International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246/ ICRTCSIT-111
Eureka Vision Publication | ICRTCSIT'25-Conference Proceeding

Original Article

Enhancing Time Series Forecasting using DRL with attention-based neural architectures

Srinivasa Kalyan Vangibhurathachhi
Solution Architect, Texas, USA

Abstract - Time series forecasting (TSF) is key to decision-making in finance, retail, supply chain management, healthcare, climate among others where accurate predictions inform resource allocation, risk management and strategic planning. While traditional statistical models such as ARIMA handle linear dependencies, they fail to capture the nonlinear and multivariate complexities of modern datasets. Deep learning models such as RNNs, LSTMs, GRUs, CNNs, and Transformers, have advanced forecasting accuracy by capturing temporal patterns and cross-variable interactions. However, these models are static and unable to adapt dynamically to regime shifts, shocks or evolving trends once trained. In addressing this gap, deep reinforcement learning (DRL) offers adaptivity by treating forecasting as sequential decision-making where agents iteratively refine predictions through reward feedback. Attention mechanisms further enhance interpretability and accuracy by highlighting critical time steps and features. This white paper critically reviewed DL and DRL models for multivariate TSF and evaluated their application in finance, retail, supply chains, climate forecasting and healthcare using research studies and datasets. Case studies demonstrate that attention-LSTM and Transformer variants outperform traditional deep models while hybrid DRL–DL approaches achieve greater adaptability. A proposed hybrid architecture integrates attention-based forecasting with DRL agents to combine predictive accuracy, adaptive learning, and interpretability. Although challenges on data, model structure and tasks remain, the approach has the potential to transform TSF into adaptive and decision-support systems across domains.

***Keywords** - TSF, Deep Learning, DRL, Attention Based Neural Architectures.*

I. Introduction

In many industries such as finance, retail, supply chain, healthcare and climate, time series forecasting (TSF) is a critical tool for estimating future values based on recorded historical data of their relevance (Arushana et al., 2024). Precisely, time series are used to study how certain measures like air pollution data, electricity consumption or ozone concentration evolve over time. Accurate forecasts enable informed decision-making in terms of predicting stock prices, product demands, weather patterns, patient vitals, or energy loads (Casolaro et al., 2023). With accurate forecasting, businesses and organisations can make effective decisions, maximise resource utilisation and develop effective strategies. According to Thota (2025), TSF mainly relies on traditional statistical models like exponential smoothing and Auto Regressive Integrated Moving Average (ARIMA). However, these models struggle with non-linear and complex, large multidimensional datasets. In recent years, Madhulatha and Ghori (2025) observed that deep learning models specifically Recurrent Neural Networks (RNNs) and their variants Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs) have contributed immensely to TSF by capturing complex temporal dependencies with superior accuracy. However, even these latest deep learning models are trained in static supervised manner and lack mechanisms to adapt their predictions as new data arrives or conditions change (Madhulatha & Ghori, 2025). This is a significant limitation in dynamic environments where statistical properties change over time. Figure 1 below summarises classical and advanced deep learning models for TSF.

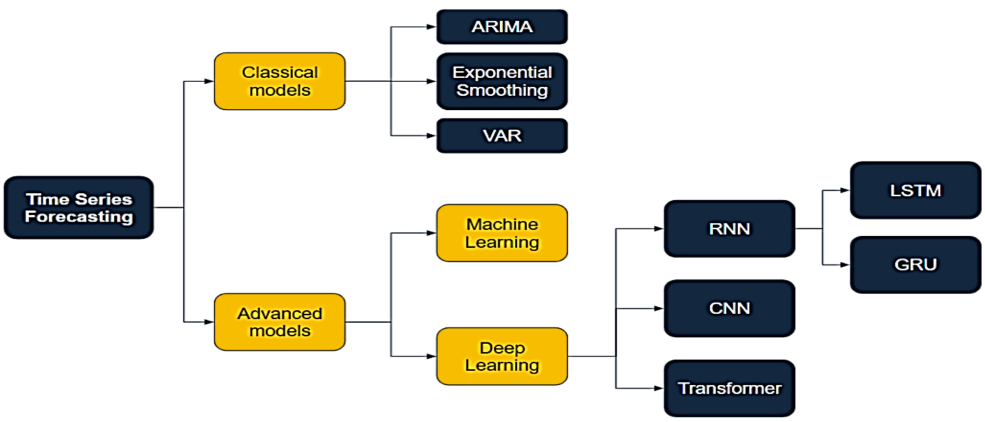

Figure 1. Advanced TSF Methods (Andres, 2023)

Due to limitations of deep learning models above, scholars have proposed deep reinforcement learning (DRL) which promises to introduce adaptivity into forecasting models. In DRL, an agent learns to make decisions through trial-and-error interactions with an environment and guided by reward feedback (Moreira et al., 2020). If forecasting is treated as sequential decision-making problem, learning agent can iteratively update its forecasting strategy based on reward signals such as forecast accuracy**.** Apart from DRL models, attention-based mechanisms represent another breakthrough in sequence modelling that could further enhance forecasting. By allowing neural networks to focus on the most relevant parts of input sequence, Zhang et al. (2024) noted that attention models have improved performance in natural language processing and vision tasks and is increasingly applied in TSF. Transformer architectures which rely entirely on self-attention have achieved incredible results in many forecasting benchmarks as they can capture long-range dependencies effectively (Vaswani, 2017).

Against this background, this whitepaper critically explores how DRL combined with attention-based neural architectures can enhance TSF. The research is guided by the following key objectives;

- To identify various deep learning and DRL models for multivariate TSF;
- To evaluate the performance of deep learning architectures, DRL models and attention based models using real-world case examples from financial, retail, climate and healthcare domains;
- To design a hybrid architecture that integrates attention mechanisms into DRL models for multivariate TSF.

2. Problem Statement

According to Oluwagbade (2025), the main challenge when analysing time-series data is its forecasting of future values based on past measurements. While modern deep learning methods have greatly improved forecasting accuracy over classical models, they often operate as one-shot predictors that do not adapt once trained. In practical scenarios, time series data can exhibit regime shifts, sudden shocks, or evolving patterns (market crashes, demand surges, climate anomalies). No matter how complex a static model is, it may fail under such changing conditions because its parameters are fixed after training (Arushana et al., 2024). The lack of iterative refinement based on new data presents key limitation. For instance, an LSTM trained on past data will continue to make the same errors if underlying pattern changes, unless it is retrained from scratch which is costly and slow for real-time applications To address these issues, DRL introduces feedback loop between predictions and outcomes. By formulating forecasting as sequential decision problem, an RL agent can receive a reward for each prediction (Smith, 2021). For example, a reward could be a forecast error so that higher reward corresponds to lower error.

The agent's goal is then to maximize cumulative reward that aligns with minimizing forecast error over time. Unlike traditional training, the agent continues to learn during deployment, updating its policy as new rewards come in. This enables adaptive learning where the model can correct itself and improve in the face of evolving data patterns. Research studies in financial forecasting have shown the promise of this approach. For instance, the hybrid LSTM-DQN model by Madhulatha and Ghori (2025) learned to adjust its predictions in response to real-time market fluctuations thus yielding significantly lower error than pure LSTM and other non-adaptive models. However, integrating DRL into forecasting is non-trivial. Designing the state representation, action space, and reward function requires care as the agent must effectively perceive time series history (state) and output forecasts or related decisions (actions) in a way that leads to meaningful learning as shown in figure 1 (Terven, 2025). Besides, DRL are sample-inefficient and can be unstable when issues like balancing exploration versus exploitation arise. These challenges call for adoption of attention mechanisms that can help focus DRL agent's capacity on the most informative parts of large multivariate input, potentially reducing noise, improving learning efficiency and interpretability (Ke, 2020).

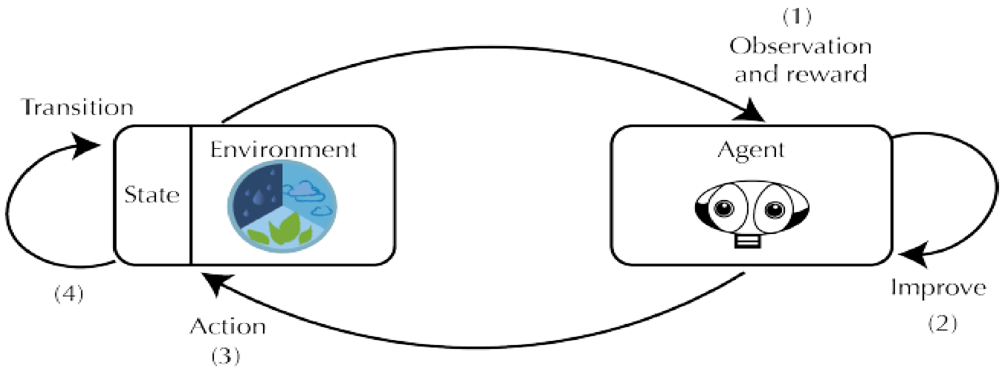

Figure 1. The reinforcement learning loop involves a series of steps: (1) Initially, the agent observes the environment. (2) This observation, along with the received reward, is utilized by the agent to enhance task performance. (3) Subsequently, the agent dispatches an action to the environment, aiming to exert positive control. (4) The environment transitions, altering its internal state based on the agent's action and its preceding state. This loop then initiates once more. The figure is inspired by [2].

Figure 2. components of DRL (Terven, 2025)

3. Review of Deep learning and DRL models for TSF

3.1. Deep learning Models

Existing research have explored several categories of deep learning models for TSF. Here, a few prominent DL models are explored and how they can be applied to forecasting.

3.1.1. Recurrent Neural Networks (RNNs)

RNNs aim to explore relations between current time series and past ones. Accordingly, RNN maintains hidden state that evolves over time and is updated based on current input and previous state (Casolaro et al., 2023). However, simple RNNs suffer from vanishing and exploding gradients when modelling long sequences. This challenge is addressed by Long short-term memory (LSTM) and gated recurrent unit (GRU) networks which introduces gating mechanisms that regulate the flow of information. Alzubaidi et al. (2021) noted that LSTM cells contain input, output and forget gates while GRUs use update and reset gates to reduce parameters and training time. Bidirectional variants (Bi-LSTM) process data in forward and backward directions to capture past and future context. These RNNs have wide applications in time-series forecasting in finance, energy and other sectors.

3.2. Convolutional Neural Networks

As noted by Alzubaidi et al. (2021), CNNs are associated with image processing but have been adapted for time series by applying one dimensional convolutions across temporal and variable dimensions. CNNs capture local patterns, robust to noise and can be efficient because convolutional filters are shared across time. However, standard CNNs have difficulty modelling non stationary signals and long term dependencies (Bu & Cho, 2020). As with RNNs, attention mechanisms and residual connections can mitigate these issues. Some CNN models apply multiple convolutions across different variables and use cross-channel attention to weight each feature (Liu & Wang, 2024). Point-wise and patch-wise convolutional attention strategies have been proposed to reduce complexity in transformer models thus highlighting the interplay between CNNs and attention. Arushana et al. (2024) summarized deep learning models as shown in table 1 below.

Table 2. Comparison of Advanced Deep Learning Models for Time Series Forecasting

Method	Description	Strengths	Limitations
RNN	Recurrent Neural Networks that handle sequential data by maintaining a hidden state.	Suitable for dynamic datasets, captures temporal and long-term dependencies.	Faces issues with vanishing gradients problem, issues with long-term dependencies, high computational cost.
ES-RNN	Exponential Smoothing RNN, a hybrid model combining traditional exponential smoothing methods with RNNs for enhanced performance.	Leverages strengths of both traditional and deep learning methods, improves accuracy and robustness. Effectively captures seasonal and trend components.	Challenges in hyperparameter tuning.
LSTM	Advanced type of RNN designed to learn long-term dependencies using special gates.	Addresses vanishing gradient problem, capable of learning long-term dependencies.	Computationally intensive, requires high-capacity resources.
Attention-LSTM	Attention-based LSTM model that improves efficiency and forecasting accuracy by incorporating attention mechanisms.	Sequence forecasting with long time steps. Nonlinearity and long memory of time series data.	Higher computational cost than basic LSTM. Large number of parameters.
CNN-LSTM	Hybrid model combining Convolutional Neural Networks (CNN) and LSTM to leverage spatial and temporal dependencies.	Captures both spatial and temporal features, enhances forecasting accuracy.	High error rate. Less reliable.
GRU	Simplified version of LSTM with fewer gates, designed to achieve similar performance with less computational complexity.	Efficient, less computationally intensive than LSTM, capable of handling long-term dependencies.	Capturing long-term dependencies, but not as effectively as LSTM.
Attention-GRU	Attention-based GRU model that improves efficiency and forecasting accuracy by	Improves long-range dependency capture.	Higher computational cost than basic GRU.

	incorporating attention mechanisms.		
Transformer	Attention-based model initially developed for natural language processing, captures long-range dependencies without recurrence.	Self-attention — past instances influence future outcomes. Efficient, captures long-range dependencies, balances relevance of input sequence segments.	High computational cost. Difficulty in capturing temporal dynamics. Overfitting. Impractical for real-world scenarios.
Hybrid Models	Combines statistical and deep learning techniques (e.g., ARIMA + LSTM, LSTM + CNN, Transformer + RNN).	Leverages strengths of multiple methods, enhances accuracy. Better interpretability. More robust to noise and missing data.	Complexity in model design, risk of overfitting.

3.3. Deep Reinforcement Learning

Terven (2025) highlights that DRL formalises sequential decision making as an agent interacts with environment to maximise reward. The environment is typically modelled as Markov decision process defined by states, actions, transition probabilities and reward functions. Key DRL models are discussed as follows;

3.3.1. Deep Q-Network (DQN)

Value-based DRL method where neural network learns by estimating expected reward (Q-value) of taking each action in a given state (Terven, 2025). Usually, DQN is used for discrete action spaces. In forecasting contexts, one approach is to discretize forecast outcomes or decisions and use DQN to pick best option. For example, DQN can be used to decide among a set of predictive models or to issue categorical predictions such as predicting if demand will rise, fall, or stay flat. In case study on currency exchange forecasting, Madhulatha and Ghori (2025) used LSTM to predict the next exchange rate and DQN agent took an action based on that prediction such as adjusting forecast or making trading decision. Thereafter, the agent received a reward based on the subsequent market movement. This LSTM-DQN hybrid agent was able to iteratively improve its policy thus achieving higher prediction accuracy than other baseline models. The authors reported the hybrid model is significantly better than CNN or RNN on the same data. Such results underscore that adding reinforcement learning layer on top of deep sequence models can enhance accurate predictions.

3.3.2. Policy-Gradient and Actor-Critic Methods (PPO, A3C)

Policy-based DRL algorithms well-suited for continuous or high-dimensional action spaces which aligns with forecasting real-valued quantities. Proximal Policy Optimization (PPO) and Advantage Actor-Critic (A3C) have been considered for problems like algorithmic trading where the agent outputs continuous buy and sell signals based on time series inputs (Terven, 2025). The continuous-action setting of policy-gradient methods is a natural fit for predicting quantities. One challenge, however, is training stability where algorithms like PPO are often preferred for their robustness in training which could be beneficial when learning from noisy time series.

3.3.3. Deep Recurrent Q-Learning (DRQN)

Variant of DQN integrates an LSTM into Q-network to allow the agent to maintain an internal state and handle partial observability in sequential data (Chen et al., 2024). This is useful for time series with long memory as LSTM can carry information from prior time steps when estimating Q-values. In water flow runoff forecasting, DRQN has predicted reservoir releases by observing rainfall-runoff time series through combining LSTM's sequence modelling with Q-learning's decision-making (Amin, 2024).

Terven (2025) provided summarized overview of DRL models and their usefulness in TSF as shown in table 2 below.

Table 2. Comparison of Deep Reinforcement Learning (DRL) Models for Time Series Forecasting

DRL Model	Core Idea	Use in Time Series Forecasting	Advantages	Limitations
DQN & Variants (Double, Dueling, Rainbow)	Value-based learning via Q-function approximation	Discrete decision forecasting (finance: buy/sell; supply chain: stock levels)	Stable, well-tested; strong in discrete domains	Poor for continuous actions; needs discretization
Policy Gradients (REINFORCE)	Directly parameterize policy and optimize with gradient ascent	Continuous control forecasting (e.g., adjusting time-varying demand levels)	Handles continuous spaces naturally	High variance, low sample efficiency
Actor–Critic (A2C,	Combines value learning	Parallelized multivariate	Reduces variance;	Sensitive to

A3C)	(critic) and policy optimization (actor)	forecasting tasks (real-time retail, traffic flow)	efficient parallelism	hyperparameters
PPO (Proximal Policy Optimization)	On-policy with clipped objective for stability	Healthcare, retail, energy demand prediction	Stable, robust, industry-adopted	Slower per iteration
DDPG (Deep Deterministic Policy Gradient)	Deterministic policy for continuous action spaces	High-dimensional forecasting (multi-factor energy/climate models)	Effective in continuous, high-dim spaces	Brittle, tuning-sensitive
Hierarchical RL (Option-Critic, FeUdal)	Decomposes tasks into subtasks/policies	Multi-scale forecasting (short vs. long horizon trends)	More interpretable, reusable skills	Complex credit assignment
Evolutionary Strategies (ES, CMA-ES)	Black-box optimization of policy parameters	Hyper parameter tuning, irregular forecasting tasks	Parallelizable, gradient-free	Sample-inefficient, resource-intensive

4. Performance Evaluation of Deep Learning and DRL with Attention Based Mechanisms in Real-World Forecasting Applications

This section compares performance of deep learning, DRL and attention-based models using publicly available datasets and empirical research in finance, retail and supply chain, climate and healthcare are evaluated.

4.1. Finance

Forecasting stock prices has been difficult due to high volatility and noisy signals. Pan et al. (2024) argues that traditional models like ARIMA and GARCH capture linear dependencies but fail to incorporate sentiment and complex interactions. This was exemplified by Alharbi et al. (2025) who found that attention-based LSTM models significantly outperform ARIMA for exchange rate and stock price prediction by capturing non-linear temporal dependencies. In the same context, Du and Shen (2024) explored DRL method using Q-learning combined convolutional neural networks and sentiment analysis from social media to predict Chinese stock prices. The model processed historical closing prices, volumes and comment texts and the DRL agent produced trading actions that achieved superior returns on two test sets compared with traditional methods and other deep learning model. Results showed that DRL variants like DRQNs provided better results than classic Q-learning because they could handle sequential data.

4.2. Retail and Supply Chain

In retail, demand forecasting is a time-series problem due to variations in seasonality, promotions and external factors. In a study by Gu et al. (2022), attention-LSTM showed strong performance on supply chain demand data which supports operational decision-making with accuracy better than more complex hybrid models. Precisely, the attention component enabled the model to pinpoint key past demand like recent spikes or seasonal events that should inform the next prediction. This model has been applied in retail sales forecasting (to capture promotions or holidays effects) and in supply chain for inventory demand (focusing on recent changes in demand trends) (Bhuiyan et al., 2025). Accordingly, attention models not only boost accuracy but also yield insights. Overall, attention-LSTM architectures generally outperform plain LSTMs on multivariate time series where certain observations have outsized importance on the forecast.

4.3. Climate and environmental forecasting

Climate systems exhibit spatiotemporal dependencies across multiple scales. Standard time-series models struggle to integrate spatial heterogeneity and dynamic interactions among variables like temperature, precipitation and land use. A recent deep learning approach for climate resilience combines graph neural networks (GNNs) with spatiotemporal attention (Chen & Dong, 2025). The model learns dynamic graphs which represents interactions between climate variables and regions, and uses attention to focus on relevant spatial–temporal dependencies. Multi-task learning helps in predicting short- and long-term outcomes, enabling early warning of extreme events. The authors note that transformers and attention mechanisms such as the Temporal Fusion Transformer provide global and local interpretability, addressing the limitations of RNNs. In hydrology application, Pölz et al. (2024) compared Transformer versus LSTM to forecast karst spring water discharge. They found that for a spring with long memory and slow dynamics, the Transformer achieved significantly better accuracy than the LSTM (about 9% lower error on average). However, on spring with very short-term response, the LSTM slightly outperformed the Transformer by approximately 4% error difference. This indicates that Transformers outperform when long-range dependencies are present but may fail to automatically win on every dataset if the data is limited.

4.4. Healthcare

Forecasting patient trajectories is critical for early detection of deterioration and personalised treatment. Forghani and Forouzanfar (2024) used Transformer-based model to forecast heart rate from ECG data and compared it to LSTM. From the findings, Temporal Fusion Transformer achieved 3.8 beats/min predicting heart rate 2 minutes ahead as compared to 4.3 beats/min with LSTM. The Transformer not only had lower error but also trained faster and captured subtle patterns in heart rate variability that LSTM missed. On this account, it can be deduced that in complex biomedical signals, the ability to attend to long-term patterns like circadian rhythms or accumulated sleep debt effects gave Transformers an edge.

Additionally, the Digital Twin–Generative Pretrained Transformer (DT-GPT) leverages large language models to forecast patient health trajectories without requiring data imputation or normalisation (Makarov et al., 2025). DT-GPT outperformed other models on datasets covering non-small cell lung cancer, intensive care unit stays and Alzheimer's disease thus reducing scaled mean absolute error by 1.3–3.4%. Overall, the interpretability of attention weights in transformers helps clinicians understand which variables drive model predictions.

5. Proposed hybrid architecture integrating attention mechanisms and DRL

To fully exploit the strengths of deep learning and DRL, a hybrid architecture that integrates attention-based forecasting module with DRL agent for multivariate TSF and decision-making is a proposed. The design emphasises interpretability and modularity and can be adapted to various sectors such as retail, climate, environment, energy and health. The proposed architecture comprises three components as shown in figure 3 below. First, an attention-based forecasting module ingests historical series, static covariates (demographics, customer attributes), and dynamic covariates (weather, calendars), transforming them into embeddings. Models like TFT or Informer variants deploy multi-head attention to capture long-range dependencies and highlight critical time steps and variables. This module outputs probabilistic multi-step forecasts along with interpretable attention weights. Second, DRL agent receives the forecasted trajectories and current state of the system. Using algorithms like PPO or DQN, the agent selects actions such as order quantities, energy dispatch, or treatment plans. Attention layers within the policy network further refine focus on salient aspects of the state. Third, the environment evolves based on the agent's actions and returns rewards aligned with domain goals (profit, service levels, patient outcomes). Forecasting and control can be trained sequentially or jointly, with reinforcement signals refining both. Overall, this integration leverages supervised forecasting accuracy, DRL's adaptive decision-making, and attention's interpretability, though challenges remain in computational efficiency, differentiable joint training, and reward design.

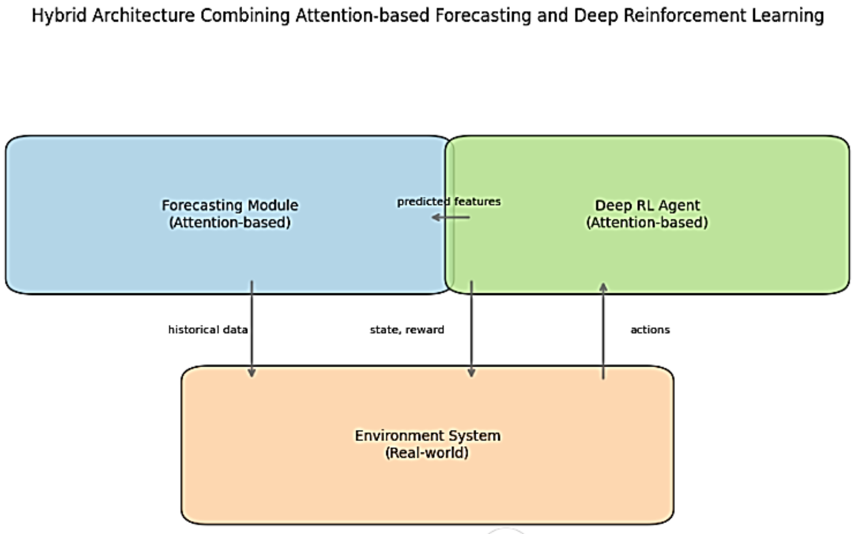

Figure 3. Proposed hybrid structure (Author)

6. Challenges and Future Research Directions

While deep learning and DRL with attention architectures have shown promise for time-series forecasting, several challenges remain. Kong et al. (2025) categorised the challenges into data-related, model structure and task-related issues as shown in figure 3 below.

- **Data-related issues**: can arise from mining data, anomalous data, noise data and data privacy leakage (Kong et al., 2025). High-quality, labelled time-series data are needed for training complex models. Models like DT-GPT demonstrate that transformers can handle missing data without imputation but further research is needed to generalise these techniques (Cheng et al., 2025). Also, forecasting models may introduce biases or amplify inequalities in healthcare or financial services. future researchers should ensure fairness and privacy preservation through techniques like differential privacy, federated learning or fairness constraints.
- **Model structure issues:** can arise from non-interpretability, non-continuity and computing resource (Kong et al., 2025). Attention weights provide some interpretability, but they do not necessarily correspond to causal importance. Future researchers should combine attention with causal inference or counterfactual analysis to provide more meaningful explanations.
- **Task-related issues:** can arise from parallel computing and variable types. Kong et al. (2025) argues that parallel

computing in TSF faces challenges of resource constraints, scalability limits, high GPU costs and inefficiency in achieving real-time online forecasting which future researchers should address using DRL models and attention mechanisms

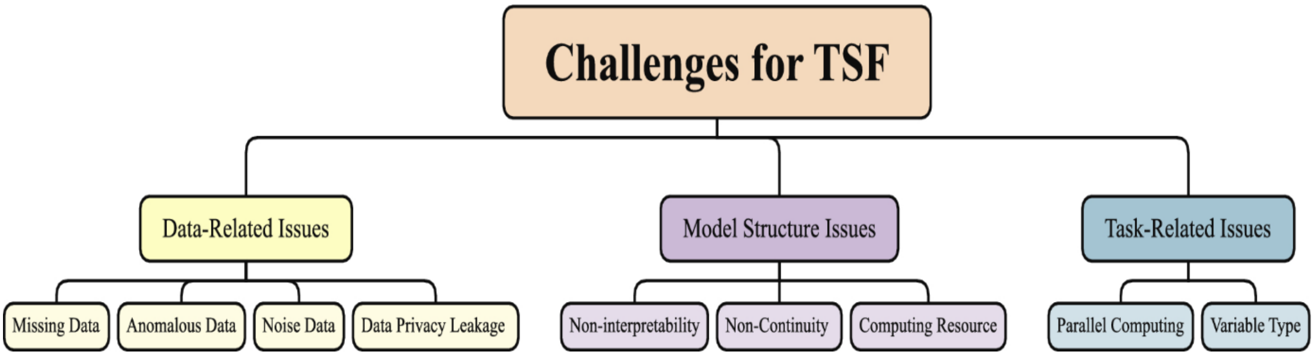

Figure 4. Challenges for TSF (Kong et al., 2025)

7. Conclusion

As seen above, TSF is a major concern with far-reaching implications across finance, supply chains, climate science and healthcare. Deep learning models like RNNs, CNNs, transformer-based architectures and emerging DRL variants have dramatically advanced the state of the art by learning complex temporal patterns and cross-variable interactions. Arguably, attention mechanisms enable models to focus on relevant features and time steps thus improving accuracy and interpretability. DRL extends these advances by enabling agents to convert forecasts into actions that maximise long-term reward. Case studies across finance, retail and supply chain, climate, and healthcare demonstrate that combining attention-based forecasting with DRL leads to significant improvements over traditional approaches.

The proposed hybrid architecture integrates an attention-based forecasting module with DRL agent to create modular pipeline that can be adapted to various domains. While challenges exist, future researchers can overcome these obstacles through innovations in model design, training strategies and ethical guidelines. Overall, researchers can build intelligent systems that not only predict the future but also make informed decisions for businesses and society at large.

References

[1] Casolaro, V. Capone, G. Iannuzzo and F. Camastra, "Deep Learning for TSF: Advances and Open Problems," *Information*, vol. 14, no. 11, 598, 2023. https://doi.org/10.3390/info14110598

[2] A. Pölz, A.P. Blaschke, J. Komma, A.H. Farnleitner and J. Derx, "Transformer versus LSTM: A comparison of deep learning models for karst spring discharge forecasting," *Water Resources Research*, vol. 60, no. 4, pp. 1-19, 2024. DOI:10.1029/2022WR032602.

[3] Vaswani, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, pp. 5998–6008, 2017.

[4] Chen and J. Dong, "Deep learning approaches for time series prediction in climate resilience applications", *Sec. Big Data, AI, and the Environment*, vol. 13, 2025, https://doi.org/10.3389/fenvs.2025.1574981

[5] Chen, Y. Vincent and D. Laird, "Deep q-learning with recurrent neural networks," *Stanford cs229 course report,* vol. 4, no. 3, 2024. https://cs229.stanford.edu/proj2016/report/ChenYingLaird-DeepQLearningWithRecurrentNeuralNetwords-report.pdf

[6] Andres, "Advanced Time Series Forecasting Methods," *Machine learning pills,* https://mlpills.dev/time-series/advanced-time-series-forecasting-methods/ [Accessed: 3rd October, 2025]

[7] Oluwagbade, "Reinforcement Learning for Time-Series Data Analysis," *University of Cape Coast,* 2025. https://www.researchgate.net/publication/395269263

[8] Ke, "An Efficient and Accurate DDPG-Based Recurrent Attention Model for Object Localization," *IEEE Access*, vol. 99, pp. 1-12, 2020. DOI:10.1109/ACCESS.2020.3008171

[9] Arushana, E.A.A. Priyadarshanie, R.T. Weliwatta and S.A.H. Amarasena, "Deep Learning Architectures for TSF,", pp. 1-15, 2024 DOI:10.13140/RG.2.2.19904.75524

[10] Pan, Y. Tang and G. Wang, "A Stock index futures price prediction approach based on the MULTI-GARCH-LSTM mixed model," *Mathematics*, vol. 12, no. 11, pp. 1677-1694, 2024.

[11] Moreira, J. Rivas, F. Cruz, R. Dazeley, A. Ayala and B. Fernandes, "DRL with interactive feedback in a human–robot environment" *Applied Sciences*, vol. 10, no. 16, pp. 5574-5593, 2020. https://doi.org/10.3390/app10165574

[12] Terven, "DRL: A Chronological Overview and Methods," *AI,* vol. 6, no. 3, pp. 46-54. https://doi.org/10.3390/ai6030046

[13] Alzubaidi, J. Zhang, A.J. Humaidi, A. Al-Dujaili, Y. Duan… and L. Farhan, "Review of deep learning: concepts, CNN architectures, challenges, applications, future directions," *Journal of Big Data,* vol. 8, no. 53, 2021. https://doi.org/10.1186/s40537-021-00444-8

[14] Cheng, Z. Liu, X. Tao, Q. Liu, J. Zhang, T. Pan, ... and E. Chen, "A comprehensive survey of TSF: Concepts, challenges, and future directions," *Authorea Preprints*, PP. 1-29, 2025. https://www.techrxiv.org/doi/full/10.36227/techrxiv.174430535.53879341

[15] M.D. Pra, "TSF with Deep Learning and Attention Mechanism," *Medium*, https://medium.com/data-science/time-series-forecasting-with-deep-learning-and-attention-mechanism-2d001fc871fc [Accessed: 2nd October 2025]

[16] M.H. Alharbi, "Prediction of the Stock Market Using LSTM, ARIMA, and Hybrid of LSTM-ARIMA Models," *Journal of Knowledge Management Application and Practice: An International Journal*, vol. 7, no. 1, pp. 15-22, 2025.

[17] M.S.A. Bhuiyan, J. Maua, S.R. Sultana, M.N.H. Mir, S.M. Sanchary and K. Nur, "Hybrid LSTM-CNN with Attention Mechanism for Accurate and Scalable Grocery Sales Forecasting," In *2025 International Conference on Electrical, Computer and Communication Engineering (ECCE)*, pp. 1-6, 2025. DOI:10.1109/ECCE64574.2025.11012951

[18] N. Forghani and M. Forouzanfar, "Performance Comparison of Transformer, LSTM, and ARIMA TSF Models: A Healthcare Application," In *the 1st International Online Conference on Bioengineering*, 16–18 October 2024, MDPI: Basel, Switzerland.

[19] N. Makarov, M. Bordukova, P. Quengdaeng, D. Garger, R. Rodriguez-Esteban… and M.P. Menden, "Large language models forecast patient health trajectories enabling digital twins," *NPJ Digit. Med.* Vol. 8, no. 588, 2025. https://doi.org/10.1038/s41746-025-02004-3

[20] N. Smith, "A review of Reinforcement learning for financial time series prediction and portfolio optimization", *Medium*, https://medium.com/journal-of-quantitative-finance/a-review-of-reinforcement-learning-for-financial-time-series-prediction-and-portfolio-optimisation-4cb2e92a23f3, 2021, [Accessed: 3rd October 2025]

[21] S. Amin, *Part I: Understanding Long Short-Term Memory (LSTM): LSTM Implementation from Scratch*, https://medium.com/@samina.amin/understanding-lstms-lstm-implementation-from-scratch-18965a150eca, 2024, Accessed [2nd October 2025].

[22] S. Du and H. Shen, "A Stock Prediction Method Based on DRL and Sentiment Analysis," *Applied Sciences*, vol. 14, no. 19, pp. 8747-8762. https://doi.org/10.3390/app14198747

[23] S.-J. Bu and S.-B. Cho, "TSF with Multi-Headed Attention-Based Deep Learning for Residential Energy Consumption," *Energies*, vol. 13, no. 18, pp. 4722-4743, 2020. https://doi.org/10.3390/en13184722

[24] T.S. Madhulatha and A.S. Ghori, "Deep neural network approach integrated with reinforcement learning for forecasting exchange rates using time series data and influential factors," *Scientific reports*, vol. 15, no. 29009, 2025. https://doi.org/10.1038/s41598-025-12516-3

[25] V.R. Thota, "*Comparative Study of TSF for Trucking Shipments: Evaluating Arima, LSTM, and Hybrid Arima-LSTM Models*," *Master's thesis*, State University of New York at Binghamton.

[26] X. Kong, Z. Chen, W. Liu, K. Ning, L. Zhang... and F. Xia, "Deep learning for TSF: a survey," *Int. J. Mach. Learn. & Cyber*, vol.16, pp. 5079–5112, 2025, https://doi.org/10.1007/s13042-025-02560-w

[27] X. Liu and W. Wang, "Deep TSF Models: A Comprehensive Survey," *Mathematics*, vol.12, no. 10, pp. 1504-1523, 2024. https://doi.org/10.3390/math12101504

[28] X. Zhang, P. Li, X. Han, Y. Yang and Y. Cui, "Enhancing Time Series Product Demand Forecasting with Hybrid Attention-Based Deep Learning Models," in *IEEE Access*, vol. 12, pp. 190079-190091, 2024, doi: 10.1109/ACCESS.2024.

[29] Y.H. Gu, D. Jin, H. Yin, R. Zheng, X. Piao and S.J. Yoo, "Forecasting agricultural commodity prices using dual input attention LSTM," *Agriculture*, vol. 12, no. 2, 256. https://doi.org/10.3390/agriculture12020256

[30] Y. Liu, C. Gong, L. Yang and Y. Chen, "DSTP-RNN: A dual-stage two-phase attention-based recurrent neural network for long-term and multivariate time series prediction," *Expert Systems with Applications*, vol. 143, no. 113082, 2020. https://doi.org/10.1016/j.eswa.2019.113082

[31] Thirunagalingam, A. (2024). Transforming real-time data processing: the impact of AutoML on dynamic data pipelines. Available at SSRN 5047601.

[32] Maroju, P. K. (2024). Advancing synergy of computing and artificial intelligence with innovations challenges and future prospects. FMDB Transactions on Sustainable Intelligent Networks, 1(1), 1-14.

[33] Priscila, S. S., Celin Pappa, D., Banu, M. S., Soji, E. S., Christus, A. T., & Kumar, V. S. (2024). Technological Frontier on Hybrid Deep Learning Paradigm for Global Air Quality Intelligence. In P. Paramasivan, S. Rajest, K. Chinnusamy, R. Regin, & F. John Joseph (Eds.), Cross-Industry AI Applications (pp. 144-162). IGI Global Scientific Publishing. https://doi.org/10.4018/979-8-3693-5951-8.ch010

[34] Panyaram, S. (2024). Utilizing quantum computing to enhance artificial intelligence in healthcare for predictive analytics and personalized medicine. FMDB Transactions on Sustainable Computing Systems, 2(1), 22-31.

[35] Mohanarajesh Kommineni. (2023/6). Investigate Computational Intelligence Models Inspired By Natural Intelligence, Such As Evolutionary Algorithms And Artificial Neural Networks. Transactions On Latest Trends In Artificial Intelligence. 4. **P**30. Ijsdcs.