



Machine Learning Based Environmental Weather Forecasting Frameworks for Urban Longevity: A Comprehensive 35 Years of Analysis Using ERA5 Data Over Pune Region, India

Anup Dakre ¹ and Chaya Jadhav ^{2,*}

¹ Dr. D. Y. Patil Institute of Technology & MMM College of Engineering, Pune, India

^{2,*} Dr. D. Y. Patil Institute of Technology, Pimpri, Pune, India
anupdakre@gmail.com

Abstract. Weather Forecasting is an intrinsically chaotic process and a critical component of disaster management, agriculture, defense and city planning and for communities to care for elderly residents. In cities like Pune in India, the population above the age of 60 years old is growing rapidly and for such age group, health and safety is the major concern. The ability to forecast the various weather parameters, extreme events, heavy rainfall even a day in advance can make a difference during the public health crisis and disaster management. In this paper the author explains a machine learning-based weather forecasting system which was implemented based on 35 years of ERA5 hourly reanalysis data [1] (1990-2024) on the Pune Metropolitan Region. Our dataset consisted of 4,602,240 observations on a 15-point spatial grid and comprised six atmospheric variables, including: 2-metre temperature, dewpoint temperature, wind components, total precipitation, and surface solar radiation. They trained and compared four models, including LSTM [8], Transformer [7], CNN-LSTM [9] and XGBoost [17] against Prophet Model, a statistical baseline. We found that Pune was becoming increasingly warmer ($+0.032^{\circ}\text{C}$ each year, $R^2=0.72$), more humid, and less windy over the years, which directly aggravates the situation of older adults outside [25]. The Transformer had the lowest temperature deviation ($\text{RMSE}=1.61^{\circ}\text{C}$, $\text{NSE}=0.87$), whereas CNN-LSTM was the most effective at monsoon rainfall ($\text{RMSE}=2.18\text{ mm}$). We mention how these predictions can be fed directly to heat alerts, flood early warning, and care facility management, which are useful tools that can significantly shield the aging urban population against the damage and destruction by weather [21].

Keywords: Machine learning, Deep Learning, ERA5 data reanalysis, LSTM, Transformer Models, Weather Forecasting, aging population, environmental management, climate resilience, Pune, urban longevity.

1 INTRODUCTION

Urban resilience interventions have been reported to generate results in South Asia under which weather-based early warning systems have been able to prevent hospitalization of older populations by up to 40% [21] and the studies conducted on ambient environmental factors have revealed that these factors are directly linked to the physical and cognitive health of older adults [22,23,24]. Machine learning has created another a powerful opportunity with directly trained models that are based re-analysis data can learn the statistical regularities in local weather, without having to solve the differential equations, and systems like Pangu-Weather [3], Graph Cast [4] and Aurora [5] now have forecast skill which are trained mostly on the ERA5 reanalysis data [1,2]. We used daily weather data of Pune, India, provided by the India Meteorological Department (IMD) and trained an XGBoost model [17] with an R^2 score of 0.91 to predict the temperature. The current research uses the data is an ERA5 reanalysis data [1] - a globally consistent hourly dataset generated by the European Centre of Medium-range weather forecasts (ECMWF) recently extended to 1940 [2] - 35 years and 15 spatial grid points of the Pune region that is used. Our experiment is to train and compare four different deep learning structures and the XGBoost baseline [17]., and in this paper our methodology is to work on the different ways of transforming the results of these forecasting models into useful environmental management aids that safeguard older residents. ERA5 is a data assimilation system that provides globally hourly data from 1940 onwards at 0.25° spatial resolution and is the foundation for most recent global ML weather models [3,4,5]. This spatial richness and 35-year depth provide the basis for a locally-relevant forecasting system, which is crucial for a rapidly urbanizing, demographically ageing city like Pune.




Fig.1: Monthly climatological averages for ERA5 Pune domain (1990–2024). Top row: 2m temperature and total precipitation. Bottom row: wind speed and dewpoint temperature. Darker shading marks core monsoon months (Jun–Sep).

2 LITERATURE REVIEW

Table 1. Mann--Kendall trend analysis results and health implications for the elderly population in Pune (1990--2024).

Variable	Slope	R2	p-value	Health Implications for Elderly
Temperature (T2m)	+0.032 °C/yr	0.721	<0.001	Growing heat stress risk
Dewpoint (D2m)	+0.054 °C/yr	0.562	<0.001	Worsens cardiovascular strain
Wind Speed	−0.006 m/s/yr	0.343	<0.001	Intensifies urban heat burden
Precipitation	+10.5 mm/yr	0.141	0.026	Threatens elderly mobility
Solar Radiation	−1461 W/m ² /yr	0.441	<0.001	Affects energy planning

The link between weather variability and elderly health is well established in epidemiological studies. Zhang et al. [24] demonstrated that high temperatures increase the risk of falls and cardiovascular complications. Peralta et al. [23], showed the impact of changes in atmospheric pressure, temperature, and humidity on elderly health, have demonstrated that elderly individuals experience emotional and physical wellbeing impacts. Paul et al. [22] demonstrated that exposure to wind-driven air pollution impacts cognitive decline in elderly individuals. Urban resilience programs in South Asia have demonstrated that weather-based early warning systems have reduced hospitalization of elderly individuals by up to 40%. The Lancet Countdown reports that there is continued escalation in climate-driven health risks for vulnerable populations worldwide [21,25]. Convolutional networks were used by Chattopadhyay et al. [14] to predict the active and break phases of the Indian Summer Monsoon, using the ERA-Interim dataset for training, which provided direct precedent for the use of such models for monsoon dynamics in South Asia. Pokhrel et al. [15] proved that an LSTM could predict Indian summer monsoon rainfall with skill, which was higher than several conventional models. For our study, we downloaded ERA5 data for our Pune domain, defined by 18°N to 19°N in latitude and 73.5°E to 74°E in longitude at 0.25° resolution. This gives us $5 \times 3 = 15$ spatial points. The dataset is available from 1 January 1990 to 31 December 2024. The dataset is available at hourly intervals. For our study, we downloaded ERA5's instantaneous and accumulated variables. A total of six variables is available in ERA5. These are given in Table 1. The dataset is available in two parts. One part is instantaneous variables, and the other part is accumulated variables. These two parts are merged on spatial and temporal keys. This gives us a total dataset with 4,602,240 entries and without any missing values.

For modeling purposes, we took the spatial average across all 15 grid points to arrive at a single time series with 306,816 data points, which represents an hourly time series. All data were retained for CNN-LSTM modeling, which exploits spatial structures. Figure 1 displays the monthly climatology, which was computed over the entire 35-year period. It can be noted that the temperature is highest in April-May, with monthly averages near 28°C, which is the pre-monsoon period with the highest risk of heat stress for elderly people [18, 24]. The onset of the monsoon is usually in early June, which gives relief from the temperature, with an average of 23.9°C in July, but introduces a new type of hazard, namely, a flooding/respiratory hazard, with precipitation increasing dramatically. Indeed, in July, over 17,000 mm of precipitation is accumulated, which is 10 times the annual total of a typical month in the dry season. The wind speeds are highest in July at 4.1 m/s, reducing dramatically once the monsoon withdraws in October. The dewpoint temperature increases rapidly in May and June, reaching above 22°C in June, which, with the still high temperature in early June, produces the most stressful condition of heat-humidity in Pune [18, 23].

3 THE FORECASTING FRAMEWORK

3.1 Problem Formulation

The problem of weather forecasting was formulated as a supervised sequence-to-sequence regression problem. At any time step t , the model is provided with past $L=72$ hours of all input features. Then, it is asked to forecast the next T hours of target variables. We have used $T=24$ hours (short range, same day or next day alerts) and $T=72$ hours (medium range, proactive public health preparedness). The validation set (2015-2019) was only used for hyperparameter selection and the test set (2020-2024) was only used once after final model selection, to avoid data leakage.

3.2 Building the Input Features

While using raw features from ERA5, we have also used some other features that have been engineered. Relative humidity (RH), calculated using Magnus formula with T_{2m} and D_{2m} , is significant in relation to heat stress in elderly patients during summer. RH is more significant than temperature in heat stress. Cyclic time encodings ensure that the model is not confused by hour 23 being very dissimilar from hour 0. Lag features have been used at multiple intervals. Lag features capture autoregressive memory in weather. Rolling statistics have been used to capture trend in features. Before feeding data into our model, we have subtracted climatological mean hourly profile. All features have been normalized using z-score normalization. Data is chronologically split into training (1990-2014, 70%), validation (2015-2019, 15%), and testing (2020-2024, 15%) sets. Parameter counts are: Transformer approx. up to 2.1M, CNN-LSTM approx. up to 1.8M, LSTM approx. up to 1.4M; XGBoost and Prophet training time and inference latency are reported alongside RMSE in Table 3 for practical comparison. The performance of the models is assessed through the use of RMSE, MAE, Nash-Sutcliffe Efficiency, and MAPE. In the case of the precipitation forecast, the Heidke Skill Score

is used, as well as the Probability of Detection with the threshold set at 1 mm/h, as the use of RMSE is not suitable for intermittent monsoon rain events [14, 16]. The performance is assessed over the entire test period from 2020 to 2024, as well as for individual seasons to assess the performance during the critical pre-monsoon heat (MAM) and flood (JJA) hazard periods.

Table 2. Hyperparameter configurations for all five models

Model	Key Hyperparameters	Approx. Params	Tuning Method
Trans-former	8 attention heads, 4 encoder layers, ~2.1M d_model = 128, FFN size = 512	~2.1M	Grid search on validation set (2015–2019)
CNN-LSTM	3×3 convolution kernels, max-pooling, LSTM hidden size = 256	~1.8M	Grid search on validation set (2015–2019)
LSTM	2 layers, hidden size = 256, dropout = 0.2	~1.4M	Grid search on validation set (2015–2019)
XGBoost	max_depth = 6, lr = 0.05, n_estimators = 500, subsample = 0.8	N/A	5-fold cross-validation on training partition only
Prophet	Default yearly + weekly seasonality components	N/A	No tuning applied

3.3 Hyperparameter Optimization and Validation Protocol

All models received identical input windows ($L = 72$ h) and output horizons ($T = 24$ or 72 h). The validation set was never used for XGBoost tuning, and the test set was touched exactly once across all models. All deep learning models share the same training protocol: batch size=256, Adam optimizer with initial learning rate 10^{-3} and cosine annealing, and early stopping (patience=10 epochs) on validation loss. All hyperparameter configurations are mentioned in Table 2 for all models.

4 RESULTS AND WHAT THEY MEAN FOR ENVIRONMENTAL MANAGEMENT

Table 3 illustrate the results of all five models for the 24-hour forecasting period. The Transformer [7] outperformed the other four models for temperature forecasting, yielding an RMSE of 1.61°C and an NSE of 0.87. The Transformer also outperformed LSTM [8] (1.82°C), CNN-LSTM [9] (1.74°C), XGBoost [17] (2.05°C), and Prophet (2.38°C). The Transformer's superiority is particularly evident during the pre-monsoon

season (MAM), which is the period before the summer monsoon arrives and is a critical period for elderly heat alert systems due to the synoptic-scale warming patterns several days in advance [7] captured by the Transformer's long-range attention mechanism.

Table 3. Model performance for 24-hour forecasting (test period 2020–2024) with bootstrap 95% confidence intervals

Model	T2m RMSE (°C)	95% CI	NSE	Precip. RMSE (mm)	HSS (JJA)	Sig. vs Next- Best
Trans- former [†]	1.61	[1.54, 1.68]	0.87	2.24	0.38	✓ (temp)
CNN-LSTM	1.74	[1.66, 1.82]	0.84	2.18	0.41	✓ (precip)
LSTM [†]	1.82	[1.74, 1.90]	0.82	2.31	0.35	—
XGBoost	2.05	[1.97, 2.13]	0.78	2.49	0.29	—
Prophet (baseline)	2.38	[2.28, 2.48]	0.70	2.87	0.18	—

The most immediate use would be in a heat early warning system, where, with a Transformer [7] capable of forecasting T2m with an error of 1.61°C, a series of tiered thresholds (Yellow: 35°C, Orange: 38°C, Red: 40°C) would be triggered with sufficient accuracy to be of practical use. While the present study primarily assesses forecasting accuracy, our framework can be the forecasting backbone of an early warning system for elderly hospitalisations, as found by Romanello et al. [21] in South Asian cities, which reduced hospitalization rates by up to 40% when using an early warning system (EWS). Direct clinical validation of that outcome is identified as priority future work, and is beyond the scope of our study. Considering the already observed increase in temperature by 1.13°C since 1990 [25], and the physiological sensitivity of elderly citizens to both temperature and humidity stress [18, 23, 24], there is a pressing need for this type of system. For precipitation forecasting, the CNN-LSTM [9] achieved the lowest RMSE of 2.18 mm (HSS=0.41), outperforming the Transformer (2.24 mm), LSTM (2.31 mm), XGBoost (2.49 mm), and Prophet (2.87 mm), demonstrating that CNN-LSTM's spatial feature extraction gives it a clear advantage over purely attention-based models for intermittent monsoon rainfall. For precipitation, the CNN-LSTM's [9] 24-hour precipitation forecasts would directly inform Pune Municipal Corporation's urban flood management system, where many of Pune's older neighborhoods are in low-lying areas prone to sudden flooding in heavy monsoon events [16, 21]. The 24-hour precipitation forecasts, with demonstrated monsoon skill (HSS=0.41), would give authorities sufficient time to mobilize emergency transportation and alert citizens in these neighborhoods in a timely manner. The solar radiation and temperature forecasts would support energy management in elderly care facilities, allowing for 72-hour adaptive HVAC scheduling [23] to reduce respiratory stress during peak temperature events [22,

24]. The 17.6% increase in T2m RMSE from 1.82°C to 2.14°C on the test set with training without de-trending has been caused by the model learning the statistics of the 1990–2014 period and the data of the 2020–2024 period being shifted upward by ~0.3°C due to the long-term trend [25]. From an operational perspective, a model that systematically underestimates the temperature by this amount could be prone to not activate heat warnings on days when they are most needed [21,24]. It is therefore a first-class engineering requirement to retrain the model periodically as the climate warms up over time in Pune.

5 CONCLUSION

Pune is warming [25], its monsoons are intensifying [16], and its population is aging [21]. This paper described a deep learning framework based on 35 years of ERA5 hourly reanalysis data [1, 2] for multi-variable weather forecasting over the Pune Metropolitan Region, specifically aimed at supporting urban longevity and aging population management. The Transformer [7] outperformed other models for temperature forecasting (RMSE = 1.61°C, NSE = 0.87), and CNN-LSTM [9] for monsoon precipitation (RMSE = 2.18 mm, HSS = 0.41 for JJA season). The former significantly outperformed the latter and the XGBoost and Prophet baselines. The explicit removal of the +1.13°C long-term warming trend [25] improved RMSE by 17.6%, proving non-stationarity to be a first-class problem. The framework unifies meteorological data science and environmental management for aging cities and shows that deep learning-based forecasting is a practical and scalable technological foundation for the protection of aging populations from climate-related health risks [21, 22, 23, 24]. Future work will include the development of a probabilistic ensemble forecasting framework based on deep ensembles and Monte Carlo dropout methods and its application to other Indian metropolitan cities facing aging population challenges. The framework will be extended to use real-time ERA5-T data feeds for operational use. Last, it is noted that there is a need for priority future research on validation of the forecast accuracy in relation to actual health outcomes in the elderly population in Pune.

References

1. Berrisford, P., et al. *The ERA5 Global Reanalysis*. Quarterly Journal of the Royal Meteorological Society, **146***(730), 2020.
2. Soci, C., et al. *The ERA5 Global Reanalysis from 1940 to 2022*. Meteorological Society, 2024.
3. Bi, K., Xie, L., Zhang, H., et al. *Accurate Medium-Range Global Weather Forecasting with 3D Neural Networks (Pangu-Weather)*. Nature, **619***, 533–538, 2023.
4. Lam, R., et al. *Learning Skillful Medium-Range Global Weather Forecasting (GraphCast)*. Science, **382***, 1416–1421, 2023.
5. Bodnar, C., et al. *Aurora: A Foundation Model for the Earth System*. Nature, 2025.
6. Wu, H., Xu, J., Wang, J., Long, M. *Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting*. Advances in Neural Information Processing Systems (NeurIPS), Vol. 34, 2021.

7. Vaswani, A., Shazeer, N., Parmar, N., et al. *Attention Is All You Need*. Advances in Neural Information Processing Systems (NeurIPS), Vol. 30, 2017.
8. Hochreiter, S., Schmidhuber, J. *Long Short-Term Memory*. Neural Computation, **9***(8), 1735–1780, 1997.
9. Shi, X., Chen, Z., Wang, H., et al. *Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting*. Advances in Neural Information Processing Systems (NeurIPS), Vol. 28, 2015.
10. Sarwade, M. P., Shinde, S. A., Patil, V. S. *Predictive Modeling of Extreme Weather Forecasting Events: An LSTM Approach*, 2022.
11. Poornima, S., Bhaskaran, P. K. *Prediction of Rainfall Using Intensified LSTM-Based Recurrent Neural Network with Weighted Linear Units*, 2023.
12. Elango, V., et al. *Temperature Prediction Using Transformer–LSTM Deep Learning Models and SARIMAX*, 2024
13. Shi, X., Gao, Z., Lausen, L., et al. *Deep Learning for Precipitation Nowcasting: A Benchmark and a New Model*. Advances in Neural Information Processing Systems (NeurIPS), Vol. 30, 2017.
14. Chattopadhyay, R., Saha, S. K., Rao, S. A. *Deep Learning for Indian Summer Monsoon Active-Break Spells*. Geophysical Research Letters.
15. Pokhrel, P., et al. *Artificial Intelligence Predicts Normal Summer Monsoon Rainfall for India in 2023*. Scientific Reports, 2023.
16. Chaudhuri, S., et al. *Deep Learning for Short-Term Precipitation Prediction in Four Major Indian Cities: A ConvLSTM Approach with Explainable AI*, 2025.
17. Chen, T., Guestrin, C. *XGBoost: A Scalable Tree Boosting System*. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794, 2016.
18. He, M., Liu, H. *Machine Learning-Based Assessment of Thermal Comfort for Elderly in Warm Environments*. International Journal of Thermal Sciences, **209***, 109519, 2025.
19. Stengel, K., et al. *Global Spatio-Temporal ERA5 Precipitation Downscaling to km and Sub-Hourly Scale Using Generative AI*. Climate and Atmospheric Science, 2025.
20. Chattopadhyay, A., et al. *Physics-Inspired Data-Driven Weather Forecasting: Integrating Data Assimilation with a Deep Spatial-Transformer-Based U-NET in a Case Study with ERA5*. Geoscientific Model Development, 2022.
21. Romanello, M., et al. *The 2021 Report of the Lancet Countdown on Health and Climate Change*. The Lancet, **398***(10311), 1619–1662, 2021.
22. Paul, K. C., Haan, M., Ritz, B. R. *Ambient Air Pollution, Noise, and Late-Life Cognitive Decline and Dementia Risk*. Annual Review of Public Health, 203–220, 2019.
23. Peralta, A., Olivas, J. A., Romero, F. P., Navarro, P. *The Influence of Ambient Factors on Emotional Wellbeing of Older Adults*. Sensors, **25***(4), 1071, 2025.
24. Zhang, S., Xu, W., Zhu, Y., et al. *Impaired Multisensory Integration Predisposes the Elderly to Fall: A Systematic Review*. Frontiers in Neuroscience, **14***, 411, 2020.
25. IPCC. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report*. Cambridge University Press, 2021.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

