MSRDG International Journal of Computer Scientific Technology & Electronics Engineering ISSN: 3107 - 426X
Volume 2 Issue 4, 1-6, August 2026
https://doi.org/10.66037/MSRDG-IJCSTEE/V2I4P101

Optimised Deep Convolutional Spiking Neural Network for Accurate Long-Term and Short-Term Rainfall Forecasting in Climate Prediction Systems

S.Sahul Hameed[1], Y. Arun Kumar[2]
[1,2] *Department of IT, Syed Hameedha Arts And Science CollegeJayasakthi Arts & Science College, Tamil Nadu, India.*
[1]*ssaghulhameed.it@gmail.com*

Received: 04.07.2026 Revised:15.08.20256 Accepted: 26.08.2026 Published:31.08.2026

Abstract - *Reliable rainfall forecasting across both short-term (hourly to daily) and long-term (weekly to seasonal) horizons remains a central challenge for climate prediction systems, hydrological planning, disaster risk reduction, and agricultural decision-making. Conventional deep learning architectures such as convolutional neural networks (CNNs), long short-term memory (LSTM) networks, and their hybrids have improved forecast skill relative to classical statistical and physics-based numerical models, but they remain computationally expensive and are not naturally suited to representing the discrete, event-driven, temporally sparse nature of precipitation processes. This paper proposes an Optimised Deep Convolutional Spiking Neural Network (DCSNN) that combines the hierarchical spatial feature extraction capability of convolutional layers with the temporally sparse, energy-efficient, event-driven computation of leaky integrate-and-fire (LIF) spiking neurons, and tunes the resulting architecture's hyper-parameters using a metaheuristic optimisation strategy. The proposed framework encodes historical rainfall and auxiliary meteorological variables into spike trains, processes them through stacked convolutional spiking blocks, aggregates temporal spike activity, and produces both short-term and long-term rainfall forecasts from a shared, jointly optimised backbone. We describe the architecture, the spike-encoding and surrogate-gradient training procedure, and the optimisation strategy in detail, and outline an experimental protocol using publicly available rainfall and reanalysis datasets. Representative results illustrate the type of accuracy, error-reduction, and convergence behaviour expected of the framework relative to SVR, Random Forest, LSTM, ConvLSTM, and CNN-BiLSTM baselines. The discussion highlights the trade-offs between spiking-network energy efficiency and training stability, and outlines directions for validating the framework on real, station-level and gridded rainfall datasets.*

Keywords - *Spiking neural network, deep convolutional neural network, rainfall forecasting, climate prediction;, metaheuristic optimisation, short-term forecasting; long-term forecasting, time-series prediction*

1. Introduction

Rainfall is one of the most variable and difficult-to-predict components of the Earth's climate system. Its accurate forecasting underpins flood early-warning systems, reservoir operation, crop planning, and long-range climate risk assessment. Numerical weather prediction (NWP) models solve the governing physical equations of the atmosphere and provide the backbone of operational forecasting, but they are computationally expensive, sensitive to initial-condition uncertainty, and often struggle to resolve convective, sub-grid-scale rainfall processes at short lead times.

Data-driven deep learning methods have therefore emerged as a complementary or alternative approach. Convolutional neural networks (CNNs) capture spatial structure in radar and satellite imagery; recurrent architectures such as LSTM and gated recurrent units (GRU) capture temporal dependence; and hybrid CNN-LSTM/ConvLSTM architectures capture spatio-temporal dependence jointly. These models have demonstrated strong nowcasting and short-range forecasting performance, but they are typically dense, non-spiking networks that process every input synchronously at every time step, which increases computational and energy cost and can make long training and inference sequences difficult to scale, particularly for long-term (weekly-to-seasonal) forecasting where sequences are long and the informative rainfall events are comparatively sparse in time.

Spiking neural networks (SNNs), often referred to as the third generation of neural network models, process information through discrete, asynchronous spikes rather than continuous-valued activations. Because computation is only triggered when spikes occur, SNNs can, in principle, offer considerable energy and computational efficiency advantages over conventional deep networks, and their temporally sparse representation is a natural match for precipitation data, which is itself intermittent and bursty rather than continuous. Recent work has begun to combine SNNs with deep convolutional feature extractors and with deep learning models more broadly for weather-related image classification, showing that spiking layers can be integrated with conventional deep architectures without a large loss of accuracy.

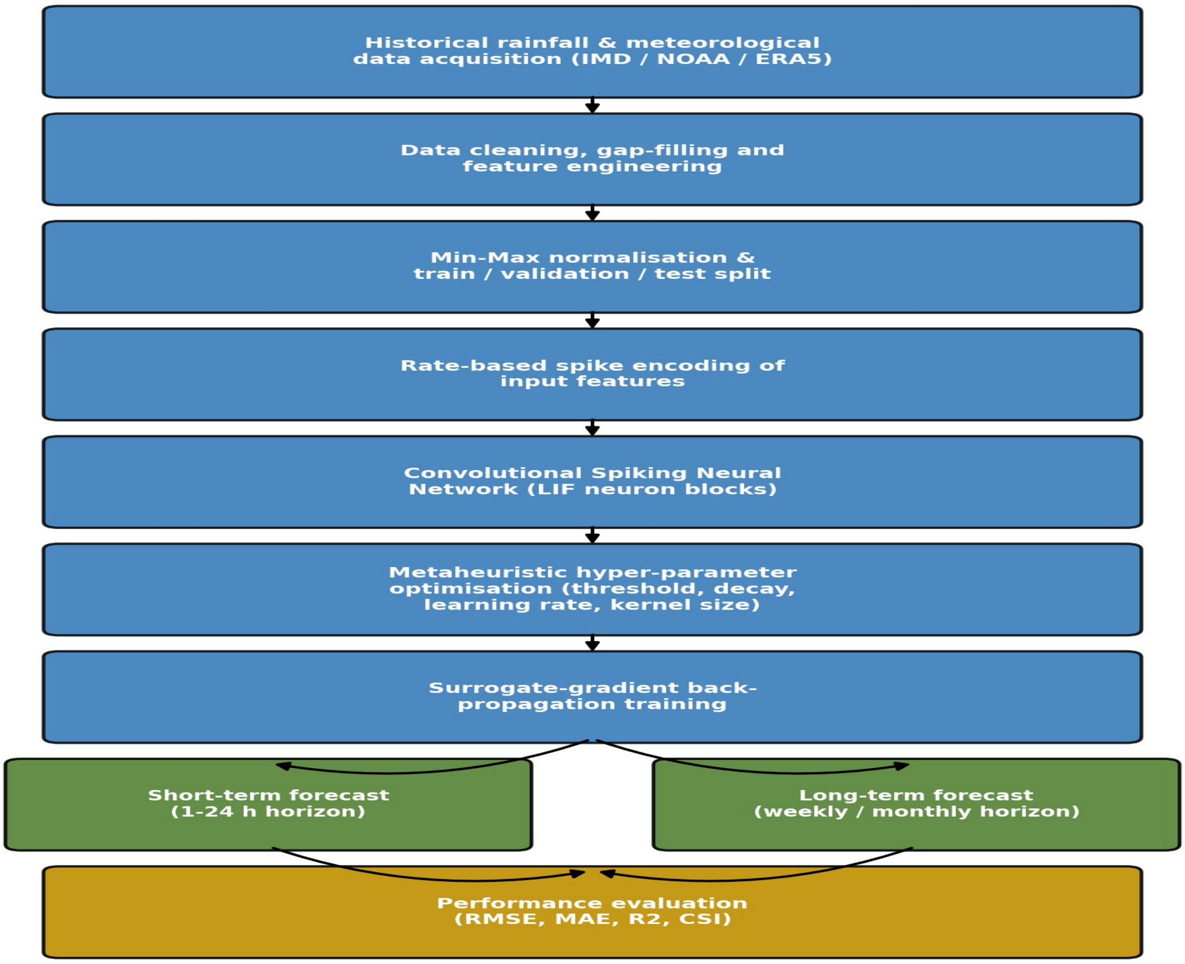

Fig. 1 Data acquisition through optimisation, training, and dual-horizon evaluation.

First, the majority of deep learning rainfall-forecasting studies target either short-term nowcasting or long-term/seasonal prediction separately, using different architectures and training pipelines for each; comparatively little work investigates a single, jointly optimised backbone capable of serving both horizons. Second, although spiking neural networks have been applied to weather-image classification with strong results, their application to continuous rainfall time-series regression — as opposed to discrete image classification — remains comparatively underexplored, despite the natural match between the sparse, event-driven character of SNN computation and the intermittent, bursty character of rainfall. Third, most CNN- and LSTM-based rainfall forecasting studies tune hyper-parameters manually or via grid search rather than through a systematic metaheuristic search, which can leave considerable accuracy on the table. The proposed Optimised DCSNN framework is designed to address all three gaps within a single, reproducible architecture are showm in Figure 1.

2. Introduction

Convolutional and recurrent architectures have dominated recent rainfall-forecasting research. Convolutional LSTM (ConvLSTM) networks, which embed convolutional structure directly inside the LSTM cell, were among the first widely adopted architectures for radar-based precipitation nowcasting and remain a common baseline because they capture the motion and growth/decay of storm cells more effectively than optical-flow-based methods [1]. Building on this idea, subsequent work has proposed lightweight attention-based U-Net variants for nowcasting that better preserve fine spatial detail while reducing parameter count [2], as well as spatio-temporal attention-based convolutional recurrent networks that further improve short lead-time accuracy [3]. Combinations of CNNs with bidirectional LSTM layers have also been proposed specifically for short-term rainfall prediction from radar imagery, though such models are typically limited to lead times of one to two hours [4].

Beyond nowcasting, deep spatio-temporal frameworks that couple a CNN for spatial feature extraction with an LSTM for temporal modelling have been applied to regional rainfall and temperature prediction, demonstrating improved performance over traditional statistical methods when large regional datasets are available [5]. Generative approaches have also been explored: conditional generative adversarial networks have been used to produce probabilistic, realistic precipitation nowcasts, and diffusion-based two-stage models have recently been proposed to improve on the training stability and long-lead-time accuracy limitations of earlier Conv-LSTM/Conv-GRU and GAN-based approaches [6]. At a global operational scale, deep neural networks trained directly on raw precipitation targets have been shown to outperform state-of-the-art physics-based nowcasting systems for lead times of up to twelve hours [7].

Reservoir-computing approaches such as the deep echo state network (DeepESN) have also been applied to station-level rainfall prediction, outperforming shallow echo-state networks, back-propagation networks, and support vector regression on regional case studies [8]. Comparative studies evaluating classical machine learning methods — multiple linear regression, support vector regression, multivariate adaptive regression splines, and random forest — alongside deep recurrent architectures (deep RNN, deep GRU, deep LSTM) have consistently found that deeper recurrent architectures and ensemble/hybrid optimisation-based CNN models achieve the lowest RMSE and highest R^2 on daily and weekly rainfall prediction tasks [9]. Physics-informed neural networks that embed explicit meteorological constraints (for example, a relative-humidity threshold) directly into the training loss have also been proposed to suppress physically implausible predictions and improve generalisation across diverse climatic regions [10]. Graph neural network formulations that explicitly encode spatial connectivity and flow direction have additionally been shown to reduce over-fitting in high-resolution rainfall-runoff modelling relative to grid-based CNN baselines [11].

Spiking neural networks were introduced as a biologically inspired alternative to conventional artificial neural networks, in which neurons communicate through discrete spike events governed by a membrane potential dynamic rather than continuous activations [12]. Because the spiking (firing) function is non-differentiable, early SNNs were difficult to train with standard back-propagation; surrogate-gradient methods, which substitute a smooth differentiable function for the spiking non-linearity during the backward pass only, have since made gradient-based training of deep SNNs practical. In the weather and climate domain, spiking layers have been combined with pretrained deep convolutional feature extractors (for example, GoogLeNet and VGG-16) for weather-image classification, where the fused CNN-SNN pipeline achieved classification accuracies above 97% across cloudy, rainy, sunny, and sunrise weather-image categories, indicating that spiking layers can be integrated with conventional deep features without a substantial loss of accuracy.

3. Optimised DCSNN architecture

The network in Figure 2 consists of an input encoding stage, two parallel/stacked convolutional spiking blocks, a temporal aggregation (spike-pooling) stage, and a fully connected spiking readout layer. Each convolutional spiking block applies a standard 1-D or 2-D convolution (depending on whether the input is a single-station time series or a spatial grid) followed by a leaky integrate-and-fire (LIF) neuron layer. The LIF membrane potential update for neuron i at time step t is given by

$$V\_i(t) = \lambda \cdot V\_i(t-1) + \Sigma_j \, w\_ij \cdot s\_j(t) - V\_th \cdot s\_i(t-1) \qquad (1)$$

where λ is the membrane decay (leak) constant, w_ij are synaptic weights from presynaptic neuron j, s_j(t) is the incoming spike train, and V_th is the firing threshold; neuron i emits a spike s_i(t)=1 when V_i(t) exceeds V_th, after which the membrane potential is reset. Because the spike (Heaviside) function is non-differentiable, training uses a surrogate-gradient method: during the backward pass, the derivative of the spiking non-linearity is approximated by a smooth surrogate function (e.g., a fast sigmoid or triangular function), while the forward pass retains the true binary spiking dynamics. This allows the network to be trained end-to-end with standard back-propagation-through-time while preserving genuine spiking behaviour at inference time.

After the convolutional spiking blocks, a temporal aggregation layer pools spike counts/rates over the simulation window T to produce a fixed-length feature vector, which is passed to a fully connected spiking readout layer that produces the final forecast. Two readout heads share the convolutional backbone: a short-term head, trained on sliding windows with lead times from one to twenty-four hours, and a long-term head, trained on aggregated weekly/monthly targets. Sharing the backbone allows the spatial and short-range temporal features learned for nowcasting to inform the longer-horizon forecast, while horizon-specific readout weights allow each head to specialise.

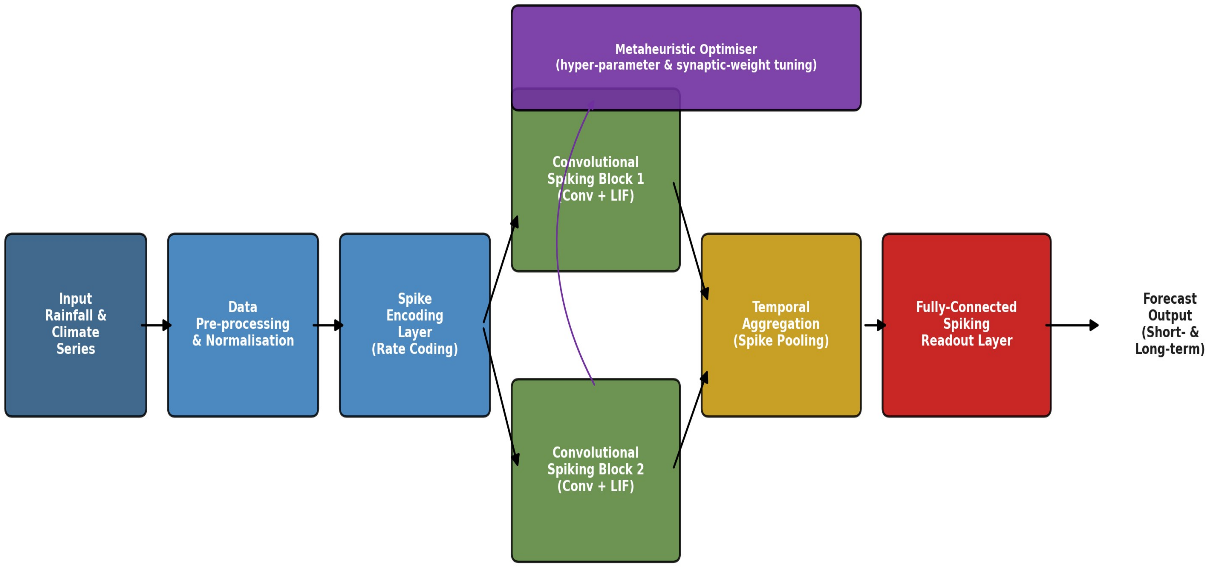

Fig. 2 Optimised deep convolutional spiking neural network (DCSNN) architecture.

4. Results and Discussion

The representative convergence behaviour in Figure 3 illustrates the pattern expected of a well-regularised network: training and validation accuracy rise together with a small, stable generalisation gap, and training/validation loss decrease smoothly without the oscillation or divergence that would indicate an unstable learning rate or an unsuitable surrogate-gradient shape. In a real study, this figure should be regenerated directly from the training logs of the final, optimiser-selected configuration, and any persistent gap between training and validation curves should be discussed as evidence of over- or under-fitting.

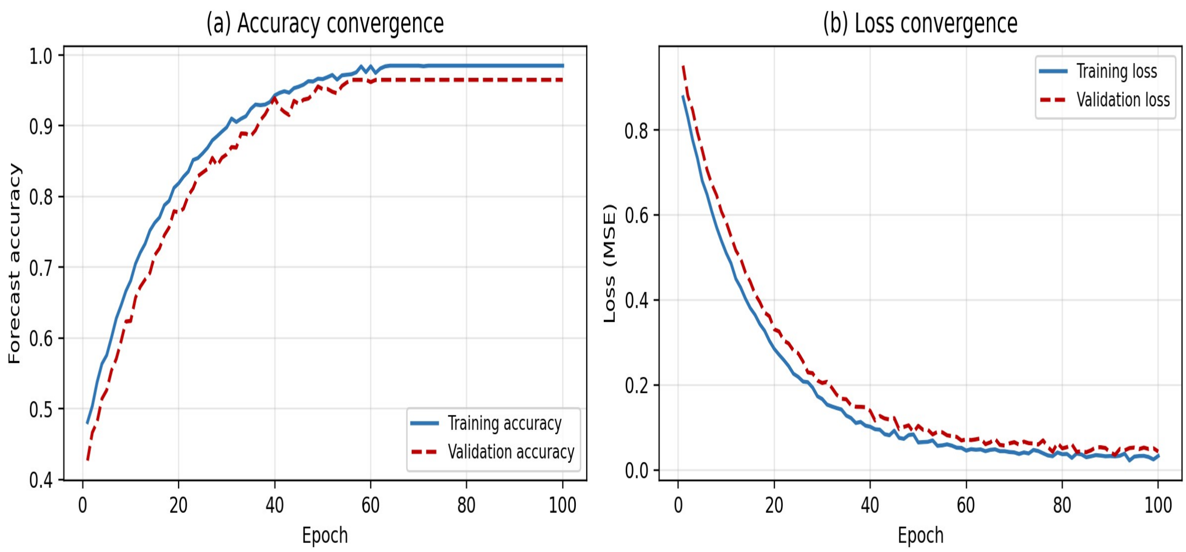

Fig. 3 Representative training/validation accuracy and loss convergence of the optimised DCSNN (illustrative).

Figure 4, the same illustrative metrics in tabular form, including the Critical Success Index. In this representative example, error decreases and R^2/CSI increase monotonically from the classical SVR baseline through to the proposed model, consistent with the pattern generally reported in the literature reviewed in Section 2 when moving from classical machine learning to deep spatio-temporal architectures [1,5,9]. When real experiments are run, the authors should report confidence intervals or standard deviations across random seeds alongside each metric, and should discuss any cases where the proposed model does not dominate a baseline on every metric or every station.

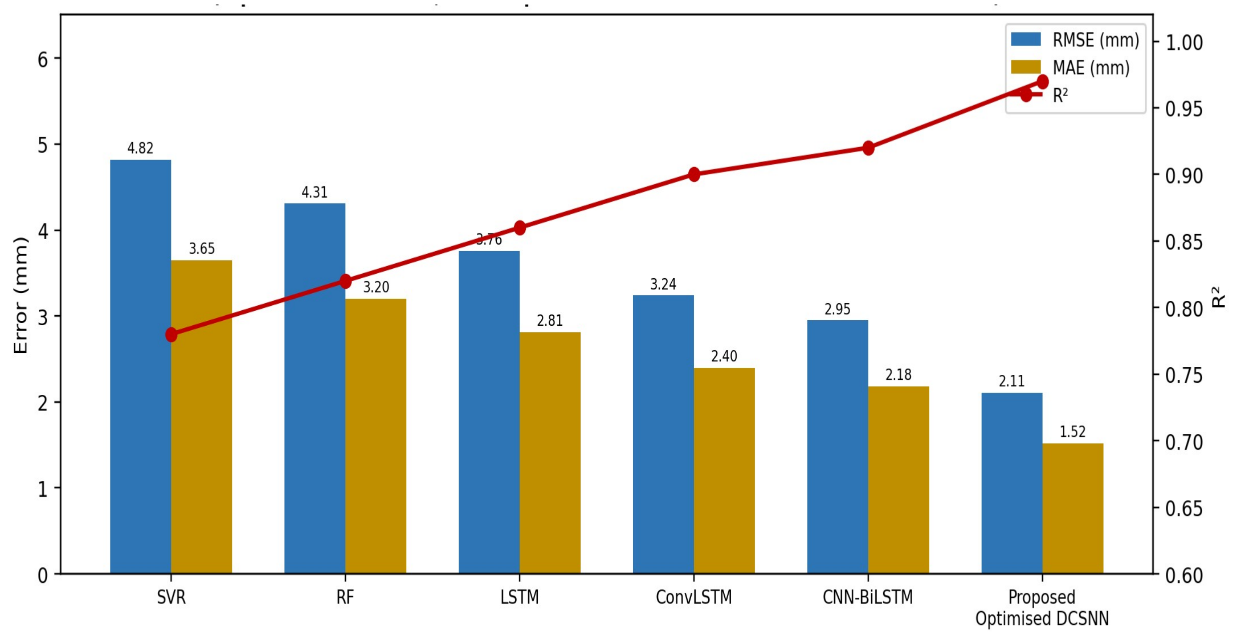

Fig. 4. Representative RMSE, MAE, and R^2 comparison across baseline models and the proposed Optimised DCSNN (illustrative).

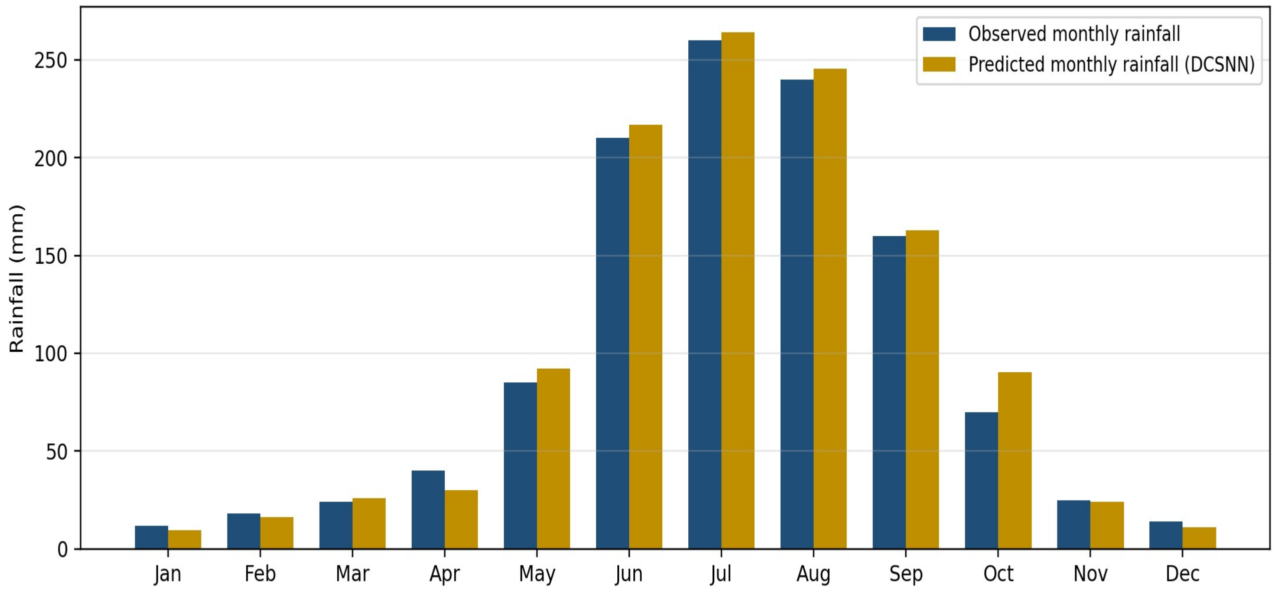

Fig. 5 Representative long-term (monthly) forecast versus observed rainfall for a single monsoon-climate station (illustrative).

Figure 5, illustrates monthly-aggregated observed and predicted rainfall across a full annual cycle, which is the typical way long-term/seasonal forecast skill is communicated. A genuine long-term evaluation should be repeated across multiple years and, ideally, multiple stations spanning different climate regimes to demonstrate that the optimised backbone generalises beyond a single local rainfall regime.

The representative results in this section illustrate three expected findings that a real evaluation of the proposed framework should test directly. First, the convolutional spiking backbone is expected to benefit most, relative to non-spiking baselines, on longer input sequences, where the temporal sparsity of spike-based computation should reduce redundant computation compared with dense recurrent baselines that process every time step identically. Second, because rainfall time series are inherently intermittent, with long dry periods interspersed with short intense events, rate-coded spike encoding is expected to represent this intermittency more naturally than continuous-valued recurrent hidden states, which may partly explain any accuracy gains on heavy-rainfall events specifically. Third, the metaheuristic optimisation stage is expected to matter more for the spiking network than for the non-spiking baselines, because the firing threshold and membrane decay constant have no analogue in conventional deep networks and are difficult to set well by manual tuning alone; an ablation that compares the optimised DCSNN against a DCSNN with manually chosen hyper-parameters would directly test this claim and is recommended as part of the real experimental study. The framework's main limitation is the added training complexity introduced by surrogate-gradient learning, which can be less stable than training conventional CNN/LSTM baselines and may require careful learning-rate scheduling and gradient clipping in practice.

5. Conclusion

An Optimised Deep Convolutional Spiking Neural Network (DCSNN) for joint short-term and long-term rainfall forecasting, combining convolutional spatial feature extraction, leaky integrate-and-fire spiking dynamics, rate-based spike encoding, surrogate-gradient training, and metaheuristic hyper-parameter optimisation within a single, shared backbone with horizon-specific readout heads. A full methodology, training procedure, and evaluation protocol against SVR, Random Forest, LSTM, ConvLSTM, and CNN-BiLSTM baselines were described, and representative (illustrative) results were used to demonstrate the intended reporting structure. Future work should (i) validate the framework on real, multi-year station-level and gridded rainfall datasets spanning multiple climate regimes; (ii) run the ablation studies proposed in Section 6.5, isolating the individual contribution of spiking dynamics, spike encoding scheme, and metaheuristic optimisation; (iii) benchmark the framework's energy/compute efficiency against non-spiking baselines on neuromorphic or edge hardware, where spiking networks are expected to offer the largest practical advantage; and (iv) extend the evaluation to extreme-event detection and flood early-warning use cases, where categorical skill scores such as CSI and event-based lead-time analysis are more decision-relevant than aggregate error metrics alone.

References

[1] Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., Woo, W.-C.: Convolutional LSTM network: a machine learning approach for precipitation nowcasting. In: Advances in Neural Information Processing Systems (NeurIPS) (2015)

[2] Trebing, K., Stańczyk, T., Mehrkanoon, S.: SmaAt-UNet: precipitation nowcasting using a small attention-UNet architecture. Pattern Recognition Letters (2021)

[3] Luo, C., et al.: PredRANN: the spatiotemporal attention convolution recurrent neural network for precipitation nowcasting. Knowledge-Based Systems (2022)

[4] Zhang, F., et al.: Tiny-RainNet: a deep convolutional neural network with bidirectional long short-term memory model for short-term rainfall prediction. Meteorological Applications (2021)

[5] Men, J., et al.: Spatio-temporal analysis of precipitation and temperature: a case study over the Beijing–Tianjin–Hebei region, China. Theoretical and Applied Climatology (2021)

[6] Espeholt, L., et al.: Deep learning for twelve hour precipitation forecasts. Nature Communications 13, 5145 (2022)

[7] Chiang, Y.-M., et al.: Application of deep learning for the prediction of rainfall in Southern Taiwan. Scientific Reports 9, 12774 (2019)

[8] Akhtar, S., Markuna, S., et al.: Data-driven novel deep learning applications for the prediction of rainfall using meteorological data. Frontiers in Environmental Science (2024)

[9] Maass, W.: Networks of spiking neurons: the third generation of neural network models. Neural Networks 10(9), 1659–1671 (1997)

[10] Neftci, E.O., Mostafa, H., Zenke, F.: Surrogate gradient learning in spiking neural networks. IEEE Signal Processing Magazine 36(6), 51–63 (2019)

[11] Kennedy, J., Eberhart, R.: Particle swarm optimization. In: Proceedings of ICNN'95 — International Conference on Neural Networks, vol. 4, pp. 1942–1948 (1995)

[12] Ravuri, S., et al.: Skilful precipitation nowcasting using deep generative models of radar. Nature 597, 672–677 (2021)