Mamba–LSTM for Long-Horizon Shoreline Prediction using Satellite Imagery

Zhen Tian, Christos Anagnostopoulos, Zhiwei Gao, Qiyuan Wang
University of Glasgow

Kostas Kolomvatsos
University of Thessaly

Presenter: Ke Xiao

ICMLT BERLIN MAY 20-22, 2026

terra An Intelligent Platform for Integrating Climate Services

ICMLT BERLIN MAY 20-22, 2026

Research Background & Motivation

The Environmental Imperative

Coastal erosion fundamentally threatens infrastructure, ecosystems, and local communities. Effective risk management requires predictive, uncertainty-aware tools capable of continuous monitoring and proactive planning.

The Observational Paradigm

Satellite remote sensing enables repeated, large-scale coastal observa-observation. Sentinel-2 provides high-revisit imagery, allowing for the extraction of highly accurate shoreline time-series data using standardized remote sensing pipelines.

The TERRA Digital Twin

Within the EU Horizon TERRA project, this research contributes to a Coastal Digital Twin. The system integrates observational data streams with predictive algorithms to map local morphological changes and long-range sediment transport processes.

Problem Formulation: Shoreline Time Series

Let $\{C(t)\}$ from t=1 to N denote a shoreline time series derived from satellite imagery. Given a history window of length L, the objective is to predict the continuous coastline at the subsequent time step:

$$\mathbf{X} = [C(t-L+1), \ldots, C(t)] \in \mathbb{R}^{L\times P\times 2} \rightarrow \hat{C}(t+1) \in \mathbb{R}^{P\times 2}$$

Notation Glossary

$C(t)$: The coastline state at time t.

P: A fixed number of ordered, two-dimensional coordinate pairs (x_i, y_i) representing the physical boundary.

L: The temporal history window (number of preceding observations).

$\hat{C}(t+1)$: The predicted sequential coordinates of the future shoreline.

The Modeling Bottleneck & Paradigm Comparison

	Standard Recurrent Networks (LSTMs)	State Space Models (Pure Mamba)	The Proposed Hybrid (Mamba-LSTM)
Processing	Step-by-step sequential processing.	Linear-time sequence modeling with selective state spaces.	Selective state space dynamics followed by sequential aggregation.
Long-Range Capacity	Struggles with extended sequence lengths; prone to memory bottlenecks.	Highly efficient processing of extended temporal histories via content-aware filtering.	Efficiently encodes infinite-horizon history into rich tokens, decoded via LSTM.
Shoreline Application	Captures local point-to-point physical continuity well, but misses multi-year seasonal erosion cycles.	Captures macro temporal trends but lacks the explicit, refined sequential decoding required for exact spatial coordinates.	Simultaneously models long-range storm/seasonal dependencies and local spatial coordinate refinement.

Proposed Architecture: Mamba-LSTM Workflow

Stage 1: Input Token Embedding

Each coastline snapshot $C(t-i)$ from the history window X is flattened and linearly projected into a d-dimensional token. Learnable spatial and temporal positional encodings are injected to preserve structural geometry.

↓

Stage 2: Selective State Space Encoding

The token sequence is passed through a stack of K Mamba blocks. Using depthwise 1D convolutions and input-dependent selectivity, the encoder filters the historical sequence, adaptively propagating critical erosion events while forgetting irrelevant noise.

↓

Stage 3: LSTM Aggregation & Decoding

The temporally rich, Mamba-encoded sequence is processed by LSTM layers. The network aggregates the filtered long-range dependencies and applies a lightweight neural decoder to regress the final $P \times 2$ coordinate matrix for the predicted shoreline.

Methodology I: Selective State Space Encoding

$$\mathbf{B}(t) = \text{Linear}_B(\mathbf{x}(t))$$

$$\mathbf{C}(t) = \text{Linear}_C(\mathbf{x}(t))$$

$$\Delta(t) = \text{softplus}(\text{Linear}_\Delta(\mathbf{x}(t)))$$

Input-Dependent Selectivity

Unlike classical State Space Models (SSMs) which utilize static transition matrices, Mamba introduces input-dependent parameters. The input projections (B,C) and the discretization step size (Δ) are formulated as functions of the input feature embedding $\mathbf{x}(t)$. The softplus function guarantees a positive continuous-time discretization step. This mechanism enables content-aware filtering, allowing the architecture to adaptively select which coastal morphological changes to remember and which to discard.

Methodology II: Sequential Aggregation & Decoding

$$\mathbf{o}_1, \ldots, \mathbf{o}_L = \mathrm{LSTM}(\mathbf{Z}^{(K)}), \ \mathbf{o}_l \in \mathbb{R}^h$$
$$\hat{C}_{t+1} = \Phi_{out}(\mathbf{o}_L) \in \mathbb{R}^{P \times 2}$$

From Latent State to Physical Coordinates

Following the K stacked Mamba blocks, the refined sequence $\mathbf{Z}^{(K)}$ requires translation back into a physical coordinate space.

The LSTM network parses the selective long-range dependencies, producing an aggregated hidden state vector o_L. A final continuous mapping function Φ_{out} reshapes this vector into the precise two-dimensional coordinate sequence mapping the future vegetation-water boundary.

Comparative Baseline: Mamba-Diffusion Benchmark

To rigorously benchmark the deterministic Mamba-LSTM framework, a conditional Denoising Diffusion Probabilistic Model (DDPM) was designed to serve as an advanced comparative baseline.

Condition Encoding

A parallel Mamba encoder processes the historical shoreline coordinates to produce a conditioning sequence C_seq. This captures the temporal distribution of the coastline without deterministic regression.

Classifier-Free Guidance

During training, the conditioning sequence is randomly dropped (p_drop). At inference, unconditional and conditional predictions are merged using a guidance scale (s = 2.0) to stabilize the generated spatial coordinates.

Cross-Attention Denoising

The forward process injects Gaussian noise into the true future shoreline. The reverse process (denoiser) projects the noisy target into a latent space, applying residual blocks and Cross-Attention tied to C_seq.

Experimental Design & Evaluation Metrics

Sentinel-2 Dataset Constraints

97 | Monthly coastal observations (May 2017 to May 2025) covering a highly dynamic site in Scotland.

256 | Uniformly-spaced coordinate pairs (**P**) per observation, extracted via LabelMe polygon annotations targeting the vegetation-water boundary.

5 | Temporal history window length (**L**), generating 92 distinct (history, target) training samples.

80/20 | Chronological train/test split.

Evaluation Metrics

Mean Squared Error (MSE) / RMSE / MAE: To quantify global point-wise coordinate displacement.

Hausdorff Distance: To measure the maximum geometric discrepancy between the curves (worst-case boundary deviation).

MaxError: To isolate extreme local prediction failures.

Quantitative Results: Performance Across Metrics

The deterministic Mamba-LSTM architecture consistently outperforms both pure LSTM recurrent models and advanced linearized/probabilistic diffusion baselines.

The hybrid approach secures the lowest error rates across RMSE, MAE, and the Hausdorff distance, proving its superiority in maintaining strict spatial geometric boundaries while predicting long-term coordinate shifts.

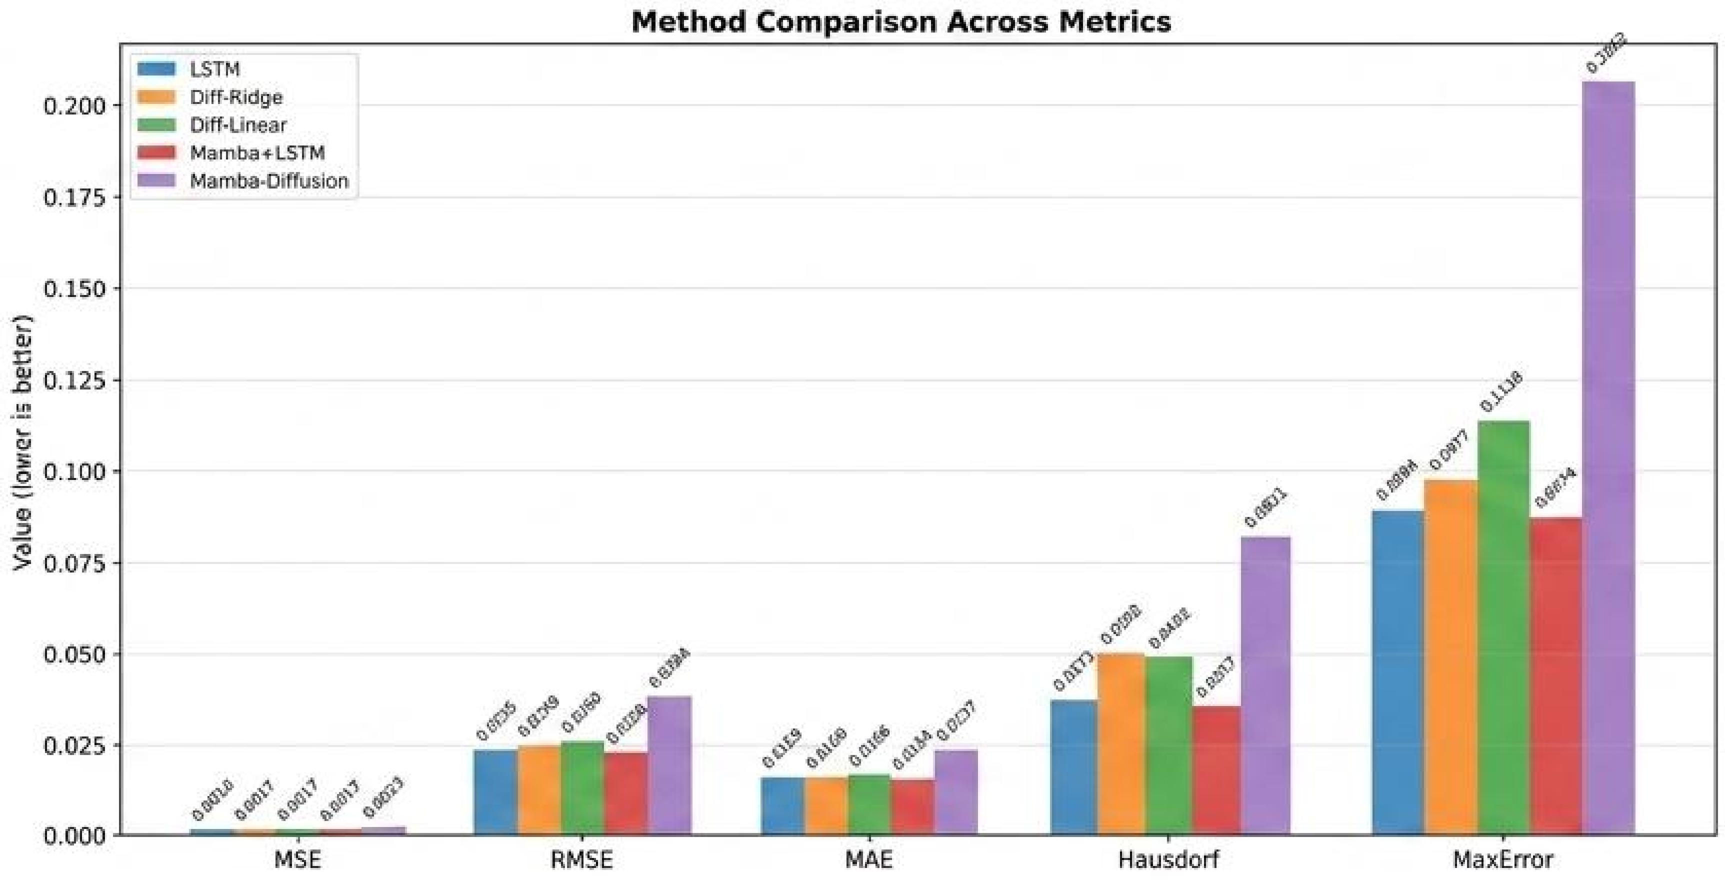

Qualitative Results: Visualizations of Future Horizons

Evaluating predictions at three distinct future time points (2024_01, 2024_09, 2025_05). While all models track the macro curve of the ground-truth coastline (black line), the pure diffusion models exhibit high variance and instability. The Mamba-LSTM framework (red line) maintains high-fidelity adherence to the complex morphological changes, specifically in regions exhibiting rapid, sharp geometric alterations.

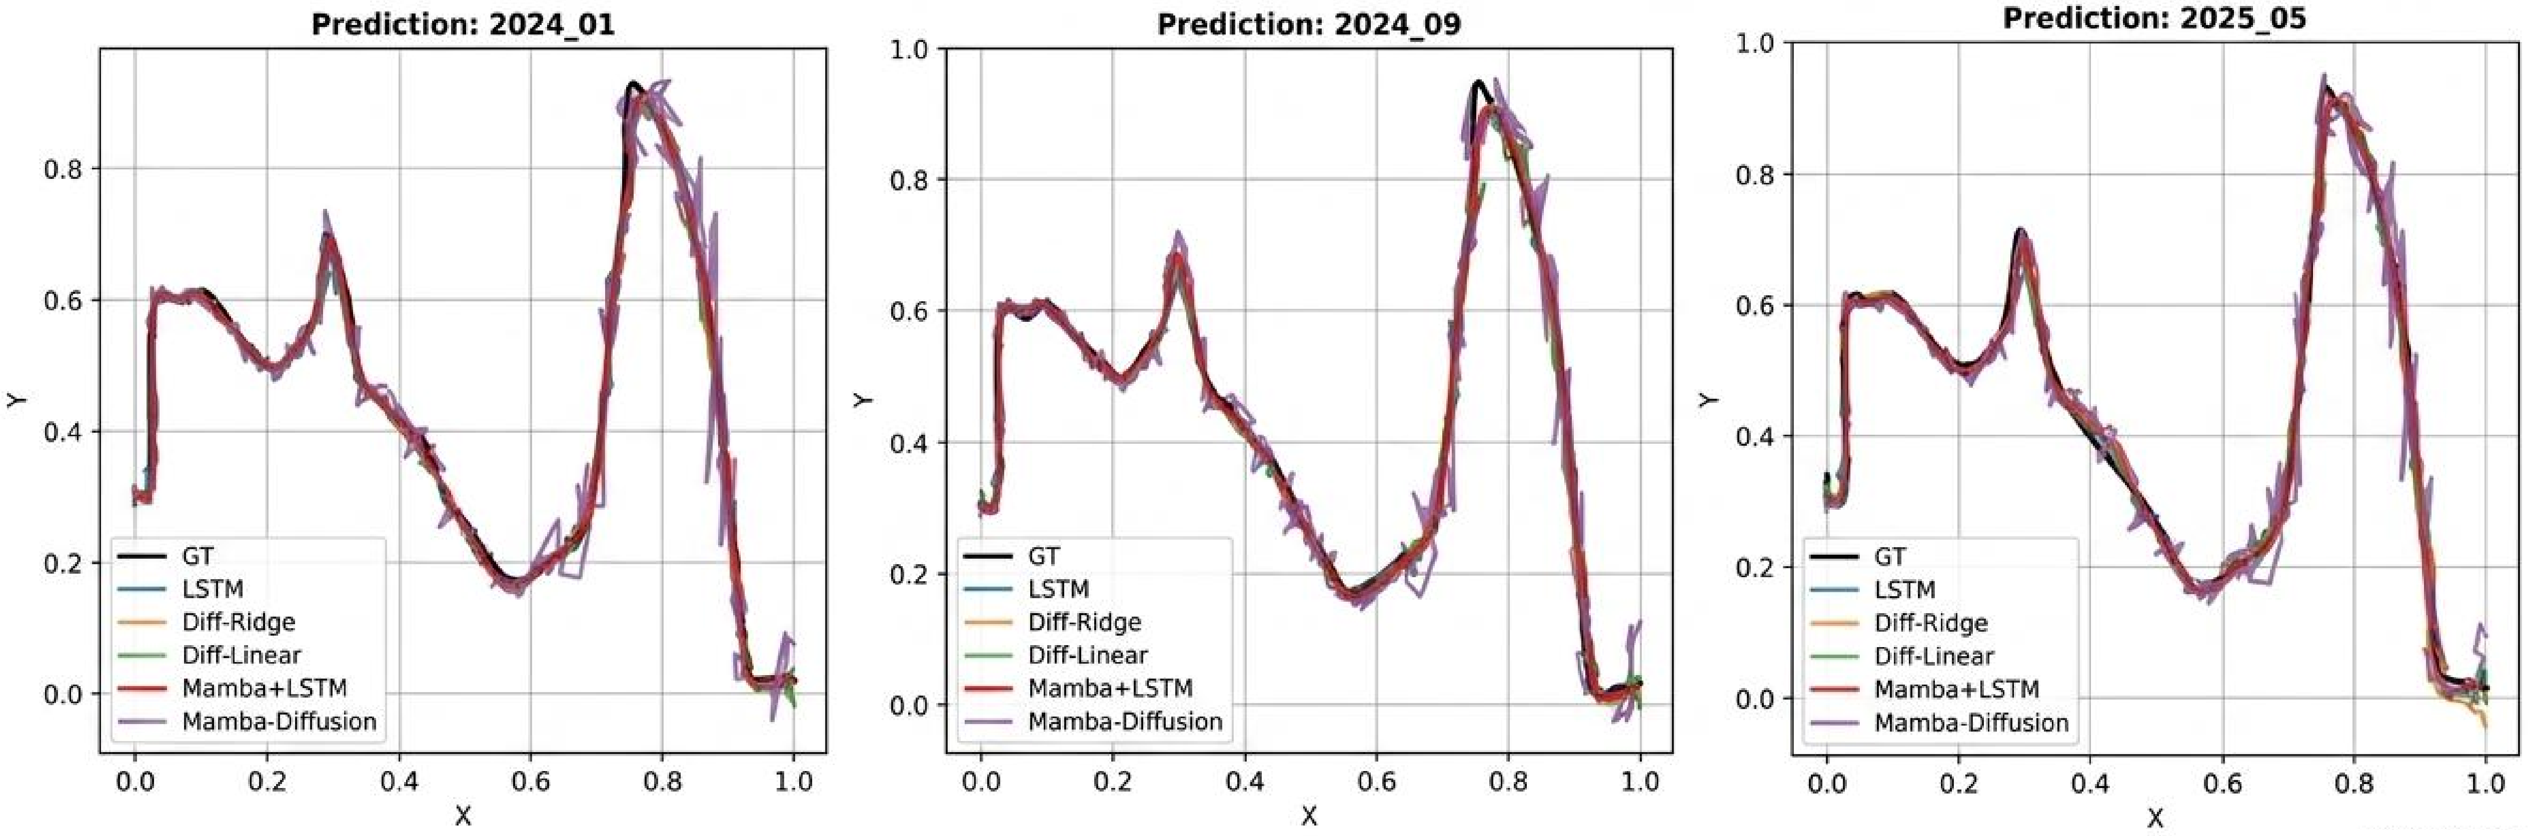

Deep Dive: Global Stability & Point-wise Error Analysis

Maximized Performance Envelope

The radar chart captures the global superiority of the Mamba-LSTM model, maximizing the enclosed area across all five competitive metrics.

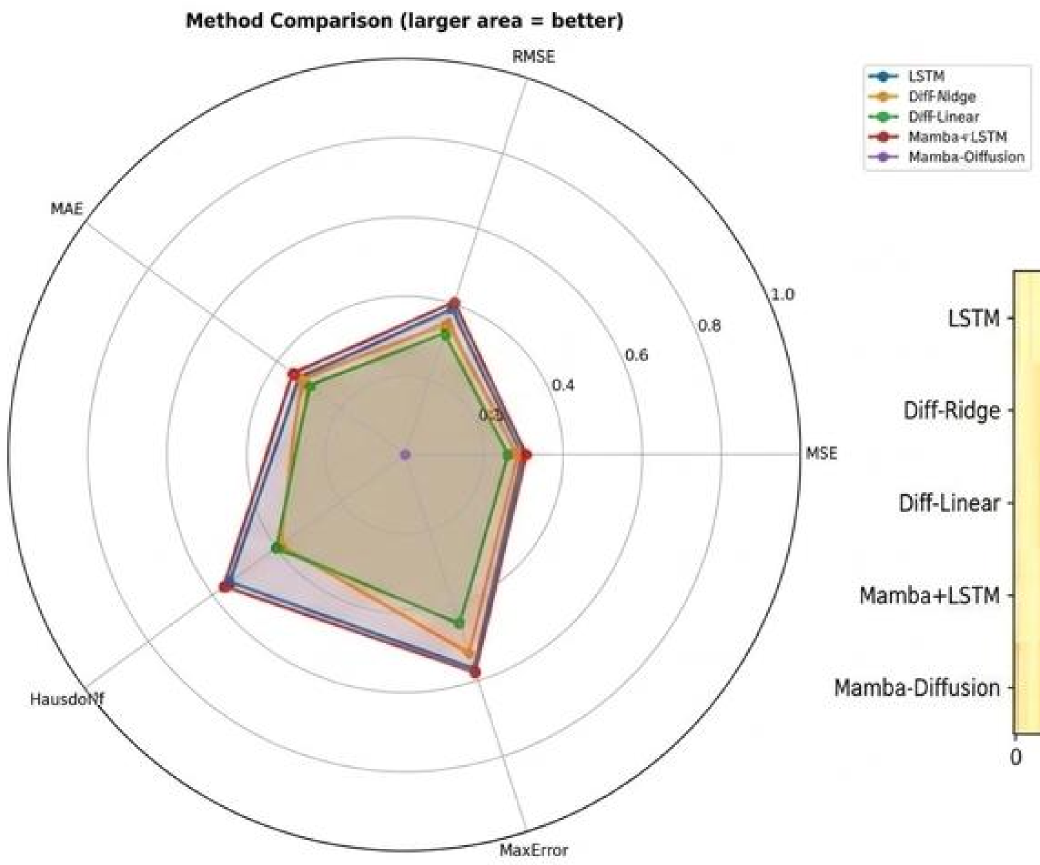

Isolating Local Divergences

The point-wise error heatmap isolates specific coordinate indices along the coastline. The Mamba-LSTM model maintains a consistently pale profile (minimal Euclidean error), starkly contrasting with the extreme localized failures (red spikes) generated by the diffusion baselines.

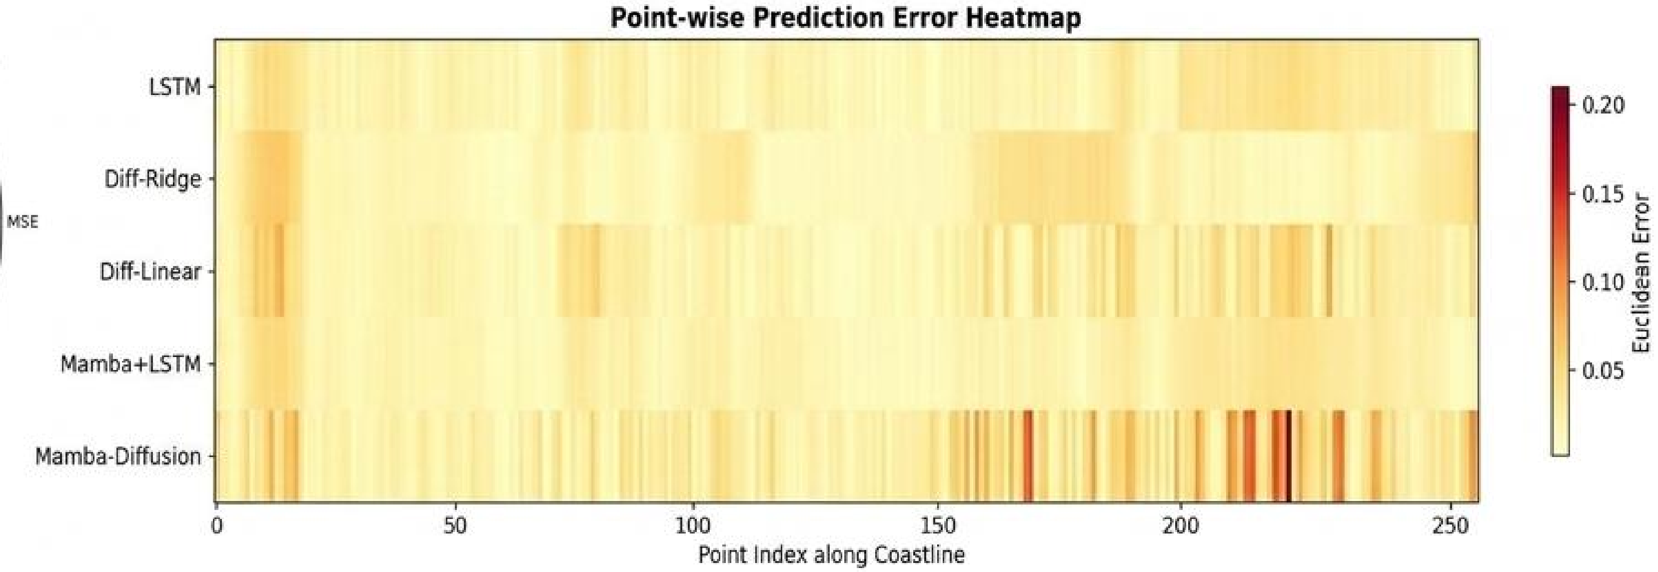

terra
An Intelligent Platform for Integrating Climate Services

University of Glasgow

Knowledge & Data Engineering Systems

Conclusion & Future Trajectories

The hybrid Mamba-LSTM architecture solves the temporal bottleneck in coastal morphodynamic prediction, marrying infinite-horizon selective filtering with precise sequential coordinate decoding.

- Superior Accuracy
Consistently outperforms LSTMs and conditional DDPMs in tracking complex vegetation-water boundaries.

- Operational Viability
Linear-time sequence modeling proves highly efficient for integration into continuous-stream satellite observation pipelines.

- Next Steps
Scaling the architecture to multi-site global datasets, extending forecast horizons

tian.zhen@glasgow.ac.uk

ICMLT BERLIN MAY 20-22, 2026