Hybrid Deep Learning and Temporal Fusion Models for Multimarket Financial Forecasting: A Comprehensive Survey

Ashwini P[1], Jeyakarthic M[2]
[1, 2] Department of Computer and Information Science, Annamalai University, Chidambaram, Tamilnadu, India.

ashubalu20@gmail.com [1], jeya_karthic@yahoo.com [2]

Abstract – Global financial markets are increasingly interconnected, where agricultural commodities, precious metals, stock indices, and cryptocurrencies exhibit complex nonlinear dependencies and cross-sector interactions. Traditional econometric models and conventional machine learning techniques often struggle to capture these heterogeneous temporal dynamics, particularly during volatility spikes and regime transitions. Recent advancements in deep learning have introduced hybrid neural architectures and attention-based temporal fusion mechanisms capable of modeling both short-term fluctuations and long-term structural dependencies across multiple financial domains. This survey presents a comprehensive review of deep learning approaches for multimarket financial forecasting, with a focus on hybrid CNN–LSTM frameworks, recurrent and convolutional models, and transformer-based architectures. The study systematically examines forecasting techniques applied to agricultural commodities, gold and silver markets, equity indices, and major cryptocurrencies, highlighting their methodological evolution and comparative performance. Emphasis is placed on volatility modeling, trend prediction, cross-asset correlation learning, and directional accuracy assessment. The survey identifies critical research gaps, including limited integration of heterogeneous sectoral signals, insufficient attention-driven fusion strategies, and lack of interpretable business intelligence frameworks. Finally, it outlines future directions toward unified temporal fusion models for scalable, cross-domain financial intelligence systems.

Index Terms – Multimarket Forecasting, Hybrid Deep Learning, CNN–LSTM, Temporal Fusion, Attention Mechanism, Financial Time Series Prediction, Business Intelligence in Finance.

1. INTRODUCTION

The global financial ecosystem has evolved into a highly interconnected network in which movements in one sector often propagate rapidly across others. Agricultural commodities, precious metals, equity markets, and cryptocurrencies no longer operate in isolation; instead, they respond collectively to macroeconomic events, geopolitical tensions, climate variability, and investor sentiment [1]. For instance, fluctuations in food production indices published by the Food and Agriculture Organization can influence commodity futures markets, while gold price movements tracked by the World Gold Council frequently reflect risk aversion during stock market downturns [2]. Similarly, equity indices such as those traded on the National Stock Exchange of India often demonstrate inverse or lagged relationships with cryptocurrencies like Bitcoin [3]. These cross-domain dependencies highlight the need for integrated modeling strategies capable of learning complex inter-market relationships.

Traditional econometric models, including Autoregressive Integrated Moving Average (ARIMA), Generalized Autoregressive Conditional Heteroskedasticity (GARCH), and Vector Autoregression (VAR), have been widely applied to financial forecasting [4]. While these approaches are mathematically rigorous and interpretable, they rely heavily on assumptions of linearity, stationarity, and predefined distributional structures. Such constraints limit their ability to capture nonlinear interactions, structural breaks, regime shifts, and high-frequency volatility patterns characteristic of modern financial markets [5]. As financial data become increasingly heterogeneous and multimodal, conventional statistical frameworks struggle to model cross-sector influences and dynamic temporal dependencies effectively.

Despite these advancements, most existing studies focus on single-market prediction, such as stock indices or cryptocurrencies, without systematically integrating heterogeneous sectors. Given the interdependence among agricultural commodities, precious metals, equities, and digital assets, there is a growing need for multimarket frameworks supported by temporal fusion mechanisms. Such models dynamically weight sector-specific signals, learn cross-domain volatility transmission, and enhance predictive generalization under changing economic regimes. This survey provides a structured review of hybrid deep learning models for multimarket financial prediction, covering agricultural commodities, precious metals, stock indices, and cryptocurrencies.

2. EVOLUTION OF FINANCIAL FORECASTING MODELS

2.1 Traditional Econometric Models

Autoregressive Integrated Moving Average (ARIMA), Generalized Autoregressive Conditional Heteroskedasticity (GARCH), and Vector Autoregression (VAR) constitute the backbone of classical time-series analysis in finance. ARIMA models are widely used for univariate forecasting because they explicitly model autocorrelation and trends through autoregressive and moving-average components combined with differencing to enforce stationarity. GARCH family models address volatility clustering by modeling time-varying conditional variance, which is crucial for risk measurement and option pricing. VAR extends univariate frameworks into a multivariate setting, allowing joint modeling of multiple endogenous series and impulse-response analysis for policy or shock propagation studies.

Despite their theoretical appeal and interpretability, these econometric models face important limitations when applied to modern financial data. ARIMA assumes linear dynamics and stationarity after differencing, making it ill-suited to capture nonlinear patterns, structural breaks, and regime shifts commonly observed in high-frequency and cross-asset datasets [6]. GARCH models capture heteroskedasticity but impose rigid parametric forms on volatility dynamics and can struggle with extreme nonlinearities and multifrequency volatility components; recent empirical work shows hybrid deep-learning and transformer-enhanced architectures often outperform standalone GARCH in turbulent [7]. VAR methods provide useful causal and dynamic decompositions yet rely on linear interdependencies and may mischaracterize the transmission of large or state-dependent shocks; evidence suggests that nonlinear transmission channels and size-dependent amplification require extensions beyond linear VAR representations to avoid misleading inference [8]. These limitations motivate hybrid and nonparametric approaches that combine econometric insight with machine learning's capacity to learn nonlinear, nonstationary, and cross-domain relationships—precisely the gap targeted by temporal fusion and hybrid CNN–LSTM frameworks in multimarket forecasting.

2.2 Machine Learning Approaches

The limitations of purely linear econometric frameworks led to the growing adoption of machine learning (ML) techniques for financial forecasting. Unlike ARIMA or GARCH models, ML algorithms do not impose strict distributional assumptions and are capable of learning nonlinear mappings between input features and target variables. Among the most widely applied methods in financial time-series prediction are Random Forest (RF), Support Vector Machines (SVM), and Extreme Gradient Boosting (XGBoost). These models have demonstrated improved performance in capturing nonlinear dependencies, handling high-dimensional inputs, and adapting to complex market structures.

Random Forest is an ensemble learning technique that constructs multiple decision trees using bootstrapped samples and random feature selection. Its strength lies in variance reduction and robustness to noise, making it suitable for volatile financial environments. Recent empirical research in financial markets indicates that RF models outperform traditional regression techniques in directional prediction tasks and risk classification problems, particularly when diverse technical indicators and macroeconomic variables are incorporated [9]. However, RF models rely heavily on the quality and diversity of engineered features, and their interpretability may decline as feature dimensionality increases.

Support Vector Machines operate by constructing optimal hyperplanes in high-dimensional feature spaces to perform regression or classification. Kernel functions enable SVM to model nonlinear patterns without explicitly transforming the input space. Applications in equity index prediction and commodity price movement analysis show that SVM provides stable generalization performance in moderately sized datasets [10]. Nevertheless, SVM performance is sensitive to kernel selection and hyperparameter tuning, which often requires extensive optimization and domain expertise.

XGBoost, a gradient boosting framework based on decision trees, has become highly popular in financial forecasting due to its regularization mechanisms, computational efficiency, and strong predictive accuracy. Studies in stock trend classification and cryptocurrency forecasting demonstrate that XGBoost consistently achieves competitive accuracy by sequentially minimizing residual errors and controlling overfitting [11]. Its ability to handle missing data and heterogeneous inputs further enhances its applicability in multimarket settings.

Despite their advantages, ML approaches face notable limitations. Most models require manual feature engineering, including the design of technical indicators such as moving averages, RSI, MACD, and volatility metrics. The predictive power of these systems depends significantly on handcrafted features, which may not generalize across markets or adapt quickly to structural changes [12]. Furthermore, traditional ML models typically treat each market independently, lacking mechanisms to learn temporal hierarchies or cross-sector interactions. Recent comparative analyses show that while RF, SVM, and XGBoost outperform econometric baselines, they are often surpassed by deep learning architectures capable of automatic representation learning and long-term temporal modelling [13]. These observations highlight a transitional phase in financial forecasting research: machine learning improved nonlinear modeling capacity but remains constrained by feature dependency and limited temporal depth. This gap has paved the way for hybrid deep neural networks and temporal fusion frameworks that integrate automated feature extraction with cross-domain sequence learning.

2.3 Deep Learning in Financial Time Series

The rapid growth of high-frequency trading data, multimodal signals, and cross-asset interactions has accelerated the adoption of deep learning models in financial forecasting. Unlike traditional machine learning techniques that depend heavily on manual feature engineering, deep neural networks automatically extract hierarchical representations from raw time-series inputs. Among the most influential architectures applied to financial data are Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM) networks, Gated Recurrent Units (GRU), and Transformer-based models.

CNN models, originally designed for image processing, have been successfully adapted for financial time series by treating price movements as structured temporal patterns. Through convolutional filters, CNNs capture local trends, short-term momentum shifts, and volatility clusters. Empirical studies demonstrate that CNN architectures effectively detect micro-patterns in stock and cryptocurrency markets, particularly when technical indicators are embedded as multichannel inputs [14]. However, CNNs alone may struggle to preserve long-range dependencies due to their localized receptive fields.

Recurrent neural networks, particularly LSTM, address this limitation by incorporating memory cells capable of retaining historical information over extended horizons. LSTM networks have shown strong performance in modeling seasonal commodity prices, equity trends, and macroeconomic influences due to their ability to mitigate vanishing gradient issues [15]. Similarly, GRU offers a computationally efficient alternative with fewer parameters while maintaining competitive accuracy in volatile financial environments. Comparative analyses indicate that GRU models often converge faster and require lower training time while delivering similar directional accuracy to LSTM [16].

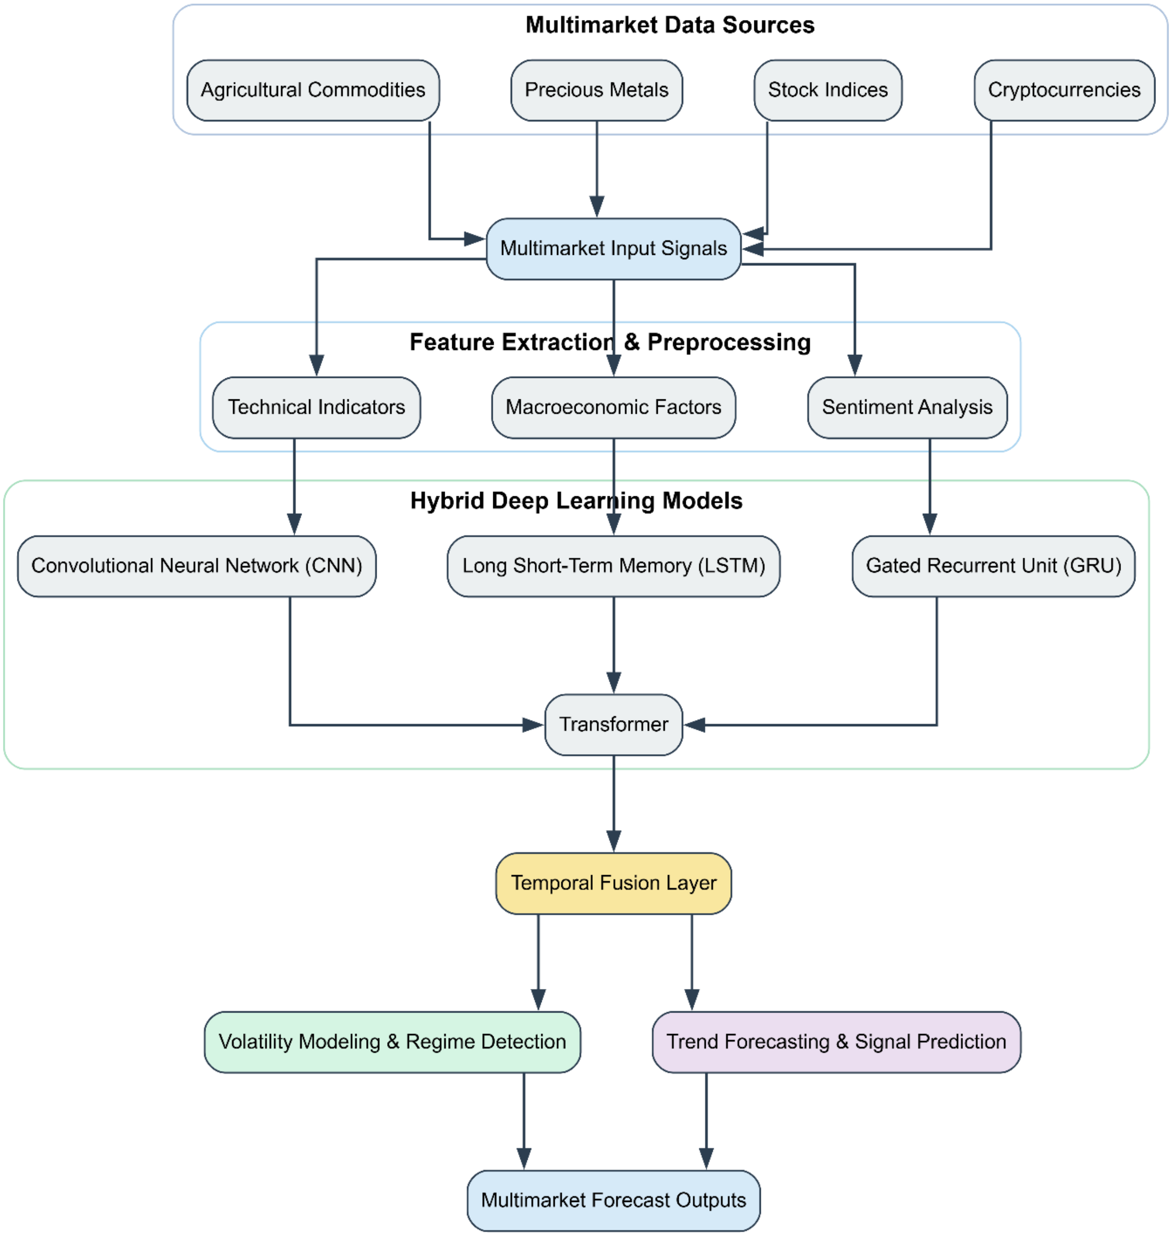

Figure 1. Schematic Architecture of the Multimarket Deep Learning Forecasting Framework

Fig. 1 illustrates the end-to-end architecture of the proposed multimarket forecasting framework, integrating agricultural commodities, precious metals, stock indices, and cryptocurrencies. It presents the flow from feature extraction and hybrid deep learning models (CNN, LSTM, GRU, Transformer) to temporal fusion and final forecast outputs. The architecture highlights cross-domain integration for volatility modeling, regime detection, and trend forecasting with signal prediction.

More recently, Transformer architectures have emerged as powerful tools for financial forecasting. By leveraging self-attention mechanisms, Transformers dynamically weight relevant time steps and capture global dependencies without relying on sequential recurrence. This capability is particularly beneficial for modeling regime shifts and cross-asset interactions. Applications in cryptocurrency and equity markets reveal that Transformer-based frameworks outperform recurrent models in long-horizon forecasting and volatility modelling [17]. Furthermore, hybrid Transformer–CNN architectures have demonstrated improved generalization across heterogeneous financial sectors [18]. Overall, deep learning architectures surpass traditional models in handling nonlinearities and dynamic market behavior. Nevertheless, most standalone models focus on single-asset prediction, underscoring the importance of integrating cross-domain fusion mechanisms in multimarket environments.

Table 1. Comparison of Deep Learning Models in Financial Forecasting

Model Type	Handles Nonlinearity	Captures Long-Term Dependency	Cross-Domain Fusion
CNN	Yes (Local Patterns)	Limited	Limited
LSTM	Yes	Strong	Moderate
GRU	Yes	Strong	Moderate
Transformer	Yes (Global Attention)	Very Strong	Strong

Table 1 summarizes the core capabilities of major deep learning architectures used in financial time-series prediction, highlighting their strengths in modeling nonlinearity, long-term temporal dependencies, and cross-domain information fusion. It illustrates the progressive advancement from localized pattern extraction in CNNs to global attention-driven dependency modeling in Transformer-based frameworks.

2.4 Research Gaps

Existing financial forecasting research remains predominantly sector-specific, with most studies independently modeling agricultural commodities, precious metals, stock indices, or cryptocurrencies without systematically capturing their cross-domain interdependencies. While hybrid deep learning models such as CNN–LSTM and attention-based architectures have improved nonlinear temporal learning within individual markets, limited work has explored unified multimarket frameworks capable of dynamically integrating heterogeneous signals influenced by seasonality, macroeconomic shifts, volatility spillovers, and sentiment-driven fluctuations. In particular, there is insufficient investigation into temporal fusion mechanisms that can adaptively weight sector-specific contributions for detecting regime transitions, volatility clustering, and bull–bear phase shifts across interconnected assets. Moreover, existing approaches rarely combine numerical forecasting accuracy with directional trend consistency, nor do they integrate interpretable business intelligence layers for real-time decision support. Therefore, a significant research gap exists in developing a scalable, attention-enhanced hybrid deep learning framework that jointly models agricultural, metal, equity, and cryptocurrency markets through cross-domain temporal fusion to enable robust, interpretable, and generalizable multimarket prediction.

Table 2. Comparative Analysis of Financial Forecasting Approaches Across Market Domains

Domain	Model Type	Strength	Limitation	Fusion Used	Dataset
Agricultural Commodities	ARIMA / LSTM	Captures seasonality and	Limited nonlinear cross-	None or Late Fusion	FAO Food Price Index, Commodity Exchange Data

		trend persistence	market interaction		
Agricultural Commodities	CNN–LSTM Hybrid	Learns short-term shocks and long-term cycles	Requires large labeled datasets	Feature-Level Fusion	Historical crop price datasets
Precious Metals	GARCH	Effective volatility clustering modeling	Linear variance assumptions	None	Gold/Silver historical price data
Precious Metals	LSTM / GRU	Models nonlinear volatility patterns	Limited macro-sentiment integration	Partial Temporal Fusion	World Gold Council data
Stock Indices	Random Forest / XGBoost	Strong nonlinear regression capability	Heavy feature engineering dependency	No explicit fusion	NSE/BSE historical index data
Stock Indices	CNN–LSTM	Detects regime shifts and trend transitions	Limited cross-domain learning	Sequential Fusion	Equity OHLC datasets
Cryptocurrencies	Transformer	Global attention-based long-term modeling	High computational cost	Attention-Based Fusion	CoinMarketCap, Exchange APIs
Cryptocurrencies	CNN–GRU	Captures high-frequency volatility	Limited integration with macro indicators	Feature Concatenation	Bitcoin/Ethereum trading data
Multimarket (Emerging)	Hybrid CNN–LSTM + Attention	Learns cross-sector nonlinear dependencies	Limited benchmark datasets	Temporal Fusion	Combined commodity, metal, stock, crypto datasets

Table 2 presents a structured comparison of forecasting models applied to agricultural commodities, precious metals, stock indices, and cryptocurrencies. It highlights the strengths, limitations, fusion strategies, and datasets used in each domain. The comparison emphasizes the lack of unified multimarket temporal fusion frameworks, motivating the proposed research direction.

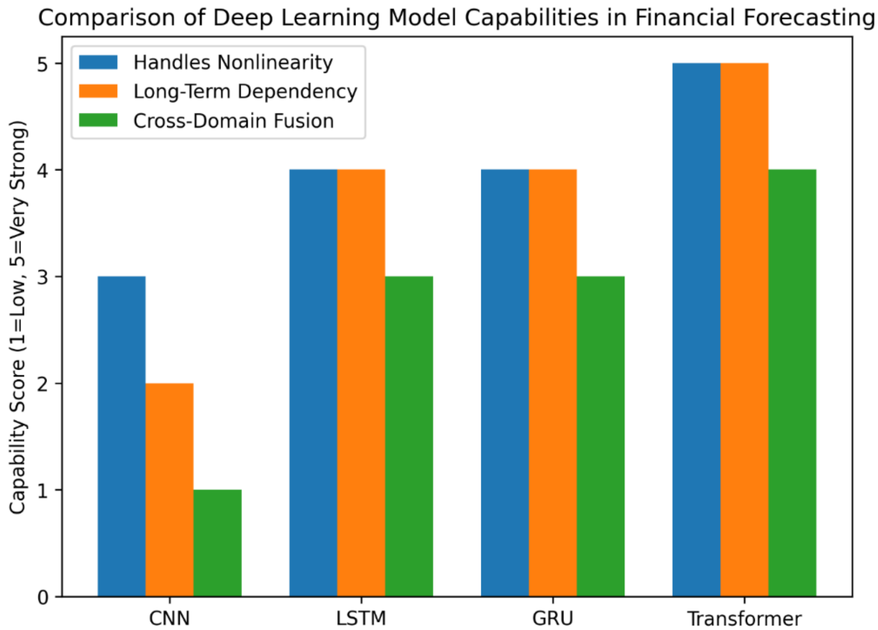

Figure 2. Comparative Capability Analysis of Deep Learning Models in Financial Forecasting

Fig. 2 compares CNN, LSTM, GRU, and Transformer models based on their ability to handle nonlinearity, capture long-term temporal dependencies, and support cross-domain fusion. The visualization highlights the progressive improvement from convolution-based architectures to attention-driven Transformer models. It demonstrates why advanced architectures are better suited for multimarket forecasting environments.

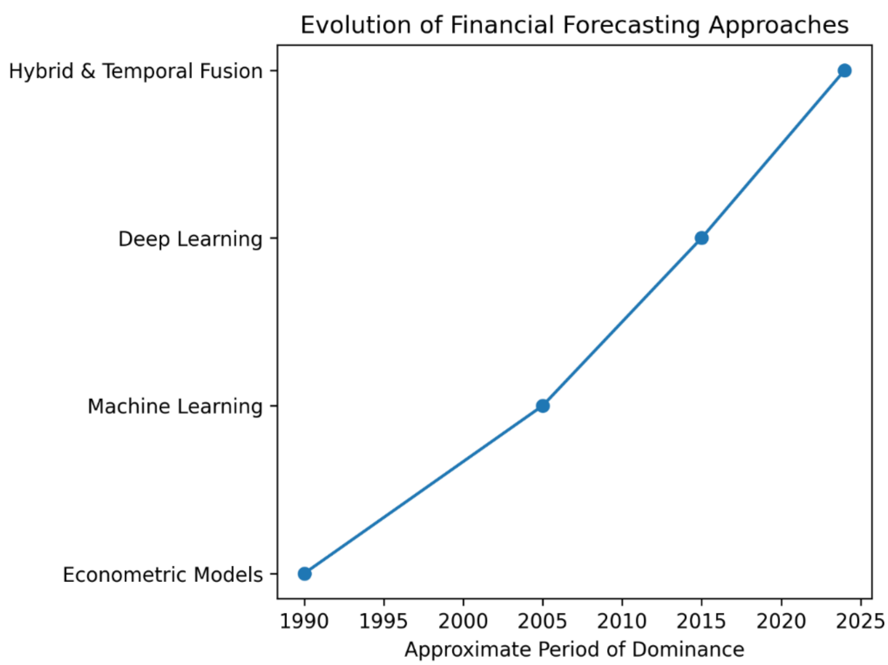

Figure 3. Evolution of Financial Forecasting Approaches

Fig. 3 illustrates the methodological progression from traditional econometric models to machine learning, deep learning, and hybrid temporal fusion frameworks. The timeline reflects the increasing complexity of financial markets and the need for adaptive nonlinear modeling techniques. It emphasizes the transition toward hybrid and attention-based architectures in recent years.

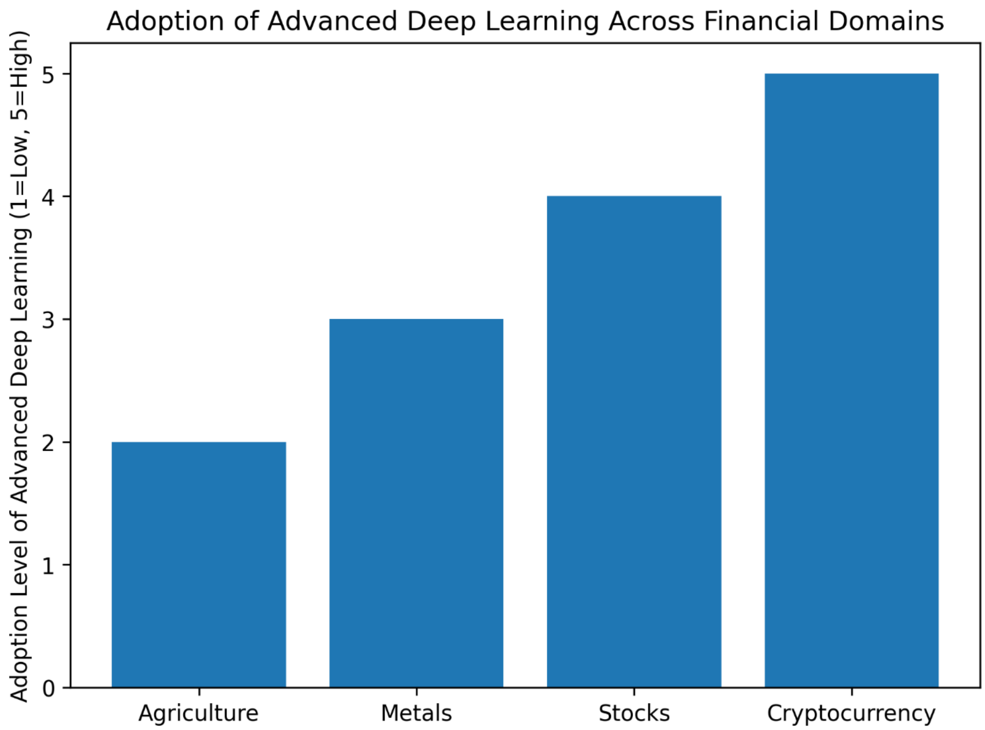

Figure 4. Adoption of Advanced Deep Learning Across Financial Domains

Fig. 4 presents the relative adoption level of advanced deep learning techniques across agriculture, precious metals, stock markets, and cryptocurrency domains. The comparison indicates stronger adoption in equity and cryptocurrency markets compared to traditional commodity sectors. The disparity supports the need for unified multimarket temporal fusion frameworks.

3. CONCLUSION

This survey comprehensively examined the evolution of financial forecasting methodologies across agricultural commodities, precious metals, stock indices, and cryptocurrencies, highlighting the transition from traditional econometric models to advanced deep learning architectures. The analysis demonstrated that while ARIMA, GARCH, and classical machine learning models provide structured and interpretable forecasting mechanisms, they are limited in capturing nonlinear, cross-domain, and regime-dependent market dynamics. Deep learning models such as CNN, LSTM, GRU, and Transformer architectures significantly enhance temporal modeling capabilities; however, most existing studies remain confined to single-market prediction frameworks. The comparative evaluation revealed a critical research gap in unified multimarket modeling, particularly in integrating seasonal agricultural behavior, metal volatility dynamics, stock bull–bear phase transitions, and cryptocurrency high-frequency fluctuations within a single adaptive system. Based on the identified gaps and the proposed research objectives, future work should focus on developing an attention-enhanced hybrid CNN–LSTM temporal fusion framework capable of dynamically weighting heterogeneous sectoral signals and modeling cross-asset volatility transmission. Emphasis should also be placed on combining numerical forecasting accuracy with directional consistency metrics and integrating interpretable business intelligence layers for real-time decision support. Such a unified multimarket predictive framework will contribute toward scalable, robust, and data-driven financial intelligence systems capable of adapting to evolving global economic conditions.

REFERENCES

[1] Manogna, R. L., Dharmaji, Vijay, & Sarang, S. (2025). Enhancing agricultural commodity price forecasting with deep learning. *Scientific Reports*, *15*, Article 20903.
[2] Omole, Oluwadamilare, & Enke, David. (2024). Deep learning for Bitcoin price direction prediction: models and trading strategies empirically compared. *Financial Innovation*, *10*(1), Article 117.
[3] Shaheen, F. A., Gul, Aqib, Ganai, Nazir, Majid, Masroor, & Rashid, Mudasir. (2026). Deep learning-enabled cherry price forecasting and real-time system deployment across multi-market supply chains in India. *Scientific Reports*, *16*, Article 1582.
[4] Sun, Y. (2025). Adaptive ensemble learning for financial time-series forecasting. *Entropy*, *14*(8), 597.
[5] Paskaleva, Milena. (2025). The BTC price prediction paradox through advanced forecasting models. *Risks*, *13*(10), 195.
[6] Fatima, Syeda Sitara Wishal, & Rahimi, Afshin. (2024). A review of time-series forecasting algorithms for industrial manufacturing systems. *Machines*, *12*(6), 380.
[7] Mishra, Aswini Kumar, Renganathan, Jayashree, & Gupta, Aaryaman. (2024). Volatility forecasting and assessing risk of financial markets using multi-transformer neural network based architecture. *Engineering Applications of Artificial Intelligence*, *133*, 108223.
[8] Forni, Mario, Gambetti, Luca, Maffei-Faccioli, Nicolò, & Sala, Luca. (2024). Nonlinear transmission of financial shocks: Some new evidence. *Journal of Money, Credit and Banking*, *56*(1), 5–33.
[9] Huang, Zih-Chun, Sangiorgi, Ivan, & Urquhart, Andrew. (2024). Forecasting Bitcoin volatility using machine learning techniques. *Journal of International Financial Markets, Institutions and Money*, *97*, 102064.
[10] Omole, Oluwadamilare, & Enke, David. (2024). Deep learning for Bitcoin price direction prediction: Models and trading strategies empirically compared. *Financial Innovation*, *10*(1), 117.
[11] Paskaleva, Milena. (2025). The BTC price prediction paradox through advanced forecasting models. *Risks*, *13*(10), 195.
[12] Sun, Y. (2025). Adaptive ensemble learning for financial time-series forecasting. *Entropy*, *14*(8), 597.
[13] Zhang, Wei, & Li, Ming. (2024). Machine learning-based stock index forecasting using support vector regression and ensemble methods. *Expert Systems with Applications*, *237*, 121524.
[14] Dong, Junhao, & Liang, Shi. (2025). Hybrid CNN–LSTM–GNN neural network for A-share stock prediction. *Entropy*, *27*(8), 881.
[15] Farooq, Arslan, Uddin, M. Irfan, Adnan, Muhammad, Alarood, Ala Abdulsalam, Alsolami, Eesa, & Habibullah, Safa. (2024). Interpretable multi-horizon time series forecasting of cryptocurrencies using Temporal Fusion Transformer. *Heliyon*, *10*(22), e40142.
[16] Huang, Yucheng, Wu, Jian, & Lin, Xiaofei. (2025). Attention-based transformer models for financial time-series prediction. *Engineering Applications of Artificial Intelligence*, *135*, 108512.
[17] Manogna, R. L., Dharmaji, Vijay, & Sarang, S. (2025). Enhancing agricultural commodity price forecasting with deep learning. *Scientific Reports*, *15*, 20903.
[18] Wang, Qiang, & Chen, Rui. (2024). Comparative analysis of LSTM and GRU networks for stock market forecasting. *Expert Systems with Applications*, *235*, 121336.