BlockGPT - Spatio-Temporal Modelling of Rainfall via Frame-Level Autoregression

Cristian Meo, Varun Sarathchandran, Avijit Majhi, Shao-Hsuan Hung, Carlo Saccardi, Ruben Imhoff, Roberto Deidda, Remko Uijlenhoet & Justin Dauwels

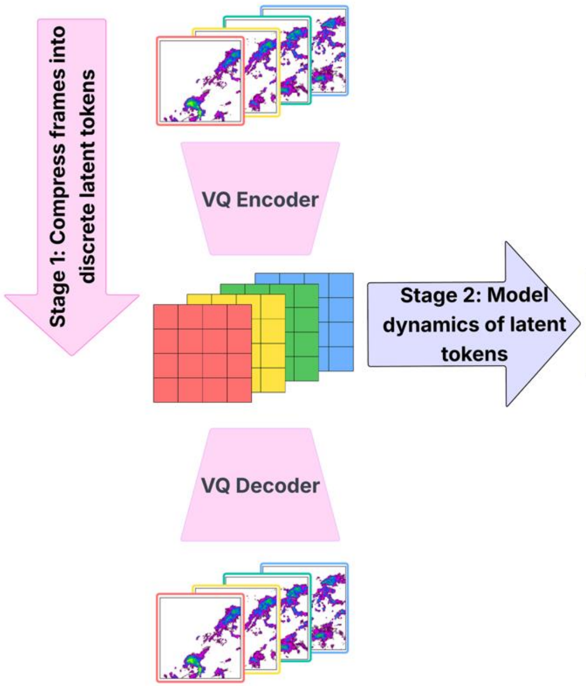

PRESENTED BY
Avijit Majhi

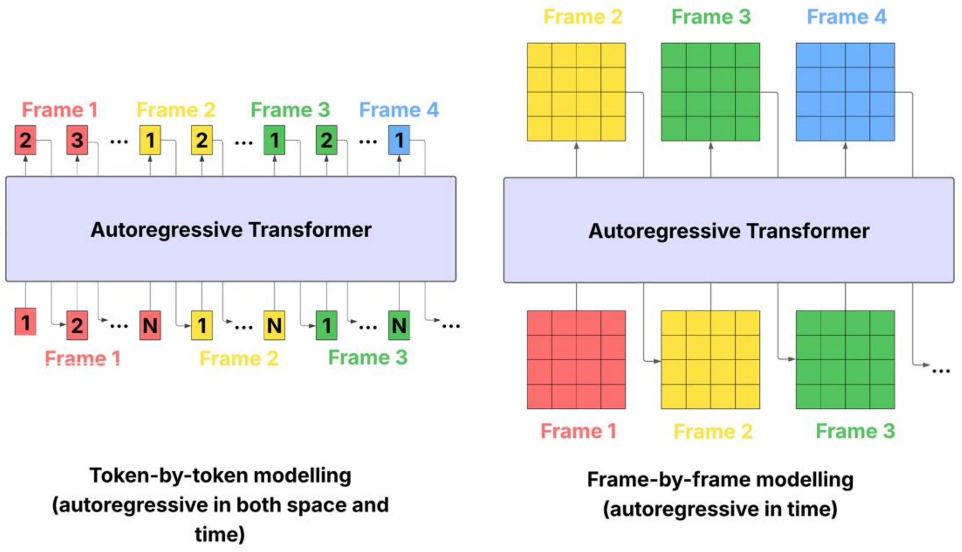

LatentWorlds AI

Delft University of Technology

UNICA
UNIVERSITÀ DEGLI STUDI DI CAGLIARI

Deltares

Nowcasting

Short-term weather forecasting, typically from a few minutes up to six hours (WMO).

Why not NWP?

- Low latency
- Coarse resolution

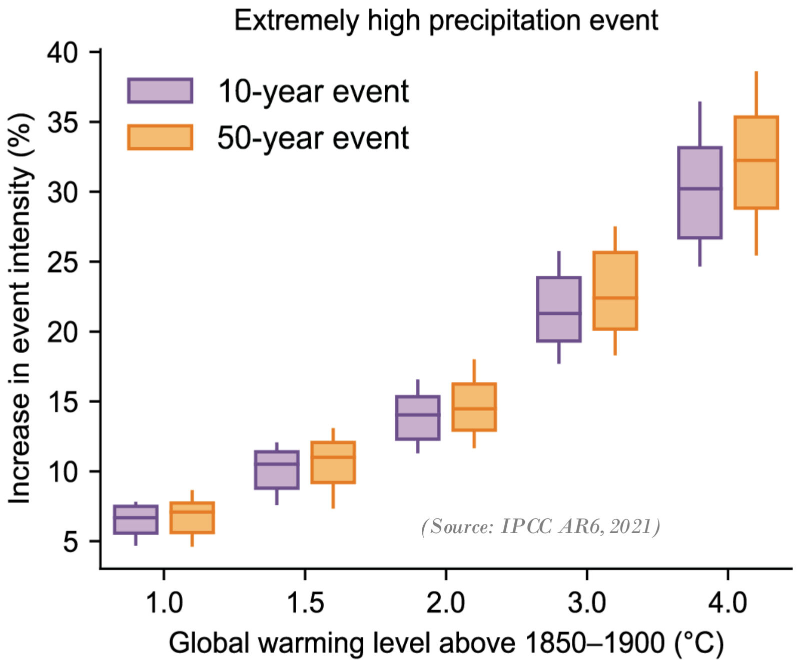

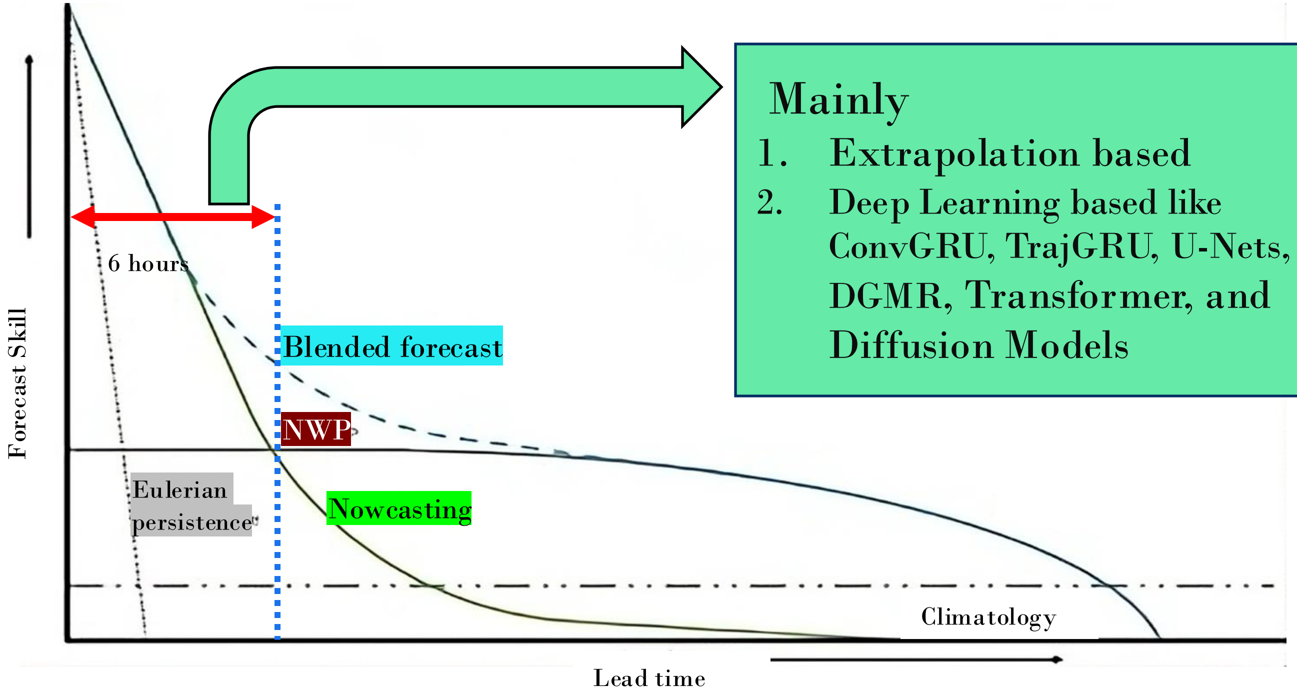

Mainly

1. Extrapolation based
2. Deep Learning based like ConvGRU, TrajGRU, U-Nets, DGMR, Transformer, and Diffusion Models

(Source: Scale-dependent blending of ensemble rainfall nowcasts and numerical weather prediction in the open-source pysteps library, Imhoff et al., 2023)

LatentWorlds AI TU Delft Delft University of Technology UNICA Università degli Studi di Cagliari Deltares NeurIPS Climate Change AI

Limitations of the previous models

Paradigm & Model	Critical Limitations
Token-Based Autoregressive e.g., NowcastingGPT	• Suffer from a flawed inductive bias. • They flatten 2D precipitation maps into a 1D sequence of tokens and predict them one by one. • This fundamentally misunderstands the spatial nature of weather, • Leading to fragmented, less coherent predictions and • Extremely slow, token-by-token inference.
Diffusion-Based (e.g.,PreDiff or DiffCast+PhyDnet)	• Foundational diffusion models like PreDiff are computationally prohibitive, with training times exceeding 30 days. • SOTA methods like DiffCast+PhyDnet mitigate this by modelling only the forecast residual; they still introduce significant architectural complexity and potential latency compared to a single, end-to-end generative model.

Concepts behind NowcastingGPT

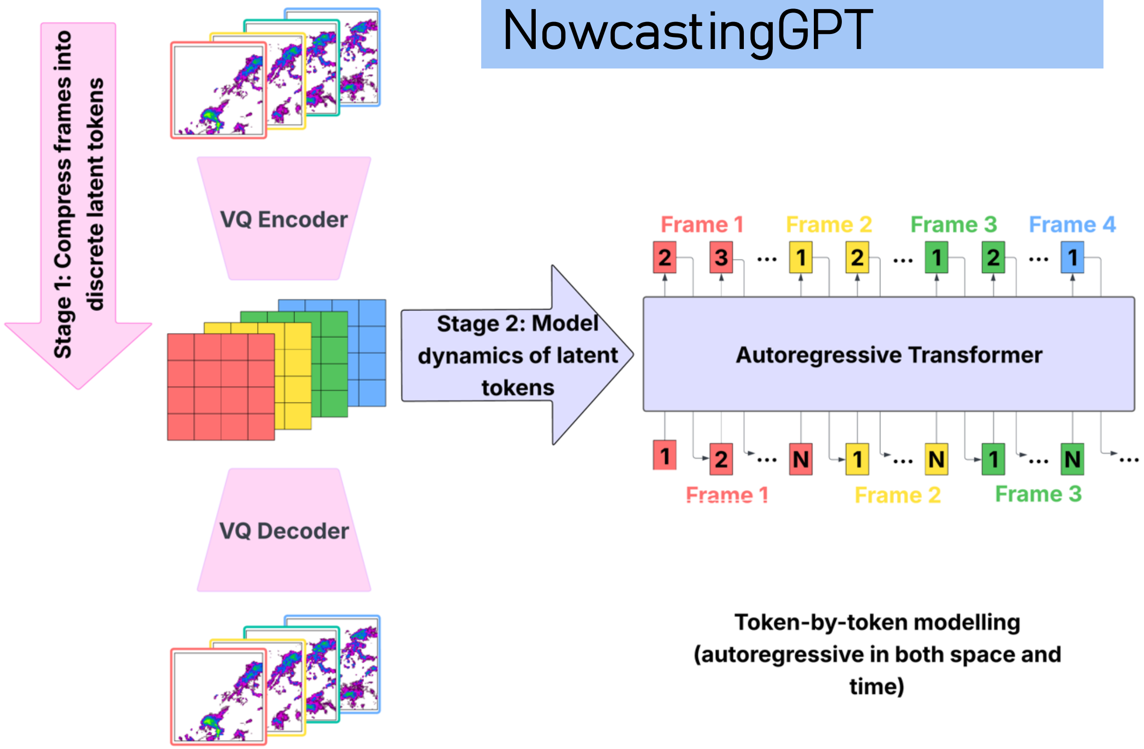

LatentWorlds AI

TU Delft Delft University of Technology

UNICA UNIVERSITÀ DEGLI STUDI DI CAGLIARI

Deltares

NeurIPS

Climate Change AI

The BlockGPT Methodology: A Paradigm Shift

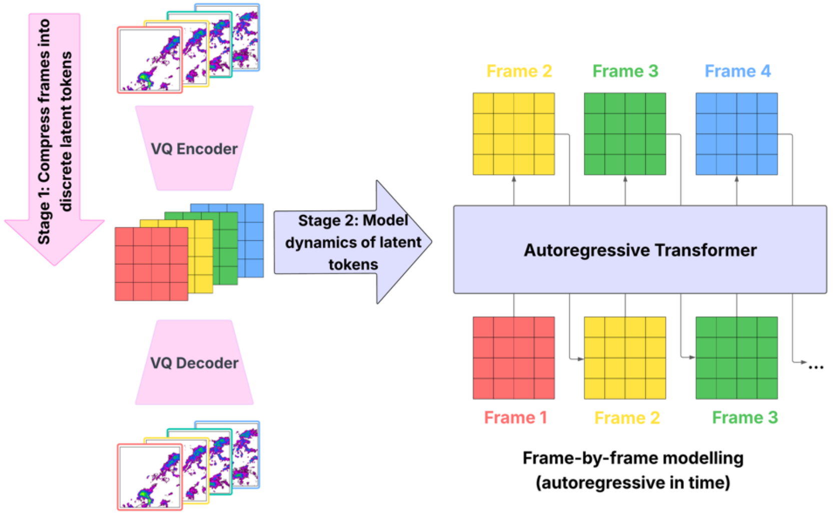

Given $\mathcal{X}_{\text{context}} = \{\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_{T_c}\}$, $\mathbf{X}_t \in \mathbb{R}^{H \times W}$, predict $\mathcal{X}_{\text{target}} = \{\mathbf{X}_{T_c+1}, \ldots, \mathbf{X}_T\}$.

Learn $f_{compress} : \mathcal{X} \mapsto \mathcal{T}$ **and** $f_{predict} : \mathcal{T}_{context} \mapsto \mathcal{T}_{target}$

NowcastingGPT flattens 2D fields,

$\mathbf{z} = \text{flatten}(\mathcal{T})$, and models $p(\mathbf{z})$ as

$$p(\mathbf{z}) = \prod_{i=1}^{T \cdot H'W'} p\left(z^{(i)} \mid z^{(1)}, z^{(2)}, \ldots, z^{(i-1)}\right).$$

This treats an video like a 1D structure.

BlockGPT instead directly models $\mathcal{T} = \{\mathbf{T}_1, \mathbf{T}_2, \ldots, \mathbf{T}_T\}$ **as**

$$p(\mathcal{T}) = \prod_{t=1}^{T} p\left(\mathbf{T}_t \mid \mathbf{T}_1, \ldots, \mathbf{T}_{t-1}\right)$$

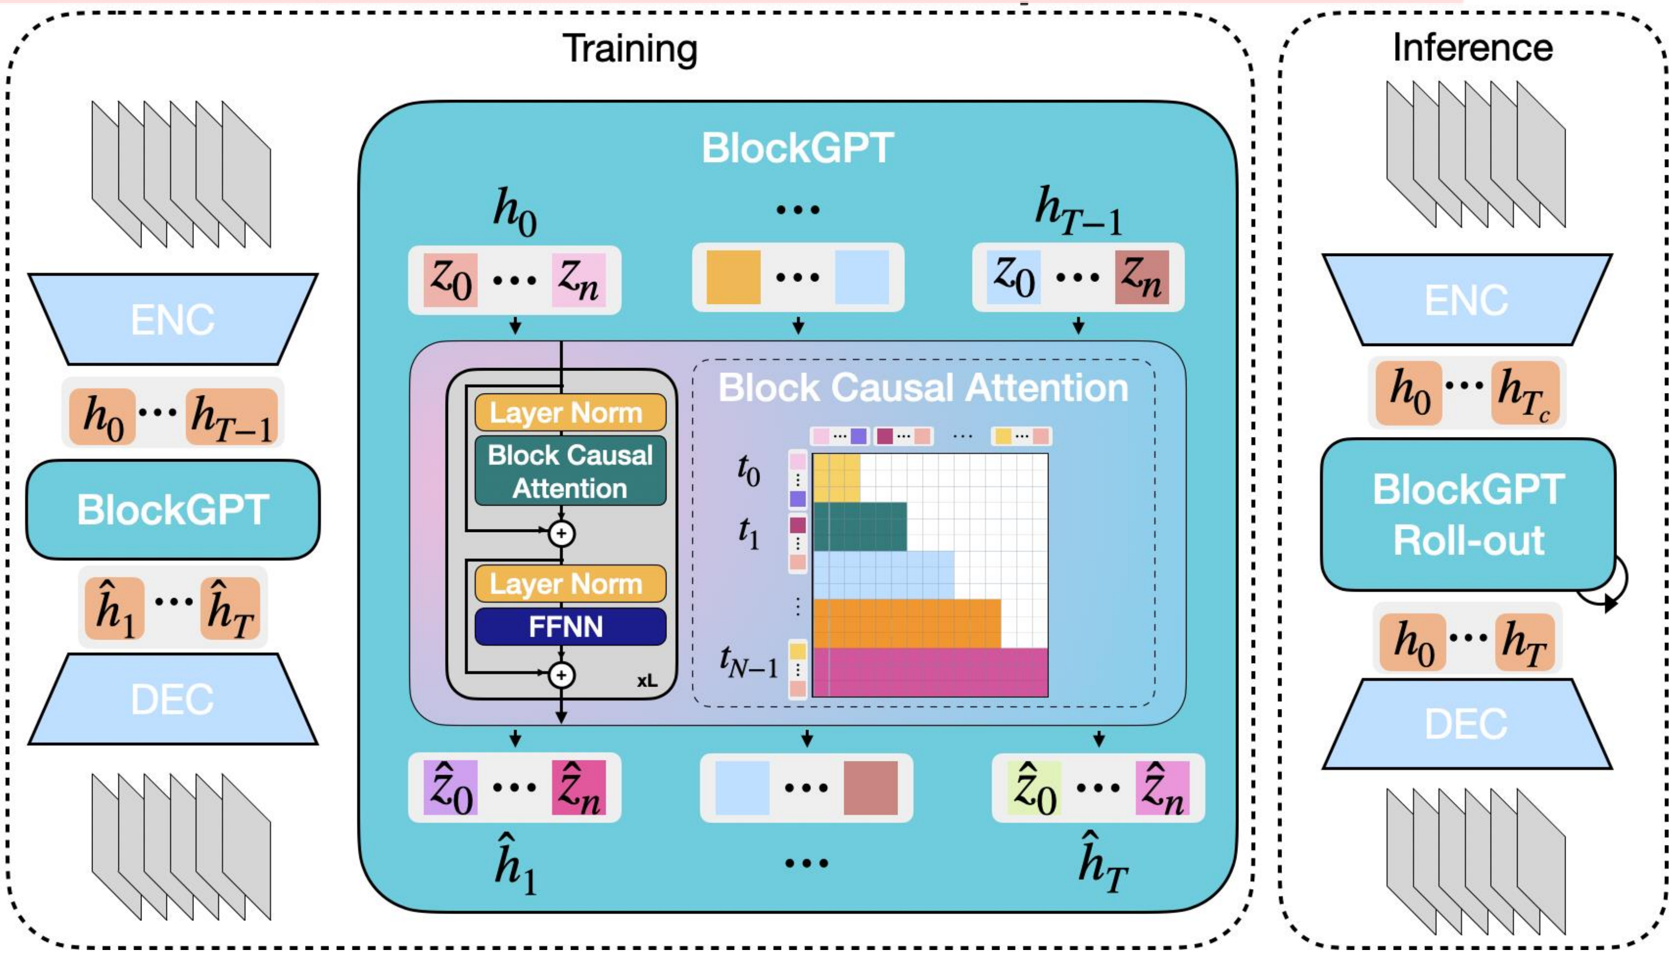

LatentWorlds AI · TU Delft – Delft University of Technology · UNICA – Università degli Studi di Cagliari · Deltares · NeurIPS · Climate Change AI

Datasets and Experimental Design

We used two distinct and challenging radar datasets to ensure our model's robustness:

- KNMI: A high-resolution, 1 km^2, 5-minute resolution dataset from the Netherlands.
- SEVIR: A large-scale storm event dataset from the United States with 1 km^2, 5-minute resolution

Forecasting Task: The primary task was to predict the next **6 radar precipitation fields** given **3 context fields**. With each field representing a 30-minute interval.

KNMI

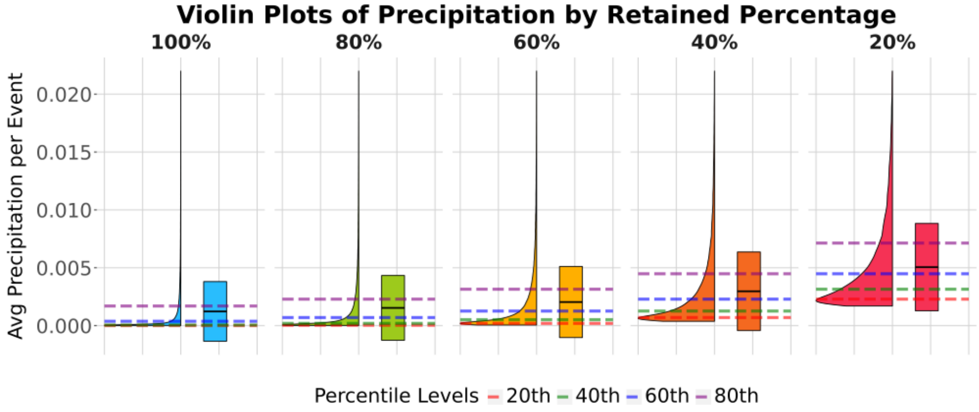

SEVIR

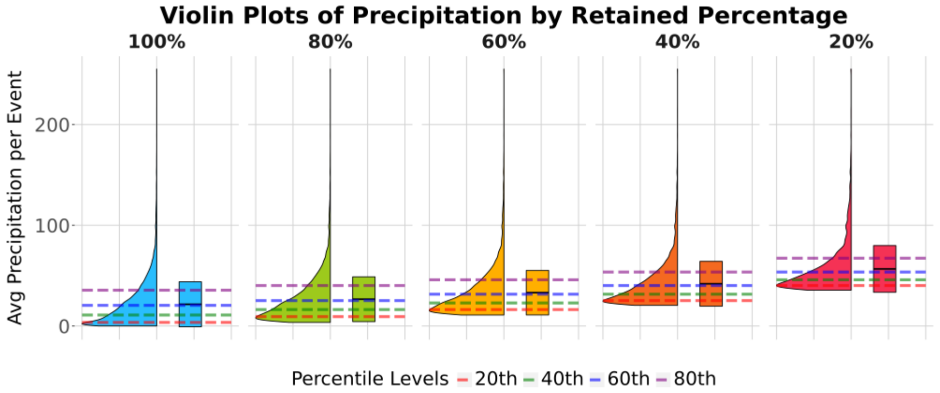

LatentWorlds AI TU Delft Delft University of Technology UNICA Università degli Studi di Cagliari Deltares NeurIPS Climate Change AI

Results and Analysis: Demonstrating State-of-the-Art Performance

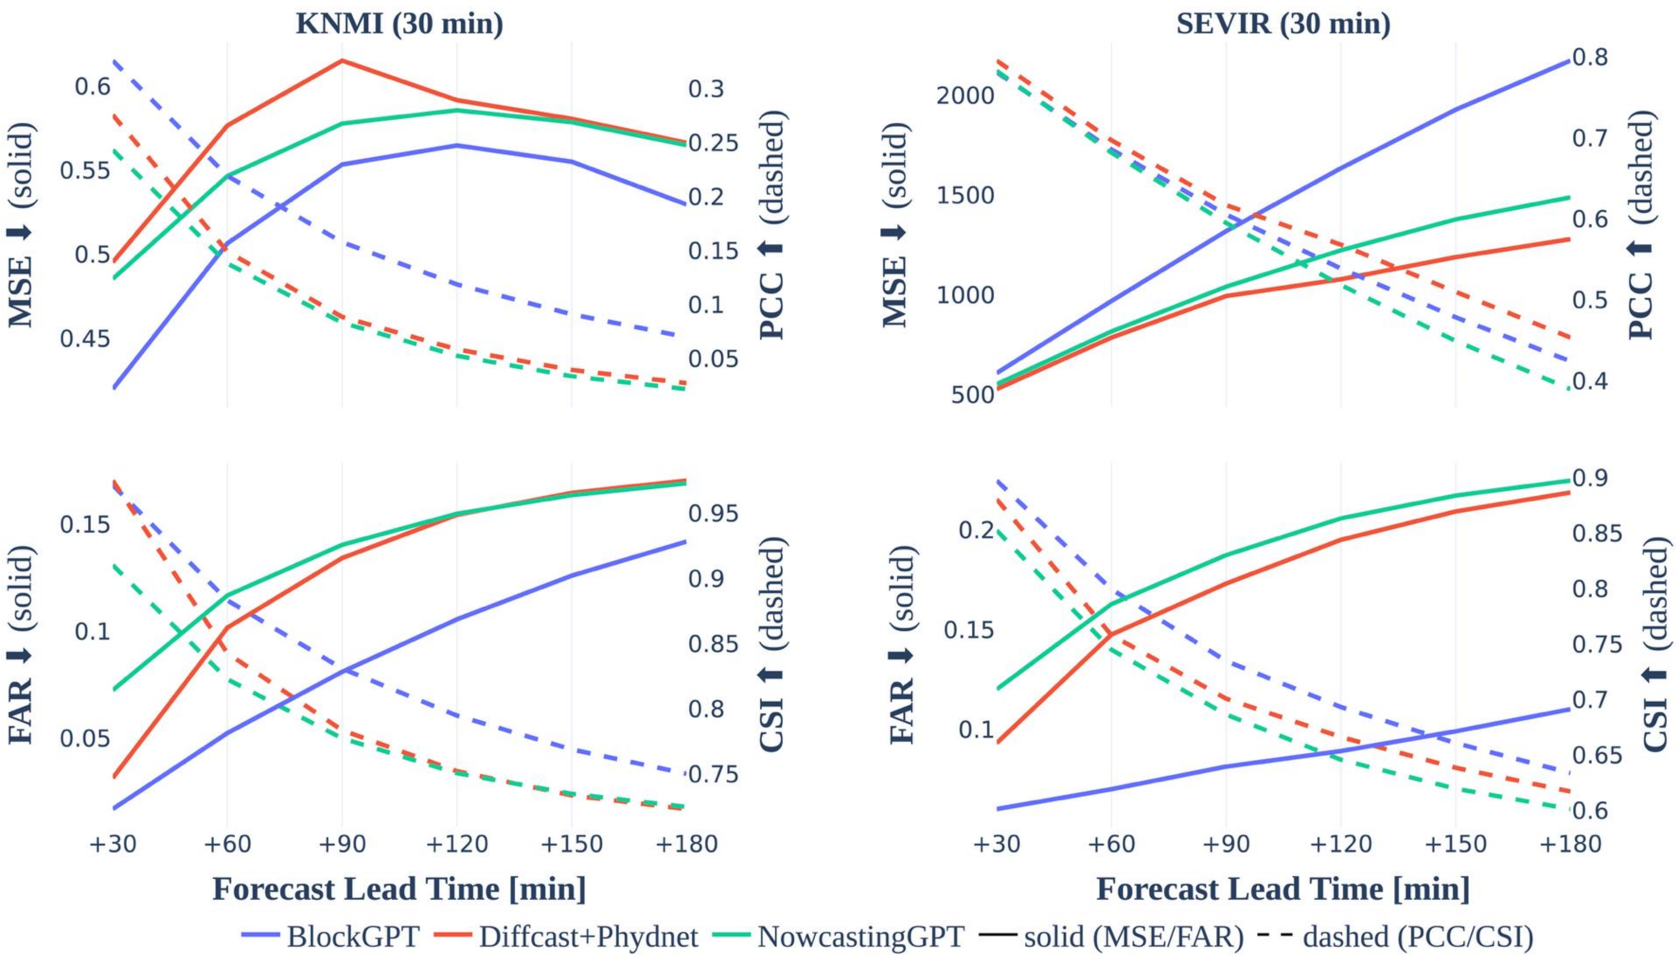

- Across both the KNMI and SEVIR datasets, BlockGPT consistently outperformed the baselines on the critical categorical metrics of CSI and FAR
- DiffCast+PhyDnet achieved lower MSE on the SEVIR dataset – which shows the double penalty effect

LatentWorlds AI

TU Delft Delft University of Technology

UNICA UNIVERSITÀ DEGLI STUDI DI CAGLIARI

Deltares

NeurIPS

Climate Change AI

Enhanced Qualitative Coherence (SEVIR)

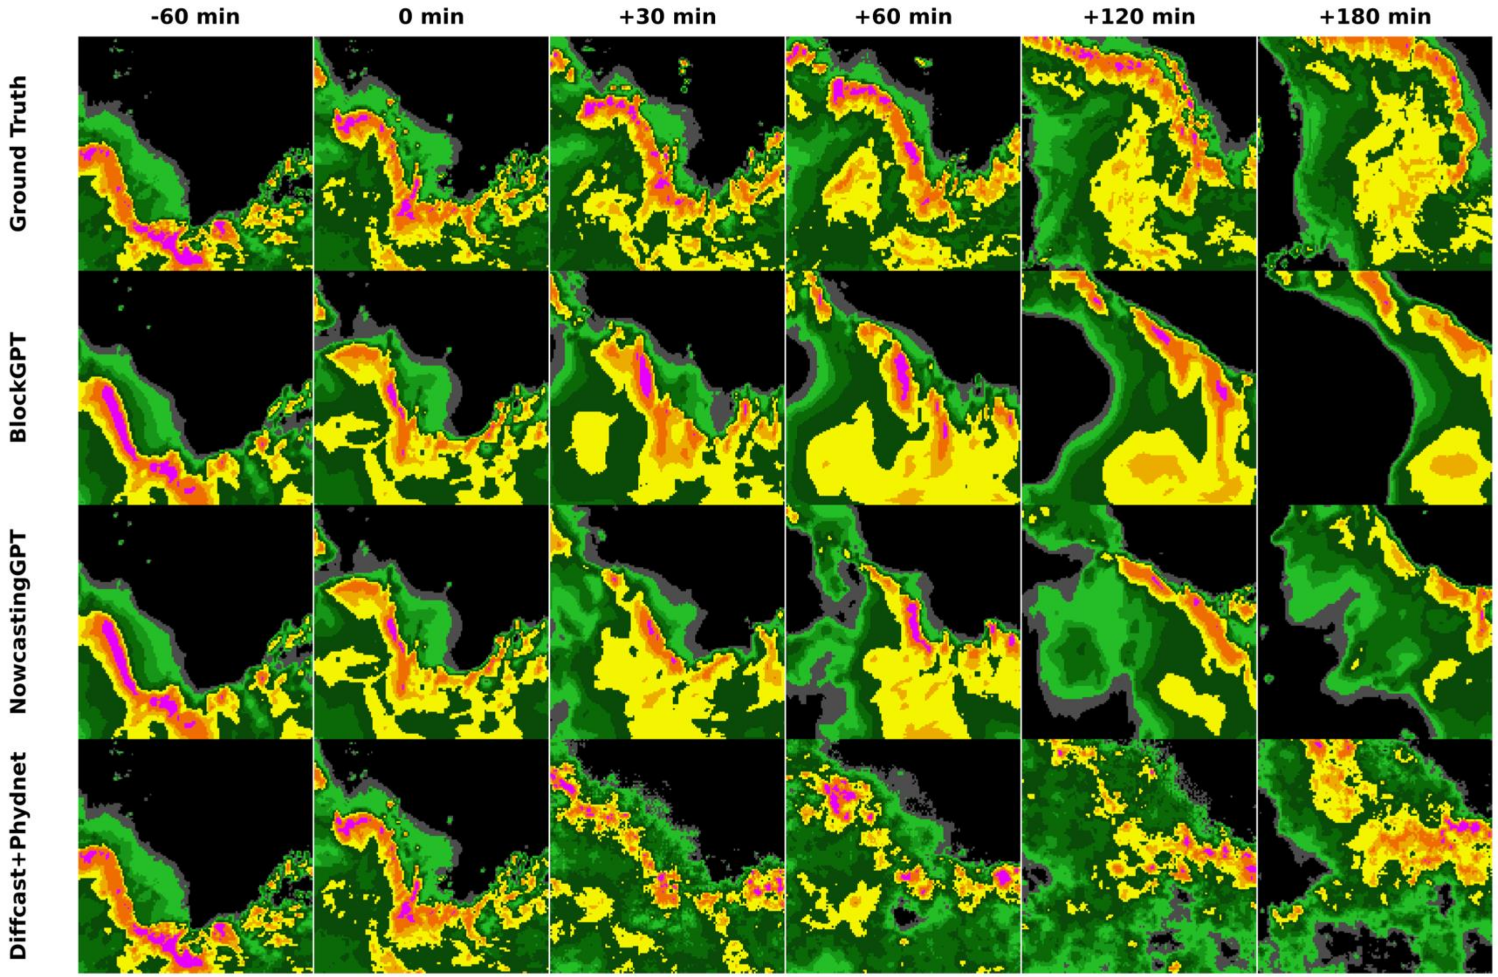

BlockGPT preserves the morphology and displacement of storm systems much more faithfully, generating predictions that look realistic and physically plausible.

In contrast, outputs from NowcastingGPT often appeared fragmented,

DiffCast+PhyDnet exhibit blobby artefacts and fail to recover the linear organisation of storm systems.

Enhanced Qualitative Coherence (KNMI)

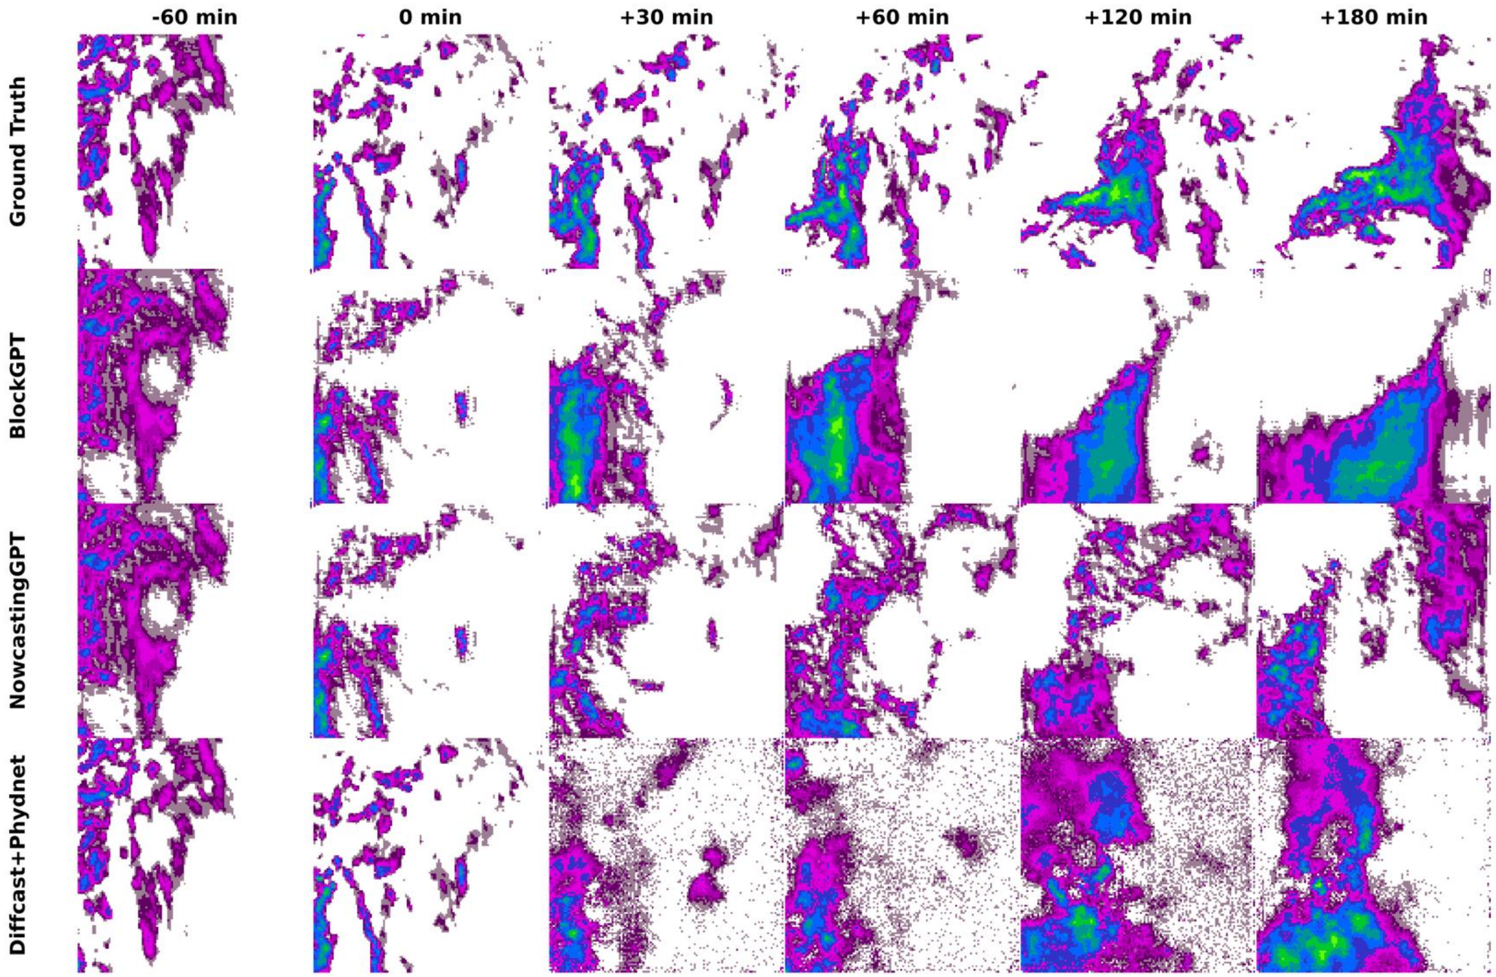

BlockGPT preserves the morphology and displacement of storm systems, which is realistic and physically plausible.

NowcastingGPT appeared fragmented.

DiffCast+PhyDnet exhibits artefacts and fails to recover the linear organisation of storm systems.

Robustness in High-Impact Scenarios – Catchment based analysis

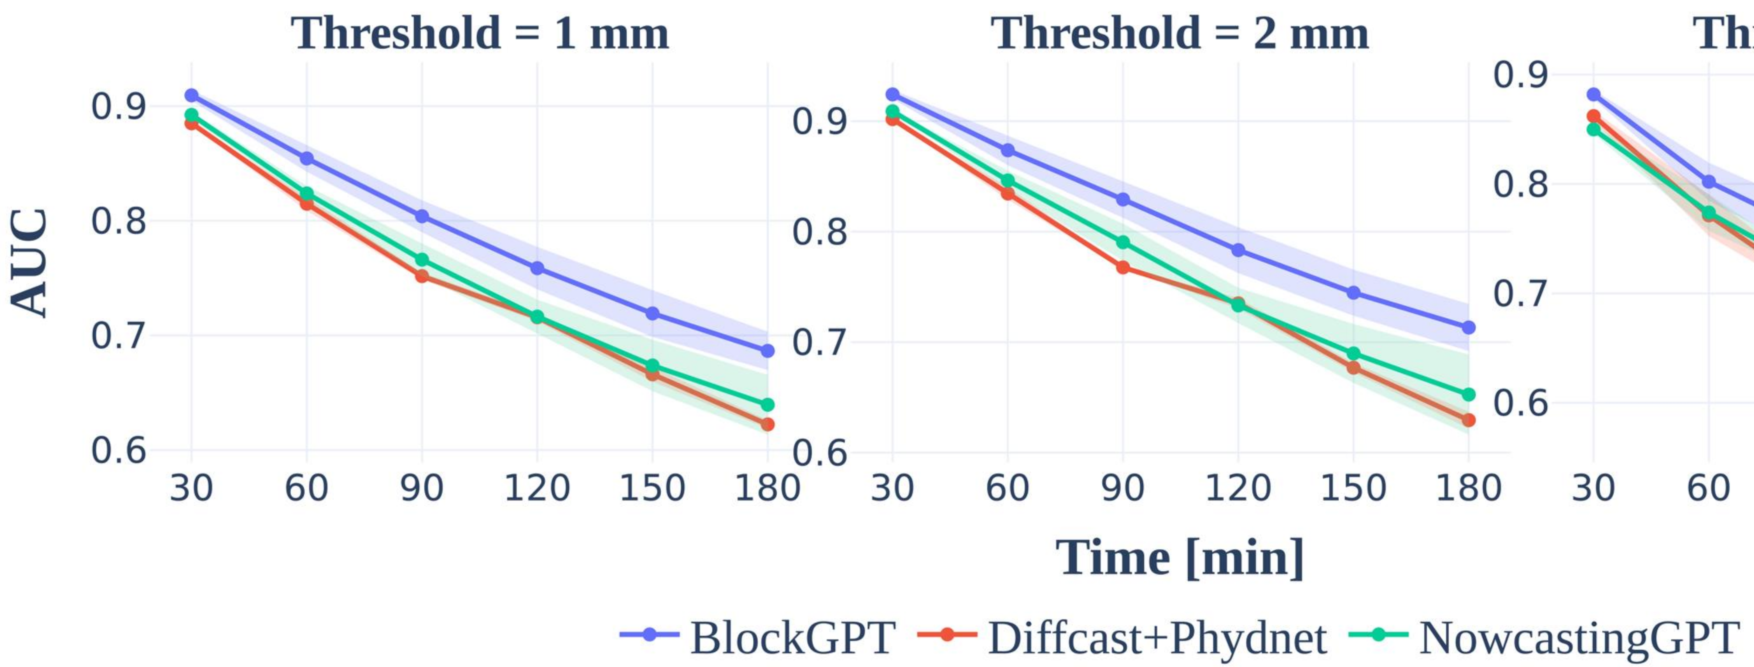

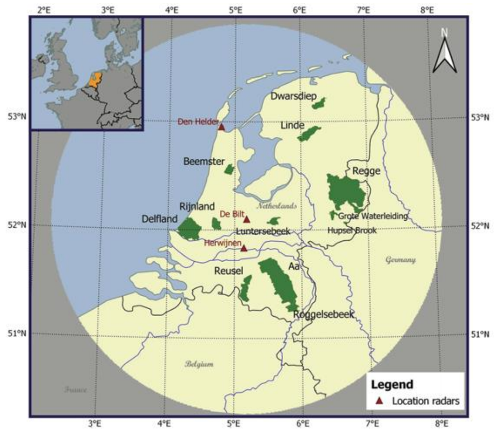

BlockGPT consistently achieved the highest Area Under the ROC Curve (AUC) at all rainfall intensity thresholds

LatentWorlds AI
TU Delft Delft University of Technology
UNICA UNIVERSITÀ DEGLI STUDI DI CAGLIARI
Deltares
NeurIPS
Climate Change AI

Unprecedented Computational Efficiency

Model	Inference Time (in Seconds)
DiffCast+PhyDnet	8.17
NowcastingGPT	7.09
BlockGPT	0.26

BlockGPT delivered its state-of-the-art forecasts with inference speeds up to 31x faster than our baselines

LatentWorlds AI

Conclusion and Future Directions

- Shifted the modelling paradigm from token-level to frame-level autoregression
- BlockGPT achieves superior performance on the metrics most critical for event localisation
- With inference speeds up to 31 times faster than its peers, this demonstrates its practical value for real-world applications.
- Explore the integration of BlockGPT as a powerful backbone within a residual diffusion framework.

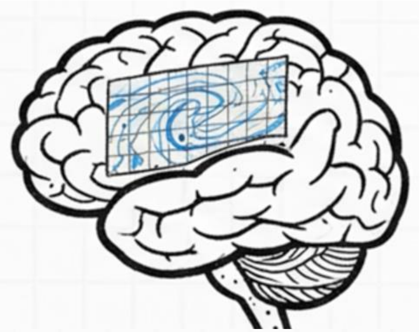

NeurIPS Climate Change AI

BlockGPT - Spatio-Temporal Modelling of Rainfall via Frame-Level Autoregression

Cristian Meo, Varun Sarathchandran, **Avijit Majhi**, Shao-Hsuan Hung, Carlo Saccardi, Ruben Imhoff, Roberto Deidda, Remko Uijlenhoet & Justin Dauwels

Thank you for your attention!!!

Scan below for more info!

Contact us:
C.Meo@tudelft.nl
V.Sarathchandran@tudelft.nl
A.Majhi@tudelft.nl

LatentWorlds AI

TU Delft Delft University of Technology

UNICA UNIVERSITÀ DEGLI STUDI DI CAGLIARI

Deltares