio, Maximum Drawdown
Volatility Forecasting	10-30 years daily	Medium-High	4-12 hours	VaR Accuracy, ES Performance
High-Frequency Prediction	1-3 years intraday	Very High	12-48 hours	Tick-by-tick accuracy, Latency

6.2 Performance Comparison Across Architectures

Empirical results consistently demonstrate the superiority of deep learning approaches over traditional econometric models, though performance variations exist across different architectures and applications. **LSTM models** show robust performance across diverse forecasting tasks, with improvements of 15-25% in RMSE compared to ARIMA baselines. The ability to capture long-term dependencies proves particularly valuable for monthly and quarterly forecasting horizons. (Kanungo, 2025b; Varadharajan, Smith, Kalla, Kumar, et al., 2024)

Transformer architectures excel in multi-horizon forecasting scenarios, achieving 20-35% improvements in directional accuracy compared to traditional approaches. The self-attention mechanism enables these models to identify relevant historical patterns without the sequential processing constraints that limit RNN-based approaches. (S. Li et al., 2025c; Y. Li et al., 2024; Mozaffari & Zhang, 2024)

Hybrid CNN-LSTM models demonstrate superior performance in capturing both local patterns and global trends, with improvements of 25-40% in risk-adjusted returns compared to single-architecture approaches. The combination of convolutional feature extraction and recurrent temporal modeling proves particularly effective for complex market dynamics. (Y. Li & Pan, 2022; Zeng et al., 2023)

Ensemble methods consistently outperform individual models across multiple evaluation metrics, with variance reduction of 20-35% and bias improvements of 10-15%. The diversity of ensemble components proves crucial for achieving these improvements, with heterogeneous architectures outperforming homogeneous ensembles. (Y. Li & Pan, 2022; Seseri, 2023a)

6.3 Dataset and Market Characteristics

The choice of datasets and markets significantly influences model performance and generalizability. **Developed market studies** using US and European data typically report higher accuracy levels due to greater market efficiency and data quality. These markets benefit from extensive historical data, high liquidity, and sophisticated market infrastructure. (Kanungo, 2025; S. Li et al., 2025b)

Emerging market applications face additional challenges including data scarcity, higher volatility, and structural breaks. Models trained on developed market data often exhibit poor transfer performance to emerging markets, highlighting the importance of market-specific model development.(Borrageiro, 2023)

Cryptocurrency studies represent a growing subset of the literature, with unique challenges including extreme volatility, market manipulation, and regulatory uncertainty. The 24/7 trading nature of cryptocurrency markets provides continuous data streams but also introduces additional complexity in feature engineering and model training.(Shubham Ghelani, n.d.)

Asset class variations reveal differential model effectiveness across equities, bonds, commodities, and currencies. Equity prediction models generally achieve the highest accuracy levels, while commodity and currency forecasting prove more challenging due to fundamental economic drivers and geopolitical influences. (Z. Li & Huang, 2025)

6.4 Temporal Horizon Analysis

Forecasting accuracy varies significantly across different prediction horizons, with distinct patterns emerging from the reviewed literature. **Intraday prediction** (minutes to hours) achieves the highest accuracy levels but faces significant practical challenges, including data quality, transaction costs, and market microstructure effects. (Linné, 2024)

Daily prediction represents the most common forecasting horizon in the literature, offering a reasonable balance between prediction accuracy and practical utility. Most deep learning models demonstrate clear superiority over traditional approaches at this horizon. (Kanungo, 2025b; S. Li et al., 2025c; Varadharajan, Smith, Kalla, Kumar, et al., 2024)

Weekly and monthly prediction faces increased challenges due to the growing influence of fundamental economic factors and reduced signal-to-noise ratios. Deep learning models maintain advantages over traditional approaches, albeit with diminishing margins. (Zeng et al., 2023)

Long-term prediction (quarterly and annual) approaches the limits of statistical predictability, with all models showing degraded performance. The efficient market hypothesis suggests that long-term predictability should be limited, and empirical results generally support these theoretical expectations. (Seseri, 2023b)

6.5 Evaluation Methodology Consistency

The reviewed literature reveals significant inconsistencies in evaluation methodologies that complicate direct performance comparisons. **Walk-forward validation** is adopted by approximately 60% of studies, while the remaining 40% use various forms of train-test splits that may introduce look-ahead bias. (Ferrari, 2022)

Benchmark selection varies widely, with some studies comparing against naive random walk models while others use sophisticated econometric baselines. This variation makes it difficult to assess the true economic value of deep learning improvements. (Tam, 2021)

Statistical significance testing is performed in fewer than 30% of reviewed studies, limiting confidence in reported performance improvements. The absence of significance testing is particularly problematic given the noisy nature of financial data and the potential for spurious correlations.(Liladhar Rane et al., 2024)

Transaction cost consideration is addressed in only 25% of studies, despite its crucial importance for practical trading applications. Studies that ignore transaction costs may overestimate the economic value of forecasting improvements, particularly for high-frequency trading strategies. (C. W. S. Chen et al., 2008; Cipollini et al., 2017; Sadon et al., 2024)

7. Challenges and Limitations

7.1 Overfitting and Generalization

Overfitting represents one of the most significant challenges in applying deep learning to financial time series. The high-dimensional parameter spaces of neural networks, combined with the noisy nature of financial data, create substantial risks of fitting to spurious patterns rather than genuine market dynamics. This problem is exacerbated by the relatively small sample sizes available for financial modeling compared to other deep learning domains. (Kong, Chen, Liu, Ning, Zhang, Muhammad Marier, et al., 2025)

Data snooping bias emerges when researchers iteratively test multiple model configurations on the same dataset, effectively using test data for model selection. This practice leads to inflated performance estimates that fail to generalize to true out-of-sample conditions. The limited availability of truly independent financial datasets makes this problem particularly severe in financial applications. (Awalee Consulting, 2023)

Temporal overfitting occurs when models learn patterns specific to particular time periods or market regimes that do not persist in future data. Financial markets exhibit structural breaks and regime changes that can render historically optimal models suboptimal in new environments. The 2008 financial crisis and COVID-19 pandemic provide stark examples of how market dynamics can shift fundamentally. (Seseri, 2023b)

Cross-validation challenges in time series contexts require specialized approaches that respect temporal dependencies. Traditional k-fold cross-validation can introduce look-ahead bias by using future information to predict past events. Time series cross-validation techniques, while more appropriate, often yield smaller training sets and higher variance in performance estimates. (Mohammed & Kora, 2023)

7.2 Model Interpretability and Explainability

The **black box nature** of deep learning models poses significant challenges for financial applications where regulatory compliance and risk management require transparent decision-making processes. Traditional econometric models provide clear parameter interpretations and established hypothesis testing frameworks, while deep learning models operate through complex nonlinear transformations that resist easy interpretation. (Milvus, n.d.)

Regulatory requirements in financial services increasingly demand model explainability, particularly for applications affecting consumer credit decisions or systemic risk assessment. The European Union's GDPR "right to explanation" and similar regulations worldwide create legal obligations for model transparency that many deep learning approaches struggle to satisfy. (Černevičienė & Kabašinskas, 2024)

Risk management applications require understanding not just what a model predicts, but why it makes specific predictions and how sensitive those predictions are to input changes. Traditional risk models enable stress testing and scenario analysis through parameter modification, while deep learning models typically require complex perturbation studies to assess sensitivity. (Karadaş et al., 2025)

Stakeholder communication becomes challenging when model predictions cannot be easily explained to non-technical decision-makers, including portfolio managers, traders, and senior executives. The intuitive appeal of simple models often outweighs the superior accuracy of complex approaches in practical decision-making contexts. (H. Wu, 2024)

7.3 Data Quality and Availability

Data preprocessing complexity in financial applications requires sophisticated handling of missing values, outliers, and structural breaks. Unlike image or text data where preprocessing is relatively standardized, financial data preprocessing involves domain-specific considerations including dividend adjustments, stock splits, and corporate actions. (Černevičienė & Kabašinskas, 2024; Yeo et al., 2025)

Survivorship bias affects many financial datasets where delisted companies or failed assets are excluded from historical records. This bias can lead to overly optimistic performance estimates and models that fail to account for the full range of possible outcomes. (Karadaş et al., 2025; J. Liu, 2025)

Look-ahead bias can inadvertently creep into feature engineering processes when using information that would not have been available at the time of prediction. This is particularly problematic when incorporating external data sources with different reporting delays or revision schedules. (Hollis et al., 2018; Varadharajan, Smith, Kalla, Kumar, et al., 2024; Wen & Li, 2023)

Alternative data challenges include irregular update frequencies, variable data quality, and potential for structural breaks when data sources change their collection methodologies. Social media data, satellite imagery, and web scraping initiatives face constant evolution that can disrupt model stability.(Mohsin & Nasim, 2025)

7.4 Computational Demands and Scalability

Training computational requirements for deep learning models can be substantial, particularly for transformer architectures and ensemble methods. The financial industry's emphasis on real-time decision-making conflicts with the extended training times required for sophisticated models. (J. Liu, 2025; Mondal et al., 2025)

Memory constraints become significant when processing high-frequency data or large cross-sections of assets. Transformer models with quadratic attention complexity face particular challenges when applied to long financial time series or broad market coverage. (Kumar et al., 2023)

Inference latency requirements vary dramatically across financial applications, from microsecond constraints in high-frequency trading to daily batch processing for portfolio optimization. Models that perform well in offline evaluation may prove impractical for latency-sensitive applications.(Yeo et al., 2025)

Infrastructure costs associated with deep learning deployment can be substantial, including specialized hardware requirements, cloud computing expenses, and ongoing maintenance costs. Cost-benefit analysis becomes crucial when comparing sophisticated models against simpler alternatives. (J. Wang et al., 2023)

7.5 Market Regime Changes and Nonstationarity

Structural breaks in financial markets can render historical patterns obsolete, causing previously successful models to fail dramatically. The transition from fixed to floating exchange rates, financial deregulation, and algorithmic trading adoption represent examples of structural changes that altered market dynamics fundamentally. (Bailey, n.d.-b)

Regime switching behavior in financial markets means that relationships between variables can change systematically over time. Models trained during one regime may perform poorly when market conditions shift, requiring adaptive approaches that can detect and respond to regime changes. (L. Li et al., 2023; S. Li & Tang, 2024; W. Zhang et al., 2024)

Concept drift occurs when the underlying relationships being modeled evolve gradually over time. Unlike abrupt structural breaks, concept drift is often subtle and difficult to detect until model performance has already degraded significantly. (Literal Labs, 2025)

Non-stationarity in financial time series violates many of the assumptions underlying traditional statistical models and can also affect deep learning approaches. While deep learning models are generally more robust to non-stationarity than linear models, they still face challenges when the data generating process changes fundamentally. (Che et al., 2024)

8. Research Gap Identification

8.1 Explainable AI in Financial Applications

Despite growing regulatory pressure for model transparency, **explainable AI (XAI) techniques** specifically designed for financial deep learning models remain underdeveloped. Current XAI approaches, primarily developed for computer vision and natural language processing, often prove inadequate for the temporal and multivariate nature of financial data. (Forough & Momtazi, 2020; Tanikonda et al., 2025)

Financial-specific explanation frameworks are needed that can provide economically meaningful interpretations of model decisions. Traditional attribution methods like LIME and SHAP may identify which features are important but fail to explain the economic rationale behind predictions. New approaches should connect model outputs to established financial theories and market mechanisms. (Tavakoli et al., 2024)

Regulatory-compliant explanation standards require development to bridge the gap between technical model capabilities and legal requirements. These standards should specify minimum levels of explainability for different types of financial decisions and provide guidelines for documentation and validation procedures. (H. Wang et al., 2025)

Dynamic explainability represents an underexplored area where explanation quality may vary across different market conditions or time periods. Models may rely on different features during normal versus stressed market conditions, requiring explanation frameworks that can adapt to changing market dynamics. (Baudry, 2019; Takahashi et al., 2019)

8.2 Transfer Learning and Domain Adaptation

Cross-market transfer learning remains largely unexplored despite its potential for emerging market applications where local data may be scarce. Developed market models could potentially transfer useful patterns to emerging markets, but current approaches lack sophisticated domain adaptation techniques that account for structural differences between markets.(Jiajie & Liu, 2025; Jiang et al., 2024; Joshi et al., 2022)

Cross-asset transfer learning offers opportunities to leverage patterns learned from liquid assets to improve predictions for less liquid securities. The relationships between different asset classes could enable transfer learning approaches that improve performance for assets with limited historical data. (Clements, 2023)

Temporal transfer learning could enable models trained on historical data to adapt more quickly to new market regimes. This approach could prove particularly valuable during crisis periods when rapid adaptation is crucial but limited data is available for the new regime. (Guo, 2025)

Multi-resolution transfer learning remains underexplored, where models trained on high-frequency data could inform lower-frequency predictions and vice versa. The hierarchical nature of market dynamics across different time scales suggests that cross-frequency transfer learning could yield significant benefits. (Choudhary et al., 2025)

8.3 Online Learning and Adaptive Systems

Real-time model adaptation capabilities are crucial for maintaining performance in dynamic financial markets, yet most current approaches rely on periodic retraining rather than continuous learning. Online learning algorithms that can update model parameters incrementally as new data arrives could provide significant advantages in rapidly changing markets. (Doosti et al., 2024)

Catastrophic forgetting in continual learning contexts poses particular challenges for financial applications where models must adapt to new conditions while retaining knowledge of recurring patterns. Financial markets exhibit both trend changes and mean reversion, requiring models that can distinguish between permanent shifts and temporary deviations. (Vuletić et al., 2023)

Adaptive ensemble methods that can dynamically adjust the weights of component models based on recent performance represent an underexplored area. These approaches could automatically detect when specific models are performing poorly and adjust the ensemble composition accordingly. (Song et al., 2025)

Online feature selection in streaming data environments remains challenging, particularly when dealing with high-dimensional alternative data sources. Algorithms that can identify relevant features in real-time while discarding irrelevant or redundant information could significantly improve model efficiency and performance. (Pistoia et al., 2021)

8.4 Multimodal and Heterogeneous Data Integration

Cross-modal attention mechanisms specifically designed for financial applications remain underdeveloped. While attention mechanisms have proven successful in computer vision and NLP, their adaptation to financial contexts where different data modalities have varying predictive power and temporal characteristics requires specialized approaches. (Kwon & Lee, 2024)

Data fusion architectures that can handle the diverse characteristics of financial data sources, including different frequencies, missing data patterns, and quality levels, represent a significant research gap. Current approaches often require extensive preprocessing to align different data sources, limiting their practical applicability. (Joshi et al., 2022)

Causal multimodal learning could enable models to understand not just correlations between different data sources but actual causal relationships. This capability would be particularly valuable for understanding how news events, social media sentiment, and other alternative data sources actually influence asset prices. (Mironowicz et al., 2024)

Quality-aware fusion techniques that can assess and weight different data sources based on their current reliability and relevance remain largely unexplored. Financial markets are subject to data quality issues, and models that can automatically detect and adjust for poor-quality inputs could significantly improve robustness. (Dixit, 2024)

8.5 Generalization Across Market Conditions

Out-of-distribution detection for financial models remains an underexplored area despite its critical importance for risk management. Models should be able to identify when current market conditions are significantly different from their training data and adjust their confidence accordingly. (Jiajie & Liu, 2025; Song et al., 2025)

Domain shift detection techniques could help identify when market regimes have changed sufficiently to warrant model retraining or architecture modification. Early detection of such shifts could prevent significant performance degradation and enable proactive model management. (Johri & Zhu, 2025)

Robust optimization approaches specifically designed for financial deep learning models could improve performance during market stress periods. Traditional robust optimization focuses on worst-case scenarios, but financial applications require approaches that balance robustness with performance during normal market conditions. (Ericson et al., 2024)

Meta-learning approaches for financial forecasting could enable models to quickly adapt to new market conditions or asset classes. These approaches could learn general patterns about how financial models should adapt rather than focusing on specific prediction tasks. (Jiajie & Liu, 2025)

9. Future Directions

9.1 Transfer Learning and Cross-Market Applications

Domain-adaptive architectures represent a promising direction for extending successful models across different markets and time periods. These approaches could enable models trained on developed markets to transfer effectively to emerging markets by explicitly accounting for distributional differences and structural variations. (Doosti et al., 2024; Pistoia et al., 2021)

Few-shot learning techniques could address data scarcity issues in emerging markets or newly listed securities. These approaches would enable models to make reasonable predictions with minimal historical data by leveraging patterns learned from similar assets or markets. (Takahashi et al., 2019)

Progressive transfer learning could enable gradual adaptation from source to target domains, potentially reducing negative transfer effects that can occur when domains are too dissimilar. This approach could be particularly valuable for cross-country applications where economic structures differ significantly. (Mironowicz et al., 2024)

Causal transfer learning represents an emerging area that could enable transfer of causal relationships rather than just correlations. Understanding why certain patterns exist in source domains could improve transfer effectiveness to target domains with different characteristics. (Borrageiro, 2023)

9.2 Reinforcement Learning and Strategic Applications

Multi-agent reinforcement learning could model the strategic interactions between different market participants, providing insights into market dynamics that single-agent approaches cannot capture. These models could help understand how algorithmic trading strategies interact and potentially lead to improved trading performance. (X. Zhang et al., 2024)

Hierarchical reinforcement learning could enable models to operate across multiple time scales simultaneously, making both short-term tactical decisions and long-term strategic allocations. This capability could prove particularly valuable for institutional investors with complex multi-horizon objectives. (Dixit, 2024; Ericson et al., 2024; R. Wu, 2024)

Risk-constrained reinforcement learning could incorporate sophisticated risk management requirements directly into the learning process. Traditional reinforcement learning optimizes expected returns, but financial applications require explicit consideration of downside risk and regulatory constraints. (Tanaka et al., 2025)

Offline reinforcement learning using historical market data could enable strategy development without the risks associated with live trading. These approaches could learn effective trading strategies from past market data while accounting for the challenges of deploying such strategies in live markets. (K. Xu et al., 2022)

9.3 Explainable AI and Regulatory Compliance

Causal explanation frameworks specifically designed for financial applications could provide economically meaningful explanations of model decisions. These frameworks should connect model predictions to established financial theories and provide insights into the economic mechanisms driving predictions. (Nettey & Ansong, 2025)

Counterfactual explanation systems could help stakeholders understand how different market conditions or input values would affect model predictions. This capability would be particularly valuable for scenario analysis and stress testing applications. (Bastos et al., 2025; Noguer I Alonso & Pereira Franklin, 2024; M. Xiao et al., 2025)

Model-agnostic explanation techniques that work across different deep learning architectures could provide consistent explanation standards for financial institutions using diverse modeling approaches. These techniques should scale efficiently to the high-dimensional feature spaces common in financial applications. (Kekana et al., 2024)

Regulatory-compliant documentation systems could automatically generate the explanations and documentation required by financial regulators. These systems should translate technical model information into forms accessible to non-technical stakeholders and regulatory reviewers.(Xia, 2019)

9.4 Multimodal Fusion and Alternative Data

Cross-modal attention mechanisms specifically designed for financial data could enable more effective integration of diverse information sources. These mechanisms should account for the different temporal characteristics and reliability levels of various data sources. (Fang et al., 2023)

Real-time multimodal learning could enable models to incorporate streaming alternative data sources as they become available. This capability would be particularly valuable for news-based trading and social media sentiment analysis applications. (Nithish Kumar et al., 2024)

Quality-aware fusion architectures could automatically assess and weight different data sources based on their current reliability and relevance. These approaches could detect when particular data sources are providing low-quality information and adjust accordingly. (Muallim, 2018)

Causal multimodal understanding could enable models to distinguish between genuine causal relationships and spurious correlations across different data modalities. This capability would be crucial for building robust models that generalize across different market conditions. (W. Zhang et al., 2022)

9.5 Online Learning and Adaptation

Continual learning architectures specifically designed for financial time series could enable models to adapt continuously to changing market conditions without forgetting useful historical patterns. These approaches must balance plasticity for new learning with stability for existing knowledge. (Neloy & Turgeon, 2024)

Meta-learning for rapid adaptation could enable models to quickly adjust to new market regimes or structural breaks. These approaches could learn general adaptation strategies that apply across different types of market changes. (Uddin et al., 2023)

Federated learning for financial institutions could enable collaborative model development while preserving data privacy and competitive advantages. This approach could be particularly valuable for risk management applications where data sharing could improve system-wide stability. (Benadict & Raju, 2024)

Adaptive ensemble methods that automatically adjust model combinations based on changing market conditions could provide more robust performance across different regimes. These methods should detect when individual models are performing poorly and adjust ensemble weights accordingly.

10. Key findings

From our analysis, we concluded:

Architectural superiority: Transformer-based models and hybrid CNN-LSTM architectures achieve 20-35% improvements in directional accuracy compared to traditional approaches, with particularly strong performance in multi-horizon forecasting scenarios. The self-attention mechanism proves especially effective for capturing long-range dependencies that characterize financial time series.

Performance consistency: Deep learning models demonstrate robust performance across diverse market conditions, with ensemble methods showing 20-35% variance reduction compared to individual models. This consistency proves particularly valuable during market stress periods where traditional models often fail.

Integration capabilities: Modern deep learning frameworks successfully incorporate multimodal data sources including news sentiment, social media data, and alternative information sources, achieving 15-25% performance improvements over single-modality approaches.

Practical limitations: Despite superior statistical performance, deep learning models face significant challenges in interpretability, computational requirements, and regulatory compliance that limit their adoption in many practical applications.

Research gaps identified through our analysis highlight critical areas requiring continued investigation:

Explainable AI development: The financial industry's regulatory requirements demand model transparency that current deep learning approaches struggle to provide. Future research must develop explanation frameworks specifically designed for financial applications that can satisfy both technical and regulatory requirements.

Transfer learning applications: The potential for cross-market and cross-asset transfer learning remains largely unexplored despite obvious practical benefits for emerging markets and data-scarce scenarios. Sophisticated domain adaptation techniques could enable broader application of successful models.

Online adaptation systems: The dynamic nature of financial markets requires models that can adapt continuously to changing conditions while retaining useful historical knowledge. Current approaches rely primarily on periodic retraining rather than true online learning.

Multimodal integration: While promising results exist for incorporating alternative data sources, optimal fusion architectures and quality assessment mechanisms remain underdeveloped.

Future research priorities should focus on:

1. **Developing financial-specific XAI techniques** that provide economically meaningful explanations while satisfying regulatory requirements
2. **Advancing transfer learning methodologies** for cross-market and cross-asset applications
3. **Creating robust online learning systems** that can adapt to regime changes without catastrophic forgetting
4. **Improving multimodal fusion architectures** with quality-aware data integration capabilities
5. **Establishing standardized evaluation protocols** that include transaction costs, statistical significance testing, and regime-specific performance analysis

11. Conclusion

This comprehensive review of deep learning applications in financial time series analysis reveals a field characterized by rapid innovation and significant achievements, yet facing substantial challenges that demand continued research attention. Our systematic analysis of multiple studies published between 2014-2025 demonstrates clear evidence that deep learning approaches consistently outperform traditional econometric models across multiple evaluation criteria.

The convergence of deep learning and financial forecasting represents a fundamental shift in how we approach market prediction and risk management. While significant challenges remain, the continued advancement of these techniques holds promise for more accurate, robust, and practically useful financial forecasting systems. Success in this endeavor will require continued collaboration between machine learning researchers, financial practitioners, and regulatory authorities to develop approaches that are technically sophisticated, economically valuable, and societally responsible.

The path forward demands not just algorithmic innovation but also careful consideration of practical implementation challenges, regulatory requirements, and ethical implications. As deep learning continues to reshape financial markets, ensuring that these powerful tools serve broader economic welfare while maintaining market stability and fairness will remain paramount considerations for future research and development efforts.

Bibliography

Abbasimehr, H., & Paki, R. (2022). Improving time series forecasting using LSTM and attention models. Journal of Ambient Intelligence and Humanized Computing, 13(1), 673–691. https://doi.org/10.1007/S12652-020-02761-X

AbdElnapi, N. M. M., Darwish, A., & Hassanien, A. E. (2024). Artificial Intelligence in Finance Sector for Risk Prediction. Studies in Systems, Decision and Control, 542, 297–312. https://doi.org/10.1007/978-3-031-63451-2_17

Abu Jamie, N. H., Abu-Jamie, T. N., & Al-Absy, M. S. M. (2024). Advances in AI and Their Effects on Finance and Economic Analysis. Studies in Systems, Decision and Control, 524, 507–523. https://doi.org/10.1007/978-3-031-54379-1_44

Alzahrani, A. (2024). Explainable AI-based Framework for Efficient Detection of Spam from Text using an Enhanced Ensemble Technique. Engineering, Technology & Applied Science Research, 14(4), 15596–15601. https://doi.org/10.48084/ETASR.7901

Ang, A., & Timmermann, A. (2011). Regime Changes and Financial Markets (17182). https://doi.org/10.3386/W17182

Awalee Consulting. (2023, September 6). Why And How To Use Ensemble Methods in Financial Machine Learning? Awalee Consulting. https://www.awaleeconsulting.com/why-and-how-to-use-ensemble-methods-in-financial-machine-learning/

Bailey, K. (n.d.-a). Forecasting Model Comparison: When You Should Use Each. Corporate Finance Institute. Retrieved January 1, 2025, from https://corporatefinanceinstitute.com/resources/financial-modeling/forecasting-model-comparison/

Bailey, K. (n.d.-b). Why Explainable AI is Critical for Financial Decision Making. Corporate Finance Institute. Retrieved January 1, 2025, from https://corporatefinanceinstitute.com/resources/artificial-intelligence-ai/why-explainable-ai-matters-finance/

Bao, T., Corgnet, B., Hanaki, N., Okada, K., Riyanto, Y. E., & Zhu, J. (2025). Financial forecasting in the lab and the field: Qualified professionals vs. smart students. Journal of

Behavioral and Experimental Finance, 46, 101051. https://doi.org/10.1016/J.JBEF.2025.101051

Bastos, Z., Freitas, J., Franco, J., & Caminha, C. (2025). Prompt-Driven Time Series Forecasting with Large Language Models. Conference: 27th International Conference on Enterprise Information Systems, 309–316. https://doi.org/10.5220/0013363800003929

Baudry, D. (2019). Generating financial time series with generative adversarial networks ([Source code]). github. https://github.com/DBaudry/Generating_Financial_Time_Series_with_GAN

Benadict, A., & Raju, A. (2024). Federated Learning for Privacy-Preserving Financial Data Sharing. IJIRMPS - International Journal of Innovative Research in Engineering & Multidisciplinary Physical Sciences, 12(3), 2349–7300. https://doi.org/10.5281/ZENODO.14209123

Bhambu, A., Bera, K., Natarajan, S., & Suganthan, P. N. (2025). High frequency volatility forecasting and risk assessment using neural networks-based heteroscedasticity model. Engineering Applications of Artificial Intelligence, 149, 110397. https://doi.org/10.1016/J.ENGAPPAI.2025.110397

Borrageiro, G. F. (2023). Online learning in financial time series. In Doctoral thesis, UCL (University College London).

Botte, A., & Bao, D. (2021). A Machine Learning Approach to Regime Modeling - Two Sigma. https://www.twosigma.com/articles/a-machine-learning-approach-to-regime-modeling/

Bruederle, A., & Hodler, R. (2018). Nighttime lights as a proxy for human development at the local level. PLoS ONE, 13(9), e0202231. https://doi.org/10.1371/JOURNAL.PONE.0202231

Černevičienė, J., & Kabašinskas, A. (2024). Explainable artificial intelligence (XAI) in finance: a systematic literature review. Artificial Intelligence Review 2024 57:8, 57(8), 1–45. https://doi.org/10.1007/S10462-024-10854-8

Che, W., Wang, Z., Jiang, C., & Abedin, M. Z. (2024). Predicting financial distress using multimodal data: An attentive and regularized deep learning method. Information Processing & Management, 61(4), 103703. https://doi.org/10.1016/J.IPM.2024.103703

Chen, C. W. S., Gerlach, R., & Lin, E. M. H. (2008). Volatility forecasting using threshold heteroskedastic models of the intra-day range. Computational Statistics & Data Analysis, 52(6), 2990–3010. https://ideas.repec.org/a/eee/csdana/v52y2008i6p2990-3010.html

Chen, Q. (2025). Comparing Different Transformer Model Structures for Stock Prediction. https://arxiv.org/pdf/2504.16361v1

Chen, W., Hussain, W., Cauteruccio, F., & Zhang, X. (2023). Deep Learning for Financial Time Series Prediction: A State-of-the-Art Review of Standalone and Hybrid Models. Computer Modeling in Engineering & Sciences, 139(1), 187–224. https://doi.org/10.32604/CMES.2023.031388

Choudhary, H., Orra, A., Sahoo, K., & Thakur, M. (2025). Risk-Adjusted Deep Reinforcement Learning for Portfolio Optimization: A Multi-reward Approach. International Journal of Computational Intelligence Systems, 18(1), 1–19. https://doi.org/10.1007/S44196-025-00875-8/TABLES/3

Cipollini, F., Gallo, G. M., & Otranto, E. (2017). On Heteroskedasticity and Regimes in Volatility Forecasting. SSRN Electronic Journal. https://doi.org/10.2139/SSRN.3037550

Clements, W. (2023). How quantum can enhance machine learning in finance . ORCA Computing. https://www.financederivative.com/how-quantum-can-enhance-machine-learning-in-finance/

Confraria, H., Godinho, M. M., & Mendonça, S. (2024). Patents as Indicators for Strategic Management. Reference Module in Social Sciences. https://doi.org/10.1016/B978-0-443-13701-3.00527-2

Corporate Finance Institute. (n.d.). Ensemble Methods - Overview, Categories, Main Types. Retrieved January 1, 2025, from https://corporatefinanceinstitute.com/resources/data-science/ensemble-methods/

Deng, D., & Lindauer, M. (2024). Optimizing Time Series Forecasting Architectures: A Hierarchical Neural Architecture Search Approach. https://arxiv.org/pdf/2406.05088

Di-Giorgi, G., Salas, R., Avaria, R., Ubal, C., Rosas, H., & Torres, R. (2025). Volatility forecasting using deep recurrent neural networks as GARCH models. Computational Statistics, 40(6), 3229–3255. https://doi.org/10.1007/S00180-023-01349-1/METRICS

Dixit, S. (2024). Advanced Generative AI Models for Fraud Detection and Prevention in FinTech: Leveraging Deep Learning and Adversarial Networks for Real-Time Anomaly Detection in Financial Transactions. Journal of Artificial Intelligence Research, 4(2), 51–77. https://thesciencebrigade.com/JAIR/article/view/402

Doosti, M., Wallden, P., Hamill, C. B., Hankache, R., Brown, O. T., & Heunen, C. (2024). A Brief Review of Quantum Machine Learning for Financial Services. https://arxiv.org/pdf/2407.12618

Érdi, P., Makovi, K., Somogyvári, Z., Strandburg, K., Tobochnik, J., Volf, P., & Zalányi, L. (2013). Prediction of Emerging Technologies Based on Analysis of the U.S. Patent Citation Network. Scientometrics, 95(1), 225–242. https://doi.org/10.1007/s11192-012-0796-4

Ericson, L., Zhu, X., Han, X., Fu, R., Li, S., Guo, S., & Hu, P. (2024). Deep Generative Modeling for Financial Time Series with Application in VaR: A Comparative Review. https://arxiv.org/pdf/2401.10370

Fang, Z., Ma, X., Pan, H., Yang, G., & Arce, G. R. (2023). Movement forecasting of financial time series based on adaptive LSTM-BN network. Expert Systems with Applications, 213, 119207. https://doi.org/10.1016/J.ESWA.2022.119207

Feng, W., Li, Y., Zhang, X., Feng, W., Li, Y., & Zhang, X. (2023). A mixture deep neural network GARCH model for volatility forecasting. Electronic Research Archive 2023 7:3814, 31(7), 3814–3831. https://doi.org/10.3934/ERA.2023194

Ferrari, E. F. (2022). Volatility forecasting with GARCH models and recurrent neural networks [Master's, Nova School of Business & Economics]. https://run.unl.pt/handle/10362/154317

Forough, J., & Momtazi, S. (2020). Ensemble of deep sequential models for credit card fraud detection. Applied Soft Computing, 99. https://doi.org/10.1016/J.ASOC.2020.106883

Gao, Z., & Kuruoğlu, E. E. (2024). Attention based hybrid parametric and neural network models for non-stationary time series prediction. Expert Systems, 41(2). https://doi.org/10.1111/EXSY.13419

Guo, Y. (2025). Application of Generative Adversarial Networks in Financial Time Series Data Augmentation. https://doi.org/10.2139/SSRN.5142078

Hollis, T., Viscardi, A., & Yi, S. E. (2018). A Comparison of LSTMs and Attention Mechanisms for Forecasting Financial Time Series. https://arxiv.org/pdf/1812.07699

Ji, G., Yu, J., Hu, K., Xie, J., & Ji, X. (2022). An adaptive feature selection schema using improved technical indicators for predicting stock price movements. Expert Systems with Applications, 200, 116941. https://doi.org/10.1016/J.ESWA.2022.116941

Jiajie, W., & Liu, L. (2025). Portfolio Optimization through a Multi-modal Deep Reinforcement Learning Framework. Authorea Preprints, 14(8). https://doi.org/10.36227/TECHRXIV.174123926.62412968/V1

Jiang, Y., Olmo, J., & Atwi, M. (2024). Deep reinforcement learning for portfolio selection. Global Finance Journal, 62, 101016. https://doi.org/10.1016/J.GFJ.2024.101016

Johri, S., & Zhu, E. (2025). Generative Quantum Machine Learning for Finance. https://ionq.com/resources/generative-quantum-machine-learning-for-finance

Joshi, D., Patil, A., & Mutyalwar, S. (2022). Portfolio Optimization using Reinforcement Learning – IJSREM. INTERANTIONAL JOURNAL OF SCIENTIFIC RESEARCH IN

ENGINEERING AND MANAGEMENT, 10(6). https://doi.org/https://doi.org/10.55041/ijsrem16672

Kanungo, P. (2025). Time series forecasting in financial markets using deep learning models. World Journal of Advanced Engineering Technology and Sciences, 15(1), 709–719. https://doi.org/10.30574/WJAETS.2025.15.1.0167

Karadaş, F., Eravcı, B., & Özbayoğlu, A. M. (2025). Multimodal Stock Price Prediction. International Conference on Agents and Artificial Intelligence, 3, 687–694. https://doi.org/10.5220/0013174500003890

Kekana, M., Sumbwanyambe, M., & Hlalele, T. (2024). Advanced Machine Learning in Quantitative Finance Using Graph Neural Networks. Journal of Advances in Information Technology, 15(9), 1025–1034. https://doi:10.12720/jait.15.9.1025-1034

Kong, X., Chen, Z., Liu, W., Ning, K., Zhang, L., Marier, S. M., Liu, Y., Chen, Y., & Xia, F. (2025). Deep Learning for Time Series Forecasting: A Survey. International Journal of Machine Learning and Cybernetics, 16(7–8), 5079–5112. https://doi.org/10.1007/s13042-025-02560-w

Kumar, I., Tripathi, B. K., & Singh, A. (2023). Attention-based LSTM network-assisted time series forecasting models for petroleum production. Engineering Applications of Artificial Intelligence, 123, 106440. https://doi.org/10.1016/J.ENGAPPAI.2023.106440

Kwon, S., & Lee, Y. (2024). Can GANs Learn the Stylized Facts of Financial Time Series? ICAIF 2024 - 5th ACM International Conference on AI in Finance, 126–133. https://doi.org/10.1145/3677052.3698661

Li, J., Sun, L., Wu, X., & Tao, C. (2024). Enhancing Financial Time Series Forecasting with Hybrid Deep Learning: Ceemdan-Informer-Lstm Model. Available at SSRN: https://ssrn.com/abstract=4923415

Li, J., Zhou, L., & Song, L. (2023). Attention-Ordering LSTM for Financial Time Series Forecasting. 2023 7th International Conference on Electrical, Mechanical and Computer Engineering, ICEMCE 2023, 990–994. https://doi.org/10.1109/ICEMCE60359.2023.10490775

Li, L., Chang, T.-Y., & Wang, H. (2023). Multimodal Gen-AI for Fundamental Investment Research. https://arxiv.org/pdf/2401.06164

Li, S., Schulwolf, Z. B., & Miikkulainen, R. (2025). Transformer Based Time-Series Forecasting for Stock. https://arxiv.org/pdf/2502.09625v1

Li, S., & Tang, H. (2024). Multimodal Alignment and Fusion: A Survey. https://arxiv.org/pdf/2411.17040

Li, X., Yu, W., Luwang, T., Zheng, J., Qiu, X., Zhao, J., Xia, L., & Li, Y. (2017). Transaction Fraud Detection Using GRU-centered Sandwich-structured Model. Proceedings of the 2018 IEEE 22nd International Conference on Computer Supported Cooperative Work in Design, CSCWD 2018, 761–766. https://doi.org/10.1109/CSCWD.2018.8465147

Li, Y., & Pan, Y. (2022). A novel ensemble deep learning model for stock prediction based on stock prices and news. International Journal of Data Science and Analytics, 13(2), 139–149. https://doi.org/10.1007/S41060-021-00279-9,

Li, Y., Wang, X., Guo, Y., Li, Y., Wang, X., & Guo, Y. (2024). CNN-Trans-SPP: A small Transformer with CNN for stock price prediction. Electronic Research Archive 2024 12:6717, 32(12), 6717–6732. https://doi.org/10.3934/ERA.2024314

Li, Z., & Huang, Q. (2025). Option Pricing Using Ensemble Learning. https://arxiv.org/pdf/2506.05799v1

Liladhar Rane, N., Choudhary, S. P., & Rane, J. (2024). Ensemble deep learning and machine learning: applications, opportunities, challenges, and future directions. Studies in Medical and Health Sciences, 1(2), 18–41. https://doi.org/10.48185/SMHS.V1I2.1225

Linné, A. (2024). Neural Networks in Volatility Forecasting: A comparative analysis of GARCH and neural networks on the Swedish, Polish and Slovenian stock indexes [Karlstad University]. https://urn.kb.se/resolve?urn=urn:nbn:se:kau:diva-101489

Literal Labs. (2025). Interpretability vs. Explainability in AI and Machine Learning. https://www.literal-labs.ai/explainable-ai/interpretability-vs-explainability.html

Liu, J. (2025). Multimodal Data-Driven Factor Models for Stock Market Forecasting. Journal of Computer Technology and Software, 4(2), 2025. https://doi.org/10.5281/ZENODO.14984969

Liu, S., & Motani, M. (2022). Improving Mutual Information based Feature Selection by Boosting Unique Relevance. Journal of Artificial Intelligence Research, 82, 1267–1292. https://doi.org/10.1613/jair.1.17219

Liu, S., Yao, J., Zhou, C., & Motani, M. (2019). Improving Mutual Information based Feature Selection by Boosting Unique Relevance. Proceedings - 2018 IEEE International Conference on Bioinformatics and Biomedicine, BIBM 2018, 687–692. https://doi.org/10.1109/BIBM.2018.8621163

Mateo-Sanchis, A., Piles, M., Muñoz-Marí, J., Adsuara, J. E., Pérez-Suay, A., & Camps-Valls, G. (2019). Synergistic integration of optical and microwave satellite data for crop yield estimation. Remote Sensing of Environment, 234. https://doi.org/10.1016/j.rse.2019.111460

Milvus. (n.d.). How do attention mechanisms enhance time series forecasting models? Retrieved January 1, 2025, from https://milvus.io/ai-quick-reference/how-do-attention-mechanisms-enhance-time-series-forecasting-models

Mironowicz, P., H., A. S., Mandarino, A., Yilmaz, A. E., & Ankenbrand, T. (2024a). Applications of Quantum Machine Learning for Quantitative Finance. https://arxiv.org/pdf/2405.10119

Mishra, A. K., Renganathan, J., & Gupta, A. (2024). Volatility forecasting and assessing risk of financial markets using multi-transformer neural network based architecture. Engineering Applications of Artificial Intelligence, 133, 108223. https://doi.org/10.1016/J.ENGAPPAI.2024.108223

Moghar, A., & Hamiche, M. (2020). Stock Market Prediction Using LSTM Recurrent Neural Network. Procedia Computer Science, 170, 1168–1173. https://doi.org/10.1016/J.PROCS.2020.03.049

Mohammed, A., & Kora, R. (2023). A comprehensive review on ensemble deep learning: Opportunities and challenges. Journal of King Saud University - Computer and Information Sciences, 35(2), 757–774. https://doi.org/10.1016/J.JKSUCI.2023.01.014

Mohsin, M. T., & Nasim, N. Bin. (2025). Explaining the Unexplainable: A Systematic Review of Explainable AI in Finance. https://arxiv.org/pdf/2503.05966

Mondal, M., Khayati, M., Sandlin, H. Â., & Cudré-Mauroux, P. (2025). A survey of multimodal event detection based on data fusion. VLDB Journal, 34(1), 1–25. https://doi.org/10.1007/S00778-024-00878-5/TABLES/4

Moodi, F., Jahangard-Rafsanjani, A., & Zarifzadeh, S. (2023). Feature selection and regression methods for stock price prediction using technical indicators. https://arxiv.org/pdf/2310.09903

Mozaffari, L., & Zhang, J. (2024, May). Predictive Modeling of Stock Prices Using Transformer Model. 2024 9th International Conference on Machine Learning Technologies (ICMLT), . https://doi.org/10.17148/iarjset.2015.2305

Muallim, T. (2018). Anomaly Detection using Variational Autoencoders. Github. https://github.com/tarekmuallim/Anomaly-Detection-using-Variational-Autoencoders/blob/master/Anomaly_Detection_using_VAE.ipynb

Neloy, A. A., & Turgeon, M. (2024). A comprehensive study of auto-encoders for anomaly detection: Efficiency and trade-offs. Machine Learning with Applications, 17, 100572. https://doi.org/10.1016/J.MLWA.2024.100572

Nettey, D. A., & Ansong, E. D. (2025). Anomaly Detection with Variational Autoencoders. https://doi.org/10.21203/RS.3.RS-5809487/V1

Nithish Kumar, P., Umeorah, N., & Alochukwu, ; Alex. (2024). Dynamic graph neural networks for enhanced volatility prediction in financial markets. https://arxiv.org/pdf/2410.16858

Noguer I Alonso, M., & Pereira Franklin, R. (2024). Large Language Models for Financial Time Series Forecasting. https://doi.org/10.2139/SSRN.4988022

Ozturk, O. (2020). Stock Price prediction by simple RNN and LSTM (Source code). Kaggle Notebooks. https://www.kaggle.com/code/ozkanozturk/stock-price-prediction-by-simple-rnn-and-lstm

Pabuccu, H., & Barbu, A. (2024). Feature selection with annealing for forecasting financial time series. Financial Innovation, 10(1), 1–26. https://doi.org/10.1186/S40854-024-00617-3/METRICS

Paras Varshney. (n.d.). RNN-based-Stock-Price-Prediction-using-LSTM [Source code] ([Source code]). GitHub. Retrieved January 1, 2025, from https://github.com/blurred-machine/RNN-based-Stock-Price-Prediction-using-LSTM

Pistoia, M., Ahmad, S. F., Ajagekar, A., Buts, A., Chakrabarti, S., Herman, D., Hu, S., Jena, A., Minssen, P., Niroula, P., Rattew, A., Sun, Y., & Yalovetzky, R. (2021). Quantum Machine Learning for Finance. IEEE/ACM International Conference on Computer-Aided Design, Digest of Technical Papers, ICCAD, 2021-November. https://doi.org/10.1109/ICCAD51958.2021.9643469

Sadon, A. N., Ismail, S., Khamis, A., & Tariq, M. U. (2024). Heteroscedasticity effects as component to future stock market predictions using RNN-based models. PLOS ONE, 19(5), e0297641. https://doi.org/10.1371/JOURNAL.PONE.0297641

Seseri, R. (2023a, March 13). AI Atlas #15: Ensemble methods in deep learning. Glasswing Ventures. https://glasswing.vc/blog/ai-atlas-15-ensemble-methods-in-deep-learning/

Seseri, R. (2023b, September 13). Ensembles in Machine Learning, Applications in Finance. PyFi. https://pyfi.com/blogs/articles/ensembles-in-machine-learning-applications-in-finance

Shali, Y., Brahma, B., Wadhvani, R., & Gyanchandani, M. (2021). Attention LSTM for Time Series Forecasting of Financial Time Series Data. 74–84. https://doi.org/10.1007/978-3-030-76736-5_8

Shubham Ghelani. (n.d.). Financial Time Series Forecasting using Deep Learning (Source code). GitHub. Retrieved January 1, 2025, from https://github.com/ShubhamG2311/Financial-Time-Series-Forecasting

Song, G., Zhao, T., Ma, X., Lin, P., & Cui, C. (2025). Reinforcement learning-based portfolio optimization with deterministic state transition. Information Sciences, 690, 121538. https://doi.org/10.1016/J.INS.2024.121538

Takahashi, S., Chen, Y., & Tanaka-Ishii, K. (2019a). Modeling financial time-series with generative adversarial networks. Physica A: Statistical Mechanics and Its Applications, 527, 121261. https://doi.org/10.1016/J.PHYSA.2019.121261

Tam, A. (2021, November 20). Using CNN for financial time series prediction. Machine Learning Mastery. https://machinelearningmastery.com/using-cnn-for-financial-time-series-prediction/

Tanaka, H., Fujimoto, A., & Sakamoto, R. (2025). FINANCIAL TIME SERIES FORECASTING USING ADAPTIVE RISK METRICS AND TRANSFORMER MODELS. Journal of Trends in Financial and Economics, 2(1). https://doi.org/10.61784/JTFE3034

Tanikonda, A., Peddinti, S. . R., & Katragadda, S. . R. (2025). Deep Learning for Anomaly Detection in E-commerce and Financial Transactions: Enhancing Fraud Prevention and Cybersecurity. Journal of Information Systems Engineering and Management, 10(30). https://jisem-journal.com/index.php/journal/article/view/4776/2232

Tavakoli, M., Chandra, R., Tian, F., & Bravo, C. (2024). Multi-Modal Deep Learning for Credit Rating Prediction Using Text and Numerical Data Streams. Applied Soft Computing, 171. https://doi.org/10.1016/j.asoc.2025.112771

Tu, T. (2025). Bridging Short- and Long-Term Dependencies: A CNN-Transformer Hybrid for Financial Time Series Forecasting. https://arxiv.org/pdf/2504.19309

Uddin, A., Tao, X., & Yu, D. (2023). Attention based dynamic graph neural network for asset pricing. Global Finance Journal, 58, 100900. https://doi.org/10.1016/J.GFJ.2023.100900

Ung, D., Chawla, K., & Miklasewski, J. (2025). Decoding Market Regimes Machine Learning Insights into US Asset Performance Over The Last 30 Years. https://www.ssga.com/us/en/intermediary/insights/machine-learning-insights

Varadharajan, V., Smith, N., Kalla, D., Kumar, G. R., Samaah, F., & Polimetla, K. (2024). Stock Closing Price and Trend Prediction with LSTM-RNN. Journal of Artificial Intelligence and Big Data, 4(1), 1–13. https://doi.org/10.31586/JAIBD.2024.877

Vuletić, M., Cucuringu, M., & Prenzel, F. (2023, July 13). Fin-GAN: Forecasting and Classifying Financial Time Series via Generative Adversarial Networks. Man Group. Retrieved January 1, 2025. https://www.man.com/insights/forecasting-via-generative-adversarial-networks

Wang, H., Xie, Z., Chiu, D. K. W., & Ho, K. K. W. (2025). Multimodal market information fusion for stock price trend prediction in the pharmaceutical sector: Multimodal market information fusion for stock price trend prediction in the pharmaceutical sector: H. Wang et al. Applied Intelligence, 55(1), 1–27. https://doi.org/10.1007/S10489-024-05894-0/FIGURES/2

Wang, J., Yang, S., Zhao, H., & Yang, Y. (2023). Social media popularity prediction with multimodal hierarchical fusion model. Computer Speech & Language, 80, 101490. https://doi.org/10.1016/J.CSL.2023.101490

Wang, M., & Hirsa, A. (2024). Advanced Generative Adversarial Networks for Financial Time Series: Incorporating Stochastic Masking, Multi-Frequency Data, Hidden Markov Models and Vector Quantization. In SSRN (SSRN Scholarly Paper No. 490954). https://papers.ssrn.com/abstract=4909544

Wen, X., & Li, W. (2023). Time Series Prediction Based on LSTM-Attention-LSTM Model. IEEE Access, 11, 48322–48331. https://doi.org/10.1109/ACCESS.2023.3276628

Wilson, C.-A. (2025). Explainable AI in Finance: Addressing the Needs of Diverse Stakeholders. https://doi.org/10.56227/25.1.25

Wu, H. (2024). Revisiting Attention for Multivariate Time Series Forecasting. Proceedings of the AAAI Conference on Artificial Intelligence, 39(20), 21528–21535. https://doi.org/10.1609/aaai.v39i20.35455

Wu, R. (2024). Financial Fraud Detection System Combining Generative Adversarial Networks and Deep Learning. 2024 International Conference on Industrial IoT, Big Data and Supply Chain (IIoTBDSC), 105–110. https://doi.org/10.1109/IIOTBDSC64371.2024.00029

Xia, M. (2019). Fraud Detection with Variational Autoencoder. Kaggle. https://www.kaggle.com/code/hone5com/fraud-detection-with-variational-autoencoder

Xiao, M., Chen, Z., Qian, L., Jiang, Z., He, Y., Xu, Y., Jiang, Y., Li, D., Weng, R.-L., Huang, J., Peng, M., Ananiadou, S., Nie, J.-Y., & Xie, Q. (2025). Retrieval-augmented Large Language Models for Financial Time Series Forecasting. https://arxiv.org/pdf/2502.05878

Xiao, Y., Yin, H., Zhang, Y., Qi, H., Zhang, Y., & Liu, Z. (2021). A dual-stage attention-based Conv-LSTM network for spatio-temporal correlation and multivariate time series prediction. International Journal of Intelligent Systems, 36(5), 2036–2057. https://doi.org/10.1002/INT.22370

Xu, K., Wu, Y., Xia, H., Sang, N., & Wang, B. (2022). Graph Neural Networks in Financial Markets: Modeling Volatility and Assessing Value-at-Risk. Journal of Computer Technology and Software, 1(2), 2022. https://ashpress.org/index.php/jcts/article/view/85

Xu, Z., Liechty, J., Benthall, S., Skar-Gislinge, N., & McComb, C. (2024). GARCH-Informed Neural Networks for Volatility Prediction in Financial Markets. https://arxiv.org/pdf/2410.00288v1

Yao, J. (2019). Illuminating Economic Growth: Satellite images of the earth at night reveal the pace of economic growth andmuch more. Finance & Development, 56(003). https://doi.org/10.5089/9781498316040.022.A015

Yeo, W. J., Van Der Heever, W., Mao, R., Cambria, E., Satapathy, R., & Mengaldo, G. (2025). A comprehensive review on financial explainable AI. Artificial Intelligence Review, 58(6), 1–49. https://doi.org/10.1007/S10462-024-11077-7/FIGURES/7

Zeng, Z., Kaur, R., Siddagangappa, S., Rahimi, S., Balch, T., Veloso, M., & Morgan, J. P. (2023). Financial Time Series Forecasting using CNN and Transformer. https://arxiv.org/pdf/2304.04912

Zhang, W., Chen, Z., Miao, J., & Liu, X. (2022). Research on Graph Neural Network in Stock Market. Procedia Computer Science, 214(C), 786–792. https://doi.org/10.1016/J.PROCS.2022.11.242

Zhang, W., Zhao, L., Xia, H., Sun, S., Sun, J., Qin, M., Li, X., Zhao, Y., Zhao, Y., Cai, X., Zheng, L., Wang, X., & An, B. (2024). A Multimodal Foundation Agent for Financial Trading: Tool-Augmented, Diversified, and Generalist. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 4314–4325. https://doi.org/10.1145/3637528.3671801/SUPPL_FILE/RTP0724-2MIN-PROMO.MP4

Zhang, X., Liang, X., Zhiyuli, A., Zhang, S., Xu, R., & Wu, B. (2019). AT-LSTM: An Attention-based LSTM Model for Financial Time Series Prediction. IOP Conference Series: Materials Science and Engineering, 569(5). https://doi.org/10.1088/1757-899X/569/5/052037

Zhang, X., Xu, Z., Liu, Y., Sun, M., Zhou, T., & Sun, W. (2024). Robust Graph Neural Networks for Stability Analysis in Dynamic Networks. https://arxiv.org/pdf/2411.11848

Zhou, K., Wang, W. Y., Hu, T., & Wu, C. H. (2020). Comparison of Time Series Forecasting Based on Statistical ARIMA Model and LSTM with Attention Mechanism. Journal of Physics: Conference Series, 1631(1). https://doi.org/10.1088/1742-6596/1631/1/012141

.