

Accelerating HEC-RAS: A Recurrent Neural Operator for Rapid River Forecasting

Edward Holmberg
GS Center for Env. Informatics
University of New Orleans,
Louisiana, USA
 eholmber@uno.edu

Pujan Pokhrel
GS Center for Env. Informatics
University of New Orleans,
Louisiana, USA
 ppokhrel@uno.edu

Elias Ioup
Center for Geospatial Sciences
Naval Research Laboratory,
Mississippi, USA
 elias.z.ioup.civ@us.navy.mil

Ken Pathak
US Army Corps of Engineers
Vicksburg District,
Mississippi, USA
 ken.pathak@usace.army.mil

Steven Sloan
US Army Corps of Engineers
Vicksburg District,
Mississippi, USA
 steven.sloan@usace.army.mil

Kendall Niles
US Army Corps of Engineers
Vicksburg District,
Mississippi, USA
 kendall.niles@usace.army.mil

Jay Ratcliff
GS Center for Env. Informatics
University of New Orleans,
Louisiana, USA
 jratclif@uno.edu

Maik Flanagan
US Army Corps of Engineers
New Orleans District,
Louisiana, USA
 maik.c.flanagan@usace.army.mil

Christian Guetl
CoDiS-Lab ISDS
Graz University of Technology,
Graz, Austria
 c.guetl@tugraz.at

Julian Simeonov
Ocean Sciences Division
Naval Research Laboratory,
Mississippi, USA
 julian.a.simeonov.civ@us.navy.mil

Mahdi Abdelguerfi
GS Center for Env. Informatics
University of New Orleans,
Louisiana, USA
 gulfsceidirector@uno.edu

Abstract—Physics-based solvers such as the Hydrologic Engineering Center’s River Analysis System (HEC-RAS) provide high-fidelity river forecasts but are too slow for on-the-fly decision-making during floods. We present a machine learning (ML) surrogate that treats HEC-RAS as a data generator and couples a Gated Recurrent Unit (GRU) for short-term memory with a geometry-aware Fourier Neural Operator (Geo-FNO) for long-range spatial coupling. Trained on 71 reaches of the Mississippi River Basin and evaluated on a year-long hold-out, the surrogate achieves a median absolute stage error of 0.28 ft. For a full 71-reach ensemble forecast, it reduces wall-clock time from 139 to 40 minutes (3.5×). By reading native HEC-RAS files and operating on a compact eight-channel feature interface, the model delivers operational speed while preserving fidelity, enabling rapid “what-if” ensemble guidance.

Index Terms—Fourier Neural Operator; Surrogate Modeling; HEC-RAS; Gated Recurrent Units.

I. INTRODUCTION

During a flood, the U.S. Army Corps of Engineers (USACE) must make critical, time-sensitive decisions—from issuing evacuation orders to scheduling gate operations—within minutes. This operational tempo is fundamentally at odds with the hours-long wall-clock times required by physics-based solvers such as the Hydrologic Engineering Center’s River Analysis System (HEC-RAS) to simulate unsteady flow [7], [8]. While pre-

computed scenario libraries or reduced-order models offer one workaround, they are often too coarse to capture the specific hydrograph that unfolds in real time [9], [11]. The central challenge, therefore, is to deliver the fidelity of an HEC-RAS simulation at a speed that enables rapid, on-the-fly ensemble forecasting.

We address this challenge by reframing the HEC-RAS workflow itself. Instead of relying on its iterative solver, we treat its native project files as a direct source of training data for a deep-learning surrogate. We propose an autoregressive model that couples a Gated Recurrent Unit (GRU) for short-term temporal memory with a geometry-aware Fourier Neural Operator (Geo-FNO) for long-range spatial dependencies. This hybrid architecture learns the coupled spatio-temporal dynamics of river flow by ingesting a minimal eight-channel vector representing dynamic state, static geometry, and boundary forcings, and then advances the system state hour by hour.

Our primary contribution is a true plug-in surrogate that requires no re-meshing or data conversion, reading native HEC-RAS files directly. This is enabled by a minimalist, reusable interface—a compact feature set sufficient for stable, multi-day forecasts. Evaluated across 71 reaches of the Mississippi River Basin, the model achieves a 3.45× end-to-end speedup while maintaining a median absolute stage error of 0.28 ft on

a year-long, unseen hold-out, elevating autoregressive neural operators from academic prototypes to operationally promising engines for rapid ensemble flood guidance.

The rest of the paper is structured as follows. Section II details the HEC-RAS data pipeline. Section III reviews related work. Section IV presents the model architecture. Section V outlines the experimental setup, followed by results in Section VI. Section VII discusses findings and limitations, and Section VIII concludes.

II. BACKGROUND: HEC-RAS AS A DATA PREPROCESSOR

A. HEC-RAS: The Industry-Standard Solver

HEC-RAS, the U.S. Army Corps of Engineers' River Analysis System, is widely regarded as the industry-standard platform for river hydraulics [8]. Under the hood, it solves the one-dimensional Saint-Venant equations[7] using an implicit Newton–Raphson finite-difference scheme, with several inner iterations per global time step to balance continuity and momentum [8]. This strategy delivers high numerical accuracy but at a steep computational cost: full-reach unsteady-flow simulations typically require hours to days of wall-clock time [8], [11].

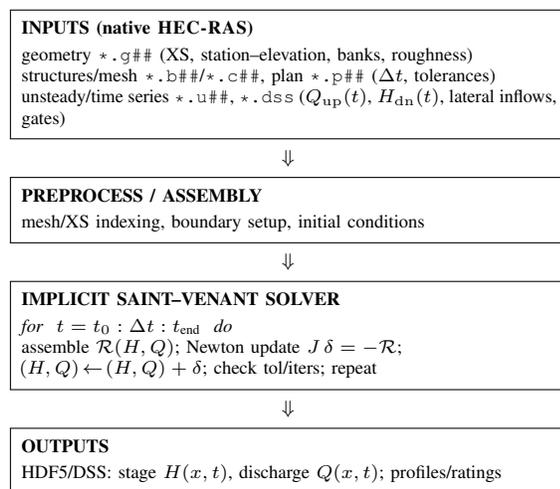


Figure 1. HEC-RAS numerical pipeline (robust non-TikZ rendering).

B. Novel Use Case: From Solver to Pre-Processor

This work treats HEC-RAS not as an end-to-end simulation tool, but as a powerful data-generation engine. By leveraging its mature GIS and project-management capabilities [8], we can assemble consistent geometries, meshes, and boundary hydrographs directly from the native project bundle [8], [12], [25], [26]. We then export this curated file set into a machine-learning pipeline. The surrogate ingests these inputs, learns the hydraulic relationships and returns reach-scale forecasts in seconds rather than hours [11]. In this workflow HEC-RAS becomes a build tool for high-quality training data, while the surrogate supplies the speed needed for rapid what-if analyses.

C. The HEC-RAS File Ecosystem

The key to this approach is the structured, information-rich file bundle that constitutes a standard HEC-RAS project. These files contain all the static, quasi-static, and dynamic information required to train a robust surrogate model, as summarized in Table I.

TABLE I. HEC-RAS FILE BUNDLE ORGANISED BY INFORMATION TYPE. ‘##’ DENOTES VERSION INDICES.

Files	Class	Key contents
<i>Static geometry</i>		
*.g##	XS/1-D	Station–elevation pairs, banks, centre-line
*.c##	2-D mesh	Cell polygons, bed elevation, roughness zones
*.b##	1-D structs	Bridge and culvert shapes, pier spacing
<i>Quasi-static metadata</i>		
*.p##	Plan	Geometry/flow linkage, solver tolerances
<i>Dynamic time-series</i>		
*.u##	Unsteady flow	Hydrograph pointers, gate schedules, run window
*.dss	DSS	Upstream $Q(t)$, downstream $H(t)$, lateral inflows

D. Key Hydraulic Terminology

To interpret the model inputs and outputs, we define the following core terms:

Reach

A contiguous channel segment between two network break-points (e.g. a confluence or control structure)[8]. Our model operates on a single reach at a time, advancing from an *upstream node* (boundary inflow Q_{up}) to a *downstream node* (boundary stage H_{dn}).

Stage (H)

The water-surface elevation at a cross-section, referenced to a project datum such as NAVD 88[7]. Units: metres.

Discharge (Q)

The volumetric flow rate through a cross-section, defined as positive in the downstream direction[7]. Units: $m^3 s^{-1}$.

III. RELATED WORK

Our work builds on advances in three key areas: data-driven hydraulic modeling, autoregressive sequence prediction, and neural operators for scientific computing.

A. Data-Driven Surrogates for River Hydraulics

Early data-driven surrogates for river hydraulics often relied on feed-forward neural networks or polynomial meta-models to emulate one or two cross-sections at a time [11]. More recent studies have scaled to full reaches by coupling convolutional encoders with graph neural networks [10], and physics-informed neural networks have now been demonstrated for single-reach stage prediction [24]; yet many approaches remain restricted to steady-flow conditions or simplified rectangular channels [9], [11]. In contrast, our study targets the entire unsteady-flow

regime of the Mississippi River model, encompassing 71 distinct reaches and thousands of irregularly spaced, natural-geometry cross-sections.

B. Autoregressive Models for Temporal Dynamics

Autoregressive (AR) models, which forecast the next state by feeding back their own previous outputs, form the backbone of classical time-series analysis [20]. The closed-loop structure is computationally efficient for long-horizon roll-outs, but a known weakness is *error accumulation*: small mistakes are recycled and amplified, ultimately drifting the forecast away from reality [21].

To mitigate this, modern hydrology has shifted from classical ARMA models to Recurrent Neural Networks (RNNs) such as Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs) [17], [22]. GRUs use update and reset gates to regulate information flow, capturing temporal dependencies while remaining parameter-efficient. When applied to river networks these RNNs typically predict each gage independently, failing to capture the spatial physics that connect them [22]. Our work addresses this by embedding a GRU within a spatial operator, allowing the recurrence to span both time and space. Furthermore, we anchor the AR loop at every step with the true boundary hydrographs (Q_{up} , H_{dn}), providing a strong physical constraint that drastically reduces long-term drift.

C. Neural Operators for Spatial Dependencies

To model the spatial physics, we turn to the Fourier Neural Operator (FNO), which learns mappings between function spaces via global convolutions in the spectral domain [2]. By modulating Fourier modes directly, FNOs capture long-range spatial dependencies with high efficiency and are essentially discretisation invariant [2], [4]. The Geo-FNO variant extends this concept to irregular meshes by injecting coordinate information into the spectral block, making it well suited to the non-uniform cross-section spacing found in river models [3]. Previous studies have already employed two-dimensional FNOs for rapid flood-inundation mapping [23]; here we adopt a one-dimensional Geo-FNO specifically tailored to the chain-like topology of a river reach.

D. Positioning This Work

Combining recurrent networks with neural operators is an emerging and powerful tool for modeling complex spatio-temporal systems [18].

A key aspect of our work is its training methodology. We show that the network learns the underlying hydraulic behavior implicitly from the data itself. This is achieved through a carefully engineered eight-channel feature vector that encodes the system's essential physical drivers: the channel geometry (z_{bed} , z_{bank}), frictional properties (n_{man}), and the mass and energy constraints imposed by boundary hydrographs (Q_{up} , H_{dn}).

The success of this approach, using a standard mean-squared error objective with a smoothness regularizer, demonstrates that meticulous feature engineering is a powerful and efficient tool for instilling physical consistency in a data-driven surrogate.

IV. METHODOLOGY

We build a one-hour-ahead, autoregressive surrogate that advances using the last $L=12$ hours of state. At each step, the network consumes this history and outputs the next-hour stage and discharge (\hat{H}_{t+1} , \hat{Q}_{t+1}); the prediction is appended to the history and the loop repeats over the forecast horizon (Figure 2).

A. Input Feature Vector

At hour t and cross-section i , we form a per-section feature vector $x_t(i)$ composed of three groups:

$$x_t(i) = \left[\underbrace{H_t(i), Q_t(i)}_{\text{dynamic (2)}} \mid \underbrace{z_{bed}(i), z_{bank}(i), n_{man}(i), x_{coord}(i)}_{\text{static (4)}} \mid \underbrace{Q_{up}(t), H_{dn}(t)}_{\text{boundary (2, broadcast over } N)}} \right]. \quad (1)$$

This *base* interface has $C_{in}=8$ channels. In our implementation, we also include two lightweight auxiliaries: depth $D_t(i) = \max\{H_t(i) - z_{bed}(i), 0\}$ and a seasonal phase, yielding $C_{in}=10$ total channels. We assemble training tensors of shape $[B, L, N, C_{in}]$ with $L=12$ hours. The 1-D coordinate is also passed as a positional input to the encoder, so its first layer receives $(C_{in}+1)$ inputs.

B. Network Architecture: A Recurrent Neural Operator

The surrogate employs a hybrid architecture that couples a GRU for temporal feature extraction with a Geometry-Aware 1-D Fourier Neural Operator (FNO) for spatial dependencies.

- 1) **Encoder**: A linear layer lifts the 10-channel input vector and its spatial coordinate x_{coord} (11 total inputs) to a 96-dimensional latent space.
- 2) **Temporal Block**: A single-layer GRU (hidden size 96) processes the 12-hour encoded sequence at each cross-section, capturing temporal dynamics and outputting its final hidden state.
- 3) **Spatial Block**: The resulting tensor of final hidden states (shape $[B, N, 96]$) is processed by a 1-D FNO. The FNO applies a global convolution in the frequency domain across the spatial dimension (N), efficiently modeling long-range dependencies. We use up to 48 Fourier modes, proportional to N .
- 4) **Decoder**: A final linear layer maps the 96-dimensional FNO output to the two target variables: the predicted stage (\hat{H}_{t+1}) and discharge (\hat{Q}_{t+1}) for the next hour.

C. Rationale for the GRU-Geo-FNO Architecture

Figure 3 highlights that river hydraulics can be viewed as two coupled 1-D signals: (i) a *spatial* profile along chainage at a fixed time (bed and stage across all cross-sections), and (ii) a *temporal* trace at a fixed cross-section (stage/flow through time). We therefore split modeling duties accordingly.

Spatial coupling (Geo-FNO). At each step, we form an ordered vector over the N cross-sections and apply a 1-D

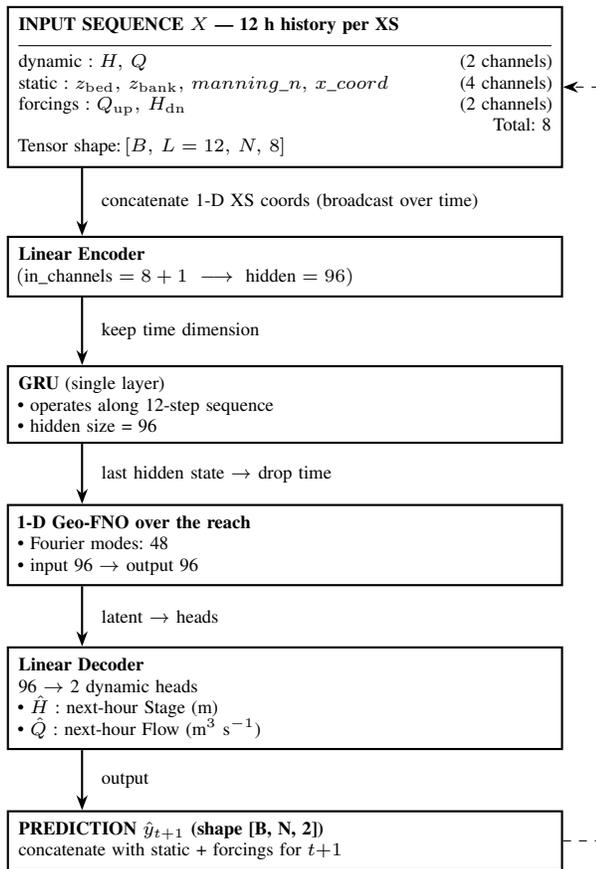


Figure 2. Autoregressive GRU-GeoFNO surrogate architecture. Vertical arrows share a common inset; the dashed loop feeds predictions back as inputs for the next step.

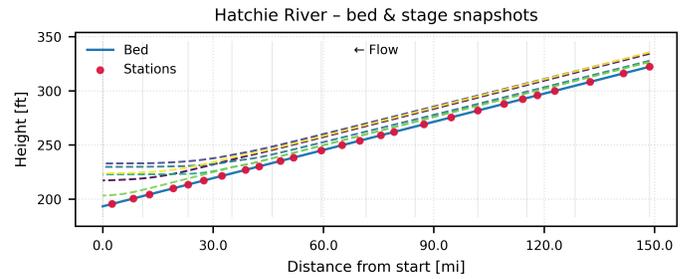
Fourier neural operator across this spatial axis. Because the FNO performs a global spectral convolution, every output depends on all cross-sections simultaneously, enabling upstream–downstream interactions (e.g., backwater/attenuation) to be learned in a single pass. Injecting the 1-D coordinate (x_{coord}) makes the operator geometry-aware, accommodating irregular cross-section spacing without re-meshing.

Temporal memory (GRU). For each cross-section, short-term dynamics are encoded by a single-layer GRU that processes a 12-hour window and returns a compact latent state. This summarizes the recent hydrograph (rising/falling limbs, lags) and supplies the per-section features that the FNO then exchanges across space.

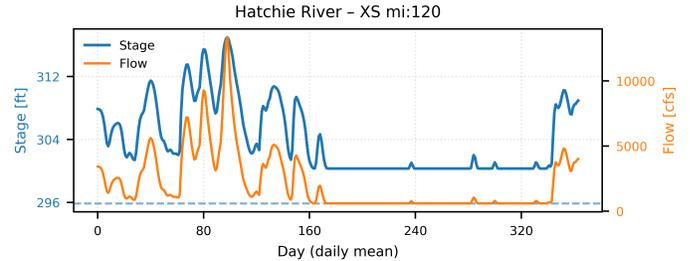
Result. The GRU provides local temporal context; the Geo-FNO propagates that context globally along the reach. This division cleanly matches the physics of 1-D hydraulics while remaining efficient and discretization-robust.

D. Training Objective and Inference

Training Objective. The network minimizes a composite loss that combines a data-fidelity term with physics-informed regularizers to promote stable, realistic predictions. The total



(a) Spatial signal at selected times: bed and stage snapshots vs. distance along the reach.



(b) Temporal signal at one cross-section: daily stage (fill shows depth) and discharge.

Figure 3. River hydraulics as 1-D signals in space and time: space is modeled by a Geo-FNO, time by a GRU.

loss \mathcal{L} is

$$\begin{aligned} \mathcal{L} = & W_H \mathcal{L}_{\text{Huber}}(\hat{H}_n, H_n) + W_Q \text{MSE}(\hat{Q}_n, Q_n) \\ & + \lambda_n \|\Delta^2 \hat{H}_n\|_2^2 + \lambda_t \|\Delta_t^2 \hat{H}_t\|_2^2 \\ & + \lambda_\Delta \|\Delta_t \hat{H} - \Delta_t H\|_2^2 + \lambda_{\text{bias}} \|\bar{\hat{H}} - \bar{H}\|_2^2. \end{aligned} \quad (2)$$

Here, $(\cdot)_n$ denotes normalization by training-set statistics. The stage (H) term uses a Huber loss (with $\beta=0.02$); discharge (Q) uses the *mean squared error* (MSE). The regularizers act on stage in physical units: spatial smoothness (Δ^2), temporal smoothness (Δ_t^2), change-in-stage matching (Δ_t), and mean-bias control ($\bar{\cdot}$).

Here, $(\cdot)_n$ denotes variables normalized by training-set statistics. The primary data-fidelity loss is computed in this normalized space, using a Huber loss ($\beta=0.02$) for stage (H) and MSE for discharge (Q). The four regularizers, which act only on stage, are computed in real physical units to enforce consistent behavior:

- **Spatial Smoothness:** Penalizes the second spatial difference (Δ^2) of the normalized stage prediction to discourage noisy outputs.
- **Temporal Smoothness:** Penalizes the second temporal difference ($\Delta_t^2 \hat{H} = \hat{H}_t - 2\hat{H}_{t-1} + \hat{H}_{t-2}$) to reduce high-frequency oscillations over time.
- **Change in Stage:** Encourages the predicted one-hour change in stage ($\Delta_t \hat{H} = \hat{H}_t - \hat{H}_{t-1}$) to match the true change.
- **Mean Bias:** Penalizes deviations in the spatial mean of the predicted stage ($\bar{\hat{H}}$) from the true mean (\bar{H}).

We use weights $W_H=1$, $W_Q=0.3$, and regularization coefficients $\lambda_n=3e-5$, $\lambda_t=3e-5$, $\lambda_\Delta=0.35$, and $\lambda_{\text{bias}}=8e-4$. An Exponential Moving Average (EMA) of the model weights (decay 0.995) is maintained and used for all evaluations.

Inference. For multi-step forecasting, we use a 12-hour-seeded autoregressive loop. The model’s prediction ($\hat{H}_{t+1}, \hat{Q}_{t+1}$) is used to construct the input for the next step by combining it with the derived depth $\hat{D}_{t+1} = \max(\hat{H}_{t+1} - z_{\text{bed}}, 0)$, the static geometric features, the seasonal phase, and the *true* boundary forcings for hour $t+1$. This complete feature vector is then re-normalized before being passed to the model.

V. EXPERIMENTAL SETUP

A. Study Area and Data Sources

We use the U.S. Army Corps of Engineers (USACE) HEC-RAS model of the Mississippi River Basin, which contains 71 distinct one-dimensional river reaches. Data is extracted from the project’s HDF5 files, which provide static geometry (cross-section shape, roughness) and hourly simulation results (Stage H , Flow Q) for three major flood years. All data were converted from imperial to SI units (metres, m³/s).

Year	Primary Flood Event	# Hourly Snapshots
2002	June–Sept. Moderate Flood	8,737
2008	May 50-Year Flood	8,783
2011	April Historic Flood	8,737

B. Training and Evaluation Protocol

A separate surrogate model is trained for each of the 71 river reaches using a strict temporal data split.

a) *Data Splits:* The 2002 and 2008 simulations (17.5k hours) form the **training set**. The first quarter of 2011 (2.2k hours) serves as the **validation set** for early stopping. The entire **2011 year** is the final **test set**.

b) *Implementation Details:* Each model is trained for up to 120 epochs using the AdamW optimizer ($lr = 2 \times 10^{-4}$) and a batch size of 64.

c) *Evaluation via Autoregressive Rollout.:* Final performance is measured via an autoregressive rollout on the year-long 2011 test set, mimicking a real-world forecast. After seeding the model with an initial 12-hour history of true data, it iteratively predicts the next 8,725 hourly steps. At each step, the model’s prediction is combined with the true boundary forcings ($Q_{\text{up}}, H_{\text{dn}}$) to form the input for the subsequent step.

d) *Evaluation Protocol.:* The primary evaluation metric is the **Mean Absolute Error (MAE)** in predicted stage, as it provides a direct, interpretable measure of the average error in feet, which is most relevant for operational flood guidance.

$$\text{MAE} = \frac{1}{T} \sum_{t=1}^T |y_t^{\text{pred}} - y_t^{\text{true}}| \quad (3)$$

Reporting convention. Unless noted otherwise, we report *per-reach* MAE in feet (Table III).

VI. RESULTS

We evaluate each of the 71 per-reach surrogates via a full-length autoregressive rollout on the unseen 2011 hold-out year. Results are reported in wall-clock time for the full ensemble and in Mean Absolute Error (MAE, feet) for stage.

A. Computational Speedup

A primary goal of this work is to accelerate forecasting. Table II summarizes the end-to-end wall-clock time required for a 1-year (8,737-hour) ensemble forecast across all reaches. The surrogate completes this task in **40 minutes** compared to **139 minutes** for the HEC-RAS solver, a **3.45×** speedup. This acceleration is operationally significant, enabling minute-scale “what-if” analysis.

It is important to note that this benchmark represents a conservative estimate of potential gains. While the neural network itself is highly parallelizable, our current rollout implementation is a single-threaded Python loop running on a CPU. **Porting this autoregressive data-handling pipeline to a GPU would unlock substantial further acceleration**, representing a key avenue for future optimization.

TABLE II. INFERENCE TIME FOR 1-YEAR, 71-REACH FORECAST.

Model	Wall-Clock Time
HEC-RAS 5.0.1	139 minutes
Recurrent FNO Surrogate	40 minutes
Speedup Factor	3.45×

B. Predictive Accuracy

We report accuracy using Mean Absolute Error (MAE) in feet, as this metric directly answers the operational question: “Is the predicted stage within a tolerable deviation of the HEC-RAS result?”

a) *Per-Reach Performance:* Table III details the MAE for each of the 71 reaches. The performance is strong across the majority of the basin, with a median MAE of **0.28 ft** and an InterQuartile Range (IQR) of **0.06–1.04 ft**. Overall, **91.5%** of reaches (65 of 71) achieve a MAE of 2.0 ft or less, meeting a key fidelity target for stage guidance.

b) *Error Analysis:* While most reaches perform well, a smaller subset of primarily smaller tributaries drives a long tail in the error distribution. Only 6 of 71 reaches (8.4%) exceed a MAE of 2.0 ft. These challenging cases, such as the ‘Hatchie River’ and ‘St. Francis’, are known to exhibit more complex hydraulics or have sparser data representation, and they are the focus of ongoing improvements.

TABLE III. PER-REACH MEAN ABSOLUTE STAGE ERROR (MAE) FOR THE 2011 HOLD-OUT YEAR, WITH THE NUMBER OF CROSS-SECTIONS (N) AND REACH LENGTH (MILES). SORTED ASCENDING BY MAE

Reach ID (River—Reach)	XS	Len [mi]	MAE [ft]
CouleeDesGrues—1	56	5.47	0.0066
Black River—R1	3	0.24	0.0131
OldRiver6—1	2	0.12	0.0131
BayouBourdeaux—1	71	6.23	0.0164
OldRiver2—1	14	6.95	0.0164
Ouachita River—R2	79	24.16	0.0164
OldRiver4—1	78	6.06	0.0230
Ouachita River—R1	237	50.89	0.0295
BayouJeansonne—1	42	3.27	0.0328

Continues on next column/page

Table III continued from previous column/page

Reach ID (River—Reach)	XS	Len [mi]	MAE [ft]
BigCanal—1	40	3.07	0.0361
OldRiverOutflow—SidneyMurray	8	19.39	0.0427
Mississippi—Below Loosahatch	2	0.60	0.0459
OldRiverOutflow—OverBank-Aux	6	1.51	0.0459
OldRiverOutflow—RedRiv-SidMur	10	3.91	0.0492
YazooRiver—Reach1	202	59.60	0.0525
OldRiver3—1	43	3.28	0.0558
DUMMY—1	3	0.04	0.0558
outlet—1	2	100.00	0.0591
OldRiver1—1	24	3.22	0.0623
Mississippi—Below Wolf	13	18.90	0.0722
Mississippi—Below Hatchie	14	7.67	0.0755
RedRiver—BelowBlack	9	8.24	0.0755
Mississippi—Below Nonconnah	11	6.88	0.0820
OldRiverOutflow—LowSill	4	0.45	0.0984
Boeuf River—R1	78	24.17	0.0984
Mississippi—Upper Miss	63	29.63	0.0984
OldRiverOutflow—Auxiliary	5	10.10	0.1017
Ohio River—Lower SOHS	2	0.20	0.1083
Ohio River—LD 52-53	10	10.23	0.1115
BayouJosen—1	70	25.79	0.1411
Mississippi—Below Obion	30	24.48	0.1509
Mississippi—Below St. Fran	31	39.52	0.1509
Tensas River—R1	16	7.89	0.2395
Cumberland River—Cumberland River	56	10.17	0.2461
Forked Deer—Forked Deer	42	4.62	0.2493
Tennessee River—Tennessee River	138	28.06	0.3117
Forked Deer—North Fork	58	6.10	0.3150
RedRiver—AboveBlackRiver	7	7.19	0.3314
Mississippi—Below Big Muddy	51	76.37	0.4197
RedRiver—BelowNatch	9	24.66	0.4364
Straight Slough—Straight Slough	5	1.22	0.4495
Mississippi—Below Cairo	72	112.65	0.4495
Obion River—Below Forked	83	12.34	0.4724
White River—White River	90	24.36	0.5229
Mississippi—Below Arkansas	81	100.58	0.5528
Atchafalaya—BelowOldRiver	172	30.99	0.6201
St. Francis—Above SS	16	3.54	0.6412
OldRiver5—1	7	0.20	0.6537
Little River—R1	11	9.86	0.7874
Big Muddy—Reach-1	29	25.65	0.8333
Black River—R3	35	23.70	0.8825
Nonconnah Cr—Nonconnah Cr	57	14.62	1.0335
Cat Oua R—R1	41	8.21	1.0542
YazooRiver—Reach2	53	12.65	1.2172
Mississippi—Below Vicksburg	42	137.20	1.2412
Arkansas River—Arkansas River	31	28.06	1.2967
YazooRiver—Reach1.5	43	10.63	1.3419
Wolf River—Wolf River	140	15.40	1.3484
BayNatch—1	160	12.32	1.4633
White River—Below Cache	77	15.34	1.4699
W-Wit—1	7	0.07	1.8209
Morganza Outlet—To Atchafalaya	26	25.08	1.9127
Forked Deer—South Fork	56	5.18	2.0046
Black River—R2	70	26.39	2.0965
Cache River—Cache River	72	49.58	2.1428
Loosahatchie—Loosahatchie	146	35.33	2.1555
Ohio River—OHS	3	0.13	2.2605
Obion River—Reach_1	198	33.77	3.6059
St. Francis—Below SS	21	5.61	3.9633
Hatchie River—Hatchie River	161	41.16	5.7316

C. Ablation Studies: Validating Architectural Choices

To validate our design, we conducted ablation studies on the ‘Cache River—Cache River’ reach, a case spanning **49.6 miles** and discretized into **71 cross-sections**. We systematically removed key architectural blocks and feature groups from our full proposed model to quantify their contribution. The results, summarized in Table IV, demonstrate that each component is critical for achieving high fidelity.

TABLE IV. ABLATION RESULTS ON THE CACHE RIVER.

Model Configuration	Component Removed	Stage MAE (ft)
Full Proposed Model	(All components included)	2.14
<i>Architecture Ablations</i>		
No FNO Block	(Replaced with pointwise MLP)	9.12
No GRU Block	(Replaced with last-frame encoding)	11.49
<i>Feature Ablations</i>		
No Boundary Forcings	Q_{up}, H_{dn} channels	10.45
No Static Geometry	$z_{bed}, z_{bank}, n_{man}, x_{coord}$	12.62
No Derived Depth	$D = \max(H - z_{bed}, 0)$	17.82
No Seasonal Phase	s_t feature	16.04

a) *Impact of Architecture*: Both the temporal and spatial blocks of the network are essential. Removing the FNO and using a simple multilayer perceptron (MLP) head (‘noFNO’) degrades performance by over 4x (MAE 9.12 ft vs. 2.14 ft), confirming that a global spatial operator is necessary to capture long-range hydraulic dependencies. Similarly, removing the GRU’s temporal memory (‘noGRU’) and encoding only the last known time step increases error by over 5x (MAE 11.49 ft), validating the need to process a sequence history.

b) *Impact of Features*: The ablation results confirm that physically-informed feature engineering is vital. Removing the derived auxiliary channels had the most severe impact: omitting the derived depth channel increased MAE by over 8x to 17.82 ft, while removing the seasonal phase increased it by over 7x to 16.04 ft. This highlights that providing the model with features that encapsulate non-obvious physical context is critical. Removing the core boundary conditions or static geometry also caused a catastrophic drop in performance, confirming that every channel in our proposed feature vector contributes meaningfully to the final accuracy.

D. Qualitative Case Study: Error Propagation

To understand the nature of the errors in challenging cases, we examine the full-year autoregressive rollout for the ‘Hatchie River—Hatchie River’ reach, which had the highest MAE. Figure 4 shows the forecast hydrographs at twelve evenly spaced cross-sections (XS) along the reach.

A clear spatial pattern emerges. At the upstream end (e.g., XS 0, XS 7), the surrogate tracks the HEC-RAS ground truth with high fidelity, capturing the primary flood waves accurately. However, performance degrades progressively downstream. At the midpoint (e.g., XS 42), minor deviations appear. By the downstream end (e.g., XS 63, XS 70, XS 77), the model becomes unstable, introducing large, high-frequency oscillations and diverging significantly from the ground truth. This suggests a pattern of spatial error propagation, where small inaccuracies from upstream are amplified as they are passed downstream by the model.

VII. DISCUSSION

Our results show that a recurrent neural operator can emulate year-long HEC-RAS runs with operational fidelity while accelerating an ensemble forecast by 3.45x. The remaining failure modes concentrate in a small set of hydraulically

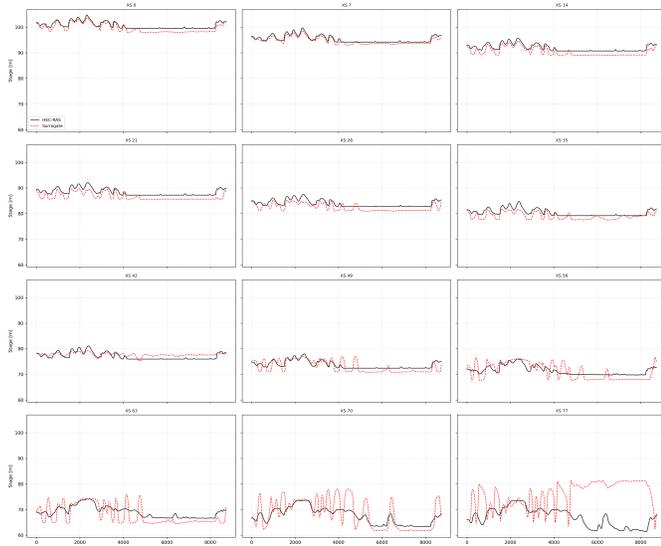


Figure 4. Year-long autoregressive rollout on the ‘Hatchie River’ reach at 12 evenly spaced cross-sections. The surrogate (red, dashed) tracks the HEC-RAS ground truth (black, solid) well at the upstream end (top-left panels) but accumulates error and develops instabilities downstream (bottom-right panels).

complex reaches, yielding a heavy-tailed error distribution (Table III). We summarize the key lessons and scope.

a) Exposure bias despite boundary anchoring.: The Hatchie River case (Fig. 4) illustrates classic autoregressive drift: small upstream errors accumulate and are propagated downstream by the FNO’s global spatial coupling. Clamping the true boundary hydrographs (Q_{up} , H_{dn}) at every step acts as a strong physical prior mirroring HEC-RAS inputs and prevents catastrophic divergence in most reaches. The residual instabilities on Hatchie indicate that internal dynamics can overwhelm this anchoring. This suggests two complementary remedies: (i) training-time strategies that reduce exposure bias (e.g., scheduled sampling, noise injection on inputs) and (ii) richer boundary/forcing information or data assimilation for reaches with complex internal hydraulics.

b) Hydraulic regime, not size, drives difficulty.: Geometric scale (reach length, cross-section count) shows no reliable relationship with MAE in Table III. Long, well-gauged main-stem reaches (e.g., *Mississippi—Below Vicksburg*) are modeled accurately, whereas shorter tributaries with backwater effects or prolonged low-flow spells (e.g., *Hatchie River, St. Francis—Below SS*) dominate the tail. In practice, the governing factor is the frequency and persistence of regimes under-represented in training, not the number of cross-sections.

c) Operational scope and limits.: Today, the surrogate is best used as a *scenario-analysis accelerator* for known hydrologic regimes: (i) models are trained and evaluated per reach, so network-scale feedbacks across confluences are not yet represented; (ii) performance depends on the hydrologic diversity seen in training generalization to far out-of-distribution events is fragile; and (iii) evaluation assumes true boundary

forcings; operational deployment will inherit uncertainty from boundary forecasts. Addressing these gaps will likely require graph neural operators for topology-aware coupling, targeted data augmentation to balance low-flow/backwater regimes, and experiments with perturbed or forecast boundary conditions to quantify skill degradation.

d) Takeaway.: Accuracy at minute-scale cost came from three ingredients: (1) an architecture that separates temporal memory (GRU) from global spatial coupling (Geo-FNO), (2) physics-aware features including derived depth and seasonal phase and (3) a stabilizing loss. The ablations substantiate each ingredient’s contribution and explain where the current model fails, thus charting a concrete path to basin-scale, autonomous forecasting.

VIII. CONCLUSION

We presented a recurrent neural-operator surrogate for 1-D HEC-RAS that delivers year-long, reach-wide forecasts at operational fidelity. On the unseen 2011 hold-out across 71 reaches, the model achieves a **3.45×** end-to-end speedup (40 vs. 139 minutes) while maintaining a **median stage MAE of 0.28 ft**, with **91.5% of reaches \leq 2 ft**.

Accuracy at minute-scale cost follows from three ingredients validated by ablations: (i) a GRU for short-term memory coupled to a Geo-FNO for global spatial coupling; (ii) physics-aware features, especially derived depth and seasonal phase; and (iii) a stabilizing loss together reducing Cache River MAE from 12.9 ft (plain MSE) to 2.14 ft and preventing drift.

Current scope is intentionally conservative: models are trained per reach (no network-scale feedbacks yet), skill degrades for out-of-distribution regimes, and rollout stability assumes true boundary forcings. These constraints point to clear upgrades: topology-aware operators for basin coupling, data/augmentation to balance challenging hydraulic regimes, assimilation or forecasted boundaries to quantify resilience, and a GPU rollout pipeline to unlock further wall-clock gains.

The surrogate is a practical *scenario-analysis accelerator* today and a potentially viable path toward basin-scale, near-real-time flood guidance with some further enhancements.

REFERENCES

- [1] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.
- [2] Z. Li, *et al.*, “Fourier neural operator for parametric partial differential equations,” in *Proc. ICLR*, 2021.
- [3] Z. Li, D. Z. Huang, B. Liu, and A. Anandkumar, “Fourier neural operator with learned deformations for PDEs on general geometries,” *arXiv preprint arXiv:2207.05209*, 2022.
- [4] N. Rahaman, *et al.*, “On the spectral bias of neural networks,” in *Proc. 36th Int. Conf. Machine Learning (ICML)*, vol. 97, pp. 5301–5310, 2019.
- [5] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 770–778.
- [6] M. Zoch, *et al.*, “Physics-informed neural network surrogate models for river stage prediction,” *arXiv:2503.16850*, 2025.
- [7] H. Chanson, *Hydraulics of Open Channel Flow*. Butterworth-Heinemann, 2004.

- [8] G. W. Brunner, "HEC-RAS river analysis system, hydraulic reference manual," U.S. Army Corps of Engineers, Hydrologic Engineering Center, Davis, CA, Tech. Rep. CPD-69, 2021.
- [9] P. Benner, S. Gugercin, and K. Willcox, "A survey of projection-based model reduction methods for parametric dynamical systems," *SIAM Review*, vol. 57, no. 4, pp. 483–531, 2015.
- [10] H. Bi, J. Lin, and X. Zhao, "Data-driven surrogate modeling in computational fluid dynamics: A review," *Computers and Fluids*, vol. 258, p. 104764, 2023.
- [11] Z. Takbiri-Borujeni, J. T. Smith, and M. D. White, "Challenges of using surrogate models for environmental simulations: A case study in hydrology," *Environmental Modelling and Software*, vol. 134, p. 104882, 2020.
- [12] M. Flanagan, *et al.*, "Hydraulic splines: A hybrid approach to modeling river channel geometries," *Computing in Science and Engineering*, vol. 9, no. 5, pp. 4–15, 2007.
- [13] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 652–660.
- [14] M. Abdelguerfi, *3D Synthetic Environment Reconstruction*. Springer, 2012.
- [15] M. Chung, *et al.*, "The geospatial information distribution system (GIDS)," in *Succeeding with Object Databases*. John Wiley & Sons, 2001, pp. 357–378.
- [16] Z. Li, *et al.*, "Graph neural operator for PDEs," *arXiv:2010.08895*, 2020.
- [17] K. Cho, *et al.*, "Learning phrase representations using RNN Encoder–Decoder for statistical machine translation," in *Proc. EMNLP*, Doha, Qatar, 2014, pp. 1724–1734.
- [18] J. Brandstetter, D. E. Worrall, and M. Welling, "Message passing neural PDE solvers," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2022. [Online]. Available: <https://arxiv.org/abs/2202.03376>
- [19] J. Pathak, *et al.*, "FourCastNet: A global data-driven high-resolution weather model using adaptive Fourier neural operators," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. [Online]. Available: <https://arxiv.org/abs/2202.11214>
- [20] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control*. San Francisco, CA: Holden–Day, 1970.
- [21] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer, "Scheduled sampling for sequence prediction with recurrent neural networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 28, Montréal, Canada, 2015, pp. 1171–1179.
- [22] F. Kratzert, D. Klotz, J. Herrnegger, G. Hochreiter, and S. Nearing, "Toward improved predictions in ungauged basins: Exploiting the power of machine learning," *Water Resources Research*, vol. 55, no. 12, pp. 11 344–11 364, 2019.
- [23] A. Stenta, J. Smith, and L. Hennig, "FloodFNO: Fourier neural operators for real-time inundation mapping," in *Proc. IAHR World Congress*, 2023.
- [24] M. Zoch, *et al.*, "Physics-informed neural network surrogate models for river stage prediction," in *Proc. ALLDATA 2025 – 11th Int. Conf. on Big Data, Small Data, Linked Data and Open Data*, 2025.
- [25] R. Wilson, *et al.*, "Geographical data interchange using XML-enabled technology within the GIDB system," in *XML Data Management: Native XML and XML-Enabled Database Systems*. Morgan Kaufmann, 2003, pp. 235–261.
- [26] M. Abdelguerfi (ed.), *3D Synthetic Environment Reconstruction*. Springer, 2001, vol. 611 of the NATO Science Series.