Deep Learning for Climate Downscaling: Generating high-resolution gridded temperature projections over India from low-resolution CMIP6 data

Danish Ansari[1] and Tabish Ansari[2*]

[1] *Independent Researcher, Bangalore, India*

[2] *Research Institute for Sustainability - Helmholtz Centre Potsdam (RIFS), Potsdam, Germany*

**Correspondence to: tabish.ansari@rifs-potsdam.de*

1. Abstract

Global climate change leads to distinct regional impacts. Future climate projections are available through global model outputs but they lack the spatial detail needed to inform local decision making for adaptation. Regional- and local-scale climate information can be generated by "downscaling" these global model outputs through regional climate model simulations which require enormous supercomputing resources thereby limiting the availability of local-scale projections especially in underdeveloped regions. Here, we build and test four different computationally-efficient machine learning models as an alternative to regional climate models for downscaling daily temperature data over India. We progressively improve our model design by utilising spatial learning and then including temporal learning and optimising the target region for downscaling. We systematically evaluate all the four models and select the best variant to project daily temperatures over India for 2030. We present seasonal maps, frequency distributions as well as city-specific time-series of daily temperatures for the downscaled data vis-a-vis the original global model data. We find a slight overall increase in the annual averaged downscaled temperatures (0.29°C) and an even larger increase in the most frequent (modal) temperature (1.05°C) but also a decrease over certain regions, particularly in the Indo-Gangetic plain, in summer and monsoon as compared to the coarse-scale temperatures in the global model data. Our machine learning model compares favourably with reference data from the ERA5 reanalyses and runs much faster, with much lesser computational resources than a typical physics-based regional climate model and therefore opens up the possibility of democratising the field of climate downscaling.

2. Introduction

Global climate change represents one of the most pressing challenges of the 21st century (IPCC, 2021). The consequences of this phenomenon are increasingly evident, ranging from rising temperatures (Hansen et al., 2006; Allen et al., 2018; Samset et al., 2023) and sea levels (Nicholls et al., 2010; Fox-Kemper et al., 2021) to more frequent and intense extreme weather events (IPCC, 2021, Scott, 2016; Frame et al., 2020). While global climate models provide invaluable insights into the broader trends and patterns of climate change, they often lack the spatial resolution required to capture the intricacies of regional climate variability (Seneviratne et al., 2020). Understanding how these global changes manifest at the regional scale is essential for assessing their impacts on local ecosystems, communities, and economies (Jacob et al., 2014;

Diffenbaugh et al., 2018). Consequently, the process of downscaling global climate projections to regional scales has emerged as a crucial area of research, aiming to bridge the gap between coarse-resolution global models and the finer-scale information needed for localised decision-making (Pal et al., 2007; Maraun et al., 2010).

Local-scale climatic information is increasingly available through in-situ ground observations, but they are geographically sparse, leaving large regions of the planet unobserved. Remote Earth observations made by satellites cover more ground than in-situ observations but they are still spatially and temporally incomplete. The complete picture is provided by merging ground and remote observations with simulation outputs from physical science-based models in a gridded format - this is called a 'reanalysis' product (e.g., Gelaro et al. 2017, Hersbach et al., 2020) which yields the most holistic climatic information but only until the present day.

In addition to past and present climatic conditions, local bodies also require future projections of climate in their region to inform local strategies. Physical science-based global climate models are excellent tools which provide future projections for the entire globe under various socioeconomic pathways which essentially boil down to different emission scenarios (Riahi et al. 2017). Multi-model merged datasets, or ensembles, are publicly available to support decision-making throughout the world (Taylor et al., 2012; Eyring et al., 2016). However, these global model outputs, while they are effective in shaping broad international and national-scale policy, lack spatial detail to effectively represent fine-scale changes, which limits their usefulness for local action. In order to obtain fine-scale climate information, these global model outputs need to be 'downscaled' to finer scales, typically of the order of a few kilometres grid cells.

The conventional process of generating fine-scale climate information from coarse-scale information is essentially the same as that used to generate the coarse-scale information from fixed initial conditions. It entails defining the entire globe on a 2D map using some map projection (e.g., cylindrical, conic or polar) and dividing the map into several gridcells where distinct climatic information is desired. Then, fundamental equations of physics, concerning conservation of mass, heat and momentum are solved for each gridcell utilising known local initial conditions (Warner, 2010a). Particularly, these are coupled partial differential equations involving both time and space derivatives in 3D cartesian space of key meteorological variables such as air velocities (u, v, w), air density (ρ), temperature (T), heat (H), and moisture (q) which lack a direct analytical solution. Therefore, they are solved using approximate numerical methods such as finite differencing by transforming them into a system of linear equations, which requires several iterations to converge to a final solution (Dudhia, 1993; Skamarock & Klemp, 2008; Warner, 2010b). This requirement makes the process inherently computationally heavy. If the gridcells are smaller (and therefore more numerous to cover the entire globe), they permit better representation of geographical features such as mountain topography, land-use and soil type, lakes and rivers etc. as well as better treatment of local moisture and cloud formation which enable a more realistic climatic description at a finer scale. However, solving the physical equations iteratively over a large number of gridcells in a short amount of time demands impractical computational capacity. To circumvent this problem, the global models are run at a coarser spatial resolution (which also permits longer time-stepping iterations due to a mathematical criterion and therefore quicker calculations; Warner, 2010b) with fewer coarse gridcells covering the globe. The global model output is then used to define initial conditions over a finer limited-area grid to perform similar physics-based calculations, which furnishes a fine-scale gridded climatic information over a chosen area of interest (e.g., Takle et al. (1999) over the USA, Christensen & Christensen (2007) over Europe, and Fu et al. (2005) over Asia). This process is called "dynamical downscaling" and suffers from the same computational disadvantages as the original method of global climate modelling; it demands heavy computational resources, typically high-performance supercomputers only available in large national-scale laboratories in the developed world. Naturally, this has led to an inequity in the availability of fine-scale climate information across the world. For example, the developing regions in the Global South, (e.g., South Asia) currently do not have publicly available long-term high-resolution downscaled gridded products from the latest CMIP6 global climate data.

Machine learning techniques, especially deep learning methods, are garnering increasing interest in the Earth and Environmental sciences (Reichstein et al., 2019; Ham et al., 2019; Liu et al., 2020; Ansari et al., 2023). Deep learning techniques, with their neural network-based model designs can be efficiently used for climate downscaling due to their low computational demand, especially after the initial model training (Lin et al., 2023; Oyama et al., 2023). Once a machine learning model has been constructed, or 'trained', it can be run on a personal computer to output downscaled climate information, typically in a matter of seconds, as opposed to days and weeks taken by physics-based models running on high-performance supercomputers with several CPUs. This is because running machine learning models does not involve iterative numerical solutions of differential equations but rather a multidimensional input-output mapping based on regression or classification algorithms. Some studies have used simple feed-forward artificial neural networks that rely on temporally stationary predictors-predictand relationships to downscale a coarse-resolution gridded model output to specific point locations with some success (e.g., Sachindra et al., 2018; Nourani et al., 2019) but have concluded that including temporal learning can improve downscaling results. Nourani et al. (2019) also used a larger geographic area for predictors than the geographic extent of the predictand to maximise spatial learning and potentially boost downscaling performance. This idea was tested more thoroughly by Kim et al. (2022) where they identified most relevant regions for plausible teleconnections (climatic influences from remote regions) and included them into their training dataset. More recent studies have tested convolutional neural networks (CNNs) that rely on spatial learning and can read climatic data in form of 2D images (Vandal et al., 2019; Liu et al., 2020; Baño-Medina et al., 2021; Damiani et al., 2024) for climate downscaling. Some studies have also explored inclusion of both spatial and temporal learning between the predictors and predictand by using different variations of Recurrent Neural Network (RNN) - based model design (Miao et al., 2019; Li et al., 2020; Wang et al., 2020; Kimura et al., 2021) and have shown improved model performance. Rampal et al. (2022) performed high resolution downscaling of rainfall over New Zealand using CNN-based model architecture and also explored the learned relationships in the CNN using explainable AI tools and discussed the physical plausibility of the revealed relationships, thereby providing a rich precedent for future studies.

The use of deep learning for climate downscaling is relatively new in the context of climate research in India. At the time of writing this manuscript, we are only aware of two deep learning-based downscaling studies over India: Sharma & Mitra (2022) who used a neural network architecture to downscale rainfall from 1°x1° to 0.25°x0.25° and Lin et al. (2023) who used a CNN-based Unet model architecture to downscale key meteorological variables primarily over East Asia but also included India within their prediction region. India presents a unique and intricate climatic landscape, offering an ideal testing ground for the development of advanced downscaling techniques. The nation's geography encompasses diverse topographical features, from the towering Himalayan mountain range to the expansive coastal plains and the arid deserts of the northwest. The Indian subcontinent is influenced by a multitude of climatic systems, including the Indian Summer Monsoon, which brings vast amounts of rainfall to the subcontinent from June to September (Gadgil and Joseph, 2003). Additionally, the complex interplay of local weather phenomena, such as the Western Disturbances and the Indian Ocean Dipole, significantly impacts regional climate variability (Dimri et al., 2015; Ashok et al., 2001). The country's coastline, spanning over 7,500 kilometres, exposes it to the maritime influences of both the Arabian Sea and the Bay of Bengal, resulting in distinct regional climatic patterns. Therefore, to effectively downscale climate data over India, it is necessary to capture the intricate interplay between these diverse climatic systems, their seasonality, and their interactions with the varied terrain at high spatial resolution.

Here, we have utilised neural networks to perform grid-to-grid downscaling of daily mean temperatures over India. We have tested four different neural architectures and two distinct target regions to perform temperature downscaling from 1.87° x 1.87° grid resolution to 0.1° x 0.1° grid resolution specifically over the Indian region. We describe the design of our experiments and models in section 3, present our model evaluation in section 4, discuss the results in section 5 and finally present our conclusions along with future outlook in section 6.

3. Experimental design

We conducted four experiments using neural networks to predict daily surface temperatures at 0.1° x 0.1° horizontal resolution over India using coarse resolution data from CMIP6 as input and fine resolution from ERA5 as target. We utilised ERA5-Land reanalysis data (https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-land?tab=form) over India (37°N 68°W 8°S 97.5°E) and CMIP6 data from the model MPI-ESM1-2-LR Germany (https://cds.climate.copernicus.eu/cdsapp#!/dataset/projections-cmip6?tab=form) over a larger region including India (50°N 50°W -20°S 100°E), generated for the SSP5-8.5 scenario. Both ERA5-Land and CMIP6 provide gridded data, with ERA5-Land having high spatio-temporal resolution of hourly 0.1° x 0.1° grids while CMIP6 having a daily resolution with 1.87° x 1.87° grids.

Due to computational and storage constraints, we only used the data for the recent five years (2017-2022) from both CMIP6 and ERA5-Land in order to build our models. Another year of data (2030) was used from CMIP6 in order to generate the predictions using the built models. We utilised the 2-metre temperature variable from the ERA5-Land dataset as the target variable, while four different predictor variables were prescribed from the CMIP6 dataset, namely, near surface air temperature, sea level pressure, near surface wind speed, and near surface specific humidity.

The CMIP6 variables were fed into the model in their original form except a change of units for near surface air temperature from Kelvin to degree Celsius. For the ERA5-Land dataset we converted the hourly 2m temperature to daily mean temperature by taking an average of all 24 values for each day and converted them from Kelvin to degree Celsius.

We defined two categories of experiments: first, which consider only spatial correlations between input and target data, and second, which consider both spatial and temporal correlations. We performed two experiments for each of these categories depending on the spatial extent of target dataset and therefore the output that they generate, i.e., the entire ERA5-Land region as target (37° N 68°W 8°S 97.5°E), and a cropped portion of the ERA5-Land region following the national boundary of India (see, Figure 1 and Table 1).

The wider region of CMIP6 used as input data ensures that the neural networks capture not only the local correlations between the input and target datasets but also learn more distant physical relationships including teleconnections. Similarly, the two experiments with a two-week historical input can allow the neural networks to learn the chain of events that may lead to a particular temperature at a location.

In experiments 1 and 2 we treated the daily ERA5 2m temperature values and the CMIP6 predictors values as paired time series and built convolutional neural network based models which can take the CMIP6 predictor gridded values for each time step and produce a corresponding ERA5 2m temperature equivalent grid. In experiments 3 and 4 we took into account the spatial as well as temporal correlations and built LSTM-based Recurrent neural network models which can take a 14 day history of CMIP6 predictor gridded values and produce the current ERA5 2m temperature equivalent grid.

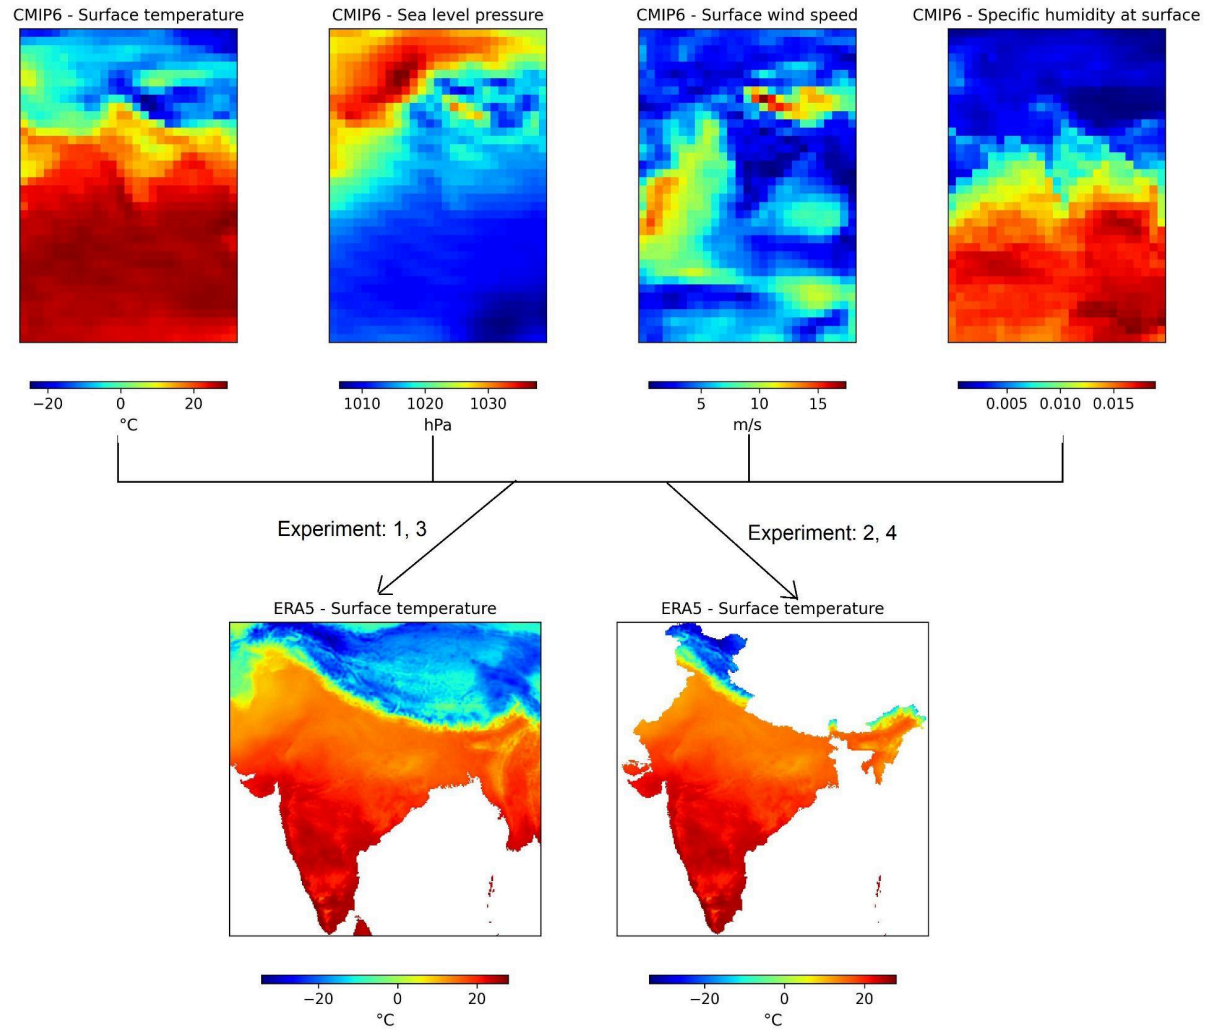

Figure 1: Contour plots of model input variables a) surface temperature b) sea-level pressure c) surface wind speed and d) surface specific humidity from CMIP6 member model at 1.87° x 1.87° horizontal resolution and target variable e) 2m temperature from ERA-5 Land dataset at 0.1°x0.1° horizontal resolution. All variables were used at a daily average level in the model construction.

Exp. No.	Correlations learnt	Input	Output	ML model architecture
Experiment 1	Spatial	CMIP6 predictors for day N	Complete ERA5 2m temperature for day N	Convolutional Neural Network
Experiment 2	Spatial	CMIP6 predictors for day N	Cropped ERA5 2m temperature for day N	Convolutional Neural Network
Experiment 3	Spatial and temporal	CMIP6 predictors - 14 day history	Complete ERA5 2m temperature for day N	Recurrent Neural Network (ConvLSTM)
Experiment 4	Spatial and temporal	CMIP6 predictors - 14 day history	Cropped ERA5 2m temperature for day N	Recurrent Neural Network (ConvLSTM)

Table 1: Details of Experiments performed.

3.1 Transformation of gridded data and treatment of ocean grid cells

For each time frame, the neural networks built in our experiments (via the TensorFlow library) can only handle 1D target data and therefore furnish 1D output data. This necessitated a grid transformation of the ERA5-Land data from 2D to 1D arrays before feeding it into the models. However, a major problem was the treatment of missing values over the ocean grids. Since neural networks cannot be fed with missing values, we tested filling the ocean grid cells with zeros but this biased the model results over land towards colder temperatures and lower variability. We also tested using ocean values from another dataset which contains ocean temperatures but were limited by a mismatch in grid resolution which would introduce further errors into our model and still bias the model output towards lower temperature variability since the temperature fluctuations over ocean are milder than those over land. Therefore, we shortened the 1D array by removing the oceanic data points. Our models were eventually trained on these shortened 1D arrays and consequently produced a 1D array of the same shortened length as output. We stored unique location identifiers in this shortened 1D array which enabled us to redistribute the output values back onto the ERA5-like 2D grid to visualise our results.

3.2 Building and training the models

We build a CNN/RNN-based model architecture, train and select the best performing model by comparing training loss against the validation loss, and perform the predictions on validation and test sets. We then compare the predictions against the unseen ERA5-Land 2m temperature values to determine the model accuracy. A more detailed evaluation was performed later which is discussed in section 4.

Experiment 1 (No History - Uncropped Target)

We created a convolutional neural network model which could take as an input the data points, each of which were of the shape (38, 27, 4) where (38, 27) are the spatial dimensions denoting latitude/longitude grids, while 4 corresponds to the four CMIP6 variables. Although this is not an image data which is usually what is processed through CNNs, it is structurally similar to it where the first two dimensions can be treated as a 2 dimensional plane consisting of pixel values, while the third dimension can be treated as the channels of the image. We can visualise this as a stack of four 2D sheets, each sheet containing the grid values against one CMIP6 variable. CNN treats this as a 2D image with four channels. The output of the CNN was set to the shape of the shortened 1D array (length = 57617) mappable to the 2D ERA5-Land 2m temperature data as discussed earlier. Activation function used for each Conv2D layer, and the dense layer, was ReLU (Rectified Linear Unit). While the activation function used for the output layer was Linear. We applied L2 type activity regularisation on both Conv2D layers and the Dense layer, with the l2 parameter for the first Conv2D layer as 0.01, while for others 0.001. The dropout rate used in both the dropout layers was 0.2. We set the optimizer to Adam with a learning rate of 0.01, the loss function as 'Mean squared error', and metrics to be collected during training as 'mean squared error' and 'mean absolute error'. We set the batch size to 32 and epochs to 100. The model summary is included in the Appendix.

The training and validation period was 01-JAN-2018 to 01-AUG-2022 while the out-of-sample test dataset spanned from 01-JAN-2017 to 31-DEC-2017. We prescribe a train-validation split of 80:20 for both CMIP6 and ERA5-Land data which resulted in 1339 training data points and 335 validation data points. Here, one data point refers to one day which includes all spatial grid cell values. We used the standard scaler to scale the CMIP6 training data, preserved the scaler mean and variance and used the same parameters to scale the CMIP6 validation data, test data and 2030 CMIP6 input data later used for climate projections.

With these training and validation datasets, we initiated the training of the model while monitoring the training and validation losses. We saved the trained models at every 10 epochs, and chose the best model based on the loss difference, and made predictions using the chosen model over the unseen test data set. We obtained a high correlation coefficient (R^2=0.97). Scatter plot between the model predictions versus test data is shown in Figure 2.

Experiment 2 (No History - Cropped Target)

Experiment 2 was similar in all respects to Experiment 1 except that the target was changed from the downloaded ERA5-Land data (37° N 68°W 8°S 97.5°E) to cropped data limited to the national boundary of India (see Figure 1). The motivation to crop the target data was to eliminate the Tibetan plateau region which is a distinct climatic zone and would presumably be affecting the model efficacy due to its colder temperatures compared to the temperatures within India. The CNN model in this experiment resembled the model structure as the one from Experiment 1, except that its output layer had a smaller number of neurons (28750) due to a smaller target array. The model summary is included in the Appendix. For this experiment we also obtained a high correlation coefficient between the predicted and test values (R^2=0.96; see Figure 2).

Experiment 3 (14 Day History - Uncropped Target)

In this experiment we wanted to exploit both spatial and temporal correlations for downscaling, we built a ConvLSTM-based Recurrent Neural Network (RNN). To learn any possible temporal correlations we supplied sequences of 14 CMIP6 time frames as data point input instead of a single time frame as in experiments 1 and 2. We chose the sequence length of 14 somewhat arbitrarily but it was motivated by an expectation that correlations must exist between the fine-grid temperatures of today and the four coarse-grid variables up to the past two weeks.

Data used for model training, validation, and testing was same as mentioned in experiment 1, except the fact that we created CMIP6 data sequences of 14 days history as input to the model while the target remained the same - current time step temperature grids of ERA5-Land. The input data points had a shape of (14, 38, 27, 4) where the first dimension is the 14 days, second and third denote the 2D grids, and the fourth represents the four predictor variables of CMIP6. The target data 1D array had the dimensions (57617) which is the same as that from Experiment 1. The model summary is included in the Appendix. Activation function used for each ConvLSTM2D layer and the dense layer, was Rectified Linear Unit (ReLU), while the activation function for the output layer was Linear. The dropout rate in both the ConvLSTM2D layers and after the dense layer was 0.2. We set the optimizer to Adam with a learning rate of 0.01, the loss function as 'Mean Squared Error', and metrics to be collected during training as 'Mean Squared Error' and 'Mean Absolute Error'. We set the batch size to 32 and epochs to 100. For this experiment we obtained a slightly poorer correlation coefficient between the predicted and test values as compared to previous experiments (R^2=0.94; see Figure 2).

Experiment 4 (14 Day History - Cropped Target)

Finally, we performed experiment 4 with a similar design as experiment 3 except that we used the cropped target data as in experiment 2. The model summary is included in the Appendix. For this experiment we also obtained a high correlation coefficient between the predicted and test values (R^2=0.96; see Figure 2). A more comprehensive evaluation of all experiments is presented in section 4.

3.3 Climate projections

Using the models built in each of the four experiments described above, we generated downscaled high-res gridded data for the future year 2030 using the coarse-res CMIP6 data as input. Therefore, we obtained four different realisations of daily fine-scale 0.1°x0.1° gridded temperatures for the entire year. ERA5-Land data is only available until the present day and therefore our ML-model results fill a crucial gap in climate information for the near-future. We performed further analyses with these downscaled projections to derive new insights which are discussed in section 5.

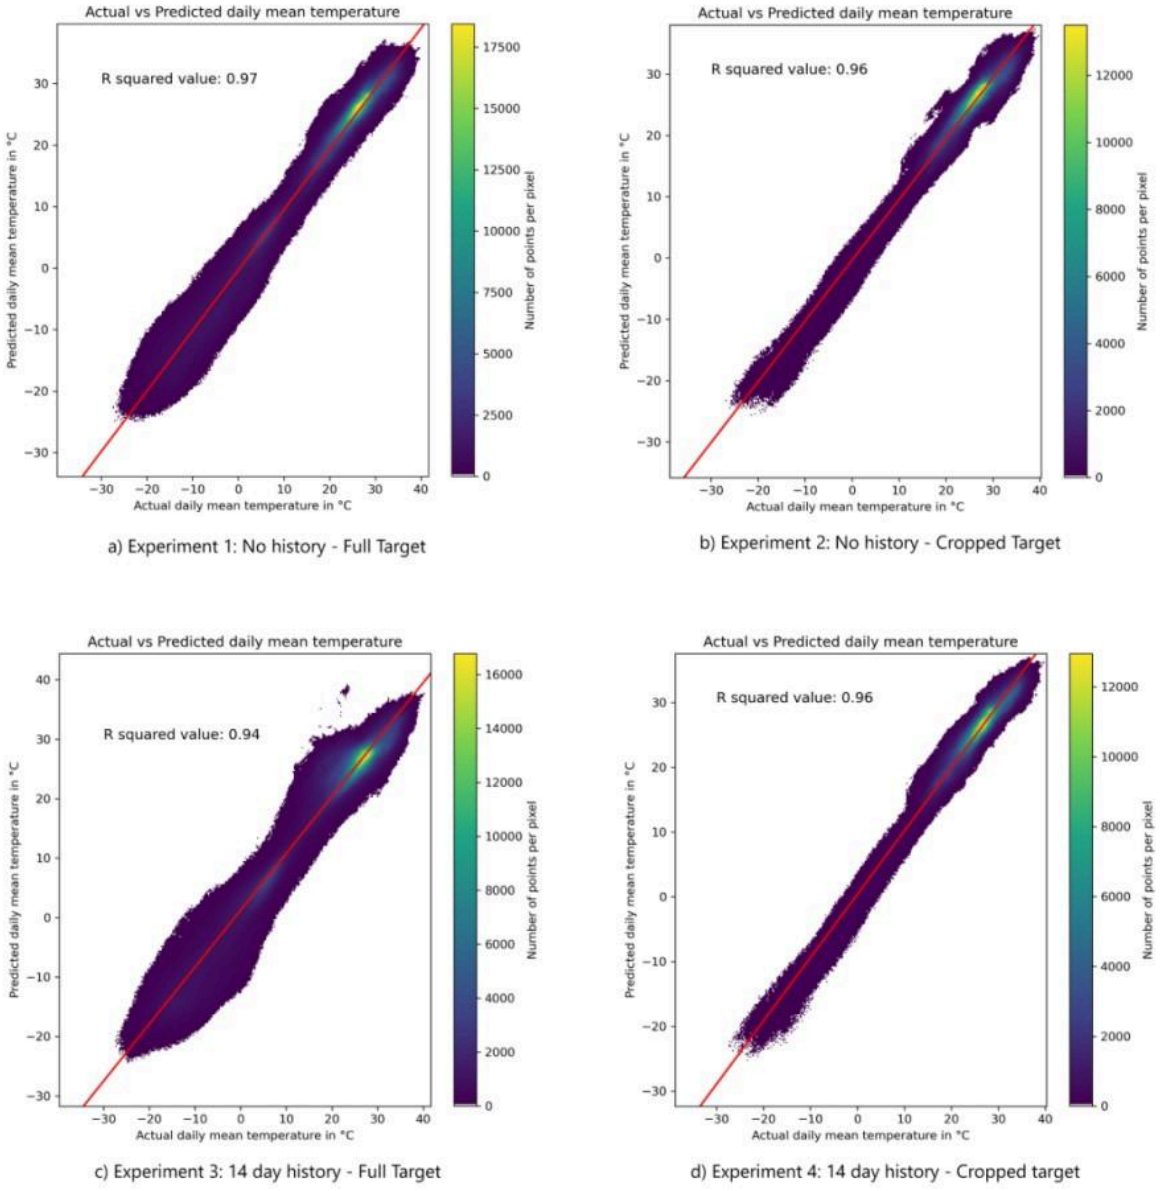

Figure 2: Model prediction vs Actual ERA5-Land 2m temp for test data

4. Model evaluation

In this section we systematically evaluate all the four experiments by comparing the ML model daily temperatures against ERA5 daily temperatures over a one year period (2017) which was not utilised in building the ML model. We begin by comparing the conventional metrics such as root mean squared error (RMSE), mean absolute error (MAE) and model bias but also present the model performance in terms of its geographic and seasonal variability as well as in terms of frequency distribution.

Figure 3a shows the Root Mean Squared Error (RMSE) time series for the four experiments. We find that all experiments show better results during the monsoon season (July - September) and relatively less good performance during post-monsoon, winter and summer. Experiment 3 shows the poorest performance overall with RMSE reaching upwards of 6°C during post-monsoon. Other experiments show a relatively less fluctuating RMSE time series. Clearly, the best performance is shown by Experiment 4 with a large portion

of the RMSE time series around 1ºC. Figure 3b shows the Mean Absolute Error (MAE) time series for the four experiments. We find that all experiments show similar performance as in the case of RMSE. Experiment 4 shows the best performance with MAE being close to or under 1ºC for a significant part of the year-long testing period. It is worth noting that again all experiments show their best performance during the monsoon season.

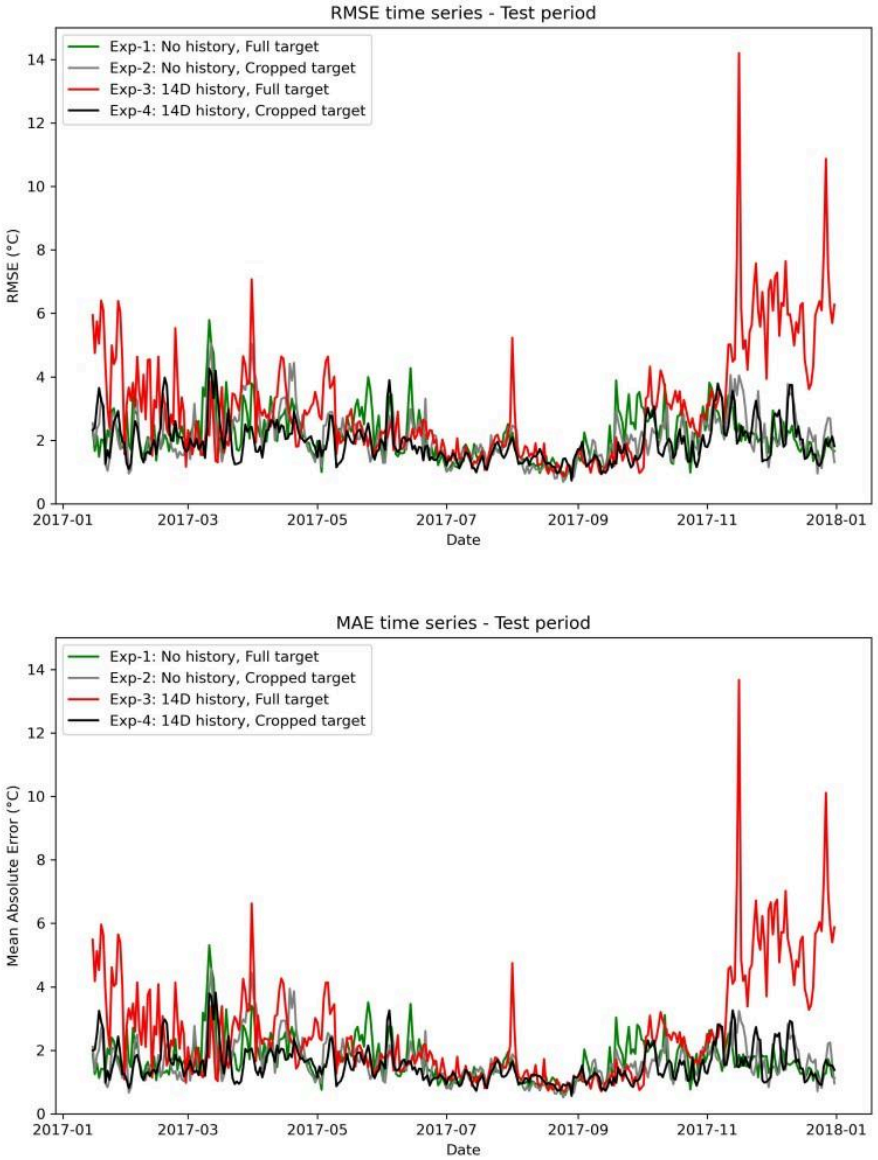

Figure 3: a) Root mean squared error and b) Mean absolute error for the four experiments evaluated against ERA5 data over the 1-year period (2017) which was not used to build the model.

To gain further insight into the model deviation from the reference ERA5 temperatures in terms of its geographic performance, we plotted a time series that provides the percentage of grid points that are within an error range of 1ºC, 2ºC and 3ºC (Figure 4). Here, it is unveiled that the model performance for experiments 1, 2 and 4 are in fact quite comparable in terms of their predictions, while experiment 3 is clearly less accurate. We find some common seasonal features in terms of model predictions for all these experiments which are worth noting: all experiments show their best performance during the monsoon season (as also noted in earlier metrics) but they show the poorest performance during the transition from winter to pre-monsoon. This fact merits further scientific attention - the drastic change of season

accompanied the reversal of winds during pre-monsoon season over India, as discussed in other studies (e.g., Gadgil 2006, 2018; Ansari et al. 2016) isn't skillfully captured by the neural network. Overall, experiment 4 shows the best performance with on average 43.25% data points predicting a temperature that is within 1°C of the reference ERA5 temperature and a very good performance during the monsoon season with ~80% grid points predicting a temperature within 1°C and ~95% grid points within a 2°C temperature of the reference dataset.

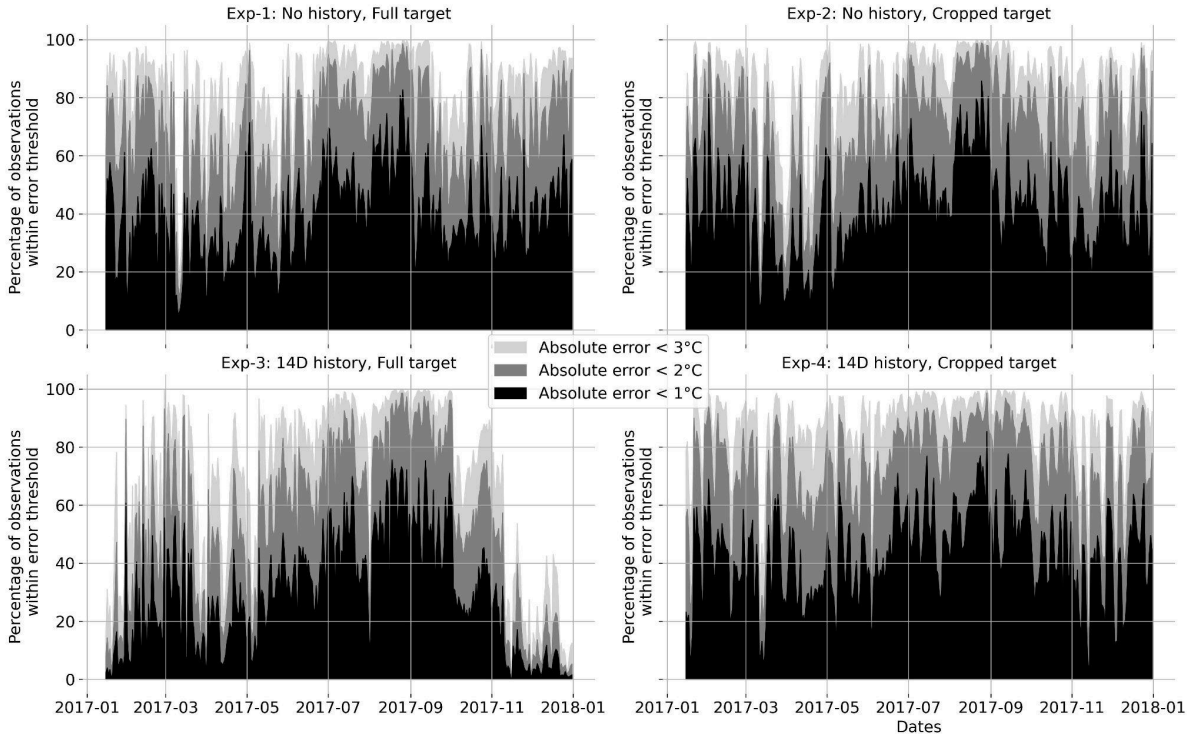

Figure 4: Time series of percentage of predicted grid-point values for each experiment that are within 1°C (black), 2°C (mid grey) and 3°C (light grey) of the reference grid-point values from the ERA5 dataset.

To further understand the model performance in terms of its regional performance on an annual basis, we plot the percentage grid point predictions that are within a 1°C error from the reference dataset as a geographic map (see, Figure 5). Again, experiment 3 shows the least good performance for most regions, experiments 1 and 2 show reasonable performance for large parts of India while the best performance is shown by experiment 4. In general, the experiments show best predictions for coastal, central and southern India with around 60-80% grid points falling within a 1°C error from the reference ERA5 dataset, which we consider a very satisfactory performance given the meagre amount of training data provided to the neural networks. Of course, it is desirable to further investigate the days when these regions exceed this error threshold which we have done in the next analysis shown in Figure 6. All model variants struggle to accurately predict the daily temperatures particularly in two regions: the western region characterised by the Thar desert and the northern border region of Kashmir which are generally very hot and cold respectively as compared to the other regions within the country. This provides us with a hint that the neural network architecture utilised here (i.e., CNN and RNN) struggles with representing sharp temperature gradients and extreme temperatures especially with contiguous regions with less extreme temperatures. This result also provides a motivation to try out other neural network architectures that may be better suited to represent extreme values than mean values.

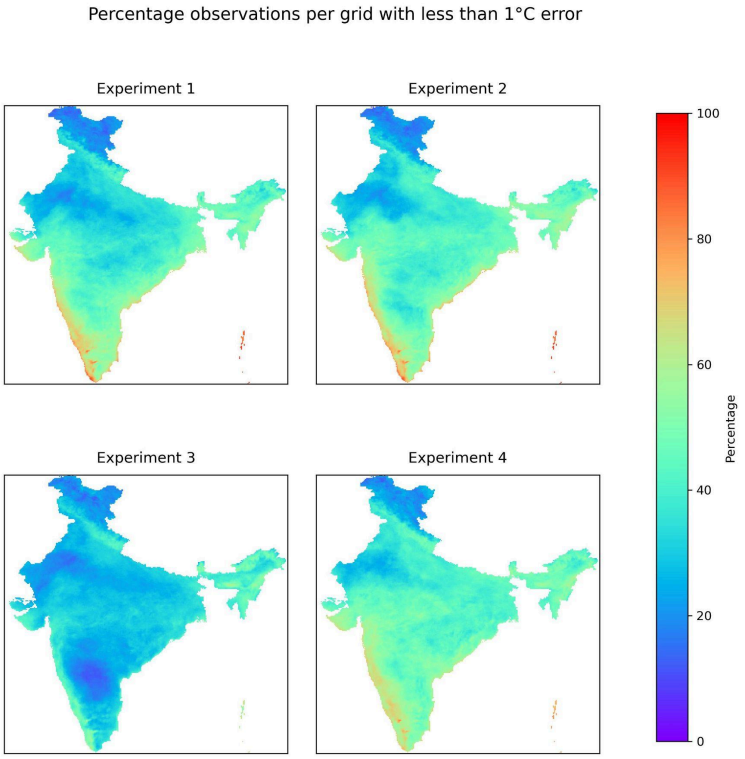

Figure 5: Geographical distribution of percentage of predicted data points (out of 365 daily predictions) that are within a 1°C difference from the reference data points from the ERA5 dataset over the year 2017.

To further understand the model performance in both regional and temporal terms, we decomposed the geographic error analysis as shown in Figure 5 into four distinct seasons of the Indian subcontinent: Winter (Dec-Feb), Pre-monsoon/Summer (Mar-May), Monsoon (Jun-Sep) and Post-Monsoon (Oct-Nov). The results are shown in Figure 6. In Winter, experiments 1, 2 and 4 show good performance with some regions, particularly in peninsular India, showing under 1°C error for up to 90% of the days of the year. The model performance in all experiments is slightly worse in the pre-monsoon season - experiment 4 showing the best results with under 1°C error for large parts of the country for up to 60% of the days of the year and a particularly good accuracy along the west coast with values reaching upwards of 80%. During the monsoon season, all experiments show good performance as seen in previous analyses. Here, experiments 1, 2 and 4 show comparable regional accuracy with values consistently above 70% for large parts of India. Finally, in post-monsoon, again experiments 1, 2, and 4 show good performance. While experiments 1 and 2 show a high temporal accuracy in coastal regions, experiment 4 shows a uniformly good accuracy for large parts of the country covering most regions (except the Thar desert and Kashmir) where the predictions are within 1°C of the reference values for almost 60% of the days of the year.

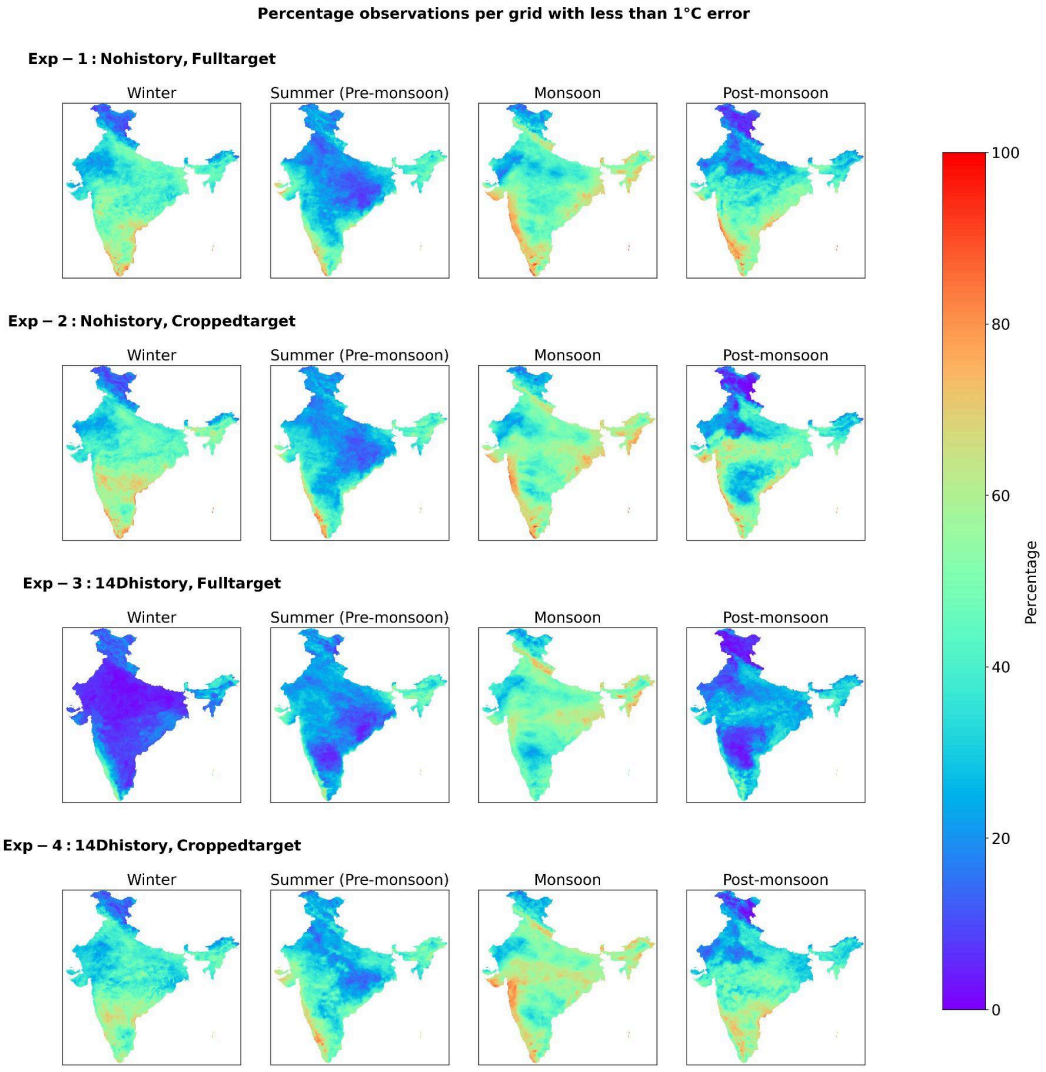

Figure 6: Same as Figure 5 but for different seasons: Winter (Dec-Feb), Summer (Mar-May), Monsoon (Jun-Sep) and Post Monsoon (Oct-Nov)

After getting a clear picture of model performance in terms of regional and temporal accuracy of the daily predictions, we investigate the country-wide frequency distribution of daily temperatures for a) Course resolution CMIP6 model data, b) Fine-resolution ERA5 data and c) ML model downscaled data in order to understand the overall temperature biases and distribution across these datasets. This is crucial from a climate adaptation policy perspective as the course resolution dataset might have large regional biases and the key value of downscaling is to identify these biases at the regional level. Figure 7 presents the frequency distribution of daily mean temperatures across India as derived from CMIP6 data alongside the physics-based downscaled and observation-assimilated ERA5 data and the ML model downscaled data over the year 2017. Since India is a tropical country, here we are mainly concerned with the maximum rather than minimum extremes particularly in relation to public health related to heat strokes which is becoming common in the country in recent years (Kovats & Koppe, 2005; Rohini et al., 2016; Sharma et al., 2017; Nori-Sarma, 2019). We find that CMIP6 data has a hot bias with 6.25% data above 35°C. This is reduced to only 1.88% data in the downscaled (and likely more accurate due to observational assimilation) ERA5 data. The ML model downscaled data successfully mirrors this temperature distribution with only 0.34% data points exceeding 35°C. This is an encouraging result both in terms of the efficacy of the neural network to

transform the coarse CMIP6 data into a finer resolution while also changing its frequency distribution but also in terms of a public health emergency. We would like to point out that the 35°C number here may not seem too severe but it is only the daily mean temperature and not the daily maximum temperature. When the daily average temperatures reach 35°C, the maximum temperatures might well surpass 45°C or more, causing a huge public health emergency, loss of lives and economy.

Another striking feature of the temperature frequency distribution across the three datasets is the most probable, or frequent, daily temperature, i.e. the modal temperature which shows a significant increase in the downscaled datasets. Since we have identified experiment 4 as the best model from our previous analyses, here we only focus on the nuances of the distribution from that experiment. The modal temperature for the CMIP6 dataset is 24.5°C which is increased by more than 2°C to 26.8°C in the ERA5 downscaled dataset. This is well captured by the ML model (experiment 4) at 25.7°C, still showing an increase of 1.2°C over the CMIP6 dataset.

While so far we have gained a clear picture of the downscaled data in terms of its overall geographic and temporal/seasonal distribution, we also extracted time-series data at five megacities of India (Delhi, Mumbai, Kolkata, Bangalore and Chennai) from the corresponding grid points of the CMIP6 dataset, the ERA5 dataset and the ML model output dataset from all the four experiments over the testing period of 2017 as shown in Figure 8. As previously established, experiment 4 shows the most accurate results with closest temperatures to the ERA5 data for each of the cities. It is encouraging to see that the ML model is able to reproduce the large and pronounced seasonal change in daily temperatures over Delhi which experiences a continental climate as well as the modest seasonal changes of the coastal cities of Mumbai, Kolkata and Chennai. It also captures the early pre-monsoon peak in Bangalore followed by a pronounced dip in monsoon. Overall, the ML model successfully predicts the annual temperature cycle at a city level for all the major cities which is notably quite different from the coarse resolution CMIP6 model output data whereby the warm biases are corrected in pre-monsoon in Delhi, Kolkata and Bangalore and cold biases are corrected in monsoon in Mumbai and Chennai. These results are promising in terms of downscaling future projections from CMIP6 data using our ML model which will likely also correct the built-in biases for a future date and allow us to obtain a more plausible, accurate and localised temperature data. We will explore this in the next section.

After deciding that experiment 4 shows the best overall performance across regions and seasons, we plotted the ML model results against the ERA5 data for four representative days for each season as shown in Figure 9. The difference between the ML model output and the ERA5 output is also plotted alongside the results for each season. We find that the ML model shows a good performance with large parts of the country lying within a 1°C difference from the ERA5 data (shown in white). These low error regions are the largest during the monsoon as previously noted. In winter, the model shows overprediction over Kashmir region and underprediction over the Northeast states with some overprediction over central and southern regions. In summer, the ML model output is colder by ~2°C around the Thar desert region and 1-2°C warmer around the eastern coastal regions but within a 1°C to the ERA5 data for other regions. In monsoon, most of the habitable parts of the country are within 1°C of the ERA5 data. In winter, most of the heavily populated Indo-Gangetic Plain and southern India is within 1°C of the ERA5 data while there is a cold bias in the Thar desert region and a warm bias in eastern India. Please note that these are only four individual days chosen to represent the four distinct seasons, a full day-to-day comparison plot for 2017 is available as a GIF file from https://drive.google.com/file/d/1vpGollnHtK2iK-wb5rVL8pGd5tEtbgLE/view?usp=sharing where the errors might be larger over different regions.

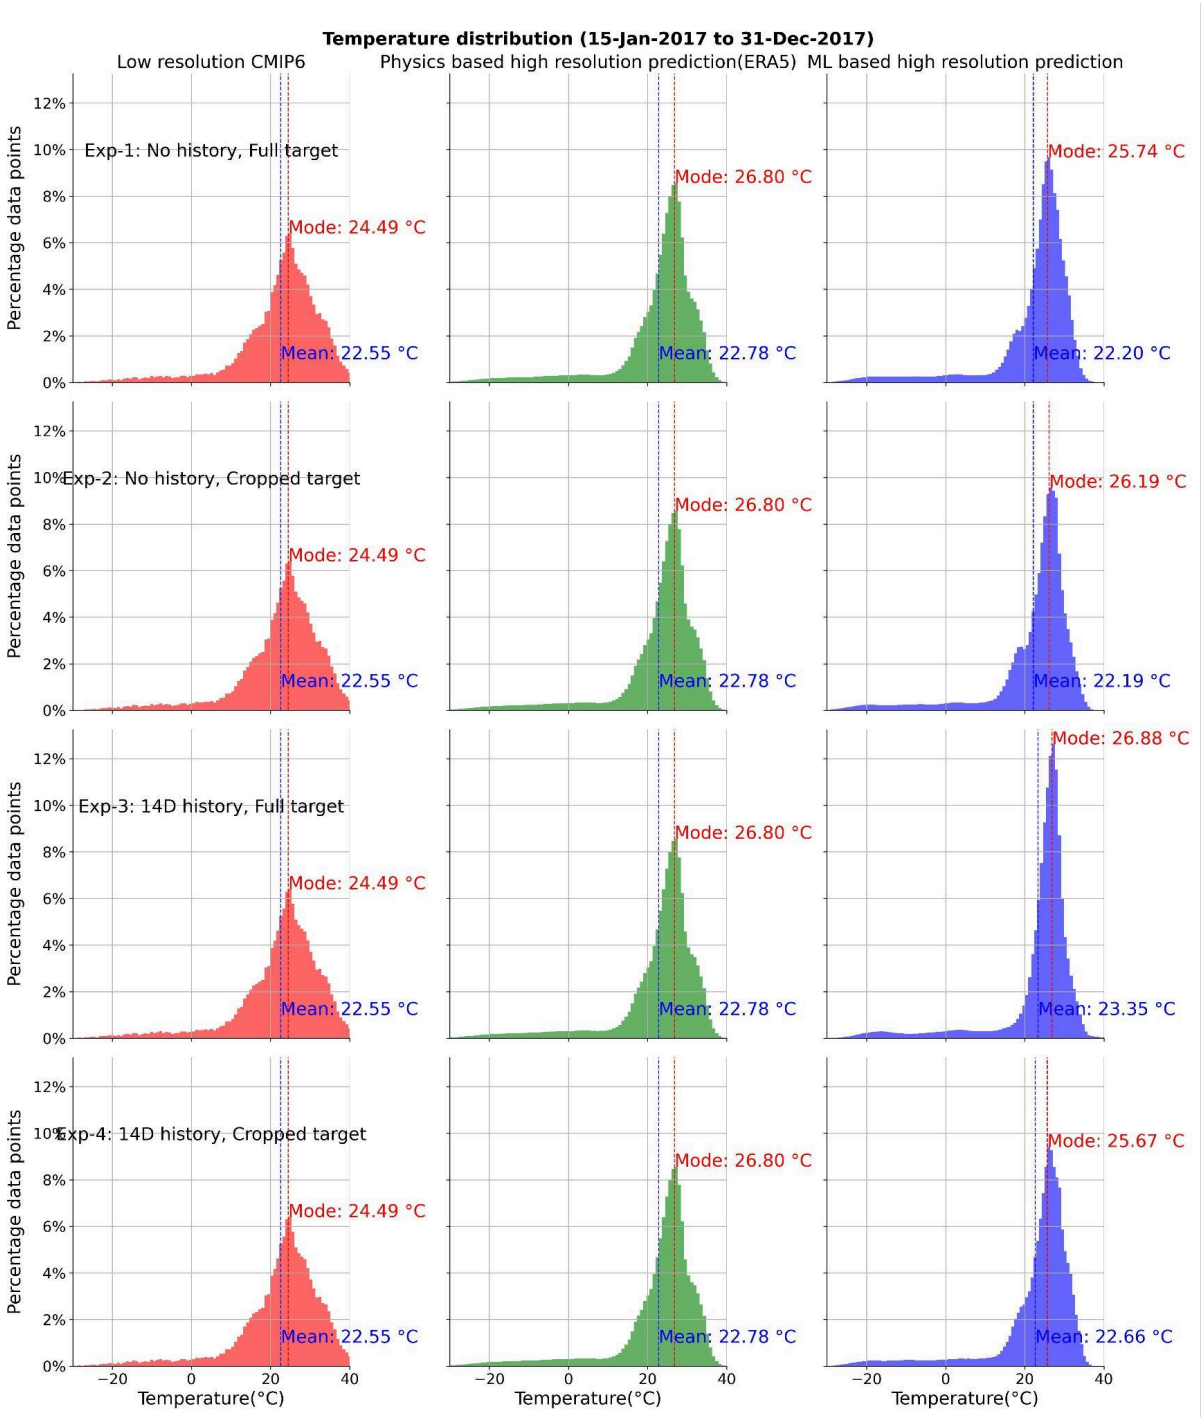

Figure 7: Frequency distribution of daily temperatures over the test period (2017) for the coarse input dataset (CMIP6), fine resolution ERA5 dataset and fine resolution ML model predicted dataset for all experiments.

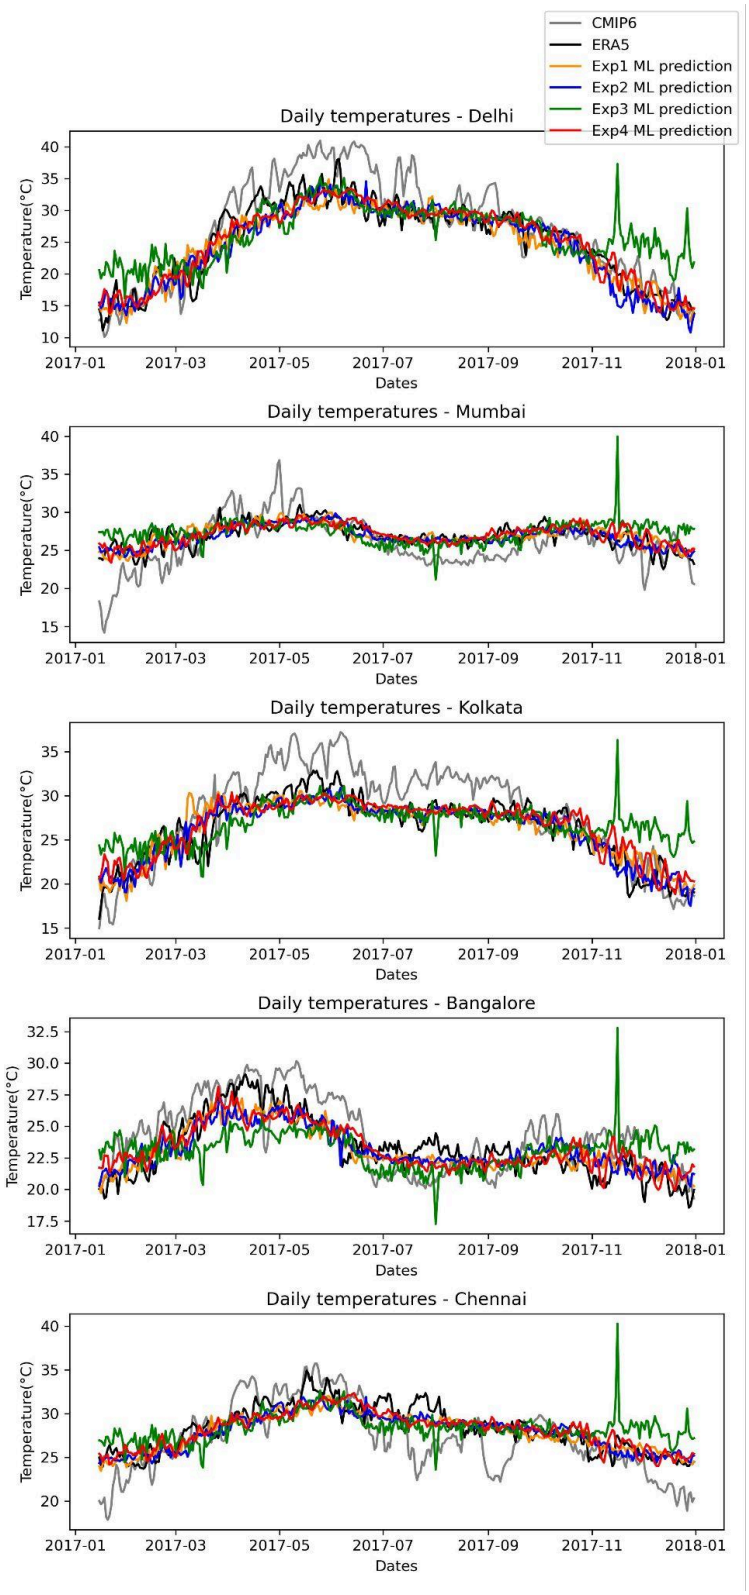

Figure 8: Comparison of ML model predicted daily temperatures for each experiment for five major Indian cities against daily temperatures from physics-based downscaled data (ERA5) and coarse GCM data (CMIP6) over the test period (2017).

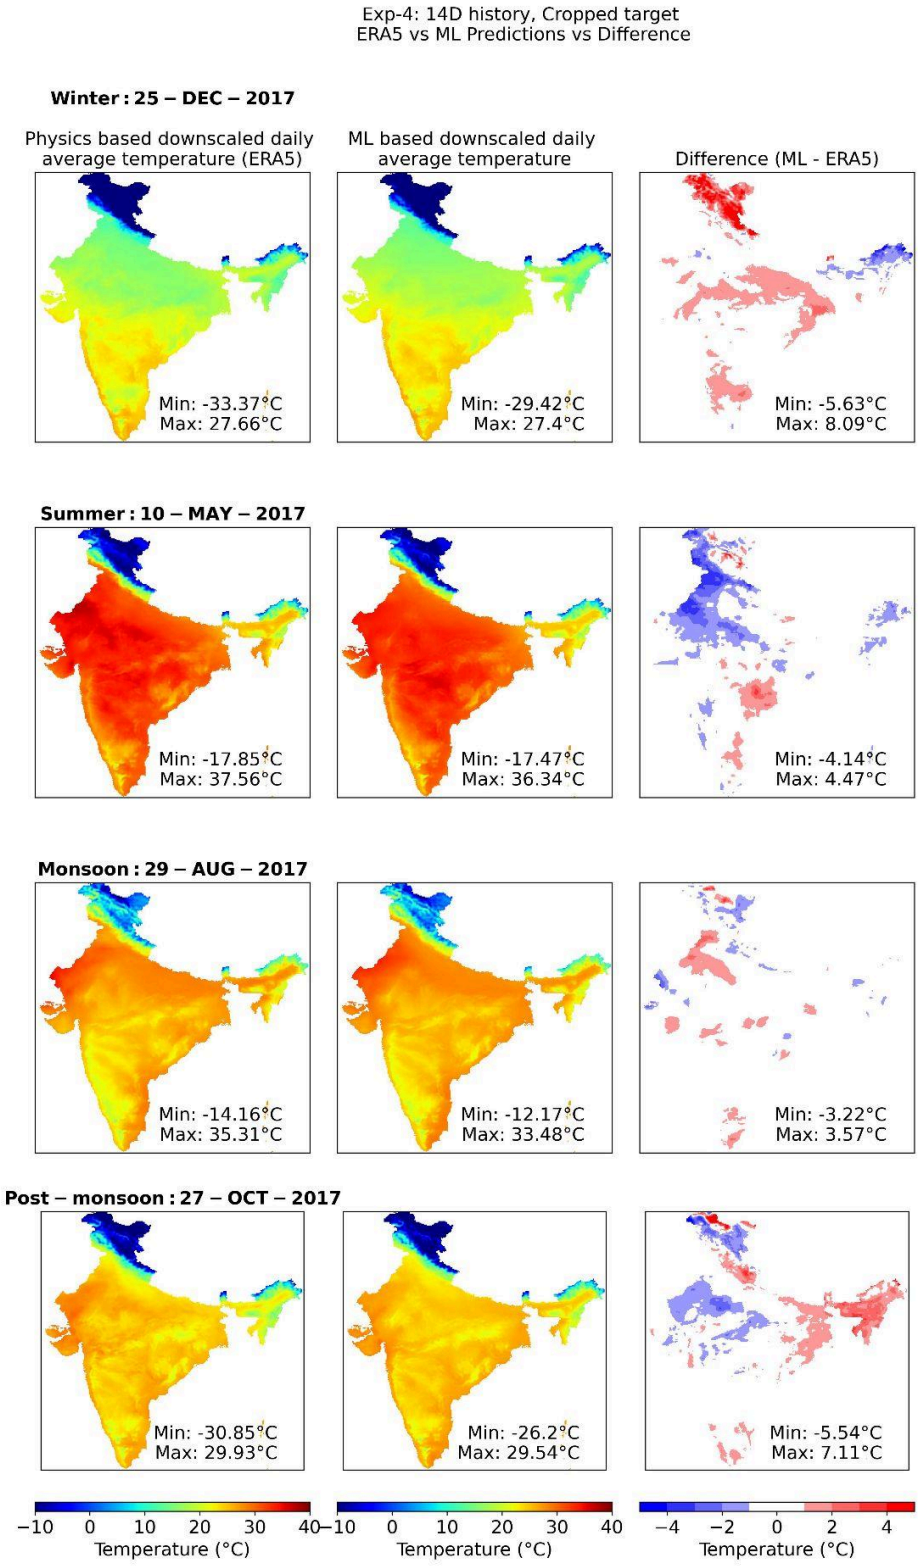

Figure 9: Geographical map of surface temperature from the reference data ERA5 (left column), ML model (experiment 4) predicted surface temperature (middle column) and the difference between the two (right column) for four representative days for each season.

5. Results and discussion

After a robust model evaluation process as described in section 4, we concluded that the model design of experiment 4 provides the best overall regional and temporal prediction. In this section, we therefore focus on key results from this selected model design as per experiment 4. Figure 10 presents our principal result showing seasonal average daily temperature maps of India for a future year (2030) from CMIP6 model data at 1.87° x 1.87° resolution and the ML model downscaled data at 0.1°x0.1° horizontal resolution. Here, for both datasets, the daily temperatures have been averaged over different months (winter, Dec-Feb; summer, Mar-May; monsoon, Jun-Sep; and post-monsoon, Oct-Nov).

The general spatial profile of temperature is similar for both coarse and downscaled data: cold temperatures in the northern region of Kashmir, higher temperatures across the Indo-Gangetic Plains for all seasons except winter, and relatively moderate temperatures over peninsular India. However, there are crucial and important differences too - and this is precisely where the downscaling provides added value. After downscaling, in winter, we see higher temperatures in central India (~20°C) as compared to the native CMIP6 data (~15°C). In general the country is warmer in winters after downscaling. In the summer season, the Indo Gangetic Plain region becomes less warm with temperatures around 25°C after downscaling as compared to CMIP6 data which shows temperatures around 30°C for that region. The downscaled data also successfully captures the relatively low temperatures in mountainous areas such as the Western Ghats in southern India which remains unresolved in the CMIP6 data. For the monsoon season also, the downscaled data shows a significantly lower temperature for much of the Indo Gangetic Plain as compared to CMIP6 data except the Thar desert region. We further investigated the source of this difference and found that the principal difference arises in the month of June which still resembles the summer conditions more than monsoon conditions in northern India. In post-monsoon, the downscaled data shows a relatively colder Indo Gangetic Plain and a relatively warmer central India as compared to the CMIP data.

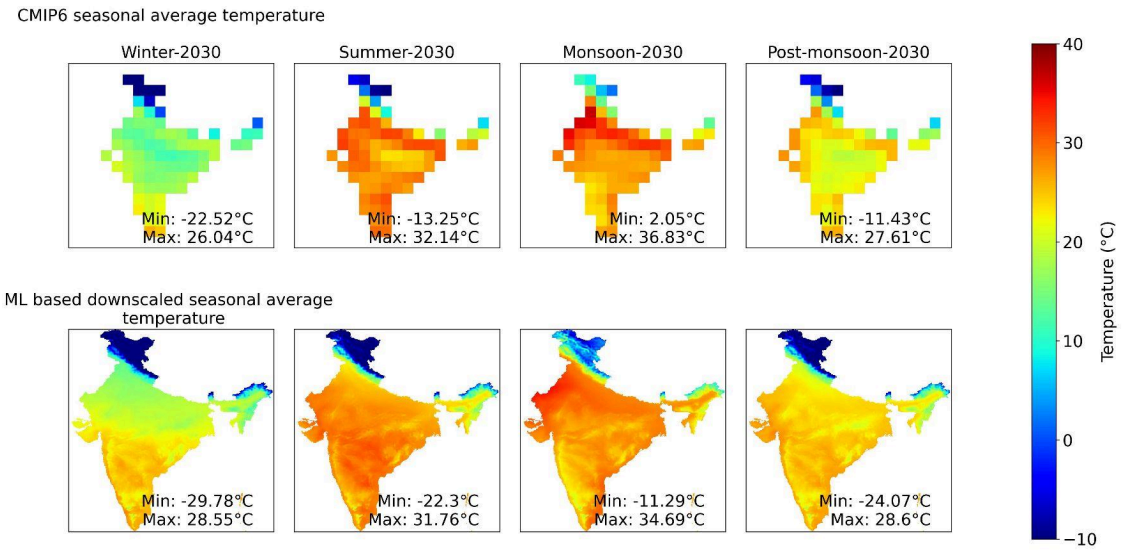

Figure 10: Geographical maps showing seasonal average surface temperatures for coarse resolution input data from CMIP6 model (top row) and ML model downscaled fine resolution surface temperatures (bottom row) for the year 2030.

Here, we have only summarised the broad differences between the coarse and downscaled data. Clearly, the downscaled data from the ML model provides a lot more local detail than the CMIP6 data. Already, these

broad-pattern temperature differences before and after downscaling denote a significant shift in local implications of these future temperatures. Such detailed downscaled temperature information can empower local authorities and communities to take appropriate decisions at the local and regional level. An important result is that it shows that the highly populated Indo-Gangetic Plain will not warm up as drastically by 2030 as shown in the coarse resolution CMIP6 data. However, we remain very cautious to over-interpret this result. This is because we have only utilised a short recent period of 2018 to mid-2022 in order to build our ML model. Perhaps a much longer training dataset (which we could not utilise due to computational and storage constraints) might lead to answer-changing differences in the final downscaled result. This is also the reason why we did not choose to downscale a much farther future period such as 2050 or 2060 when many countries have pledged to reach national net-zero carbon emissions. Our ML model can easily furnish downscaled data for a much farther future period such as 2050 or 2100 or until whenever the CMIP6 data is available however, we believe that the relatively short training period of the ML model might not be sufficient to incorporate the long-term climatic trends such as the El Nino Southern Oscillation (ENSO), the Quasi-biennial Oscillation (QBO) etc. which ought to be included when downscaling temperatures for the far-future.

After analysing the spatial distribution of the ML-model downscaled data for different seasons for the year 2030, we further explored the overall frequency distribution of the downscaled temperature data as shown in Figure 11. Here, we note that the general transformation of the temperature histogram from CMIP6 to downscaled data is similar to the transformation seen during the test period (2017; Figure 7) where the modal value is slightly increased and the overall distribution is peaked higher than the CMIP6 distribution. We note some key differences between the input data and the downscaled data here: the mean temperature over the entire country across all seasons increases from 21.99°C to 22.28°C, an increase of almost 0.3°C and the modal temperature increases by more than 1°C (24.52°C to 25.57°C). This is a significant result and shows the importance of spatial downscaling on not only local temperatures but also national average temperatures which often become the drivers for climate policies.

Furthermore, we also note that the frequency of the upper extreme values reduces after downscaling. For example, the number of data points above a threshold of, say, 35°C are much lower in the downscaled dataset as compared to the coarse dataset. This mostly is due to the relatively lower temperatures in the downscaled data along the Indo Gangetic Plain in the summer months.

There is no true way of perfectly evaluating the validity of these downscaled results for the future - we will have to wait for another 7 years until the actual ERA5 dataset is available for the year 2030 for a true performance evaluation of our ML-model results for the future. Still, we argue that these downscaled results are an improvement over the coarse CMIP6 model output since they incorporate a similar data transformation from coarse to fine grid as learnt from the present-day datasets.

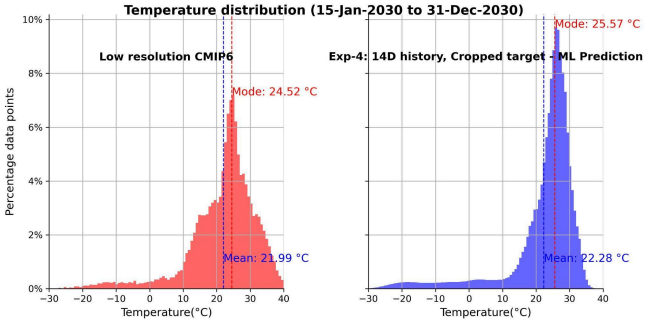

Figure 11: Frequency distribution of surface temperatures for the year 2030 from low resolution CMIP6 model and high resolution ML downscaled model output.

Finally, we analyse future temperatures of five major Indian cities from the downscaled data sampled from the nearest grid point and plot them against the corresponding temperatures from the CMIP6 data. Figure 12 shows the time series plots of the daily temperatures in these cities for the year 2030 before and after downscaling. We find large differences in the annual temperature cycles of all the cities when compared to the coarse dataset values. In general, the annual temperature range is shrunk for all the cities after downscaling. For Delhi and Kolkata, the downscaled pre-monsoon and summer temperatures are 5-10°C lower than the CMIP6 temperatures, which is a hopeful result signifying lesser prevalence of heatwaves. For Mumbai, the downscaled temperatures are significantly (3-8°C) warmer than the CMIP6 temperatures for all seasons but the Summer, when they are 1-2°C colder. For Bangalore, the downscaled temperatures are 2-4°C cooler than CMIP6 temperatures until the beginning of the monsoon season after which they are 1-2°C warmer, signifying milder weather in general which is similar to the present Bangalore climate. A similar pattern is seen for Chennai, where the downscaled temperatures are cooler than CMIP6 temperatures in the summer and warmer than CMIP6 temperatures in the winter, representing the maritime climate of the coastal city of Chennai more realistically.

Overall, these analyses (national-spatial, frequency distribution, and city-temporal) highlight the crucial value addition by the downscaling process to the coarse resolution CMIP6 model dataset.

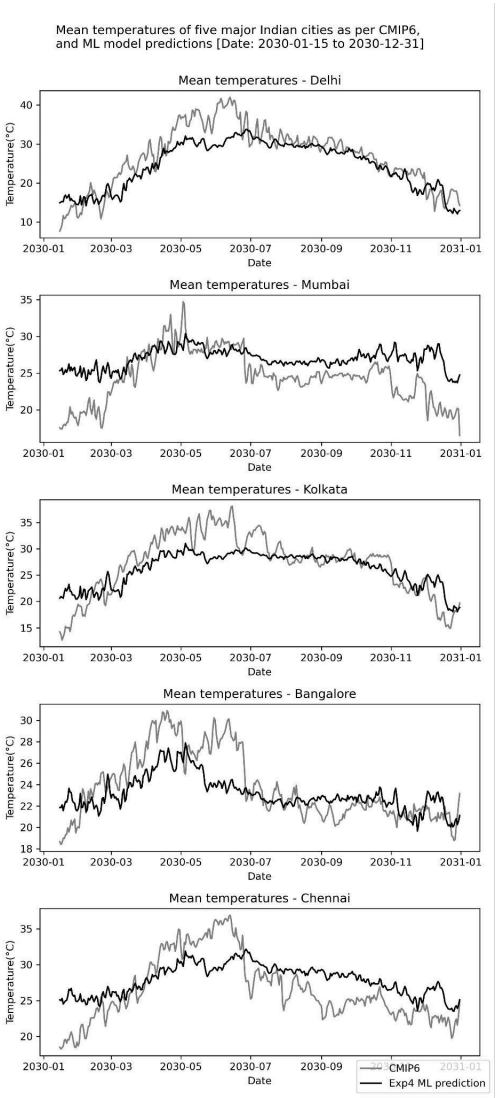

Figure 12: Comparison of daily temperatures in major cities of India as derived from coarse resolution CMIP6 model output and fine resolution ML model downscaled output.

6. Conclusion, limitations and future outlook

In this study we employed various neural network architectures to build a deep learning model for climate downscaling. Choosing India as our study region, we utilised coarse-grid model output from a CMIP6 candidate model as input feed to map fine-grid model outputs from the ERA5-Land dataset as target in the neural network models. We tested several model designs differing in 1) the type of correlations that could be learnt: same-day spatial correlations, and spatiotemporal correlations with historical data, and 2) the size of the target grid by including and excluding the climatically distinct Tibetan Plateau region from the rest of the Indian subcontinent. We performed a robust model evaluation for all model architectures which consisted of a range of tests including simple scatterplot analyses for all designs, to RMSE and MAE time-series, and more detailed temperature error threshold-based tests, performed both temporally, and spatially (annually and season-based). We also analysed the temperature frequency distributions of the predicted data against those of input data. We further evaluated the model performance for individual grid points representing major cities. We found that the model design with 14-day historical data input from the CMIP6 grid and a cropped target grid without the Tibetan Plateau (Experiment 4) showed the most accurate results.

Using the trained model from our Experiment 4, which showed best performance overall, we generated future downscaled temperatures over India for the year 2030 by inputting future Shared Socioeconomic Pathway 5 scenario (SSP5: business-as-usual fossil-fuel driven world) data from CMIP6. It is important to note that these downscaled future temperatures do not represent simply a downscaled version of the SSP5 scenario over our study region, but they represent the most plausible conditions on ground in 2030 reflecting a continuation of current emissions trajectory until 2030. This is because during the model training, the SSP5 scenario CMIP6 data were related to the actual present-day data of ERA5 over the 2017-2022 period. Had we trained our model using historical CMIP6 data without any SSP forcing against the ERA5 data and then supplied the SSP5 CMIP6 data for 2030 to the trained model, we would then expect to obtain downscaled temperatures reflecting SSP5 conditions over the country. This nuance is very important to appreciate before interpreting our results. We trained our models over a relatively short period of 4 years mainly due to practical and hardware constraints as our primary motivation was to simply test the general applicability of neural networks to the problem of grid-to-grid downscaling. The primary goal of this study was to demonstrate the application of CNN/RNN models for climate downscaling, particularly over a developing region of the Global South where high-res physics-based climate downscaling products are still lacking. This study adds to the growing body of literature on climate downscaling using deep learning and is one of the first to tackle grid-to-grid downscaling over India. Through this work, we have elucidated the process of applying deep learning and selecting the most optimal neural network model architecture for accurate downscaling of daily temperatures. We hope that other groups, particularly those with more hardware resources, will take up this challenge and further improve ML-based downscaling by utilising more training data (ideally full 70 years of the ERA5 dataset) as well as testing new machine learning algorithms particularly suited to downscaling extremes rather than means. Our study is a first step in the direction of providing computationally efficient yet accurate and reliable climate services to the citizens of the Global South.

Appendix

A.1 Details of deep learning models built

Model name	CNN (Exp 1: No history - Uncropped target)	CNN (Exp 2: No history - Cropped target)
Trainable parameters	147,315,409	75,119,042
Input size	(38, 27, 4)	(38, 27, 4)
Output size	(57617)	(28750)
Regularization	Activity regularization, Dropout	Activity regularization, Dropout
Batch Normalization	Yes	Yes
Loss function	MSE	MSE
Dense activation	ReLU, Linear	ReLU, Linear
Dense layers	[2500]	[2500]
Convolutional filters	16, 32	16, 32
Kernel size	(3,3)	(3,3)
Convolutional activations	ReLU	ReLU
Pooling	Max Pooling 2D	Max Pooling 2D

Table A.1 : Details of model designs for Experiment 1 and 2 - CNN based models

Model name	ConvLSTM (Exp 3: 14 Day History - Uncropped Target)	ConvLSTM (Exp 4: 14 Day History - Cropped Target)
Trainable parameters	59,022,721	30,126,854
Input size	(14, 38, 27, 4)	(14, 38, 27, 4)
Output size	(57617)	(28750)
Regularization	Recurrent dropout, Dropout	Recurrent dropout, Dropout
Batch Normalization	Yes	Yes
Loss function	MSE	MSE
Dense activation	ReLU, Linear	ReLU, Linear
Dense layers	[1000]	[1000]
Convolutional LSTM filters	16, 32	16,32
Kernel size	(3,3)	(3,3)
Convolutional activations	ReLU	ReLU
Pooling	Max Pooling 3D, Max Pooling 2D	Max Pooling 3D, Max Pooling 2D

Table A.2: Details of model designs for Experiment 3 and 4 - ConvLSTM based models

A.2 Mathematical explanation for the sequence learning approach

Each input data point to predict the output at timestep t representing the current day can be written as $\{x_{t-13}, x_{t-12}, x_{t-11}, x_{t-10}, \ldots\ldots, x_{t-1}, x_t\}$. Each element x_i within this data point is 3-dimensional (LatX × LonX × CMIPVar). Each input sequence data point is therefore a tensor with dimensions (Timestep × LatX × LonX × CMIPVar). We can represent these tensor data points as:

$X_t = \{x_{t-13}, x_{t-12}, x_{t-11}, x_{t-10}, \ldots\ldots, x_{t-1}, x_t\}$

The problem can be mathematically represented as finding a function F that maps the sequence input to a single predicted output at timestep t, y_t as:

$F : \{x_{t-13}, x_{t-12}, x_{t-11}, x_{t-10}, \ldots\ldots, x_{t-1}, x_t\} \rightarrow y_t$

y_t is a high resolution temperature grid with dimensions (LatY × LonY), where (LatY × LonY) are at a higher resolution than (LatX × LonX).

The ConvLSTM model is a function approximator trying to find:

$y_t = F(x_{t-13}, x_{t-12}, x_{t-11}, x_{t-10}, \ldots\ldots, x_{t-1}, x_t) = F(X_t)$

A.3 General equations for a Convolutional-LSTM

Equations describing a typical Convolutional-LSTM are mentioned below. More details can be found in Shi et al., (2015).

$$i_t = \sigma(W_{xi} * X_t + W_{hi} * H_{t-1} + W_{ci} \circ C_{t-1} + b_i)$$

$$f_t = \sigma(W_{xf} * X_t + W_{hf} * H_{t-1} + W_{cf} \circ C_{t-1} + b_f)$$

$$C_t = f_t \circ C_{t-1} + i_t \circ tanh(W_{xc} * X_t + W_{hc} * H_{t-1} + b_c)$$

$$o_t = \sigma(W_{xo} * X_t + W_{ho} * H_{t-1} + W_{co} \circ C_t + b_o)$$

$$H_t = o_t \circ tanh(C_t)$$

Inputs	:	$X_1, X_2, \ldots\ldots X_t$
Cell outputs	:	$C_1, C_2, \ldots\ldots C_t$
Hidden states	:	$H_1, H_2, \ldots\ldots H_t$
Input, forget, and output gates	:	i_t, f_t, o_t
Convolutional operations	:	$*$
Hadamard product	:	$\circ$

Libraries and Packages used

We worked with the NetCDF climate data using the Python package Xarray. We built and trained the CNN and ConvLSTM models using the python deep learning API Keras that runs on top of Tensorflow. We used the python library Numpy for numerical computations and Matplotlib for making plots.

Data and Code Availability

The trained ML model files as well as the model output will be made publicly available through a public link upon final peer-reviewed publication of the study. They can currently be procured from the authors upon reasonable request.

Author contributions

TA conceived the idea. DA elucidated and designed the experiments. DA performed the model construction, evaluation, post-processing and visualisation with inputs from TA. DA and TA together interpreted the results and wrote the manuscript.

Acknowledgements

DA and TA would like to thank the Copernicus Atmosphere Data Store for making the model training data publicly available and Google Collaboratory for providing free cloud GPU access for model training.

References

Allen, M.R., Dube, O.P., Solecki, W., Aragón-Durand, F., Cramer, W., Humphreys, S., Kainuma, M., Kala, J., Mahowald, N., Mulugetta, Y., Perez, R., Wairiu, M., & Zickfeld, K. (2018). Masson-Delmotte, V., Zhai, P., Pörtner, H.-O., Roberts, D., Skea, J., Shukla, P.R., Pirani, A., Moufouma-Okia, W., Péan, C., Pidcock, R., Connors, S., Matthews, J.B.R., Chen, Y., Zhou, X., Gomis, M.I., Lonnoy, E., Maycock, T., Tignor, M., & Waterfield, T. (Eds.), Framing and Context. In: Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty (pp. 49-92). Cambridge University Press. https://doi.org/10.1017/9781009157940.003

Ansari, D., Verma, P., & Ansari, T. (2023).Towards retrieving aerosol chemical composition from temporal variations of total PM mass concentrations: Theoretical approach, insights, and the promise of machine learning technique. Journal of Research in Atmospheric Science (JRAS). http://dx.doi.org/10.29228/resatmsci.71191

Ansari, T.U., Ojha, N., Chandrasekar, R. et al. Competing impact of anthropogenic emissions and meteorology on the distribution of trace gases over Indian region. J Atmos Chem 73, 363–380 (2016). https://doi.org/10.1007/s10874-016-9331-y

Ashok, K., Guan, Z., & Yamagata, T. (2001). Impact of the Indian Ocean dipole on the relationship between the Indian monsoon rainfall and ENSO. Geophysical Research Letters, 28(23), 4499-4502. https://doi.org/10.1029/2001GL013294

Baño-Medina, J., Manzanas, R. & Gutiérrez, J.M. On the suitability of deep convolutional neural networks for continental-wide downscaling of climate change projections. Clim Dyn 57, 2941–2951 (2021). https://doi.org/10.1007/s00382-021-05847-0

Christensen, J.H., Christensen, O.B. A summary of the PRUDENCE model projections of changes in European climate by the end of this century. Climatic Change 81 (Suppl 1), 7–30 (2007). https://doi.org/10.1007/s10584-006-9210-7

Damiani, A., Ishizaki, N.N., Sasaki, H. et al. Exploring super-resolution spatial downscaling of several meteorological variables and potential applications for photovoltaic power. Sci Rep 14, 7254 (2024). https://doi.org/10.1038/s41598-024-57759-8

Dimri, A. P., D. Niyogi, A. P. Barros, J. Ridley, U. C. Mohanty, T. Yasunari, and D. R. Sikka (2015), Western Disturbances: A review, Rev. Geophys., 53, 225–246. doi:10.1002/2014RG000460.

Dudhia, J., 1993: A Nonhydrostatic Version of the Penn State–NCAR Mesoscale Model: Validation Tests and Simulation of an Atlantic Cyclone and Cold Front. Mon. Wea. Rev., 121, 1493–1513, https://doi.org/10.1175/1520-0493(1993)121<1493:ANVOTP>2.0.CO;2.

Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., and Taylor, K. E.: Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization, Geosci. Model Dev., 9, 1937–1958, https://doi.org/10.5194/gmd-9-1937-2016, 2016.

Fox-Kemper, B., H.T. Hewitt, C. Xiao, G. Aðalgeirsdóttir, S.S. Drijfhout, T.L. Edwards, N.R. Golledge, M. Hemer, R.E. Kopp, G. Krinner, A. Mix, D. Notz, S. Nowicki, I.S. Nurhati, L. Ruiz, J.-B. Sallée, A.B.A. Slangen, and Y. Yu, 2021: Ocean, Cryosphere and Sea Level Change. In Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 1211–1362, doi: 10.1017/9781009157896.011.

Frame, D.J., Rosier, S.M., Noy, I. et al. Climate change attribution and the economic costs of extreme weather events: a study on damages from extreme rainfall and drought. Climatic Change 162, 781–797 (2020). https://doi.org/10.1007/s10584-020-02729-y

Fu, C., and Coauthors, 2005: Regional Climate Model Intercomparison Project for Asia. Bull. Amer. Meteor. Soc., 86, 257–266, https://doi.org/10.1175/BAMS-86-2-257.

Gadgil, S., Joseph, P.V. On breaks of the Indian monsoon. J Earth Syst Sci 112, 529–558 (2003). https://doi.org/10.1007/BF02709778

Gadgil, S. The Indian monsoon. Reson 11, 8–21 (2006). https://doi.org/10.1007/BF02855775

Gadgil, S. The monsoon system: Land–sea breeze or the ITCZ?. J Earth Syst Sci 127, 1 (2018). https://doi.org/10.1007/s12040-017-0916-x

Gelaro, R., and Coauthors, 2017: The Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2). J. Climate, 30, 5419–5454, https://doi.org/10.1175/JCLI-D-16-0758.1.

Ham, YG., Kim, JH. & Luo, JJ. Deep learning for multi-year ENSO forecasts. Nature 573, 568–572 (2019). https://doi.org/10.1038/s41586-019-1559-7

Hansen, J., Sato, M., Ruedy, R., Lo, K., Lea, D. W., & Medina-Elizade, M. (2006). Global temperature change. Proceedings of the National Academy of Sciences, 103(39), 14288-14293. https://doi.org/10.1073/pnas.0606291103

Hersbach H, Bell B, Berrisford P, et al. The ERA5 global reanalysis. Q J R Meteorol Soc. 2020; 146: 1999–2049. https://doi.org/10.1002/qj.3803

Jacob, D., Petersen, J., Eggert, B. et al. EURO-CORDEX: new high-resolution climate change projections for European impact research. Reg Environ Change 14, 563–578 (2014). https://doi.org/10.1007/s10113-013-0499-2

Kim J, Lee M, Han H, Kim D, Bae Y, Kim HS. Case Study: Development of the CNN Model Considering Teleconnection for Spatial Downscaling of Precipitation in a Climate Change Scenario. Sustainability. 2022; 14(8):4719. https://doi.org/10.3390/su14084719

Kimura N, Ishida K, Baba D. Surface Water Temperature Predictions at a Mid-Latitude Reservoir under Long-Term Climate Change Impacts Using a Deep Neural Network Coupled with a Transfer Learning Approach. Water. 2021; 13(8):1109. https://doi.org/10.3390/w13081109

Kovats, R. S., & Koppe, C. (2005). Heatwaves: past and future impacts on health. Integration of public health with adaptation to climate change: lessons learned and new directions, 136-160.

Li, X., Li, Z., Huang, W. et al. Performance of statistical and machine learning ensembles for daily temperature downscaling. Theor Appl Climatol 140, 571–588 (2020). https://doi.org/10.1007/s00704-020-03098-3

Lin, H., Tang, J., Wang, S. et al. Deep learning downscaled high-resolution daily near surface meteorological datasets over East Asia. Sci Data 10, 890 (2023). https://doi.org/10.1038/s41597-023-02805-9

Liu, Y., Ganguly, A. R., & Dy, J. (2020). Climate Downscaling Using YNet: A Deep Convolutional Network with Skip Connections and Fusion. Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20), 3145–3153. https://doi.org/10.1145/3394486.3403366

Maraun, D., et al. (2010), Precipitation downscaling under climate change: Recent developments to bridge the gap between dynamical models and the end user, Rev. Geophys., 48, RG3003, doi:10.1029/2009RG000314.

Miao Q, Pan B, Wang H, Hsu K, Sorooshian S. Improving Monsoon Precipitation Prediction Using Combined Convolutional and Long Short Term Memory Neural Network. Water. 2019; 11(5):977. https://doi.org/10.3390/w11050977

Noah S. Diffenbaugh et al. ,Unprecedented climate events: Historical changes, aspirational targets, and national commitments.Sci. Adv.4,eaao3354(2018).DOI:10.1126/sciadv.aao3354

Nori-Sarma, A., Anderson, G. B., Rajiva, A., ShahAzhar, G., Gupta, P., Pednekar, M. S., Son, J.-Y., Peng, R. D., & Bell, M. L. (2019). The impact of heat waves on mortality in Northwest India. Environmental Research, 176, 108546. https://doi.org/10.1016/j.envres.2019.108546

Nourani, V., Razzaghzadeh, Z., Baghanam, A.H. et al. ANN-based statistical downscaling of climatic parameters using decision tree predictor screening method. Theor Appl Climatol 137, 1729–1746 (2019). https://doi.org/10.1007/s00704-018-2686-z

Oyama, N., Ishizaki, N.N., Koide, S. et al. Deep generative model super-resolves spatially correlated multiregional climate data. Sci Rep 13, 5992 (2023). https://doi.org/10.1038/s41598-023-32947-0

Pal, J. S., and Coauthors, 2007: Regional Climate Modeling for the Developing World: The ICTP RegCM3 and RegCNET. Bull. Amer. Meteor. Soc., 88, 1395–1410, https://doi.org/10.1175/BAMS-88-9-1395.

Peter Stott ,How climate change affects extreme weather events. Science 352, 1517-1518(2016). DOI:10.1126/science.aaf7271

Reichstein, M., Camps-Valls, G., Stevens, B. et al. Deep learning and process understanding for data-driven Earth system science. Nature 566, 195–204 (2019). https://doi.org/10.1038/s41586-019-0912-1

Riahi K, van Vuuren DP, Kriegler E, Edmonds J, O'Neill BC, Fujimori S, Bauer N, Calvin K, Dellink R, Fricko O, Lutz W, Popp A, Crespo Cuaresma J, KC S, Leimbach M, Jiang L, Kram T, Rao S, Emmerling J, Ebi K, Hasegawa T, Havlik P, Humpenöder F, Da Silva LA, Smith S, Stehfest E, Bosetti V, Eom J, Gernaat D, Masui T, Rogelj J, Strefler J, Drouet L, Krey V, Luderer G, Harmsen M, Takahashi K, Baumstark L, Doelman JC, Kainuma M, Klimont Z, Marangoni G, Lotze-Campen H, Obersteiner M, Tabeau A, Tavoni M. The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. Global Environmental Change. 2017;42:153-168. ISSN 0959-3780. https://doi.org/10.1016/j.gloenvcha.2016.05.009.

Robert J. Nicholls, Anny Cazenave ,Sea-Level Rise and Its Impact on Coastal Zones. Science 328, 1517-1520 (2010). DOI:10.1126/science.1185782

Rohini, P., Rajeevan, M. & Srivastava, A. On the Variability and Increasing Trends of Heat Waves over India. Sci Rep 6, 26153 (2016). https://doi.org/10.1038/srep26153

Sachindra, D. A., Ahmed, K., Rashid, M. M., Shahid, S., & Perera, B. J. C. (2018). Statistical downscaling of precipitation using machine learning techniques. Atmospheric Research, 212, 240-258. ISSN: 0169-8095. https://doi.org/10.1016/j.atmosres.2018.05.022

Samset, B.H., Zhou, C., Fuglestvedt, J.S. et al. Steady global surface warming from 1973 to 2022 but increased warming rate after 1990. Commun Earth Environ 4, 400 (2023). https://doi.org/10.1038/s43247-023-01061-4

Seneviratne, S. I., & Hauser, M. (2020). Regional climate sensitivity of climate extremes in CMIP6 vs CMIP5 multi-model ensembles. Earth's Future. 8, e2019EF001474. https://doi.org/10.1029/2019EF001474

Sharma, S., Mujumdar, P. Increasing frequency and spatial extent of concurrent meteorological droughts and heatwaves in India. Sci Rep 7, 15582 (2017). https://doi.org/10.1038/s41598-017-15896-3

Skamarock, W. C., & Klemp, J. B. (2008). A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. Journal of Computational Physics, 227(7), 3465-3485. https://doi.org/10.1016/j.jcp.2007.01.037.

Takle, E. S., et al. (1999), Project to Intercompare Regional Climate Simulations (PIRCS): Description and initial results, J. Geophys. Res., 104(D16), 19443–19461, doi:10.1029/1999JD900352.

Taylor, K. E., R. J. Stouffer, and G. A. Meehl, 2012: An Overview of CMIP5 and the Experiment Design. Bull. Amer. Meteor. Soc., 93, 485–498, https://doi.org/10.1175/BAMS-D-11-00094.1.

Vandal, T., Kodra, E., Ganguly, S., Michaelis, A., Nemani, R., & Ganguly, A. R. (2017). DeepSD: Generating High Resolution Climate Change Projections through Single Image Super-Resolution. Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17), 1663-1672. https://doi.org/10.1145/3097983.3098004

Wang, Q., Huang, J., Liu, R., Men, C., Guo, L., Miao, Y., Jiao, L., Wang, Y., Shoaib, M., & Xia, X. (2020). Sequence-based statistical downscaling and its application to hydrologic simulations based on machine learning and big data. Journal of Hydrology, 586, 124875. ISSN: 0022-1694. https://doi.org/10.1016/j.jhydrol.2020.124875

Warner TT. Numerical solutions to the equations. In: Numerical Weather and Climate Prediction. Cambridge University Press; 2010b:17-118.

Warner TT. The governing systems of equations. In: Numerical Weather and Climate Prediction. Cambridge University Press; 2010a:6-16.

Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. 2015. Convolutional LSTM Network: a machine learning approach for precipitation nowcasting. In Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1 (NIPS'15). MIT Press, Cambridge, MA, USA, 802–810.