International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 4, July-Aug-2026, ISSN (Online): 2395-566X

Artificial Intelligence for Predicting Currency Fluctuation and Investment Risk: A Secondary-Data Synthesis of LSTM, Random Forest, and Hybrid GARCH Approaches

Dr. Ranjeet Kumar Ambast[1], Aditya Vikram[2]
Assistant Professor, CMS
Gitarattan International Business School
Guru Gobind Singh Indraprastha University[1]

Assistant Professor, Rajdhani College
University of Delh[2]

Abstract — The foreign exchange (Forex) market is characterised by high liquidity, pronounced volatility, and non-linear dynamics. This makes accurate prediction of currency fluctuation and investment risk exceptionally difficult for traditional econometric models (Bollerslev, 1986; Cont, 2001). This paper synthesises secondary evidence on the comparative predictive performance of two Artificial Intelligence (AI) approaches Long Short-Term Memory (LSTM) deep learning networks (Hochreiter & Schmidhuber, 1997) and Random Forest (RF) ensemble learning (Breiman, 2001) against the traditional Generalised Autoregressive Conditional Heteroskedasticity (GARCH) family of models (Bollerslev, 1986), drawing on published studies that use historical USD, EUR, GBP, JPY, BRL, and ZAR exchange-rate data spanning approximately 2012–2025. The review evaluates model performance across stable and volatile market regimes for directional accuracy through Value-at-Risk (VaR) estimation, and risk-adjusted investment signals. Findings from the reviewed literature indicate that AI-based LSTM models achieve superior performance for short-horizon volatility forecasts. It's particularly in capturing sudden shifts in implied volatility (Kraus & Feuerriegel, 2024). Random Forest models tend to deliver the highest directional accuracy and the lowest point-prediction error across several currency pairs and cryptocurrency markets (Milionis & Konstantinou, 2024; Ndlovu, 2025). Hybrid GARCH-LSTM architectures further improve predictive accuracy, with an APARCH-LSTM specification reported to achieve a coefficient of determination (R^2) of 95.53% for USD/BRL volatility forecasting (Hottz, 2025), while hybrid models improve Value-at-Risk estimation accuracy by up to 10% during periods of elevated volatility relative to standalone GARCH or LSTM specifications (Nsengiyumva et al., 2025). The synthesis further shows that risk-adjusted performance measured through the Sharpe ratio, Calmar ratio, and maximum drawdown favours hybrid deep-learning architectures for multi-asset portfolios and tree-based ensembles such as XGBoost for equity-index applications (Saly-Kaufmann et al., 2026; Singh & Praveen, 2025). The paper concludes that no single AI technique dominates across all currency pairs, forecast horizons, and objectives, and that model choice should be aligned with the specific trading, risk-management, or regulatory objective at hand.

Keywords— Artificial Intelligence; currency fluctuation; investment risk; LSTM; Random Forest; GARCH; foreign exchange; volatility forecasting; Value-at-Risk; secondary data.

I. INTRODUCTION

The foreign exchange market is the largest and most liquid market in the world, with USD 7.5 trillion of daily trading volume (Bank for International Settlements, 2022). Because currency fluctuations influence international commerce, cross-border investments, and sovereign debt assessment, the ability to predict exchange rate fluctuations and measure the respective investment risk is one of many continuous challenges of financial econometrics (Cont, 2001; Meese & Rogoff, 1983). Currency returns display that they have heavy tails, volatility clustering, leverage effects, and other statistical phenomena, - thus preventing classical forecasting models from successfully predicting currency market behaviour (Cont, 2001). The most common forecasting models used for forecasting risks and volatility in currency trading are Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) (Box & Jenkins, 1970; Bollerslev, 1986). Although these models have proven valid and still represent industry standards in predicting volatility clustering, they operate under the assumptions of linear relationships and normal distribution, which are often violated in real currency markets (Cont, 2001; Chebbi & Hedhli, 2022).

In contemporary times, AI (Artificial Intelligence) and machine learning have come up as robust substitutes for financial time-series forecasting. This research employs two different AI methods. The one is Long Short-Term Memory (LSTM) network, which is a deep learning architecture developed to address the problem of vanishing gradient in conventional recurrent neural networks and to understand long-range time dependencies in sequential data (Hochreiter & Schmidhuber, 1997). The other one is Random Forest (RF), which is a machine learning method that systematically builds a number of decision trees while the learning process (Breiman, 2001). An important note has to be made that collecting primary data for high frequency multiple currency market data is an expensive and, sometimes, impossible job to accomplish. Therefore, the research will be based on secondary data – it will focus on literature available in peer-reviewed journals and those papers that have not yet been published but are very much useful for the present research (Kraus & Feuerriegel, 2024; Ndlovu, 2025; Nsengiyumva et al., 2025; Saly-Kaufmann et al., 2026; Singh & Praveen, 2025). This method allows for the comparative assessment of model performance across multiple studies rather than a single data source being used.

II. LITERATURE REVIEW

1. Traditional Econometric Models: ARIMA and GARCH

The foundation of time series forecasting in finance rests on Box-Jenkins ARIMA technique, which provides a structured approach to model identification, estimation and diagnostic testing (Box & Jenkins, 1970). Bollerslev (1986)12 introduced GARCH model, which built on Archie's (1982) Autoregressive Conditional Heteroskedasticity theory and has become the conventional standard in the foreign exchange market. It defines volatility clusters through the means of letting what happens in time series evolve under the influence of past squared error remnants and past nonconditional variances. GARCH specifications, however, suggest the existence of linear relationships between variables, which might limit their applications especially in times of financial crises (Cont, 2001). Research conducted by Chebbi and Hedhli (2022) re-examined how accurate conventional Value-at-Risk approaches are and showed that conventional GARCH models tend to underestimate tail risk and might cause insufficient capital reserves in times of market turmoils. In order to counter these drawbacks, the authors develop ARIMAX-GARCH model that adds certain exogenous variables and demonstrates remarkable directional accuracy. For instance, it shows Mean Directional Accuracy (MDA) of over 95% for certain pairs using Rand currency (Ndlovu, 2025). This means that GARCH model, even though it is linear in its essence, should be regarded as very efficient in terms of predicting direction of currency movement instead of its volume.

2. Long Short-Term Memory (LSTM) Networks

LSTM networks resolve the vanishing gradient issue that afflicts traditional recurrent neural networks and hence have the possibility for learning long-term dependencies. They do so by using a mechanism of input, forget and output gates, which control the way information flows through their memory cell (Hochreiter & Schmidhuber, 1997). Kraus & Feuerriegel (2024) performed a detailed comparative analysis of LSTM, Random Forest and AR-GARCH models with regard to EUR/USD implied volatility forecasting, showing that LSTM outperformed other models on a short-term horizon proving to be good at catching abrupt and short-term shifts in the implied volatility.

In contrast, Ndlovu (2025) revealed that LSTM showed significantly lower in-sample results for South African Rand currency pairings, recording Mean Directional Accuracy (MDA) rates nearly equal to 50%, which is close to random chance. This indicates that LSTM has a hard time coping with abrupt shifts in regimes within highly volatile emerging-market currencies and represents a critical contradiction in research: the problem of LSTM's prowess in capturing the short-term behavior of high-frequency volatility patterns (Kraus & Feuerriegel, 2024) not being able to guarantee exceptional results in the area of directional predictions for structurally volatile emerging-markets currencies (Ndlovu, 2025).

3. Random Forest and Ensemble Methods

Random Forest is an ensemble learning technique that builds multiple decision trees and uses bagged data sets and random feature subsets in training, according to Breiman's (2001) conception. This bagging approach helps to lower the variance in data compared to a single decision tree and increases resistance to the noise of financial data. It was established by Ndlovu (2025) that Random Forest had the smallest value for RMSE and MAE in most currency pairs considered, including EUR/ZAR RMSE of only 0.0037 and MAE of 0.0025, which corresponds to 88% directional accuracy in the analysis of all data.

In the same way, Milionis and Konstantinou (2024) established that Random Forest reached the best directional accuracy in cryptocurrency markets, while LSTM scored lower prediction errors on the same data. The difference between point accuracy (how close a prediction is to the actual value) and directional accuracy (whether the direction of the movement is being predicted correctly, exists in all the literature and comes with specific implications: those who trade, do care about directional accuracy while risk managers interested in knowing how much capital they need will pay more attention to the point accuracy (Milionis & Konstantinou, 2024; Ndlovu, 2025).

4. Hybrid GARCH-LSTM Models

An increasing number of studies indicate that hybridisation of traditional econometric models and deep learning leads to better results than any of the methods used separately. Hottz (2025) created a hybrid GARCH-LSTM model to predict USD/BRL exchange rate volatility and discovered that the Asymmetric Power ARCH-LSTM (APARCH-LSTM) model had 95.53% of the R^2 coefficient, which is significantly better than GARCH-type models. Nsengiyumva et al. (2025) utilised the hybrid GARCH-LSTM model for the Rwandan foreign exchange market by using the standardised residuals of the GARCH(1,1) model as the input to the LSTM network. They demonstrated that the hybrid model achieved an up to 10% better accuracy in the Value-at-Risk function estimation compared to the standard models in certain instances, especially when the markets experience abnormal volatility. These results serve as an evidence of the valid argument about the effectiveness of hybridisation – GARCH format allows to obtain the proven base levels of volatility, while LSTM is able to identify the non-linearity features in the data, which GARCH cannot express (Hottz, 2025; Nsengiyumva et al., 2025).

5. Value-at-Risk (VaR) Forecasting with AI

In their research, Nsengiyumva et al. (2025) studied and evaluated the effectiveness of GARCH, LSTM, and hybrid GARCH-LSTM methodologies for the estimation of the value-at-risk (VaR) in the forex market of Rwanda employing the data registered during the period of 2012–2024 concerning USD, EUR, and GBP currencies. The results indicate that GARCH methodology has shown the highest stability and low values of the risk indicators and estimates in comparison with the other traditional model, LSTM. The described trend corresponds to the prior conclusion made by Chebbi and Hedhli (2022) regarding the fact that the historical data applied by GARCH methodology do produce underestimations of the level of the tail risks in bad times. In contrast to the GARCH, LSTM estimated the values of risks higher reacting to changes in the market more rapidly.

6. Risk-Adjusted Investment Signals

Saly-Kaufmann et al. (2026) compared different machine learning systems across various types of financial instruments such as commodities, stock indexes, bonds, and currency markets from the period of 2010 up until 2025 in order to determine the performance of those techniques. Variable Selection Network combined with LSTM (VSN+LSTM) was found to have the best overall Sharpe ratio (2.40), with a corresponding CAGR of 24.0%, while a combination of LSTM and PatchTST performed better in the downside-adjusted metrics of performance such as CAGR (25.5%) and Calmar ratio (1.47). One more architecture called xLSTM has shown that it is characterized by the largest buffer for breakeven transaction costs in comparison to the other models being investigated.

In an accompanying study which focussed on equities rather than multi-asset portfolios, Singh and Praveen (2025) indicate that XGBoost, a gradient boosting tree, comparable to Random Forest (Breiman, 2001), produced superior performance among NIFTY 50 stocks and obtained the highest mean Sharpe ratio (1.72) and the lowest maximum drawdown among the three techniques. This is consistent with the indication of directional accuracies by Milionis and Konstantinou (2024) and Ndlovu (2025) that tree-based ensemble methods have a better risk-adjusted performance than pure sequence models across asset classes, despite deep-learning hybrids having an advantage with return maximisation of multi-assets (Saly-Kaufmann et al., 2026).

III. RESEARCH GAP

While previous research has shown the success of the AI models independently, the existing gaps still need to be closed. First, many studies have relied on LSTM vs GARCH comparison without introducing Random Forest in the context of directional accuracy evaluation using the same data (Kraus & Feuerriegel, 2024; Ndlovu, 2025). Second, very few studies have examined the effectiveness of AI models in their capacity to forecast Value-at-Risk of currency markets under various confidence levels and in many developing and developed countries (Nsengiyumva et al., 2025). Third, there are few options for comparing the risk-adjusted investment signals Sharpe ratio, Calmar ratio, maximum drawdown for AI models in the context of currency trading instead of stocks and commodities (Saly-Kaufmann et al., 2026; Singh & Praveen, 2025). To close these gaps, the paper uses data from different secondary sources to provide a comprehensive and comparative analysis of research involving LSTM, Random Forest, GARCH, and their combinations in the prediction of currency fluctuations, Value-at-Risk estimation, and generation of risk-adjusted signals.

Research Objectives

This study pursues three specific objectives:

- To compare the predictive accuracy of LSTM and Random Forest models against the traditional GARCH model for short-term (daily) currency price direction.
- To evaluate the ability of AI models to forecast Value-at-Risk (VaR) for currency markets relative to standalone GARCH benchmarks.
- To assess which AI technique yields superior risk-adjusted investment signals for currency-related trading strategies.

Significance of the Study

- For Forex Traders and Fund Managers: the findings provide evidence-based guidance on selecting appropriate AI models for different currency pairs and market conditions, potentially reducing unexpected trading losses.
- For Risk Analysts and Financial Institutions: the comparison of VaR estimation accuracy across models offers insights for improving capital-reserve calculations and complying with regulatory requirements.
- For Academic Researchers: the synthesis of comparative performance data across multiple studies establishes a benchmark for future primary research on AI applications in financial forecasting.
- For Quantitative Developers: the analysis of risk-adjusted performance metrics Sharpe ratio, maximum drawdown, and CAGR provides practical guidance for model selection in automated trading systems.

IV. RESEARCH METHODOLOGY

This study employs a systematic secondary-data (desk-research) synthesis design to evaluate the predictive accuracy of AI models LSTM and Random Forest against traditional GARCH models for currency-fluctuation prediction, Value-at-Risk forecasting, and risk-adjusted investment-signal generation. Rather than estimating new models on primary exchange-rate data, the methodology integrates quantitative findings reported in peer-reviewed journal articles and working papers, each utilising distinct datasets, currency pairs, and evaluation periods (Ndlovu, 2025; Nsengiyumva et al., 2025; Saly-Kaufmann et al., 2026).

Sources were purposively selected on the basis of three inclusion criteria: (i) the study empirically compares at least one AI technique (LSTM, Random Forest, XGBoost, or a hybrid architecture) against a traditional econometric benchmark such as ARIMA or GARCH; (ii) the study reports quantitative performance metrics RMSE, MAE, MAPE, sMAPE, MDA, VaR estimates, Sharpe ratio, Calmar ratio, or maximum drawdown that permit cross-study comparison; and (iii) the study applies these methods to a foreign-exchange, cross-asset, or closely related financial time series. Reported figures, tables, and effect sizes were extracted directly from the original sources and are reproduced or summarised, with full attribution, in Section 7 and Table 4.

V. DATA ANALYSIS AND FINDINGS

Finding 1: Comparative Predictive Accuracy for Currency Price Direction

Table 1 presents the comparative performance of ARIMAX-GARCH, Random Forest, and LSTM models across five currency pairs against the South African Rand, as reported by Ndlovu (2025).

Table 1: Comparative Predictive Accuracy of ARIMAX-GARCH, Random Forest, and LSTM Models

Model	Currency	RMSE	MAE	MAPE	sMAPE	MDA (%)
ARIMAX-GARCH	EUR/ZAR	0.0805	0.0218	0.1951	0.2081	87.62
ARIMAX-GARCH	USD/ZAR	0.0749	0.0084	0.0821	0.0952	96.86
ARIMAX-GARCH	JPY/ZAR	0.0060	0.0003	0.3504	0.2466	85.30
ARIMAX-GARCH	GBP/ZAR	0.1321	0.0308	0.2080	0.2227	87.71
ARIMAX-GARCH	CNY/ZAR	0.0024	0.0011	0.0698	0.0700	95.91
Random Forest	EUR/ZAR	0.0037	0.0025	0.1063	0.1063	91.42
Random Forest	USD/ZAR	0.0030	0.0017	0.0747	0.0747	94.90
Random Forest	JPY/ZAR	0.0057	0.0035	0.1495	0.1495	92.27
Random Forest	GBP/ZAR	0.0060	0.0043	0.1588	0.1588	88.83
Random Forest	CNY/ZAR	0.0038	0.0020	1.5038	1.3768	94.35
LSTM	EUR/ZAR	0.0434	0.0355	1.4373	1.4261	50.28

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 4, July-Aug-2026, ISSN (Online): 2395-566X

Model	Currency	RMSE	MAE	MAPE	sMAPE	MDA (%)
LSTM	USD/ZAR	0.0550	0.0465	2.0453	2.0176	50.35
LSTM	JPY/ZAR	0.0590	0.0484	2.0761	1.5060	47.06
LSTM	GBP/ZAR	0.0358	0.0285	1.0471	1.0547	50.17
LSTM	CNY/ZAR	0.0255	0.0207	24.2565	13.3565	50.69

RMSE (Root Mean Square Error) measures average prediction error while heavily penalising large mistakes, so lower values indicate better accuracy; MAE (Mean Absolute Error) is the simple average absolute prediction error, for which lower values are likewise preferred. As Table 1 shows, Random Forest achieved the lowest RMSE and MAE across most currency pairs, with an EUR/ZAR RMSE of just 0.0037 and MAE of 0.0025 (Ndlovu, 2025). ARIMAX-GARCH, by contrast, excelled in directional accuracy (MDA), achieving 96.86% for USD/ZAR and 95.91% for CNY/ZAR, while LSTM showed markedly weaker in-sample performance, with MDA values near 50% essentially random-guess level across all five pairs.

Finding 2: Value-at-Risk Estimation Across Models

AI models including LSTM and hybrid GARCH-LSTM approaches have been used to forecast Value-at-Risk (VaR) at the 95% confidence level, compared against traditional standalone GARCH models. Table 2 summarises VaR estimates reported for the Rwandan forex market by Nsengiyumva et al. (2025).

Table 2: VaR Estimates for the Rwandan Forex Market (95% Confidence, One-Day Horizon, 100 Million FRW Investment)

Currency	Model	VaR Estimate (FRW)	Interpretation
USD	GARCH(1, 1)	2,109,637.22	Lowest risk estimate (most conservative)
USD	LSTM	4,522,130.21	Highest risk estimate (most sensitive to volatility)
USD	Hybrid GARCH-LSTM	3,866,771.26	Moderate estimate (balances both approaches)
EUR	GARCH(1, 1)	33,529,520.73	Baseline (conservative) estimate

The GARCH model consistently produced the lowest risk estimates across all currencies, while the LSTM model yielded significantly higher estimates, reflecting its greater responsiveness to short-term fluctuations and sensitivity to market volatility (Nsengiyumva et al., 2025). The hybrid model produced a VaR estimate that fell between the two standalone models, improving estimation accuracy by up to 10%, especially during periods of elevated market volatility (Nsengiyumva et al., 2025).

Finding 3: Risk-Adjusted Performance of AI Techniques

Table 3: Risk-Adjusted Performance of AI Techniques Across Asset Classes

AI Technique	Sharpe Ratio	Max. Drawdown	CAGR / Annual Return	Key Context	Source
VSN + LSTM	2.40	Not specified	24.0%	Highest overall Sharpe ratio across commodities, equity indices, bonds, and FX (2010–2025)	Saly-Kaufmann et al. (2026)
LSTM + PatchTST	2.31	Not specified	25.5%	Superior downside-adjusted characteristics; Calmar ratio of 1.47	Saly-Kaufmann et al. (2026)
xLSTM	1.80	Not specified	19.3%	Largest breakeven transaction-cost buffer; Calmar ratio of 1.35	Saly-Kaufmann et al. (2026)
XGBoost	1.72	Lowest (best)	Not specified	Outperformed Random Forest and LSTM across 10 NIFTY 50 stocks (2014–2024)	Singh & Praveen (2025)
Random Forest	Not reported	Higher than XGBoost	Not specified	Competitive but outperformed by XGBoost	Singh & Praveen (2025)
LSTM	Lower than XGBoost	Higher variability	Not specified	Weaker overall performance, likely due to limited dataset size	Singh & Praveen (2025)

VSN+LSTM achieved the highest overall Sharpe ratio (2.40) across multi-asset portfolios including foreign exchange (Saly-

Kaufmann et al., 2026), while LSTM+PatchTST exhibited superior downside-adjusted characteristics, with a 25.5% CAGR (Saly-Kaufmann et al., 2026).

Interpretation of Findings

The results illustrate the relationship between the accuracy of the point prediction over the directional accuracy. In most cases, Random Forest has shown the lowest value of RMSE and MAE for most currency pairs. The RMSE is only 0.0037 and MAE 0.0025 for EUR/ZAR currency pair, and the directional accuracy remains high above 88% for all different currency pairs (Ndlovu, 2025). This means that the Random Forest model is capable of understanding the complicated nature of the movements of various currencies because it is based on the ensemble method that is also known as the bootstrapping method (Breiman, 2001). On the other hand, ARIMAX-GARCH showed great results in the directional accuracy with 96.86% level for USD/ZAR and 95.91% for the CNY/ZAR currency pairs. Therefore, it can be concluded that traditional econometric models can be very effective in predicting if the currency becomes stronger or weaker against other currencies (Ndlovu, 2025).

The most remarkable discovery is that LSTM performs poorly in terms of directional accuracy, with MDA values close to 50% in terms of all currency pairs, which is little better than random selection. This correlates with the finding that LSTM does not perform well in turbulent times and sudden market changes, especially for emerging markets' currencies (Ndlovu, 2025), while the same approach works quite successfully for predicting short-term implied volatility rates in more stable and highly liquid market pairs such as EUR/USD (Kraus and Feuerriegel, 2024). As a result, the main take-home message for currency market traders is that one cannot use LSTM without thinking twice about it as far as volatile market currencies are concerned.

The results of the study into GARCH, LSTM, and hybrid GARCH-LSTM methodology for calculating Value at Risk showed very striking differences in risk behaviours. The GARCH model indicated consistently the smallest risk values across all studied currencies. For example, the VaR was 2,109,637.22 FRW for a 100-million-FRW investment in USD at a confidence level of 95 percent (Nsengiyumva et al., 2025). This conservative approach may lead to an underestimation of risks in the tail point during volatile periods and could be considered a limitation of GARCH models (Chebbi & Hedhli, 2022). Unlike GARCH, the LSTM model provided much higher values of the VaR (4,522,130.21 FRW for USD), which is a result of the LSTM design making it more responsive to fluctuations in the short-term and sensitive to volatility. However, it must be noted that although this contributes to the cautiousness of LSTM, it may still lead to a high level of kept capital reserves and low efficiency of trading. The hybrid model GARCH-LSTM produced a moderate VaR (3,866,771.26 FRW for USD) meaning that it enabled the researchers to combine both methodologies in one model and to improve the accuracy of the VaR calculation by 10 percent during turbulent market periods (Nsengiyumva et al., 2025).

The study conducted by Saly-Kaufmann et al. (2026) offers an extensive analysis of various deep learning models used for gaining returns adjusted for risks in different asset classes, including different currencies. Empirical proof of the effective use of VSN+LSTM models for securing the optimal Sharpe ratio of 2.40 and 24.0% CAGR from 2010 to 2025 is significant as the Sharpe ratio is one of the most prominent indicators for measuring returns earned against the risk. These findings indicate that hybrid models that combine the variable selection networks with LSTMs have proven to be effective in finding and utilizing the patterns of relationship across various asset classes. In addition, LSTM + PatchTST model has received credit for the highest downside adjusted performance characteristics with 25.5% CAGR and a good Calmar ratio of 1.47, which proves to be crucial for investors who seek to protect their capital in exchange for the lowest returns (Saly-Kaufmann et al., 2026). The fact that xLSTM has proved to be the most flexible model in terms of transaction costs is significantly important for using it in currency trading.

VI. SYNTHESIS TABLE OF REVIEWED SECONDARY SOURCES

Table 4 consolidates the full set of secondary sources reviewed in this paper, summarising each study's focus, data and sample, method(s) applied, and principal finding, to give readers a single point of reference for the evidence base underpinning Sections 2, 7, and 8.

Table 4: Comprehensive Summary of Secondary Sources Reviewed

Author(s) & Year	Focus / Topic	Data / Sample	Method(s)	Principal Finding	Outlet
Bank for International Settlements [BIS] (2022)	Global FX market scale	Triennial central bank survey	Survey / descriptive statistics	Daily FX turnover exceeds US$7.5 trillion, underscoring market liquidity	BIS Triennial Survey

Author(s) & Year	Focus / Topic	Data / Sample	Method(s)	Principal Finding	Outlet
				and systemic relevance	
Bollerslev (1986)	Volatility modelling	Theoretical / methodological	GARCH model derivation	Conditional variance depends on past squared residuals and variances; foundational volatility benchmark	Journal of Econometrics
Box & Jenkins (1970)	Time-series forecasting	Theoretical / methodological	ARIMA methodology	Established systematic framework for model identification, estimation, and diagnostics	Book (Holden-Day)
Breiman (2001)	Ensemble learning	Theoretical / methodological	Random Forest algorithm	Bagged decision-tree ensembles reduce variance and improve predictive robustness	Machine Learning
Chebbi & Hedhli (2022)	VaR model accuracy	Multi-market VaR backtests	GARCH-based VaR evaluation	Standard GARCH specifications can underestimate tail risk during market crises	Journal article (risk management)
Cont (2001)	Stylised facts of returns	Cross-market empirical review	Empirical / statistical review	Financial returns exhibit heavy tails,	Quantitative Finance
				volatility clustering, and leverage effects violating linear-model assumptions	
Hochreiter & Schmidhuber (1997)	Deep learning architecture	Theoretical / methodological	LSTM network derivation	Gated memory-cell architecture overcomes vanishing gradients, enabling long-range sequence learning	Neural Computation
Hottz (2025)	Hybrid volatility forecasting	USD/BRL exchange-rate series	APARCH-LSTM hybrid model	APARCH-LSTM achieves R^2 of 95.53% for USD/BRL volatility forecasting	Working paper / journal article
Kraus & Feuerriegel (2024)	Implied-volatility forecasting	EUR/USD implied volatility, multiple option maturities	LSTM vs. Random Forest vs. AR-GARCH	LSTM performs best for shorter option maturities, capturing immediate volatility shifts	Journal article
Meese & Rogoff (1983)	Exchange-rate model evaluation	Major currency pairs, out-of-sample	Empirical exchange-rate model comparison	Structural exchange-rate models fail to outperform a random walk out-of-sample	Journal of International Economics

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 4, July-Aug-2026, ISSN (Online): 2395-566X

Author(s) & Year	Focus / Topic	Data / Sample	Method(s)	Principal Finding	Outlet
Milionis & Konstantinou (2024)	Cryptocurrency forecasting	Cryptocurrency price series	Random Forest vs. LSTM	Random Forest achieves highest directional accuracy despite LSTM's lower point-prediction error	Journal article
Ndlovu (2025)	ZAR exchange-rate forecasting	5 Rand-denominated currency pairs, daily data	ARIMAX-GARCH vs. Random Forest vs. LSTM	Random Forest lowest RMSE/MAE; ARIMAX-GARCH highest MDA; LSTM near-random MDA (~50%)	Journal article
Nsengiyumva, Mung'atu, & Ruranga (2025)	Hybrid VaR forecasting	USD, EUR, GBP vs. Rwandan Franc, 2012–2024	GARCH (1,1) vs. LSTM vs. hybrid GARCH-LSTM	Hybrid model improves VaR estimation accuracy by up to 10% during volatile periods	FinTech (MDPI)
Saly-Kaufmann et al. (2026)	Risk-adjusted multi-asset signals	Commodities, equity indices, bonds, FX, 2010–2025	VSN+LSTM, LSTM+PatchTST, xLSTM benchmarking	VSN+LSTM highest Sharpe ratio (2.40); LSTM+PatchTST best downside-adjusted return	Working paper / journal article
Singh & Praveen (2025)	Equity risk-adjusted signals	10 NIFTY 50 stocks, 2014–2024	XGBoost vs. Random Forest vs. LSTM	XGBoost highest Sharpe ratio (1.72) and lowest maximum drawdown	Journal article

VII. CONCLUSION

The secondary data demonstrates that there is no AI model that can be regarded as universally the best across different currency pairs, time frames, or tasks. For directional trading purposes, Random Forest is preferred as its prediction error is consistently low, and it has good ability to predict directions of movements. As for risk management, hybrid GARCH-LSTM models have been shown to be effective in improving Value-at-Risk estimation accuracy. Model choice must be commensurate with the specific investment objectives, time frame, and currency pair. Also, it should be understood that AI brings benefits compared to conventional GARCH models for currency prediction purposes but only if the choice of the model is made to fit the task; the hybrid model that combines econometric techniques with deep learning is better than the other models. The main conclusion is not about which model is "the best" but which one is most appropriate for a specific currency pair, time frame, and risk preferences.

Limitations of the Study

This research employs the use of secondary data amalgamation, obtaining outcomes from diverse databases used in different periods of time, foreign exchange pairs, and means of evaluation, which may limit the applicability of the resulting comparative findings. One particular source reports fairly low LSTM success rate (just under 50% of MDA) in obvious contrast to other research proving LSTM efficiency with stable currencies, which probably measures the extreme volatility level of ZAR currency pairs rather than the architecture disadvantages of LSTM. Furthermore, being based on narrative rather than statistical evidence, comparative conclusions are not confirmed with the help of meta-analysis as effect sizes as well as standard errors are not reported uniformly in the various sources, hence, the uncertainty level remains high. There are also some limitations in terms of practicality including computational costs, problems with real-time implementation, and model interpretability because of the "black-box" nature of the deep learning models in relation to Basel requirements as well as finding out ways to compute risk-adjusted performance metrics that can be backtested and have been applied in the research. It is, thus, necessary to interpret the reported Sharpe ratios being rather careful due to the potential market impact, as well as considering the viability dimension of the findings in terms of out-of-sample use.

Recommendations for Future Research

- Future studies should develop unified frameworks that directly compare LSTM, Random Forest, GARCH, and hybrid models using the same dataset, currency pairs, and evaluation period to eliminate confounding factors arising from cross-study comparison.
- Researchers should investigate the specific conditions under which LSTM succeeds or fails for currency prediction, potentially identifying volatility thresholds or currency characteristics that determine model suitability.
- Future work should examine the real-time implementation challenges of hybrid deep-learning models for currency trading, including computational cost, inference latency, and transaction-cost management.
- Research should explore interpretable AI techniques, such as attention mechanisms or SHAP values, to address the "black-box" concern for regulatory compliance.
- Longitudinal studies tracking the out-of-sample performance of these models over extended periods (ten or more years) would provide more robust evidence of their practical value for currency trading and risk management.

REFERENCES

1. Bank for International Settlements. (2022). OTC foreign exchange turnover in April 2022 (Triennial Central Bank Survey). https://www.bis.org/statistics/rpfx22_fx.htm
2. Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1
3. Box, G. E. P., & Jenkins, G. M. (1970). Time series analysis: Forecasting and control. Holden-Day.
4. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
5. Chebbi, A., & Hedhli, A. (2022). Revisiting the accuracy of Value-at-Risk methods: Evidence from GARCH-family models [Citation details as provided in the source manuscript; verify journal, volume, and page numbers before submission].
6. Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. Quantitative Finance, 1(2), 223–236. https://doi.org/10.1080/713665670
7. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735
8. Hottz, [Initial]. (2025). Hybrid APARCH-LSTM volatility forecasting for the USD/BRL exchange rate [Citation details as provided in the source manuscript; verify journal, volume, and page numbers before submission].
9. Kraus, M., & Feuerriegel, S. (2024). Comparing LSTM, Random Forest, and AR-GARCH models for EUR/USD implied-volatility forecasting [Citation details as provided in the source manuscript; verify journal, volume, and page numbers before submission].
10. Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? Journal of International Economics, 14(1–2), 3–24. https://doi.org/10.1016/0022-1996(83)90017-X
11. Milionis, [Initial], & Konstantinou, [Initial]. (2024). Random Forest and LSTM for directional accuracy in cryptocurrency markets [Citation details as provided in the source manuscript; verify journal, volume, and page numbers before submission].
12. Ndlovu, T. (2025). Forecasting exchange rate dynamics: A comparative study of traditional econometric models and machine learning models [Citation details as provided in the source manuscript; verify journal, volume, and page numbers before submission].
13. Nsengiyumva, E., Mung'atu, J. K., & Ruranga, C. (2025). Hybrid GARCH-LSTM forecasting for foreign exchange risk. FinTech, 4(2), Article 22. https://doi.org/10.3390/fintech4020022
14. Saly-Kaufmann, [Initial], et al. (2026). Benchmarking deep-learning architectures for risk-adjusted multi-asset investment signals [Citation details as provided in the source manuscript; verify journal, volume, and page numbers before submission].
15. Singh, [Initial], & Praveen, [Initial]. (2025). XGBoost, Random Forest, and LSTM for risk-adjusted equity trading signals: Evidence from NIFTY 50 [Citation details as provided in the source manuscript; verify journal, volume, and page numbers before submission].