

Article **Open Access**

Portfolio Risk Management in the Age of AI: Analyzing NVIDIA's Market Performance and Predictive Models for High-Tech Investments

Ran Duan ^{1,*}

¹ University of International Business and Economics, Beijing, China

* Correspondence: Ran Duan, University of International Business and Economics, Beijing, China

Abstract: This study addresses the evolving challenges of portfolio risk management in the era of Artificial Intelligence (AI), using NVIDIA Corporation (NVDA) as a representative high-tech asset. We propose an integrated AI-augmented framework that combines a hybrid Long Short-Term Memory (LSTM)-GARCH (1,1) model with alternative data, including macroeconomic indicators and financial sentiment derived from news and social media. Trained on daily data from 2018 to 2022 and evaluated out-of-sample in 2023, the hybrid model significantly outperforms ARIMA, standalone LSTM, and GARCH benchmarks in forecasting both NVDA returns (MAE: 1.72%) and conditional volatility (Mincer-Zarnowitz R^2 : 0.68). Risk analysis based on model outputs reveals NVDA's pronounced exposure to semiconductor supply chain disruptions, regulatory shifts, and tech-sector sentiment. A dynamic portfolio strategy that adjusts NVDA allocation according to predicted volatility achieves a higher Sharpe ratio (3.38 vs. 2.95) and lower maximum drawdown (-28.5% vs. -42.1%) than a static benchmark. Our findings demonstrate that adaptive, AI-driven risk models are essential for managing the non-linear dynamics and tail risks characteristic of AI-centric investments.

Keywords: portfolio risk management; artificial intelligence; NVIDIA; machine learning; LSTM; GARCH; volatility forecasting; high-tech investments; CVaR; sentiment analysis

Received: 07 November 2025

Revised: 29 December 2025

Accepted: 13 January 2026

Published: 18 January 2026



Copyright: © 2025 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The dawn of the Artificial Intelligence (AI) era has catalyzed a profound transformation across global industries, with the financial sector experiencing one of its most significant disruptions. AI is no longer merely a tool for automation; it has become a primary driver of corporate strategy, market valuation, and systemic risk [1]. Nowhere is this more evident than in the meteoric rise of companies positioned at the epicenter of the AI hardware and software ecosystem. NVIDIA Corporation (NVDA) stands as a quintessential case study, a company that transitioned from a graphics processing unit (GPU) specialist to the de facto engine powering the global AI revolution [2]. Its stock price, reflecting immense investor optimism about AI's future, has exhibited extraordinary growth coupled with pronounced volatility, presenting unique challenges and opportunities for portfolio managers [3].

Traditional portfolio risk management frameworks, often rooted in Modern Portfolio Theory (MPT) and reliant on historical variance-covariance matrices and linear factor models (e.g., CAPM, Fama-French), struggle to adequately capture the dynamics of such AI-centric assets [4]. These models frequently assume normally distributed returns, linear relationships, and stable correlations, assumptions demonstrably violated in the context

of disruptive technological innovation characterized by hype cycles, rapid paradigm shifts, network effects, and sentiment-driven price surges [5]. The "black swan" events associated with technological breakthroughs or regulatory crackdowns can induce tail risks that standard Value-at-Risk (VaR) measures fail to anticipate [6].

Consequently, there is an urgent need to develop and validate next-generation risk management methodologies that leverage the very technologies, AI and Machine Learning (ML), driving the market changes they seek to manage. ML algorithms, particularly deep learning architectures like Recurrent Neural Networks (RNNs) and their variants (e.g., LSTMs), excel at identifying complex, non-linear patterns and temporal dependencies in high-dimensional data, making them ideally suited for modeling the intricate behavior of high-tech stocks like NVDA [7]. Furthermore, integrating alternative data sources, such as real-time news sentiment and social media chatter, provides crucial forward-looking signals that traditional financial data often lags.

This paper addresses this critical gap by presenting a comprehensive investigation into portfolio risk management for AI-driven high-tech investments, using NVIDIA as our primary empirical subject. Our core contributions are threefold: (1) Framework Development: We design and implement an integrated risk analytics framework that synergistically combines deep learning for price prediction with econometric models for volatility forecasting, augmented by alternative data. (2) Empirical Validation: We rigorously test our hybrid LSTM-GARCH model against established benchmarks using out-of-sample forecasting, demonstrating superior predictive power for both NVDA's price levels and its conditional volatility. (3) Practical Risk Application: We translate these predictive insights into actionable portfolio risk metrics (CVaR, stress testing) and construct illustrative portfolios to showcase how dynamic allocation based on our model outputs can optimize the risk-return profile for investors exposed to the AI sector.

2. Literature Review

The intersection of AI, finance, and risk management has garnered substantial academic and practitioner interest. Early applications of AI in finance focused on algorithmic trading and credit scoring using simpler models such as decision trees and support vector machines (SVMs) [8,9]. The advent of deep learning, however, has unlocked new capabilities for time-series forecasting. A key development in this area was the introduction of the long short-term memory (LSTM) architecture, which proved particularly effective for financial time series due to its ability to mitigate the vanishing gradient problem and capture long-term dependencies, features essential for modeling asset prices shaped by persistent trends and regime shifts [10].

Concurrently, financial econometrics has long addressed the challenge of modeling volatility clustering, a well-documented stylized fact in financial returns whereby periods of high volatility tend to be followed by further high volatility. This led to the development of the autoregressive conditional heteroscedasticity (ARCH) model and its generalized extension (GARCH), which became foundational tools for volatility forecasting and risk measurement, particularly in value-at-risk (VaR) estimation [11,12]. Despite their widespread use, these models are inherently parametric and may have limitations in capturing the complex, non-linear volatility dynamics triggered by exogenous shocks, such as major technological announcements or geopolitical disruptions affecting global semiconductor supply chains.

Recent research has sought to bridge the gap between these paradigms. Several studies have explored hybrid models, combining the pattern recognition prowess of neural networks with the statistical rigor of GARCH-type models [13]. For instance, some approaches use neural networks to predict the residuals or the conditional variance directly, while others employ them to forecast the mean equation, feeding into a GARCH volatility filter [14]. However, the application of such sophisticated hybrids specifically to the unique risk profile of dominant AI enablers like NVIDIA remains underexplored.

Furthermore, the role of sentiment analysis in financial markets has gained prominence. Studies have shown that textual data from news articles, earnings call transcripts, and social media platforms (e.g., Twitter, StockTwits) contain valuable predictive signals about future price movements and volatility, often preceding moves reflected in traditional price and volume data [15]. Integrating these unstructured data sources into quantitative models represents a significant advancement in capturing market psychology, which is especially potent in hype-driven sectors like AI.

Our work builds upon this foundation but focuses explicitly on the practical challenge of managing portfolio risk in the current AI boom, using a state-of-the-art hybrid model applied to a real-world, high-impact asset, and translating the outputs into concrete risk management actions.

3. Methodology

3.1. Research Framework

Our analytical approach is encapsulated in the research framework depicted in Figure 1. The process begins with multi-source data acquisition, encompassing structured financial data, macroeconomic indicators, and unstructured textual data. This raw data undergoes extensive preprocessing and feature engineering to create a unified, model-ready dataset. The core of our framework is the Hybrid LSTM-GARCH model, which simultaneously learns the conditional mean (price trend) and conditional variance (volatility) of NVDA returns. The model's outputs, predicted prices and predicted volatilities, are then fed into a risk analytics module that calculates key portfolio risk metrics (VaR, CVaR) and performs scenario-based stress tests. Finally, these risk insights inform dynamic portfolio allocation strategies.




Figure 1. Integrated AI-Augmented Portfolio Risk Management Framework.

3.2. Data Collection and Description

Our analysis draws on a comprehensive dataset spanning January 1, 2018, to December 31, 2023, a period that encompasses NVIDIA's strategic evolution into a leading AI infrastructure provider. This interval includes pivotal developments such as the rollout of its data center GPUs, major collaborations with cloud service providers, and notable

market downturns. Daily Open, High, Low, Close (OHLC) prices and trading volume for NVIDIA (NVDA) were obtained from Yahoo Finance, while the S&P 500 Index (SPX) served as the broad market benchmark. Macroeconomic indicators, including the U.S. Federal Funds Rate, 10-Year Treasury Yield, Consumer Price Index, and the Semiconductor Equipment Billings (SEMI) index, were sourced from the Federal Reserve Economic Data (FRED) database. To incorporate market sentiment, a daily sentiment score for NVIDIA was constructed by aggregating financial news headlines and social media posts from sources like Reuters, Bloomberg, StockTwits, and Twitter. Each text item mentioning "NVIDIA" or "NVDA" was processed using a pre-trained FinBERT model to generate a sentiment polarity score between -1 and +1, which were then combined into a volume-weighted average daily sentiment metric.

3.3. Feature Engineering

The raw data was transformed into model-ready features through systematic engineering. Daily log-returns for NVIDIA and the S&P 500 were computed, alongside technical indicators including 10- and 30-day simple moving averages, RSI, and Bollinger Bands. To account for delayed market reactions, returns and sentiment scores were lagged by 1, 2, and 5 days. A volatility proxy was derived as the 5-day rolling standard deviation of returns, while changes in macroeconomic variables (e.g., $\Delta\text{FEDFUNDS}$) captured shifting economic conditions. Finally, all numerical features were standardized to have zero mean and unit variance to ensure consistent model performance.

3.4. Model Specification: Hybrid LSTM-GARCH

Our core predictive model is a hybrid LSTM-GARCH framework that jointly models the conditional mean and variance of NVIDIA's daily returns. The conditional mean is captured by an LSTM network that processes a 30-day window of historical features, including returns, technical indicators, macroeconomic variables, and sentiment scores, and outputs a predicted return $\hat{\mu}_t$. The architecture comprises two hidden layers (64 and 32 units) with dropout regularization (rate = 0.2). The conditional variance follows a GARCH (1,1) process: $\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$, where $\varepsilon_{t-1} = r_{t-1} - \hat{\mu}_{t-1}$ is the residual from the LSTM-predicted mean. Parameters ω , α , β , and LSTM weights are estimated jointly by minimizing a composite loss function: $\text{Loss} = \text{MSE}(\hat{\mu}_t, r_t) + \lambda \cdot \text{MSE}(\hat{\sigma}_t^2, \text{RV}_t)$, with $\lambda=0.5$ and RV_t denoting the 5-day realized volatility proxy. The hybrid model is benchmarked against ARIMA(1,1,1), a standalone LSTM, and a standard GARCH(1,1) with constant mean.

3.5. Risk Metrics and Portfolio Analysis

Using the hybrid model's predicted conditional return distribution, assumed normal for tractability, we compute 1-day 95% and 99% Value-at-Risk (VaR) and Conditional VaR (CVaR), with CVaR providing a more robust measure of tail risk by estimating expected losses beyond the VaR threshold. We also conduct scenario-based stress testing under three AI-sector-specific shocks: a semiconductor supply chain disruption (20% drop in the SEMI index plus 50% volatility spike), a regulatory crackdown (sentiment falling by two standard deviations alongside a 100 bps rise in rates), and a broad tech correction (15% Nasdaq-100 decline). For portfolio analysis, we compare a static 60% NVIDIA / 40% S&P 500 allocation against a dynamic strategy that reduces NVIDIA exposure to 40% when predicted 30-day volatility exceeds its 6-month moving average, otherwise maintaining 60% weight, with the balance always in the S&P 500.

3.6. Model Training Protocol and Validation Strategy

To ensure robustness and mitigate the risk of overfitting, a critical concern when deploying deep learning models on financial time series, we implemented a rigorous training and validation protocol. The dataset was partitioned chronologically: data from

January 2018 to December 2021 served as the training set, January 2022 to December 2022 as the validation set for hyperparameter tuning, and the full year of 2023 was reserved exclusively for out-of-sample testing. This temporal split prevents look-ahead bias and reflects real-world deployment conditions where future data is inaccessible during model development.

Hyperparameter optimization was conducted using Bayesian optimization over a predefined search space, targeting the composite loss function described in Section 3.4. Key hyperparameters included the LSTM layer sizes (32-128 units), dropout rate (0.1-0.5), learning rate ($1e-4$ to $1e-2$), and the volatility weighting coefficient λ (0.1-1.0). Early stopping was employed with a patience of 15 epochs based on validation loss to prevent overtraining. All models were implemented in Python using TensorFlow/Keras for the neural components and arch library for GARCH estimation, with training performed on a GPU-accelerated computing environment to handle the computational load efficiently.

Furthermore, we conducted a rolling-window backtest over the 2022-2023 period, re-estimating the hybrid model every 30 trading days. This simulates a practical investment setting where models are periodically refreshed with new data. The consistency of performance across these rolling windows provided additional confidence in the model's stability and generalizability, particularly during periods of heightened market turbulence such as the rapid interest rate hikes of 2022 and the AI-driven rally of 2023.

4. Results and Analysis

4.1. Descriptive Statistics and Preliminary Analysis

Table 1 presents key descriptive statistics for NVDA and the SPX over our sample period. NVDA exhibits dramatically higher average returns (0.28% daily vs. 0.04%) but also significantly higher volatility (standard deviation of 3.52% vs. 1.35%). Its kurtosis (12.8) is far above 3, indicating fat tails and a higher probability of extreme events compared to the SPX (kurtosis=8.1). The maximum drawdown for NVDA (-67.2%) was also much more severe than for the SPX (-33.8%), highlighting its elevated risk profile despite superior returns.

Table 1. Descriptive Statistics of Daily Log-Returns (Jan 2018 - Dec 2023).

Statistic	NVIDIA (NVDA)	S&P 500 (SPX)
Mean (Daily %)	0.28%	0.04%
Standard Deviation	3.52%	1.35%
Skewness	-0.42	-0.78
Kurtosis	12.8	8.1
Minimum Return	-22.1%	-12.0%
Maximum Return	+24.4%	+11.5%
Max Drawdown	-67.2%	-33.8%
Sharpe Ratio (Annual)	1.28	0.48

A correlation analysis revealed that NVDA's correlation with the SPX was moderate (0.65) over the full period but exhibited significant time-variation, spiking during market-wide sell-offs (e.g., >0.8 during the 2022 bear market) and decoupling during AI-specific rallies. This time-varying beta underscores the limitation of static correlation assumptions in MPT.

4.2. Model Performance Evaluation

We conducted an out-of-sample forecast evaluation for the year 2023 (see Table 2). The models were trained on data from 2018-2022 and used to generate daily 1-step-ahead forecasts for 2023. Performance was assessed using Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) for return prediction, and the Mincer-Zarnowitz regression for volatility forecast accuracy.

Table 2. Out-of-Sample Forecast Performance (2023).

Model	Return Prediction	Volatility Prediction (Mincer-Zarnowitz)		
	MAE	RMSE	Slope (β)	R ²
ARIMA(1,1,1)	2.15%	2.87%	0.42	0.28
Standalone LSTM	1.98%	2.65%	N/A	N/A
GARCH(1,1)	2.30%	3.10%	0.78	0.52
Hybrid LSTM-GARCH	1.72%	2.31%	0.95	0.68

As shown in Table 2, the Hybrid LSTM-GARCH model consistently outperforms all benchmarks. It achieves the lowest MAE and RMSE for return prediction, demonstrating its superior ability to capture the complex price dynamics of NVDA. More importantly, its volatility forecasts are remarkably accurate. The Mincer-Zarnowitz regression slope of 0.95 is very close to the ideal value of 1.0, indicating that the forecasted volatility is an unbiased predictor of realized volatility, and the R² of 0.68 shows it explains a substantial portion of the variance in future volatility. The standalone GARCH model, while decent for volatility, performs poorly on returns, while the standalone LSTM cannot provide a formal volatility forecast.

A visual inspection of the forecasts (not shown here for brevity) confirmed that the hybrid model was particularly adept at anticipating large price swings following major events, such as NVIDIA's quarterly earnings announcements or significant AI-related news, largely due to the integration of the sentiment feature.

4.3. Risk Analysis: VaR, CVaR, and Stress Testing

Using the hybrid model's conditional distribution forecasts for 2023, we calculated daily VaR and CVaR. On average, the 99% 1-day VaR for a 1M position in NVDA was 78,500, while the CVaR was 112,000. This means that on the worst 112,000, significantly higher than the VaR threshold, emphasizing the severity of tail losses.

The scenario stress tests yielded alarming results, as summarized below: (1) Supply Chain Shock: This scenario resulted in a projected portfolio loss of -32% for a 100% NVDA position. The model indicated that NVDA's beta to the semiconductor equipment sector would surge to over 2.0 during such a shock. (2) Regulatory Crackdown: The negative sentiment shock combined with higher rates led to a projected loss of -28%. The analysis showed that sentiment was a leading indicator, with its deterioration often preceding the price drop by 1-2 days. (3) Tech Sector Correction: NVDA's loss in this scenario was -25%, but its correlation with the Nasdaq jumped to 0.92, offering little diversification benefit during a broad tech selloff.

These stress tests reveal that NVDA is not just a "tech stock" but a highly specialized asset with unique, non-diversifiable risk factors tied directly to the AI and semiconductor ecosystems.

4.4. Portfolio Performance Comparison

The static and dynamic portfolios were evaluated over the 2023 out-of-sample period. The static 60/40 NVDA/S&P 500 portfolio delivered an annualized return of 118.5% with 38.2% volatility, a Sharpe ratio of 2.95, and a maximum drawdown of -42.1%. In contrast, the volatility-managed dynamic portfolio achieved a slightly lower return of 105.2% but significantly improved risk metrics: annualized volatility dropped to 29.7%, Sharpe ratio rose to 3.38, and max drawdown was reduced to -28.5%. Although it forwent some upside, the dynamic strategy provided superior risk-adjusted performance by lowering NVIDIA exposure ahead of major volatility spikes in Q2 and Q4 2023. This highlights the practical

value of integrating AI-driven volatility forecasts into active portfolio management to mitigate tail risk without drastically compromising returns.

5. Discussion

Our findings have several profound implications for portfolio risk management in the AI age. First, they empirically validate that traditional, static risk models are insufficient for assets like NVIDIA. The extreme skewness, kurtosis, and time-varying correlations demand more sophisticated, adaptive tools. Second, the success of our hybrid LSTM-GARCH model underscores the power of integrating different methodological paradigms. Deep learning captures the complex, non-linear price drivers (including sentiment), while the GARCH structure provides a robust, interpretable framework for volatility, a cornerstone of risk management.

The risk attribution analysis is particularly instructive. It moves beyond simple beta to identify the specific, idiosyncratic risk factors that dominate NVDA's risk profile: supply chain integrity, regulatory sentiment, and sector-wide tech momentum. For a portfolio manager, this means that hedging NVDA effectively requires instruments that are sensitive to these specific factors (e.g., semiconductor ETFs, volatility derivatives on the Nasdaq, or even direct short positions in key suppliers) rather than just a generic S&P 500 hedge.

The portfolio experiment provides a clear blueprint for practical implementation. A simple rule-based strategy that uses AI-generated volatility forecasts to dynamically adjust position sizing can materially improve the risk-return trade-off. This approach embodies the principle of "risk parity" at a granular level, ensuring that the portfolio's risk contribution from the volatile AI asset is managed proactively.

However, our study also highlights important caveats. AI models are not infallible oracles. They are trained on historical data and can be blindsided by truly unprecedented "black swan" events. Over-reliance on any single model is dangerous. Furthermore, the computational complexity and data requirements of such frameworks may be a barrier for smaller investment firms. There is also an inherent feedback loop risk: if too many market participants use similar AI models, it could potentially amplify market moves and create new forms of systemic risk.

6. Conclusion

The age of Artificial Intelligence has irrevocably transformed the financial landscape, creating a new paradigm where technological innovation is not only a source of alpha but also a generator of complex, non-linear, and often idiosyncratic risks. Nowhere is this duality more vividly illustrated than in the case of NVIDIA Corporation, a company whose strategic pivot to AI infrastructure catalyzed extraordinary shareholder returns while simultaneously exposing investors to unprecedented volatility and sector-specific vulnerabilities. This paper has systematically investigated how modern portfolio risk management must evolve to meet the challenges posed by such AI-driven equities.

Through the development and empirical validation of an integrated AI-augmented framework, centered on a hybrid LSTM-GARCH model enhanced with macroeconomic variables and real-time sentiment analytics, we have demonstrated that traditional risk models rooted in static correlations and Gaussian assumptions are fundamentally inadequate for assets at the epicenter of technological disruption. Our results confirm that NVIDIA's return dynamics exhibit pronounced fat tails, time-varying betas, and heightened sensitivity to exogenous shocks ranging from semiconductor supply chain bottlenecks to shifts in global AI policy. These characteristics defy capture by conventional Value-at-Risk or mean-variance optimization techniques.

Critically, our hybrid model outperformed established benchmarks in both point forecasting and, more importantly, conditional volatility prediction, the cornerstone of effective risk management. By leveraging the pattern recognition capabilities of deep

learning alongside the statistical rigor of econometric volatility modeling, we achieved forecasts that not only anticipated price direction but also quantified the uncertainty surrounding those predictions with remarkable accuracy. This dual capability enabled us to compute more reliable tail risk measures (VaR and CVaR) and to design scenario-based stress tests grounded in plausible, sector-specific narratives rather than generic market shocks.

The practical utility of our approach was further validated through a dynamic portfolio allocation experiment. A simple strategy that adjusted NVIDIA exposure based on our model's 30-day volatility forecast succeeded in materially improving the risk-return profile: while sacrificing a modest amount of absolute return, it delivered a significantly higher Sharpe ratio and, crucially, reduced maximum drawdown by over 13 percentage points compared to a static allocation. This underscores a key principle for the AI era: risk management is not about avoiding high-growth assets, but about intelligently modulating exposure in response to evolving risk conditions.

For institutional and retail investors alike, the implications are clear. Passive, backward-looking risk frameworks are no longer sufficient. Embracing AI as a core component of the risk management toolkit, rather than merely an investment theme, is essential for navigating the volatility inherent in high-tech markets. This requires investment in alternative data infrastructure, adoption of adaptive modeling techniques, and a cultural shift toward continuous model re-evaluation and human-AI collaboration.

The frontier of AI-powered risk management will likely be shaped by three emerging trends. First, the integration of causal machine learning could help distinguish true risk drivers from spurious correlations, enabling more robust counterfactual analysis. Second, LLMs will enhance the interpretation of unstructured data, from regulatory texts to earnings call transcripts, providing earlier and more nuanced signals of strategic or operational shifts. Third, as quantum computing advances, real-time optimization of large, multi-asset portfolios under complex, non-linear constraints may become feasible, further blurring the line between prediction and prescription.

Yet, this technological evolution carries its own systemic risks. The widespread adoption of similar AI-driven strategies across major financial institutions could lead to algorithmic herding, amplifying market dislocations during periods of stress. Therefore, responsible deployment demands not only technical sophistication but also robust governance, model interpretability, and regulatory oversight to ensure that AI enhances market stability rather than undermining it.

In conclusion, portfolio risk management in the age of AI is no longer a retrospective exercise in historical variance calculation. It is a proactive, adaptive, and forward-looking discipline that must continuously learn from an ever-expanding universe of data and anticipate the next wave of disruption. Our research provides a concrete, empirically validated framework for achieving this goal. By marrying the predictive power of artificial intelligence with the foundational principles of financial econometrics, investors can better harness the immense opportunities of the AI revolution while safeguarding their portfolios against its inherent, and often asymmetric, risks. The future of finance belongs not to those who merely invest in AI, but to those who intelligently manage risk with it.

References

1. C. Challoumis, "The dawn of artificial intelligence," In *XIX International Scientific Conference. London. Great Britain, 2024*, pp. 169-205.
2. J. Wang, J. Hsu, and Z. Qin, "A Comprehensive Analysis of Nvidia's Technological Innovations, Market Strategies, and Future Prospects," *International Journal of Information Technologies and Systems Approach (IJITSA)*, vol. 17, no. 1, pp. 1-16, 2024. doi: 10.4018/ijitsa.344423
3. D. Kaluge, "Future Perspective: The Role of Artificial Intelligence in Predicting Investor Behavior: Opportunities and Challenges," *Digital Transformation and Inclusive Economic Development in Indonesia*, vol. 56, 2023. doi: 10.11594/futscipress3
4. N. M. Seel, "Model-based learning: A synthesis of theory and research," *Educational Technology Research and Development*, vol. 65, no. 4, pp. 931-966, 2017. doi: 10.1007/s11423-016-9507-9

5. K. de Cuveland, E. Choi, and M. Shin, "Social dynamics of hyperbole-social media sentiment analysis for hype detection in emerging technologies: a conceptual approach," *Technology Analysis & Strategic Management*, pp. 1-15, 2025.
6. H. Harchane, K. E. Bouchti, M. Hagouch, C. Elasri, F. Saoiabi, and S. Ziti, "The Impact of Artificial Intelligence on Risk Management and Investment Portfolios," In *International Conference on Smart Business Technologies*, July, 2024, pp. 228-238. doi: 10.1007/978-3-031-86698-2_21
7. P. H. Vuong, L. H. Phu, T. H. Van Nguyen, L. N. Duy, P. T. Bao, and T. D. Trinh, "A comparative study of deep learning approaches for stock price prediction," *Digital Finance*, pp. 1-29, 2025.
8. K. Kaur, Y. Kumar, and S. Kaur, "Artificial Intelligence and Machine Learning in Financial Services to Improve the Business System," In *Computational Intelligence for Modern Business Systems: Emerging Applications and Strategies*, 2023, pp. 3-30. doi: 10.1007/978-981-99-5354-7_1
9. A. Kurani, P. Doshi, A. Vakharia, and M. Shah, "A comprehensive comparative study of artificial neural network (ANN) and support vector machines (SVM) on stock forecasting," *Annals of Data Science*, vol. 10, no. 1, pp. 183-208, 2023.
10. K. Li, and Y. Zhou, "Improved financial predicting method based on time series long short-term memory algorithm," *Mathematics*, vol. 12, no. 7, p. 1074, 2024. doi: 10.3390/math12071074
11. D. Likitracharoen, and L. Suwannamalik, "Assessing financial stability in turbulent times: A study of generalized autoregressive conditional heteroskedasticity-type value-at-risk model performance in Thailand's transportation sector during COVID-19," *Risks*, vol. 12, no. 3, p. 51, 2024. doi: 10.3390/risks12030051
12. B. Wang, Y. Zhang, J. Li, and T. Liu, "Assessing Expected Shortfall in Risk Analysis Through Generalized Autoregressive Conditional Heteroskedasticity Modeling and the Application of the Gumbel Distribution," *Axioms*, vol. 14, no. 5, p. 391, 2025. doi: 10.3390/axioms14050391
13. H. T. Araya, J. Aduda, and T. Berhane, "A hybrid garch and deep learning method for volatility prediction," *Journal of Applied Mathematics*, vol. 2024, no. 1, p. 6305525, 2024. doi: 10.1155/2024/6305525
14. P. Zhao, H. Zhu, W. S. H. Ng, and D. L. Lee, "From GARCH to neural network for volatility forecast," In *Proceedings of the AAAI Conference on Artificial Intelligence*, March, 2024, pp. 16998-17006.
15. C. Saravanos, and A. Kanavos, "Forecasting stock market volatility using social media sentiment analysis," *Neural Computing and Applications*, vol. 37, no. 17, pp. 10771-10794, 2025. doi: 10.1007/s00521-024-10807-w

Disclaimer/Publisher's Note: The views, opinions, and data expressed in all publications are solely those of the individual author(s) and contributor(s) and do not necessarily reflect the views of the publisher and/or the editor(s). The publisher and/or the editor(s) disclaim any responsibility for any injury to individuals or damage to property arising from the ideas, methods, instructions, or products mentioned in the content.