

Artificial Intelligence for Water Quality Prediction: From Data-Driven Models to Explainable, Physics-Informed and Generalizable Hydrological Intelligence

B Srikanth*

Yash Technologies B.V, Amsterdam, Netherlands

*Corresponding author email ID: srikanth.bhr@gmail.com

HIGHLIGHTS

- Water-quality AI is shifting from site-specific accuracy to generalizable prediction.
- Spatiotemporal and transfer-learning models address connectivity and data scarcity.
- Explainability must be separated from causal or process interpretation.
- Physics-informed learning can constrain predictions using hydrological knowledge
- Uncertainty, external validation and reproducibility are priorities for deployment.

ARTICLE INFO

Article History:

Received: 1 August 2026

Accepted: 21 September 2026

Published: 01 October 2026

Keywords:

Explainability

Artificial Intelligence

Remote sensing

Water Quality

ABSTRACT

Artificial intelligence (AI) and machine learning (ML) are increasingly used to predict water-quality dynamics from high-frequency sensors, routine monitoring, hydrometeorological observations, remote sensing and process-model outputs. However, the field has progressed faster in algorithmic complexity than in hydrological validation. This focused review synthesizes recent developments in conventional ML, deep learning, hybrid and ensemble methods, graph-based spatiotemporal learning, transfer learning, explainable AI (XAI), probabilistic prediction and physics-informed learning for surface-water quality prediction. Rather than ranking algorithms by reported accuracy, we evaluate their suitability against seven criteria: predictive skill, temporal generalization, spatial transferability, physical consistency, interpretability, uncertainty representation and reproducibility. Tree ensembles remain strong baselines for structured monitoring data; recurrent and attention-based networks are advantageous for sequential prediction; graph neural networks explicitly represent river-network connectivity; and transfer learning provides a route from data-rich to data-poor sites. Yet common practices—random data splitting, preprocessing before partitioning, single-site evaluation and deterministic point forecasts—can overstate operational skill. XAI improves model interrogation but does not establish causality, while physics-informed approaches are promising only when constraints are hydrologically defensible and independently tested. We propose a validation framework that couples chronological and spatial holdouts, event-stratified evaluation, uncertainty calibration and physical diagnostics. The next generation of water-quality AI should therefore be judged not by marginal reductions in RMSE alone, but by its ability to remain accurate, physically plausible, interpretable and uncertainty-aware under hydrological conditions that differ from the training data.

1. INTRODUCTION

Water quality is an emergent property of interacting hydrological, biogeochemical and

anthropogenic processes. Streamflow controls dilution and residence time; rainfall and runoff mobilize sediments, nutrients and microorganisms; temperature regulates reaction kinetics and oxygen solubility; and land use, point sources and in-stream

transformations determine pollutant loading. These controls vary across space and time, making prediction of dissolved oxygen (DO), turbidity, suspended solids, nutrients, chlorophyll-a, salinity, biochemical and chemical oxygen demand, microbial indicators and composite water-quality indices intrinsically nonlinear and nonstationary.

Process-based water-quality models remain essential because they encode transport, reaction and mass-balance relationships, but they can be demanding to parameterize, calibrate and operate at high temporal resolution. Data-driven AI offers a complementary strategy: relationships are learned directly from observations, sensor surrogates, hydrometeorological forcing, remote sensing and model outputs. Machine-learning applications have consequently expanded from artificial neural networks (ANNs) and support vector regression (SVR) to random forests (RF), boosting, long short-term memory (LSTM) networks, gated recurrent units (GRU), attention mechanisms, Transformers and graph neural networks (GNNs). LSTM emulators have, for example, reproduced spatiotemporal outputs of process-based reservoir models at substantially lower computational cost (Mohammed et al., 2022), while recent Journal of Hydrology studies combine sensors, remote sensing and self-attention to predict nutrients in data-scarce urban rivers (Wang et al., 2026).

The rapid growth of the literature creates a second problem: methodological novelty is often claimed from small gains in a single accuracy statistic. Existing reviews have comprehensively catalogued ML methods (Yan et al., 2024), deep-learning applications (Li et al., 2024), inland-water LSTM studies (Zhi et al., 2023) and reservoir assessment using ML, remote sensing and multivariate statistics

(Nikoo et al., 2025). A further algorithm inventory would therefore add limited value. The more consequential question is whether water-quality AI can generalize across seasons, events, monitoring stations and changing environmental regimes while respecting hydrological constraints.

This focused review addresses that gap. We synthesize recent progress through seven dimensions: predictive skill, temporal generalization, spatial transferability, physical consistency, interpretability, uncertainty representation and reproducibility. The review concentrates on surface-water prediction at river, lake, reservoir and catchment scales, where hydrological connectivity and nonstationarity are central. The objective is not to identify a universally superior algorithm, but to define what constitutes trustworthy hydrological intelligence for water-quality prediction.

2. REVIEW SCOPE AND EVIDENCE SYNTHESIS

The review was structured as a focused critical synthesis rather than a formal meta-analysis. Literature was identified using combinations of “water quality prediction/forecasting”, “machine learning”, “deep learning”, “LSTM”, “graph neural network”, “transfer learning”, “explainable AI”, “uncertainty” and “physics-informed” with emphasis on 2022–2026 peer-reviewed studies. Priority was given to hydrology and water-quality journals, studies with explicit temporal or spatial validation, and papers addressing generalization, uncertainty or physical constraints. Foundational algorithm papers were retained where necessary. Because reported performance is not directly comparable across datasets, variables, horizons and split strategies, numerical model rankings were deliberately avoided.

Table 1. Water-quality prediction as a hydrological learning problem

Target/problem	Dominant controls	Suitable model classes	Key validation risk
DO / temperature	heat exchange, flow, reaeration, metabolism	RF/XGBoost, LSTM/GRU, Transformer	seasonal persistence leakage
Turbidity / TSS	rainfall, runoff, erosion, antecedent wetness	boosting, LSTM, hybrid decomposition	event imbalance; peak underprediction
Nutrients (TN, TP, NH ₃ -N)	runoff, sources, retention, uptake	ensemble ML, attention, GNN	site transfer and load nonstationarity
Chlorophyll-a / blooms	nutrients, light, temperature, residence time	boosting, probabilistic ML, LSTM	rare-event calibration
Salinity / EC	discharge, tides, mixing, abstraction	ANN/LSTM, PINN, hybrid models	boundary-condition shift
WQI / classes	multiple physicochemical indicators	RF, XGBoost, AutoML + XAI	index circularity and class imbalance

3. MODEL FAMILIES: CAPABILITY DEPENDS ON THE PREDICTION PROBLEM

3.1 Conventional machine learning remains a necessary baseline

ANN, SVR, RF and gradient-boosting models remain highly relevant because much water-quality monitoring is tabular, irregular and modest in sample size. RF is robust to nonlinear interactions and mixed

predictor importance, while boosting methods such as XGBoost often perform strongly when lagged hydrometeorological variables and water-quality surrogates are engineered explicitly. Their principal advantage is not simplicity alone: compared with deep networks, they are easier to tune, cheaper to retrain and more compatible with feature-attribution workflows. Recent uncertainty-aware WQI studies

continue to show that relatively simple tree models can compete with more elaborate ensembles when preprocessing and optimization are rigorous (Mengistu et al., 2026).

The limitation is temporal representation. A tree model supplied only with contemporaneous predictors cannot infer memory processes such as

antecedent rainfall, delayed nutrient delivery or thermal persistence. Lags, rolling statistics and event descriptors therefore become part of the model specification. This makes feature engineering scientifically consequential: poorly selected lags can leak future information or produce apparently strong performance through autocorrelation rather than genuine forecast skill.




Figure 1. Evolution of water-quality AI from local prediction to hydrologically trustworthy intelligence

3.2 Deep learning for temporal dependence

LSTM and GRU networks dominate sequential water-quality prediction because gated recurrence can retain information over multiple time steps. They are particularly useful when high-frequency sensors produce long multivariate sequences. Multivariate LSTM models have successfully linked climatic and water-quality variables to TN, DO and WQI dynamics (Gao et al., 2023), and hybrid decomposition-recurrent architectures can isolate multi-scale signals before forecasting (Yu et al., 2022). Attention mechanisms extend recurrent models by allowing the network to weight informative time steps or variables, and self-attention LSTM has recently been used with sensor and land-use information for nutrient prediction in urban rivers (Wang et al., 2026).

Transformers offer a more direct representation of long-range dependence through self-attention and parallel sequence processing. Their theoretical attraction is substantial for water-quality systems influenced by multiple antecedent time scales, but their practical advantage is conditional on data volume and careful regularization. In small

datasets, a larger architecture can simply increase variance. For hydrological applications, deep-learning studies should therefore compare against persistence, climatology and strong tree-based baselines rather than only against older neural networks.

3.3 Spatiotemporal learning and river-network connectivity

Single-station forecasting treats monitoring sites as independent even though river systems are connected by directed flow, travel time and tributary exchange. GNNs address this mismatch by representing stations as nodes and hydrological or distance-based relations as edges. Wan et al. (2025) developed a graph-based framework that extracted temporal and spatial features across multiple river-network sites, while recent probabilistic representation networks combine graph convolutions with latent environmental-state modelling (Zhang et al., 2026). These approaches are conceptually better aligned with basin-scale prediction, but graph construction is critical. Euclidean distance alone may be hydrologically misleading; directed connectivity, discharge-dependent travel time, upstream area and source pathways are more defensible edge definitions.

**4.THE CENTRAL CHALLENGE:
GENERALIZATION UNDER HYDROLOGICAL
NONSTATIONARITY**

Water-quality time series are nonstationary because climate variability, land-use change, infrastructure, regulation, ecological succession and episodic pollution alter predictor-response relationships. Randomly partitioning temporally autocorrelated records allows near-neighbour observations from the same regime to appear in training and testing sets, producing optimistic estimates of operational skill. Chronological holdouts, blocked cross-validation and rolling-origin evaluation are therefore preferable when the stated task is forecasting. Preprocessing must also respect the split: normalization, imputation, decomposition and feature selection performed using the complete series can leak test information.

Spatial generalization is equally important. A model validated only at the site used for training demonstrates interpolation within that site's data distribution, not basin-scale transferability. Transfer

learning provides a practical response to data scarcity by pretraining on information-rich sites and fine-tuning at data-poor locations. Chen et al. (2024) coupled wavelet analysis, LSTM and transfer learning to address nonstationarity and limited river data, while feature- and model-based transfer has also been shown to move information from data-rich to data-poor sites for river DO prediction (Zhang et al., 2024). Recent Journal of Hydrology studies extend the principle to data-scarce phosphorus prediction and multimodal forecasting, with explicit tests at independent streams or fine-tuning with limited local data (Kumar and Zhang, 2026; Wang et al., 2026).

Generalization should therefore be treated as a design objective, not a post-hoc statistic. A convincing study should ask three separate questions: can the model predict later periods at the same site; can it transfer to unseen sites; and can it remain calibrated during rare or shifted hydrological regimes? These tests correspond to temporal, spatial and regime generalization and should not be conflated.

Table 2. Model families and their hydrological strengths and failure modes

Family	Main strength	Hydrological weakness	Best use
RF / boosting	strong tabular nonlinear baseline; efficient XAI	memory must be engineered	routine monitoring; surrogate prediction
SVR / GPR	effective with modest datasets; GPR can quantify uncertainty	scaling and kernel sensitivity	data-limited regression
LSTM / GRU	learns sequential memory	site-specific overfitting; opaque states	high-frequency time series
Attention / Transformer	long-range multivariate dependence	data hungry; complexity may not add skill	large multivariate sensor records
GNN / graph-temporal	explicit inter-site connectivity	graph definition can be physically weak	river networks; multi-site prediction
Transfer learning	reuses knowledge across sites/regimes	negative transfer when domains differ	data-poor monitoring sites
Physics-informed ML	constrains solutions using process knowledge	incorrect constraints can bias prediction	transport/reaction-informed prediction
Probabilistic ML	decision-relevant uncertainty	calibration can degrade under shift	risk-based early warning

5. EXPLAINABILITY IS USEFUL, BUT IT IS NOT CAUSALITY

The adoption of SHapley Additive exPlanations (SHAP), permutation importance and related XAI methods has improved transparency in water-quality modelling. XAI can identify which variables a fitted model relies upon globally and why individual predictions differ locally. Recent studies use SHAP to interpret WQI predictions (e.g., 2024–2026 applications), and explainable AutoML has identified reduced indicator sets while retaining

classification skill across spatial and temporal validation schemes (Qin et al., 2026).

The interpretation boundary is nevertheless important. A large SHAP value means that a variable influenced the model prediction conditional on the fitted model and data distribution; it does not prove that the variable physically caused the observed water-quality response. Correlated predictors can redistribute attribution, and proxy variables may receive importance because they encode season, flow regime or unmeasured processes. Hydrological XAI should therefore be paired with sensitivity analysis,

process knowledge and independent evidence. Maier et al. (2024) similarly cautioned against treating XAI as intrinsically explanatory in hydrology. The strongest use of XAI is diagnostic: detecting implausible dependencies, revealing regime-specific controls and generating hypotheses that can subsequently be tested.

6. PHYSICS-INFORMED AND HYBRID MODELLING

Purely data-driven models can produce numerically accurate yet physically impossible predictions, particularly outside the training range. Physics-informed learning attempts to reduce this risk by embedding governing equations, conservation constraints, bounds, transport relationships or process-model information into model architecture, loss functions or feature design. Salinity PINNs have demonstrated how physical information can constrain estuarine prediction (Water, 2023), while physics-informed ensemble ML has incorporated advection, dispersion and reaction information into dynamic water-quality prediction (Moeini and Abokifa, 2025).

A second pathway is process emulation rather than equation-constrained learning. Mohammed et al. (2022) trained LSTM models to emulate a calibrated hydrodynamic-water-quality model, providing rapid predictions of microbial and metal concentrations. This approach transfers process-model behaviour into a computational surrogate, but it also transfers the assumptions and structural errors of the teacher model. Hybrid systems should therefore be validated against observations, not solely against process-model output. Physics-informed does not automatically mean physically correct. A penalty based on an oversimplified equation can force a model toward the wrong solution; poorly observed boundary conditions can dominate apparent physical consistency; and multiple processes may be equifinal. Physical constraints should be justified at the spatial and temporal scale of application and evaluated with diagnostics independent of the training loss.

7. UNCERTAINTY AND DECISION RELEVANCE

Most water-quality AI produces deterministic point estimates, yet management decisions depend on the probability of threshold exceedance and the cost of false alarms or missed events. Uncertainty arises from measurement error, sparse sampling, model parameters, structural assumptions and distribution shift. Ensemble spread, Gaussian processes, Bayesian

neural networks, quantile regression and conformal prediction provide different routes to prediction intervals. Recent studies explicitly combine ML with Bayesian or probabilistic uncertainty for WQI and near-term water-quality forecasting (Mengistu et al., 2026; Shukla et al., 2026).

Uncertainty intervals must themselves be validated. Coverage probability, interval width, reliability diagrams and threshold-event calibration are more informative than visually plausible bands. Moreover, uncertainty estimated under the training distribution can become overconfident after a hydrological regime shift. Operational systems should therefore distinguish aleatoric uncertainty from epistemic/model uncertainty where possible and flag out-of-distribution inputs rather than silently extrapolating.

8. A VALIDATION FRAMEWORK FOR TRUSTWORTHY WATER-QUALITY AI

The field needs a common minimum validation standard. Table 3 proposes a framework that separates accuracy from scientific credibility. No single metric can demonstrate trustworthiness: high R^2 can coexist with biased peaks, physically implausible responses or poor transfer. Conversely, a modest reduction in average RMSE may be operationally valuable if it improves event detection or uncertainty calibration.

Three methodological errors deserve explicit attention. First, all transformations that learn from data—including scaling, imputation, feature selection and decomposition—must be fitted inside the training fold. Second, model selection should not repeatedly inspect the final test set; a separate validation period or nested procedure is needed. Third, persistence and climatological baselines should be reported for forecasting. Without them, a sophisticated model can appear accurate simply because water quality is strongly autocorrelated (Figure 2).

Evaluation should also reflect the management objective. For eutrophication or pollution warnings, peak timing, threshold exceedance, recall and lead time may matter more than average RMSE. For long-term planning, bias and transfer under altered flow regimes may dominate. Model evaluation is therefore inseparable from the decision context (Table 3).

Three pathways for integrating hydrology and AI in water-quality prediction




Figure 2. Three pathways for integrating hydrology and AI in water-quality prediction.

Table 3. Proposes a framework that separates accuracy from scientific credibility

Dimension	Minimum evidence	Stronger evidence
Predictive skill	MAE/RMSE plus R ² or NSE	event- and season-stratified errors; baseline skill
Temporal generalization	strict chronological holdout	rolling-origin or blocked validation
Spatial transfer	independent monitoring site	leave-site/basin-out validation
Physical consistency	bounds and directional checks	mass/transport/process diagnostics
Interpretability	feature importance	stable XAI + sensitivity + process corroboration
Uncertainty	prediction intervals	coverage, sharpness, reliability, threshold calibration
Reproducibility	data period, hyperparameters	code/ data, seeds, full preprocessing pipeline

9. RESEARCH PRIORITIES

Hydrologically structured spatiotemporal models. Future GNNs should encode directed river topology, travel time and dynamic discharge rather than generic spatial adjacency.

Transfer under nonstationarity. Domain adaptation and transfer learning should be evaluated for negative transfer and tested across climate, land-use and pollution regimes.

Multimodal prediction. High-frequency sensors, laboratory measurements, discharge, meteorology, remote sensing and catchment attributes should be fused with explicit treatment of different temporal resolutions.

Probabilistic early warning. Forecasts should provide calibrated exceedance probabilities and decision thresholds, not only deterministic concentration estimates. Process-guided learning. Physical constraints should target well-supported processes

and be tested independently; hybrid process-AI models should quantify both observational and structural uncertainty.

Benchmarking and reproducibility. Shared multi-basin datasets, standardized temporal/spatial splits and transparent preprocessing are needed to

determine whether architectural advances genuinely generalize.




Figure 3. A Seven-Axis Validation Framework for Hydrologically Credible and Trustworthy Water-Quality AI

10. CONCLUSIONS

AI has become a powerful component of water-quality prediction, but the scientific frontier is no longer the substitution of one algorithm for another. Tree ensembles remain competitive for structured monitoring data; recurrent and attention-based networks exploit temporal dependence; GNNs introduce river-network structure; transfer learning addresses uneven monitoring density; probabilistic models expose predictive uncertainty; and physics-informed approaches offer a route toward physically constrained extrapolation. The critical limitation across these families is generalization. Models trained and tested within one site or one hydrological regime can report excellent accuracy while failing when conditions change.

For hydrological science, trustworthy water-quality AI should therefore satisfy more than predictive accuracy. It should be tested on future periods and unseen locations, remain physically plausible, communicate uncertainty, support interpretable diagnostics and be reproducible. XAI should be used to interrogate models rather than claim causality, and physics-informed constraints should be independently validated rather than assumed to guarantee realism. Progress will be best measured by robust performance under hydrological change and by demonstrable value for monitoring and management—not by marginal improvements in a single error metric.

DECLARATIONS

Conflicts of Interest: The author declares no conflicts of interest.

Funding: No external funding was reported in the source thesis.

Author Contributions: The sole author was responsible for the conceptualization, methodology, data collection and curation, analysis, visualization, interpretation of results, writing of the original draft, and review and editing of the manuscript. The author has read and approved the final version of the manuscript.

AI Disclosure: The authors declare that artificial intelligence (AI) tools were used solely to assist in language refinement, organization, and editing of the manuscript. AI was not used to generate original research data, conduct analyses, or draw scientific conclusions. All technical content, interpretations, and final decisions were developed and verified by the authors, who take full responsibility for the accuracy, integrity, and originality of the work.

REFERENCES

- Abbas, F., Cai, Z., Shoaib, M., Iqbal, J., Ismail, M., Arifullah, Alrefaei, A.F., Albeshr, M.F., 2024. Machine learning models for water quality prediction: A comprehensive analysis and uncertainty assessment in Mirpurkhas, Sindh, Pakistan. *Water* 16, 941. <https://doi.org/10.3390/w16070941>.
- Breiman, L., 2001. Random forests. *Machine Learning* 45, 5–32. <https://doi.org/10.1023/A:1010933404324>.
- Chen, S., Huang, J., Wang, P., Tang, X., Zhang, Z., 2024. A coupled model to improve river water quality prediction towards addressing non-stationarity and data limitation. *Water Research* 248, 120895. <https://doi.org/10.1016/j.watres.2023.120895>.
- Chen, T., Guestrin, C., 2016. XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. <https://doi.org/10.1145/2939672.2939785>.
- Maier, H.R., Taghikhah, F.R., Nabavi, E., Razavi, S., Gupta, H., Wu, W., Radford, D.A.G., Huang, J., 2024. How much X is in XAI: Responsible use of “Explainable” artificial intelligence in hydrology and water resources. *Journal of Hydrology X* 25, 100185. <https://doi.org/10.1016/j.hydroa.2024.100185>.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Computation* 9, 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Kumar, A., Zhang, K., 2026. Improving total phosphorus prediction and monitoring in data scarce streams using a hybrid machine learning framework. *Journal of Hydrology* 665, 134636. <https://doi.org/10.1016/j.jhydrol.2025.134636>.
- Li, W., Zhao, Y., Zhu, Y., Dong, Z., Wang, F., Huang, F., 2024. Research progress in water quality prediction based on deep learning technology: a review. *Environmental Science and Pollution Research* 31, 26415–26431. <https://doi.org/10.1007/s11356-024-33058-7>.
- Mengistu, T.D., Chung, I.-M., Chang, S.W., 2026. Machine learning for water quality prediction and uncertainty assessment. *Physics and Chemistry of the Earth* 143, 104319. <https://doi.org/10.1016/j.pce.2026.104319>.
- Moeini, M., Abokifa, A.A., 2025. Physics-informed ensemble machine learning for dynamic water quality prediction in water distribution systems. *Journal of Water Process Engineering* 79, 108841. <https://doi.org/10.1016/j.jwpe.2025.108841>.
- Mohammed, H., Tornyeviadzi, H.M., Seidu, R., 2022. Emulating process-based water quality modelling in water source reservoirs using machine learning. *Journal of Hydrology* 609, 127675. <https://doi.org/10.1016/j.jhydrol.2022.127675>.
- Nikoo, M.R., Al Aamri, A., Etri, T., Al-Rawas, G., 2025. A review of machine learning, remote sensing, and statistical methods for reservoir water quality assessment. *Journal of Hydrology* 659, 133323. <https://doi.org/10.1016/j.jhydrol.2025.133323>.
- Pang, Z., Zhou, Z., Fu, J., Jiang, W., Qin, X., Sun, M., 2025. Deep learning-based remote sensing retrieval of inland water quality: A review. *Journal of Hydrology: Regional Studies* 61, 102759. <https://doi.org/10.1016/j.ejrh.2025.102759>.
- Qin, W., Zhong, H., Hu, J., Luo, W., Shao, Z., Hou, H., Liu, Z., Zhang, X., Liang, S., Yang, J., 2026. Explainable AutoML-driven surface water quality classification with key indicators identification. *Journal of Hydrology* 665, 134634. <https://doi.org/10.1016/j.jhydrol.2025.134634>.
- Shukla, R., Carey, C.C., Breef-Pilz, A., et al., 2026. Advancing near-term water quality forecasting with explainable machine learning and probabilistic uncertainty quantification. *ACS ES&T Water*. <https://doi.org/10.1021/acsestwater.6c00757>.
- Thi, H.P., Thanh, P.N., Dinh, V.N., Truong, P., Anh, D.T., Le, N.A., Mai, T.S., Van, S.P., Nguyen, T.B.H., Van Tung, P., 2026. Uncertainty quantification in water quality prediction using Optuna-optimized machine learning: A case study in the Quan Lo-Phung Hiep irrigation system, Vietnamese Mekong Delta. *Journal of Water Process Engineering* 86, 109858. <https://doi.org/10.1016/j.jwpe.2026.109858>.
- Wan, H., Xiang, L., Cai, Y., Xie, Y., Xu, R., 2025. Temporal and spatial feature extraction using graph neural networks for multi-point water quality prediction in river network areas. *Water Research* 281, 123561. <https://doi.org/10.1016/j.watres.2025.123561>.
- Wang, Z., Osmond, P., Shi, B., Janmohammadi, M., Zhang, K., 2026. Next-generation water quality monitoring: sensor-based deep learning prediction and calibration optimization in urban rivers. *Journal of Hydrology* 666, 134715. <https://doi.org/10.1016/j.jhydrol.2025.134715>.
- Wang, Z., Shi, B., Osmond, P., Zhang, K., 2026. Multimodal deep-learning-driven water-quality forecasting and scenario assessment: shifting from reactive monitoring to proactive basin management. *Journal*

- of Hydrology 675, 135554.
<https://doi.org/10.1016/j.jhydrol.2026.135554>.
- Yan, X., Zhang, T., Du, W., Meng, Q., Xu, X., Zhao, X., 2024. A comprehensive review of machine learning for water quality prediction over the past five years. Journal of Marine Science and Engineering 12, 159.
<https://doi.org/10.3390/jmse12010159>.
- Yao, J., Li, B., Liu, Z., Chen, S., Ruan, X., 2025. Explainable deep learning for dual goals: Predicting total phosphorus concentrations and identifying pollution sources in a tidal river network. Journal of Hydrology 663, 134203.
<https://doi.org/10.1016/j.jhydrol.2025.134203>
- Yuan, M., Li, Y., Zhang, L., Zhao, W., Zhang, X., Li, J., 2025. Rapid prediction approach for water quality in plain river networks: A data-driven water quality prediction model based on graph neural networks. Water 17, 2543.
<https://doi.org/10.3390/w17172543>.
- Zhang, Q., Chen, L., Chen, H., 2026. PRO-Net: Probabilistic representation network for spatio-temporal water quality prediction. Journal of Hydrology 664, 134341.
<https://doi.org/10.1016/j.jhydrol.2025.134341>
- Zhi, W., Feng, D., Tsai, W.-P., Sterle, G., Harpold, A., Shen, C., Li, L., 2023. Long short-term memory models of water quality in inland water environments. Water Research X 21, 100207.
<https://doi.org/10.1016/j.wroa.2023.100207>.