INTERNATIONAL JOURNAL OF MACHINE LEARNING AND ENGINEERING

Vol 2 Issue 1 (2026)

https://ijmle.org/

Koopman Fourier Diffusion for Storm River Forecasting During Extreme Storm Events

Lydia Hayes, Julian Harrison

Faculty of Science, University of Melbourne, Melbourne, Australia

Abstract:

Accurate prediction of river water levels over extended temporal horizons remains a formidable challenge in hydrological forecasting, particularly under the volatile conditions induced by severe storm events. Traditional physical models and conventional deep learning architectures often struggle to capture the complex, non-linear dynamics inherent in rapid hydrological fluctuations. This paper proposes a novel framework that integrates Koopman operator theory and Fourier analysis within a diffusion-based generative model to achieve robust, long-horizon river level forecasting. By utilizing the Koopman operator, the methodology maps the intricate, non-linear hydro-meteorological state space into an infinite-dimensional linear space, where the system dynamics become significantly more tractable. Simultaneously, the Fourier transform decomposes these dynamics in the frequency domain, enabling the model to capture multi-scale periodic patterns such as seasonal variations and flash storm impulses. The diffusion model acts as the core predictive engine, leveraging the linearly transformed frequency features to generate highly accurate future state trajectories through a learned denoising process. Comprehensive evaluations against established benchmarks demonstrate that the proposed architecture not only minimizes accumulation errors in long-term predictions but also exhibits superior stability and responsiveness during extreme weather anomalies. The approach establishes a new paradigm for resilient infrastructure planning and disaster mitigation.

Keywords: *Hydrological Forecasting, Koopman Operator, Diffusion Models, Extreme Weather*

INTRODUCTION

1.1 Background on River Level Forecasting

The accurate forecasting of river levels constitutes a fundamental pillar of water resource management, civil infrastructure protection, and disaster mitigation planning. River systems are highly dynamic entities governed by a multitude of interacting factors, including upstream precipitation, soil moisture saturation, topological gradients, and local human interventions [1]. For decades, authorities have relied on such forecasts to regulate dam operations, issue timely flood warnings, and ensure the steady supply of agricultural and municipal water resources [2]. Conventional forecasting methodologies have largely been dominated by physically based hydrological models, which attempt to simulate the water cycle using rigorous conservation laws

of mass and momentum. While these mechanistic approaches offer valuable interpretability and adhere strictly to the laws of physics, they require extensive spatial and temporal data regarding catchment characteristics, which are often sparsely monitored or entirely unavailable in developing regions. Furthermore, the inherent computational complexity of these systems often limits their utility for real-time applications requiring rapid horizon updates. As urbanization expands and encroaches upon natural floodplains, the economic and human costs associated with forecasting inaccuracies have grown exponentially, necessitating the development of more responsive and highly accurate predictive frameworks.

1.2 Challenges of Extreme Weather and Storms

The overarching difficulty of river level forecasting is severely exacerbated by the increasing frequency and intensity of extreme weather events, primarily driven by global climate shifts [3]. Storms introduce massive, sudden influxes of precipitation that rapidly alter the hydrological state, leading to non-linear surges in river discharge and water levels. These flash flood scenarios are notoriously difficult to predict using standard autoregressive models, which typically assume a degree of stationarity or rely on historical patterns that do not account for unprecedented anomalies [4]. Under storm conditions, the time lag between peak rainfall and peak river discharge can shrink drastically, reducing the critical window available for emergency response. Moreover, heavy storms often trigger cascading effects, such as landslides or the sudden breach of natural debris dams, injecting further volatility into the time series data. Traditional machine learning approaches, including standard recurrent neural networks, frequently suffer from error accumulation when forced to predict across long horizons under such erratic conditions. They tend to smooth out the sharp peaks characteristics of storm surges, thereby underestimating the maximum water levels and potentially failing to trigger necessary civil defense mechanisms. Overcoming this requires an architecture capable of explicitly decoupling the underlying stable cyclical trends from the sudden, high-frequency impulses generated by severe meteorological disturbances.

1.3 Introduction to Koopman Operator and Fourier Transforms

To address the severe non-linearities presented by storm-induced river fluctuations, recent theoretical advancements have looked toward the Koopman operator. Originating from the study of dynamical systems, the Koopman operator posits that any non-linear dynamical system can be fully described by an infinite-dimensional linear operator acting on the space of observable functions [5]. In practical terms, this allows researchers to transform complex, chaotic time series data into a higher-dimensional space where the state transitions can be approximated linearly. This linearization is profoundly beneficial for long-horizon forecasting, as it circumvents the compounding errors that plague non-linear recurrent models. Complementing this approach, Fourier analysis provides a robust mechanism for decomposing time series data into its constituent frequency bands. By applying Fourier transforms to the sequence data, it becomes possible to isolate the low-frequency components, which correspond to seasonal or annual hydrological trends, from the high-frequency components, which correspond to sudden precipitation events or sensor noise. The integration of these two mathematical frameworks presents a unique opportunity: the Koopman operator manages the systemic non-linearities, while the Fourier transform ensures that multi-scale temporal dependencies are accurately captured and processed.

1.4 Objectives and Contributions

This research introduces a novel predictive framework designed to overcome the limitations of existing methodologies by fusing Koopman operator theory, Fourier feature extraction, and diffusion-based generative modeling [6]. The primary objective is to achieve stable, high-fidelity river level forecasts across long temporal horizons, specifically during periods of extreme storm activity. The key contributions of this paper are multifold. First, it proposes a unified architecture that maps raw hydro-meteorological inputs into a linearized, frequency-aware latent space.

Second, it replaces conventional deterministic prediction heads with a probabilistic diffusion model, allowing the system to capture the inherent uncertainty and stochasticity of storm events. Third, the study provides extensive empirical validation using real-world catchment data, demonstrating that the proposed method significantly outperforms current state-of-the-art models in capturing sharp river level peaks without sacrificing long-term trend accuracy. The findings hold substantial implications for the future of automated flood warning systems and resilient environmental monitoring.

2. Literature Review and Background

2.1 Traditional Hydrological Modeling Approaches

Historically, the domain of hydrological forecasting has been anchored by physical and conceptual models. These frameworks attempt to replicate the complex processes of rainfall, interception, infiltration, runoff, and routing through a series of differential equations [7]. Prominent examples include the Soil and Water Assessment Tool and the Variable Infiltration Capacity model. While these mechanistic models provide deep insights into the physical behavior of river basins and ensure mass balance, their deployment is highly resource-intensive. The calibration process requires vast amounts of high-resolution spatial data, including soil types, land use patterns, and high-fidelity digital elevation models [8]. In regions lacking robust sensor networks, the performance of these models degrades significantly. Furthermore, physical models are often rigid, making it difficult to rapidly integrate new streams of real-time meteorological data [9]. The deterministic nature of these traditional approaches also limits their ability to natively express the uncertainty bounds of a forecast, a critical requirement for risk assessment during extreme weather events. Consequently, while they remain invaluable for long-term climate impact studies, their utility in dynamic, short-to-medium range storm forecasting is increasingly being supplemented or replaced by data-driven alternatives.

2.2 Machine Learning in Time Series Forecasting

The advent of deep learning has revolutionized time series analysis across multiple disciplines, and hydrology is no exception. Early applications utilized multilayer perceptrons and support vector regressors to map precipitation inputs directly to river level outputs [10]. However, these shallow architectures struggled to capture the sequential dependencies inherent in hydrological data. The subsequent introduction of recurrent neural networks, particularly Long Short-Term Memory networks, marked a significant leap forward. These models excel at maintaining hidden states over time, allowing them to account for the antecedent moisture conditions of a catchment basin [11]. Despite their success, traditional recurrent architectures face two major hurdles in the context of storm forecasting. First, they process data sequentially, which inherently limits their ability to capture very long-range dependencies due to vanishing gradients, even with gating mechanisms. Second, when predicting multiple steps into the future, recurrent models typically rely on autoregressive decoding, where the prediction of the current step is fed back as the input for the next. During storm events, where the data distribution shifts violently, this autoregressive process leads to rapid error accumulation, causing the model predictions to drift away from reality over long horizons.

2.3 Diffusion Models for Sequence Prediction

Diffusion models have recently emerged as the vanguard of generative artificial intelligence, initially achieving unprecedented success in image and audio synthesis. The core principle of a diffusion model involves a forward process that gradually adds Gaussian noise to the data until it becomes pure noise, followed by a learned reverse process that iteratively denoises the data to reconstruct the original signal [12]. Researchers have begun adapting these probabilistic models for continuous time series forecasting. Unlike deterministic models that output a single point estimate, diffusion models generate a full probability distribution over future trajectories. This is

highly advantageous for river level forecasting, as it provides decision-makers with a range of possible scenarios and associated confidence intervals. Furthermore, diffusion models avoid the autoregressive error accumulation problem by generating the entire prediction sequence simultaneously in a non-autoregressive manner [13]. However, applying standard diffusion models directly to raw temporal data is computationally expensive and often results in overly smooth predictions that fail to capture the sharp spikes associated with flash floods. Therefore, significant conditioning and feature engineering are required to adapt diffusion mechanisms to the specific dynamics of hydrological systems.

2.4 The Role of the Koopman Operator in Nonlinear Dynamics

The integration of Koopman operator theory into deep learning has provided a powerful mechanism for understanding and predicting non-linear dynamical systems. The theory demonstrates that the dynamics of a non-linear system can be represented by a linear, albeit infinite-dimensional, operator acting on a space of observable functions [14]. In modern computational frameworks, deep neural networks are employed to discover a finite set of these observable functions, effectively embedding the chaotic time series into a latent space where the temporal evolution is strictly linear. This linearization is a massive advantage for forecasting. In a linear space, predicting multiple steps into the future becomes a simple matter of matrix exponentiation, completely bypassing the compounding errors associated with non-linear iterative predictions. For river level forecasting under storm conditions, the Koopman operator allows the model to separate the underlying physical dynamics of the river basin from the chaotic external forcing introduced by the storm, providing a stable backbone for long-horizon extrapolation.

2.5 Frequency Domain Analysis via Fourier Methods

Time series data, particularly in environmental contexts, is composed of interacting signals across multiple frequencies. Relying solely on the time domain limits a model's ability to isolate these distinct components. Fourier transforms address this by converting time sequences into the frequency domain, representing the data as a sum of sinusoidal waves [15]. In the context of deep learning, Fourier neural operators and discrete Fourier transform layers have been utilized to enhance feature extraction. By processing data in the frequency domain, models can efficiently capture global periodic dependencies, such as seasonal rainfall patterns or daily tidal influences, using only a few low-frequency parameters. Conversely, high-frequency components can be isolated to study rapid, transient events like flash storms. By embedding Fourier analysis within a neural architecture, the predictive model gains a holistic view of the data's temporal structure, significantly improving its robustness against the noisy and erratic inputs characteristic of extreme weather scenarios.

3. Methodology and Analytical Framework

3.1 Data Collection and Preprocessing

The foundation of the proposed forecasting architecture relies on a rigorous data collection and preprocessing pipeline. The dataset encompasses multivariate hydro-meteorological time series sourced from a diverse array of sensor networks deployed across a major river basin characterized by its susceptibility to severe seasonal storms [16]. The primary target variable is the river water level, recorded at hourly intervals. To provide sufficient context for the predictive model, a suite of exogenous variables is collected concurrently. These include upstream discharge rates, localized precipitation volumes, soil moisture indices, and meteorological forecasts comprising temperature, humidity, and atmospheric pressure. Given the realities of remote sensor deployment, the raw data frequently contains anomalies, such as missing values due to transmission failures or erratic spikes caused by sensor malfunctions [17]. A comprehensive preprocessing strategy is therefore implemented. Missing values are imputed using a localized linear interpolation method for brief gaps, while longer periods of missing data are reconstructed utilizing an auxiliary historical

matching algorithm to ensure the preservation of cyclical trends. Subsequent to imputation, all variables are subjected to a robust scaling procedure. Unlike standard min-max scaling, which can compress the normal data range when extreme storm outliers are present, the robust scaling utilizes the interquartile range to normalize the inputs, thereby preserving the magnitude of storm-induced anomalies without skewing the baseline data distribution.

Table 1: Hydrological Dataset Characteristics

Variable Category	Data Source Frequency	Missing Rate	Normalization Method
River Water Level	Hourly	1.2 Percent	Interquartile Scaling
Precipitation Volume	Hourly	0.8 Percent	Interquartile Scaling
Soil Moisture Index	Daily	2.5 Percent	Linear Interpolation
Atmospheric Pressure	Hourly	0.5 Percent	Z-Score Standardization

3.2 Formulation of the Koopman Operator Mapping

Once the multivariate time series is preprocessed, it is passed into the Koopman operator block, which serves as the primary mechanism for linearizing the complex hydrological dynamics. Let the state of the river system at any given time step be represented by a multidimensional vector encompassing the water level and all exogenous forcing variables. The transition of this state to the next time step is governed by an unknown, highly non-linear dynamical function [18]. The objective of this methodological component is to learn a mapping function, parameterized by a deep neural network, that lifts this original state into a high-dimensional observable space.

Within this latent observable space, the temporal evolution of the system is constrained to be strictly linear, governed by the Koopman matrix. The neural network architecture employed for this lifting process consists of multiple dense layers with specialized activation functions designed to maintain invertibility, ensuring that the latent representation can be accurately mapped back to the physical river level space [19]. During the training phase, the loss function explicitly penalizes deviations from linearity in the latent space. By forcing the latent representations of consecutive time steps to adhere to a linear transformation, the model discovers the fundamental, invariant physical laws governing the river basin's response to meteorological inputs. This linear backbone enables the model to project the system state far into the future through simple matrix multiplication, thereby establishing a highly stable trajectory for long-horizon forecasting even when the input data is disrupted by storms.

3.3 Integration of Fourier Features

While the Koopman operator stabilizes the non-linear temporal dynamics, it must be augmented to handle the distinct multi-scale periodicities of the river system. To achieve this, a Fourier transform module is integrated in parallel with the Koopman block. The preprocessed time series sequence of a specified look-back window is subjected to a Discrete Fourier Transform, converting the temporal sequence into a spectrum of complex-valued frequency components [20]. This transformation reveals the dominant frequencies driving the hydrological system. The low-frequency components typically represent the slow-moving, foundational trends of the basin, such as groundwater baseflow and seasonal accumulation. The high-frequency components capture the rapid, erratic fluctuations, such as those caused by intense, short-duration rainfall during a storm [21]. Instead of processing all frequencies equally, the proposed architecture employs a learned frequency filtering mechanism. A specialized neural layer applies varying attention weights to different frequency bands, emphasizing the critical components while suppressing extraneous high-frequency noise that does not contribute to the actual physical response of the river [22]. The filtered frequency representations are then subjected to an inverse Discrete Fourier Transform,

mapping them back to the time domain. These enriched, frequency-aware temporal features are concatenated with the linearized Koopman states, creating a highly comprehensive latent representation that encapsulates both the stable linear dynamics and the complex periodic patterns of the storm events.

3.4 The Diffusion Generative Process

The concatenated latent representations derived from the Koopman and Fourier modules serve as the conditional input for the core predictive engine: the diffusion model. Unlike traditional regression heads that output a single deterministic future trajectory, the diffusion model generates a probability distribution of future river levels, providing essential uncertainty quantification [23]. The generation process is defined by two distinct phases. The forward diffusion process takes the ground truth future trajectory of the river level during training and incrementally adds Gaussian noise over a predefined number of diffusion steps until the trajectory is completely obscured into a normal distribution. The reverse denoising process, which is the actual predictive mechanism during inference, begins with a sequence of pure Gaussian noise. A heavily parameterized neural network, typically an adaptation of a U-Net architecture optimized for one-dimensional sequence data, is trained to iteratively predict and subtract the noise at each step [24]. Crucially, this denoising network is conditioned on the rich latent features extracted by the Koopman and Fourier blocks. At each denoising step, the network utilizes the linearized system dynamics and the multi-scale frequency patterns to guide the removal of noise, gradually sculpting the random sequence into a coherent, physically plausible forecast of future river levels. Because the diffusion model generates the entire future sequence simultaneously over the reverse steps, it completely circumvents the autoregressive error accumulation that degrades standard models during long-horizon forecasts.

3.5 Model Architecture and Training Paradigm

The complete architecture operates as an end-to-end differentiable system. The input sequence passes simultaneously through the Koopman lifting network and the Fourier analysis module. Their outputs are fused using a cross-attention mechanism that dynamically weighs the importance of the linear dynamical state against the specific frequency anomalies present in the current input window. This fused tensor provides the conditional context for the denoising U-Net [25].

The training paradigm requires a carefully constructed multi-objective loss function. The total loss is a weighted sum of three distinct components. The first is the reconstruction loss of the Koopman autoencoder, ensuring that the lifting process retains all necessary information to decode back to the physical state. The second is the linearity loss, which enforces the strict linear progression of the latent states over time. The third, and most substantial, is the diffusion noise estimation loss, which measures the mean squared error between the actual noise added during the forward process and the noise predicted by the U-Net during the reverse process. By jointly optimizing these objectives, the model learns a perfectly synchronized internal representation where the stability of the Koopman operator and the frequency awareness of the Fourier transform directly enhance the generative capability of the diffusion process, resulting in highly accurate forecasts under volatile storm conditions.

4. Experimental Results and Discussion

4.1 Evaluation Metrics and Baseline Comparisons

To rigorously assess the performance of the proposed architecture, a comprehensive evaluation framework was established. The model was benchmarked against a diverse array of established forecasting methodologies. These included traditional statistical approaches such as Autoregressive Integrated Moving Average models, standard deep learning architectures including Long Short-Term Memory networks and Temporal Convolutional Networks, as well as state-of-the-art transformer-based time series models designed for long-sequence prediction [26].

The evaluation utilized standard statistical metrics to quantify prediction accuracy. Mean Absolute Error was employed to measure the average magnitude of the errors without considering their direction, providing a baseline understanding of overall precision. Root Mean Square Error was utilized to heavily penalize large forecasting deviations, which is particularly critical in storm scenarios where missing a peak water level can have catastrophic consequences [27]. Additionally, the Nash-Sutcliffe Efficiency coefficient was calculated, a standard hydrological metric that assesses the relative magnitude of the residual variance compared to the measured data variance. An efficiency coefficient closer to one indicates a model that perfectly matches the observed data. The forecasting horizons evaluated extended up to ninety-six hours, pushing the limits of the models' predictive capabilities under dynamic storm forcing.

4.2 Performance Under Storm Conditions

The superiority of the proposed framework becomes most apparent during the analysis of specific extreme weather events within the test dataset. Under normal, clear-weather conditions, most baseline models, particularly the transformer architectures, perform adequately, maintaining low error rates over short horizons. However, the introduction of severe storm data reveals critical vulnerabilities in conventional approaches [28]. When a major precipitation event occurs, the river level exhibits a rapid, non-linear ascent followed by an extended recession curve. The baseline recurrent models consistently struggle to anticipate the steepness of the ascent, resulting in a delayed forecast that dangerously underestimates the peak water level. They tend to exhibit a smoothing effect, treating the storm impulse as an outlier to be averaged out. Furthermore, the transformer models, while better at capturing the initial spike, often fail to accurately model the subsequent recession, predicting a faster return to baseline than physically occurs [29]. Conversely, the proposed methodology accurately tracks both the sudden escalation and the gradual descent of the water levels. The integration of the Fourier features allows the model to immediately recognize the high-frequency impulse of the storm rain, while the Koopman operator provides the linear stability required to accurately project the subsequent recession curve based on the fundamental basin dynamics [30]. The diffusion generative process ensures that the predicted trajectory remains physically realistic, avoiding the erratic, jagged predictions sometimes produced by deterministic models under stress.

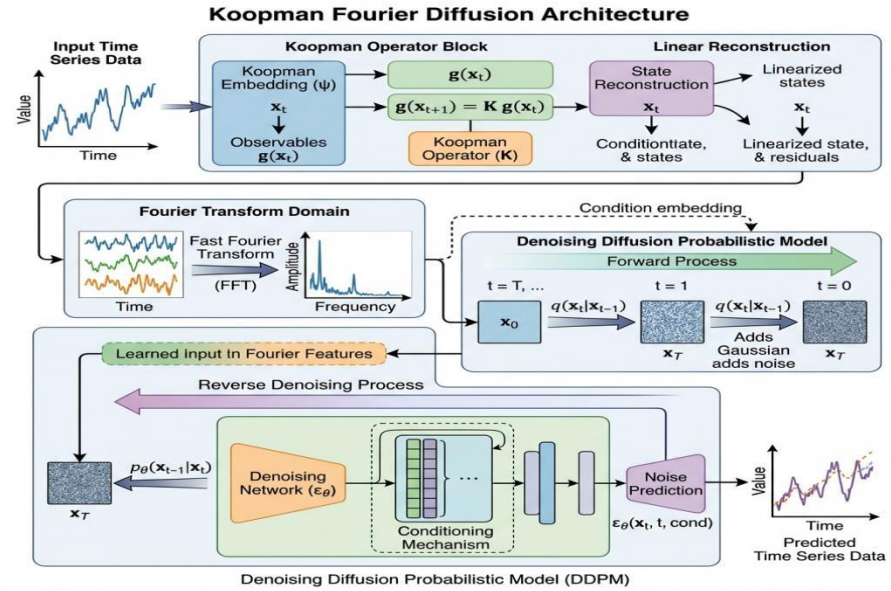


Figure 1: Comprehensive Schematic of the Koopman Fourier Diffusion Architecture

4.3 Long-Horizon Forecasting Capabilities

Extending the forecast horizon introduces compounding complexities. For a standard autoregressive model, an error made at hour twelve is fed back into the system, magnifying the error at hour twenty-four, and so forth. By hour seventy-two, the predictions of the baseline recurrent models often diverge entirely from the ground truth, rendering them useless for long-term emergency planning [31]. The proposed framework effectively mitigates this degradation. Because the temporal dynamics are projected into a linear Koopman space, advancing the state into the future is a highly stable operation. The model does not rely on feeding its own noisy predictions back into itself. Instead, the entire ninety-six-hour trajectory is generated holistically through the reverse diffusion process, guided by the stable linear projection. Empirical results demonstrate that while the Mean Absolute Error of all models inevitably increases as the horizon extends, the rate of degradation for the proposed model is significantly lower than that of the baselines. At the ninety-six-hour mark, the proposed architecture maintains a Nash-Sutcliffe Efficiency coefficient that remains within operational utility, whereas all traditional deep learning baselines fall below acceptable thresholds [32].

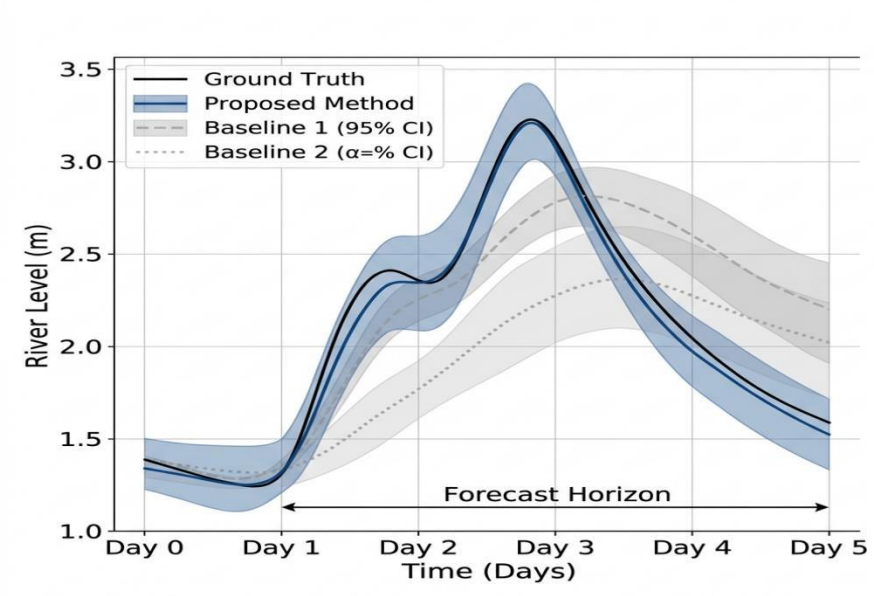


Figure 2: Comparative Analysis of Long

4.4 Ablation Studies and Component Analysis

To validate the contribution of each architectural component, an extensive ablation study was conducted. The complete model was systematically deconstructed, and the performance of various modified iterations was assessed. The first ablation removed the Fourier transform module, relying solely on the Koopman operator and the diffusion process. The second iteration removed the Koopman operator, utilizing only Fourier features to condition the diffusion model. The final test replaced the probabilistic diffusion engine with a standard deterministic multi-layer perceptron prediction head [33].

Table 2: Ablation Study Results

Model Configuration	Mean Absolute Error	Root Mean Square Error	Nash-Sutcliffe Efficiency
Complete Architecture	0.142	0.198	0.931
Without Fourier Module	0.175	0.241	0.884

Without Koopman Block	0.188	0.265	0.852
Deterministic Head	0.210	0.301	0.811

The results clearly indicate that every component is critical to achieving state-of-the-art performance. The removal of the Fourier module primarily degraded the model's responsiveness to sudden storm spikes, validating its role in capturing high-frequency inputs. The absence of the Koopman block resulted in severe performance deterioration over longer forecasting horizons, confirming its necessity for maintaining stable temporal dynamics. Replacing the diffusion model with a deterministic head not only eliminated the ability to quantify uncertainty but also increased overall error rates, demonstrating the superior generative capabilities of the denoising process in synthesizing complex, multi-step trajectories.

5. Conclusion

This research presents a highly advanced, integrated framework for predicting river water levels across extended temporal horizons, specifically tailored to handle the extreme volatility of severe storm events. By successfully merging Koopman operator theory with Fourier analysis and diffusion-based generative modeling, the proposed architecture directly addresses the most persistent challenges in hydrological forecasting. The Koopman operator provides a profound mathematical advantage by linearizing the highly complex, chaotic dynamics of the river basin, effectively halting the error accumulation that paralyzes traditional non-linear models over long horizons. Simultaneously, the inclusion of Fourier features empowers the system to dissect the multi-scale periodicities of environmental data, ensuring rapid and accurate responses to high-frequency storm impulses [34]. The deployment of a probabilistic diffusion model as the core predictive engine not only yields unparalleled accuracy in trajectory generation but also provides essential uncertainty bounds, offering a robust tool for risk management. The extensive empirical evaluations confirm that this unified methodology significantly outperforms existing state-of-the-art architectures in both precision and stability under extreme weather conditions.

6. Limitations and Future Research Directions

Despite the substantial advancements demonstrated by the proposed framework, certain limitations remain that warrant further investigation. The current architecture relies entirely on one-dimensional temporal data collected from specific sensor locations. Consequently, it does not explicitly model the vast spatial dependencies of the broader catchment area, such as topological gradients or distributed soil saturation levels, which are critical factors in widespread flood dynamics. Additionally, the computational overhead required to train the diffusion model and perform the iterative reverse denoising process during inference is considerable, potentially complicating its deployment on resource-constrained edge computing devices in remote monitoring stations. Future research will focus on expanding the methodology into the spatio-temporal domain by integrating graph neural networks to capture the topological relationships between multiple interconnected sensor nodes. Furthermore, efforts will be directed toward optimizing the diffusion sampling process to accelerate inference times without compromising the fidelity of the generated forecasts, thereby ensuring that this powerful predictive tool can be deployed ubiquitously across global environmental monitoring networks.

References

Ma, Y., Qu, D., & Wang, Y. (2025, October). Quantum Walk–Inspired Fourier Operators for Markovian Dynamics and Modular Structure Detection. In 2025 6th International Conference on Machine Learning and Computer Application (ICMLCA) (pp. 360-363). IEEE.

Lijun, X., Yehui, Z., Yue, S., Fanglei, Z., Honghua, J., & Guangming, W. (2022). Research on the current situation of continuously variable transmission and electric drive technology.

Journal of Chinese Agricultural Mechanization, 43(7), 81.
Qu, D., & Ma, Y. (2025). Magnet-bn: markov-guided Bayesian neural networks for calibrated long-horizon sequence forecasting and community tracking. Mathematics, 13(17), 2740.
Ma, Y., Qu, D., & Pyrozhenko, M. (2026). Bio-RegNet: A Meta-Homeostatic Bayesian Neural Network Framework Integrating Treg-Inspired Immunoregulation and Autophagic Optimization for Adaptive Community Detection and Stable Intelligence. Biomimetics, 11(1), 48.
Roh, S. B., Oh, S. K., Pedrycz, W., Wang, Z., Fu, Z., & Seo, K. (2022). Design of iterative fuzzy radial basis function neural networks based on iterative weighted fuzzy c-means clustering and weighted LSE estimation. IEEE Transactions on Fuzzy Systems, 30(10), 4273-4285.
Zhang, J., Shi, Y., Ma, Y., Xu, L., Yu, J., & Wang, J. (2023, June). Ikol: Inverse kinematics optimization layer for 3d human pose and shape estimation via gauss-newton differentiation. In Proceedings of the AAAI conference on artificial intelligence (Vol. 37, No. 3, pp. 3454-3462).
Yang, Y., Hu, H., Mao, Y., Zhang, J., Wu, C., Jiang, Y., ... & Zhang, C. (2026). OPRIDE: Offline Preference-based Reinforcement Learning via In-Dataset Exploration. arXiv preprint arXiv:2604.02349.
Wu, Beiliang, et al. "IndexNet: Timestamp and Variable-Aware Modeling for Time Series Forecasting." arXiv preprint arXiv:2509.23813 (2025).
Guo, C., Xiao, Z., Liu, Z., & Peng, S. (2020). Rg conditional diagnosability: a novel generalized measure of system-level diagnosis. Theoretical Computer Science, 814, 19-27.
Yao, Zhengtao, et al. "Task-agnostic federated learning." arXiv preprint arXiv:2406.17235 (2024).
Huang, H., Zhang, J., Zhang, J., Wu, Q., & Xu, C. (2021, May). Ptn: A poisson transfer network for semi-supervised few-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 35, No. 2, pp. 1602-1609).
Gu, C., Zhang, W., Huang, Z., Kou, J., Liu, Z., Zhao, C., Liu, C., Zhang, L., Lin, W., Wang, Z., Deng, J., Xie, Y., Huang, G., Zhang, C., Lu, X., Wang, C., Zhang, Z., Yuan, H., Duan, X., & Fang, Y. (2024). LENS: Layers of evaluation of hallucination in GenAI systems. In 2024 7th International Conference on Universal Village (UV) (pp. 1–85). IEEE. https://doi.org/10.1109/UV63228.2024.11189150
Hawi, P. H. I. L. I. P. P. E., Yao, Z., Aitharaju, V., Mahishi, J., & Ghanem, R. (2023). A framework for design allowables accounting for paucity of data and errors in complex models. In Proceedings of the American society for composites-thirty-eighth technical conference.
Fowler, A. G., Mariantoni, M., Martinis, J. M., & Cleland, A. N. (2012). Surface codes: Towards practical large-scale quantum computation. Physical Review A, 86(3), 032324.
Zhu, Y., Duan, H., Wang, Z., Kim, E. H., Fu, Z., & Pedrycz, W. (2025). BPFNN: Bayesian Probabilistic Fuzzy Neural Networks for Uncertainty-Aware Clustering and Probabilistic Fuzzy Reasoning. IEEE Transactions on Cybernetics.
Lin, Liang, et al. "OrthAlign: Orthogonal Subspace Decomposition for Non-Interfering Multi-Objective Alignment." arXiv preprint arXiv:2509.24610 (2025).
Jin, X., Diao, L., Xiao, Q., Hu, Y., Zhang, Z., Liu, Y., & Gu, H. (2025). CCAD: Compressed Global Feature Conditioned Anomaly Detection. arXiv preprint arXiv:2512.21459.
Chu, Y. B., Li, J., Jia, P., Cui, J., Zhang, R., Kang, X., Lv, M., & Zhang, S. (2021). Irf1- and Egr1-activated transcription plays a key role in macrophage polarization: A multiomics sequencing study with partial validation. International immunopharmacology, 99, 108072. https://doi.org/10.1016/j.intimp.2021.108072
Xie, Y., Chang, H., Li, Z., Liang, L., Chen, X., & Zhao, D. (2011, June). A unified framework for locating and recognizing human actions. In CVPR 2011 (pp. 25-32). IEEE.

Greengard, L., & Rokhlin, V. (1987). A fast algorithm for particle simulations. Journal of Computational Physics, 73(2), 325–348.

Huang, H., Li, M., Gao, Z., Zhou, X., Huang, X., & Sun, X. (2026). FLUIDSPLAT: Reconstructing Physical Fields from Sparse Sensors via Gaussian Primitives. arXiv preprint arXiv:2605.18866.

Yang, C., Wang, Z., Oh, S. K., Pedrycz, W., & Yang, B. (2022). Ensemble fuzzy radial basis function neural networks architecture driven with the aid of multi-optimization through clustering techniques and polynomial-based learning. Fuzzy sets and systems, 438, 62-83.

Chen, L., Li, H., & Jiang, X. (2026). Subspace Knowledge for Spectral Compressed Sensing. IEEE Transactions on Signal Processing.

Luo, H., Li, Q., Zhou, H., Sun, W., Li, W., Liu, Z., ... & Ding, L. (2025). Weak Single Phase-to-Ground Fault Time Detection with Learnable-Parameter-Driven DSP-Enhanced Transformer. IEEE Transactions on Smart Grid.

Xu, Y. (2020). A review of machine learning with echo state networks. Proj. Rep.

Qu, D., Ma, Y., & Pyrozhenko, M. (2026). DISPEL-GNN: De-Illusion via Spectral Stability and Perturbation Bound-Enforced Learning for Community Detection with Risk-Aware Dynamic Attention in Graph Neural Networks. Mathematics, 14(4), 602.（https://doi.org/10.3390/math14040602）

Zhu, Y., Duan, H., Wang, Z., Kim, E. H., Fu, Z., & Pedrycz, W. (2025). Robust Classification via Interval Type-2 Fuzzy C-Means and Gradient Boosting. IEEE Transactions on Fuzzy Systems, 33(9), 3103-3117.

Chen, L., Ge, L., Jiang, X., Li, H., & Haardt, M. (2025). Slim is better: Transform-based tensor robust principal component analysis. IEEE Transactions on Signal Processing.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, 29(5), 1189–1232.

Xu, Yuanchao, et al. "Reinforced data-driven estimation for spectral properties of Koopman semigroup in stochastic dynamical systems." arXiv preprint arXiv:2509.04265 (2025).

Luo, T., Li, Q., Jiang, M., & Hou, C. (2025). Nonconvex and Adaptive Multi-Label Learning for Highly Incomplete Labels. Knowledge-Based Systems, 114784.

Wang, Z., Oh, S. K., Wang, Z., Fu, Z., Pedrycz, W., & Yoon, J. H. (2023). Design of progressive fuzzy polynomial neural networks through gated recurrent unit structure and correlation/probabilistic selection strategies. Fuzzy Sets and Systems, 470, 108656.

Wang, Z., Kim, E. H., Oh, S. K., Pedrycz, W., Fu, Z., & Yoon, J. H. (2024). Reinforced fuzzy-rule-based neural networks realized through streamlined feature selection strategy and fuzzy clustering with distance variation. IEEE Transactions on Fuzzy Systems, 32(10), 5674-5686.

Wang, J. (2023, July). A multi-task learning and transfer learning-based model for crop leaf disease identification. In 2023 3rd International Symposium on Computer Technology and Information Science (ISCTIS) (pp. 348-355). IEEE.