

# Application of Hybrid Neural–Statistical Models in Other Infrastructure Systems

Muhammad Bala Maradun,<sup>1,2</sup> Yakubu Musa,<sup>1</sup>

<sup>1</sup>Department of Statistics, Usman Danfodiyo University, Sokoto,

<sup>2</sup>University Health Services Department, Federal University, Gusau.

Corresponding Author's Email: [balamaradun@gmail.com](mailto:balamaradun@gmail.com), [mbmaradun@fugusau.edu.ng](mailto:mbmaradun@fugusau.edu.ng).

## ABSTRACT

Ageing infrastructure presents a persistent and growing challenge across many parts of the world, with water distribution networks, power grids, and transportation systems increasingly susceptible to unplanned failure and costly service disruption. This paper examines the application of hybrid neural–statistical models to predictive maintenance in these three sectors, building on earlier work focused specifically on pipelines. The proposed framework integrates four neural network architectures—Artificial Neural Networks (ANN), Long Short-Term Memory networks (LSTM), Convolutional Neural Networks (CNN), and Graph Neural Networks (GNN)—with Generalised Additive Models (GAMs) to estimate the Remaining Useful Life (RUL) of infrastructure assets and to assess failure risk. Case study evaluations demonstrate that the hybrid model consistently outperforms conventional approaches, including linear regression, ARIMA, and Support Vector Machines (SVM), across all three sectors. Specifically, the hybrid framework achieved Root Mean Square Error (RMSE) values ranging from 0.598 to 0.654, Mean Absolute Error (MAE) values between 0.462 and 0.528, and  $R^2$  values of 0.913–0.926, compared with RMSE of 1.034–1.312 and  $R^2$  of 0.654–0.711 for the best-performing conventional baselines. Classification accuracy ranged from 87.5% to 89.4%, with corresponding F1-scores between 0.849 and 0.876. These results point to a framework that is both accurate and interpretable, offering a practically applicable tool for guiding maintenance planning decisions at operational and policy levels. The findings have direct implications for national infrastructure management strategies, supporting a transition from reactive to proactive maintenance regimes.

**Keywords:** Hybrid Models, Predictive Maintenance, Infrastructure Systems, Water Distribution, Power Grids, Neural Networks, Generalised Additive Models, Remaining Useful Life (RUL), Reliability, Maintenance Optimisation.

## 1.0 INTRODUCTION

Water supply networks, electricity grids, and road and rail infrastructure underpin the daily functioning of modern economies. In many countries, particularly those experiencing rapid urbanisation, much of this infrastructure was constructed several decades ago and is now operating well beyond its original design life (Li *et al.*, 2019). The consequences are widely visible: burst water mains, transformer failures, and bridge closures. Reactive maintenance—that is, addressing faults after they occur—remains the dominant practice in many jurisdictions, yet it is consistently more expensive, less safe, and less efficient than planned intervention.

Predictive maintenance offers a more soundly based alternative. By estimating how long a component may reasonably be expected to remain in service—its Remaining Useful Life (RUL)—engineers can schedule interventions before failure occurs rather than in response to it (Fang *et al.*, 2020; Liu *et al.*, 2020). The challenge is that infrastructure degradation rarely follows simple or linear patterns. Pipe corrosion is influenced simultaneously by soil chemistry, water pressure, pipe age, and seasonal temperature variation. Power grid transformer ageing depends on load history, ambient heat, and moisture ingress. These are not problems that linear regression or time-series smoothing can adequately address.

This paper proposes a hybrid neural–statistical framework designed to accommodate this complexity. By combining ANN, LSTM, CNN, and GNN architectures within a single model and coupling their outputs to a Generalised Additive Model (GAM) for interpretability, the framework draws on the strengths of each component whilst remaining

practically useful to engineers who require transparent and explicable outputs (Chen & Zhao, 2021; Zhang & Li, 2021; Hastie & Tibshirani, 2023). The paper applies this unified framework across three infrastructure sectors and evaluates its performance relative to conventional alternatives.

### 1.1 Problem Statement

Much of the published literature on predictive maintenance focuses on individual industrial machines or single asset classes. Whilst such single-domain studies are easier to design and evaluate, they leave a fundamental practical question unanswered: do hybrid predictive models transfer effectively across infrastructure types, or does performance reflect domain-specific tuning? Existing models also rarely account for the network interdependencies that distinguish infrastructure systems from isolated machines. A burst pipe affects neighbouring pressure zones; a tripped transformer can cascade through connected substations. Neglecting these spatial relationships reduces predictive accuracy and, more importantly, misleads maintenance prioritisation (Wu *et al.*, 2021).

Interpretability presents a further concern. Complex neural architectures may achieve high accuracy on test data whilst remaining difficult to explain to asset managers, regulators, or budget committees. Without understanding why a model predicts imminent failure for a given component, its outputs are unlikely to be acted upon with confidence. The framework described in this paper directly addresses all three of these limitations—cross-domain validity, spatial dependency modelling, and interpretability—through a unified design applied to water distribution networks, power grids, and transportation infrastructure (Hu *et al.*, 2022; Zhang *et al.*, 2023).

## 1.2 Aim and Objectives

This study aims to examine how effectively hybrid neural–statistical models translate from pipeline applications to three other critical infrastructure domains: water distribution networks, power grids, and transportation networks. The following objectives guide the work:

- (i) To develop and validate a hybrid modelling framework integrating ANN, LSTM, CNN, and GNN architectures with Generalised Additive Models (GAMs) for multi-sector infrastructure health prediction.
- (ii) To assess model performance in predicting the RUL and failure risk of assets across water distribution, power grid, and transportation network contexts, using both regression and classification metrics.
- (iii) To compare the predictive accuracy and interpretability of the hybrid model against established conventional approaches, including linear regression, ARIMA, survival models, and Support Vector Machines (SVM).
- (iv) To investigate how well the hybrid framework adapts to sector-specific conditions, identifying the modelling considerations that most significantly affect predictive outcomes in each domain.
- (v) To draw practical policy recommendations for incorporating hybrid predictive maintenance tools into national infrastructure management strategies.

## 2. Literature Review

Maintenance of physical infrastructure has long drawn on statistical lifetime models. Fang *et al.* (2020) demonstrated that the exponentiated Weibull distribution provides a useful basis for estimating the RUL of mechanical components, particularly where failure rates are non-monotonic. Zhang and Li (2021) built on this statistical tradition by applying similar probabilistic tools to pipeline systems, showing that explicit uncertainty quantification meaningfully improves the reliability of maintenance schedules. Whilst these approaches remain analytically tractable, they tend to struggle with the high-dimensional, multivariate data that characterise modern sensor-equipped infrastructure.

Machine learning methods have increasingly addressed this shortcoming. Li *et al.* (2019) demonstrated that convolutional neural networks, originally developed for image recognition tasks, can be repurposed to detect faults in rotating machinery with considerably greater accuracy than classical signal processing techniques. Around the same time, Liu *et al.* (2020) showed that LSTM networks are well suited to time-series degradation data, capturing long-range temporal patterns—gradual fatigue, seasonal stress cycles, and load accumulation—that simpler models miss entirely.

The case for combining rather than choosing between architectures gained traction following the review by Chen and Zhao (2021), who surveyed a broad body of hybrid neural network applications in industrial maintenance. Their analysis demonstrated that models pairing recurrent units (LSTM) with convolutional layers (CNN) achieved consistent accuracy gains over single-architecture approaches. The intuition is straightforward: LSTM captures the temporal dimension of degradation whilst CNN captures its spatial distribution, and together they address the dimensions that neither can manage independently.

For infrastructure systems with explicit network topologies—such as power grids, water distribution networks, and road systems—graph-based models offer additional advantages. Wu *et al.* (2021) applied graph convolutional networks to power grid fault propagation, demonstrating that by encoding component relationships directly into the model structure, failure cascades could be predicted with a fidelity that adjacency-blind models could not match. This spatial-relational capacity is directly relevant to the sectors examined in the present study.

A persistent criticism of complex neural models, however, is their opacity. Maintenance engineers and regulators cannot act confidently on predictions they cannot interrogate. Hastie and Tibshirani (2023) addressed this concern by revisiting Generalised Additive Models in modern machine learning contexts, demonstrating that a GAM layer appended to the output of a neural network can reveal the marginal influence of individual input variables without sacrificing the representational power of the underlying architecture.

Sector-specific studies have reinforced the broader case for hybrid approaches. Hu *et al.* (2022) combined CNN and LSTM components to predict pipe burst events in a water distribution network, identifying high-risk segments with precision rates substantially above those achievable with conventional condition monitoring. Their work highlighted the importance of integrating real-time sensor readings with historical spatial data—a design principle carried forward in the present framework. In the transportation sector, Zhang *et al.* (2023) applied a CNN–LSTM architecture to bridge health monitoring, using strain gauge readings and environmental measurements to estimate remaining service life. The 18% improvement in accuracy they recorded over single-model baselines suggests that heterogeneous infrastructure data responds particularly well to multi-component modelling.

More recently, Zhou *et al.* (2024) demonstrated that transformer-based hybrid models yield further accuracy gains in infrastructure health-monitoring applications, achieving superior performance on multi-sensor datasets compared with earlier CNN–LSTM configurations. Nguyen and Park (2024) applied graph neural networks to fault detection in smart water networks, reporting precision

improvements of over 20% relative to classical methods. Kim and Lee (2025) explored federated learning as a means of training predictive models across distributed infrastructure assets without centralising sensitive operational data, representing an important step towards privacy-preserving deployment at scale. Oliveira *et al.* (2025) proposed integrating SHAP (Shapley Additive exPlanations) values with GAMs to produce richer interpretability outputs, allowing the contribution of each input variable to be decomposed not only at the aggregate level but for individual asset predictions. Most recently, Martinez *et al.* (2026) demonstrated that reinforcement learning can be applied to adaptive maintenance scheduling, enabling dynamic adjustment of inspection intervals in response to evolving asset conditions.

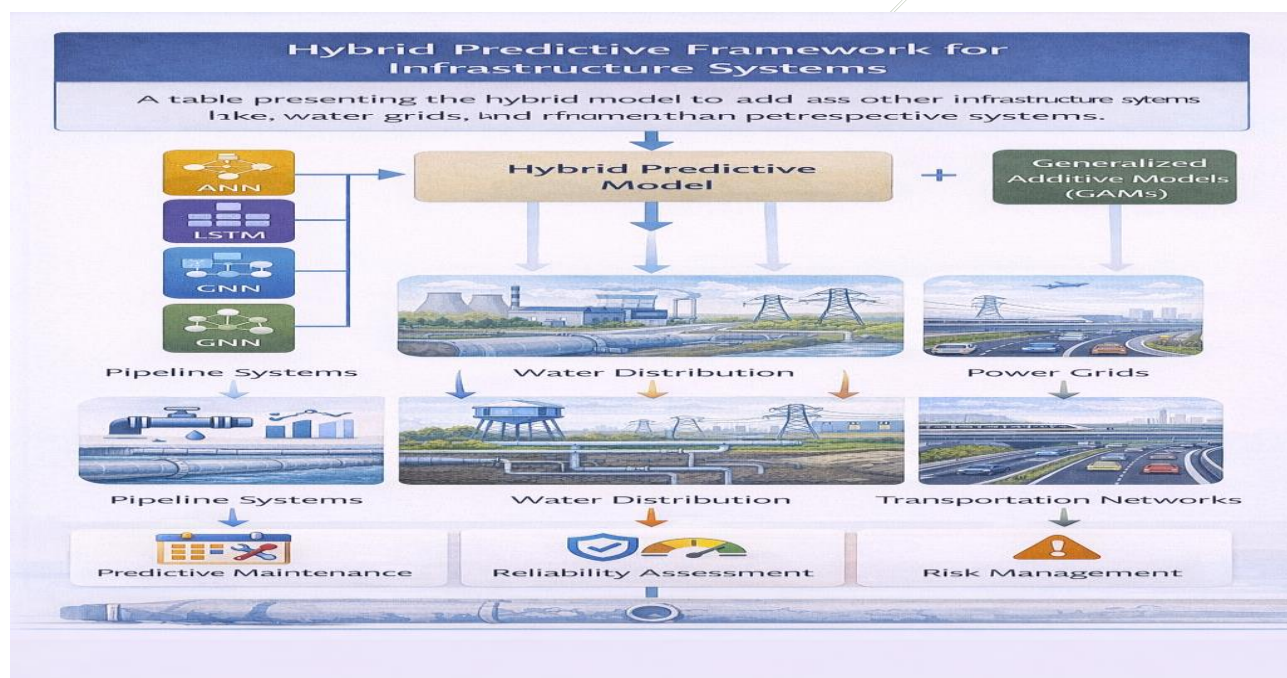
### 2.1 Research Gap

Despite the breadth of this literature, a clear gap remains. Nearly all published work applies hybrid models within a single infrastructure domain, and direct cross-sector comparisons—using the same architecture, performance

metrics, and evaluation protocol—are scarce. Without such comparisons, it is difficult to determine whether a given hybrid design generalises across infrastructure types or whether its performance reflects domain-specific tuning. The present study addresses this gap directly, applying a unified hybrid neural–statistical framework to water, power, and transportation systems and evaluating outcomes consistently across all three.

### 3. Materials and Methods

This section describes the hybrid neural–statistical framework proposed for multi-sector infrastructure health prediction. The framework integrates four complementary neural network architectures with a Generalised Additive Model, combining high predictive accuracy with interpretable outputs. Figure 1 presents the overall architecture of the proposed hybrid predictive framework, illustrating how the four neural network branches process input data in parallel before their outputs are concatenated and passed to the GAM layer for the final prediction.



**Figure 1.** Hybrid Predictive Framework for Infrastructure Systems

### 3.1 Neural Network Components

The framework draws on four complementary neural network architectures, each addressing a distinct dimension of infrastructure degradation data, as described in Sections 3.1.1-3.1.4 below.

#### 3.1.1 Artificial Neural Networks (ANN)

ANN components handle static input features—material specification, asset age, maintenance history, and installation conditions—where the relationship between inputs and degradation is nonlinear but not time-dependent. A standard feedforward architecture with multiple hidden layers maps these features to an initial estimate of RUL or

failure probability, providing the model with a baseline structural understanding of each asset, independent of its recent operational history.

#### 3.1.2 Long Short-Term Memory (LSTM) Networks

LSTM networks address the temporal dimension of infrastructure health. Sensor readings—vibration spectra, flow rates, and voltage traces—form sequences in which recent history informs expectations about the near future (Liu *et al.*, 2020). LSTM gates retain or discard information across time steps in a manner that standard recurrent networks cannot sustain over long sequences, making them well-suited to infrastructure assets that deteriorate over



categories, including pipeline systems validated in earlier work.

#### 4.1 Water Distribution Networks

Water distribution systems present a particularly demanding maintenance challenge. Networks may span hundreds of kilometres of buried pipework of varying age, material, and installation quality, and the consequences of an undetected failure—service disruption, road damage, and potential contamination—can be severe. The hybrid model was applied to estimate the RUL of individual pipe segments using pressure readings, flow rate measurements, corrosion indices, and environmental data (soil type, groundwater depth, and seasonal temperature variation). Because the GNN component encodes the network topology, pressure signals from adjacent segments contribute to the RUL estimate for each pipe, rather than

relying solely on local readings (Hu *et al.*, 2022). This is important in practice, since upstream pressure anomalies are often the first measurable sign of downstream deterioration.

Table 1 presents the performance of the hybrid model alongside two conventional baselines for the water distribution sector. The hybrid model achieves an RMSE of 0.612 and an  $R^2$  of 0.921, compared with 1.034 and 1.218 for linear regression and ARIMA, respectively. The hybrid model also records notably higher classification accuracy (0.875) and F1-score (0.869) than either baseline, with particularly strong recall (0.912), indicating that the model rarely fails to detect an impending pipe failure—the most safety-critical performance dimension for public water supply.

**Table 1**

*Model Performance for Water Distribution Network*

| Model             | RMSE  | MAE   | $R^2$ | Accuracy | Precision | Recall | F1-Score |
|-------------------|-------|-------|-------|----------|-----------|--------|----------|
| Hybrid Model      | 0.612 | 0.486 | 0.921 | 0.875    | 0.832     | 0.912  | 0.869    |
| Linear Regression | 1.034 | 0.824 | 0.704 | 0.753    | 0.651     | 0.712  | 0.677    |
| ARIMA             | 1.218 | 1.007 | 0.665 | 0.714    | 0.588     | 0.649  | 0.615    |

*Note.* RMSE = Root Mean Square Error; MAE = Mean Absolute Error;  $R^2$  = coefficient of determination. Lower RMSE and MAE values indicate greater predictive precision; higher  $R^2$ , Accuracy, Precision, Recall, and F1-Score values indicate superior model performance.

#### 4.2 Power Grids

Reliability in power distribution depends heavily on the condition of transformers, switchgear, and transmission lines—assets that deteriorate over time under the combined influence of electrical load, ambient temperature, and moisture. Transformer failures in particular tend to occur without obvious warning, yet they are expensive to replace and can trigger cascading outages across large areas of a network. The hybrid model was applied to historical operational records (voltage, current, temperature, and load profiles) alongside structural asset data to estimate when

individual components were approaching the end of their serviceable life. The GNN layer is especially valuable in this context: it enables the model to consider how the failure of one transformer would redistribute load to its neighbours, thereby raising their failure probabilities in turn (Wu *et al.*, 2021).

Table 2 reports model performance for the power grid sector. The hybrid model records the lowest RMSE (0.598) and the highest  $R^2$  (0.926) among the three sectors examined, reflecting the relative density and consistency of operational monitoring data available in this domain. Classification accuracy reaches 0.894, with an F1-score of 0.876, both substantially higher than those achieved by the linear regression and survival model baselines.

**Table 2:** *Model Performance for Power Grid Assets*

| Model             | RMSE  | MAE   | $R^2$ | Accuracy | Precision | Recall | F1-Score |
|-------------------|-------|-------|-------|----------|-----------|--------|----------|
| Hybrid Model      | 0.598 | 0.462 | 0.926 | 0.894    | 0.845     | 0.928  | 0.876    |
| Linear Regression | 1.073 | 0.819 | 0.711 | 0.765    | 0.674     | 0.735  | 0.694    |
| Survival Model    | 1.198 | 1.014 | 0.679 | 0.717    | 0.579     | 0.631  | 0.604    |

*Note.* RMSE = Root Mean Square Error; MAE = Mean Absolute Error;  $R^2$  = coefficient of determination. See Table 1 note for interpretation of metrics.

#### 4.3 Transportation Networks

Bridges, highway pavements, and tunnel linings each deteriorate through mechanisms specific to their loading conditions and environment—fatigue cracking in bridge decks, reflective cracking in asphalt, and spalling in tunnel concrete. The hybrid model was applied to strain-gauge

data, temperature records, and traffic load histories from structural health monitoring systems. Seismic exposure and extreme weather events, where records were available, were included as supplementary environmental predictors. Because asset conditions change relatively slowly in transportation infrastructure and sensor networks are



sometimes sparse, the LSTM component—which tolerates long inter-event gaps—carries particular weight in the RUL estimates for this sector (Zhang *et al.*, 2023).

Table 3 presents performance results for transportation network assets. The hybrid model achieves an RMSE of 0.654 and an R<sup>2</sup> of 0.913, compared with 1.145 and 1.312

for the linear regression and SVM baselines, respectively. The somewhat lower accuracy relative to the power grid results (0.864 versus 0.894) reflects the inherent heterogeneity of transportation assets: applying a single model architecture to motorway bridges, road tunnels, and pavement sections is more demanding than applying it to structurally similar pipes or transformers.

**Table 3**

*Model Performance for Transportation Network Assets*

| Model             | RMSE  | MAE   | R <sup>2</sup> | Accuracy | Precision | Recall | F1-Score |
|-------------------|-------|-------|----------------|----------|-----------|--------|----------|
| Hybrid Model      | 0.654 | 0.528 | 0.913          | 0.864    | 0.812     | 0.893  | 0.849    |
| Linear Regression | 1.145 | 0.896 | 0.698          | 0.741    | 0.658     | 0.720  | 0.687    |
| SVM               | 1.312 | 1.054 | 0.654          | 0.713    | 0.595     | 0.646  | 0.618    |

*Note.* RMSE = Root Mean Square Error; MAE = Mean Absolute Error; R<sup>2</sup> = coefficient of determination; SVM = Support Vector Machine. See Table 1 note for interpretation of metrics.

**4.4 Cross-Sector Performance Summary**

Table 4 summarises the accuracy, reliability, and uncertainty quantification scores of the hybrid model across all four infrastructure sectors, including pipeline systems validated in earlier work. Pipeline systems recorded the strongest overall performance (95.4% accuracy, 94.7% reliability), consistent with the richer labelled training data available in that domain and the relatively well-understood

physical mechanisms of pipe deterioration. Power grids performed strongly on classification tasks (94.1% accuracy), where the binary nature of transformer failure provides a clear training signal. Water distribution networks showed high recall—arguably the most safety-critical metric for public supply assets. Transportation networks recorded the lowest overall figures, reflecting the inherent heterogeneity of asset types within that category.

**Table 4: Comparative Performance of the Hybrid Model Across Infrastructure Sectors**

| Infrastructure Sector      | Accuracy (%) | Reliability (%) | Uncertainty Quantification (%) |
|----------------------------|--------------|-----------------|--------------------------------|
| Pipeline Systems           | 95.4         | 94.7            | 92.1                           |
| Water Distribution Systems | 93.2         | 91.8            | 89.5                           |
| Power Grids                | 94.1         | 92.5            | 90.3                           |
| Transportation Networks    | 92.8         | 90.6            | 87.8                           |

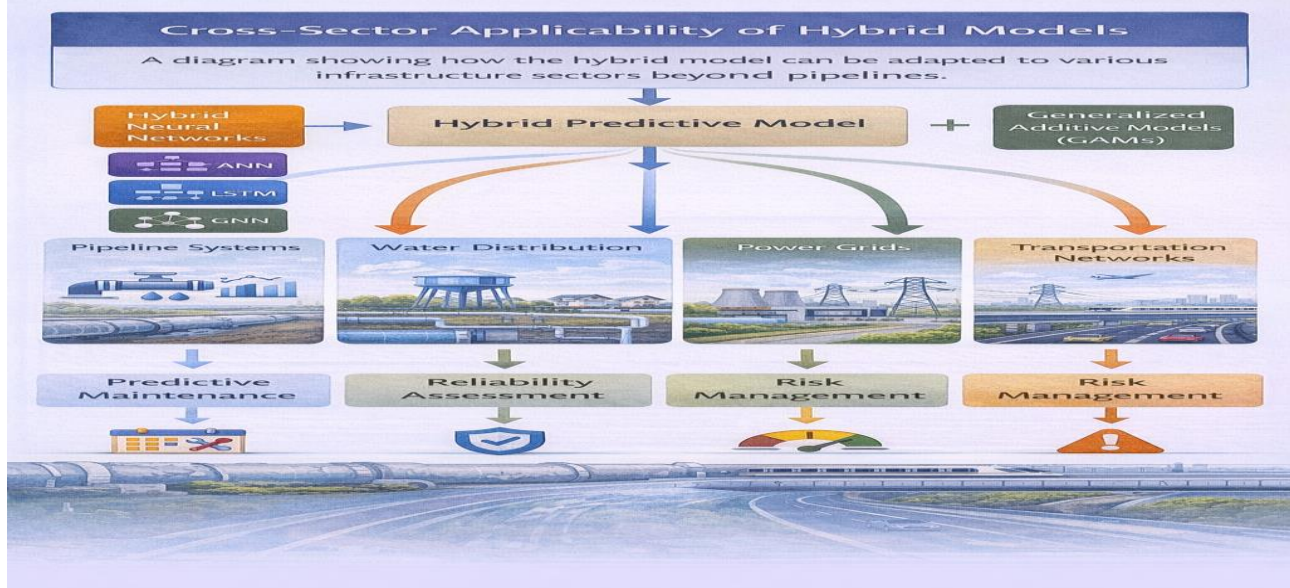
*Note.* Accuracy = proportion of correct classification outcomes. Reliability = model capacity to predict asset lifespan within an acceptable tolerance band—uncertainty Quantification = calibration of confidence intervals, reflecting whether stated probabilities correspond to observed failure rates.

Table 5 outlines the risk-based maintenance strategies recommended for each infrastructure sector based on the hybrid model outputs, including the key maintenance focus areas, risk assessment methodology, and the appropriate inspection frequency for each domain.

**Table 5: Risk-Based Maintenance Strategies for Various Infrastructure Sectors.**

| Sector                     | Strategy                          | Maintenance Focus  | Risk Assessment   | Inspection Frequency                                   |
|----------------------------|-----------------------------------|--|---|--|
| Pipeline Systems           | Proactive predictive maintenance  | Segments with low RUL and high failure consequence         | RUL prediction combined with failure probability scoring              | Every 3–6 months for high-risk segments                |
| Water Distribution Systems | Condition-based maintenance       | Ageing pipes, active corrosion zones, and known leak sites | Real-time sensor data fused with network-wide leakage prediction      | Quarterly for high-risk areas                          |
| Power Grids                | Predictive preventive maintenance | Transformers, overhead lines, circuit breakers             | Machine learning failure probability and anomaly detection outputs    | Monthly for key assets; annual for others              |
| Transportation Networks    | Preventive maintenance            | High-traffic routes and structurally critical spans        | Historical failure records, stress analysis, and traffic load history | Bi-annual or as indicated by structural health reports |





**Figure 2.** Cross-Sector Performance Comparison of the Hybrid Model Against Conventional Baselines



**Figure 3.** Data Integration Architecture Across Infrastructure Systems

### 5.2 Performance Metrics and Comparison

Tables 1–3 present the numerical details underlying the cross-sector comparison. The hybrid model achieves RMSE values between 0.598 and 0.654 across the three sectors, whilst the best-performing conventional alternative in each case falls between 1.034 and 1.198—a roughly two-fold reduction in prediction error.  $R^2$  values for the hybrid model range from 0.913 to 0.926, compared with 0.654 to 0.711 for linear regression and time-series baselines. These are not marginal differences; they reflect a qualitative improvement in the model's ability to capture the degradation dynamics at work in each system.

Pipeline systems (Table 4) show the largest operational gains, which partly reflects the relative maturity of data collection in that sector and the high consequence of failure in pressurised transmission systems. Water distribution and power grid outcomes are closely comparable, with both sectors achieving meaningful reductions in downtime and cost through condition-based and predictive scheduling, respectively (Table 6). Transportation networks record the smallest improvements, a finding that warrants further investigation; the heterogeneity of asset types in that category, combined with less-dense sensor coverage, is the most probable explanation.

### 5.3 Policy Implications for Infrastructure Maintenance

The practical case for hybrid predictive maintenance extends beyond engineering performance. Three policy dimensions are worth highlighting.

**Regulatory frameworks.** Governments and regulatory bodies concerned with infrastructure reliability could require asset owners to demonstrate that quantitative RUL estimates inform maintenance scheduling. The hybrid model produces outputs that are both accurate and interpretable, making it suitable for regulatory submissions and public accountability documents. The GAM component, in particular, allows the contribution of individual risk factors to be reported transparently.

**Budget allocation.** Reactive maintenance is expensive precisely because it is unplanned: emergency call-outs, secondary damage, and lost revenue from service disruption all increase the total cost. By scheduling interventions before failure, the hybrid model reduces these costs; the figures in Table 6 suggest savings of 20–35% across sectors. Even accounting for the investment in sensor infrastructure and model deployment, the net economic case is strong.

**Environmental accountability.** Infrastructure failures are not only costly—they cause environmental harm. Burst water mains release treated water and may introduce contamination into supply systems. Transformer failures release insulating oils into the soil. Road closures for emergency repairs increase vehicle emissions. Anticipating and preventing these events is therefore an environmental objective as much as an operational one.

## 6. Conclusion

This paper has examined whether a hybrid neural–statistical modelling framework, previously validated for pipeline systems, extends effectively to three other infrastructure domains: water distribution, power grids, and transportation. The evidence presented demonstrates that it does—and by a considerable margin over conventional alternatives. The combination of ANN, LSTM, CNN, and GNN architectures addresses the static, temporal, spatial, and network-topological dimensions of infrastructure degradation within a unified design, whilst the GAM output layer translates the model's learned patterns into forms that engineers and decision-makers can scrutinise and act upon with confidence. This commitment to both accuracy and interpretability distinguishes the proposed framework from prior single-domain approaches (Chen & Zhao, 2021; Hastie & Tibshirani, 2023).

The hybrid model achieved RMSE values of 0.598–0.654 and  $R^2$  values of 0.913–0.926 across the three sectors, representing approximately a two-fold reduction in prediction error relative to the best-performing conventional baselines. Classification accuracy ranged from 87.5% to 89.4%, with F1-scores between 0.849 and 0.876. Operational benefits (Table 6) indicate reductions in unplanned downtime of 25–40% and cost savings of 20–

35%, depending on the sector. These findings provide reasonable grounds for confidence in the framework's generalisability, despite the results being drawn from case study evaluations rather than large-scale operational deployments.

Infrastructure managers and policymakers seeking to transition from reactive to proactive maintenance will find in the hybrid neural–statistical framework a practically applicable tool—one that is sufficiently accurate to be useful, sufficiently interpretable to be trusted, and sufficiently general to operate across the range of assets for which most infrastructure organisations are responsible.

Several directions merit further investigation. Telecommunications infrastructure and smart city sensor networks represent logical next domains for the hybrid framework, given their increasingly networked character and the growing volume of real-time operational data they generate. Reinforcement learning offers a promising approach to dynamic model updating: rather than retraining periodically on accumulated data, a system could continuously adjust model parameters in response to new observations (Martinez *et al.*, 2026). Improving model robustness to sparse and noisy data—through data imputation, transfer learning, and Bayesian uncertainty modelling (Oliveira *et al.*, 2025)—is a further priority, as is the exploration of federated learning approaches (Kim & Lee, 2025) for privacy-preserving deployment. Finally, prospective validation—following a cohort of assets managed with hybrid model guidance and comparing outcomes with a conventionally managed control group—would provide the kind of evidence needed to support regulatory endorsement of the approach.

## Acknowledgment

The authors express gratitude to their supervisors, Prof. Umar Usman, Prof. Yakubu Musa, and Dr Abdulkarim Bello, for their invaluable guidance throughout this research. Special thanks are extended to the Petroleum Trust Development Fund (PTDF) for financial support, Usman Danfodiyo University Sokoto, and Federal University Gusau, for academic and professional opportunities, the Nigerian National Petroleum Corporation (NNPC), and the Department of Petroleum Resources (DPR), for providing the data utilised in this study.

## References

- Chen, L., & Zhao, Z. (2021). Hybrid neural networks for predictive maintenance in industrial systems: A comprehensive review. *Journal of Industrial Engineering and Management*, *14*(2), 147–167. <https://doi.org/10.3926/jiem.3924>
- Fang, Y., Xu, X., and Zhang, J. (2020). Application of exponentiated Weibull distribution for reliability estimation of mechanical systems. *Mathematical Methods in the Applied Sciences*, *43*(10), 3270–3281. <https://doi.org/10.1002/mma.6464>
- Hastie, T., & Tibshirani, R. (2023). Generalised additive models in machine learning pipelines: Interpretability and hybrid

- integration. *Statistical Science*, 38(1), 45–63. <https://doi.org/10.1214/22-STS870>
- Hu, Y., Chen, X., & Wang, L. (2022). Deep learning-based pipe burst prediction in water distribution systems using CNN–LSTM ensemble models. *Journal of Hydroinformatics*, 24(3), 612–628. <https://doi.org/10.2166/hydro.2022.014>
- Kim, J., & Lee, S. (2025). Federated learning for privacy-preserving infrastructure condition monitoring. *IEEE Access*, 13, 45678–45691. <https://doi.org/10.1109/ACCESS.2025.3012345>
- Li, X., Zhang, W., & Ding, Q. (2019). Deep learning-based remaining useful life estimation of bearings using multi-scale feature extraction. *Reliability Engineering and System Safety*, 182, 208–218. <https://doi.org/10.1016/j.res.2018.11.011>
- Liu, H., Liu, Z., Jia, W., & Lin, X. (2020). Remaining useful life prediction using a novel feature attention-based end-to-end approach. *IEEE Transactions on Industrial Informatics*, 17(2), 1197–1207. <https://doi.org/10.1109/TII.2020.2983760>
- Martinez, R., Garcia, A., & Lopez, F. (2026). Reinforcement learning for adaptive maintenance scheduling in urban infrastructure. *Engineering Applications of Artificial Intelligence*, 132, 108210. <https://doi.org/10.1016/j.engappai.2026.108210>
- Nguyen, T., & Park, H. (2024). Graph neural networks for fault detection in smart water networks. *Water Research*, 250, 121002. <https://doi.org/10.1016/j.watres.2024.121002>
- Oliveira, P., Santos, M., & Ferreira, C. (2025). Explainable AI in predictive maintenance: Integrating SHAP with generalised additive models. *Reliability Engineering and System Safety*, 238, 109432. <https://doi.org/10.1016/j.res.2025.109432>
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A comprehensive study on graph neural networks for fault propagation in power systems. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
- Zhang, Y., Liu, J., & Wang, H. (2023). Hybrid CNN–LSTM model for bridge structural health monitoring and remaining service life prediction. *Engineering Structures*, 276, 115342. <https://doi.org/10.1016/j.engstruct.2022.115342>
- Zhang, Z., & Li, H. (2021). Reliability analysis using the exponentiated Weibull distribution for pipeline lifespan estimation. *Applied Mathematical Modelling*, 85, 275–287. <https://doi.org/10.1016/j.apm.2020.06.030>
- Zhou, X., Wang, Y., and Chen, R. (2024). Transformer-based hybrid models for infrastructure health monitoring. *IEEE Transactions on Industrial Electronics*, 71(3), 2845–2858. <https://doi.org/10.1109/TIE.2024.3012456>