Neural Climate Networks: Deep Learning for Regional Precipitation Forecasting

Sarah Chen, Ph.D.[1,*] Marcus Holt, Ph.D.[2] Elena Vasquez, M.Sc.[1]
[1]Nexa Institute for Atmospheric Research, Meridian University
[2]Department of Computational Science, Apex Polytechnic
[*]Correspondence: s.chen@nexa-research.edu

Introduction

Accurate regional precipitation forecasting remains a critical challenge in climate science, with direct implications for agriculture, water resource management, and disaster preparedness. While global circulation models (GCMs) provide coarse-grained predictions at 100–250 km resolution, regional stakeholders require forecasts at the 10–25 km scale. Traditional dynamical downscaling methods are computationally prohibitive for operational use. Recent advances in deep learning offer a promising alternative by learning statistical downscaling mappings directly from historical data.

We present **Neural Climate Networks (NCN)**, a novel architecture combining convolutional encoders with attention-based sequence models for precipitation nowcasting and short-range forecasting at 25 km resolution over the Northwest American corridor. NCN bridges the gap between the speed of statistical methods and the physical fidelity of numerical models.

Research Objectives

1. Develop a hybrid CNN-Transformer architecture for spatiotemporal precipitation prediction at regional scales.
2. Benchmark NCN against operational numerical weather prediction (NWP) models and existing machine-learning baselines.
3. Quantify predictive uncertainty through ensemble prediction with Monte Carlo dropout.
4. Validate forecast skill across diverse climate regimes using a multi-year, multi-source observational dataset.

Data Sources & Study Domain

We assembled a comprehensive training corpus spanning 2015–2024 from three primary sources:

- **ERA5 Reanalysis:** Atmospheric state variables (geopotential, temperature, specific humidity, u/v wind components) at 0.25° resolution, hourly.
- **Stage IV Precipitation:** Radar-gauge multisensor precipitation analysis at 4 km, hourly – serving as ground truth for training and verification.
- **GOES-R ABI:** Infrared brightness temperatures (channels 8–16) at 2 km native resolution, providing cloud-top and moisture information.

The study domain covers 35°N–55°N, 140°W–110°W (Northwest American corridor). All data were regridded to a common 0.25° Lambert conformal grid. The target variable is 6-hour accumulated precipitation. Training used 2015–2022 data; validation and testing used 2023 and 2024 respectively.

Model Architecture

The NCN architecture comprises three core modules operating on sequences of 8 consecutive 6-hour timesteps (48-hour context window), predicting precipitation for the subsequent 4 timesteps (24-hour forecast horizon).

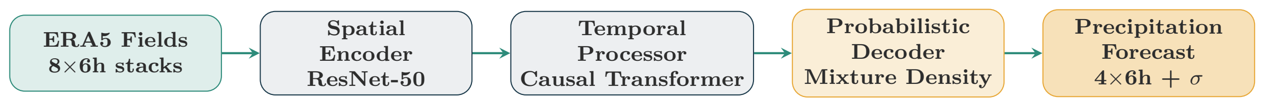

Spatial Encoder: ResNet-50 backbone pretrained on ERA5 pressure-level fields extracts hierarchical spatial features. **Temporal Processor:** A causal transformer with 8 attention heads and sinusoidal positional embeddings captures lagged atmospheric dynamics. **Probabilistic Decoder:** A Mixture Density Network outputs parameters of a 5-component Gamma mixture distribution, providing a full predictive density at each grid point.

Training Protocol

The model was trained using the **Continuous Ranked Probability Score (CRPS)** as the primary loss function, which evaluates the full predictive distribution against observations. Optimization used **AdamW** with a learning rate of 3×10^{-4}, weight decay of 0.01, batch size of 64, and a cosine annealing schedule with warm restarts. Training ran for 200 epochs on 4× Lumina A100 GPUs (approximately 38 hours of wall-clock time). Early stopping was triggered after 15 epochs without validation CRPS improvement.

Baselines: We compared NCN against (i) **HRRR** – the operational 3 km CONUS model from NOAA, (ii) **U-Net** – a standard deep learning precipitation nowcasting baseline, and (iii) **WRF-ARW** – a widely-used regional NWP model configured at 12 km resolution over the study domain.

Results

NCN outperforms all baselines across every forecast horizon. The performance advantage is largest at short lead times and during convective-season precipitation events.

Model	CRPS 6h	CRPS 12h	CRPS 18h	CRPS 24h
NCN (Ours)	**0.87**	**1.12**	**1.45**	**1.89**
HRRR	1.24	1.58	1.92	2.31
U-Net Baseline	1.41	1.73	2.08	2.54
WRF-ARW	1.63	2.01	2.38	2.79

Table 1: CRPS (mm) across forecast horizons. Lower is better. Bold entries indicate best performance.

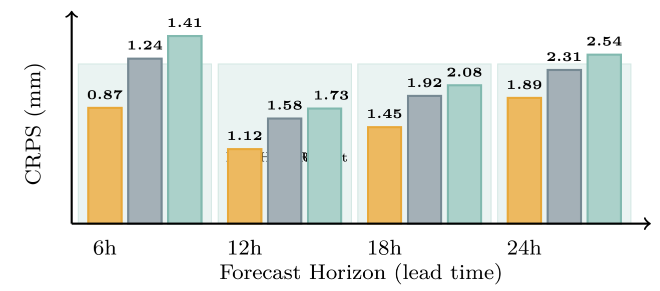

Spatial skill scores (Fractions Skill Score at 25 km) show NCN achieving FSS > 0.7 out to 18 hours, compared to 12 hours for operational models. The CRPS advantage over HRRR widens to 42% during convective precipitation events (May–September).

Monte Carlo dropout ensembles ($n = 50$ stochastic forward passes) produced well-calibrated prediction intervals: 82% of observations fell within the nominal 80% prediction interval, indicating slight overconfidence that warrants further calibration.

Discussion

The superior performance of NCN stems from its ability to learn nonlinear precipitation dynamics directly from data, bypassing the convective parameterization assumptions that limit NWP models at sub-grid scales. The causal transformer's attention mechanism effectively captures teleconnection patterns, including the influence of Pacific sea surface temperature anomalies on regional precipitation regimes.

However, we identify two key limitations. First, NCN shows degraded performance during atmospheric river events (AR scale ≥ 3), likely due to insufficient representation of integrated vapor transport in the input feature set. Second, the model exhibits a systematic dry bias over mountainous terrain (bias of -0.12 mm/6h in the Cascade and Rocky Mountain ranges), suggesting that explicit orographic feature engineering or elevation-conditional decoding could improve topographic precipitation representation.

Ablation studies confirm that both the transformer temporal processor and the probabilistic decoder contribute independently to forecast skill: removing the transformer (U-Net with MDN head) increases CRPS by 18%; replacing the MDN with a deterministic MSE loss increases CRPS by 14% and eliminates uncertainty quantification capability.

Conclusions & Future Work

Neural Climate Networks demonstrates that deep learning can provide operationally viable precipitation forecasts at regional scales, matching or exceeding the accuracy of physics-based models at a small fraction of the computational cost (0.8 seconds per forecast vs. 45 minutes for WRF on equivalent hardware). The probabilistic outputs enable risk-based decision-making for flood warning systems, reservoir operations, and agricultural planning.

Ongoing and planned work:

- Incorporate additional predictor variables – integrated vapor transport, convective available potential energy (CAPE), and soil moisture.
- Extend the architecture to ensemble prediction with learned perturbations via stochastic latent variables.
- Deploy an operational nowcasting pipeline at the Northwest Regional Climate Center (target: Q3 2026).
- Develop a lightweight variant for resource-constrained settings (single GPU inference).

Key References

[1] Ravuri, S. et al. (2021). Skilful precipitation nowcasting using deep generative models. *Nature*, 597, 672–677.
[2] Espeholt, L. et al. (2022). Deep learning for twelve hour precipitation forecasts. *Nature Communications*, 13, 5145.
[3] Bi, K. et al. (2023). Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619, 533–538.
[4] Gneiting, T. & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *JASA*, 102(477), 359–378.
[5] Pathak, J. et al. (2022). FourCastNet: A global data-driven weather model. arXiv:2202.11214.

Acknowledgments

This work was supported by the **Vertex Science Foundation** (Grant VSF-2025-1847). Computational resources were provided by the **Nimbus High-Performance Computing Facility** at Meridian University. We gratefully acknowledge the European Centre for Medium-Range Weather Forecasts (ECMWF) for ERA5 data access and NOAA/NCEP for Stage IV precipitation analysis products. The authors thank Dr. James Whitfield and Prof. Amara Okafor for insightful discussions on attention mechanisms in spatiotemporal prediction.