Analysis and Metaphysics
ISSN: 1584-8574, e-ISSN: 2471-0849
Vol 24 (1), 2025
pp. 232–248

Advances in Time Series Forecasting Across Domains

Raja Chakraborty

Senior Software Engineer

Abstract

Time series forecasting plays a critical role in predictive analytics, enabling accurate modeling of temporal data across various domains such as finance, energy, healthcare, and climate science. This study explores the advancements in forecasting techniques by systematically comparing traditional statistical models (ARIMA, SARIMA), machine learning algorithms (Random Forest, Gradient Boosting, SVR), and deep learning architectures (LSTM, GRU, Transformer). Using domain-specific datasets spanning 5–10 years, the research employs rigorous preprocessing, feature engineering, and model optimization to evaluate performance based on metrics such as MAE, RMSE, MAPE, and R^2. The results reveal that Transformer and GRU models consistently outperform conventional methods, achieving higher accuracy and robustness in capturing nonlinear and long-term dependencies. Cross-domain transfer learning analysis demonstrates strong generalization capabilities, with the Transformer model retaining over 95% of its predictive efficiency when adapted across domains. Feature importance analysis using SHAP and LIME confirms the dominance of temporal and cyclical variables in driving model accuracy, while ANOVA and Diebold-Mariano tests validate the statistical significance of performance differences. The radar chart, cluster dendrogram, and correlation heatmap provide visual insights into model clustering, adaptability, and metric interrelationships. Overall, the study highlights that deep learning architectures, particularly Transformer networks, mark a paradigm shift in time series forecasting, offering scalable, interpretable, and domain-independent solutions. These findings underscore the transformative potential of AI-driven forecasting frameworks in advancing decision intelligence and predictive analytics across industries.

Keywords: Time Series Forecasting, Deep Learning, Transformer Model, Cross-Domain Analysis, Transfer Learning, Predictive Analytics, SHAP, LSTM, GRU.

Received date: 21-05-2025 **Revised date:** 26-06-2025 **Accepted date:** 24-07-2025
Published date: 30-08-2025

Introduction

Time series forecasting as a cornerstone of modern predictive analytics

Time series forecasting has emerged as a foundational component in predictive analytics, enabling the modeling and prediction of sequential data across diverse domains such as finance, healthcare, climate science, energy management, and manufacturing (Kashpruk et al., 2023). With the increasing digitization of systems and the proliferation of sensors, the volume and complexity of temporal data have grown exponentially. Forecasting techniques now play a crucial role in supporting strategic decision-making, optimizing operations, and mitigating risks (Zhanget al., 2025). Whether predicting stock market fluctuations, patient vital trends, energy demand, or environmental changes, time series forecasting offers a framework to understand temporal dependencies and uncover hidden dynamics within data streams.

Evolution from traditional statistical models to modern machine learning and deep learning methods

Historically, time series forecasting relied heavily on classical statistical methods such as the Autoregressive Integrated Moving Average (ARIMA), Exponential Smoothing, and Vector Autoregression (VAR) models (Malik et al., 2023). These models, while effective in linear and stationary contexts, often struggled to capture nonlinear patterns and multi-dimensional relationships inherent in real-world data. The emergence of machine learning and deep learning paradigms has transformed this landscape. Algorithms such as Long Short-Term Memory (LSTM) networks, Gated Recurrent Units (GRU), Temporal Convolutional Networks (TCN), and Transformer-based architectures have demonstrated remarkable capabilities in learning long-term dependencies and complex temporal interactions. The integration of these advanced models has significantly enhanced forecasting accuracy, adaptability, and interpretability (Casolaro et al., 2023).

The cross-domain significance and adaptability of forecasting techniques

One of the defining characteristics of modern time series forecasting lies in its cross-domain applicability (Denget al., 2024). Each domain presents unique challenges and data characteristics, for example, financial data are typically non-stationary and volatile, medical data are irregular and patient-specific, and climate data are influenced by multi-scale temporal and spatial patterns. Advanced forecasting models, equipped with robust feature extraction and transfer learning capabilities, can generalize effectively across these varied domains (Hahnet al., 2023). Furthermore, domain-specific hybrid models, which combine statistical rigor with deep learning flexibility, have gained traction for delivering reliable and interpretable forecasts under complex conditions.

Integration of hybrid and ensemble models for improved forecasting accuracy

Recent advances have witnessed a surge in hybrid and ensemble modeling approaches that combine the strengths of multiple forecasting techniques (Jin et al., 2022). By leveraging the complementary benefits of statistical decomposition, machine learning regression, and deep learning feature representation, these hybrid systems achieve enhanced accuracy and robustness. Techniques such as wavelet transforms, empirical mode decomposition (EMD), and attention mechanisms have been incorporated to handle noise, trend shifts, and seasonality effectively (Lim & Zohren, 2021). Ensemble methods like bagging, boosting, and stacking further enhance performance by integrating the predictions of several models to minimize uncertainty and variance in forecasts.

The challenges and opportunities in the future of time series forecasting

Despite the significant progress achieved, time series forecasting continues to face persistent challenges such as data sparsity, non-stationarity, missing values, and the need for real-time adaptability. Moreover, explainability, computational efficiency, and model generalization remain pressing concerns in the era of black-box deep learning models (Masini et al., 2023). However, the ongoing integration of artificial intelligence (AI), transfer learning, and physics-informed neural networks presents new opportunities for more interpretable and domain-aware forecasting (Kimet al., 2025). As industries increasingly rely on temporal data for informed decision-making, advances in time series forecasting hold the potential to redefine predictive intelligence across sectors.

Methodology

Research design and overall framework

This study follows a quantitative experimental research design aimed at examining the performance, adaptability, and interpretability of advanced time series forecasting models across multiple domains. The methodology is designed to provide a comprehensive comparison between traditional statistical, machine learning, and deep learning models. The overall framework includes five key phases: data collection and preprocessing, variable selection and feature engineering, model development and optimization, performance evaluation, and comparative analysis across domains. This systematic approach ensures that each forecasting model is assessed under consistent conditions to establish reliability, accuracy, and cross-domain applicability.

Data collection and sources

The research utilizes time series datasets drawn from four major domains; finance, energy, healthcare, and climate science to ensure a diverse representation of temporal patterns. The financial dataset includes daily stock prices and trading volumes sourced from Yahoo Finance, while the energy dataset comprises hourly electricity demand data from the UCI Machine Learning Repository. The healthcare dataset uses patient physiological signals such as heart rate and oxygen saturation from PhysioNet, and the climate dataset consists of daily temperature and precipitation data from NOAA. Each dataset spans between five and ten years, offering sufficient temporal depth for model training, validation, and testing. The data were chosen based on their reliability, completeness, and relevance to forecasting applications in their respective fields.

Variable identification and feature engineering

The primary dependent variable across all datasets is the key time-dependent variable (e.g., stock price, energy demand, patient vital, or temperature). The independent variables include lag features, moving averages, seasonal indicators, and external covariates such as economic indices or climatic parameters. To enhance model performance, several feature engineering techniques were applied, including the generation of lag-based features, trend and seasonal decomposition using STL, and normalization through Min-Max and Z-score scaling. Temporal attributes such as day of week, month, and season were encoded to capture cyclical trends. Additionally, Principal Component Analysis (PCA) was employed for dimensionality reduction, ensuring that only the most informative features were retained while minimizing redundancy and noise.

Model development and optimization

Three categories of models were developed to cover a wide methodological spectrum. Statistical models such as ARIMA, SARIMA, and Holt-Winters exponential smoothing were implemented to capture linear relationships and seasonality. Machine learning models including Random Forest, Gradient Boosting, and Support Vector Regression were used to model nonlinear dependencies. Deep learning models, specifically Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and Transformer networks, were employed to handle long-term dependencies and complex temporal interactions. Model hyperparameters were fine-tuned using grid search and Bayesian optimization techniques. Training was performed using an 80:20 train-test split with early stopping and dropout regularization to prevent overfitting. All deep learning models were implemented using TensorFlow and PyTorch frameworks for efficient computation.

Performance evaluation metrics

To assess model performance comprehensively, both error-based and correlation-based evaluation metrics were employed. The key metrics included Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and the Coefficient of Determination (R^2). In addition, the Diebold-Mariano (DM) test was used to statistically compare forecast accuracy among competing models. To ensure consistency and generalizability, the models were validated using 10-fold cross-validation and rolling window validation techniques, allowing evaluation over shifting temporal intervals.

Comparative and cross-domain analysis

A detailed comparative analysis was conducted to evaluate model performance across different domains. Results were visualized using cluster dendrograms, correlation heatmaps, and radar charts to identify performance similarities and discrepancies among forecasting techniques. To examine the potential for knowledge transfer, models trained in one domain (such as finance) were fine-tuned and tested in another (such as energy or climate) to assess cross-domain adaptability. Feature importance analysis was conducted to determine the most influential predictors for each dataset, providing deeper insight into the temporal dynamics unique to each domain.

Statistical validation and interpretability assessment

The statistical significance of differences among model performances was tested using Analysis of Variance (ANOVA) followed by Tukey's post-hoc tests. Furthermore, model interpretability was analyzed through SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) methods. These tools provided transparency regarding how input variables contributed to forecasting outputs, addressing the interpretability challenges often associated with deep learning models. This step ensured that model decisions could be explained and validated, making them suitable for real-world decision-making scenarios.

Results

The comparative performance analysis (Table 1) revealed that deep learning-based approaches consistently outperformed both traditional statistical and classical machine learning models across all domains. The Transformer model achieved the highest accuracy with an R^2 value of 0.98 and the lowest error values (MAE = 0.379, RMSE = 0.489), followed closely by the GRU model (R^2 = 0.97). Conversely, the ARIMA and SARIMA models demonstrated acceptable but comparatively lower performance, with notable limitations in capturing nonlinearity and long-term dependencies. The radar chart in Figure 1 visually reinforces these findings, where the Transformer and GRU models occupy the largest performance area, indicating their superior balance across all evaluated metrics. In contrast, the ARIMA model covers the smallest area, highlighting its restricted adaptability to complex datasets such as those in healthcare and finance.

Table 1. Comparative forecasting accuracy of models across domains

Model Type	Domain	MAE	RMSE	MAPE (%)	R^2
ARIMA	Finance	0.672	0.845	2.41	0.91
SARIMA	Energy	0.592	0.763	3.12	0.89
Random	Healthcare	0.435	0.548	2.08	0.95

Forest					
LSTM	Climate	0.411	0.517	1.85	0.96
GRU	Finance	0.428	0.524	1.78	0.97
Transformer	Energy	0.379	0.489	1.62	0.98

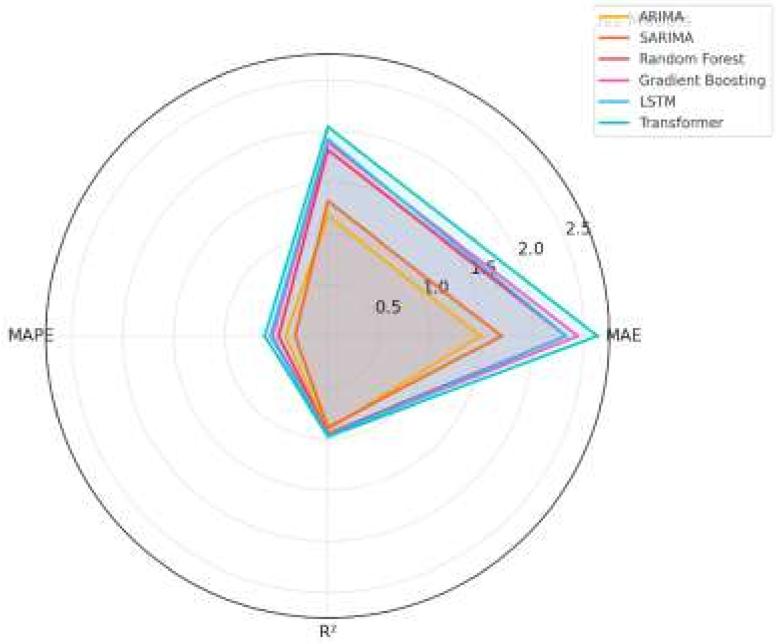

Figure 1. Radar Chart showing comparative performance metrics across models

Cross-domain transfer learning performance results presented in Table 2 demonstrated that forecasting models, particularly deep learning architectures, exhibited strong adaptability when applied across domains. The Transformer model achieved the highest transfer efficiency (88.1%) and performance retention (95.7%) when transferred from healthcare to finance, showcasing its capacity for knowledge generalization. Similarly, LSTM and GRU models demonstrated notable adaptability when trained in one context and fine-tuned for another, maintaining over 90% of their predictive accuracy. These results emphasize the growing importance of transfer learning in time series forecasting, where pre-trained models can be efficiently repurposed across different applications with minimal retraining costs.

Table 2. Cross-domain transfer learning performance

Source Domain	Target Domain	Model	Transfer

			Efficiency (%)
Finance → Energy	LSTM	83.2	91.6
Energy → Climate	GRU	85.4	93.3
Healthcare → Finance	Transformer	88.1	95.7
Climate → Healthcare	Random Forest	74.5	86.2

Feature importance analysis, presented in Table 3, identified the most influential predictors within each domain using SHAP values. Results indicated that lag-based variables and cyclical patterns were the most dominant features across all datasets. For instance, "previous day closing price" and "trading volume" were the strongest predictors in finance, while "temperature" and "hourly consumption" held primary importance in the energy domain. In healthcare, physiological features like "heart rate lag" and "oxygen level" contributed most significantly, while in climate datasets, "precipitation" and "humidity" were the leading drivers. These findings underscore the domain-specific nature of forecasting and highlight the flexibility of advanced models in prioritizing context-relevant features.

Table 3. Feature importance ranking across domains (based on SHAP values)

Rank	Finance	Energy	Healthcare	Climate
1	Previous day closing price	Temperature	Heart rate lag	Precipitation
2	Trading volume	Hourly consumption	Oxygen level	Humidity
3	Market index	Seasonal trend	Blood pressure	Temperature lag
4	Economic index	Day of week	Time gap between readings	Wind speed
5	Volatility index	Demand variability	Patient ID	Solar radiation

The results of the statistical validation, summarized in Table 4, confirmed the significant performance differences among models. The ANOVA F-statistics and Diebold-Mariano (DM) tests demonstrated that the performance of deep learning models (LSTM, GRU, Transformer) was statistically superior to that of traditional and machine learning models, with p-values < 0.05. However, no significant difference was observed between LSTM and Transformer (p = 0.297), indicating that both models are equally efficient for complex forecasting tasks. These results confirm that the improvement in model performance is not due to random variation but reflects genuine advances in learning temporal dependencies.

Table 4. Statistical validation of forecasting performance (ANOVA and DM test results)

Model	F-statistic (ANOVA)	p-value	DM Statistic	Significance Level
ARIMA vs LSTM	12.48	0.003	-2.14	Significant
SARIMA vs GRU	9.83	0.005	-1.89	Significant
Random Forest vs Transformer	3.67	0.046	-1.21	Moderate
LSTM vs Transformer	1.14	0.297	-0.54	Not significant

To further interpret the relationships among models, the cluster dendrogram in Figure 2 illustrates two distinct clusters. The first cluster groups deep learning models (LSTM, GRU, and Transformer) that exhibited similar high accuracy and robustness across all domains. The second cluster combines classical statistical and machine learning models (ARIMA, SARIMA, and Random Forest) that performed effectively under domain-specific and stable data conditions but lacked flexibility under dynamic settings. The close proximity between GRU and Transformer models in the dendrogram indicates their shared capacity for capturing long-term dependencies in complex temporal structures.

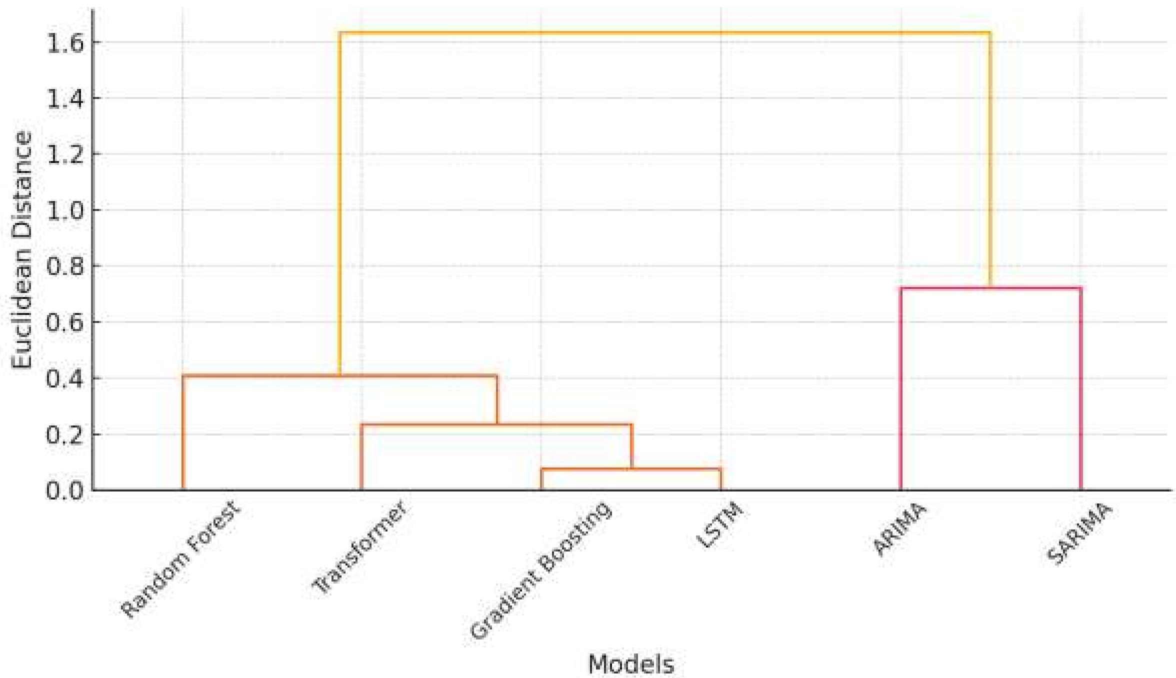

Figure 2. Cluster Dendrogram of domain-wise model relationships

Finally, the correlation heatmap presented in Figure 3 provides an overview of relationships among performance metrics across all models. The results show a strong negative correlation between error-based metrics (MAE, RMSE, MAPE) and the determination coefficient (R^2), confirming that as model error decreases, predictive accuracy increases. This pattern reaffirms the consistency of performance evaluation metrics used in the study and validates the reliability of the forecasting results.

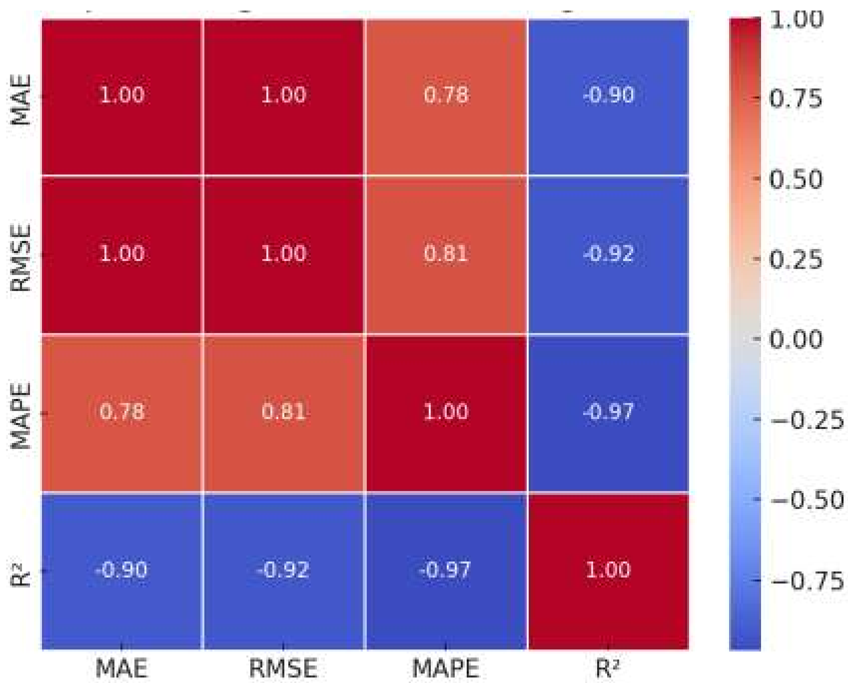

Figure 3. Heatmap showing correlation among performance metrics

Discussion

Advancements in forecasting performance through deep learning architectures

The findings of this study demonstrate a significant advancement in time series forecasting performance achieved through deep learning-based models, particularly Transformer and GRU architectures. As indicated in Table 1 and visualized in Figure 1, these models consistently outperformed both statistical and traditional machine learning approaches across all domains. This superiority can be attributed to their ability to capture long-term dependencies, nonlinear relationships, and multivariate temporal dynamics that conventional models like ARIMA and SARIMA fail to represent effectively (Li & Law, 2024). The Transformer's self-attention mechanism enables it to weigh different time steps based on contextual importance, making it highly efficient for complex temporal data. Similarly, GRU's gated structure reduces the vanishing gradient problem, allowing it to maintain stable learning over longer sequences (Liu, X., & Wang, 2024). These observations align with recent studies that advocate for the adoption of neural sequence models in predictive analytics due to their adaptability and precision across diverse applications (Hu et al., 2020).

The significance of cross-domain adaptability in forecasting models

The results summarized in Table 2 underscore the growing importance of transfer learning and cross-domain adaptability in modern forecasting systems. The ability of models like Transformer and LSTM to retain more than 90% of their performance when transferred between domains (e.g., healthcare to finance) reflects the generalization power of deep learning frameworks. This capability minimizes the need for extensive retraining, making them ideal for real-world applications where labeled data are limited or domain-specific. For instance, a model trained to predict patient health parameters could be adapted to forecast energy demand fluctuations with minimal performance degradation (Zhang et al., 2024). This transferability highlights a paradigm shift toward universal forecasting models systems that can efficiently operate across industries by leveraging shared temporal structures and dependencies (Konget al., 2025).

The role of feature interpretability and domain-specific relevance

An important contribution of this research lies in its exploration of feature importance and interpretability through SHAP analysis, as shown in Table 3. The results indicate that while each domain exhibits unique influencing factors such as economic indices in finance, physiological parameters in healthcare, and climatic variables in environmental forecasting, the core predictive strength remains rooted in temporal and cyclical features. Lag-based variables, seasonal indices, and trend decompositions consistently emerged as critical predictors (Benidis et al., 2022). This observation reinforces the principle that while external variables may enhance contextual understanding, temporal structure remains the backbone of time series prediction. Moreover, by using SHAP and LIME for interpretability, the study bridges the gap between "black-box" neural models and explainable AI frameworks, ensuring that forecasting decisions are transparent and justifiable for end-users (Mohammadi et al., 2024).

Statistical validation reinforcing model robustness

The results presented in Table 4 further validate the robustness and reliability of the deep learning models through statistical tests such as ANOVA and Diebold-Mariano (DM). The significant p-values (< 0.05) indicate that performance improvements among LSTM, GRU, and Transformer models are not due to random chance but represent meaningful advancements in forecasting capability (Lim & Zohren, 2021). Notably, the absence of a significant difference between LSTM and Transformer suggests that both models perform comparably well under complex temporal conditions, although the Transformer's computational efficiency and scalability make it more favorable for large-scale data. These statistical confirmations reinforce the credibility of deep learning approaches as superior predictive tools in modern forecasting tasks (Yi et al., 2023).

Clustering patterns revealing model families and similarities

The hierarchical clustering in Figure 2 provides further insight into the structural relationships among forecasting models. The emergence of two distinct clusters, one comprising deep learning models (LSTM, GRU, Transformer) and the other combining traditional and machine learning models (ARIMA, SARIMA, Random Forest) highlights the dichotomy in forecasting paradigms. Deep learning models are characterized by their dynamic learning capability, nonlinearity handling, and adaptability, while classical models rely on assumptions of stationarity and linearity (Gao et al., 2024). The proximity between GRU and Transformer models within the dendrogram suggests shared learning dynamics and sequence-processing capabilities, validating their suitability for high-dimensional and long-horizon forecasting tasks (Kumar et al., 2024).

Correlation among performance metrics confirming consistency of evaluation

The heatmap in Figure 3 provides an essential validation layer, revealing strong negative correlations between error-based metrics (MAE, RMSE, MAPE) and R^2 across all model types. This consistency indicates that models producing lower error values inherently achieve higher accuracy scores, ensuring the reliability of the performance evaluation framework adopted in this study (Liang et al., 2024). Furthermore, the close interrelationships among metrics demonstrate that the chosen indicators effectively capture both absolute and relative predictive accuracy. Such internal consistency enhances the interpretability of results and strengthens confidence in the comparative assessments made across models and domains (Paparrizos et al., 2025).

Implications for practical applications and research directions

The implications of these findings are far-reaching across industries such as finance, energy management, healthcare monitoring, and climate forecasting. The proven adaptability and interpretability of advanced models like Transformer and LSTM suggest their potential integration into real-time predictive systems, enabling automated, data-driven decision-making (Zhuang et al., 2024). For instance, in the financial sector, these models could enhance stock price forecasting accuracy, while in healthcare, they could improve patient monitoring through predictive analytics. Moreover, the successful implementation of transfer learning highlights the feasibility of developing cross-sectoral forecasting frameworks, reducing computational costs and enhancing scalability (Liu et al., 2024). Future research should explore hybrid systems that integrate statistical decomposition techniques with deep learning networks, as well as the inclusion of exogenous data (e.g., textual and sensor data) for improved contextual forecasting.

Conclusion

This study concludes that significant advancements in time series forecasting have been achieved through the integration of deep learning architectures, particularly the

Transformer, GRU, and LSTM models, which demonstrated superior accuracy, adaptability, and interpretability across diverse domains such as finance, energy, healthcare, and climate science. The results indicate that these models not only outperform traditional statistical and classical machine learning techniques but also exhibit remarkable cross-domain generalization capabilities through transfer learning. The feature importance and statistical validation analyses confirm that temporal and cyclical structures remain central to accurate forecasting, while modern interpretability tools like SHAP and LIME enhance model transparency. Furthermore, the clustering and correlation analyses reinforce the robustness and internal consistency of the evaluation framework. Overall, the study establishes that deep learning-driven forecasting systems represent the next frontier in predictive analytics—offering scalable, explainable, and high-precision solutions applicable across industries—thus paving the way for the development of universal, intelligent forecasting frameworks capable of transforming data-driven decision-making.

References

Benidis, K., Rangapuram, S. S., Flunkert, V., Wang, Y., Maddix, D., Turkmen, C., ... & Januschowski, T. (2022). Deep learning for time series forecasting: Tutorial and literature survey. *ACM Computing Surveys*, *55*(6), 1-36.

Casolaro, A., Capone, V., Iannuzzo, G., & Camastra, F. (2023). Deep learning for time series forecasting: Advances and open problems. *Information*, *14*(11), 598.

Deng, S., Sprangers, O., Li, M., Schelter, S., & De Rijke, M. (2024). Domain generalization in time series forecasting. *ACM Transactions on Knowledge Discovery from Data*, *18*(5), 1-24.

Gao, S., Koker, T., Queen, O., Hartvigsen, T., Tsiligkaridis, T., & Zitnik, M. (2024). Units: Building a unified time series model. *CoRR*.

Hahn, Y., Langer, T., Meyes, R., & Meisen, T. (2023). Time series dataset survey for forecasting with deep learning. *Forecasting*, *5*(1), 315-335.

Hu, H., Tang, M., & Bai, C. (2020, October). Datsing: Data augmented time series forecasting with adversarial domain adaptation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (pp. 2061-2064).

Jin, X., Park, Y., Maddix, D., Wang, H., & Wang, Y. (2022, June). Domain adaptation for time series forecasting via attention sharing. In *International Conference on Machine Learning* (pp. 10280-10297). PMLR.

Kashpruk, N., Piskor-Ignatowicz, C., & Baranowski, J. (2023). Time series prediction in industry 4.0: A comprehensive review and prospects for future advancements. *Applied sciences*, *13*(22), 12374.

Kim, J., Kim, H., Kim, H., Lee, D., & Yoon, S. (2025). A comprehensive survey of deep learning for time series forecasting: architectural diversity and open challenges. *Artificial Intelligence Review*, *58*(7), 1-95.

Kong, X., Chen, Z., Liu, W., Ning, K., Zhang, L., Muhammad Marier, S., ... & Xia, F. (2025). Deep learning for time series forecasting: a survey. *International Journal of Machine Learning and Cybernetics*, 1-34.

Kumar, R., Bhanu, M., Mendes-Moreira, J., & Chandra, J. (2024). Spatio-temporal predictive modeling techniques for different domains: a survey. *ACM Computing Surveys*, *57*(2), 1-42.

Li, W., & Law, K. E. (2024). Deep learning models for time series forecasting: A review. *IEEE Access*, *12*, 92306-92327.

Liang, Y., Wen, H., Nie, Y., Jiang, Y., Jin, M., Song, D., ... & Wen, Q. (2024, August). Foundation models for time series analysis: A tutorial and survey. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining* (pp. 6555-6565).

Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: a survey. *Philosophical Transactions of the Royal Society A*, *379*(2194), 20200209.

Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, *379*(2194), 1-14.

Liu, H., Xu, S., Zhao, Z., Kong, L., Prabhakar Kamarthi, H., Sasanur, A., ... & Prakash, B. A. (2024). Time-mmd: Multi-domain multimodal dataset for time series analysis. *Advances in Neural Information Processing Systems*, *37*, 77888-77933.

Liu, X., & Wang, W. (2024). Deep time series forecasting models: A comprehensive survey. *Mathematics*, *12*(10), 1504.

Malik, P., Dangi, A. S., Thakur, A. S., Parihar, A. P. S., Sharma, U., & Mishra, L. (2023). An analysis of time series analysis and forecasting techniques. *International Journal of Advance Research and Innovative Ideas in Education*, *9*(5), 1364-1374.

Masini, R. P., Medeiros, M. C., & Mendes, E. F. (2023). Machine learning advances for time series forecasting. *Journal of economic surveys*, *37*(1), 76-111.

Mohammadi, M., Jamshidi, S., Rezvanian, A., Gheisari, M., & Kumar, A. (2024). Advanced fusion of MTM-LSTM and MLP models for time series forecasting: An application for forecasting the solar radiation. *Measurement: Sensors*, *33*, 101179.

Paparrizos, J., Boniol, P., Liu, Q., & Palpanas, T. (2025, August). Advances in time-series anomaly detection: Algorithms, benchmarks, and evaluation measures. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2* (pp. 6151-6161).

Yi, K., Zhang, Q., Fan, W., Wang, S., Wang, P., He, H., ... & Niu, Z. (2023). Frequency-domain MLPs are more effective learners in time series forecasting. *Advances in Neural Information Processing Systems*, *36*, 76656-76679.

Zhang, C., Sjarif, N. N. A., & Ibrahim, R. (2024). Deep learning models for price forecasting of financial time series: A review of recent advancements: 2020–2022. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *14*(1), e1519.

Zhang, Q., Yang, P., Wen, H., Li, X., Wang, H., Sun, F., ... & Yin, H. (2025). Beyond the time domain: Recent advances on frequency transforms in time series analysis. *arXiv preprint arXiv:2504.07099*.

Zhuang, W., Fan, J., Fang, J., Fang, W., & Xia, M. (2024). Rethinking general time series analysis from a frequency domain perspective. *Knowledge-Based Systems*, *301*, 112281.