Journal of Oceanology and Limnology

https://doi.org/10.1007/s00343-026-6189-y

Research Paper

A dual-loop sliding-window framework for physically consistent precipitation nowcasting*

Pengda WANG[1], Nan YANG[2, **], Lingqi LI[1, **], Xiaofeng LI[2]

[1] *College of Physics Science, Center for Marine Observation and Communications, Qingdao University, Qingdao 266071, China*

[2] *Key Laboratory of Ocean Circulation and Waves, Institute of Oceanology, Chinese Academy of Sciences, Qingdao 266071, China*

Received Apr. 29, 2026; accepted in principle May 29, 2026; accepted for publication Jun. 30, 2026

© Chinese Society for Oceanology and Limnology, Science Press and Springer-Verlag GmbH Germany, part of Springer Nature 2026

Abstract Precipitation nowcasting is essential for meteorological early warning but remains challenging due to the highly nonlinear dynamics of atmospheric systems. Although recent advection-based deep learning models have improved forecasting skill, three key issues have not been sufficiently considered: 1) inaccurate intensity forecasts over longer lead times; 2) blurring of precipitation fields due to spatial misalignment; and 3) region-boundary artifacts introduced by interpolation, such as spatial striping patterns. To address these issues, we propose a Physically constrained Dual-loop Sliding-window framework (PDS), which consists of: 1) a dual-loop sliding-window architecture that can update precipitation inputs dynamically during prediction; 2) a variant Bottleneck-ConvLSTM that is used to maintain states across successive windows; and 3) a Physical Consistency Out-of-Bounds Suppression (PCOBS) mechanism for mitigating boundary artifacts. Specifically: i) the sliding-window strategy progressively updates the precipitation input used during prediction, thereby reducing long-term error accumulation and intensity underestimation; ii) the ConvLSTM module maintains cross-timestep states across successive windows, helping preserve coherent precipitation patterns and alleviating blurring as well as spatial misalignment; iii) by introducing a physically constrained mechanism, PCOBS suppresses boundary-induced artifacts caused by interpolation. Compared with some state-of-the-art models on the MeteoNet dataset, PDS achieves CSI improvements of 4.58%, 19.52%, and 268.74% at 0.1-, 1.0-, and 5.0-mm/h thresholds, respectively. Substantial gains are also observed in POD (up to 561.05%) and HSS (up to 216.82%) at higher intensity levels, indicating enhanced detection capability for heavy precipitation.

Keyword: precipitation nowcasting; deep learning; model analysis; advection-based modeling

1 INTRODUCTION

Precipitation nowcasting (0–2 h), which focuses on high-resolution forecasting of rainfall in the near term (Ravuri et al., 2021), plays a critical role in weather-dependent decision-making (Li et al., 2026), including emergency response, energy management, and flood early-warning systems (Seiler et al., 2002; Briones-Estébanez and Ebecken, 2017; Sowby and Capener, 2023). It relies heavily on high-resolution radar observations (Sun et al., 2014; Schmid et al., 2019), whose rapid advancement, together with satellite sensing (Kitzmiller et al., 2013; Funk et al., 2015; Levizzani et al., 2020), has substantially enriched precipitation datasets and enabled the development of data-driven nowcasting methods (Smith et al., 1994; Zhang et al., 2016; Yang et al., 2023). However, achieving both predictive accuracy and physical consistency remains challenging due to the intrinsically non-linear and nonstationary nature of atmospheric dynamics (Bowler et al., 2006; Sun et al., 2014).

* Supported by the National Natural Science Foundation of China (No. 42306214) and the Taishan Scholars Program (No. tsqn202312260)

Academic Papers of the 28th Annual Meeting of the China Association for Science and Technology

** Corresponding authors: yangnan@qdio.ac.cn; lingqilee@qdu.edu.cn

Traditional optical-flow-based approaches (Austin and Bellon, 1974), such as TREC (Rinehart and Garvey, 1978), interpret precipitation evolution as quasi-rigid advection under the assumption of morphological invariance. While computationally efficient, this formulation tends to break down when faced with complex structural transformations, leading to rapid performance degradation (Shi et al., 2015; Ayzel et al., 2020).

Deep learning methods have significantly advanced precipitation nowcasting by enabling data-driven modeling (Ma et al., 2023; Pan et al., 2024; Yang and Li, 2024a; Yang et al., 2024a, 2026; Wang et al., 2025). Early recurrent architectures, exemplified by ConvLSTM (Shi et al., 2015), have improved spatiotemporal dependency modeling by introducing memory mechanisms. Subsequent studies, including trajectory-aware models such as TrajGRU (Shi et al., 2017) and memory-enhanced frameworks like PredRNN and its variants (Wang et al., 2017, 2018), further improved the modeling of complex and long-range dependencies. More recently, advection-aware or physics-informed frameworks, such as NowcastNet (Zhang et al., 2023) and related approaches, have further improved performance. In parallel, Transformer-based models have also been introduced to better capture long-range spatial dependencies (e.g., MetNet; Sønderby et al., 2020; Espeholt et al., 2022). The integration of physical constraints with data-driven architectures provides a more flexible and expressive representation of precipitation dynamics.

Despite these advances, several key limitations remain: 1) most existing models rely on implicit temporal modeling and lack explicitly maintained cross-timestep state representations, a limitation also noted in recurrent-based nowcasting models (Wang et al., 2018). Without such explicit state propagation, temporal dependencies are only indirectly captured, making it difficult to preserve long-term dynamical constraints. Consequently, errors accumulate over time, leading to spatial misalignment and inaccurate intensity forecasts, while also degrading the preservation of coherent large-scale kinematic patterns over extended forecast horizons; 2) a second issue arises from the fixed dependency between motion field estimation and subsequent autoregressive advection. In many advection-based frameworks, including optical-flow-based methods (e.g., COTREC; Li et al., 1995) and recent deep learning approaches such as NowcastNet (Zhang et al., 2023), motion fields and intensity residuals are inferred from initial observations and then repeatedly reused during the entire forecasting process. This design effectively assumes that future motion and intensity evolution can be reliably extrapolated from early observations alone. However, such an assumption rarely holds in real meteorological systems, where precipitation evolves continuously and may deviate significantly from its initial trajectory. As the predicted states progressively drift away from the initial reference, the pre-estimated motion fields become increasingly inconsistent with the evolving dynamics, ultimately leading to spatiotemporal misalignment; 3) a third limitation is intrinsic to the resampling process, a well-recognized issue with differentiable resampling operators (e.g., spatial transformer-based sampling; Jaderberg et al., 2015). Continuous coordinate mapping, although convenient, inevitably introduces interpolation-induced boundary artifacts that lack physical consistency. These artifacts are often handled through heuristic cropping, which removes affected regions at the cost of discarding valid precipitation information. While functional in a limited sense, such remedies neither maintain full spatial coverage nor address the root cause of the inconsistency introduced by out-of-bound transformations.

To overcome the above challenges, we propose a Physically constrained Dual-loop Sliding-window framework (PDS) that couples an inner-loop physical recursion with an outer-loop re-estimation mechanism, enabling a progressive prediction strategy with iterative correction.

The main contributions of this work are summarized as follows:

1) We propose a dual-loop sliding-window update strategy that dynamically updates precipitation inputs during prediction, which mitigates error accumulation and improves intensity prediction at longer lead times.

2) We introduce a bottleneck-structured ConvLSTM module to maintain cross-timestep states across sliding windows, which preserves coherent precipitation patterns and reduces blurring and spatial misalignment.

3) We propose a Physical Consistency Out-of-Bounds Suppression (PCOBS) mechanism to suppress interpolation-induced boundary artifacts, improving the physical plausibility of predicted precipitation fields.

In summary, PDS tightly integrates physical evolution with dynamic feature perception. This synergy enhances long-term forecasting stability and

provides a more physically consistent and robust precipitation nowcasting solution.

2 RELATED WORK

2.1 Spatiotemporal predictive learning

Spatiotemporal predictive learning aims to infer future evolution from historical observations by modeling underlying spatial and temporal dependencies. In precipitation nowcasting, these approaches typically operate on radar echo sequences and learn data-driven representations of system dynamics to capture both structural evolution and motion patterns.

Early studies extended recurrent neural networks to spatial domains, with ConvLSTM (Shi et al., 2015) representing a seminal development. By embedding convolutional operations into the recurrent gating mechanism, ConvLSTM enables the propagation of spatially structured hidden states, thereby capturing coupled spatial-temporal dependencies.

Subsequent efforts have focused on enhancing temporal modeling capacity. PredRNN (Wang et al., 2017) introduces a spatiotemporal memory mechanism that enables cross-layer and cross-time state transitions, improving the modeling of long-term temporal dependencies. Its variants, such as PredRNN++ (Wang et al., 2018), further mitigate gradient vanishing issues through dedicated information flow pathways.

Despite these advances, recent studies have pointed out that existing models still struggle to accurately characterize fine-grained motion dynamics and high-frequency structures in precipitation fields. For instance, as observed in recent studies (Ravuri et al., 2021; Yang and Li, 2024b), many data-driven models tend to learn averaged motion patterns, resulting in motion mismatch, blurred structures, and inaccurate intensity estimation, especially for individual precipitation systems.

More broadly, RNN-based models suffer from error accumulation in long-term forecasting, CNN-based models lack temporal modeling capacity, and hybrid CNN-RNN frameworks often introduce additional complexity without fully resolving these limitations (Yang et al., 2024b).

2.2 Advection-driven nowcasting

To overcome the limitations in motion representation and temporal consistency observed in spatiotemporal predictive models, advection-driven approaches have emerged as a dominant paradigm (Rinehart and Garvey, 1978). These methods approximate atmospheric transport processes through estimated motion fields, which are then used to advect historical radar echoes. To account for non-rigid intensity variations, such as initiation and dissipation, residual components are often introduced, forming a unified framework that integrates motion-driven advection with intensity refinement.

Representative methods, such as NowcastNet (Zhang et al., 2023) and related advection-aware frameworks, jointly predict motion fields and intensity residuals from historical observations, enabling multi-step forecasting through iterative advection. This design offers clear advantages in structural preservation: as a spatial resampling process over existing echoes, advection naturally maintains sharp boundaries and fine-scale textures, often outperforming purely recurrent or implicit approaches in visual fidelity and the representation of intense precipitation (Zhang et al., 2023).

In typical advection-based nowcasting frameworks, motion fields and intensity residuals for multiple future steps are predicted in a single forward pass from historical inputs. During inference, these pre-estimated fields are applied in an autoregressive manner, where the latest predicted frame is iteratively advected to generate subsequent states. Although computationally efficient, this paradigm implicitly constrains future predictions to rely on motion fields inferred from a fixed historical context. In the absence of intermediate feedback, the pre-estimated motion remains biased toward early-stage dynamics, and the discrepancy progressively increases as precipitation systems evolve. As a result, spatial mismatches become more pronounced over time, especially when precipitation structures undergo large displacements, deformation, or approach domain boundaries, where motion fields fail to adapt accordingly.

These limitations reveal a fundamental challenge in advection-driven nowcasting: it is difficult to simultaneously maintain long-term temporal consistency while adapting to evolving local motion patterns when future motion is estimated only once from historical inputs. As a result, small errors in motion estimation can gradually accumulate over time, leading to spatial drift, distorted structures, and artifacts near domain boundaries, especially in long-range forecasts and under complex dynamical conditions. This highlights the need for nowcasting

frameworks that can continuously update motion representations while preserving structural coherence over extended lead times.

2.3 Boundary effect and limitation of advection operators

Advection operators serve as a mechanism linking motion estimation with structural evolution in spatiotemporal prediction. They are widely adopted in nowcasting and generative modeling, typically implemented through differentiable spatial sampling modules, such as those used in Spatial Transformer Networks (Jaderberg et al., 2015), or motion-field-driven resampling schemes (Zhang et al., 2023). By interpolating features over continuous spatial coordinates, these operators enable flexible spatial transformations while preserving structural continuity, playing a central role in precipitation nowcasting.

Despite their effectiveness, existing studies primarily emphasize motion estimation accuracy and generative capability, while the behavior of advection operators near domain boundaries remains insufficiently explored. In practice, when sampling locations approach or exceed image boundaries, standard implementations use predefined padding schemes (Jaderberg et al., 2015), such as zero-padding or border replication, to ensure stable computation. Although efficient, these strategies are heuristic and lack guarantees of numerical stability and physical consistency.

Several practical remedies have been explored to alleviate boundary artifacts. For instance, spatial cropping is often applied to retain only central regions, while domain padding is used to buffer boundary effects, as commonly adopted in advection-based frameworks. However, such approaches remain limited. Under iterative forecasting and long lead times, boundary-induced errors gradually accumulate and propagate through the autoregressive process, ultimately degrading overall prediction quality.

These limitations highlight a key gap in current advection-based frameworks: the lack of explicit constraints on out-of-bounds sampling at the operator level. From a physical perspective, advection should satisfy a basic requirement—mass transported beyond the spatial domain should not re-enter the system artificially. Enforcing such constraints remains challenging. Addressing this issue is therefore crucial for improving boundary robustness and developing physically consistent spatiotemporal forecasting models.

3 MATERIAL AND METHOD

3.1 Dataset and data preprocessing

3.1.1 MeteoNet dataset

This study uses the MeteoNet dataset released by Météo-France (Larvor et al., 2020), a high-quality, comprehensive meteorological benchmark. It provides long-term, high-resolution observations over northwestern and southeastern France. Since the primary task is spatiotemporal precipitation nowcasting, only radar reflectivity data (dBZ) are extracted, as they effectively capture the spatial distribution and intensity of precipitation systems. The input and output variables of the model are both dBZ. The loss function is computed in the normalized dBZ space. The thresholds of 0.1, 1.0, and 5.0 mm/h are adopted as widely used values representing low, medium, and high rainfall intensities. The Z-R relationship from the official MeteoNet documentation is used. The original radar data have a spatial resolution of 1 km and a temporal resolution of 5 min.

All experiments were conducted using the northwestern France subset of the MeteoNet radar reflectivity data. After preprocessing and sliding-window sequence construction, approximately 40 000 valid spatiotemporal samples were obtained. To ensure reproducibility and avoid information leakage between different evaluation stages, the samples were divided into training, validation, and test sets with a ratio of 70%, 20%, and 10%, respectively. Each sample consists of 24 consecutive radar frames at a 5-min interval, including 8 historical input frames and 16 future target frames.

The same data split, input-output sequence setting, preprocessing procedure, optimizer configuration, learning-rate schedule, and early-stopping criterion were used for all compared models whenever applicable.

3.1.2 Data preprocessing and normalization

Radar reflectivity data often contain missing values and anomalous echoes caused by ground clutter. To ensure data quality and facilitate model convergence, three preprocessing steps are applied: missing-value imputation, quality control via clipping, and min-max normalization.

Specifically, missing regions encoded with

predefined values (e.g., 255) are set to 0 dBZ as a proxy for no precipitation, thereby preserving spatial continuity. Reflectivity values are then constrained to the physically valid range of [0, 70] dBZ, and outliers are clipped to suppress abnormal gradients. Finally, the processed data are normalized to [0, 1] using linear scaling for stable network training.

After the above quality-control procedure, the cropped MeteoNet radar domain used in this study has an original spatial coverage of 565×784 grid points. To meet the input-size requirement of the neural network and reduce computational cost, the radar fields are resized to 128×128. We adopt nearest-neighbor interpolation for this resizing step. Compared with bilinear or area interpolation, nearest-neighbor interpolation avoids intensity averaging during resampling, thereby better preserving sharp echo boundaries and localized high-reflectivity cores. This is particularly important for threshold-based verification, because smoothing-based interpolation methods may weaken small-scale but intense convective regions and alter the distribution of heavy-precipitation pixels. Although other interpolation strategies were also considered, nearest-neighbor interpolation was finally adopted to reduce artificial smoothing of precipitation cores and to maintain compatibility with the discrete-grid characteristics of the subsequent threshold-based evaluation. It should also be noted that resizing inevitably introduces a certain degree of spatial aggregation, and very fine-scale structures below the effective grid scale may still be affected. Nevertheless, the use of nearest-neighbor interpolation helps retain localized high-reflectivity cores more explicitly than smoothing-based interpolation methods.

3.1.3 Spatiotemporal sequence construction

To meet the requirements of end-to-end sequence prediction models, the preprocessed spatial grids are resampled to a uniform 128×128 resolution using nearest-neighbor interpolation. This preserves the original structural boundaries of precipitation echoes while ensuring a consistent receptive field.

Along the temporal dimension, continuous radar streams are segmented into overlapping sequences via a sliding window. Given a 5-min interval, each sample consists of 24 consecutive frames (i.e., 2 h). Specifically, each sample (X, Y) contains an input sequence X of 8 historical frames (past 40 min) and a target sequence Y of 16 future frames (next 80 min).

3.2 Physically constrained Dual-Loop Sliding-Window Framework

In precipitation nowcasting tasks, given a historical radar echo sequence $X_{\text{in}}=\{X_{(t-T_{\text{in}}+1)}, \cdots, X_t\}$ of length T_{in}, the objective is to predict the future sequence $\hat{X}_{\text{out}}=\{\hat{X}_{(t+1)}, \cdots, \hat{X}_{(t+T_{\text{out}})}\}$ of length T_{out}. Conventional autoregressive nowcasting schemes, particularly those based on iterative advection, often suffer from spatial misalignment, error accumulation, and boundary-related artifacts as the prediction horizon increases. These issues are closely related to the limitations discussed in Section 1.

To address these challenges in a unified manner, we develop a framework that decomposes the nowcasting problem into three complementary components: long-term spatiotemporal dependency modeling, adaptive motion-intensity re-estimation, and physically consistent boundary handling. Each component is explicitly designed to target one of the aforementioned limitations.

Recognizing the critical need for dynamic motion adaptation, we propose a PDS, as illustrated in Fig.1. The proposed framework is built upon a baseline Evolution Network and extends it with a dual-loop prediction paradigm, where inner-loop autoregressive advection is coupled with outer-loop sliding-window re-estimation. Within this unified architecture, two dedicated submodules are further incorporated: a spatiotemporal memory module, implemented as a Bottleneck ConvLSTM embedded at the latent bottleneck, and a PCOBS mechanism for boundary-aware resampling. These components operate as integral parts of the overall framework, enhancing long-term dependency modeling and suppressing boundary-induced artifacts, respectively.

PDS dynamically unrolls the forecasting process by decomposing the total horizon T_{out} into a sequence of sliding windows of length K.

The pipeline initializes with a fixed-capacity input buffer (length T_{in}) that is continuously updated. This buffer maintains a sequence composed of recent ground-truth observations and model-generated predictions, ensuring that the model always uses the most recent local context for prediction. This design enables a coordinated dual-loop prediction strategy:

Outer loop (Deep Feature Re-estimation): At each outer loop iteration w (spanning time steps from t to $t+K$), the model re-estimates the motion field (optical flow) V_w and the intensity variation

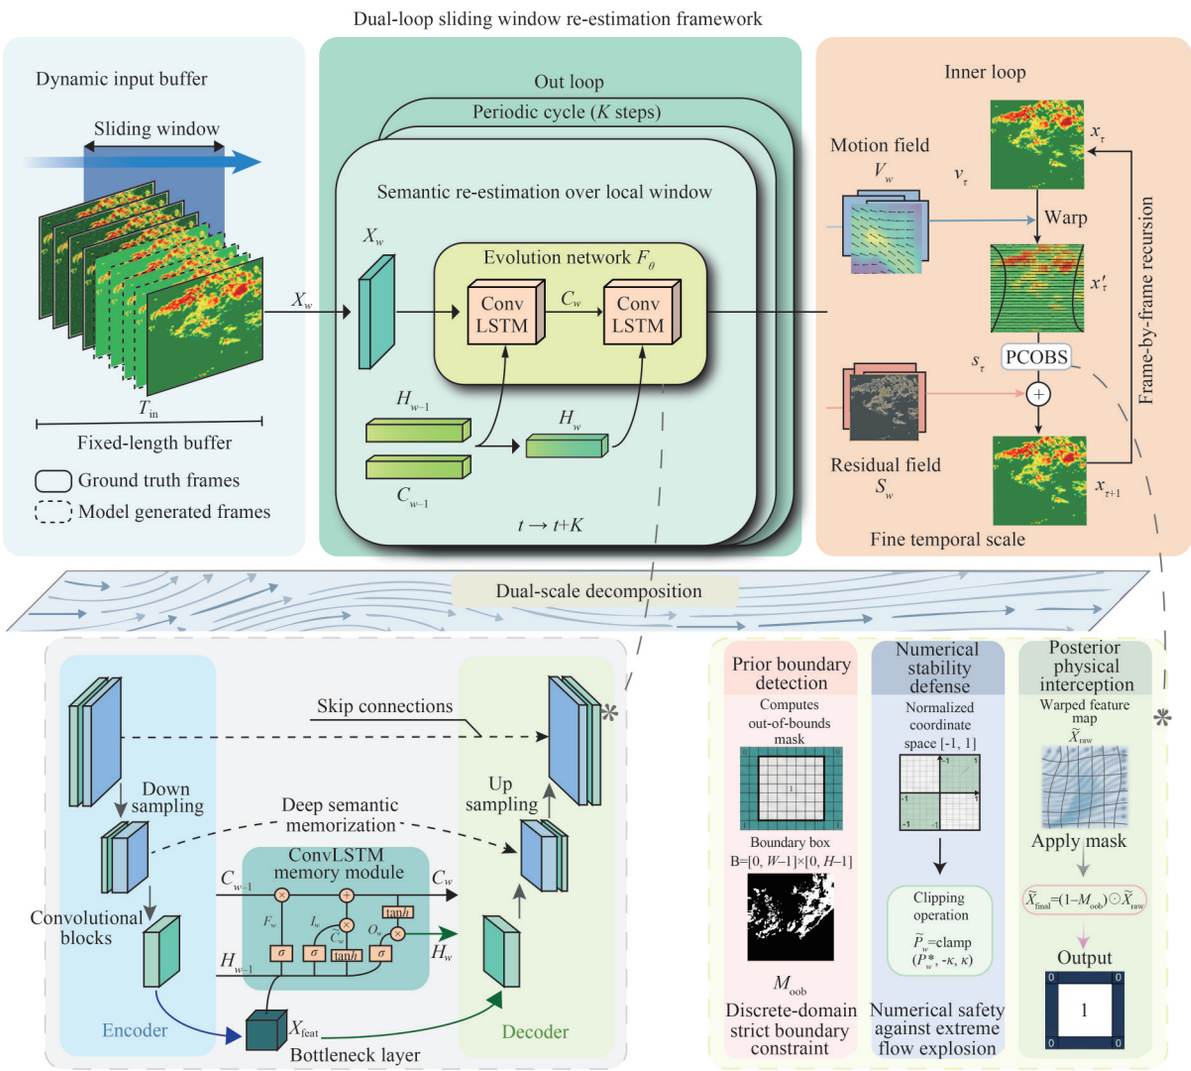


Fig.1 Overview of the proposed dual-loop sliding-window prediction framework

(residual) S_w for the localized window. This estimation is conditioned on the most recent T_{in} frames stored in the input buffer, alongside the hidden state tensors H_w and C_w from the Bottleneck ConvLSTM (detailed in Section 3.3), which encapsulate long-term evolutionary features. This process is formulated as:

$$(V_w, S_w, H_w, C_w) = F_\theta(X_w, H_{w-1}, C_{w-1}), \tag{1}$$

where F_θ denotes the memory-augmented Evolution Network parameterized by θ; w is the discrete index of the current outer sliding window; and X_w represents the recent feature sequence capturing temporal dependencies. This re-estimation mechanism enables the network backbone to dynamically update the flow field in response to spatiotemporal evolution, thereby improving its robustness against complex convection and non-rigid advection.

Inner loop (Physical Recurrence and Advection Simulation): Once the outer loop establishes the motion field and residual fields for the localized temporal window, the inner loop executes frame-by-frame update. By treating the state at the preceding frame $\hat{X}_\tau$ as the base state, the model spatially warps it using the corresponding predicted optical flow vector $v_\tau \in V_w$ and subsequently adds the predicted residual $s_\tau \in S_w$:

$$\hat{X}_{\tau+1} = \text{Warp}(\hat{X}_\tau, v_\tau) + s_\tau. \tag{2}$$

Formally, this operation can be interpreted as a discrete approximation of fluid advection. The motion-driven spatial transformation is implemented through motion-guided spatial resampling, with an additive residual term introduced to explicitly

account for intensity fluctuations. This stepwise advection scheme decomposes long-term dynamics into continuous short-term transport and residual-based intensity variation. Consequently, it preserves the structural consistency of spatial structures while mitigating shape distortion and spatiotemporal misalignment during long-horizon forecasting.

3.3 Deep bottleneck ConvLSTM for long-term spatiotemporal memorization

In the PDS, the outer loop performs re-estimation within local temporal windows. While this design improves adaptability, it also limits the effective temporal range that the model can capture. Without explicitly maintaining information across longer time steps, the model tends to rely mainly on spatial features, making it difficult to describe long-term processes such as the development and decay of large-scale convective systems.

To address this limitation, we introduce a ConvLSTM-based temporal modeling module (Shi et al., 2015) into the bottleneck layer of the U-Net backbone (Ronneberger et al., 2015), as illustrated in Fig.2. Instead of applying recurrent operations on high-resolution feature maps, we place the module at the deepest layer of the network. This design is motivated by the following considerations:

1) Large receptive field: The bottleneck layer covers the largest spatial context, allowing the model to capture large-scale precipitation patterns more effectively.

2) Reduced sensitivity to local noise: Performing temporal modeling at a high-level feature space helps reduce the influence of small-scale variations, enabling the model to focus on dominant structures.

3) Computational efficiency: Compared with applying recurrent operations at high-resolution layers, this design significantly reduces memory consumption and computational cost.

As the input observation tensor X_w descends through the encoder hierarchy, successive multi-level spatial downsampling significantly attenuates high-frequency sensor noise and localized micro-shifts. Arriving at the bottleneck layer, the aggregated deep feature tensor, denoted as X_{feat}, resides in a highly compact latent space, encoding large-scale precipitation patterns and their temporal evolution.

Positioned within this latent space, the ConvLSTM unit serves as a temporal connection

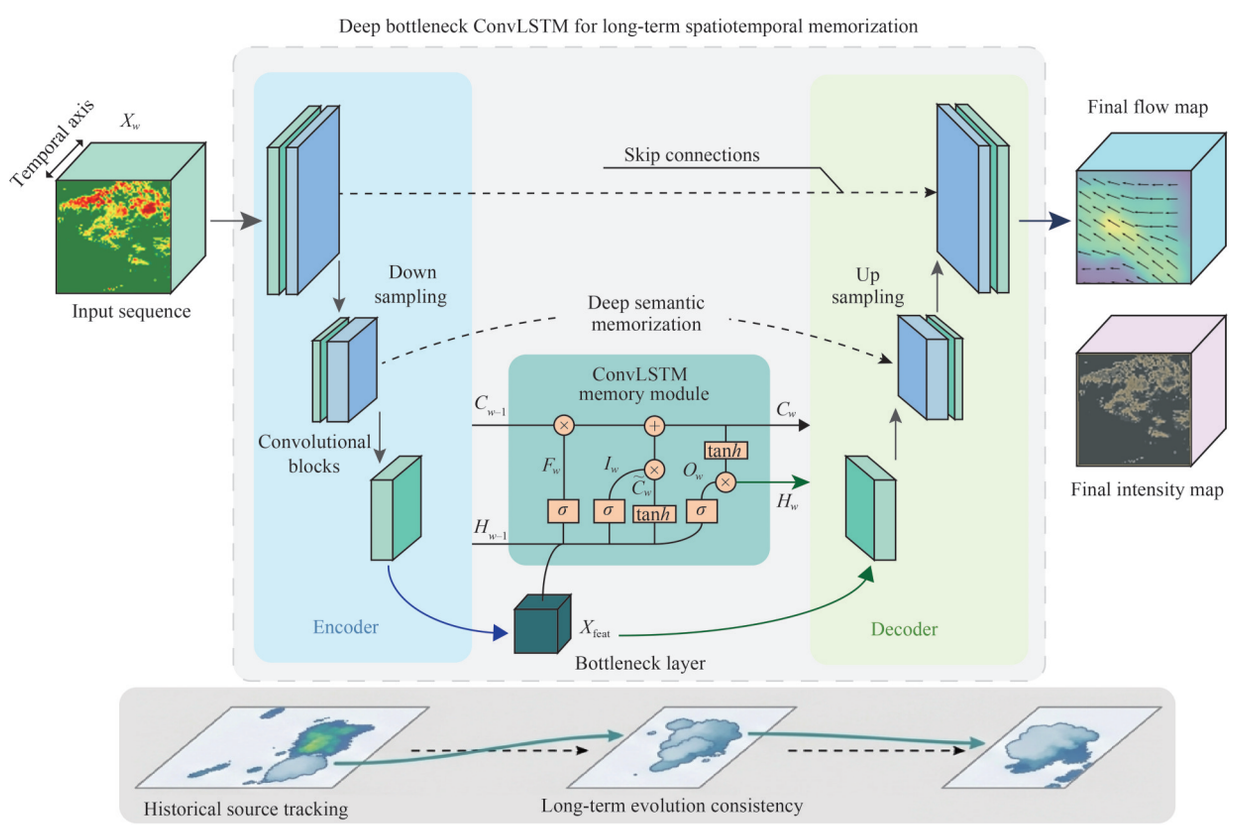


Fig.2 Architecture of the deep bottleneck ConvLSTM module within the U-Net backbone

across different time steps. Leveraging its internal gating mechanisms, it combines short-term changes with long-term information of the current window w with the continuous long-term memory spanning multiple historical windows:

$$\begin{aligned} I_w &= \sigma\left(W_{xi} * X_{\text{feat}} + W_{hi} * H_{(w-1)} + b_i\right) \\ F_w &= \sigma\left(W_{xf} * X_{\text{feat}} + W_{hf} * H_{(w-1)} + b_f\right) \\ \tilde{C}_w &= \tan h\left(W_{xc} * X_{\text{feat}} + W_{hc} * H_{(w-1)} + b_c\right) \\ C_w &= F_w \odot C_{(w-1)} + I_w \odot \tilde{C}_w \\ O_w &= \sigma\left(W_{xo} * X_{\text{feat}} + W_{ho} * H_{(w-1)} + b_o\right) \\ H_w &= O_w \odot \tan h\left(C_w\right) \end{aligned}. \tag{3}$$

The above equations describe the spatiotemporal memory mechanism established at the deep bottleneck. Here, I_w, F_w, O_w, and $\tilde{C}_w$ denote the input, forget, and output gates, along with the candidate cell state, respectively. $W_{(x\cdot)}$, $W_{(h\cdot)}$, and b represent the learnable convolutional kernel weight matrices and bias vectors inherent to each gating unit. The operator * indicates 2D convolution, $\odot$ denotes the Hadamard product (element-wise multiplication), and $\sigma(\cdot)$ identifies the non-linear sigmoid activation function, which constrains the soft-attention gate outputs into the (0, 1) interval. Notably, the core memory cell C_w and the hidden state H_w serve as a persistent latent representation of long-term dynamics, remaining largely invariant to the explicit spatial transformations induced by the inner-loop advection process. This property enhances the robustness of the memory against coordinate distortions introduced during iterative spatial resampling.

Following the latent-state update, the hidden state is directly propagated to the decoder, enriching the original feedforward representations with temporally accumulated context. This design introduces a key advantage: subsequent motion field estimation and residual prediction within the current outer-loop window are not only conditioned on instantaneous precipitation evolution (e.g., growth and dissipation), but are also guided by a stable latent kinematic space encoded in $H_{(w-1)}$. Even when precipitation structures undergo significant displacement or morphological variation during the inner-loop advection process, the deep memory representation preserves their historical context, enabling consistent tracking across time.

Ultimately, this deep integration ensures that fields inferred in the newly updated sliding coordinate frame maintain physical consistency with historical advection trends. By anchoring stable state propagation within a constrained latent space, the framework precisely captures the temporal evolution of complex precipitation systems.

3.4 Physical Consistency Out-of-Bounds Suppression (PCOBS)

In advection-based precipitation nowcasting, models rely on spatial resampling operators, such as differentiable grid sampling equipped with bilinear or nearest-neighbor interpolation, to simulate the advection processes of atmospheric fluids mathematically. Within deep autoregressive architectures, particularly those based on a dual-loop sliding window, standard resampling mechanisms can nonetheless introduce noticeable streaking and trailing artifacts. To mitigate such recurrent artifacts, we further introduce a lightweight, fully differentiable PCOBS mechanism as illustrated in Fig.3.

In a standard spatiotemporal mapping operator, assuming the absolute coordinate system of the target pixel grid is $\Omega \in R^{H\times W\times 2}$ and the model-predicted motion field is $V_w \in R^{H\times W\times 2}$, the source sampling coordinates can be formulated as $P_w = \Omega + V_w$. Theoretically, when the resampling boundary is handled using zero padding, any sampling coordinates exceeding the physical image boundary B=[0, W–1]×[0, H–1] should return a zero value. However, in practical implementations (e.g., GPU-based tensor libraries), due to floating-point truncation and grid alignment, sub-pixel coordinates that are minutely out of bounds are often not strictly classified as illegal. Instead, they trigger a forced boundary alignment, in which the sampling kernel is forcibly retracted and anchored to the nearest physical boundary pixel.

The error induced by such sub-pixel perturbations is negligible in a single forward mapping pass. However, within high-frequency autoregressive decoding iterations of the inner loop, edge pixels are repeatedly dragged inward into the interior field. As a result, this infinitesimal error gradually accumulates across both spatial and temporal dimensions. Over successive iterations, precipitation cloud systems that should have advected out of the radar field of view (FOV) fail to fully exit the domain. Instead, they result in elongated, streak-like artifacts near the image boundaries, which are inconsistent with physical

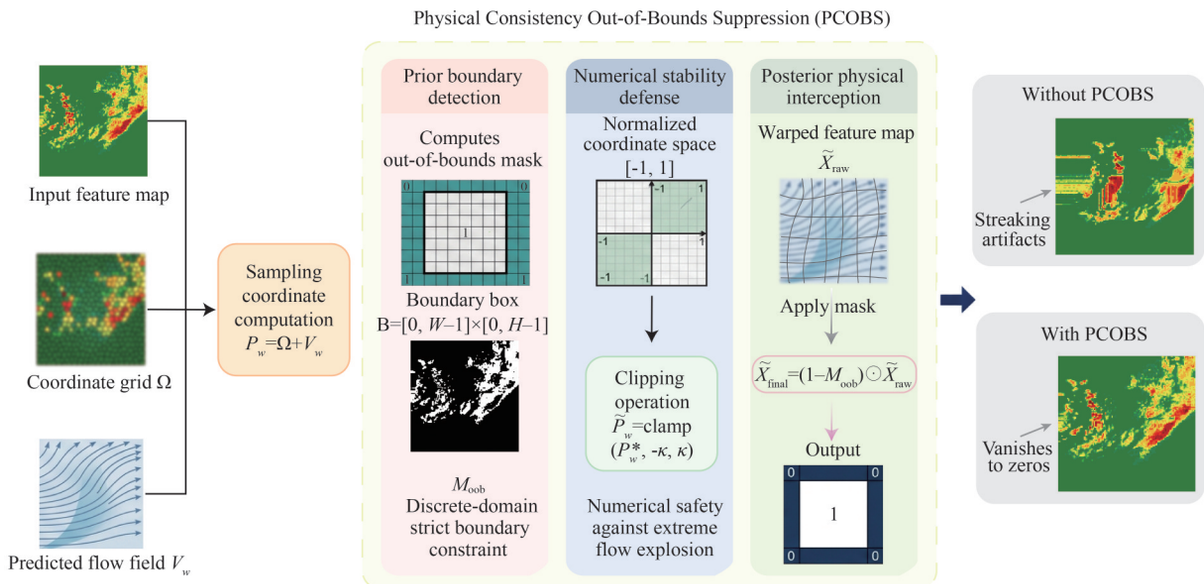


Fig.3 The proposed Physical Consistency Out-of-Bounds Suppression (PCOBS) mechanism

motion patterns. This not only disrupts the structural consistency of the flow field but also negatively affects the downstream outer-loop motion-intensity re-estimation network due to degraded observational priors.

To suppress such physical inconsistencies while preserving the computation graph of deep operators, PCOBS introduces a three-stage defensive strategy consisting of three steps: boundary detection, numerical stabilization, and output masking:

First, we construct a discrete spatial prior boundary mask. Before normalizing the sampling grid P_w into the [-1, 1] continuous coordinate space, the system performs an explicit out-of-bounds check directly on the high-precision original discrete pixel coordinate system. This yields a Boolean prior out-of-bounds mask $M_{oob}\in\{0, 1\}^{(H\times W)}$:

$$M_{oob}(x,y)=\begin{cases}1, & \text{if } P_w^{(x)}(x,y)\notin[0,W-1]\vee P_w^{(y)}(x,y)\notin[0,H-1]\\ 0, & \text{otherwise}\end{cases}. \tag{4}$$

The essence of this step is to establish rigidly defined physical constraints prior to the continuous spatial mapping, thereby mitigating precision ambiguity introduced by subsequent floating-point normalization.

Second, we implement a numerical stabilization defense. To enhance numerical robustness under extreme weather conditions or during early training, we introduce a truncation mechanism for numerical stability. This prevents ultra-large outlier coordinates from triggering kernel-level exceptions or gradient explosions:

$$\tilde{P}_w=\max(\min(P_w^*,\kappa),\ -\kappa), \tag{5}$$

where $\tilde{P}_w$ denotes the normalized coordinate field, and κ is the clipping threshold parameter. This defensive layer ensures numerical safety for the entire forward computation throughout non-linear iterations.

Finally, post-processing masking and consistency calibration are performed. After the differentiable resampling operator W completes the warping (as commonly used in optical flow-based motion estimation; Adil et al., 2024) process and yields the transitional field $\hat{X}_{raw}$, PCOBS activates the pre-established prior boundary mask M_{oob}. It applies conditional zeroing on all feature values designated as out-of-bounds, thereby accomplishing the final physical calibration at the system output terminal:

$$\hat{X}_{final}=(1-M_{oob})\odot W(\hat{X}_r,\tilde{P}_w). \tag{6}$$

The PCOBS mechanism enforces consistency with fluid advection laws under open boundary conditions. When a precipitation mass is driven outside the observational boundary by a divergent flow field, the proposed system ceases to artificially retain it via non-physical boundary-induced artifacts. Instead, the precipitation mass is assumed to leave the observable radar domain and is therefore not represented within the prediction field. Accordingly, PCOBS serves as a built-in mechanism during the forward pass. It alleviates boundary-

related errors in standard grid sampling. This mechanism requires no invasive modifications to the core network. By restricting the masking operation to invalid regions, it preserves gradient propagation within valid regions. As a result, it exhibits favorable scalability and training stability.

3.5 Training loss

The overall loss function in our model is based on the NowcastNet formulation to ensure a fair comparison. Specifically, we retain the same set of loss components, including motion reconstruction loss, flow smoothness regularization, content loss, adversarial loss (Goodfellow et al., 2020), and auxiliary regularization terms.

However, rather than directly adopting the original configuration, the weighting coefficients for individual loss terms are adjusted to account for differences in dataset characteristics and model architecture, thereby ensuring stable training dynamics and balanced optimization.

Importantly, by maintaining the loss formulation at the structural level while performing only dataset-aware coefficient calibration, we avoid introducing additional optimization benefits from loss engineering. This design isolates the effect of the proposed architectural modifications, ensuring that any observed performance gains are likely attributed to the enhanced network design.

For completeness, the adopted loss terms are summarized as follows.

3.5.1 Motion branch loss

The motion branch aims to predict the motion field by warping past frames toward future frames to compute reconstruction errors. The total motion loss L_{motion} consists of two components:

$$L_{\text{motion}} = \text{weighted } L_1 + \lambda_{\text{motion}} \cdot J_{\text{smooth}}, \tag{7}$$

where λ_{motion}=0.01. Since the model inputs and outputs are normalized to the range [0, 1] (approximately corresponding to a maximum rainfall intensity of 70 dBZ), the ground truth R_{real} is first rescaled to the physical precipitation rate ($R\approx$ target×70). Accordingly, a rainfall-dependent weighting function $w(R)$ is defined as:

$$w(R)=\begin{cases} 1.0, & R_{\text{real}} < 24 \text{ mm/h} \\ \min\left(\dfrac{R_{\text{real}}}{24.0},\ 6.0\right), & R_{\text{real}} \geq 24 \text{ mm/h} \end{cases}. \tag{8}$$

The weighted reconstruction loss is then formulated as:

$$\text{Weighted } L_1 = \frac{1}{T}\sum \left| P_{\text{warp}} - R_{\text{real}} \right| \cdot w(R_{\text{real}}). \tag{9}$$

To encourage large-scale physical consistency of the predicted flow field (u, v), particularly in regions of intense precipitation, a precipitation-weighted smoothness regularization is introduced. Spatial gradients along the x- and y-directions are extracted using Sobel operators (∇x, ∇y), and the flow magnitude is defined as:

$$\text{Mag}_{\text{flow}} = \sqrt{\left(\nabla_x u\right)^2 + \left(\nabla_y u\right)^2 + \left(\nabla_x v\right)^2 + \left(\nabla_y v\right)^2 + 10^{-8}}. \tag{10}$$

The corresponding smoothness loss is given by:

$$J_{\text{smooth}} = \text{mean}\left(\text{Mag}_{\text{flow}} \cdot \left(1.0 + R_{\text{real}}\right)\right). \tag{11}$$

3.5.2 Intensity branch loss

The intensity branch is responsible for refining predictions and generating realistic precipitation textures. In this stage, it is modeled as a residual generator. The overall loss is defined as:

$$L_{\text{intensity}} = J_{\text{content}} + \lambda_{\text{sparsity}} \cdot J_{\text{sparsity}} + \beta \cdot J_{\text{adv}} + \gamma \cdot J_{\text{pool}}. \tag{12}$$

The content loss measures the discrepancy between the refined prediction and the ground truth using the same weighted L_1 formulation:

$$J_{\text{content}} = \text{weighted } L_1\left(P_{\text{final}}, R_{\text{real}}\right), \tag{13}$$

where $\lambda_{\text{sparsity}}$=0.01, β=6, and γ=20. These coefficients were selected according to the relative magnitudes of the corresponding loss terms during preliminary validation experiments, with the objective of preventing any single auxiliary term from dominating the optimization while maintaining stable convergence. In particular, a small coefficient was assigned to the flow smoothness and sparsity terms because they serve as regularizers rather than primary reconstruction objectives, whereas larger weights were used for the adversarial and pooled consistency losses to encourage realistic precipitation structures and preserve large-scale rainfall patterns.

Even with generated fine-scale details, radar echoes should remain consistent with the large-scale structure of observations. Therefore, both the predicted output P_{final} and the ground truth R_{real} are downsampled using 2D max-pooling operator (kernel size=5, stride=2), followed by recomputation of the weighted L_1 loss:

$$J_{\text{pool}} = \text{weighted } L_1\left(\text{MaxPool 2d}\left(P_{\text{final}}\right), \text{MaxPool 2d}\left(R_{\text{real}}\right)\right). \tag{14}$$

The adversarial loss adopts the standard hinge loss formulation for GAN generators, aiming to fool the discriminator D:

$$J_{\text{adv}} = \text{-mean}\left(D\left(P_{\text{final}}\right)\right). \tag{15}$$

To prevent the intensity branch from generating spurious radar echoes, an additional L_1 penalty is imposed on the isolated residual component P_{isolated}:

$$J_{\text{sparsity}} = \text{mean}\left[\left|P_{\text{isolated}}\right|\right]. \tag{16}$$

3.6 Evaluation metrics

To comprehensively evaluate precipitation nowcasting performance, we adopt a multi-metric framework including the Critical Success Index (CSI), Heidke Skill Score (HSS), and Probability of Detection (POD), following standard meteorological verification practices (Sun et al., 2014). Predictions and ground truth from different models are binarized using multiple rainfall intensity thresholds. Let TP, FP, FN, and TN denote the number of true positives, false positives, false negatives, and true negatives, respectively. The evaluation metrics are defined as follows:

$$\begin{aligned}
\text{CSI} &= \frac{\text{TP}}{\text{TP}+\text{FP}+\text{FN}},\\
\text{POD} &= \frac{\text{TP}}{\text{TP}+\text{FN}},\\
\text{HSS} &= \frac{2(\text{TP}\times\text{TN}-\text{FP}\times\text{FN})}{(\text{TP}+\text{FN})(\text{FN}+\text{TN})+(\text{TP}+\text{FP})(\text{FP}+\text{TN})}.
\end{aligned} \tag{17}$$

Among these indicators, CSI is adopted as the primary metric to reflect overall prediction accuracy by jointly accounting for misses and false alarms. HSS evaluates the skill of the forecast relative to random chance, providing a more balanced assessment under class imbalance. POD measures the ability of the model to correctly detect precipitation events, with particular emphasis on reducing missed detections.

4 RESULT

4.1 Experimental setting

For the proposed PDS, the outer loop update interval is set to K=4, meaning that a latent feature and motion field re-estimation is performed every four forward prediction steps.

Within the PCOBS Configuration, the extreme threshold for physical out-of-bounds detection is set to κ=10^4. A fully differentiable mask-based interception strategy is employed to ensure stable gradient propagation during training.

The model is optimized using the AdamW optimizer with an initial learning rate of 2×10^{-4} and a Piecewise Constant Decay schedule. Training is conducted in parallel on NVIDIA A100 GPUs for up to 100 epochs, with early stopping to prevent overfitting.

4.2 Comparative experiment and performance analysis

4.2.1 Quantitative comparison across precipitation threshold

Quantitative comparisons across different precipitation intensity thresholds using identical training settings are summarized in Table 1. The evaluated models include representative approaches from different methodological paradigms, namely the recurrent-based model PredRNN (Wang et al., 2017), the physics-aware model PhyDNet (Le Guen and Thome, 2020), the advection-based baseline NowcastNet (Zhang et al., 2023), as well as AlphaPre (Lin et al., 2025) and a re-implemented baseline model for fair comparison.

All evaluated models exhibit a downward trend in performance as the precipitation intensity threshold increases. However, the proposed model demonstrates improved robustness and higher metric scores under medium-to-high rainfall conditions.

At the 5.0-mm/h threshold, the strongest competing model, AlphaPre, achieves a CSI of 0.138 7 and a POD of 0.165 9, whereas the proposed model achieves a CSI of 0.225 3 and a POD of 0.454 8. These results indicate that the proposed method mitigates the underestimation of peak rainfall, capturing intense precipitation events that are often missed by baseline models.

At the 1.0-mm/h threshold, the proposed model maintains competitive performance. It achieves a POD of 0.778 2, exceeding PredRNN (0.662 4) and PhyDNet (0.521 6), indicating improved sensitivity to developing convective structures under moderate rainfall conditions. Meanwhile, PredRNN remains competitive in terms of CSI, reflecting its overall strong performance in capturing precipitation patterns.

At the 0.1-mm/h threshold, PredRNN yields the highest scores across all metrics (e.g., a CSI of 0.737 5), demonstrating its effectiveness in representing widespread weak precipitation. Its

Table 1 Comparison of spatiotemporal forecasting performance across different precipitation thresholds

Method	0.1-mm/h CSI ↑	0.1-mm/h POD ↑	0.1-mm/h HSS ↑	1.0-mm/h CSI ↑	1.0-mm/h POD ↑	1.0-mm/h HSS ↑	5.0-mm/h CSI ↑	5.0-mm/h POD ↑	5.0-mm/h HSS ↑
AlphaPre	0.566 2	**0.889 0**	0.659 4	0.450 8	0.626 3	0.603 8	0.138 7	0.165 9	0.234 7
Baseline	0.646 7	0.706 8	0.731 8	0.464 1	0.560 2	0.617 2	0.061 1	0.068 8	0.114 1
PhyDNet	0.652 1	0.727 9	0.756 7	0.455 3	0.521 6	0.609 8	0.064 6	0.068 4	0.120 7
PredRNN	**0.737 5**	0.782 4	**0.816 8**	**0.580 2**	0.662 4	**0.721 7**	0.108 5	0.119 1	0.194 7
PDS	0.676 3	0.747 4	0.762 5	0.554 7	**0.778 2**	0.694 2	**0.225 3**	**0.454 8**	**0.361 5**

Bold values indicate the highest score for each metric at each precipitation threshold.

relatively higher scores at low thresholds suggest a tendency to emphasize broad precipitation coverage. In contrast, the proposed model achieves comparable performance—outperforming AlphaPre, the Baseline, and PhyDNet—while maintaining a more balanced representation. Rather than emphasizing broad coverage, it better preserves intense rainfall structures, contributing to its improved performance at the 5.0-mm/h threshold.

4.2.2 Continuous-threshold performance analysis

To further examine model behavior under varying precipitation intensities, a continuous-threshold comparison among the proposed model, PredRNN, and the baseline is provided in Fig.4, spanning from 0.1 to 10 mm/h.

At the low-intensity end, PredRNN achieves the highest CSI values, while the proposed model and the baseline exhibit slightly lower but comparable performance. As the threshold increases, the performance gap gradually narrows, and the proposed model becomes increasingly competitive. In the medium-to-high intensity range, the proposed model consistently surpasses both PredRNN and the baseline, while maintaining more stable CSI values.

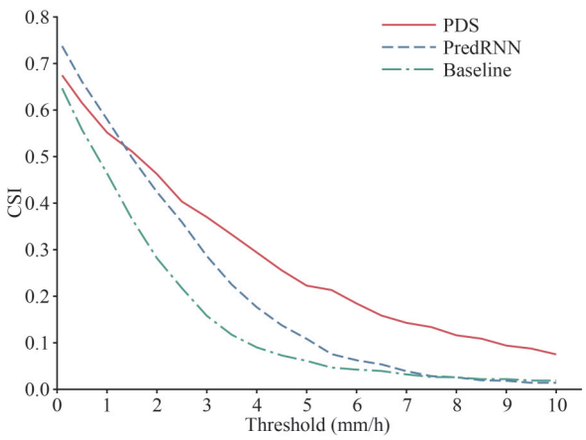


Fig.4 CSI comparison among the proposed model, PredRNN, and the baseline across different precipitation thresholds

This divergence becomes more pronounced at higher thresholds, where both PredRNN and the baseline exhibit faster performance degradation. In contrast, the proposed model shows a more gradual decline, indicating stronger robustness under strict evaluation criteria. The observed trend visually confirms the aforementioned trade-off: while PredRNN tends to favor widespread weak precipitation, PDS achieves competitive performance for weak precipitation and shows substantial advantages under moderate-to-heavy precipitation thresholds.

4.2.3 Temporal stability analysis across lead times

To investigate temporal stability across different precipitation intensities, Fig.5 illustrates the CSI evolution of all models across lead times under three precipitation thresholds (0.1, 1.0, and 5.0 mm/h), enabling a joint analysis of temporal stability and intensity-dependent performance.

A consistent decline in CSI is observed for all models as lead time increases, reflecting the accumulation of recursive prediction errors. However, the relative performance dynamics vary significantly across thresholds. At the low threshold (0.1 mm/h), PredRNN achieves slightly higher CSI values over most lead times. The proposed model attains comparable performance with only minor differences. Importantly, it achieves this without sacrificing its representational accuracy at higher intensities, further indicating its adaptability.

At the moderate threshold (1.0 mm/h), the proposed model maintains competitive CSI values throughout the prediction horizon and demonstrates stronger stability at later lead times. While PredRNN shows slightly higher scores at early stages, its performance degrades more rapidly, whereas the proposed model preserves more consistent medium-intensity precipitation structures as temporal errors accumulate.

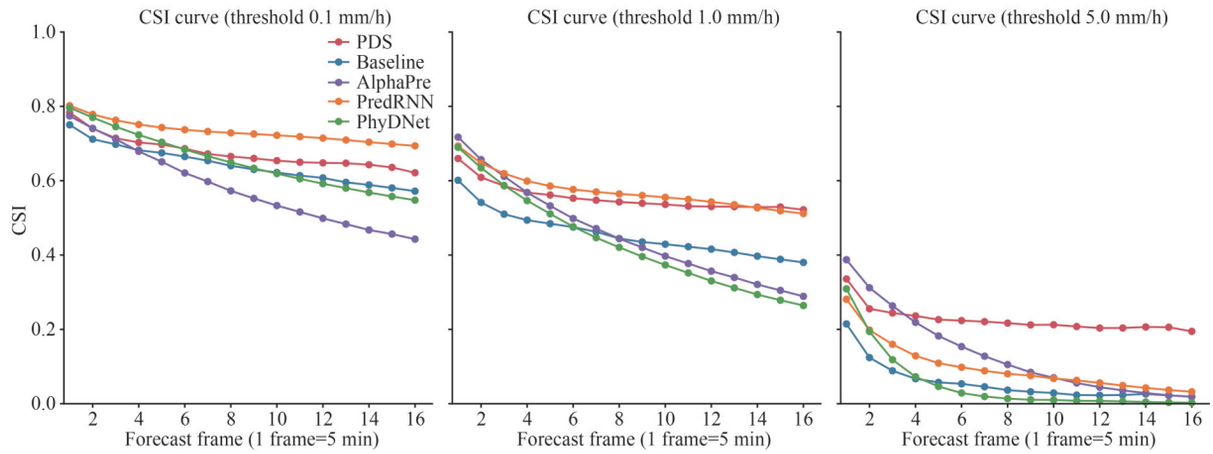


Fig.5 CSI curves of all models as a function of lead time

At the high threshold (5.0 mm/h), the proposed model consistently outperforms all baselines across all lead times, with a substantially slower decaying trend. In comparison, other methods, including PredRNN, show rapid performance decay and approach very low CSI values at extended lead times, highlighting their limitations in preserving heavy rainfall patterns.

These observations suggest that the proposed model achieves a more favorable balance between temporal stability and intensity sensitivity. While it does not prioritize weak precipitation to the same extent as PredRNN, it constrains error propagation and preserves high-intensity structures during iterative forecasting.

4.2.4 Qualitative comparison of forecasting results

Qualitative comparisons further support these findings. Figure 6 depicts that other models exhibit spatial blurring and structural misalignment as the forecast horizon extends, accompanied by a systematic underestimation of intense precipitation. This manifests as a progressive attenuation of high-intensity regions, leading to weakened convective cores and the loss of localized extreme features in later frames.

In contrast, the proposed method preserves sharper structural boundaries, more coherent high-intensity regions, and more consistent large-scale motion patterns. These improvements are consistent with the combined effects of the proposed components: the Bottleneck ConvLSTM enhances deep spatiotemporal memory and stabilizes feature evolution; the PDS mechanism alleviates long-horizon error accumulation; and the PCOBS suppresses physically implausible intensity artifacts and promotes more physically consistent precipitation distributions. Together, these mechanisms may contribute to reducing cumulative prediction errors and improving the retention of physically meaningful precipitation structures.

4.3 Ablation study

To rigorously evaluate the contribution of each component, we conducted ablation studies on the dual-loop mechanism, the Bottleneck ConvLSTM, and the PCOBS module.

4.3.1 Effect of the dual-loop sliding-window mechanism

To evaluate the dual-loop mechanism, we replace it with the standard NowcastNet configuration. As shown in Figs.7 & 8, the baseline model exhibits a faster decline in CSI as the lead time increases. In later prediction stages, it shows a significantly reduced capability to capture cloud systems with highly non-linear advection, accompanied by severe blurring and loss of structural coherence. In contrast, the proposed dual-loop sliding-window design achieves a better balance between stability and adaptability, resulting in improved long-range forecasting performance.

4.3.2 Effect of the Bottleneck ConvLSTM module

To investigate the role of the Bottleneck ConvLSTM, we remove it from the PDS framework, retaining only feedforward feature propagation. As illustrated in Figs.7 & 8, the model without ConvLSTM is able to capture advection-dominated translational motion but fails to model localized intensity variations.

This limitation becomes evident in regions with complex growth and decay processes, where the

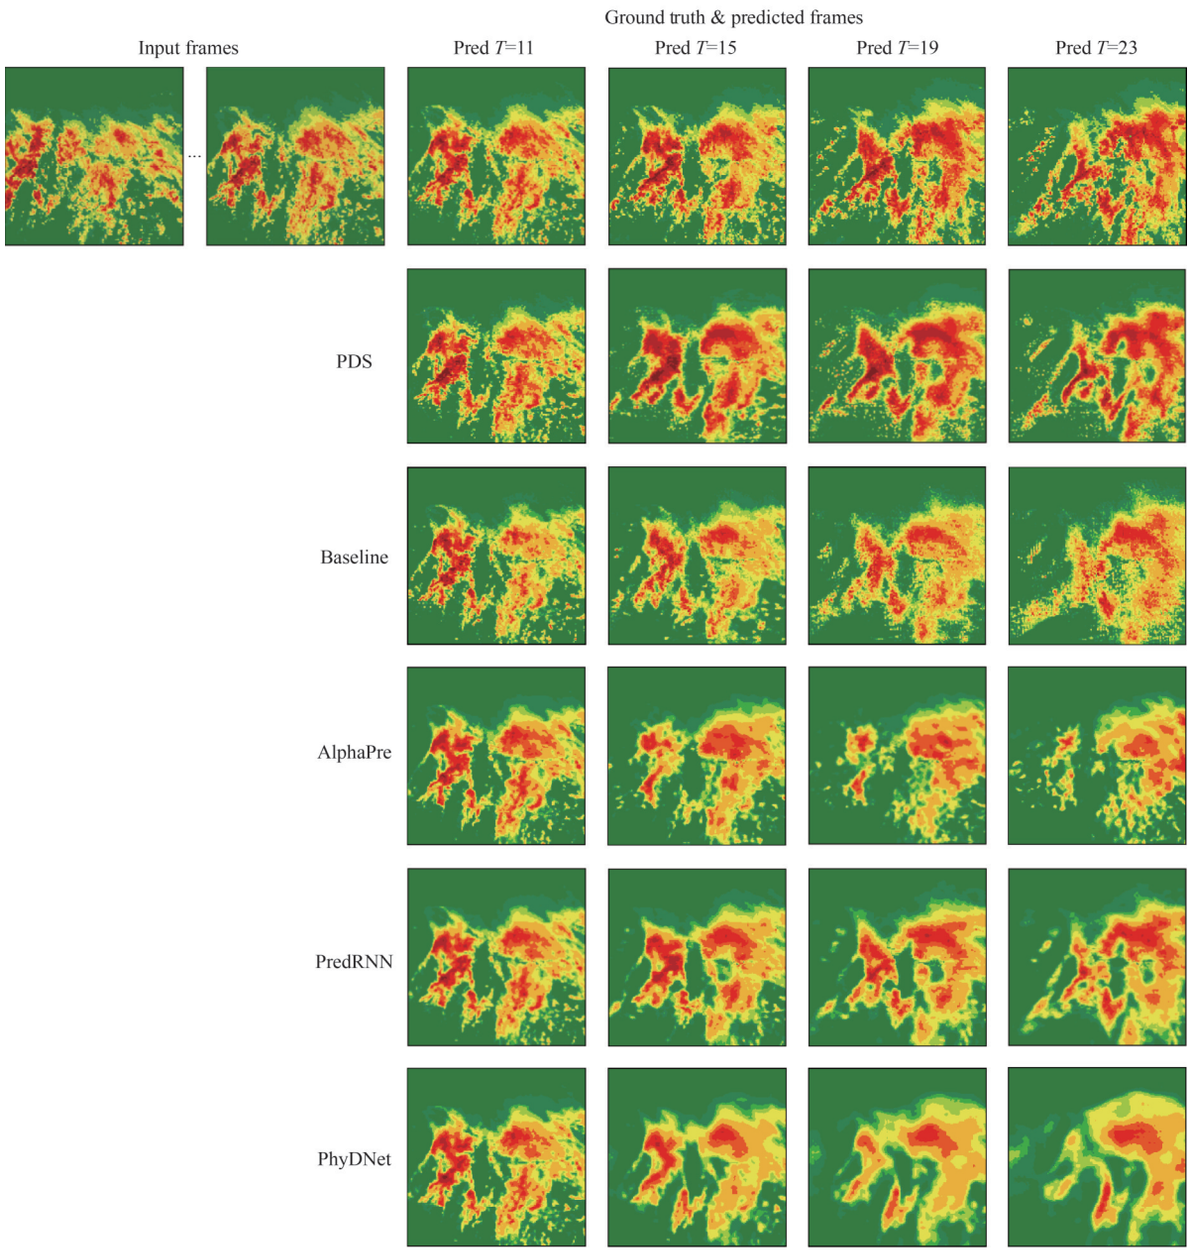


Fig.6 Performance comparison across different lead times

model produces weaker convective cores and degraded structural coherence. In some cases, its performance even falls below that of the baseline model. These effects are visually reflected by the attenuation of high-intensity regions and the loss of fine-scale structures, indicating reduced capability in representing strong precipitation dynamics.

4.3.3 Effect of the PCOBS mechanism

Effect of the PCOBS mechanism. Figure 9 compares the predictions with and without the proposed PCOBS module. Without PCOBS, standard differentiable grid sampling introduces sub-pixel interpolation artifacts near image boundaries, resulting in noticeable streaking and filament-like distortions. These artifacts not only degrade boundary quality but also propagate into subsequent time steps. By incorporating PCOBS, out-of-bound pixels are identified and suppressed prior to spatial resampling, thereby preventing invalid sampling coordinates from participating in interpolation. Consequently, the predicted frames exhibit more physically consistent boundary behavior, with reduced artifacts and improved structural integrity. Notably, this improvement is achieved without

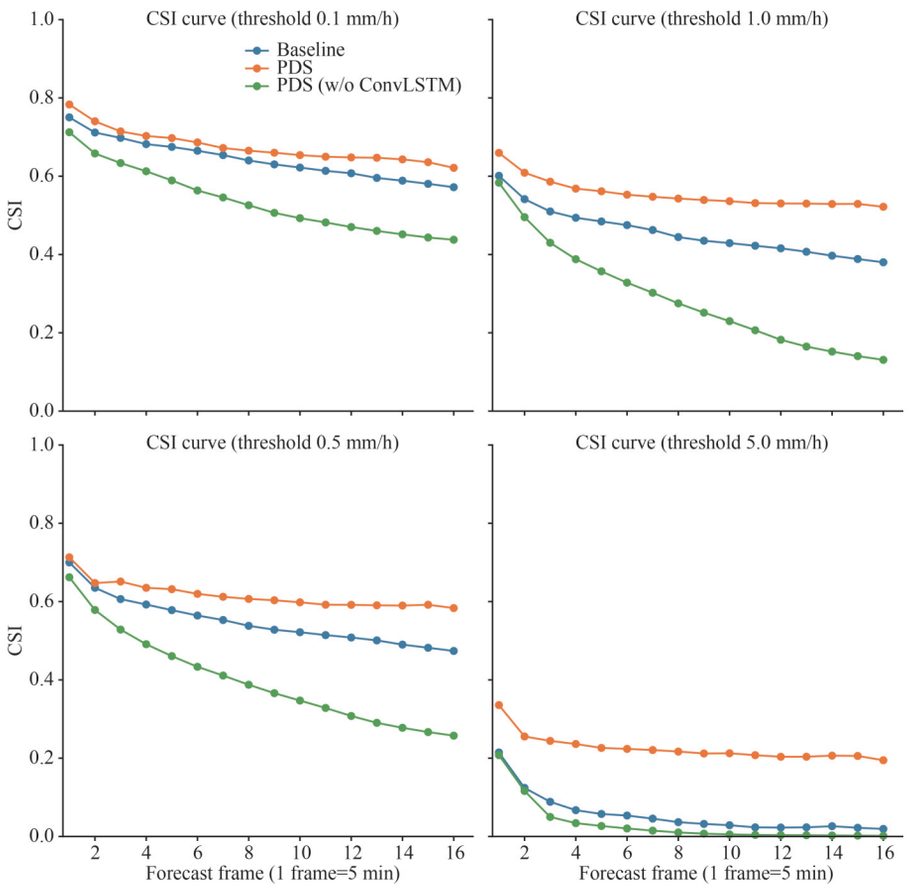


Fig.7 CSI evolution of different models across lead times under various precipitation thresholds

introducing additional learnable parameters, while reducing perceptual errors in boundary regions.

5 DISCUSSION AND CONCLUSION

The experimental results demonstrate that the proposed Physically constrained Dual-loop Sliding-window framework exhibits strong capability in addressing several critical limitations in advection-based precipitation nowcasting, including blurring, spatial misalignment, and boundary-induced artifacts, while improving the preservation of high-intensity precipitation patterns over extended lead times.

First, the dual-loop design avoids relying on early-stage observations for long-horizon prediction. Instead of depending on a fixed initial input, the outer loop dynamically updates the precipitation input used for motion and intensity estimation, thereby mitigating the accumulation of positional drift during inner-loop stepwise prediction and maintains consistency with the actual evolution of precipitation systems.

Second, the Bottleneck ConvLSTM module enhances temporal modeling by maintaining cross-timestep states across successive sliding windows. By performing spatiotemporal modeling in a compact latent space, the module stabilizes feature evolution and preserves coherent precipitation patterns, thereby reducing blurring and spatial misalignment at longer lead times.

Finally, the effectiveness of the PCOBS mechanism highlights the importance of explicitly handling boundary conditions during spatial resampling. By constraining out-of-bound values before sampling, PCOBS suppresses boundary-induced artifacts and prevents unrealistic reappearance of precipitation outside the observation domain, improving the physical plausibility of predicted fields without introducing additional model complexity.

Comprehensive experiments on the MeteoNet dataset demonstrate that the proposed framework achieves consistent improvements across multiple evaluation metrics, particularly in capturing high-intensity convective events while reducing structural degradation in later prediction stages. These results

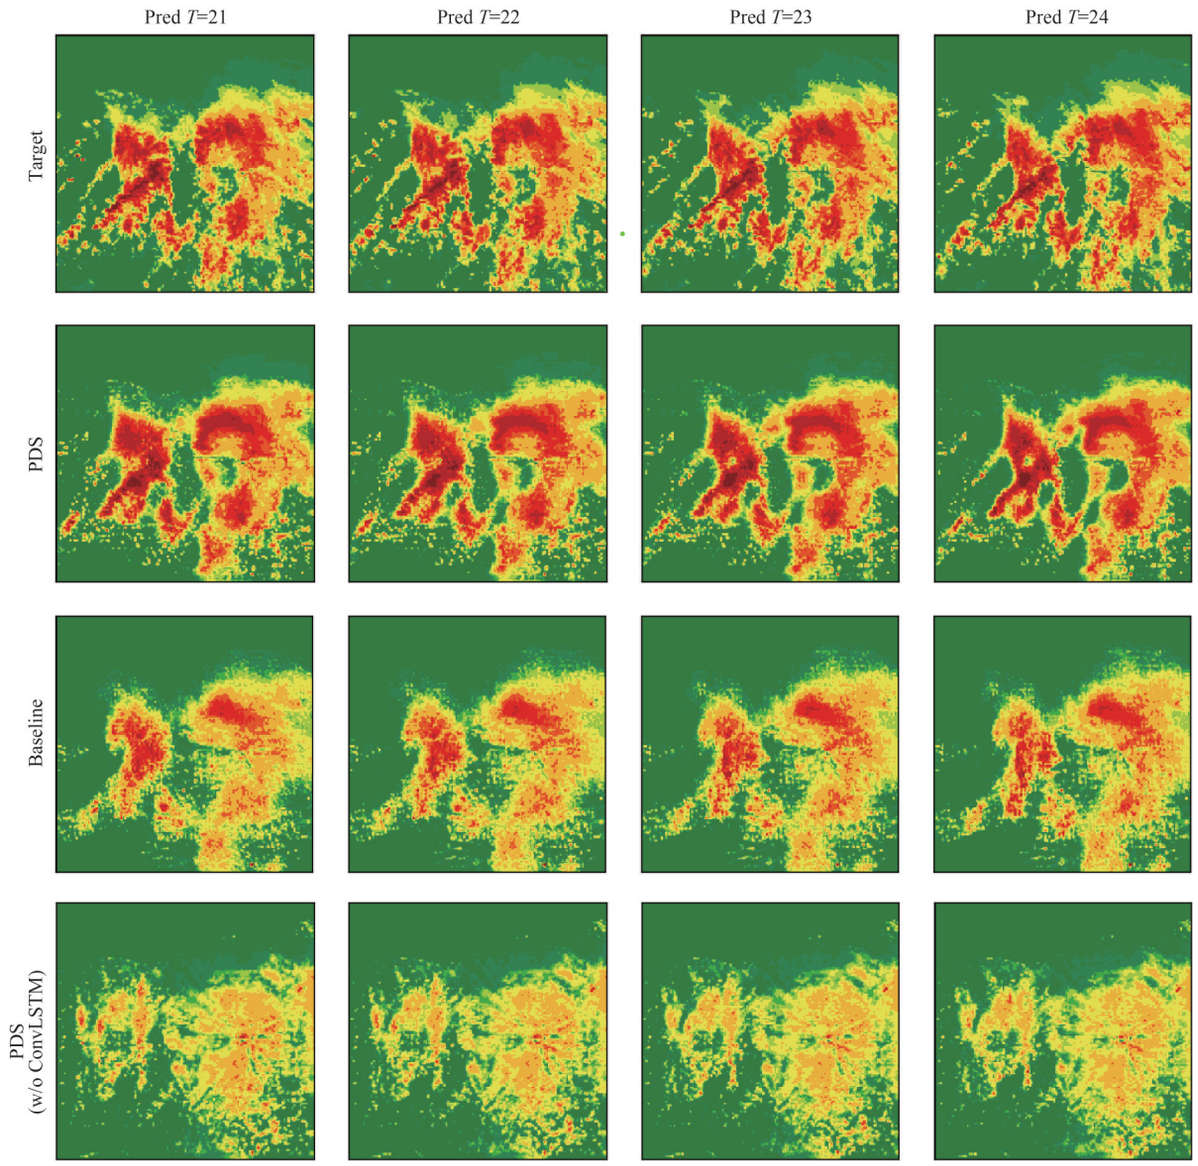


Fig.8 Visual comparison of late-stage prediction frames across different models

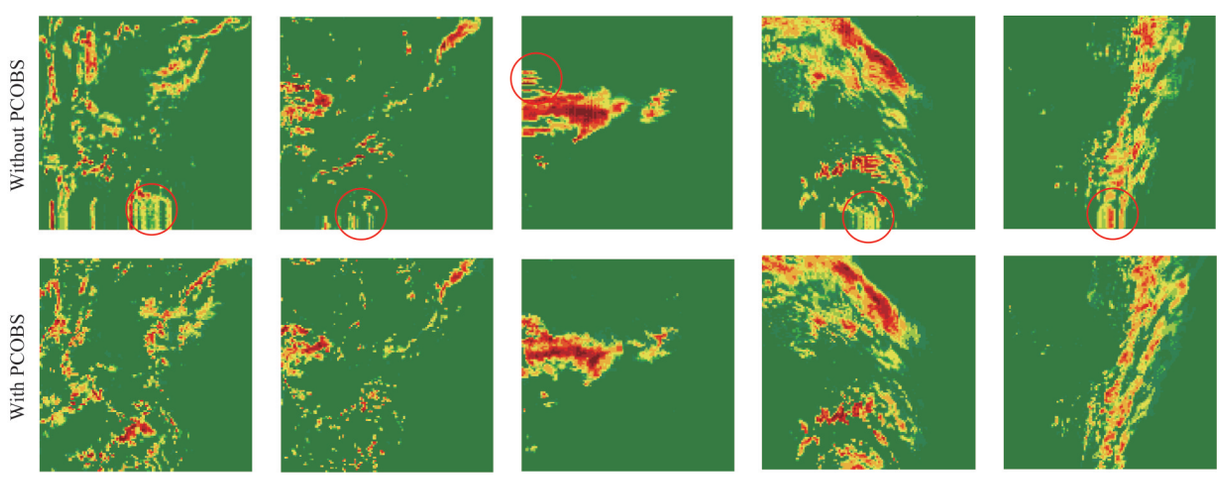


Fig.9 Visual comparison of predictions with and without the PCOBS mechanism

Red circles highlight boundary-induced streaking artifacts in the predictions without PCOBS.

indicate that the integration of physically constrained modeling with adaptive deep learning mechanisms leads to more robust and temporally consistent nowcasting performance.

However, the current framework still relies solely on radar-based inputs, which may limit its applicability under more complex meteorological conditions. Therefore, future work will focus on extending the proposed framework to a broader range of scenarios and incorporating multi-source atmospheric observations to further enhance its generalization capability across diverse spatial and temporal scales.

6 DATA AVAILABILITY STATEMENT

The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

References

Adil O, Mahraz M A, Riffi J et al. 2024. Advancements in deep learning-based optical flow estimation: a comprehensive review of models and techniques. *In*: 2024 Sixth International Conference on Intelligent Computing in Data Sciences (ICDS). IEEE, Marrakech, Morocco. p.1-7, https://doi.org/10.1109/ICDS62089.2024.10756420.

Austin G L, Bellon A. 1974. The use of digital weather radar records for short-term precipitation forecasting. *Quarterly Journal of the Royal Meteorological Society*, **100**(426): 658-664, https://doi.org/10.1002/qj.49710042612.

Ayzel G, Scheffer T, Heistermann M. 2020. RainNet v1.0: a convolutional neural network for radar-based precipitation nowcasting. *Geoscientific Model Development*, **13**(6): 2631-2644, https://doi.org/10.5194/gmd-13-2631-2020.

Bowler N E, Pierce C E, Seed A W. 2006. STEPS: a probabilistic precipitation forecasting scheme which merges an extrapolation nowcast with downscaled NWP. *Quarterly Journal of the Royal Meteorological Society*, **132**(620): 2127-2155, https://doi.org/10.1256/qj.04.100.

Briones-Estébanez K M, Ebecken N F F. 2017. Occurrence of emergencies and disaster analysis according to precipitation amount. *Natural Hazards*, **85**(3): 1437-1459, https://doi.org/10.1007/s11069-016-2635-z.

Espeholt L, Agrawal S, Sønderby C et al. 2022. Deep learning for twelve hour precipitation forecasts. *Nature Communications*, **13**(1): 5145, https://doi.org/10.1038/s41467-022-32483-x.

Funk C, Verdin A, Michaelsen J et al. 2015. A global satellite-assisted precipitation climatology. *Earth System Science Data*, **7**(2): 275-287, https://doi.org/10.5194/essd-7-275-2015.

Goodfellow I, Pouget-Abadie J, Mirza M et al. 2020. Generative adversarial networks. *Communications of the ACM*, **63**(11): 139-144, https://doi.org/10.1145/3422622.

Jaderberg M, Simonyan K, Zisserman A et al. 2015. Spatial transformer networks. *In*: Advances in Neural Information Processing Systems 28. Curran Associates, Inc.: 2017-2025. Accessed 7 September 2026, https://doi.org/10.5555/2969442.2969465.

Kitzmiller D, Miller D, Fulton R et al. 2013. Radar and multisensor precipitation estimation techniques in national weather service hydrologic operations. *Journal of Hydrologic Engineering*, **18**(2): 133-142, https://doi.org/10.1061/(ASCE)HE.1943-5584.0000523.

Larvor G, Berthomier L, Chabot V et al. 2020. MeteoNet, an open reference weather dataset by METEO FRANCE. Dataset. METEO FRANCE, https://meteonet.umr-cnrm.fr/. Accessed 7 September 2026.

Le Guen V, Thome N. 2020. Disentangling physical dynamics from unknown factors for unsupervised video prediction. *In*: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, Seattle, USA, p.11471-11481, https://doi.org/10.1109/CVPR42600.2020.01149.

Levizzani V, Kidd C, Kirschbaum D B et al. 2020. Satellite Precipitation Measurement. Springer, Cham, https://doi.org/10.1007/978-3-030-24568-9.

Li L, Schmid W, Joss J. 1995. Nowcasting of motion and growth of precipitation with radar over a complex orography. *Journal of Applied Meteorology and Climatology*, **34**(6): 1286-1300, https://doi.org/10.1175/1520-0450(1995)034<1286:NOMAGO>2.0.CO;2.

Li X F, Zhang X D, Ren Y B et al. 2026. Artificial intelligence for ocean phenomena forecasting. *Science Bulletin*, **71**(10): 2632-2652, https://doi.org/10.1016/j.scib.2026.04.002.

Lin K H, Zhang B Q, Yu D M et al. 2025. AlphaPre: amplitude-phase disentanglement model for precipitation nowcasting. *In*: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, Nashville, USA, p.17841-17850, https://doi.org/10.1109/CVPR52734.2025.01662.

Ma Z F, Zhang H, Liu J. 2023. MM-RNN: a multimodal RNN for precipitation nowcasting. *IEEE Transactions on Geoscience and Remote Sensing*, **61**: 4101914, https://doi.org/10.1109/TGRS.2023.3264545.

Pan Z F, Hang R L, Liu Q S et al. 2024. A short-long term sequence learning network for precipitation nowcasting. *IEEE Transactions on Geoscience and Remote Sensing*, **62**: 4106814, https://doi.org/10.1109/TGRS.2024.3424250.

Ravuri S, Lenc K, Willson M et al. 2021. Skilful precipitation nowcasting using deep generative models of radar. *Nature*, **597**(7878): 672-677, https://doi.org/10.1038/s41586-021-03854-z.

Rinehart R E, Garvey E T. 1978. Three-dimensional storm motion detection by conventional weather radar. *Nature*, **273**(5660): 287-289, https://doi.org/10.1038/273287a0.

Ronneberger O, Fischer P, Brox T. 2015. U-net: convolutional networks for biomedical image segmentation. *In*: 18th International Conference Medical Image Computing and Computer-Assisted Intervention-MICCAI 2015. Springer, Munich, Germany, p.234-241, https://doi.org/10.1007/

978-3-319-24574-4_28.

Schmid F, Wang Y, Harou A. 2019. Nowcasting guidelinesA summary. *WMO Bulletin*, **68**(2): 63-68, https://wmo.int/media/magazine-article/nowcasting-guidelines-summary. Accessed 7 September 2026.

Seiler R A, Hayes M, Bressan L. 2002. Using the standardized precipitation index for flood risk monitoring. *International Journal of Climatology*, **22**(11): 1365-1376, https://doi.org/10.1002/joc.799.

Shi X J, Chen Z R, Wang H et al. 2015. Convolutional LSTM network: a machine learning approach for precipitation nowcasting. *In*: Advances in Neural Information Processing Systems 28. Curran Associates, Inc., Montreal, Canada, p.802-810, https://doi.org/10.48550/arXiv.1506.04214.

Shi X J, Gao Z H, Lausen L et al. 2017. Deep learning for precipitation nowcasting: a benchmark and a new model. *In*: Proceedings of the 31st International Conference on Neural Information Processing Systems. Curran Associates Inc., Long Beach, USA. p.5617-5627, https://doi.org/10.48550/arXiv.1706.03458.

Smith E A, Kummerow C, Mugnai A. 1994. The emergence of inversion-type profile algorithms for estimation of precipitation from satellite passive microwave measurements. *Remote Sensing Reviews*, **11**(1-4): 211-242, https://doi.org/10.1080/02757259409532266.

Sønderby C K, Espeholt L, Heek J et al. 2020. Metnet: a neural weather model for precipitation forecasting. arXiv preprint arXiv: 2003.12140, https://doi.org/10.48550/arXiv.2003.12140. Accessed 7 September 2026.

Sowby R B, Capener A. 2023. The influence of precipitation on the energy footprint of Denver's water supply: a 20-year analysis and implications for climate change. *Energy Nexus*, **9**: 100166, https://doi.org/10.1016/j.nexus.2022.100166.

Sun J Z, Xue M, Wilson J W et al. 2014. Use of NWP for nowcasting convective precipitation: recent progress and challenges. *Bulletin of the American Meteorological Society*, **95**(3): 409-426, https://doi.org/10.1175/BAMS-D-11-00263.1.

Wang C, Yang N, Li X F. 2025. Advancing forecasting capabilities: a contrastive learning model for forecasting tropical cyclone rapid intensification. *Proceedings of the National Academy of Sciences of the United States of America*, **122**(4): e2415501122, https://doi.org/10.1073/pnas.2415501122.

Wang Y B, Gao Z F, Long M S et al. 2018. Predrnn++: towards a resolution of the deep-in-time dilemma in spatiotemporal predictive learning. *In*: Proceedings of the 35th International Conference on Machine Learning. PMLR, Stockholm, Sweden. p.5123-5132, https://doi.org/10.48550/arXiv.1804.06300.

Wang Y B, Long M S, Wang J M et al. 2017. PredRNN: recurrent neural networks for predictive learning using spatiotemporal LSTMs. *In*: Proceedings of the 31st International Conference on Neural Information Processing Systems. Curran Associates Inc., Long Beach, USA, p.879-888.

Yang N, Li X F. 2024a. Lightweight AI-powered precipitation nowcasting. *The Innovation Geoscience*, **2**(2): 100066, https://doi.org/10.59717/j.xinn-geo.2024.100066.

Yang N, Li X F. 2024b. An individual motion-driven artificial intelligence method for precipitation forecasting using radar image sequencing. *IEEE Transactions on Geoscience and Remote Sens*ing, **62**: 4208516, https://doi.org/10.1109/TGRS.2024.3439871.

Yang N, Wang C, Li X F. 2024a. Improving tropical cyclone precipitation forecasting with deep learning and satellite image sequencing. *Journal of Geophysical Research*: *Machine Learning and Computation*, **1**(2): e2024JH000175, https://doi.org/10.1029/2024JH000175.

Yang N, Wang C, Li X F. 2024b. Evaluation of precipitation forecasting methods and an advanced lightweight model. *Environmental Research Letters*, **19**(9): 094006, https://doi.org/10.1088/1748-9326/ad661f.

Yang R K, Jiao S J, Yang N. 2026. Deep neural network based on adversarial training for short-term high-resolution precipitation nowcasting from radar echo images. *Journal of Oceanology and Limnology*, **44**(1): 85-98, https://doi.org/10.1007/s00343-025-4354-3.

Yang W S, Chen H N, Han L et al. 2023. Multitask learning for precipitation estimation using satellite observations from the GOES-R series. *IEEE Geoscience and Remote Sens*ing *Letters*, **20**: 1002205, https://doi.org/10.1109/LGRS.2023.3313150.

Zhang J, Howard K, Langston C et al. 2016. Multi-radar multi-sensor (MRMS) quantitative precipitation estimation: initial operating capabilities. *Bulletin of the American Meteorological Society*, **97**(4): 621-638, https://doi.org/10.1175/BAMS-D-14-00174.1.

Zhang Y C, Long M S, Chen K Y et al. 2023. Skilful nowcasting of extreme precipitation with NowcastNet. *Nature*, **619**(7970): 526-532, https://doi.org/10.1038/s41586-023-06184-4.