An LSTM-based Downscaling Framework for Australian Precipitation Projections

Matthias Bittner *TU Wien Vienna, Austria matthias.bittner@tuwien.ac.at

Muhammad Zawish Walton Institute, South East Technological University, Ireland muhammad.zawish@waltoninstitute.ie Sanaa Hobeichi The University of New South Wales Sydney, Australia s.hobeichi@unsw.edu.au

> Samo Diatta Assane Seck University of Ziguinchor, Sénégal samo.diatta@univ-zig.sn

Remigious Ozioko
University of Nigeria
Nsukka, Nigeria
remigius.ozioko@unn.edu.ng

Sharon Xu Indigo AG Boston, United States sxu@indigoag.com Axel Jantsch TU Wien Vienna, Austria axel.jantsch@tuwien.ac.at

Abstract

Understanding potential changes in future rainfall and their local impacts on Australian communities can inform adaptation decisions worth billions of dollars in insurance, agriculture, and other sectors. This understanding relies on downscaling a large ensemble of coarse Global Climate Models (GCMs), our primary tool for simulating future climate. However, the prohibitively high computational cost of downscaling has been a significant barrier. In response, this study develops a costefficient downscaling framework for daily precipitation using Long Short-Term Memory (LSTM) models. The models are trained with ERA5 reanalysis data and a customized quantile loss function to better capture precipitation extremes. The framework is employed to downscale precipitation from a GCM member of the CMIP6 ensemble. We demonstrate the skills of the downscaling models to capture spatial and temporal characteristics of precipitation. We also explore regional future changes in precipitation extremes projected by the downscaled GCM. In general, this framework will enable the generation of a large ensemble of regional future projections for Australian rainfall. This will further enhance the assessment of likely climate risks and the quantification of their uncertainties.

1 Introduction

There is a growing demand for regional precipitation *data* for the near, medium, and far future. This data is crucial for understanding the potential impacts of changes in precipitation on various human and natural sectors and for developing adaptation strategies tailored to the unique climate risks of specific regions. In particular, understanding the spatial patterns of changes in extreme rainfall events is essential to enhancing flood or drought resilience and mitigating unexpected economic shocks. Reliable projections of trends and variability in future rainfall are fundamental for assessing climate risks, water resources scarcity, and linked systems such as hydropower systems. These systems play a

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2023.

^{*}Christian Doppler Laboratory for Embedded Machine Learning

pivotal role in achieving Australia's greenhouse gas emission reduction targets set out in the Climate Change Act 2022 [1].

The current provision of regional climate projection data over Australia relies on using Regional Climate Models (RCMs) to downscale Global Climate Models (GCMs). They use outputs from the multi-model ensemble of the Coupled Model Intercomparison Project version 6 (CMIP6) [2] from a native coarse resolution of 50-250km to produce a much finer 25km resolution. However, due to the prohibitively high computational costs associated with running RCMs, current initiatives have primarily focused on downscaling a narrow selection of potential GCMs and only two out of five Shared Socioeconomic Pathways (SSPs) [3]. As a result, the downscaled projections will only partially represent the range of potential future climates. Additionally, the constrained size of the downscaled ensemble hampers the ability to accurately quantify uncertainties in projected changes, a critical component for thorough climate risk assessment [4].

Empirical downscaling, which includes statistical methods and machine learning techniques, has emerged as a cost-effective method for downscaling GCMs [5]. The most common approach, known as *perfect prognosis*, establishes an empirical relationship between coarsened reanalysis/observation (matched to the GCM's resolution) and a finer reanalysis/observation of the target variable [6]. This relationship is then applied to downscale GCMs. Though many studies highlight machine learning's superior downscaling over traditional statistical methods [7], no specific machine learning algorithm has been identified as the best. Empirical downscaling using machine learning has employed a diverse range of algorithms, from simple linear regression to foundation models [8], encompassing both conventional [9–11] and deep learning techniques [12–18].

This work develops the first machine learning-based downscaling framework for Australia, capable of generating a large ensemble of regional future precipitation projections. We demonstrate its effectiveness by downscaling a single GCM from the CMIP6 ensemble and plan to extend this to all available GCMs, covering all five SSPs. The new framework creates a grid-specific recurrent Long Short-Term Memory (LSTM) model [19] for each grid cell across the continent while also considering neighboring grid information. The LSTMs are driven by ERA5 reanalysis [20] and trained using a custom loss function that combines conventional loss, i.e. MSE, with a quantile loss optimized for capturing extreme precipitation events [21]. This design addresses four criteria reflecting the distinct physical and topographic nuances of Australian precipitation: (i) Incorporating atmospheric memory into the model to effectively capture recurring weather patterns and persistent weather systems; (ii) Adding spatial awareness from neighboring regions to ensure the model considers influences from nearby topography and shared atmospheric circulation patterns; (iii) Excluding climatic information from distant regions, since this may introduce variability imposed by different climate systems and fail to capture local factors; (vi) Ensuring the framework retains enough flexibility to enhance downscaling in specific regions where performance may be lacking.

2 Data and Methods

Spatial Domain We build a separate downscaling model for each of the 546 coarse grid cells (150km) in the spatial domain (Fig. 8). These models are then used to downscale daily precipitation output from a CMIP6 model to 25km. In this study, we have chosen to downscale CNRM-ESM2-1 [22], an Earth System Model developed by the Centre National de Recherches Météorologiques denoted hereafter by 'GCM'. GCM's data used in this study is listed in Tab. 1 and consists of projections of daily fields under the highest emission scenario SSP5-8.5 and covering the period 1979 to 2070. Predictors from both ERA5 and GCM are regridded to 150km before downscaling. Throughout this paper, a coarse resolution refers to 150km, while a fine resolution denotes 25km.

Predictors Training the downscaling models is driven by ERA5 reanalysis. The predictor set encompasses a range of coarse climate variables listed in Tab. 1. These include daily atmospheric variables at both the surface level and the 850hPa pressure level from the current and all neighboring grid cells (up to 8). Additionally, we add spatial encodings, to help the model capture spatial dependencies between fine grids enclosed within the coarse grids, by adding latitude and longitude in the form of Cartesian coordinates $\mathbf{c}_k = [\cos(|at_k|) \times \cos(|on_k|), \cos(|at_k|) \times \sin(|on_k|), \sin(|at_k|)]$. Similarly, we compute two additional predictors to capture temporal patterns within a year, this enables the model to learn seasonality, which would bring value in areas where seasons have a strong effect on precipitation.



Figure 1: Performance (downscaled vs. ERA5 groundtruth) over the testing period [2010-2020].



Figure 2: Extreme precipitation index, r40mm (Tab.2) calculated over the testing period [2010-2020] based on coarse ERA5, ERA5 groundtruth (ERA5-target) and downscaled ERA5 (ERA5-prediction).

This is achieved according to this equation, $\mathbf{t}_{day} = [\sin(2\pi \times day/366), \cos(2\pi \times day/366)]$. All predictors are first coarsened to a 150km resolution. The predictand, on the other hand, is the fine-resolution daily precipitation from ERA5 at 25km. In this setting, the downscaling model computes daily precipitation fields for 36 subgrids (25km) enclosed within each coarse grid. ERA5 data is split into training [1979-2005], validation [2006-2009], and testing [2010-2020] sets.

Architecture and Loss Function We use an LSTM since it is well-suited for analyzing and predicting sequential data. Additionally, it is lightweight and performs the prediction task in an auto-regressive fashion. Standard MSE loss for such a regression task will be able to capture the main shape of the distribution but will fail to properly capture its tails (Fig. 4). These tails represent extreme precipitation events. To account for extreme events, we add a second statistical/quantile loss term. The combined loss is defined as

$$L(\hat{\mathbf{p}}_{k,1:T}, \mathbf{p}_{k,1:T}) = \alpha \text{MSE}(\hat{\mathbf{p}}_{k,T}, \mathbf{p}_{k,T}) + (1 - \alpha) \text{MSE}(\mathbf{q}(\hat{\mathbf{p}}_{k,1:T}), \mathbf{q}(\mathbf{p}_{k,1:T})),$$
(1)

where $\mathbf{q} \in \mathbb{R}^{n_q}$ is a vector containing the non-uniform spaced n_q -th quantiles of precipitation time series $\mathbf{p}_{k,1:T} \in \mathbb{R}^{1 \times T}$. The factor α is used for weighting, and was set to 0.6. The index k refers to a fine grid cell, and T is the length of the sliding window used for training.

3 Experiments and Results

Model Configuration We use PyTorch [23] for building and training the models. The neural network architecture consists of four LSTM cells with a hidden size of 50, followed by a single fully connected layer that maps the 50 hidden states of the last cell to a scalar prediction output. We train each network for 50 epochs, with a static learning rate of 10^{-3} , standard Adam optimizer, and a batch size of 36. The model is selected by the lowest validation loss. We train it on two Tesla V100S-PCIE-32GB and use sliding windows with a window size of T = 56 days (stride=1 for training and stride=2 for validation). We are able to train 6 models at once and the training for 546 models took ≈ 6.3 days. The evaluation and model inference on the ERA5 testing period and GCM data are performed in an autoregressive fashion.

Performance of downscaling Fig. 1 displays the downscaling performance based on three metrics over the entire testing period [2010-2020]. The results show a consistently high temporal correlation, with the median exceeding 0.9. This high correlation is anticipated given the aptness of LSTMs to learn patterns over time scales, especially since they were trained over a wide time window of 8 weeks.



Figure 3: Precipitation maps for two exemplary days in the testing period corresponding to coarse resolution ERA5 (coarse), fine-resolution ERA5 (targets), and downscaled ERA5 (predictions).

The Root Mean Squared Error (RMSE) is low over land areas but increases in the tropical regions in the North near the coasts, corresponding with the high precipitation in those areas. The mean bias is close to zero across most land areas, but it indicates that downscaling slightly underestimates precipitation over the ocean in the tropics. Fig. 3 illustrates downscaled precipitation (predictions) for two random days in the testing period, the corresponding coarse resolution from the coarsened ERA5 (coarse), and the target fine-resolution from ERA5 (target). Each coarse grid in the "coarse" map was downscaled using a separate LSTM. The maps reveal that downscaling captures both precipitation extremes, which were smoothed out in the coarse map, as well as the spatial variations in precipitation. Additionally, the downscaling is skillful in capturing the spatial characteristics of extreme precipitation and its magnitude across four extreme precipitation indices: rx1day, r10mm, r20mm, and proptot (see Tab. 2 for a full description). This proficiency is underscored by the close match of these extreme indices between the fine-resolution ERA5 (ERA5-target) and the downscaled ERA5 (ERA5-prediction) calculated on the testing period as illustrated in Fig. 2 and Fig.7. The only exception is the downscaling's tendency to underestimate the maximum one-day precipitation (rx1day). Moreover, when comparing coarse ERA5 to the downscaled ERA5, it becomes apparent that downscaling provides added value, revealing regional precipitation extremes that get smoothed out in the coarse resolution.



Figure 4: Probability density functions representing the distribution of precipitation in the testing set (target) and the output of three training settings *mse*, *mse*+*neigh*, and *mse*+*quantile*+*neigh*.

Effectiveness of the loss function To showcase the effectiveness of the quantile loss term and adding the neighborhood predictor set, we define three training settings: (1) *mse*: standard MSE loss, (2) *mse+neigh*: MSE loss with neighbor information, (3) *mse+quantile+neigh*: same as (2), with an additional quantile loss term. LSTM models are trained for each setting on four coarse grids ($i \in [49, 227, 301, 468]$) randomly chosen from different climate zones (Fig. 8). We compare the distribution of the predictions and the targets on the testing set (Fig. 4), and we calculate four metrics of performance as illustrated in Fig. 6. Setting *mse* exhibits the worst performance while

setting *mse+quantile+neigh* achieves the best performance overall. This is indicated by scoring the smallest bias and a low variance, and by its superior capability in capturing the tail of the distribution. Nevertheless, all three settings are able to reproduce the shape of the target distribution (Fig. 4).

Projected changes in precipitation extremes Fig. 5 illustrates the projected changes in four extreme precipitation indices (defined in Tab. 2), for the period [2031-2070] in reference to [1979-2018]. The change signal in a given index is calculated using the formula: $100 \times (\text{future} - \text{historical})/\text{historical}$. In this equation, "future" refers to the precipitation index calculated for the future based on the GCM, while "historical" refers to the precipitation index calculated for the historical period using groundtruth ERA5 data.



Figure 5: Future [2031 - 2070] projected changes by the downscaled GCM in four extreme precipitation indices (Tab. 2) calculated in reference to ERA5 [1979-2018].

4 Discussion

The downscaling framework was designed to cater to a broad spectrum of applications and end-users (e.g., agricultural land managers, and catastrophe modelers). Some users might need data for specific regions to assess particular risks, such as in flood-prone areas, and may require a larger ensemble than what can be obtained by downscaling each GCM member in the CMIP6 ensemble (\approx 50 members). In such instances, LSTMs developed for similar regions can be adapted to collectively generate a more expansive ensemble. The method for identifying similar regions can follow the approach described in [7]. Furthermore, we can incorporate concepts from transfer learning by selecting LSTMs trained for one region and fine-tuning them to downscale another. This latter strategy is worth investigating. If successful, it could facilitate using the LSTMs developed here to downscale new regions outside the Australian domain where computational resources are scarce.

The implemented downscaling approach is lightweight. A single grid cell model has $\approx 96k$ parameters and needs 97.65k FLOPs (Floating Point Operations) for predicting the precipitation for a single day and fine grid cell. Assuming single precision (32 Bytes per parameter), all 546 models occupy 200MB of storage. Considering the case of downscaling the GCM data (1979 - 2070) and solely focusing on the computational requirements occupied by our algorithm, we can downscale the whole continent of Australia in 4.55 seconds (on a Tesla V100S-PCIE with 14×10^{12} FLOPs per second).

The completion of this project (i.e. downscaling all CMIP6 members) will offer the Australian policy-making communities the data needed for climate risk assessment across a broad range of plausible scenarios and future climates, along with a thorough quantification of the uncertainties associated with climate risks. At present, generating such data using RCMs is not feasible.

5 Acknowledgement

This work is supported by the Austrian Federal Ministry for Digital and Economic Affairs, the National Foundation for Research, Technology, and Development, the Christian Doppler Research Association, the Australian Research Council Centre of Excellence for Climate Extremes (CLEX; CE170100023). This research was undertaken with the assistance of resources and services from the National Computational Infrastructure (NCI), which is supported by the Australian Government.

References

- [1] Climate Change Act 2022. [Online]. Available: https://www.legislation.gov.au/Details/ C2022A00037
- [2] V. Eyring, S. Bony, G. A. Meehl, C. A. Senior, B. Stevens, R. J. Stouffer, and K. E. Taylor, "Overview of the coupled model intercomparison project phase 6 (cmip6) experimental design and organization," *Geoscientific Model Development*, vol. 9, no. 5, pp. 1937—1958, 2016.
- [3] K. Riahi, D. P. Van Vuuren, E. Kriegler, J. Edmonds, B. C. O'neill, S. Fujimori, N. Bauer, K. Calvin, R. Dellink, O. Fricko *et al.*, "The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: An overview," *Global environmental change*, vol. 42, pp. 153–168, 2017.
- [4] V. Masson-Delmotte, P. Zhai, S. Pirani, C. Connors, S. Péan, N. Berger, Y. Caud, L. Chen, M. Goldfarb, and P. M. Scheel Monteiro, "Ipcc, 2021: Summary for policymakers. 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," 2021.
- [5] C. Huntingford, E. S. Jeffers, M. B. Bonsall, H. M. Christensen, T. Lees, and H. Yang, "Machine learning and artificial intelligence to aid climate change research and preparedness," *Environmental Research Letters*, vol. 14, no. 12, p. 124007, 2019.
- [6] S. Schubert, "Downscaling local extreme temperature changes in south-eastern australia from the csiro mark2 gcm," *International Journal of Climatology: A Journal of the Royal Meteorological Society*, vol. 18, no. 13, pp. 1419–1438, 1998.
- [7] S. Hobeichi, N. Nishant, Y. Shao, G. Abramowitz, A. Pitman, S. Sherwood, C. Bishop, and S. Green, "Using machine learning to cut the cost of dynamical downscaling," *Earth's Future*, vol. 11, no. 3, p. e2022EF003291, 2023.
- [8] T. Nguyen, J. Brandstetter, A. Kapoor, J. K. Gupta, and A. Grover, "Climax: A foundation model for weather and climate," *arXiv preprint arXiv:2301.10343*, 2023.
- [9] E. Sharifi, B. Saghafian, and R. Steinacker, "Downscaling satellite precipitation estimates with multiple linear regression, artificial neural networks, and spline interpolation techniques," *Journal of Geophysical Research: Atmospheres*, vol. 124, no. 2, pp. 789–805, 2019.
- [10] S. Hadi Pour, S. B. Harun, and S. Shahid, "Genetic programming for the downscaling of extreme rainfall events on the east coast of peninsular malaysia," *Atmosphere*, vol. 5, no. 4, pp. 914–936, 2014.
- [11] X. He, N. W. Chaney, M. Schleiss, and J. Sheffield, "Spatial downscaling of precipitation using adaptable random forests," *Water resources research*, vol. 52, no. 10, pp. 8217–8237, 2016.
- [12] J. Baño-Medina, R. Manzanas, E. Cimadevilla, J. Fernández, J. González-Abad, A. S. Cofiño, and J. M. Gutiérrez, "Downscaling multi-model climate projection ensembles with deep learning (deepesd): contribution to cordex eur-44," *Geoscientific Model Development*, vol. 15, no. 17, pp. 6747–6758, 2022.
- [13] J. González-Abad, J. Baño-Medina, and J. M. Gutiérrez, "Using explainability to inform statistical downscaling based on deep learning beyond standard validation approaches," *arXiv* preprint arXiv:2302.01771, 2023.
- [14] P. Harder, V. Ramesh, A. Hernandez-Garcia, Q. Yang, P. Sattigeri, D. Szwarcman, C. Watson, and D. Rolnick, "Physics-constrained deep learning for downscaling," Copernicus Meetings, Tech. Rep., 2023.
- [15] S. Misra, S. Sarkar, and P. Mitra, "Statistical downscaling of precipitation using long shortterm memory recurrent neural networks," *Theoretical and applied climatology*, vol. 134, pp. 1179–1196, 2018.
- [16] N. Nishant, S. Hobeichi, S. Sherwood, G. Abramowitz, Y. Shao, C. Bishop, and A. Pitman, "Comparison of a novel machine learning approach with dynamical downscaling for australian precipitation," *Environmental Research Letters*, vol. 18, no. 9, p. 094006, 2023.
- [17] N. Rampal, P. B. Gibson, A. Sood, S. Stuart, N. C. Fauchereau, C. Brandolino, B. Noll, and T. Meyers, "High-resolution downscaling with interpretable deep learning: Rainfall extremes over new zealand," *Weather and Climate Extremes*, vol. 38, p. 100525, 2022.

- [18] S. C. M. Sharma and A. Mitra, "Resdeepd: A residual super-resolution network for deep downscaling of daily precipitation over india," *Environmental Data Science*, vol. 1, p. e19, 2022.
- [19] Y. Yu, X. Si, C. Hu, and J. Zhang, "A review of recurrent neural networks: Lstm cells and network architectures," *Neural computation*, vol. 31, no. 7, pp. 1235–1270, 2019.
- [20] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers *et al.*, "The ERA5 global reanalysis," *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999–2049, 2020.
- [21] A. Blanchard, N. Parashar, B. Dodov, C. Lessig, and T. Sapsis, "A multi-scale deep learning framework for projecting weather extremes," in *NeurIPS 2022 Workshop* on *Tackling Climate Change with Machine Learning*, 2022. [Online]. Available: https://www.climatechange.ai/papers/neurips2022/65
- [22] R. Séférian, P. Nabat, M. Michou, D. Saint-Martin, A. Voldoire, J. Colin, B. Decharme, C. Delire, S. Berthet, M. Chevallier *et al.*, "Evaluation of cnrm earth system model, cnrm-esm2-1: Role of earth system processes in present-day and future climate," *Journal of Advances in Modeling Earth Systems*, vol. 11, no. 12, pp. 4182–4227, 2019.
- [23] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems* 32. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: http://papers.neurips. cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf
- [24] M. C. Peel, B. L. Finlayson, and T. A. McMahon, "Updated world map of the köppen-geiger climate classification," *Hydrology and earth system sciences*, vol. 11, no. 5, pp. 1633–1644, 2007.

A Appendix

A.1 Data Availability

The ERA5 reanalysis data used is the latest generation of reanalysis of the ECMWF and is publicly available from Copernicus Climate Change Service: Single level: Single level and Pressure level. CMIP6 data is available from any of: the ESGF Data Portal, NCAR Climate Data Gateway, the British Atmospheric Data Centre BADC, and the World Climate Research Programme WCRP Daily aggregates of these datasets are available on the National Computational Infrastructure for members.

A.2 Supplementary Tables and Figures

Table 1: List of predictors for the LSTM downscaling model. All climate variables from ERA5 and GCM (i.e., CNRM-ESM2-1), are aggregated to a daily temporal resolution. ERA5 predictors are coarsened from 25km to 150km using simple average. GCM predictors are regridded from 150km to 150km using bilinear interpolation. The atmospheric variables shown in bold represent those at the 850 hPa pressure level. The remaining variables are at the surface level.

Climate predictors (unit)	ERA5 short name	CNRM-ESM2-1 short name
Temperature (°K)	t	ta
Specific humidity (kg kg ⁻¹)	q	hus
U wind component (m s ⁻¹)	u	ua
V wind component (m s ⁻¹)	v	va
Geopotential (m ² s ⁻²)	z	zg
10 metre U wind component (m s ⁻¹)	u10	uas
10 metre V wind component (m s ⁻¹)	v10	vas
Maximum temperature at 2 metres (°K)	mx2t	Tasmax
Minimum temperature at 2 metres (°K)	mn2t	Tasmin
Cloud area fraction (0-1)	tcc	clt
2 metre temperature (°K)	2t	tas
Mean sea level pressure (Pa)	msl	psl
Total daily precipitation (mm)	Convective p	pr
Position (°, °, °)		
Time (°, °)		

Table 2: List of Extreme Precipitation Indices. These indices are calculated for each year in a time period and then averaged over all years.

_

Precipitation Index	Description
rx1day	Average annual maximum 1-day precipitation (mm)
r10mm	Number of heavy precipitation days $\geq 10 \text{ mm}$
r40mm	Number of extreme precipitation days $\geq 40 \text{ mm}$
preptot	Average annual total wet day precipitation (mm year ^{-1})



Figure 6: Boxplots illustrating the performance of each downscaling setting: (mse) MSE-based loss function, (mse_with_neigh) MSE-based loss function and neighbouring grid cells precipitation in the predictors set, (mse_quantile_with_neigh) same as the previous setting, but the loss function also includes quantile optimization. The performance metrics are calculated on the testing period by comparing yearly precipitation time series (downscaled ERA5 versus fine ERA5) from 144 fine-resolution grid cells contained in the four sample grid cells shown in Fig.8.



Figure 7: Maps illustrating four extreme precipitation indices (Tab.2) calculated over the testing period [2010-2020] based on coarse ERA5, fine-resolution ERA5 (ERA5-target) and downscaled ERA5 (ERA5-prediction).



Figure 8: Location of four sample grid cells on a Climate Classification Map generated by the Australian Bureau of Meteorology based on a modified Koppen classification system [24]. These grid cells are used to compare the performance of LSTMs across three settings (mse, mse+neigh, and mse+neigh+quantile).