From Global to Local: AI-based Climate Downscaling for Southeast Asia

Juntao Yang [1] **Xin Wang** [2] **Chen Chen** [3] **Jeff Adie** [1] **Simon See** [1] **Sandeep Sahany** [3] **Aurel Moise** [3] **Gianmarco Mengaldo** [2]

[1]*NVIDIA* [2]*National University of Singapore* [3]*Centre of Climate Research Singapore*. Correspondence to: Gianmarco Mengaldo mpegim@nus.edu.sg .

Abstract

High-resolution regional climate information is essential for impact assessment and adaptation in Southeast Asia. In this abstract, we share some preliminary results on an AI-based climate downscaling investigation for Southeast Asia. Focusing on daily precipitation and daily maximum near-surface temperature, we train a NVIDIA CorrDiff model to downscale 250km ACCESS-CM2 fields to 8km resolution regional data, achieving a downscaling factor of more than 30x. In addition, the inference speed of the AI downscaling model achieves the speed of 165 Simulated Years Per Day (SYPD) on a single GPU, e.g. 30 ensembles of downscaling for 100 years can be generated in about one day with 2 nodes of 8 GPU server.

1. Introduction

The Coupled Model Intercomparison Project Phase 6 (CMIP6) as a coordinated framework produces standardized simulations from over 100 Global Climate Models (GCMs) to improve understanding of past, present, and future climate changes, serving as a foundational resource for climate science research and providing critical data [1]. Precipitation and temperature are the two most essential meteorological variables for end users, such as policy makers, planners, engineers, and impact modelers. However, CMIP6 GCMs typically have resolutions in the range between 100 and 250 km. These models do not provide reliable information on the finer scales at which hydrological processes typically occur. In particular, the resolutions are too coarse to resolve circulation patterns to accurately capture regional climate variations and extreme weather events [2]. Downscaling of climate data from the CMIP6 data becomes a critical process that bridges the gap between global models and the finer scales required for local decision-making and impact assessments by transforming the coarse resolution climate model output into finer scales. This process is important to understand the impacts of climate change on specific regions and to develop targeted adaptation and mitigation strategies.

In Southeast Asia (SEA), the importance of understanding and addressing environmental issues cannot be overstated, as it is susceptible to climate-related challenges such as rising sea levels, extreme weather events, and loss of biodiversity. In addition, many countries in the SEA are still developing, often resulting in a lack of comprehensive data and studies needed for effective policy-making and mitigation plans. This gap not only hampers their ability to respond to climate challenges, but also limits their capacity to plan for sustainable development. Therefore, research efforts in these regions are very valuable in improving resilience and supporting vulnerable populations to adapt to changing climate.

Downscaling is particularly important for Singapore due to its small geographical size and its vulnerability to climate change impacts such as sea-level rise, extreme weather, and temperature changes. Being an urbanized island nation with limited natural resources and land space, more precise climate projections are crucial for effective urban planning, infrastructure development, and resource management. In response to these challenges, Singapore has been proactive in addressing climate change, as evidenced by the Singapore's Third National Climate Change Study (V3), which emphasizes the importance of using downscaled climate data to tailor strategies that mitigate risks and improve resilience [3].

Dynamical downscaling is a well-established method using Regional Climate Model (RCM). Numerical based RCM is driven by coarse resolution global projections to generate high resolution regional climate data. For example, the SINGV numerical model is used for dynamical downscaling for SEA [3]. However, the computational cost for dynamical downscaling can be prohibitive, limiting regional coverage and resolution. An alternative proposed in recent years is AI downscaling, which many evidences show a significant advantage in the speed of generating high resolution downscaled data. Generative models such as the Generative Adversarial Network (GAN) and Diffusion models have been found to be particularly promising in the downscaling task generating high quality downscaled simulations [4, 5, 6, 7]. However, most of the mentioned literature test their AI downscaling model with artificial coarsened dataset under the so called "perfect model framework" assumption. In this study, we investigate the capability of diffusion based AI downscaling model for the region around Singapore utilizing the Singapore V3 study data, where the AI downscaling model is trained on actual GCM input instead of artificially coarsened dataset. We focus on two variables, the daily maximum near surface air temperature and the daily total precipitation, which are of the most interest in the study of regional climate.

2. Dynamical downscaling and Data

V3 represents a landmark achievement in regional climate science, providing the world's highest-resolution climate projections for SEA [3]. V3 em-

ploys dynamical downscaling with the SINGV model to map coarse-resolution global climate projection into highly detailed regional projections, achieving unprecedented 8km spatial resolution for SEA and 2km resolution for Singapore, as depicted in Fig A1. The study examines three Shared Socioeconomic Pathways (SSPs) from IPCC AR6—SSP1-2.6 (low emissions), SSP2-4.5 (medium emissions), and SSP5-8.5 (high emissions)—to explore the range of potential climate outcomes by 2100, providing an important overview for Singapore and the whole of SEA.

The coarse historical and projection data from CMIP6 GCMs, including ACCESS-CM2, EC-Earth3, MIROC6, MPI-ESM1-2-HR, NorESM2-MM, UKESM1-0-LL as well as ERA5 data, are first downscaled to 8km resolution. The data from ACCESS-CM2, EC-Earth3, MPI-ESM1-2-HR, NorESM2-MM, UKESM1-0-LL and ERA5 data are then downscaled to a spatial resolution of 2 km.

However, dynamical downscaling is highly limited by computational challenge. The computational cost of models based on numerical methods increases exponentially, leading to prohibitively high computational cost to downscale to high spatial resolution for a large domain. Limited by computational cost, the coverage of regional climate studies is limited in terms of time and space. For example, limited by computational budget, the downscaling study from 8 to 2 km only covers the period in 1995-2014, 2040-2059, and 2080-2099. Not only in terms of time coverage, it only covers a smaller domain D2 as shown in Fig A1. In addition to temporal and spatial coverage, it also limits the number of SSP scenarios and ensembles. In terms of SSP scenarios, the V3 study only covers SSP1-2.6 (low emissions), SSP2-4.5 (medium emissions) and SSP5-8.5 (high emissions) of 8 main scenarios. For each SSP scenario in the V3 study, only one ensemble sample is downscaled for each GCM.

However, recent development of AI weather/climate models, such as DLWP, FourCastNet, and Aurora, shows the potential to significantly speed up by AI emulations [8, 9, 10] and hybrid approaches [11]. There have also been succussful demonstrations of AI downscaling for weather and climate data [4, 5]. However, such investigations have not been done for the tropical region. With the availability of data generated by the V3 study, our aim is to investigate regional climate downscaling in the Western Maritime Continent (WMC) with a generative AI model.

3. Method

NVIDIA CorrDiff is a generative AI model developed by NVIDIA for high-resolution weather forecasting and atmospheric downscaling [4]. It employs a two-step approach: first, a UNet model predicts the mean field, then a diffusion model refines the residual. This architecture allows CorrDiff to effectively downscale from coarse-resolution global weather data to fine-resolution regional predictions, synthesize new variables, and improve accuracy in predicting extreme weather events and fine-scale atmospheric phenomena. The model's U-Net structure, enhanced with attention mechanisms and skip connections, is conditioned on both the coarse-resolution input data and the initial U-Net prediction. CorrDiff has demonstrated promising results in capturing complex weather patterns, including sharpening of wind and temperature gradients, intensifying typhoons, and synthesizing rain band structures. As a faster and more energy-efficient alternative to traditional numerical weather models, CorrDiff represents a significant advance in applying generative AI to weather forecasting, offering improved accuracy, efficiency, and detail in predicting atmospheric conditions. Despite its success in weather data downscaling [4], it has not yet been implemented for regional climate downscaling.

In this study, our aim was to train an AI model for downscaling from CMIP GCMs resolution to 8km resolution. To be specific, we chose data from ACCESS-CM2 ensemble sample r4i1p1f1 historical, SSP126 and SSP585 as our training data and used the ACCESS-CM2 ensemble sample r4i1p1f1 SSP245 as evaluation data. Unlike what has been done in most literature, we investigate the real performance of AI downscaling with the actual GCM data as input. The results are important for understanding the operational potential of the AI downscaling model.

4. Results and Discussion

Due to the page limit, we present our preliminary results in Appendix B. We compared the results of the deterministic U-Net and the diffusion model on the logarithmic histogram distance metric (LHD), Radially Averaged Logarithmic Spectral Distance (RALSD), temporal variability, seasonal precipitation, Rx1day, and Consecutive Dry Day climatology in the test dataset (SSP245 2080-2099). The diffusion model consistently outperforms the deterministic U-Net in all metrics. The deterministic U-Net model in our experiment underestimates precipitation, leading to low performance in most evaluation metrics. This aligns with many prior findings in the literature that a deterministic model like U-Net tends to predict a smooth mean field with difficulties capturing small scale features and extreme values. Overall, the CorrDiff model has shown promising results on this difficult downscaling task from the CMIP 250km resolution to SINGV 8km. Such high downscaling ratio on actual climate projection data from CMIP has never been done before; however, our experiment shows that the diffusion model is able to capture the seasonal precipitation projection relatively well.

Acknowledgments

We acknowledge the support from National Supercomputing Centre Singapore for the data access/storage and compute infra support.

Appendix A. Training and Testing

In this study, we take advantage of the valuable data generated from the V3 study with dynamical downscaling. Fig A1 shows the domain coverage of the V3 dynamical downscaling project. D1 is the domain coverage for the first stage downscaling in V3 study from GCM resolution to 8km. Instead of using the data covering entire D1 domain, in this study we limit our domain between (15.94S, 18.62N) and (84.51E, 119.07E). The overall dimension of the 8km resolution data over the domain is (480, 480).

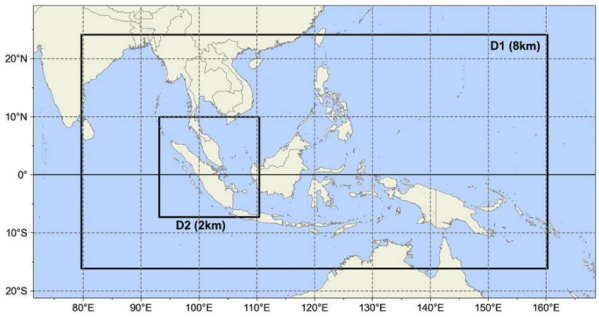

Fig. A1: Downscaling domains for V3. D1 (16.16 S–24.08 N; 79.68 E–160.248 E) is the 8km spatial resolution domain, and D2 (7.29 S–9.972 N; 93.16 E–110.422 E) is the 2km spatial resolution domain. Note: Figure adopted from V3 study [3]

There are 6 GCM data dynamically downscaled in the V3 study. We choose ACCESS-CM2 at around 250km resolution as input and corresponding SINGV downscaled 8km resolution data as target. Referring to previous work in [6, 7], we choose specific humidity (hus), zonal and meridional wind speed (ua, va), air temperature (ta) at pressure levels of 500 hPa and 850 hPa and two additional variable daily precipitation (pr) and near surface air temperature (tas) as input for our AI downscaling model. The target variable is chosen to be daily precipitation (pr) and daily maximum near surface air temperature (tasmax) which are the two most essential variables in regional downscaling. The ACCESS-CM2 GCM input data is downloaded from ESGF. We download data for historical, SSP126, SSP245 and SSP585 with ensemble number r4i1p1f1, which are the same as the data used in the V3 downscaling data.

For the AI downscaling model, we use a 5 level U-Net for the regression task as well as the diffusion model in the NVIDIA CorrDiff model with model channel size of [128, 256, 256, 256, 256] for each level and a self-attention layer at the last layer. EDM diffusion model is used in the NVIDIA CorrDiff model. In the diffusion model, the U-Net is conditioned on the high resolution prediction from the regression model and additional 4 input channels with sinusoidal positional embedding. We use the data from historical, SSP126 and SSP585 for training, the data from year 2012 to 2014 in the historical is used for validation. SSP245 is used as a test data set. Both input and target daily precipitation data (mm/day) are logarithmically transformed by $\ln(\mathrm{pr} + 0.001)$. All input and target data sets are normalized by their mean and standard deviation during training.

For both regression and diffusion model training, we have total training durations of 2000000 samples and a rampup period of 200000 samples (10% of total training duration) with final learning rate 0.0002. Batch size used is 192 with batch size per GPU being 12. The total of 16 H100 GPUs is used for the training with NVIDIA PhysicsNeMo, incorporating optimizations such as mixed precision with bf16, optimized apex group normalization, and torch compile. After training, we tested the model with the SSP245 test dataset from 2080 to 2099. Overall for deterministic U-Net model, we get an RMSE error of 8.16e-01 and MAE error of 4.42e-01 for near surface daily max temperature and RMSE error of 3.03 and MAE error of 2.12 for log daily precipitation. For the diffusion model, we get an RMSE error of 8.27e-01 and CRPS error of 3.41e-01 for near to surface daily max temperature and RMSE error of 3.09 and CRPS error of 1.40 for log daily precipitation. We show one test example on date 2099-01-01 under SSP245 in Fig. A2 the input data and Fig A3 one ensemble of the AI downscaling output against the ground truth from RCM. We also include Fig A4 to show 5 ensemble members out of 30 total ensembles generated. The generation speed of one ensemble is about 1.435s (equivalent to 165 SYPD on a single GPU, e.g. given 2 node of 8 GPU server, we can generate 30 ensemble downscaling for 100 years in about one day.) with 18 steps of stochastic sampler, which is magnitudes faster than classical dynamical downscaling.

Appendix B. Preliminary Results

2.1 Precipitation Distribution

We evaluated the AI model downscaled precipitation distribution using a histogram for all grid points and timesteps on the SSP245 dataset ranging from 2080 to 2099. Logarithmic Histogram Distance (LHD) is used following the practices in [12, 6, 7], defined as

$$LHD(dB) = \left(\frac{1}{N} \sum_i \left(10 \log_{10} \frac{\gamma_{true_i}}{\gamma_{pred_i}} \right)^2 \right)^{1/2} . \qquad \text{(A1)}$$

Here, γ_{pred} is the predicted histogram, γ_{true} is the ground truth histogram, and N is the total bin number of the histogram. Evenly spaced 50 bins are used between the range 1-1800 mm/day for the LHD computation, and only γ_{true_i} and γ_{pred_i} more than 10 counts are included in the computation. The test results

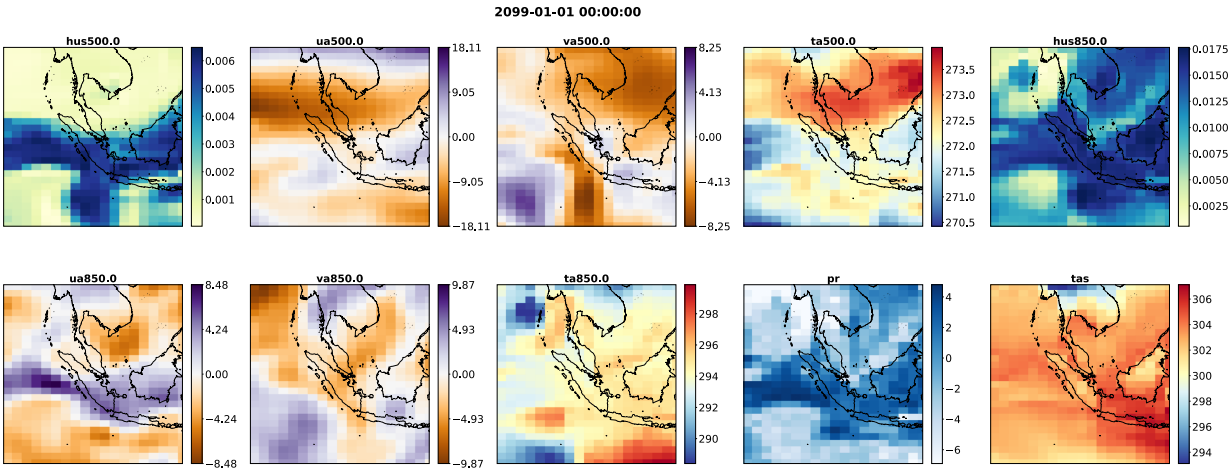

Fig. A2: Test dataset input sample at 2099-01-01

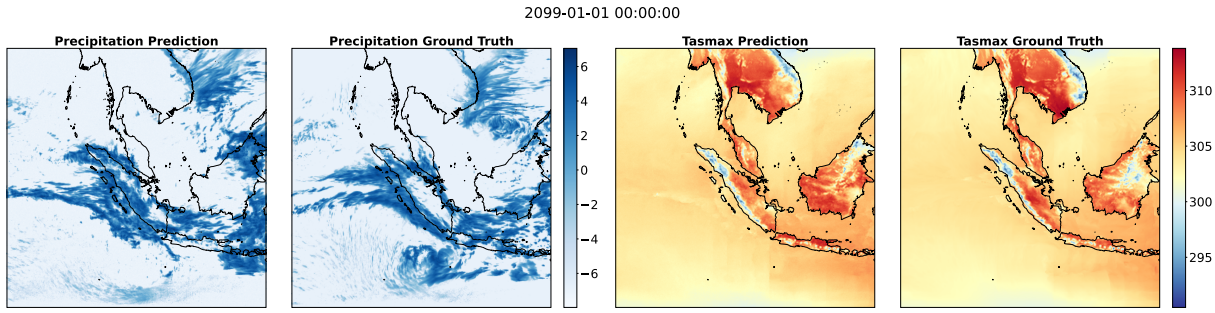

Fig. A3: Prediction and ground truth on test dataset, sample 2099-01-01

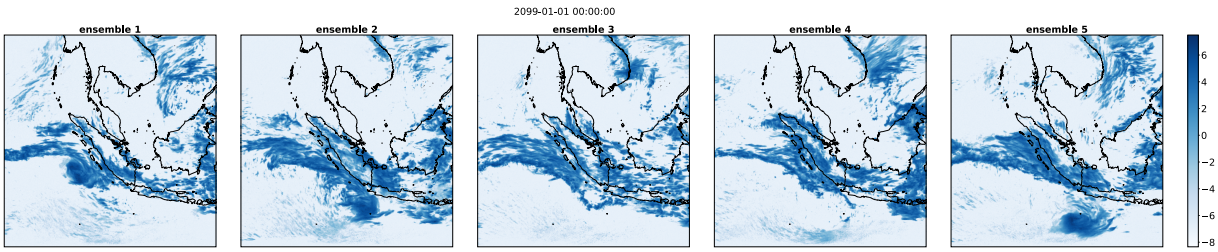

Fig. A4: Some ensemble members of the precipitation prediction on 2099-01-01

from the deterministic U-Net regression model have an LHD value of 21.35 in our investigation. The metric improves significantly with CorrDiff with a LHD value of 5.79.

2.2 Power Spectral Density

To evaluate the skill in resolving fine scale features, Radially Averaged Logarithmic Spectral Distance (RALSD) metric is used. It is computed with the following definition,

$$RALSD(dB) = \left(\frac{1}{N} \sum_i \left(10 log_{10} \frac{\hat{F}_{true_i}}{\hat{F}_{pred_i}} \right)^2 \right)^{1/2}, \quad \text{(A2)}$$

where $\hat{F}_{pred_i}$ and $\hat{F}_{true_i}$ are the power spectral density of the predicted data and ground truth data, N is the total number of spectrum. With deterministic U-Net, we get a RALSD value of 30.46. Significant improvement is seen with the CorrDiff model with a RALSD value of 2.27.

2.3 Seasonal Precipitation Sensitivity

We computed the seasonal precipitation for DJF (December, January and February) and JJA (June, July and August). We get seasonal average precipitation of 6.452 mm/day from DJF ground truth, only 0.4535 mm/day from the deterministic U-Net prediction and 8.083 mm/day from CorrDiff model 30 ensemble average. For JJA seasonal average precipiation, we get 7.1963 mm/day for ground truth data, only 1.2773 mm/day for deterministic U-Net prediction and 8.481 mm/day from CorrDiff model 30 ensemble average.

2.4 Rx1day and Consecutive Dry Day Climatology

We evaluated the highest 1-day precipitation amount (Rx1day) annually for each grid point. For the Rx1day mean across years and space, we get 218.1667 mm/day for the ground truth data, 15.4531 mm/day from the deterministic U-Net predictions and 71.173 mm/day from CorrDiff model 30 ensemble average. We evaluated the maximum number of consecutive dry days annually known as consecutive dry day climatology (CDD) with a threshold of 1 mm/day. For the CDD mean across years, we get 90.42 days from ground truth data, 202.43 days from the deterministic U-Net prediction and 170.94 days from CorrDiff model 30 ensemble average.

2.5 Continuous Ranked Probability Score (CRPS)

CRPS is a commonly used metric for generative AI models. Let X be a random variable and F be the cumulative distribution function (CDF) of X, given an observation x, CRPS is defined as

$$\text{CRPS}(F, x) = \int_{-\infty}^{\infty} (F(y) - \mathbb{I}(y - x))^2 dy, \quad \text{(A3)}$$

with $\mathbb{I}$ the Heaviside step function. In the case of deterministic prediction, the CRPS is reduced to be equivalent of Maximum Absolute Error (MAE). For deterministic U-Net, the MAE of test results are 2.12 for log scale daily precipitation. We generated 30 ensembles from CorrDiff model results in a CRPS value of 1.402.

	U-Net	CorrDiff	Dynamical
DJF	0.5-3.9 °C	0.5-4.1 °C	1.5-3.3 °C
MAM	0.9-4.7 °C	0.6-4.8 °C	1.6-4.7 °C
JJA	0.9-4.7 °C	0.6-4.9 °C	1.7-3.9 °C
SON	0.2-4.5 °C	0.5-4.8 °C	1.6-3.7 °C

Table A1: Projected changes in average seasonal maximum of daily maximum temperature in the end-century (2080-99) under SSP245 in comparison to historical 1995-2014

2.6 Case studies, Daily maximum temperature under SSP245 (Year 2080-2099)

We compare the climate change projection generated from the dynamical downscaling, and AI model during end-century (2080-99).

Fig A5 shows the original dynamical downscaling projected changes in the average seasonal maximum of daily maximum temperatures (alternatively, the hottest day of the season) in the end-century under SSP245 scenario. On average, the hottest day of the season is expected to become hotter. We summarize the projected changes in the average seasonal maximum of daily maximum temperatures from the deterministic U-Net and CorrDiff model in Table A1.

References

[1] Intergovernmental Panel on Climate Change. Climate change 2023: Synthesis report. Technical report, IPCC, Geneva, Switzerland, March 2023.

[2] S Brienen et al. Precipitation downscaling under climate change: Recent developments to bridge the gap between dynamical models and the end user. *Reviews of geophysics*, 48(3), 2010.

[3] Centre for Climate Research Singapore. Third national climate change study (v3). Technical report, Centre for Climate Research Singapore (CCRS), 2024.

[4] Morteza Mardani, Noah D. Brenowitz, Yair Cohen, Jaideep Pathak, Chieh-Yu Chen, Cheng-Chin Liu, Arash Vahdat, Karthik Kashinath, Jan Kautz, and Michael S. Pritchard. Residual corrective diffusion modeling for km-scale atmospheric downscaling. *Communications Earth & Environment*, 6, 2023.

[5] Neelesh Rampal, Sanaa Hobeichi, Peter B. Gibson, Jorge Baño-Medina, Gab Abramowitz, Tom Beucler, Jose González-Abad, William Chapman, Paula Harder, and José Manuel Gutiérrez. Enhancing regional climate downscaling through advances in machine learning. *Artificial Intelligence for the Earth Systems*, 3(2):230066, 2024.

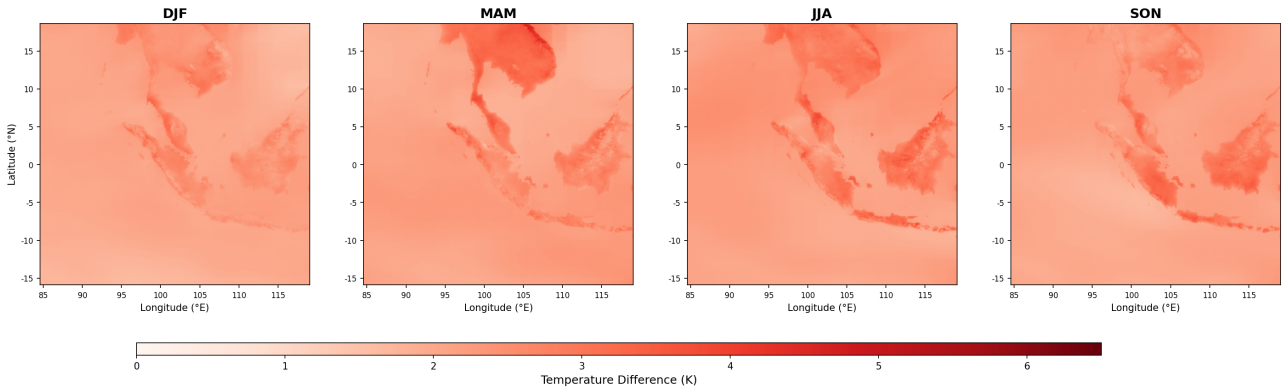

Fig. A5: Dynamical downscaling projected changes in seasonal max near-surface air temperature during end-century (2080–99) relative to the historical period (1995–2014) under scenario SSP245

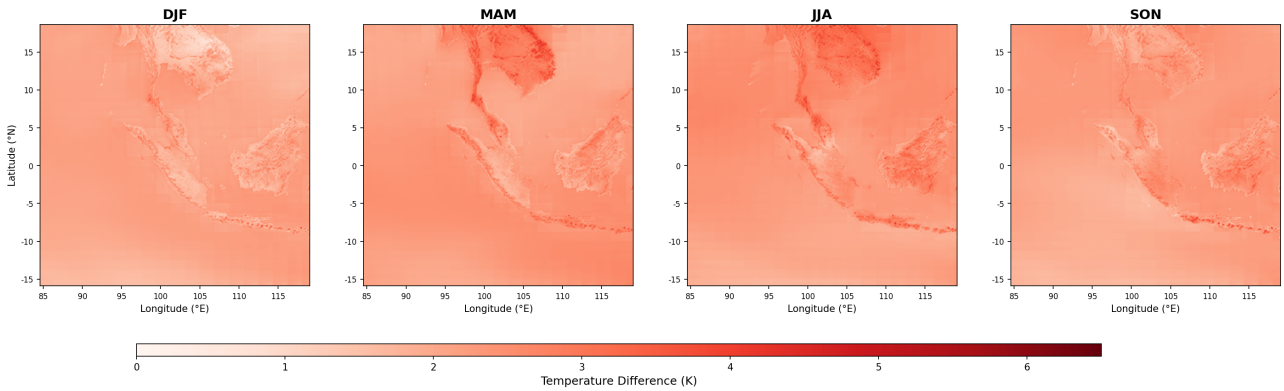

Fig. A6: CorrDiff downscaling projected changes in seasonal max near-surface air temperature during end-century (2080–99) relative to the historical period (1995–2014) under scenario SSP245

[6] Neelesh Rampal, Peter B Gibson, Steven Sherwood, Gab Abramowitz, and Sanaa Hobeichi. A reliable generative adversarial network approach for climate downscaling and weather generation. *Journal of Advances in Modeling Earth Systems*, 17(1):e2024MS004668, 2025.

[7] Bryn Ward-Leikis, Neelesh Rampal, Yun Sing Koh, Peter B. Gibson, Hong-Yang Liu, Vassili Kitsios, Tristan Meyers, Jeff Adie, Yang Juntao, and Steven C. Sherwood. An intercomparison of generative machine learning methods for downscaling precipitation at fine spatial scales, 2025.

[8] Jonathan A Weyn, Dale R Durran, Rich Caruana, and Nathaniel Cresswell-Clay. Sub-seasonal forecasting with a large ensemble of deep-learning weather prediction models. *Journal of Advances in Modeling Earth Systems*, 13(7):e2021MS002502, 2021.

[9] Jaideep Pathak, Shashank Subramanian, Peter Z. Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zong-Yi Li, Kamyar Azizzadenesheli, Pedram Hassanzadeh, Karthik Kashinath, and Anima Anandkumar. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *ArXiv*, abs/2202.11214, 2022.

[10] Cristian Bodnar. Aurora: A foundation model of the atmosphere. In *AGU Fall Meeting Abstracts*, volume 2024, pages GC21C–03, 2024.

[11] Xin Wang, Jianda Chen, Juntao Yang, Jeff Adie, Simon See, Kalli Furtado, Chen Chen, Troy Arcomano, Romit Maulik, Wei Xue, et al. Condensnet: Enabling stable long-term climate simulations via hybrid deep learning models with adaptive physical constraints. *npj Climate and Atmospheric Science*, 9(1):7, 2026.

[12] Lucy Harris, Andrew TT McRae, Matthew Chantry, Peter D Dueben, and Tim N Palmer. A generative deep learning approach to stochastic downscaling of precipitation forecasts. *Journal of Advances in Modeling Earth Systems*, 14(10):e2022MS003120, 2022.