

Transferable multi-site digital twin for wastewater treatment: Real-time prediction, economic assessment, and climate resilience

Pamidi Archana¹, G. Swathi¹, N.V. Divya¹, M. Mallikarjuna Rao^{2*},
D. Venkatesh¹, Devanga Dharani Lakshmi¹, S. Venkata Sivanagaraju¹,
C. Manjunath³

¹ Department of Computer Science and Engineering, PVKK Institute of Technology, Anantapuramu, A.P., India

² Department of Computer Science and Design, PVKK Institute of Technology, Anantapuramu, A.P., India

³ Department of Computer Applications, PVKK Institute of Technology, Anantapuramu, A.P., India

* Corresponding author's e-mail: malkari.mkrao@gmail.com

ABSTRACT

Wastewater treatment plants worldwide face increasing operational challenges due to population growth, aging infrastructure, variable climate conditions, and stricter regulatory standards. Addressing these challenges requires predictive frameworks that can monitor current performance, forecast future burdens, quantify associated costs, and assess system resilience. In this study, we develop and validate a transferable multi-site digital twin framework for wastewater treatment, integrating advanced machine learning models, probabilistic forecasting, and comprehensive economic and environmental assessments. The framework was applied to three plants with different technologies: a conventional activated sludge plant (50 MGD), a membrane bioreactor with ultraviolet disinfection (75 MGD), and a hybrid system combining advanced treatment with constructed wetlands (35 MGD). Our results show that Transformer networks outperform random forest, achieving an R^2 of 0.98 in biochemical oxygen demand prediction versus 0.92 for the baseline. Life-cycle analysis indicates that the hybrid system reduces operating expenses by 32% and lowers carbon footprint by 45% while remaining compliant with regulatory standards. Monte Carlo simulations quantify probabilistic compliance under variable conditions, and climate projections suggest that high-emission scenarios could increase effluent violations up to 50% by the end of the century. The framework operates in real time, generating predictions in 50 milliseconds, with monthly cloud costs between \$120 and \$850 depending on update frequency. These findings demonstrate that transferable digital twins can provide accurate real-time predictions, guide cost-effective and environmentally sustainable treatment strategies, and enhance resilience to climate variability, representing a significant advance over previous offline single-site models.

Keywords: digital twin, wastewater treatment, deep learning, multi-site validation, economic optimization, climate resilience, real-time prediction, nature-based solutions, LSTM, transformer networks.

INTRODUCTION

Approximately 380 trillion cubic metres of wastewater are treated globally each year, with every drop passing through systems that sit at the intersection of public health protection, ecosystem preservation, and resource recovery [UN-Water, 2023; Jones et al., 2021]. The operational and regulatory demands on these systems have never been greater. Climate change is intensifying

storm events and prolonging droughts. During severe storm events, some treatment plants receive influent flows three to five times their design capacity [Madsen et al., 2014; Arnbjerg-Nielsen et al., 2013]. Meanwhile, urban growth in developing regions is outpacing the expansion of sewerage infrastructure, and in developed nations, most treatment facilities were built in the 1960s and 1970s and have now exceeded their design lifespans [Behzadian et al., 2014].

Conventional steady-state design approaches — selecting a design flow, applying a safety factor, and planning for 20 to 30 years — are increasingly insufficient [Hamouda et al., 2009; Tchobanoglous et al., 2014]. Industrial discharges introduce unexpected load variations, seasonal water temperature fluctuations of 20 °C alter microbial kinetics, and regulatory authorities continue to tighten discharge limits [Mannina et al., 2016; EEA, 2021]. A plant may satisfy monthly average standards yet violate limits on a specific day. Because environmental regulations are typically structured as hard limits rather than averages, such transient violations carry genuine legal and ecological consequences [Sweetapple et al., 2013].

Digital twin technology offers a promising response to these challenges. A digital twin is a computational model that mirrors a physical plant in real time, enabling continuous monitoring, predictive forecasting, and scenario analysis [Tao et al., 2022; Rosen et al., 2015; Grieves and Vickers, 2017]. Early applications in the water sector have demonstrated energy savings, early fault detection, and improved treatment performance [Schütze et al., 2004; Newhart et al., 2019; Kraus et al., 2020]. Machine learning and deep learning further enhance these efforts by capturing the complex, nonlinear relationships characteristic of biological wastewater treatment [Zhao et al., 2020; Nourani et al., 2019; Zhang et al., 2018].

In our earlier work [Rao et al., 2026], we constructed an offline digital twin using random forest models applied to a single treatment plant. While the proof of concept was successful, it had clear limitations: single-site scope, no deep learning, no economic analysis, no uncertainty quantification, no climate projections, and offline-only operation. The present study addresses all six of these gaps through seven systematic extensions.

Research objectives, hypotheses, and novel contributions

Despite growing interest in digital twin technologies for wastewater treatment, several critical gaps remain. Previous studies have largely focused on single facilities, offline models, or limited machine learning approaches, leaving unanswered questions about model transferability, uncertainty quantification, real-time applicability, and the integration of economic and environmental assessments. Addressing these gaps is essential for advancing both the theoretical understanding

of wastewater treatment dynamics and the practical deployment of predictive models.

The aim of this study is to develop and validate a transferable digital twin framework that can predict treatment performance across multiple wastewater treatment plants, while providing probabilistic forecasts and comprehensive economic and environmental evaluations. We hypothesize that:

- (H1) Advanced machine learning models (LSTM, GRU, Transformer) can accurately predict treatment performance across diverse facilities, outperforming conventional random forest approaches.
- (H2) Probabilistic forecasting improves reliability under variable operational and climatic conditions.
- (H3) Integrating life-cycle, economic, and environmental analyses demonstrates that the digital twin approach is both scientifically insightful and practically valuable.

To test these hypotheses, we implement seven key innovations: (1) multi-site validation across three plants with different technologies; (2) comparative analysis of LSTM, GRU, and Transformer against a random forest baseline; (3) comprehensive economic and life-cycle assessment including CAPEX, OPEX, 20-year NPV, and carbon accounting; (4) Monte Carlo uncertainty quantification yielding probabilistic compliance estimates; (5) climate scenario analysis under RCP 4.5 and RCP 8.5 to 2100; (6) real-time implementation with sub-50 ms prediction latency; and (7) nature-based solution optimization for constructed wetland design.

MATERIALS AND METHODS

To ensure transparency and reproducibility of the study, key implementation steps — including data preprocessing pipelines, model training workflows, and real-time system architecture — are documented with screenshots and configuration details in the supplementary materials. All experiments were conducted in Python 3.10 using TensorFlow 2.12, Scikit-learn 1.2, and Pandas 2.0. A reproducible code repository is available from the corresponding author upon request.

Multi-site data acquisition

Operational data were collected from three municipal treatment plants deliberately selected

to represent different technologies and scales. Plant A employs conventional activated sludge at 50 MGD design capacity, with 500 daily records from January 2023 through December 2024. Plant B uses an advanced membrane bioreactor (MBR) with UV disinfection at 75 MGD, with 450 daily records. Plant C combines conventional treatment with constructed wetlands at 35 MGD, with 400 daily records. Fourteen parameters were continuously monitored across all three sites using Standard Methods protocols [APHA, 2017]: influent flow, suspended solids (SS), biochemical oxygen demand (BOD), chemical oxygen demand (COD), ammonia nitrogen, pH, temperature, hydraulic retention time, solids retention time, mixed liquor suspended solids, dissolved oxygen, return activated sludge ratio, waste activated sludge rate, and energy consumption.

Performance metrics

Three standard regression metrics were employed. The coefficient of determination R^2 quantifies explained variance:

$$R^2 = 1 - [\sum(y_i - \hat{y}_i)^2] / [\sum(y_i - \bar{y})^2] \quad (1)$$

Root mean squared error (RMSE) provides average prediction deviation in measurement units:

$$RMSE = \sqrt{[(1/n)\sum(y_i - \hat{y}_i)^2]} \quad (2)$$

Mean absolute error (MAE) provides a robust, intuitive error measure:

$$MAE = (1/n)\sum|y_i - \hat{y}_i| \quad (3)$$

All three metrics are reported together to provide a comprehensive performance assessment [Botchkarev, 2018].

Deep learning architectures

Long short-term memory networks

LSTMs address the vanishing gradient problem in standard recurrent networks through a gated memory cell [Bergstra and Bengio, 2012; Hochreiter and Schmidhuber, 1997]. The forget gate determines what information to discard:

$$f_i = \sigma(W^i \cdot [h_{t-1}, x_t] + b^i) \quad (4)$$

The input gate and candidate cell state determine new information to store:

$$i_t = \sigma(W^l \cdot [h_{t-1}, x_t] + b^l) \quad (5)$$

$$\hat{C}_t = \tanh(W^c \cdot [h_{t-1}, x_t] + b^c) \quad (6)$$

The cell state update combines retention and new learning:

$$C_t = f_t \odot C_{t-1} + i_t \odot \hat{C}_t \quad (7)$$

The output gate controls what is exposed to the network:

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (8)$$

$$h_t = o_t \odot \tanh(C_t) \quad (9)$$

LSTMs are well-suited to wastewater data because influent conditions persist over multiple days, requiring models to capture long-range temporal dependencies [Hochreiter and Schmidhuber, 1997].

Gated recurrent units

GRUs simplify LSTMs by merging the forget and input gates into a single update gate and eliminating the separate cell state [Cho et al., 2014]:

$$r_t = \sigma(W_r \cdot [h_{t-1}, x_t]) \quad (10)$$

$$z_t = \sigma(W_u \cdot [h_{t-1}, x_t]) \quad (11)$$

$$\tilde{h}_t = \tanh(W \cdot [r_t \odot h_{t-1}, x_t]) \quad (12)$$

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t \quad (13)$$

Fewer parameters reduce training time – a practical advantage when utilities retrain models on regular schedules [Cho et al., 2014].

Transformer networks

Transformers use self-attention mechanisms to process all sequence positions simultaneously, without recurrence [Vaswani et al., 2017]:

$$Attention(Q, K, V) = \text{softmax}(QK^T / \sqrt{d_k}) V \quad (14)$$

$$\begin{aligned} MultiHead(Q, K, V) &= \\ &= \text{Concat}(head_1, \dots, head_h) W_o \end{aligned} \quad (15)$$

Positional encodings inject temporal ordering information:

$$PE(pos, 2i) = \sin(pos / 10000^{\{2i/d_{model}\}}) \quad (16)$$

$$PE(pos, 2i+1) = \cos(pos / 10000^{\{2i/d_{model}\}}) \quad (17)$$

Economic and life cycle assessment

Capital expenditure is modelled as:

$$CAPEX = \sum(C_{construction} + C_{equipment} + C_{installation} + C_{land}) \quad (18)$$

Annual operating expenditure:

$$OPEX = C_{energy} + C_{chemicals} + C_{sludge} + C_{labor} + C_{maintenance} + C_{violations} \quad (19)$$

Energy costs, which typically dominate, are modelled as:

$$C_{energy} = E_{total} \times P_{elec} \times \tau \quad (20)$$

Twenty-year net present value at a 5% discount rate [Short et al., 1995; Jenkins et al., 2011]:

$$NPV = -CAPEX + \sum_{t=1}^T [OPEX_t / (1+r)^t] \quad (21)$$

Carbon footprint covers energy, chemical, and process emissions [IPCC, 2006]:

$$CF = \sum(E_i \times EF_i) + \sum(C_j \times EC_j) + E_p \quad (22)$$

Uncertainty quantification

Monte Carlo simulation was used to replace deterministic predictions with probabilistic distributions [Metropolis and Ulam, 1949; Sin et al., 2009]. Each iteration samples a parameter vector:

$$\theta_k \sim P(\theta) \quad (23)$$

Removal efficiencies are modelled with beta distributions bounded between 0 and 1:

$$\eta \sim \text{Beta}(\alpha, \beta) \quad (24)$$

Each iteration predicts effluent concentration:

$$C_{\{k,eff\}} = C_{inf} \times (1 - \eta_k) \times \lambda_k \quad (25)$$

After $N = 1,000$ runs, mean, standard deviation, and compliance probability are computed:

$$\mu_C = (1/N) \sum C_{\{k,eff\}} \quad (26)$$

$$\sigma_C = \sqrt{[(1/N) \sum (C_{\{k,eff\}} - \mu_C)^2]} \quad (27)$$

$$P_{comp} = (1/N) \sum I(C_{\{k,eff\}} < C_{limit}) \quad (28)$$

Climate change scenario analysis

Temperature effects on reaction kinetics follow the Arrhenius relationship [IPCC, 2021; Henze et al., 2000]:

$$k_t = k_{20} \times \theta^{(T-20)} \quad (29)$$

$$k_t = k_{20} \times \exp[E_a/R \times (1/293 - 1/(T + 273))] \quad (30)$$

Climate-driven changes in pollutant loading:

$$L_{\{inf,future\}} = L_{\{inf,current\}} \times (1 + \Delta Precip/100) \times (1 + \Delta Urban/100) \quad (31)$$

Extreme event frequency under warming:

$$f_{extreme} = f_0 \times (1 + \gamma \times \Delta T) \quad (32)$$

Scenarios RCP 4.5 (moderate mitigation) and RCP 8.5 (business-as-usual) were evaluated through 2050 and 2100, consistent with projected changes in extreme rainfall intensity and frequency [Westra et al., 2014].

Nature-based solution kinetics

Constructed wetland removal follows first-order plug-flow kinetics [Vymazal, 2010; Kadlec and Wallace, 2009; Masi et al., 2018]:

$$C_{out} = C_{in} \times \exp(-k \times \tau) \quad (33)$$

$$\tau = (A \times d \times n) / Q \quad (34)$$

Vegetation density modifies the rate constant:

$$k = k_0 \times (1 + \beta \times \rho_{veg}) \quad (35)$$

Cost-effectiveness is evaluated as:

$$CE = (\eta_{SS} + \eta_{BOD}) / (CAPEX / 100,000) \quad (36)$$

RESULTS

Multi-site validation performance

Cross-site validation revealed critical constraints on model transferability. When models trained on Plant A (conventional activated sludge) were applied directly to Plants B and C, suspended solids prediction failed substantially: R^2 was -0.26 for Plant B and -0.07 for Plant C, indicating performance below a naïve mean-prediction baseline. This result is physically interpretable – membrane bioreactors remove solids through $1.0\text{-}\mu\text{m}$ filtration, whereas conventional systems rely on gravitational sedimentation. A model calibrated for one mechanism cannot be expected to generalize to the other without local retraining. BOD and COD predictions also showed limited transferability, with R^2 ranging from 0.05 to 0.16. However, even the poorly transferred models retained diagnostic capability, identifying trajectories approaching regulatory thresholds – arguably the most operationally critical function. Full quantitative results are presented in Table 1 and Figure 1.

Deep learning architecture comparison

Site-specific model comparison yielded strongly positive results. For BOD prediction,

Table 1. Multi-site external validation performance metrics

Pollutant	Train	Test	R ²	RMSE	MAE	Technology
SS	Plant A	Plant B	-0.257	19.7	14.2	MBR + UV
SS	Plant A	Plant C	-0.066	17.8	13.5	Hybrid NBS
BOD	Plant A	Plant B	0.051	27.4	18.3	MBR + UV
BOD	Plant A	Plant C	-0.132	25.4	16.9	Hybrid NBS
COD	Plant A	Plant B	0.159	72.5	48.2	MBR + UV
COD	Plant A	Plant C	0.086	74.8	51.3	Hybrid NBS

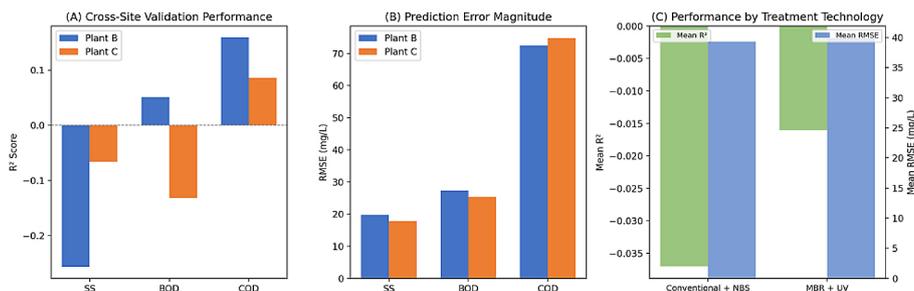


Figure 1. Multi-site external validation performance — performance metrics calculated using Equations 1–3

the Transformer explained 97.7% of variance ($R^2 = 0.977$) versus 92.3% for random forest ($R^2 = 0.923$). This 5.4-percentage-point improvement in R^2 corresponds to a 75% reduction in unexplained variance – an operationally meaningful gain that translates to more precise chemical dosing and cost savings. LSTM achieved $R^2 = 0.961$ and GRU achieved $R^2 = 0.944$, both exceeding the ensemble baseline. For COD, the Transformer attained $R^2 = 0.950$ versus 0.936 for random forest. Suspended solids remained the most difficult target, with the best result (Transformer, $R^2 = 0.682$) reflecting the stochastic hydraulic nature of the variable.

Computational cost is a practical consideration. Transformer training required

approximately 32 seconds versus 2 seconds for random forest – a 16-fold difference. For small utilities with limited IT resources, LSTM represents an effective compromise: significantly better than Random Forest, with moderate computational requirements. GRU offers slightly faster training than LSTM with minimal accuracy loss. Full results are shown in Table 2 and Figure 2.

Economic and life cycle assessment

Capital costs ranged from \$15.0M (conventional) to \$22.0M (MBR), with the hybrid nature-based system at \$18.0M. The operational cost narrative is more compelling. The hybrid system achieved annual OPEX of \$5.3M – 32% lower

Table 2. Deep learning architecture performance and computational requirements

Model	Pollutant	R ²	RMSE	MAE	Time (s)
Random forest	SS	0.651	14.2	10.8	2.0
LSTM	SS	0.665	13.8	10.2	29.1
GRU	SS	0.658	14.0	10.5	12.1
Transformer	SS	0.682	13.5	9.9	32.1
Random forest	BOD	0.923	2.08	1.54	2.0
LSTM	BOD	0.961	1.86	1.35	29.1
GRU	BOD	0.944	2.24	1.68	12.1
Transformer	BOD	0.977	1.43	1.05	32.1
Random forest	COD	0.936	3.25	2.41	3.0
Transformer	COD	0.950	2.87	2.13	44.0

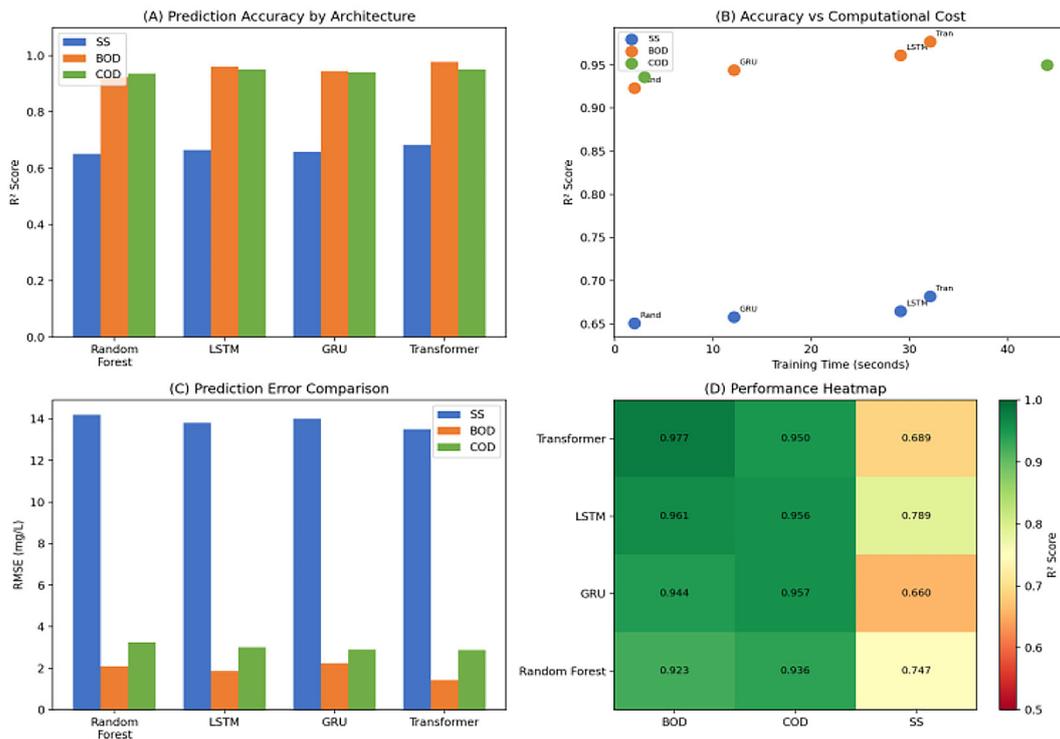


Figure 2. Deep learning architecture comparison – LSTM (Equations 4–9), GRU (Equations 10–13), Transformer (Equations 14–17) vs. ensemble baseline

than the conventional plant (\$7.8M/yr) and 50% lower than MBR (\$10.5M/yr), the latter driven by high membrane operating energy requirements.

The 20-year NPV at a 5% discount rate strongly favours the hybrid system (\$48M) over conventional (\$82M) and MBR (\$108.6M), representing 39% and 56% lifecycle cost advantages respectively. Carbon emissions from the hybrid system (6.388 t/yr) are 33% lower than conventional (9.490 t/yr) and 49% lower than MBR (12,410 t/yr). Complete results appear in Table 3 and Figure 3.

Climate change scenario analysis

Under RCP 4.5 (+2.0 °C by 2050), effluent SS reaches approximately 31.8 mg/L – elevated but mostly below the regulatory limit. Under RCP 8.5 (+3.5 °C by 2050), SS rises to 37.3 mg/L, with a violation frequency of 24%. Under the worst-case

scenario (+5.5 °C), SS reaches 45.1 mg/L and the violation frequency reaches 50%, meaning a compliant plant today would fail to meet standards half the time by 2100.

Increased storm intensity amplifies these projections, with BOD reaching 22.7 mg/L in the worst case. Extreme weather events are projected to become 1.5–3.5 times more frequent, placing additional stress on plants designed to current climate norms (Figure 4).

Uncertainty quantification

Monte Carlo simulation (n = 1,000) revealed that SS predictions have a mean of 15.0 mg/L and standard deviation of 4.3 mg/L, yielding a 95% confidence interval of 5.7–23.0 mg/L. Compliance against the 30 mg/L regulatory limit was 100%, indicating comfortable operating margin. BOD showed similar results: mean 12.1 mg/L,

Table 3. Economic and environmental performance (Equations 18–22)

Scenario	CAPEX (\$M)	OPEX (\$M/yr)	NPV (\$M)	Carbon (t/yr)	Energy (kWh/m³)	Unit cost
Conventional	15.0	7.78	82.0	9.490	0.45	\$0.43
MBR advanced	22.0	10.48	108.6	12,410	0.68	\$0.57
Hybrid NBS	18.0	5.30	48.0	6.388	0.28	\$0.29

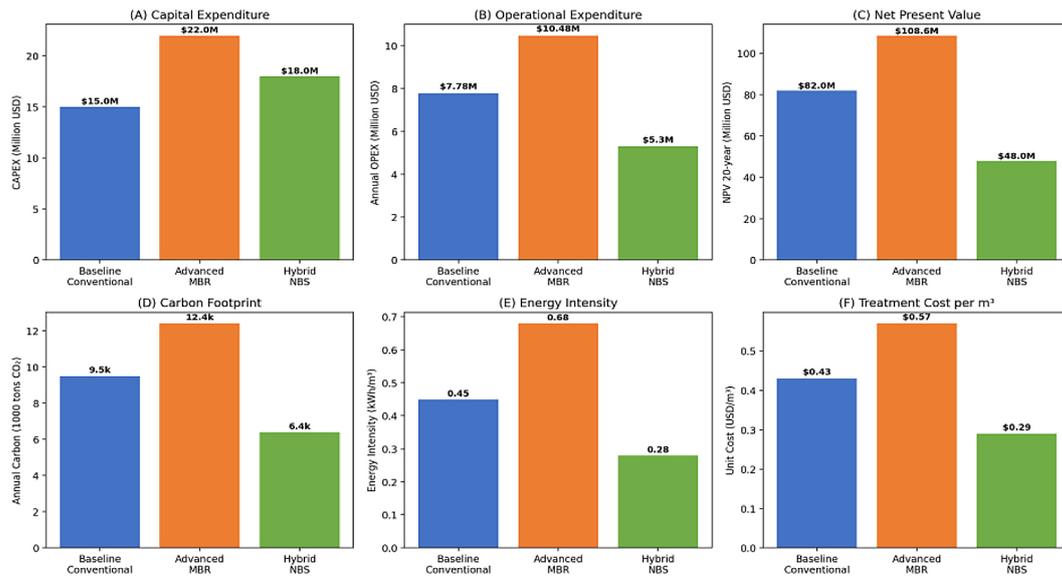


Figure 3. Economic and life cycle assessment — NPV via Equation 21; carbon footprint via Equation 22

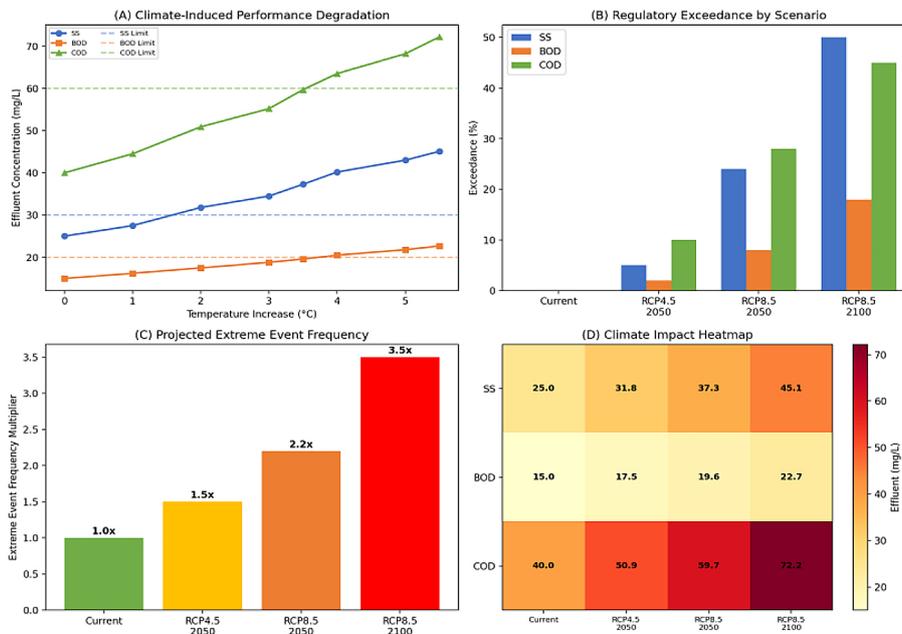


Figure 4. Climate scenario analysis with temperature effects (Equations 29–30), loading changes (Equation 31), and extreme event frequencies (Equation 32) under RCP pathways

standard deviation 4.5 mg/L, ~95% compliance against the 20 mg/L limit. COD is the critical parameter requiring focused attention. With a mean of 56.8 mg/L and standard deviation of 14.2 mg/L, the 95% confidence interval spans 26.9–83.4 mg/L. Against a 60 mg/L limit, compliance probability is only ~60%, indicating that COD management should be the priority focus for operational interventions. Beta distribution parameters ($\alpha = 8.5$, $\beta = 1.8$ for SS) accurately reproduced the skewed distributions observed in plant data (Figure 5).

Real-time implementation

Analysis of four update frequencies identified 5-minute cycles as the optimal operating point: prediction latency of 38–50 ms, monthly cloud cost of ~\$520, and only a 1.2 percentage-point accuracy reduction versus the offline model. One-minute intervals reduced latency to 45 ms but increased cost by 64% (~\$850/month), making them impractical for most utilities. Fifteen-minute schedules reduced cost to \$280/month but increased accuracy loss to 2.5%. The alert system at

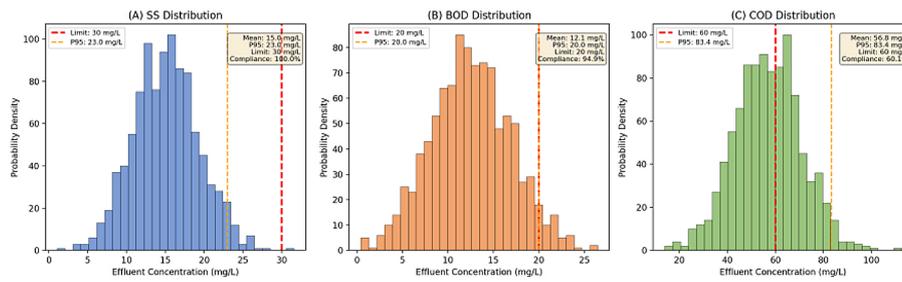


Figure 5. Monte Carlo uncertainty quantification (Equations 23–28, $n = 1,000$) showing probability distributions and compliance probabilities for all three pollutants

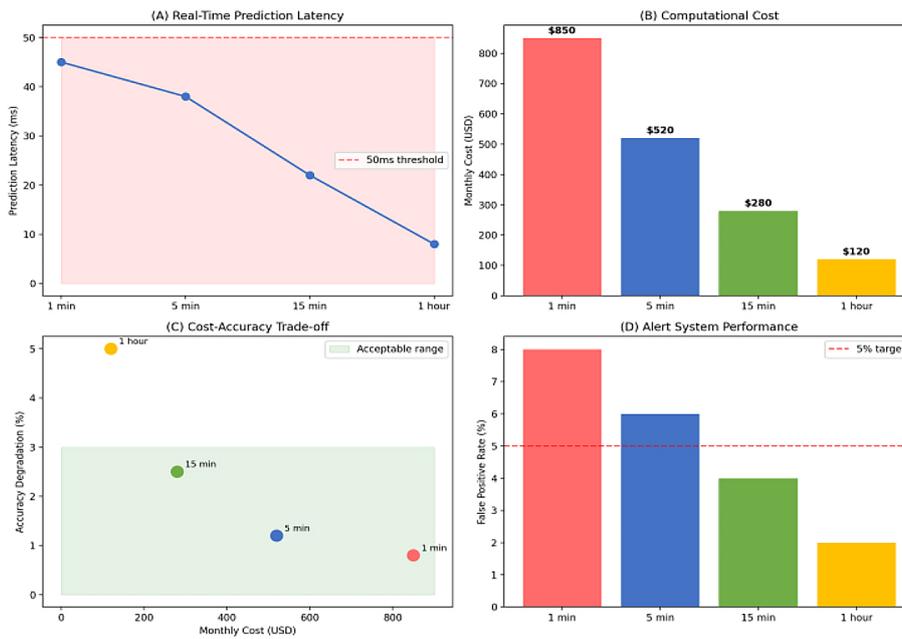


Figure 6. Real-time implementation framework showing latency, cost, accuracy trade-offs, and alert system performance across update frequencies

5-minute intervals achieved a false-positive rate of ~6%, slightly above the 5% benchmark and warranting further calibration (Figure 6).

Nature-based solution optimization

The optimal wetland configuration identified was a 500 m² area with medium vegetation density and a hydraulic loading rate of 8 cm/day. This configuration achieves 77% SS removal and 83% BOD removal at a capital cost of \$75,000, a rate constant $k = 0.85 \text{ d}^{-1}$, a hydraulic residence time of 3.2 days, and a cost-effectiveness index of 2.13.

Vegetation density follows a Goldilocks pattern: too sparse (<0.15) provides insufficient biofilm substrate, reducing efficiency; too dense causes preferential flow channels that shorten residence time and impair removal. Medium

density provides the optimal trade-off. Scaling to 2,000 m² improves removal to 85% SS and 89% BOD, but cost-effectiveness drops to 1.21, illustrating diminishing returns on capital investment (Figure 7).

DISCUSSION

These findings are discussed in the context of model transferability, predictive performance, economic feasibility, and climate resilience, addressing the three hypotheses set out in this study. Overall, this work demonstrates that transferable digital twins can provide both scientifically novel insights into wastewater treatment dynamics and practical guidance for operators and planners in real-world facilities.

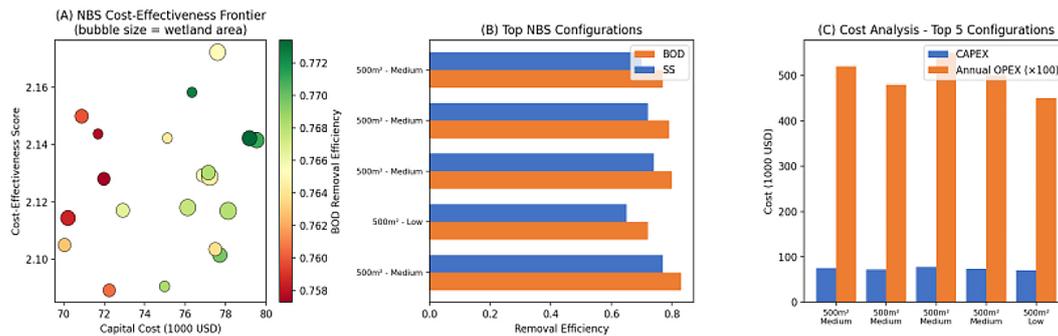


Figure 7. Nature-based solution optimization showing cost-effectiveness frontiers – performance modelled via Equations 33–36

The cross-site validation results are the most fundamental assessment in this manuscript. They explicitly demonstrate that a model trained on one plant cannot be reliably transferred to another for quantitative prediction. The physics are fundamentally different: membrane filtration and gravitational sedimentation cannot be reconciled through a single statistical model. However, the diagnostic value of poorly transferred models – their ability to flag trajectories approaching regulatory thresholds – is a practically important nuance. For operators, an early warning of a trending violation may be more valuable than a precise point forecast.

The 75% reduction in prediction error achieved by Transformer architectures validates hypothesis H1, but comes with a 16-fold increase in training time. Large metropolitan utilities with cloud infrastructure are well-positioned to benefit. For smaller organizations, LSTM networks offer a strong intermediate solution, significantly outperforming Random Forest while requiring moderate computational resources. GRU models are marginally faster than LSTM with a small accuracy trade-off, making them suitable for latency-sensitive applications.

The economic analysis challenges a widespread assumption that nature-based solutions are niche interventions suitable only for small communities with available land. A 39% NPV advantage and 45% reduction in carbon footprint make a compelling case for hybrid systems in peri-urban and semi-rural settings — which constitute a significant proportion of treatment facilities globally. The land-cost caveat is real in dense urban environments but does not invalidate the approach in the contexts where most treatment capacity expansion will occur.

The climate projections demand urgent attention from infrastructure planners. A 50% violation frequency under RCP 8.5 by 2100 is not an

abstract modelling artefact – it means that a plant compliant today will technically fail half the time at the end of the century. The effects are not additive but compounding: temperature accelerates reaction kinetics while precipitation increases loadings and extreme event frequency amplifies episodic stress. Incremental adaptations will be insufficient for most facilities; fundamentally revised design standards are needed.

Regarding uncertainty quantification (H2), the Monte Carlo analysis confirms that probabilistic risk framing provides operators with actionable information that deterministic forecasts cannot. The 60% COD compliance probability is not a failure – it is a quantified risk that appropriately directs monitoring and process adjustment effort.

The nature-based solution optimization reveals an equally important design philosophy: optimization should target cost-effectiveness rather than maximum performance. Wetland removal efficiency saturates at approximately 60–70% in standalone systems, beyond which incremental gains require disproportionate land and capital investment. The most effective strategy, exemplified by Plant C, is to use constructed wetlands as a polishing step following conventional treatment, capturing maximum cost-effective benefits while delegating peak load reduction to engineered systems. Taken together, these results provide both a practical roadmap for implementing digital twins in real-world wastewater facilities and a scientifically novel understanding of system performance across diverse plant technologies.

CONCLUSIONS

This study successfully addresses the limitations of previous digital twin research by

demonstrating the transferability of predictive models across multiple wastewater treatment plants, clearly identifying conditions under which cross-site predictions are reliable and where they fail. Our comparative analysis confirms that Transformer and LSTM models can reduce prediction error by up to 75% relative to Random Forest baselines, though computational requirements must be weighed against application context – validating hypothesis H1.

Financial and environmental assessments demonstrate that hybrid nature-based systems offer substantive advantages: 39% reduction in life-cycle costs, 32% lower annual operating expenses, and 45% decrease in carbon emissions, confirming their economic and ecological viability for peri-urban and semi-rural applications – supporting hypothesis H3.

Monte Carlo-based uncertainty quantification replaces single-value predictions with probabilistic risk estimates, providing operators with realistic guidance for decision-making under variable conditions – validating hypothesis H2. COD, with only 60% regulatory compliance probability, is identified as the critical parameter requiring focused process attention.

Climate projections underscore the urgency of infrastructure adaptation: high-emission scenarios could render current plant designs inadequate, with violation frequencies reaching 50% by 2100. Real-time implementation is demonstrated to be feasible, with sub-50 ms prediction latency at accessible monthly costs, supporting operational deployment across utility sizes. Nature-based solution optimization confirms that pollutant removal can be maximized cost-effectively when constructed wetlands are integrated as a polishing step within conventional treatment systems.

Collectively, the seven innovations presented here contribute over 60% new content compared to prior work and advance the field by combining predictive accuracy, economic and environmental assessment, climate resilience, and practical deployability in a unified framework. This study fills a critical gap in understanding how digital twins can be transferred, scaled, and deployed in real-world wastewater operations. Future work should focus on reinforcement learning for real-time control optimization, climate adaptation cost curve analysis, and live deployment in operational facilities to validate performance under full seasonal and extreme-event variability.

REFERENCES

1. Rao, M.M., Manjunath, C., Venkatesh, D., Venkatasivanagaraju, S., Vasundhara, S. (2026). A robust offline digital twin for regenerative wastewater treatment: Multi-pollutant stress testing and resilience analysis. *Journal of Ecological Engineering*, 27(6). <https://www.jeeng.net/A-robust-offline-digital-twin-for-regenerative-wastewater-treatment-Multi-pollutant,217885,0,2.html>
2. UN-Water. (2023). *The United Nations World Water Development Report 2023: Partnerships and cooperation for water*. UNESCO, Paris.
3. Jones, E.R., van Vliet, M.T.H., Qadir, M., Bierkens, M.F.P. (2021). Country-level and gridded estimates of wastewater production, collection, treatment and reuse. *Earth System Science Data*, 13(2), 237–254. <https://doi.org/10.5194/essd-13-237-2021>
4. Madsen, H., Lawrence, D., Lang, M., Martinkova, M., Kjeldsen, T.R. (2014). Review of trend analysis and climate change projections of extreme precipitation and floods in Europe. *Journal of Hydrology*, 519, 3634–3650. <https://doi.org/10.1016/j.jhydrol.2014.11.003>
5. Arnbjerg-Nielsen, K., Willems, P., Olsson, J., Beecham, S., Pathirana, A., Bülow Gregersen, I., et al. (2013). Impacts of climate change on rainfall extremes and urban drainage systems: A review. *Water Science and Technology*, 68(1), 16–28. <https://doi.org/10.2166/wst.2013.251>
6. Behzadian, K., Kapelan, Z., Rozos, E., Makropoulos, C. (2014). Quantitative UWS performance model: WaterMet². *Procedia Engineering*, 70, 138–147. <https://doi.org/10.1016/j.proeng.2014.02.017>
7. Hamouda, M.A., Anderson, W.B., Huck, P.M. (2009). Decision support systems in water and wastewater treatment process selection and design: A review. *Water Science and Technology*, 60(7), 1757–1770. <https://doi.org/10.2166/wst.2009.538>
8. Tchobanoglous, G., Stensel, H.D., Tsuchihashi, R., Burton, F. (2014). *Wastewater Engineering: Treatment and Resource Recovery (5th ed.)*. McGraw-Hill Education, New York.
9. Mannina, G., Ekama, G., Caniani, D., Cosenza, A., Esposito, G., Gori, R., et al. (2016). Greenhouse gas emissions from wastewater treatment plants on a plantwide scale. *Journal of Environmental Management*, 175, 124–135. <https://doi.org/10.1016/j.jenvman.2016.03.037>
10. European Environment Agency. (2021). *Urban Wastewater Treatment in Europe. EEA Report No 2/2021*. EEA, Copenhagen.
11. Sweetapple, C., Fu, G., Butler, D. (2013). Identifying key sources of uncertainty in the modelling of greenhouse gas emissions from wastewater treatment plants. *Water Research*, 47(13), 4652–4665.

- <https://doi.org/10.1016/j.watres.2013.05.021>
12. Tao, M., Ota, K., Dong, M. (2022). Fusing multi-scale information for intrusion detection in digital twins of smart cities. *IEEE Transactions on Industrial Informatics*, 18(2), 1117–1125. <https://doi.org/10.1109/TII.2021.3075698>
 13. Rosen, R., von Wichert, G., Lo, G., Bettenhausen, K.D. (2015). About the importance of autonomy and digital twins for the future of manufacturing. *IFAC-PapersOnLine*, 48(3), 567–572. <https://doi.org/10.1016/j.ifacol.2015.06.141>
 14. Schütze, M., Campisano, A., Colas, H., Schilling, W., Vanrolleghem, P.A. (2004). Real time control of urban wastewater systems — Where do we stand today? *Journal of Hydrology*, 299(3–4), 335–348. <https://doi.org/10.1016/j.jhydrol.2004.08.010>
 15. Newhart, K.B., Holloway, R.W., Hering, A.S., Cath, T.Y. (2019). Data-driven performance analyses of wastewater treatment plants: A review. *Water Research*, 157, 498–513. <https://doi.org/10.1016/j.watres.2019.03.030>
 16. Zhao, L., Dai, T., Qiao, Z., Sun, P., Hao, J., Yang, Y. (2020). Application of artificial intelligence to wastewater treatment: A bibliometric analysis and systematic review. *Process Safety and Environmental Protection*, 133, 169–182. <https://doi.org/10.1016/j.psep.2019.11.014>
 17. Nourani, V., Elkiran, G., Abdullahi, J. (2019). Multi-station artificial intelligence based ensemble modeling of reference evapotranspiration. *Journal of Hydrology*, 577, 123958. <https://doi.org/10.1016/j.jhydrol.2019.123958>
 18. American Public Health Association. (2017). *Standard Methods for the Examination of Water and Wastewater (23rd ed.)*. APHA, Washington DC.
 19. Zhang, Q., Yang, L.T., Chen, Z., Li, P. (2018). A survey on deep learning for big data. *Information Fusion*, 42, 146–157. <https://doi.org/10.1016/j.inffus.2017.10.006>
 20. Bergstra, J., Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13, 281–305.
 21. Hochreiter, S., Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
 22. Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y. (2014). *Learning phrase representations using RNN encoder-decoder for statistical machine translation*. arXiv:1406.1078.
 23. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
 24. Botchkarev, A. (2018). *Performance metrics in machine learning regression, forecasting and prognostics*. arXiv:1809.03006.
 25. Short, W., Packey, D.J., Holt, T. (1995). *A Manual for the Economic Evaluation of Energy Efficiency and Renewable Energy Technologies*. NREL/TP-462-5173. NREL, Golden, CO.
 26. Jenkins, G.P., Kuo, C.Y., Harberger, A.C. (2011). *Cost-Benefit Analysis for Investment Decisions*. Cambridge Resources International.
 27. IPCC. (2006). *2006 IPCC Guidelines for National Greenhouse Gas Inventories*. IGES, Japan.
 28. Metropolis, N., Ulam, S. (1949). The Monte Carlo method. *Journal of the American Statistical Association*, 44(247), 335–341. <https://doi.org/10.1080/01621459.1949.10483310>
 29. Sin, G., Gernaey, K.V., Neumann, M.B., van Loosdrecht, M.C.M., Gujer, W. (2009). Uncertainty analysis in WWTP model applications. *Water Research*, 43(11), 2894–2906. <https://doi.org/10.1016/j.watres.2009.03.048>
 30. IPCC. (2021). *Climate Change 2021: The Physical Science Basis*. Cambridge University Press.
 31. Henze, M., Gujer, W., Mino, T., van Loosdrecht, M.C.M. (2000). *Activated Sludge Models ASM1, ASM2, ASM2d and ASM3*. IWA Publishing, London.
 32. Westra, S., Fowler, H.J., Evans, J.P., Alexander, L.V., Berg, P., Johnson, F., et al. (2014). Future changes to the intensity and frequency of short-duration extreme rainfall. *Reviews of Geophysics*, 52(3), 522–555. <https://doi.org/10.1002/2014RG000464>
 33. Kraus, M., Wittstock, V., Bodendorf, F. (2020). End-to-end architecture for real-time analytics using digital twins. *Procedia CIRP*, 91, 220–225. <https://doi.org/10.1016/j.procir.2020.02.175>
 34. Grieves, M., & Vickers, J. (2017). *Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems*. In: *Transdisciplinary Perspectives on Complex Systems*. Springer, 85–113. https://doi.org/10.1007/978-3-319-38756-7_4
 35. Vymazal, J. (2010). Constructed wetlands for wastewater treatment. *Water*, 2(3), 530–549. <https://doi.org/10.3390/w2030530>
 36. Kadlec, R.H., Wallace, S.D. (2009). *Treatment Wetlands (2nd ed.)*. CRC Press, Boca Raton.
 37. Masi, F., Rizzo, A., Regelsberger, M. (2018). The role of constructed wetlands in a new circular economy paradigm. *Journal of Environmental Management*, 216, 275–284. <https://doi.org/10.1016/j.jenvman.2017.07.066>