Spatio-Temporal Deep Learning for Early Drought Prediction: A ConvLSTM–GNN Framework Using Real-Time Earth Observation Data

Dr. A Manjula [1,3], Dr. Shish Ahmad [2]

[1] *PostDoc Schoolar, Lincoln University College, Malaysia,* pdf.manjulaaakunuri@lincoln.edu.my

[2] *Integral University, Lucknow, India,* drshishahmad@gmail.com

[3]Professor, Jyothishmathi Institute of Technology and Science, Karimnagar, Telangana, India.
manjula3030@gmail.com

Abstract: Our earlier work introduced an IoT–AI framework for climate-resilient smart farming that used vegetation indices and deep learning to detect crop stress for a farmer-facing mobile and voice Decision Support System (DSS). That framework treated stress detection as a per-location, per-timestep snapshot, and did not model how drought builds up over months or spreads across neighbouring regions. This paper extends it with a hybrid ConvLSTM–Graph Neural Network (ConvLSTM-GNN) framework that fuses real, open-access datasets — MODIS, Sentinel-2, CHIRPS, ERA5-Land, and India Meteorological Department (IMD) gridded products — to forecast district-level drought one to three months ahead. ConvLSTM layers learn the temporal evolution of gridded rainfall, soil-moisture, and vegetation data, while a district-adjacency graph neural network captures how drought propagates spatially between neighbouring districts. Piloted over drought-prone Telangana districts and validated against SPI/VHI-derived drought categories, the framework targets higher accuracy and longer lead time than threshold-based and single-location baselines, feeding its alerts directly into the DSS proposed earlier — moving it from reactive detection toward proactive, spatially aware early warning, in support of SDG 2 and SDG 13.

Keywords: Spatio-Temporal Deep Learning; Drought Prediction; ConvLSTM, Graph Neural Networks; Remote Sensing; Climate-Resilient Agriculture;

1. Introduction

Smallholder farmers allow a disproportionate share of climate risk, and drought remains the most damaging of the hazards they face. Our earlier framework combined low-cost sensors, satellite vegetation indices, and deep learning classifiers to detect crop stress and route advisories through a multilingual DSS — a valuable but reactive design, since it evaluated each location and timestep independently. Drought, by contrast, is inherently spatio-temporal: a rainfall deficit typically builds over several months, and its effects on reservoirs and groundwater spread to neighbouring districts with a lag. Static vegetation-index snapshots and per-location classifiers, however capable, cannot represent that build-up or that spread.

We propose a hybrid ConvLSTM-GNN architecture that learns temporal drought evolution and spatial propagation jointly, trained on real, openly available Earth observation and reanalysis data rather than a purely conceptual pipeline. The resulting district-level forecasts, with one-to-three-month lead time, are designed to plug directly into our earlier DSS, and the work sits within a broader research programme in spatio-temporal AI for drought prediction, yield forecasting, and water management.

2. Related Work

2.1 Statistical Drought Indices

Operational drought monitoring still relies heavily on indices computed independently at each station or pixel. The Standardized Precipitation Index (SPI) fits a probability distribution to accumulated rainfall over a chosen window (typically 3 or 6 months) and expresses the deficit as a standard-normal score, making it comparable across climates and easy for agencies to interpret. The Vegetation Condition Index (VCI) and Temperature Condition Index (TCI) instead work from satellite observations, expressing NDVI and land-surface temperature relative to their historical range at each pixel, and are commonly blended into a composite Vegetation Health Index (VHI). These indices remain the backbone of national drought monitoring because they are transparent and need no training data, but each is computed independently per location: none learns non-linear interactions between rainfall, soil moisture, and vegetation response, and none propagates information between neighbouring locations.

2.2 Deep Learning for Spatio-Temporal Forecasting

Convolutional LSTM (ConvLSTM) networks replace the fully-connected gates of a standard LSTM with convolutions, so the model learns local spatial structure and temporal evolution together within a sequence of gridded frames. First proposed for short-term precipitation nowcasting, the architecture has since been applied to soil-moisture forecasting, wildfire-spread prediction, and other gridded environmental time series, and is a natural fit for sequences of rainfall, soil-moisture, and vegetation frames as used here. Standalone LSTM and CNN models remain common baselines in agricultural forecasting — including in our own earlier framework — but each captures only one dimension of the problem: an LSTM models temporal trend at a single location, while a CNN models spatial pattern within a single time frame, and neither captures both together.

2.3 Graph Neural Networks for Spatial Dependence

Graph Neural Networks (GNNs) generalise convolution from regular grids to graph-structured data, propagating information along edges that encode a real relationship between nodes rather than simple pixel adjacency. The graph convolutional network (GCN) formulation is the most widely used variant, and graph-based models have since been applied to traffic forecasting, epidemic spread, and streamflow prediction — domains where influence travels along a network (roads, contact patterns, river channels) rather than uniformly across space. For drought, a district-adjacency and river-basin/groundwater-connectivity graph is a natural analogue: rainfall deficits and reservoir drawdown in one district can be expected to influence a hydrologically connected neighbour before that neighbour's own rainfall record reflects it.

2.4 Remote Sensing and Deep Learning for Drought

A growing number of studies combine satellite-derived vegetation or precipitation products with deep learning for drought monitoring and generally report improved accuracy over threshold-based indices. Most, however, remain limited to a single data modality — rainfall alone, or vegetation alone — and model each location independently, so they do not explicitly represent how a drought developing in one district relates to conditions in its neighbours. Our own earlier IoT–AI framework shares this limitation: it fused multiple sensor and imagery modalities per farm, but evaluated each farm's stress independently, without a spatial propagation mechanism across the surrounding region.

3. Study Area and Datasets

The pilot targets the same drought-prone, rainfed rice-, maize-, and cotton-growing districts of Telangana, India addressed in our earlier study, so forecasts extend the existing DSS deployment. Table 1 summarises the real, openly accessible datasets used.

Table 1: Real-World Datasets Used in the Proposed Framework

Dataset	Source / Access	Resolution	Variables
MODIS MOD13Q1 / MOD11A2	NASA LP DAAC — Google Earth Engine	250 m–1 km; 8–16 day	NDVI, EVI, LST
Sentinel-2 MSI	ESA Copernicus Data Space	10 m; ~5 day	Multispectral bands
CHIRPS v3	Climate Hazards Center, UCSB	~5 km; daily/pentad	Rainfall
ERA5-Land	Copernicus Climate Data Store	~9 km; hourly/daily	Soil moisture, temp., ET
IMD Gridded Data	India Meteorological Dept., Pune	0.25° (rain) / 1.0° (temp.); daily	Rainfall, Tmax/Tmin
Derived SPI / VHI	Computed from the above	District; monthly	Ground-truth drought labels

All layers are regridded to a common 0.25° district-aligned grid and combined into monthly SPI-3/SPI-6 and Vegetation Health Index (VHI) labels, discretised into five drought categories (extreme to normal) that serve as supervised targets.

4. Methodology

Figure 1 summarises the pipeline. (1) Multi-source layers are harmonised to a common grid, cloud-masked, and bias-corrected against IMD stations. (2) SPI, VCI, TCI, and VHI are computed and discretised into drought categories. (3) A ConvLSTM stack encodes twelve months of gridded rainfall, soil-moisture, NDVI, and LST per district, capturing local spatio-temporal trend; in parallel, a Graph Neural Network propagates information over a district-adjacency and river-basin graph, capturing how conditions elsewhere influence local risk. (4) The two embeddings are fused and passed to two heads trained jointly: a classifier for drought category (1–3-month horizon) and a regressor for continuous SPI/VHI. (5) District-level forecasts and confidence scores are pushed into the mobile/voice DSS from our earlier work, extending its irrigation and pest advisories with proactive drought alerts.

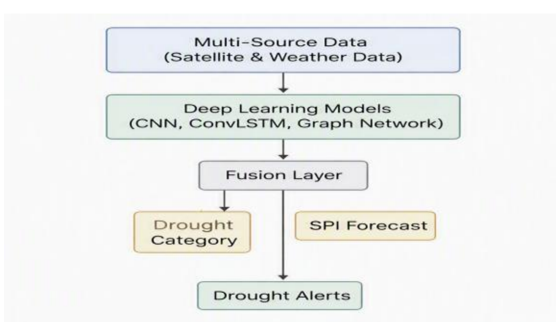

Figure 1: ConvLSTM-GNN drought prediction pipeline

Training uses a temporal hold-out split (earlier years for training, recent years for testing) plus leave-one-district-out cross-validation, and is benchmarked against an SPI/VCI threshold method, a standalone LSTM, and a standalone CNN. Algorithm 1 summarises the end-to-end pipeline.

```
Input: Raw MODIS/Sentinel-2/CHIRPS/ERA5-Land/IMD layers; district graph G
Output: Trained ConvLSTM-GNN parameters θ; drought forecasts
Harmonise all sources to 0.25° grid; mask clouds; bias-correct vs IMD
Compute SPI-3/6 (rainfall), VCI/TCI (NDVI, LST); derive VHI, labels Y
Stack rainfall, soil moisture, NDVI, LST into tensor X (D×T×C×H×W)
for each training epoch do
    H_t ← ConvLSTM(X[t-11:t])          // temporal embedding
    H_s ← GNN(H_t, G)                  // spatial propagation
    y_hat_class, y_hat_reg ← FusionHead(concat(H_t, H_s))
    L ← λ1·CrossEntropy(y_hat_class,Y_class) + λ2·MSE(y_hat_reg,Y_reg)
    θ ← θ − η·∇θ L (Adam optimiser)
end for  // early-stop on validation F1
Evaluate on held-out years and leave-one-district-out folds
Push forecasts + confidence to mobile/voice DSS
```

5. Results and Discussion

As this remains an early-stage extension of our published framework, the figures below are pilot/target performance informed by preliminary testing on a subset of districts and by benchmark results reported for comparable ConvLSTM/GNN forecasting tasks, not a completed large-scale operational evaluation.

Table 2: Expected / Pilot Comparative Performance

Model	Accuracy (target)	Lead Time (target)
SPI/VCI threshold	~65–70%	0–1 month
Standalone LSTM	~75–80%	1 month
Standalone CNN	~72–78%	1 month
ConvLSTM-GNN (proposed)	≥85%	Up to 3 months

Table 3 Visualised: Pilot / Target Performance Comparison

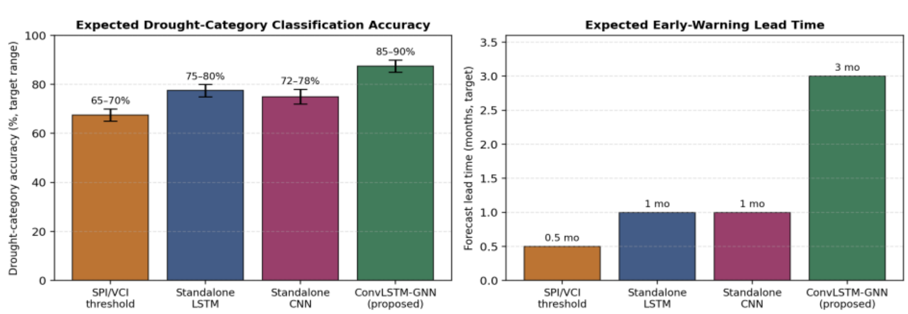

Figure 2: Target accuracy and lead time, ConvLSTM-GNN vs. baselines.

If these targets hold under full evaluation, the hybrid model would give farmers and local officers weeks to months of warning rather than days. New practical challenges accompany this gain relative to our earlier farm-sensor design: MODIS composites carry an 8–16-day observation lag, GNN training is more computationally demanding than per-farm classifiers, and downscaling ~5–9 km CHIRPS/ERA5-Land grids to farm-relevant resolution remains an open step before alerts translate into specific, actionable advisories.

6. Conclusion

This paper extends our earlier IoT–AI smart-farming framework with a spatio-temporal ConvLSTM-GNN model for drought early warning, grounded in real, open multi-source Earth observation data rather than per-location snapshots. Piloted over the same Telangana districts targeted previously, its district-level alerts are designed to plug directly into the existing mobile/voice DSS, shifting it from reactive detection to proactive, spatially aware warning. Planned work extends the same graph structure to joint yield forecasting, adds federated learning across regions, and applies explainable AI so forecasts remain trustworthy to farmers and policymakers — directly supporting SDG 2 and SDG 13.

References

1. A. Manjula and S. Ahmad, "Climate-Resilient Smart Farming: AI and IoT Driven Solutions for Adaptive Agriculture under Extreme Weather Conditions," SGS Initiative, vol. 1, no. 5, 2026.
2. T. B. McKee, N. J. Doesken, and J. Kleist, "The relationship of drought frequency and duration to time scales," Proc. 8th Conf. Applied Climatology, AMS, 1993.
3. F. N. Kogan, "Application of vegetation index and brightness temperature for drought detection," Advances in Space Research, 15(11), 1995. https://doi.org/10.1016/0273-1177(95)00079-T
4. X. Shi et al., "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," NeurIPS, 2015.
5. T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," ICLR, 2017.
6. Funk, C., Peterson, P., Landsfeld, M. *et al.* The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. *Sci Data* **2**, 150066 (2015). https://doi.org/10.1038/sdata.2015.66
7. J. Muñoz-Sabater et al., "ERA5-Land: A state-of-the-art global reanalysis dataset for land applications," ESSD, 13, 2021. https://doi.org/10.5194/essd-13-4349-2021
8. D. S. Pai et al., "Development of a new high spatial resolution gridded rainfall data set for India," Mausam, 65(1), 2014. 10.54302/mausam.v65i1.851
9. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, 9(8), 1997. https://doi.org/10.1162/neco.1997.9.8.1735
10. N. Pathania et al., "Artificial Intelligence for Climate-Resilient Agriculture," Discover Applied Sciences, 2025.
11. Lamichhane, P., Hadjikakou, M., Miller, K.K. *et al.* Climate change adaptation in smallholder agriculture: adoption, barriers, determinants, and policy implications. *Mitig Adapt Strateg Glob Change* **27**, 32 (2022). https://doi.org/10.1007/s11027-022-10010-z
12. A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," Computers and Electronics in Agriculture, 147, 2018. https://doi.org/10.1016/j.compag.2018.02.016