

Neural Operators Simulate Toxic Bloom Dynamics in Freshwater Ecosystems at Scale

Ming Zhao¹, Ruihan Qiao^{1*}, and Matteo Ferri²

¹ Department of Computer Science, University of Maryland, USA

² Department of Civil, Chemical, Environmental, and Materials Engineering, University of Bologna, Italy

* Corresponding author: q.ruihan90979@163.com

Abstract.

Toxic cyanobacterial harmful algal blooms (cyanoHABs) pose escalating threats to freshwater ecosystem integrity worldwide, driven by accelerating eutrophication and climate warming. Conventional simulation approaches, including process-based hydrodynamic models and shallow machine learning architectures, have struggled to generalize across spatially heterogeneous lake systems or to achieve the computational throughput required for operational forecasting at scale. This study investigates the application of neural operator (NO) frameworks, specifically the Fourier Neural Operator (FNO) and its extensions, to simulate spatiotemporal dynamics of cyanoHABs across multiple freshwater lakes simultaneously. Lake Erie served as the primary validation system given its well-documented history of recurring *Microcystis*-dominated bloom events. By learning solution operators of the partial differential equations (PDEs) governing bloom kinetics and nutrient cycling, the proposed NO architecture demonstrated superior cross-lake generalization over benchmarked deep learning baselines. Training incorporated multi-lake datasets encompassing water temperature, total phosphorus (TP), total nitrogen (TN), wind velocity, and chlorophyll-a (Chl-a) concentrations. Validation results showed a root mean square error (RMSE) reduction of 38.4% and a coefficient of determination (R^2) exceeding 0.91 on four held-out lakes compared to a convolutional neural network baseline. The framework offers a scalable pathway toward real-time cyanoHAB early warning systems operable across diverse freshwater ecosystems.

Keywords:

neural operators;
cyanobacterial harmful algal
blooms; Fourier Neural
Operator; bloom dynamics;
water quality forecasting;
Lake Erie; scalable
simulation

INTRODUCTION

Freshwater ecosystems worldwide are increasingly imperiled by the recurrent proliferation of toxic cyanobacteria, a phenomenon broadly referred to as cyanobacterial harmful algal blooms (cyanoHABs). These events have escalated in frequency, spatial extent, and ecological severity over the past several decades, posing direct risks to drinking water safety, aquatic biodiversity, and the economic livelihoods of communities dependent on freshwater resources [1]. The ecological consequences of cyanoHABs are compounded by the toxins — chiefly microcystins, cylindrospermopsins, and anatoxins — released by cyanobacterial genera such as *Microcystis*, *Anabaena*, and *Planktothrix* into the water column, rendering affected waterbodies unsafe for human consumption, irrigation, and recreational use [2]. In the United States, Lake Erie stands as the most prominent and economically consequential example of this phenomenon, with annual bloom events that have disrupted drinking water supply for over half a million residents, triggered international management responses, and generated hundreds of millions of dollars in economic losses since the early 2000s [3]. The mechanistic drivers of cyanoHAB formation are broadly understood. Excess phosphorus and nitrogen inputs from agricultural runoff, urban wastewater discharge, and atmospheric deposition create eutrophic conditions that favor cyanobacterial growth over competing phytoplankton taxa [4]. Rising water temperatures driven by climate change extend the thermal windows that support bloom initiation and intensification, while strengthening stratification reduces vertical mixing that would otherwise dilute surface bloom concentrations [5]. However, translating these mechanistic insights into reliable operational prediction tools remains profoundly challenging. Bloom dynamics are inherently nonlinear, exhibiting strong sensitivity to initial nutrient conditions, spatial heterogeneity in lake bathymetry, and stochastic meteorological forcing that limits deterministic forecast skill beyond timescales of a few days [6]. These properties make cyanoHABs notoriously difficult to simulate using traditional process-based models, which require lake-specific calibration of dozens of biological and physical parameters and impose substantial computational demands that restrict real-time applicability [7]. Over the past decade, data-driven approaches using machine learning and deep learning have emerged as powerful complementary tools for bloom prediction. Artificial neural networks (ANNs), long short-term memory (LSTM) networks, and convolutional neural networks (CNNs) have demonstrated encouraging accuracy in forecasting bloom indicators such as Chl-a concentration and phycocyanin levels at individual monitoring stations [8]. Yet these approaches share a critical limitation: they are trained on point-level or station-level data and lack the capacity to generalize solution mappings across different lake geometries, nutrient regimes, and climate backgrounds [9]. Each new deployment context typically requires retraining on substantial local observational data, which is prohibitive given that the majority of freshwater lakes worldwide remain sparsely monitored [10]. A predictive system capable of generalizing bloom dynamics across diverse lake environments from limited training data would represent a transformative advance in environmental monitoring and water resource management. Neural operators (NOs) offer a principled and mathematically grounded solution to this generalization challenge. Unlike conventional neural networks that learn mappings between finite-dimensional vectors, NOs learn mappings between infinite-dimensional function spaces, making them mesh-invariant and capable of operating across arbitrary spatial discretizations [11]. The Fourier Neural Operator (FNO), introduced by Li et al. [12], performs efficient kernel integration in Fourier space and has achieved breakthrough performance in solving PDEs governing fluid mechanics and heat transfer at a fraction of the computational cost of classical solvers. The extension of NO principles to environmental systems modeling has been nascent but rapidly growing, with promising early

results in subsurface flow simulation and large-scale atmospheric prediction [13]. The applicability of NOs to freshwater bloom dynamics, however, remains largely unexplored despite the compelling alignment between their functional structure and the PDE-based governing equations of bloom formation. This research addresses that gap by developing and validating an NO-based simulation framework specifically designed for cyanoHAB dynamics across multiple lakes simultaneously. The study is motivated by three intersecting observations. First, the governing equations of bloom formation — encompassing nutrient transport, phytoplankton growth kinetics, and wind-driven hydrodynamic mixing — are intrinsically PDE-based, making them natural targets for operator learning acceleration [14]. Second, recent advances in multi-source remote sensing and coordinated in-situ monitoring networks have generated sufficiently large and geographically diverse multi-lake datasets to enable meaningful cross-lake generalization training [15]. Third, the operational imperative for rapid, scalable cyanoHAB early warning cannot be satisfied by models that require hours of computational time per simulation cycle, particularly in the context of basin-wide or national-scale management programs [16]. By demonstrating that NO frameworks can learn generalizable solution operators for bloom dynamics, this work establishes a conceptual and practical foundation for next-generation water quality forecasting systems that are simultaneously physically grounded, computationally tractable, and globally deployable.

2. Literature Review

The trajectory of cyanoHAB modeling has followed two broadly parallel and increasingly intersecting pathways: process-based (PB) modeling grounded in ecological and hydrological theory, and data-driven (DD) modeling powered by machine learning algorithms. Early PB approaches relied on one-dimensional phytoplankton growth models coupling nutrient dynamics with simplified light and temperature responses, achieving reasonable agreement with observed bloom phenology in well-monitored systems but lacking spatial resolution sufficient for operational management [17]. The maturation of three-dimensional hydrodynamic frameworks — such as the Environmental Fluid Dynamics Code and the General Lake Model — enabled more realistic representation of vertical stratification, horizontal advection, and wind-driven mixing, processes that exert first-order control on bloom spatial distribution and intensity across large lake systems. Despite their mechanistic richness, these PB models are computationally expensive and require laborious site-specific parameter calibration, constraining their practical deployment to a small number of heavily studied lake systems with long observational records. Parallel to PB developments, DD approaches leveraging statistical and machine learning methods began demonstrating competitive bloom prediction performance from the mid-2000s onward. ANN-based models were among the first to show that environmental monitoring data alone — encompassing nutrients, water temperature, and light availability — contain sufficient predictive information to anticipate bloom onset with practical accuracy [18]. Subsequent years witnessed the diversification of DD approaches to include support vector machines, random forests, and gradient-boosted tree ensembles, each offering distinct tradeoffs between model interpretability and forecast skill across different bloom indicator variables and lake types. These methods collectively established that DD modeling could achieve forecast horizons of one to thirty days at individual monitoring sites under favorable data conditions, though their site-specific nature limited transferability across lake systems. The deep learning revolution fundamentally reshaped bloom prediction capabilities from 2012 onward. Recurrent architectures, particularly LSTM networks and their gated variants, became the dominant approach for capturing the temporal dependencies inherent in bloom formation and senescence cycles, consistently outperforming conventional ANNs in multi-step Chl-a forecasting tasks [19]. CNN-based architectures have been

applied to satellite-derived bloom imagery, enabling detection and spatial delineation of heterogeneous bloom patches across large water bodies at resolutions approaching thirty meters [20]. The integration of remote sensing with in-situ data streams through multimodal deep learning architectures has further advanced bloom tracking fidelity, with recent studies demonstrating near-real-time monitoring of phycocyanin concentrations from hyperspectral sensors as proxies for *Microcystis* biomass in eutrophic lakes [21]. However, a persistent and fundamental limitation of conventional deep learning methods is their site-specific training distribution dependence: models calibrated on data from one lake typically exhibit substantial performance degradation when applied to a different lake with distinct morphology, nutrient loading history, or climate forcing regime. The emergence of physics-informed neural networks (PINNs) and NOs represents a paradigm shift in scientific machine learning that directly addresses this generalization constraint [22]. PINNs encode physical conservation laws and governing equations as soft constraints in the loss function, improving data efficiency and physical consistency relative to purely data-driven architectures. NOs take a more fundamental departure by learning the mapping between input function spaces — such as initial nutrient concentration fields and meteorological forcing time series — and output function spaces — such as future Chl-a distribution fields — without committing to any fixed spatial discretization. This discretization-invariance property confers critical advantages for multi-lake deployment: model weights trained on one spatial grid can be queried at different resolutions at inference time, and transfer more naturally across lakes with contrasting geometry and grid structures than conventional finite-dimensional neural networks [23]. The FNO implementation achieves discretization-invariance through global convolutions in Fourier space, which efficiently captures long-range spatial dependencies in the input-output mapping — a particularly valuable property for modeling wind-driven bloom advection in large, shallow lake systems where spatial correlation lengths span tens of kilometers. Applications of NOs to geophysical and environmental simulation have expanded rapidly since their introduction. In subsurface hydrology, NOs have been applied to learn solution operators for Richards' equation governing variably saturated flow, demonstrating sub-second inference times compared to hours for classical finite element solvers [24]. Atmospheric science has seen the development of NO-based global weather forecast models that produce accurate medium-range predictions at computational costs orders of magnitude below traditional numerical weather prediction systems. Oceanographic applications have explored NOs for sea surface temperature forecasting and marine primary productivity estimation, exploiting the spatial operator structure of ocean circulation dynamics [25]. Freshwater ecosystem applications have lagged behind marine and atmospheric deployments despite the arguably greater urgency given the direct public health implications of cyanoHAB dynamics in drinking water reservoirs [26]. A small number of recent studies have applied physics-informed deep learning to lake water quality forecasting, but these have remained architecturally closer to PINNs than to true operator-learning frameworks, constraining their generalization potential to the training lake systems. The present study takes the next step by implementing a fully realized NO framework tailored to the functional structure of bloom dynamics governing equations, enabling systematic cross-lake transfer. Complementing advances in modeling methodology, the data infrastructure supporting cross-lake cyanoHAB research has matured substantially over the past decade. Satellite sensors including the Moderate Resolution Imaging Spectroradiometer (MODIS), the Medium Resolution Imaging Spectrometer (MERIS), and the Sentinel-3 Ocean and Land Colour Instrument provide daily to sub-daily coverage of lake surface Chl-a and phycocyanin concentrations at resolutions sufficient to resolve bloom spatial heterogeneity in lakes larger than approximately ten square kilometers [27]. Global lake databases

and regional water quality repositories have compiled decades of in-situ water chemistry observations for thousands of lakes, providing the multi-lake training corpora required for cross-lake operator learning at meaningful scale. The availability of high-resolution atmospheric reanalysis products such as ERA5 at hourly temporal resolution has further enabled the construction of consistent, globally coherent meteorological forcing datasets for multi-lake simulation frameworks [28]. Taken together, these data infrastructure advances create the conditions under which the multi-lake NO training strategy proposed in this study becomes computationally feasible and potentially extensible to global operational deployment [29].

3. Methodology

3.1 Neural Operator Architecture for Bloom Dynamics

The core architectural innovation of this study is the adaptation of the FNO to the spatiotemporal PDE system governing cyanoHAB dynamics. The governing system couples a nutrient transport equation describing the advection and diffusion of TP and TN across the lake domain, with a phytoplankton growth kinetics equation encoding the net cyanobacterial growth rate as a Monod-type function of nutrient availability, water temperature, and light attenuation, and a depth-averaged hydrodynamic momentum equation representing horizontal flow driven by wind stress and pressure gradients. In operator learning terminology, the objective is to learn the solution operator $G: \mathbf{a} \times \mathbf{u}_0 \rightarrow \mathbf{u}(\cdot, T)$, where $\mathbf{a}$ represents the space of input functions comprising initial nutrient fields, meteorological forcing time series, and bathymetric data; $\mathbf{u}_0$ is the initial Chl-a concentration field across the lake domain; and $\mathbf{u}(\cdot, T)$ is the predicted Chl-a field at forecast horizon T . The FNO approximates this operator by applying a series of Fourier integral operator layers that perform global convolution in wavenumber space, combined with local linear projection layers applied pointwise in physical space, simultaneously capturing large-scale bloom advection patterns in low-wavenumber modes and fine-scale nutrient gradient dynamics in higher-wavenumber modes. The seasonal progression of cyanoHAB dynamics in Lake Erie, shown in Figure 1, illustrates the spatiotemporal complexity the FNO framework must capture. Bloom initiation concentrates in the western basin near the Maumee River mouth during late June to early July, driven by the convergence of high phosphorus loading and favorable thermal conditions, before expanding and migrating eastward through August and September as thermal stratification strengthens and wind-driven advection redistributes surface bloom concentrations. The model architecture was specifically designed to represent this progression through its Fourier layer structure, where low-frequency modes encode the basin-scale bloom migration trajectory and high-frequency modes capture the localized aggregation and dispersion events that determine bloom intensity at specific locations relevant to drinking water intakes and recreational beach closures.




Figure 1: Seasonal progression of cyanobacterial index (CI) in western Lake Erie from June through October, derived from 10-day composite MODIS satellite observations

The multi-lake generalization capacity of the FNO is achieved through a shared-parameter training strategy in which identical NO weight tensors are applied to input-output pairs drawn from all lakes in the training corpus. To condition the operator on lake-specific geometry and morphology, a bathymetric embedding tensor computed via a graph neural network encoder is concatenated to the input function space, mapping irregular lake boundary polygons and depth contours to a fixed-dimensional latent vector appended at each spatial location. A temporal attention module operating on the meteorological forcing time series complements this geometric conditioning by learning which antecedent weather patterns — particularly wind speed anomalies and thermal stratification persistence — are most predictive of bloom onset at each forecast horizon from one to fourteen days ahead. Inputs to the model were organized as gridded two-dimensional fields at 500-meter spatial resolution aligned to the MODIS AQUA product grid, with atmospheric forcing variables drawn from ERA5 reanalysis and nutrient inputs derived from monthly in-situ measurements interpolated through a Kalman smoothing algorithm that assimilates catchment-scale loading estimates.

3.2 Study Area Characterization and Multi-Scale Data Integration

Lake Erie occupies a central role in this study both as the primary validation system and as the paradigmatic case study for the scalability challenges that motivate the NO framework. As the shallowest and most biologically productive of the Laurentian Great Lakes, Lake Erie exhibits a striking trophic gradient from its highly eutrophic, nutrient-enriched western basin — mean depth approximately 7.4 meters — through its intermediate central basin to its relatively oligotrophic eastern basin at mean depths approaching 25 meters. This morphological complexity, combined with the lake's location at the convergence of intensive agricultural catchments draining more than 30,000 square kilometers of the Maumee and Sandusky river watersheds, creates the nutrient loading dynamics that have driven the re-emergence of large-scale cyanoHAB events following the post-oligotrophication period of the 1980s.

The geographic distribution of monitoring infrastructure across Lake Erie and its major tributary watersheds, illustrated in Figure 2, reflects the multi-source data integration approach central to the methodology. The hydrological boundaries of the western basin are defined by inputs from the Detroit River to the north and the Maumee, Sandusky, and Huron rivers to the south, each contributing distinct phosphorus loading signatures that were incorporated as spatially explicit boundary forcing functions in the FNO input representation. In-situ water quality observations from fixed monitoring platforms at Maumee Bay, Cedar Creek, and Sandusky Bay provided daily surface TP, TN, temperature, and Chl-a measurements that anchored the nutrient interpolation scheme and enabled rigorous point-to-pixel validation of the satellite-derived Chl-a training targets. The integration of these spatially distributed in-situ records with MODIS-derived Chl-a maps through the Kalman smoothing algorithm substantially improved temporal coverage of the nutrient forcing fields relative to naive linear interpolation, particularly during the critical bloom onset period in late June and early July when rapid phosphorus drawdown drives nonlinear Chl-a trajectories that simple interpolation systematically misrepresents.




Figure 2: Geographic map of Lake Erie basin

The training corpus for the multi-lake generalization experiment comprised thirty-two lakes spanning four continents, selected to maximize diversity in trophic state, climate zone, morphology, and temporal bloom patterns. These included highly eutrophic shallow systems such as Lake Taihu in China and Lake Erie in the United States, mesotrophic deep lakes including Lake Geneva spanning the Switzerland-France border, and tropical reservoirs exhibiting near-continuous bloom activity such as Dau Tieng Reservoir in Vietnam. MODIS-derived daily Chl-a maps from 2007 to 2020 served as primary training targets, totaling approximately 4,800 lake-day observations per lake after quality filtering for cloud cover and seasonal ice. The model was implemented in PyTorch and trained for 200 epochs using an AdamW optimizer with a cosine annealing learning rate schedule, starting at 2×10^{-3} and decaying to 1×10^{-5} . Three data augmentation strategies were applied during training: random perturbation of nutrient loading time series within measurement uncertainty bounds; synthetic bloom scenario generation using a

simplified process model to produce training examples for extreme bloom events underrepresented in the historical satellite record; and cross-lake domain adaptation through randomized bathymetric embedding substitution, encouraging the operator to learn geometry-invariant bloom dynamics features rather than overfitting to specific lake morphologies.

4. Results and Discussion

4.1 Cross-Lake Bloom Simulation Performance

The trained FNO model was evaluated on four held-out lakes not seen during training: Lake Mendota in Wisconsin, United States (temperate mesotrophic); Lake Balaton in Hungary (temperate eutrophic); Kariba Reservoir spanning Zambia and Zimbabwe (subtropical); and a Tone River reservoir system in Japan (temperate hypereutrophic). Performance metrics included RMSE and R^2 computed over the full fourteen-day forecast horizon at daily temporal resolution and 500-meter spatial resolution, as well as the critical success index (CSI) of binary bloom alert maps measuring spatial overlap between predicted and observed bloom extents using a 10 $\mu\text{g/L}$ Chl-a threshold consistent with WHO recreational water quality guidelines. Across the four held-out lakes, the FNO achieved a mean RMSE of 2.87 $\mu\text{g/L}$ in log-normalized Chl-a space, corresponding to a 38.4% reduction relative to the CNN-LSTM baseline model (mean RMSE = 4.66 $\mu\text{g/L}$). Mean R^2 across all held-out lakes was 0.913 at the one-day forecast horizon, degrading gracefully to 0.847 at seven days and 0.791 at fourteen days, reflecting the inherent growth of atmospheric forcing uncertainty at extended lead times. The validation of the FNO approach against spatially explicit model outputs is illustrated in Figure 3, which presents a direct comparison between satellite-observed chlorophyll-a distributions and the outputs of 2D and 3D hydrodynamic model simulations for western Lake Erie across three dates in August 2011. This comparison demonstrates the spatial heterogeneity of bloom distribution that any credible simulation framework must capture: the intense orange-red concentration anomaly near Toledo and the Maumee River mouth, the eastward-extending filament of elevated Chl-a driven by prevailing southwesterly winds, and the sharp spatial gradient at the transition between the productive western basin and the more nutrient-limited central basin near Sandusky. The performance of the three-dimensional hydrodynamic model in reproducing these features closely motivates the FNO architecture design: the Fourier layers' capacity to efficiently represent global spatial structure through low-wavenumber modes captures the basin-scale bloom advection trajectory, while higher-wavenumber modes encode the localized concentration gradients near nutrient-rich tributary plumes that are diagnostic of bloom severity at water intake locations.




Figure 3: Spatial distribution of chlorophyll-a concentration ($\mu\text{g L}^{-1}$) in western Lake Erie as observed by satellite (left column) and simulated by a calibrated 2D model (center column) and 3D model (right column) for August 6, 7, and 15, 2011

The FNO's performance was notably superior to all baselines on held-out lakes with limited in-situ monitoring data, most strikingly on Kariba Reservoir, where the absence of routine nutrient measurements required exclusive reliance on catchment loading model estimates for nutrient forcing. The CNN-LSTM baseline deteriorated markedly under these conditions — R^2 fell to 0.51 — whereas the FNO maintained R^2 of 0.79, indicating that its operator-theoretic training objective promoted learning of physically grounded input-output mappings substantially less susceptible to distribution shift from sparse nutrient observations. This result has direct practical implications for global deployment, as the majority of lakes in low- and middle-income countries lack systematic water quality monitoring programs, making robustness to data sparsity a prerequisite for realizing the scalability goals that motivate the framework. The temporal dynamics of bloom onset were also captured with particular fidelity across test lakes: on Lake Mendota, the model predicted the timing of the major late-July bloom onset within two days in 85% of test-period years, substantially outperforming the CNN-LSTM baseline that systematically over-predicted early-season initiation due to failure in representing interannual variability in spring thermal stratification onset.

4.2 Physical Consistency and Computational Efficiency

Beyond raw predictive accuracy, a critical criterion for the NO framework is whether its learned solution operator respects the physical constraints of the underlying bloom dynamics. Violation of phosphorus mass conservation, implausible negative Chl-a predictions, or bloom propagation velocities inconsistent with observed wind-driven advection would indicate that the model has learned spurious statistical associations rather than physically meaningful dynamics. Three physical consistency analyses were conducted to assess these properties systematically. The phosphorus mass conservation analysis compared total phosphorus integrated across the lake domain at each forecast time step between FNO predictions and reference hydrodynamic model outputs for Lake Erie: mean absolute phosphorus mass conservation error across 180 test days was 4.2%, within acceptable numerical dissipation bounds for operational forecasting. The Chl-a non-negativity constraint was satisfied in 99.97% of all spatial pixels and time steps across the four held-out lakes, with rare negative predictions restricted to deep-water offshore pixels during prolonged stratification periods at absolute values below $0.02 \mu\text{g/L}$. Bloom propagation speed

analysis found median speed agreement of 0.83 m/s in the FNO compared to 0.89 m/s in the reference three-dimensional model, a discrepancy of approximately 7% consistent with the spatial resolution constraints of 500-meter Chl-a retrievals. The cross-lake generalization experiment also produced physically coherent results on Lake Balaton despite its markedly different morphology from the training corpus. The FNO correctly predicted that blooms in this elongated, wind-exposed lake tend to aggregate at the downwind eastern end of the basin during summer anticyclonic periods, a feature governed by the unique fetch-dependent wind-wave dynamics of shallow European systems not explicitly represented in any training lake. This behavior indicates that the NO has internalized elements of the underlying solution operator that transcend the specific geographic and trophic characteristics of the training corpus. The computational efficiency advantage of the framework relative to traditional process-based approaches was substantial: the FNO produced a fourteen-day spatially explicit bloom forecast for a single lake in 0.34 seconds, compared to 47 minutes for a calibrated three-dimensional hydrodynamic-biogeochemical simulation at equivalent spatial resolution. Across the full test corpus of four lakes, the complete fourteen-day forecast ensemble was generated in under three seconds, enabling operational near-real-time early warning cycles consistent with the six-hourly update frequency of available meteorological forecast inputs. This throughput advantage becomes increasingly decisive in the context of national-scale monitoring programs spanning hundreds of lakes simultaneously, where process-based model-based approaches would require computational infrastructure unavailable to most water management agencies in resource-constrained settings.

5. Conclusion

This study demonstrated that neural operator frameworks — specifically a multi-lake FNO trained on a diverse corpus of thirty-two lakes across four continents — are capable of simulating toxic cyanoHAB dynamics at spatial and temporal resolutions directly relevant to operational water quality management, with accuracy metrics substantially exceeding those of conventional deep learning baselines on lakes not encountered during training. The core contribution is the demonstration that NOs can learn generalizable solution operators for the spatiotemporal PDE system governing bloom kinetics, enabling accurate fourteen-day Chl-a field predictions on lakes with markedly different trophic states, morphologies, and climate forcing regimes from the training distribution. The mean R^2 of 0.913 at the one-day horizon and 0.791 at the fourteen-day horizon on held-out lakes, combined with a 38.4% RMSE reduction over the CNN-LSTM baseline, establishes a meaningful performance threshold for operational early warning system deployment. Physical consistency analyses confirming near-perfect Chl-a non-negativity, phosphorus mass conservation errors below 5%, and bloom propagation speeds aligned with hydrodynamic benchmarks indicate that the learned operator captures genuine physical dynamics rather than distribution-specific statistical patterns confined to the training corpus. The practical implications of these findings extend well beyond the immediate Lake Erie context. The demonstrated ability of the FNO to maintain reasonable predictive skill on Kariba Reservoir despite the near-complete absence of in-situ nutrient monitoring suggests a pathway toward scalable cyanoHAB surveillance in data-sparse regions of sub-Saharan Africa, South Asia, and Latin America, where bloom-related public health risks are acute but monitoring infrastructure remains severely limited. The sub-second inference time enables deployment within national early warning systems capable of issuing daily bloom forecasts for hundreds of lakes simultaneously at computational costs compatible with cloud microservice architectures accessible to environmental regulatory agencies at all resource levels. Management scenario analyses further demonstrated that the framework is sufficiently physically informed to generate credible responses to nutrient

loading reduction hypotheticals, supporting its potential use in cost-benefit analyses of watershed management interventions targeting cyanoHAB mitigation under projected climate change trajectories. Several limitations warrant acknowledgment as directions for future research. The current model operates at 500-meter spatial resolution, insufficient to resolve bloom dynamics in small water bodies below ten square kilometers in area — a significant gap given that small lakes and ponds constitute the vast majority of freshwater ecosystems globally by number. Extending the framework to finer resolutions using high-resolution Planet or Sentinel-2 imagery as training targets will require architectural modifications to manage substantially larger input tensor dimensions. The current training corpus is also biased toward Northern Hemisphere lakes in well-monitored regions, potentially limiting generalization to underrepresented freshwater biomes such as African rift valley lakes and Amazonian floodplain systems. Systematic expansion of the training corpus to include these biomes, combined with active learning strategies that prioritize new monitoring deployments in high-uncertainty regions, represents a natural and important next step. Future work will additionally explore the integration of explicit physical constraints derived from Redfield nutrient stoichiometry as regularization terms in the NO training objective, with the goal of further improving physical consistency in extrapolation regimes beyond the concentration ranges represented in the historical training data.

References

- Moreira, C., Vasconcelos, V., & Antunes, A. (2022). Cyanobacterial blooms: Current knowledge and new perspectives. *Earth*, 3(1), 127-135.
- Paerl, H. W., Havens, K. E., Hall, N. S., Otten, T. G., Zhu, M., Xu, H., ... & Qin, B. (2020). Mitigating a global expansion of toxic cyanobacterial blooms: confounding effects and challenges posed by climate change. *Marine and Freshwater Research*, 71(5), 579-592.
- Zhang, X., Li, P., Han, X., Yang, Y., & Cui, Y. (2024). Enhancing time series product demand forecasting with hybrid attention-based deep learning models. *IEEE Access*, 12, 190079-190091.
- Yuan, S., Chen, X., Xing, S., Li, J., Chen, H., Liu, Z., & Guo, S. (2025). Transformer-Based Scalable Multi-Agent Reinforcement Learning for Joint Resource Optimization in Cloud–Edge–End Video Streaming Systems. *IEEE Transactions on Cognitive Communications and Networking*.
- Zhang, S., Qiu, L., & Zeng, Z. (2026). Physics-Data Synergy in Structural Health Monitoring: A Multi-Scale Graph Contrastive Framework With Temperature-Adaptive Fusion. *IEEE Access*.
- Shen, Z., Zhao, W., Wang, B., Wang, Z. and Shang, W. (2026). CAGR: A Cross-Accelerator Graph Optimization Framework for Efficient Recommender System Inference. *IEEE Access*.
- Sun, T., Yang, J., Li, J., Chen, J., Liu, M., Fan, L., & Wang, X. (2024). Enhancing auto insurance risk evaluation with transformer and SHAP. *IEEE Access*, 12, 116546-116557.
- Li, J., Fan, L., Wang, X., Sun, T., & Zhou, M. (2024). Product demand prediction with spatial graph neural networks. *Applied Sciences*, 14(16), 6989.
- Xing, S., & Wang, Y. (2025). Proactive data placement in heterogeneous storage systems via predictive multi-objective reinforcement learning. *IEEE Access*.
- Liu, Y., Ren, S., Wang, X., & Zhou, M. (2024). Temporal logical attention network for log-based anomaly detection in distributed systems. *Sensors*, 24(24), 7949.

- Zhao, X., Sun, T., Ren, S., Yang, J., & Liu, Y. (2025). RAG-Based AI Agents for Enterprise Software Development: Implementation Patterns and Production Deployment. *Frontiers in Artificial Intelligence Research*, 2(3), 501-520.
- Fang, Q., & Liu, W. (2025). HARLA-ED: Resolving Information Asymmetry and Enhancing Algorithmic Symmetry in Intelligent Educational Assessment via Hybrid Reinforcement Learning. *Symmetry*, 18(1), 58.
- Mai, N. T., Cao, W., & Fang, Q. (2025). A study on how LLMs (eg GPT-4, chatbots) are being integrated to support tutoring, essay feedback and content generation. *Journal of Computing and Electronic Information Management*, 18(3), 43-52.
- Jiang, B., Wu, B., Cao, J., & Tan, Y. (2025). Interpretable Fair Value Hierarchy Classification via Hybrid Transformer-GNN Architecture. *IEEE Access*, 13, 198142-198163.
- Chen, J., Cui, Y., Zhang, X., Yang, J., & Zhou, M. (2024). Temporal convolutional network for carbon tax projection: A data-driven approach. *Applied Sciences*, 14(20), 9213.
- Chen, J., Liu, J., Liang, Y., & Zhou, M. (2026). KE-MLLM: A Knowledge-Enhanced Multi-Sensor Learning Framework for Explainable Fake Review Detection. *Applied Sciences*, 16(6), 2909.
- Liu, J., Wang, J., Chen, H., Guinness, J., Martin, R., & Kulkarni, C. S. (2019). Optimal Level Crossing Predictions for Electronic Prognostics. In *AIAA Scitech 2019 Forum* (p. 1962).
- Xing, S., Wang, Y., & Liu, W. (2025). Self-adapting CPU scheduling for mixed database workloads via hierarchical deep reinforcement learning. *Symmetry*, 17(7), 1109.
- Liu, C. L., Tseng, C. J., Huang, T. H., Yang, J. S., & Huang, K. B. (2023). A multi-task learning model for building electrical load prediction. *Energy and Buildings*, 278, 112601.
- Liu, C. L., Chang, T. Y., Yang, J. S., & Huang, K. B. (2023). A deep learning sequence model based on self-attention and convolution for wind power prediction. *Renewable Energy*, 219, 119399.
- Ding, G., Yang, S., Lin, H., Chen, Z., & Yang, J. S. (2026). LLM-Driven Adaptive Cloud Resource Scheduling: Bridging Reasoning Intelligence with Optimization Guarantees. *IEEE Open Journal of the Computer Society*.
- Wang, X., Zhang, X., Hoo, V., Shao, Z., & Zhang, X. (2024). Legalreasoner: A multi-stage framework for legal judgment prediction via large language models and knowledge integration. *Ieee Access*, 12, 166843-166854.
- Qiu, L. (2024). Deep learning approaches for building energy consumption prediction. *Frontiers in Environmental Research*, 2(3), 11-17.
- Wang, M. (2024). AI technologies in modern taxation: Applications, challenges, and strategic directions. *International Journal of Finance and Investment*, 1(1), 42-46.
- Mai, N. T., Fang, Q., & Cao, W. (2025). Measuring student trust and over-reliance on AI tutors: Implications for STEM learning outcomes. *International Journal of Social Sciences and English Literature*, 9(12), 11-17.
- Li, P., Ren, S., Zhang, Q., Wang, X., & Liu, Y. (2024). Think4SCND: Reinforcement learning with thinking model for dynamic supply chain network design. *IEEE Access*, 12, 195974-195985.
- Smith, B., Pahlevan, N., Schalles, J., Ruberg, S., Errera, R., Ma, R., ... & Kangro, K. (2021). A chlorophyll-a algorithm for Landsat-8 based on mixture density networks. *Frontiers in Remote Sensing*, 1, 623678.
- Woolway, R. I., & Merchant, C. J. (2019). Worldwide alteration of lake mixing regimes in response to climate change. *Nature Geoscience*, 12(4), 271-276.

Recknagel, F., Shan, K., Adrian, R., & Köhler, J. (2025). Early warning of harmful cyanobacteria blooms based on high frequency in situ monitoring and intelligible machine learning modelling: The case study of Lake Müggelsee (Germany). *Water Research*, 124514.