A Robust Generative Adversarial Network Approach for Climate Downscaling and Weather Generation

Neelesh Rampal[1,3], Peter B. Gibson[2], Steven Sherwood[3], Gab Abramowitz[3] and Sanaa Hoibechi[3]

[1]National Institute of Water and Atmospheric Research, Auckland, New Zealand.

[2]National Institute of Water and Atmospheric Research, Wellington, New Zealand

[3] Climate Change Research Centre & ARC Centre of Excellence for Climate Extremes, University of New South Wales, Sydney, Australia.

Corresponding author: Neelesh Rampal (neelesh.rampal@niwa.co.nz)

Key Points:

- Regression-based downscaling methods struggle to capture extreme events and poorly preserve important climate statistical measures.
- Generative Adversarial Networks outperform regression-based methods, but their skill is very sensitive to the selection of hyperparameters.
- Using constraints in the Generative Adversarial Network's loss function improves its robustness and skill of across a wide range of metrics.

Abstract

Anticipating climate impacts and risks in present or future climates requires predicting the statistics of high-impact weather events at fine-scales. Direct numerical simulations of fine-scale weather are computationally too expensive for many uses. While regression-based (deep-learning or statistical) downscaling of low-resolution climate simulations is several orders of magnitude faster than direct numerical simulations, it suffers from several limitations. These limitations include the tendency to regress to the mean, which produces excessively smooth predictions and underestimates the magnitude of extreme events. Additionally, they also fail to preserve statistical measures that are key for climate research. We use a conditional GAN (c-GAN) architecture to downscale daily precipitation as a Regional Climate Model (RCM) emulator. The c-GAN generates plausible residuals on top of the predictable expectation state produced by a regression-based DL algorithm. The skill of c-GANs is highly sensitive to a hyperparameter known as the weight of the adversarial loss (λ_{adv}), and the value of λ_{adv} required for accurate results varies with season and performance metric, casting doubt on the robustness of c-GANs as usually implemented. But, by applying a simple intensity constraint to the loss function, it is possible to obtain robust performance results across λ_{adv} spanning two orders of magnitude. C-GANs are considerably more skillful in capturing climatological statistics including the distribution and spatial characteristics of extreme events. We expect c-GANs with this modification to be readily transferable to other problems and time periods, making them a useful weather generator for representing extreme event statistics in present and future climates.

Plain Language Summary

Climate projections produced by Global Climate Models (GCMs) have a typical resolution of 100-200km, which is too coarse for studying climate impacts at regional scales. Dynamical downscaling involves running a Regional Climate Model (RCM) to simulate physical processes that are not resolved at the resolution of GCM, enabling high-resolution climate projections for studying localized climate change impacts. However, RCMs are computationally expensive, limiting both the number of GCMs that can be downscaled and estimates of uncertainty. Deep learning (DL) methods offer a promising, cost-effective alternative to RCMs, and recent studies have emulated certain aspects of RCMs at a fraction of the computational cost. Generative DL algorithms such as Generative Adversarial Networks (GANs) appear to show promise in accurately emulating RCMs, but their training instability and inconsistent performance across climate contexts raises concerns about their robustness for downscaling climate projections. Here we develop and introduce a simple technique to improve the stability in GAN performance across a wide range of training configurations. This improves robustness and utility in broader climate applications.

1. **Introduction**

The coarse spatial resolution of Global Climate Models (GCMs) limits their ability to simulate climate changes at regional and local scales, where the impacts of climate change are most directly experienced (Benestad, 2004, 2010; Fowler et al., 2007; Maraun, 2016; Maraun et al., 2010). Dynamical downscaling aims to address this resolution issue by capturing finer-scale aspects of mesoscale circulation and regional climate across different landscapes such as mountain ranges, valleys, and coastal boundaries (Feser et al., 2011; Gensini et al., 2023; Giorgi et al., 1994; Hoogewind et al., 2017; Jones et al., 1995; Liu et al., 2017; Prein et al., 2015; Xu et al., 2019). Dynamical downscaling typically involves running a Regional Climate Model (RCM) from the lateral boundary conditions of a GCM. A major drawback of RCMs is their high computational cost, which today limits their spatial resolution to scales of 12-50km when run operationally in Coordinated Regional climate Downscaling Experiment (CORDEX) type experiments (Giorgi et al., 2009). Additionally, the computational cost of RCMs limits the number of GCMs that can be downscaled. Consequently, this small number of downscaled GCMs means that model structural and internal variability uncertainty are under sampled in regional climate projections, despite its known importance on regional scales (Deser et al., 2012; Deser & Phillips, 2023; Gibson et al., 2024; Hawkins & Sutton, 2009, 2011).

Recently, computationally efficient statistical/empirical algorithms have been explored for RCM emulation, including simple multiple linear regression (Holden et al., 2015), multilayer perceptron (Chadwick et al., 2011; Hobeichi et al., 2023; Nishant et al., 2023), statistical analogues (Boé et al., 2023), and normalizing flows (Groenke et al., 2020). In both RCM emulation and other downscaling applications, there has been a shift towards regression-based deep learning computer vision algorithms such as CNNs (Babaousmail et al., 2021; Bano-Medina et al., 2023; Doury et al., 2022; van der Meer et al., 2023). These are better suited to the complex non-linear relationships between large-scale predictors and local-scale climate variables (Rampal et al., 2022) and have generally outperformed traditional statistical and machine learning (ML) techniques (Baño-Medina et al., 2020; Rampal, 2024; Rampal et al., 2022).

While regression-based approaches (including deep learning) are skillful in capturing the “mean-state” in instantaneous predictions (i.e. they regress to the mean), they tend to underestimate extreme events and struggle to resolve fine scale details (Harris et al., 2022; Mardani et al., 2023; Rampal, 2024; Reddy et al., 2023; Vosper et al., 2023; J. Wang et al., 2021). Unlike weather forecasting, accurate instantaneous predictions are less useful than climatological metrics (i.e. how often a given weather event occurs) in a climate projection context, as atmospheric variability is chaotic and effectively random beyond a short horizon. This may create a trade-off between accuracy of instantaneous predictions, and the skill in capturing climatological metrics and extreme events (Rampal et al., 2024). This is particularly problematic for extreme events (e.g. convective high intensity short duration rainfall events) which can have the highest societal impact. While there have been a wide variety of algorithm developments to overcome such issues in regression-based approaches, these issues persist (for recent reviews see Rampal et al., 2024; Sun et al., 2024) .

Generative Adversarial Networks (GANs) are a recent development in ML that may offer a solution to some of these shortcomings. GANs have been used in many research areas (Goodfellow et al., 2014; Isola et al., 2018; Mirza & Osindero, 2014; X. Wang et al., 2018), and have recently been adapted from the computer vision sub-field of super-resolution (which focuses on enhancing image resolution) to climate downscaling. GANs, also often described as conditional GANs (c-GANs) in this context, have significantly improved regression-based computer-vision algorithms in predicting local-scale extreme events and resolving high-resolution spatial structure in the downscaled predictions (Annau et al., 2023; Brochet et al., 2023; Izumi et al., 2022; Leinonen et al., 2021; Miralles et al., 2022; Oyama et al., 2023; Price & Rasp, 2022; Ravuri et al., 2021a; Saha & Ravela, 2022; Vosper et al., 2023; J. Wang et al., 2021).

Unlike traditional regression-based ML algorithms, which optimize for loss functions such as mean squared error (MSE), GANs are generative algorithms that incorporate an adversarial loss (Goodfellow et al., 2014; Mirza & Osindero, 2014) and use stochastic noise to generate an ensemble of predictions for a given set of large-scale predictor variables (i.e. coarse-resolution variables from GCMs).The adversarial loss function drives a competitive process between two CNNs, a generator and discriminator. The generator attempts to generate realistic samples (i.e. pseudo-RCM simulations), while the discriminator tries to distinguish between real data (i.e. RCM simulations) and the generator's output (Goodfellow et al., 2014; Mirza & Osindero, 2014). This competition leads to the generator implicitly learning through a powerful loss function that goes beyond traditional pixel-wise comparisons, encouraging the generation of outputs to be distributionally and structurally similar to the real data (Gulrajani et al., 2017).

The effectiveness of GANs for climate downscaling in present-day or future climates has not been well-assessed (Rampal, et al., 2024). Existing research mainly focuses on using traditional error metrics such as root-mean-squared error (Rampal et al., 2024; Sun et al., 2024) instead of climatological metrics. Additionally, GANs are notoriously unstable and challenging to train, where stability is often determined by selecting the correct hyperparameters (Arjovsky et al., 2017; Goodfellow et al., 2014; Gulrajani et al., 2017; Mirza & Osindero, 2014).

One particularly important hyperparameter is the weighting of the adversarial loss function (λ_{adv}) during training (refer to section 2.1 for more details), which determines the strength of the adversarial loss during training. While studies have analyzed the impact of model architecture and loss function choices on generated output quality, this research has been limited to computer vision applications (Abu-Srhan et al., 2022; Isola et al., 2018; Ledig et al., 2017; X. Wang et al., 2018). For example, Isola et al., (2018) highlighted that values too large would often hallucinate and generate artifacts (i.e. $\lambda_{adv} = 1$), and found optimal performance when $\lambda_{adv} = 0.01$ for image-to-image translation. Existing studies in downscaling applications have only conducted their research using a specific value of λ_{adv} (i.e. Annau et al., 2023; Harris et al., 2022; Leinonen et al., 2021; Vosper et al., 2023), with limited exploration of how the strength of the adversarial loss affects climate downscaling.

Our study therefore aims to focus on two aspects of evaluating GANs. Firstly, we assess whether GANs add value over regression-based RCM emulators. Secondly, we explore the robustness of GAN performance for RCM emulation, by varying the hyperparameter λ_{adv}. Our study uses a comprehensive set of evaluation metrics to ensure that GANs are useful in a variety of climate downscaling contexts. These metrics assess the emulator's ability to learn various climate statistics, such as the climatology of precipitation, extreme events, and the persistence of dry spells. We also evaluate the skill of GANs to generate ensembles, which have significant implications for uncertainty quantification in climate science and weather forecasting.

2. Materials and Methods

2.1 Training and Evaluation Data

2.1.1 Regional Climate Model Configuration

Our RCM emulator was trained using predictor and target variables from the Conformal Cubic Atmospheric Model (CCAM), a global non-hydrostatic atmospheric model with a variable-resolution cubic grid (Chapman et al., 2023; Gibson et al., 2023; McGregor & Dix, 2008; Thatcher & McGregor, 2009). In contrast to commonly used RCMs like the Weather Research and Forecasting Model (WRF), which rely on lateral boundary conditions from reanalysis or CMIP6 GCMs, CCAM is run as a global variable-resolution model (McGregor & Dix, 2008). CCAM is run globally with spectral nudging to input fields from GCM atmospheric variables. A detailed evaluation of CCAM is presented in Gibson et al. (2023) for this region, which used a very similar version of CCAM (i.e. model grid and physics configuration).

Although CCAM is a global model, our emulation efforts concentrate on the New Zealand region (165°E-184°W, 33°S-51°S) as shown in Figure 1 (target domain), where the highest resolution face of CCAM is near-uniformly 12km. Due to its diverse array of microclimates, the New Zealand region provides an ideal case study for RCM emulation. These microclimates arise due to New Zealand's complex geography, including coastlines, mountains, and its position in the mid-latitudes. New Zealand is also exposed to weather phenomena such as tropical cyclones, atmospheric rivers, and large-scale climate drivers such as the El Niño-Southern Oscillation (ENSO), and the Southern Annular Mode (SAM) (Refs). While physical processes governing New Zealand's regional climate are generally well captured by physics-based RCMs (Ackerley et al., 2012; Gibson et al., 2023), a key challenge is ensuring that RCM emulators can also learn these processes (Rampal et al., 2024).

2.1.2 Training Data

The main target variable is daily accumulated high-resolution (~12km) precipitation (pr) from CCAM output. Precipitation is logarithmically normalized ($z = \log_e(pr + 0.001)$) to reduce its distributional skewness, as implemented in various weather forecasting and downscaling studies (i.e. Rasp et al., 2020; Renwick et al., 2009). The coarse-resolution predictor variables are daily-averaged large-scale prognostic variables from CCAM, which include zonal wind (U), meridional wind (V), temperature (T) and specific humidity (Q) at two pressure levels, 500hPa and 850hPa in the atmosphere. The domain extent of the predictor variables is slightly larger than the target variable (151°E-188°W, 26°S-59°S) as illustrated in Figure 1, and was chosen to prevent information scarcity at the boundaries of the target domain (Bailie et al., 2024; Rampal et al., 2022). These predictor variables are re-gridded from 12km to a resolution of 1.5° (~150km) using conservative remapping. The predictor variables are normalized relative to the mean and standard deviation computed over the entire training dataset as implemented in Rampal et al., (2022) and Rasp et al., (2020). The rationale for using daily-aggregated predictor and target variables instead of sub-daily is to both speed up model training and inference time but also reduce CPU/GPU memory usage. Using daily input fields also ensures that the emulator can be applied to a much larger number of GCMs, since the availability of daily data is much greater than sub-daily data across the CMIP6 archive. It is important to note, that daily-aggregation also incurs a loss of temporal information, making the problem somewhat more challenging than using instantaneous fields (i.e. hourly).

Our study focuses solely on evaluating and training DL algorithms on historical RCM simulations. Our evaluation framework does not focus on out-of-distribution performance temporally (i.e. to future climates), but rather tests whether the emulator can be applied more broadly to other un-seen GCM/RCM simulations from training. All emulators were trained on 55 years of simulation (~21,000 days) from the CMIP6 ACCESS-CM2 (1960-2014). We assess the performance of all emulators using ground-truth downscaled simulations from CCAM, configured identically from two additional CMIP6 GCMs (EC-Earth3 and NorESM2-MM). This out-of-sample evaluation covers a 20-year historical period from 1986 to 2005 (~7300 days). Here, our emulator is applied to the CCAM-coarsened predictor fields (perfect framework) from these simulations. Doing so provides a true out-of-sample test of the emulator, testing the performance (and ability to generalize) on additional driving fields from GCMs which were unseen in training.

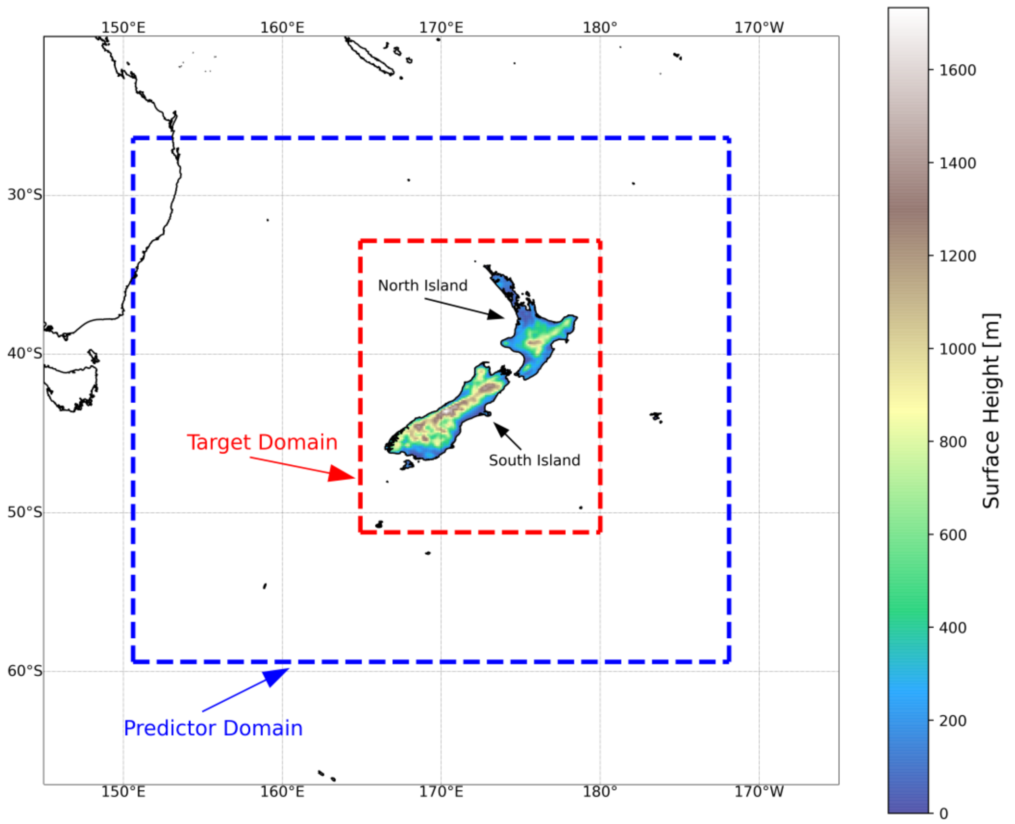

Figure 1: A depiction of the domain extent of the predictor variables (blue) and target variables (red) across the New Zealand region, with the color scale representing the region's surface elevation

2.1.3 Training framework

We have used CCAM-coarsened predictor variables as opposed to variables from the GCM directly. This training strategy is known as the perfect framework. It differs from the imperfect training framework, which uses GCM fields as predictor variables directly (Rampal, et al., 2024). Training an emulator through the imperfect framework is more challenging as the RCM's mean state can significantly deviate from the GCM (Bartók et al., 2017; Boé et al., 2020; Sørland et al., 2018). An emulator trained in the imperfect framework needs to learn both the deviations between the RCM and GCMs mean state, and the finer scales of RCM (Rampal et al., 2024), whereas the perfect framework emulator is only required to learn the latter.

Additionally, emulators trained in the imperfect framework have been shown to learn a relationship that is unique to a specific GCM/RCM pair (Bano-Medina et al., 2023; Boé et al., 2023) and thus is less portable across the wider GCM/RCM matrix. Conversely, training in the perfect framework has very little dependence on the RCM simulation used in training (as it does not have to account for differences in circulation between RCM and GCM). While there is an ongoing discussion about which framework is optimal in an out-of-sample operational setting, training and evaluation in the perfect framework is simpler and involves fewer considerations (Rampal, et al., 2024). This is advantageous for the purposes of testing different configurations of GANs in this study, and more broadly applying the emulator to downscaling multiple GCMs to generate high-resolution "pseudo" simulation.

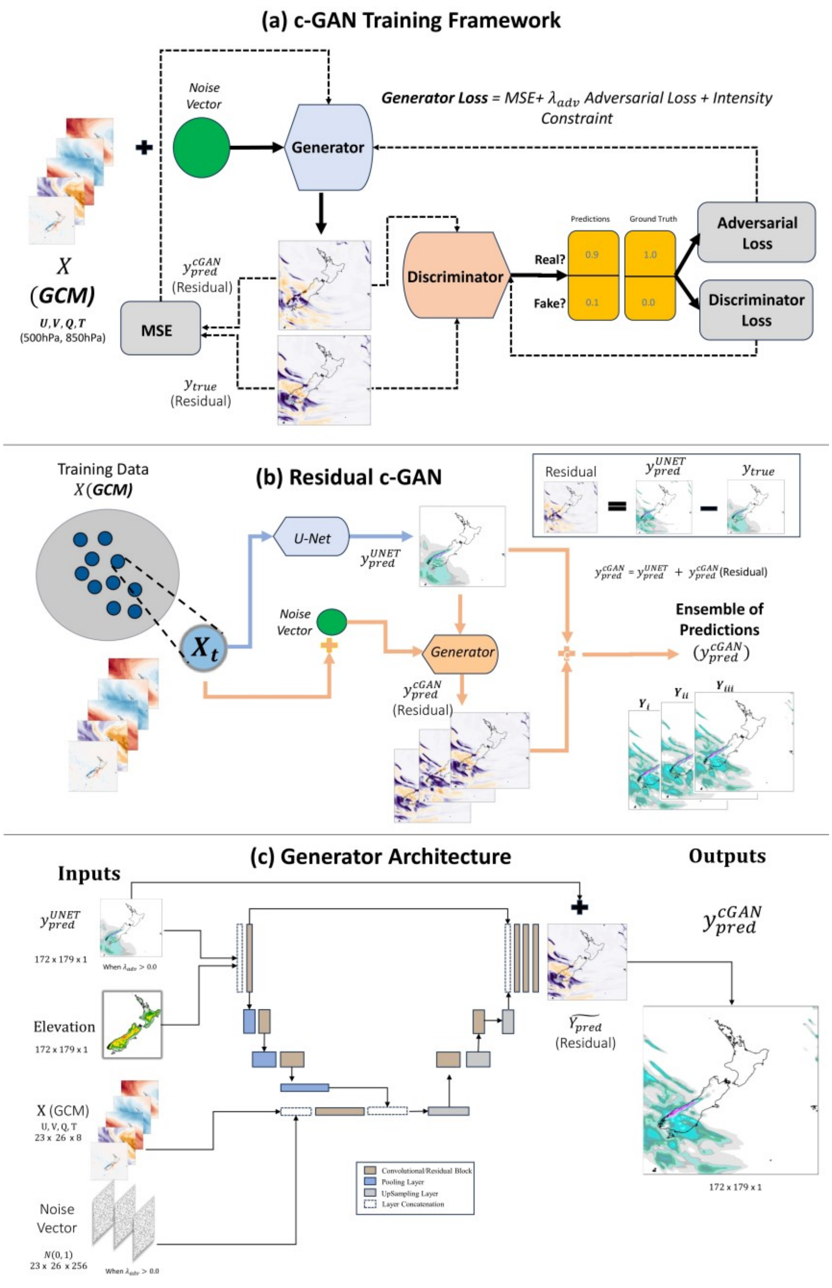

Figure 2: (a) An illustration of the training feedback loop of a GAN (c-GAN). The generator creates high-resolution predictions from low-resolution inputs and a noise vector. The discriminator then measures how realistic the generator's outputs are through a discriminator loss. (b) A residual GAN for downscaling. The residual GAN consists of two steps; first, a regression-based U-Net is trained to produce a deterministic prediction (y_{pred}^{UNET}) for a given X (predictor fields). This prediction and a noise vector are then input into the

generator, which is trained to predict the residual between the U-Net and the ground truth in logarithmic space. (c) Generator Architecture: This flowchart depicts the architecture of the generator within the GAN. It shows how multiple inputs, including a U-Net prediction (y_{pred}^{UNET}), elevation data, low-resolution GCM data (X), and a noise vector pass through a series of layers and processes to ultimately produce the high-resolution climate field prediction (y_{pred}^{cGAN}) used by the GAN.

2.2 Generative Adversarial Networks

The GAN architecture for downscaling consists of two main components: a generator and a discriminator (also known as a critic). The generator aims to create a high-resolution climate field from a "low-resolution" climate field as an input (the condition), and a discriminator evaluates whether the generated image is likely real (ground truth high-resolution simulations) or fake (synthetic high-resolution fields generated by the generator that may have characteristic artefacts). There are two main loss functions in training a GAN: the generator loss (G_{loss}) and the discriminator (critic) loss.

2.2.1 Generator Loss

In this study we train c-GANs with two different loss function configurations. In a downscaling or image super-resolution context, the generator loss usually consists of traditional loss functions such as the MSE and an adversarial loss function, G_{adv}, which is weighted by some constant factor λ_{adv}, as shown in Equation 1.

$$(1) \quad G_{loss}(y_{true}, y_{pred}) = MSE(y_{true}, y_{pred}) + \lambda_{adv} * G_{adv}(D(y_{pred})),$$

$$G_{adv}(y_{pred}) = \overline{-D(y_{pred})}$$

Here, y_{true} and y_{pred} refer to the ground truth RCM simulations and generated samples from the emulator, respectively. The adversarial loss function (G_{adv}) is calculated by taking the negative average of the discriminator's (D) output on generated samples $D(y_{pred})$. In simpler terms, the adversarial loss increases when the discriminator is not fooled by the generated images, penalizing the current weight set in the generator. The generator loss shown in Equation 1 is one of the two main loss function configurations explored in this study. It is widely used in many super-resolution and downscaling studies (i.e. Harris et al., 2022; Leinonen et al., 2021; Vosper et al., 2023). Note we use the MSE loss function as opposed to the MAE loss as it is more sensitive to errors in extreme events (not shown). It is important to note that training with an λ_{adv} too large is often unstable (Isola et al., 2018; Vosper et al., 2023), and the majority of existing studies generally use values of λ_{adv} less than 0.005 (Harris et al., 2022; Izumi et al., 2022; Leinonen et al., 2021; Vosper et al., 2023; X. Wang et al., 2018).

We also explore a second loss function configuration that incorporates an intensity constraint (IC), analogous to Ravuri et al. (2021) and Price & Rasp., (2022). The intensity constraint penalizes both the model's maximum precipitation intensity over the regional domain (Y^{max}) at each timestep, and its batch-averaged precipitation rate (Y^{mean}) for each location, as shown in Equation 2. The maximum precipitation intensity constraint prevents precipitation intensities from growing too large, and the batch-averaged precipitation (where the batch size is 32) is a proxy for conserving monthly precipitation averages. Note that during training, the batches are randomly shuffled at each epoch.

$$(2):\ G_{loss} = MSE(y_{true}, y_{pred}) + \lambda_{adv} * G_{Adv}(y_{pred}) + IC(y_{true}, y_{pred})$$

where $IC(y_{true}, y_{pred}) = MSE(Y_{true}^{max}, Y_{pred}^{max}) + MSE(Y_{true}^{mean}, Y_{pred}^{mean})$

2.2.2 Discriminator Loss

Similar to previous studies (Gulrajani et al., 2017; Harris et al., 2022; Leinonen et al., 2021; Vosper et al., 2023), we use the 1-Wasserstein distance (D_{adv}) as an discriminator or critic loss function (yielding what are often known as Wasserstein-GANs), where

$$D_{adv}(y_{true}, y_{pred}) = \overline{D(y_{true})} - \overline{D(y_{pred})}.$$

We also use a gradient penalty value of 10 (Gulrajani et al., 2017; Harris et al., 2022; Leinonen et al., 2021; Vosper et al., 2023). As implemented in these studies, we also train the discriminator three times as frequently as the generator. Overall, these refinements to the discriminator have been shown to improve training stability and a reduction of sensitivity to the choice of architecture and hyperparameters in c-GANs (Arjovsky et al., 2017).

2.2.3 Adversarial Parameter Selection

Our study examines how the solutions produced by GANs to the contribution of the adversarial loss weight (λ_{adv}). Increasing λ_{adv} allows the solutions from the GAN to diverge from the regression baseline as the adversarial loss becomes increasingly important. We explore seven different values of λ_{adv}: 0.0, 0.0001, 0.00125, 0.0025, 0.005, 0.01 and 0.1. Here, $\lambda_{adv} = 0$ refers to the regression baseline. The range of λ_{adv} was chosen to encompass the wide variety of values used in climate downscaling / weather forecasting literature.

2.3 Algorithm Architectures

In this study, we train two types of emulators: a regression baseline in which $\lambda_{adv} = 0.0$ and a residual GAN (Figure 2b). For the residual GAN, we test two different loss function configurations: with (Equation 2) and without an additional intensity constraint (Equation 1).

2.3.1 Regression Baseline

The regression baseline is based on the widely used U-Net deep learning model (Ronneberger et al., 2015), as illustrated in Figure 2c. The U-Net architecture consists of a contracting path, extracting information from the input predictor variables into a lower dimensional latent space. The expansive path involves reconstructing the high-resolution output (precipitation) from the latent space. The U-Net regression model consists of two input data streams: normalized high-resolution elevation data (12km) from CCAM and the large-scale prognostic predictor variables (1.5°). Our model uses residual convolutional layers (or residual blocks) with batch normalization, which have shown better performance than traditional convolutional layers and help address instability issues in deep-learning models (Rampal, et al., 2024; Sun et al., 2024). Following several residual convolutional blocks and pooling layers, the two input streams are concatenated and mixed to form the latent space of the model. Then, there are a series of upsampling (increasing the spatial resolution) and residual convolutional blocks until the output reaches the desired shape. Additionally, we repeated our experiments with and without batch normalization within our residual blocks, which had a minimal impact on our results.

2.3.2 Residual GAN

The residual GAN is trained to predict residuals ($r = \widetilde{y_{\lambda_{adv}=0}} - y_{true}$) between a regression baseline ($\lambda_{adv} = 0$) and the ground truth CCAM, as illustrated in Figure 2b. This residual methodology adapted from Mardani et al. (2023), who employed a similar approach in training a different type of generative model for downscaling called diffusion models. The regression baseline learns the expectation of all possible outcomes (the predictable component) from the RCM simulations, which tend to be smooth in both space and time (large-scale precipitation structures). This allows the residual GAN to focus on generating plausible fluctuations around this expectation, which include high-frequency variations and potentially some larger-scale contributions. The architecture of the generator in the residual GAN is nearly identical to the regression baseline, with two additional predictors: high-resolution prediction of precipitation from the regression baseline ($y_{\lambda_{adv}=0}$), and a stochastic noise vector, as inputs (see Figure 2c). Both the regression baseline and the residual GAN have approximately 3.5 million trainable parameters.

The discriminator or critic evaluates the perceptual realism of the residuals (either ground truth or predictors from the residual-GAN) conditioned on the regression baseline precipitation predictions ($y_{\lambda_{adv}=0}$), topography, and large-scale meteorological predictors (x), as shown in Figure 1a. The discriminator architecture features two input data streams analogous to the generator architecture: one for low-resolution fields with four convolutional layers, and another for high-resolution fields consisting of five convolutional layers. The two input data streams re subsequently concatenated in lower layers of the network. Both input data streams to the discriminator use strided convolutional layers for dimensionality reduction. To reduce model complexity and computational cost we excluded residual blocks from the discriminator architecture, which had negligible impact on our results (not shown).

In the discriminator and generator architectures, we use Leaky Rectified Linear Unit (ReLU) activation functions in all layers, as implemented in Leinonen et al. (2021). ReLU activation functions have been suggested to improve stability in training both the generator and discriminator. For the final output activation of the residual GAN, we use the LeakyRelu function ($r_\beta\ (x)$).

$$r_\beta\ (x) = \begin{Bmatrix} x & x > 0.0 \\ 0.5\,x & x \leq 0.0 \end{Bmatrix}.$$

In addition to LeakyRelu, we experimented with other output activation functions, such as a modified hyperbolic tangent (Tanh) function, which had a similar skill across all evaluation metrics used in this study.

Similar to previous studies (Gulrajani et al., 2017; Leinonen et al., 2020), both the generator and discriminator are trained with an initial learning rate of 2 x 10^{-4}, and a batch size of 32. The regression baseline is trained with a learning rate of 7 x 10^{-4}. We also explored smaller learning rates (i.e. $1\ x\ 10^{-6}$), which overall produced similar results but increased algorithm training times (not shown). To control overfitting and improve stability during training, we use learning decay for training the generator and the regression baseline with decay rates of 0.9945 and 0.989 per 1000 iterations, respectively. Each model was trained for 240 epochs, which equates to approximately 48 hours of training on a single NVIDIA A100 GPU with 80GB RAM. Additionally, learning rate decay also stabilizes GAN performance across different epochs (i.e. similar results were obtained from using 200 epochs instead of 240), addressing fluctuations in performance reported in prior studies. (Harris et al., 2022).

Predictions from the residual GAN are added to the regression baseline and inverse transformed ($pr = \exp\ (Y_{\lambda_{adv}=0..0} +\ Y_{\lambda_{adv}})$-0.001) to produce daily precipitation fields. Each experiment was repeated three times with a different random seed to ensure the consistency of results, and a separate regression baseline (U-Net) was trained with and without the intensity constraint. Generating a single simulation (one member) of 20-year daily precipitation (7300-time steps) record takes approximately 20 seconds on an A100 GPU.

2.4 Evaluation Metrics

Analogous to many RCM historical evaluation studies (i.e. Chapman et al., 2023; Di Virgilio et al., 2019, 2020; Isphording et al., 2023), we use three common climatological evaluation metrics to assess the out-of-sample performance over 20 years from 1986-2005. The first metric assesses the ability to capture seasonal averages in precipitation. This assessment for emulators is particularly important for New Zealand, where significant shifts in large-scale circulations, such as the subtropical and polar jet, occur between summer and winter and affect seasonal precipitation. Here, we use the summer and winter periods for evaluation: December-February (DJF) and June – August (JJA), respectively. We also use two other ETCCDI metrics that assess the performance of our emulator on capturing the climatology of extreme events (Isphording et al., 2023; Rampal et al., 2024; Zhang et al., 2011) the wettest day of the year (RX1Day) and the average number of consecutive dry days (CDD) per year.

3. Results

3.1 The Adversarial Effect on Local-scale Extremes

3.1.1 Case Studies of Extreme Events

To understand how the λ_{adv} affects the ability to resolve mesoscale structures and precipitation intensity; we present a case study of an extreme precipitation event simulated in EC-Earth3 over the New Zealand region. The emulator's predictions of precipitation across all λ_{adv} values (including the regression baseline) are spatially aligned with CCAM's precipitation patterns and associated low-pressure centers, as depicted in Figure 3. Overall, this demonstrates the emulator's proficiency in learning the effects of mesoscale circulation on extreme rainfall. This result is also consistent without the intensity constraint (Figure S1) and amongst other case studies (Figure S2-S3).

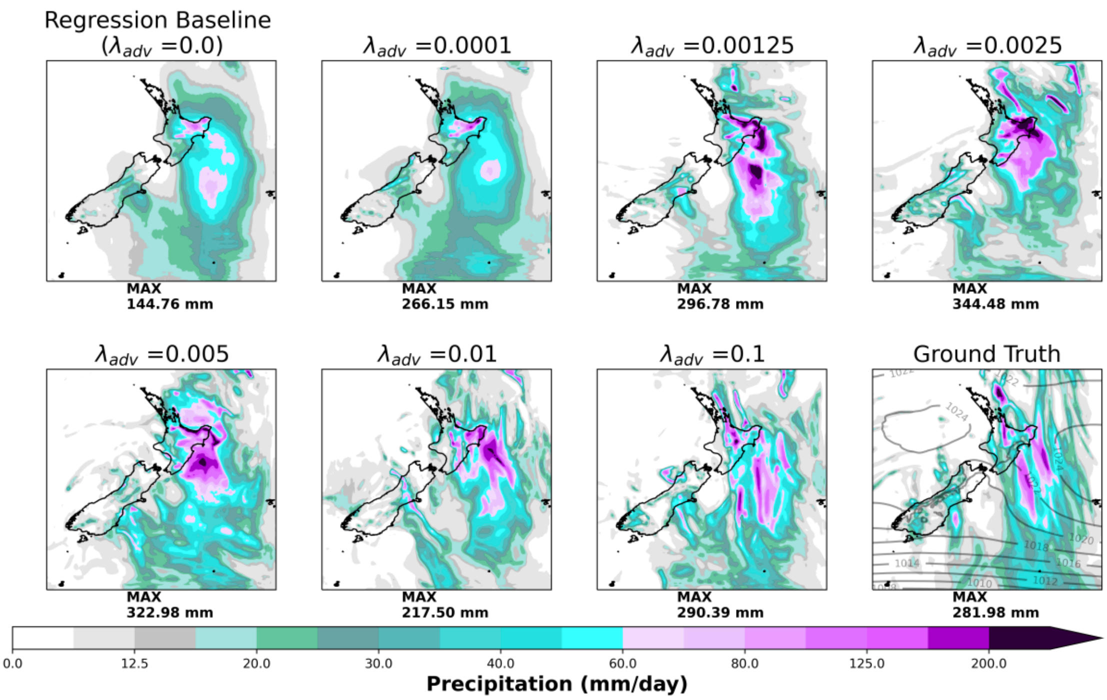

Figure 3: Example of daily precipitation predictions from GAN with the intensity constraint for a simulated extreme event from EC-Earth3 (2002-02-27), relative to the ground truth (CCAM downscaling EC-Earth3). The maximum precipitation intensity across the domain is shown in the text below the plot. The contours show CCAM's Mean Sea Level Pressure (MSLP) patterns for the same event.

The regression baseline ($\lambda_{adv} = 0$) significantly underestimates the maximum precipitation intensity and overly smooths mesoscale precipitation structures for the event depicted in Figure 3 relative to ground truth CCAM (also shown in Supplementary Figure S3-S4). However, when λ_{adv} ≥0.00125, the emulator can better resolve mesoscale structures and more accurately estimate the maximum precipitation rates across the domain. When the intensity constraint is not used, there are instances where the maximum precipitation intensity is significantly overestimated. Most notably, when $\lambda_{adv} = 0.01$ or 0.1, the intensity is overestimated by over 200% (see Supplementary Figure S1).

3.1.2 Precipitation Distribution

To quantify performance more generally, we examine the distribution of precipitation across the entire New Zealand region (including land and ocean) via a one-dimensional histogram of precipitation for all grid points and daily timesteps as shown for both loss function configurations (Figure 4).

The regression-baseline captures the mean of the precipitation distribution relatively well (the lowest intensity histogram bins) but underestimates the frequency of the most extreme events (i.e. >200mm/day), as shown in Figure 4a-4b. GANs do not always outperform regression models in capturing the precipitation distribution. Rather, their performance depends heavily on the specific loss function configuration (with or without intensity constraints) and the weighting of the adversarial loss (λ_{adv}).

Overall, varying λ_{adv} has a minimal effect on the precipitation distribution when the intensity constraint is used (Equation 2; Figure 4a). For nearly all values of λ_{adv} the precipitation distribution closely matches CCAM's – albeit slightly underestimating the most extreme precipitation events (>500mm), as illustrated in Figure 4a. In contrast, when no intensity constraint is used (Equation 1), varying λ_{adv} has a strong effect on the precipitation distribution (Figure 4b). Here, the regression baseline and $\lambda_{adv} = 0.0001$ case, both underestimate precipitation frequency at all intensities relative to CCAM, whereas when $\lambda_{adv} = 0.1$ there is a significant overestimation of precipitation frequency across all intensities, including a maximum of over 1,000,000 mm/day. Unphysically large precipitation values have also been reported in previous studies (Harris et al., 2022; Vosper et al., 2023). Further evaluation using Quantile-Quantile (Q-Q) plots is shown in Supplementary Figure S4-S6.

3.1.3 Mesoscale Variability

To evaluate the emulator's skill in resolving finer scale aspects of precipitation, we computed the Power Spectral Density (PSD) on predictions from the 200 rainiest days on average across the domain (although we obtain similar results using all days). The PSD is computed on each day's two-dimensional field of precipitation, and then averaged across all days. Here, the PSD is the integrated Fourier Transform as a function of the spatial wavelength ($K = \sqrt{(k_x^2 + k_y^2)}$), where k_x and k_y are the wavelengths in the x and y directions, respectively. We normalized each day's precipitation so that that the PSD receives equal weight from all included days.

Our results indicate that for very small values of λ_{adv} (≤0.00125) including the regression baseline, variability at small spatial is underestimated regardless of the intensity constraint. Here, GANs do not fully resolve mesoscale structures in a similar capacity to CCAM as shown in Figure 4(c-d). For all other values of λ_{adv} The emulator's PSD closely follows CCAM for both loss functions. However, there is one major exception when λ_{adv} = 0.1, and variability is overestimated across all spatial wavelengths, leading to an exaggerated representation of large-scale and mesoscale variability. It is important to note that at very small spatial scales (~1/25km), there is generally good agreement across all λ_{adv}, including the regression baseline. This agreement is primarily attributed to incorporating topography as a predictor variable (not shown), enabling the algorithm to account for the influence of orographic precipitation (Bailie et al., 2024). We thus conclude that GANs can relatively robustly capture the range of spatial scales, and that this is not much affected by the intensity constraint.

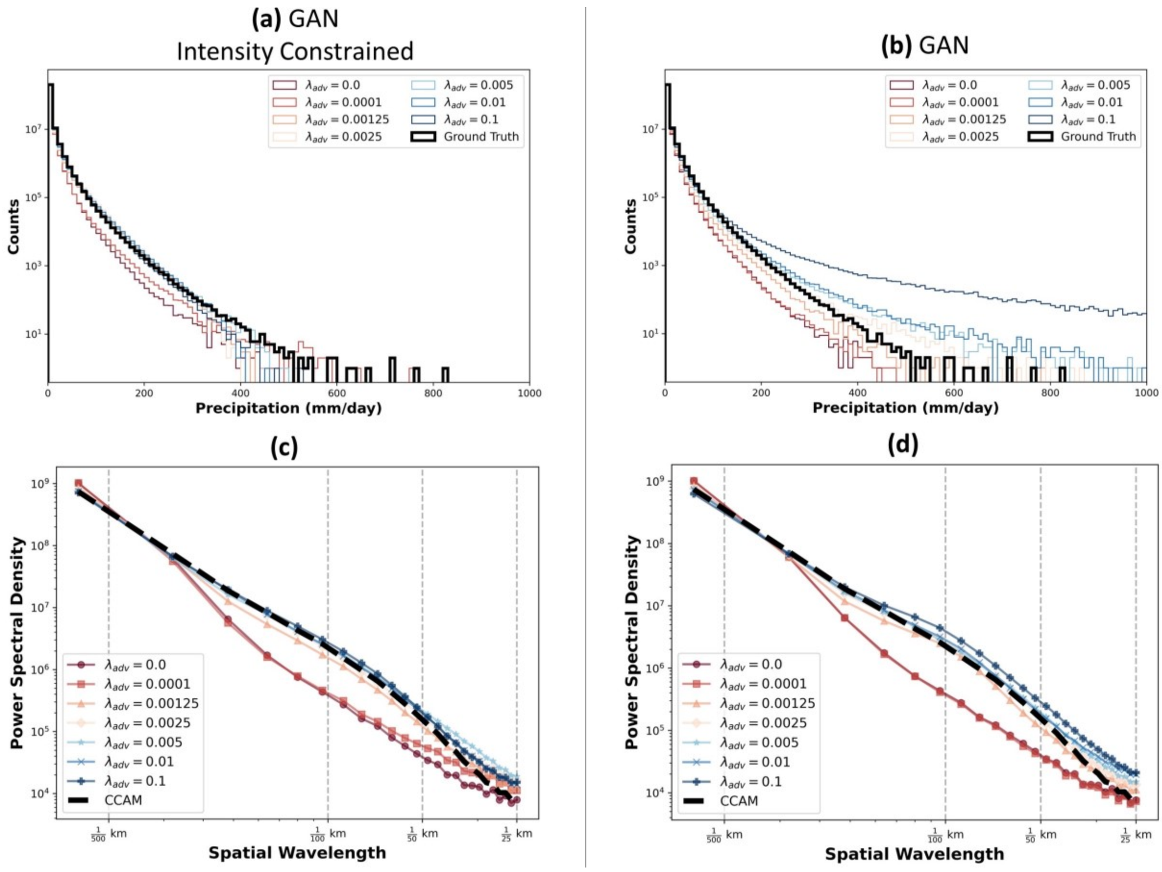

Figure 4: The precipitation distribution as a function of λ_{adv} when the RCM emulator is applied out-of-sample to EC-Earth3. (a) the histogram with the intensity constraint, and (b) without. Here, precipitation counts in the histogram are aggregated across all locations over the domain. The black curve highlights ground truth CCAM.

3.2 Temporal Variability

To evaluate performance on temporal variability (un-normalized precipitation), we compute the ratio (σ_r) in temporal standard deviation between each emulator's precipitation field (σ_e) and ground truth CCAM (σ_{CCAM}), where $\sigma_r = \frac{100 * \sigma_e}{\sigma_{CCAM}}$. Ratios less than 100% underestimate temporal variability, while values greater than 100% overestimate variability. This ratio is computed for each grid cell. Precipitation values exceeding 2000mm/day were excluded when computing standard deviation, which removes the contribution of unphysically large precipitation values, but this only affects the case when $\lambda_{adv} = 0.1$ without the intensity constraint.

The regression baseline and GAN (Figure 5a-b, left panel) substantially underestimates temporal variability by over 40%, regardless of generative loss configuration (Equation 1 or 2). However, as λ_{adv} increases further the ratio increases, as illustrated in the center panel in Figure 5a-b and Supplementary Figure S7. With the intensity constraint and when λ_{adv}=0.1, the emulator performs exceptionally well at capturing CCAM's temporal variability with an average ratio of 100%. When no intensity constraint is used, best performance is achieved when λ_{adv} =0.01, with an average 97% ratio (Supplementary Figure S7). However, when $\lambda_{adv} = 0.1$ without the intensity constraint the average ratio exceeds 110%, and in several individual grid points it exceeds 800%. Note if values exceed 2000mm, the ratio exceeds 1000%. Thus, for capturing temporal variability robustly, training with an intensity constraint appears important. However, without the intensity constraint, when λ_{adv} is large (≥ 0.1), the behavior appears unstable, likely due to the overestimation in extreme precipitation (Figure 2b) which inflates the temporal standard deviation.

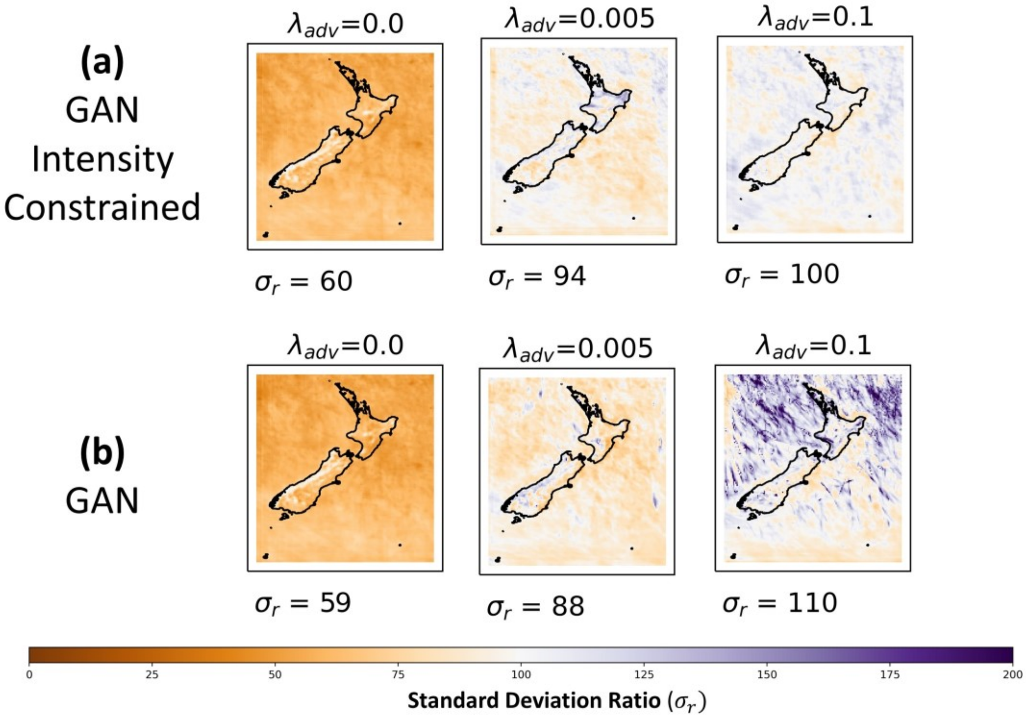

Figure 5: The percentage ratio of RCM emulated to ground truth temporal standard deviation in CCAM for the EC-Earth3 simulation. (a) shows the percentage ratio for the LeakyReLU activation with an intensity constraint applied and (b) without the constraint for three values of λ_{adv}. The variance ratio is calculated per grid pixel relative to the CCAM ground truth. The text below each Figure shows the average ratio (σ_r) across the entire domain.

3.3 Climate Statistics

This section evaluates the skill in emulating climate statistics/metrics and conventional error metrics such as MAE.

3.3.1 Seasonal Precipitation

Figure 6 shows the 10-member ensemble average out-of-sample emulator performance in representing four key climate metrics: (a) DJF and (b) JJA climatological precipitation, (c) RX1Day, and (d) CDD - each averaged the domain. The regression-baseline and the $\lambda_{adv} = 0.0001$ case have the highest MAE across all out-of-sample evaluation metrics. Increasing λ_{adv} improves the skill in reproducing the spatial patterns of seasonal precipitation (DJF and JJA) where the lowest MAEs are observed for $\lambda_{adv} \geq 0.01$. The regression baseline has an overall dry bias and increasing λ_{adv} better captures seasonal precipitation rates over the New Zealand region, as illustrated in Figure 7a(i-ii). A similar result is also shown for JJA climatological precipitation (Figure 7b(i-ii)), where the improvement is even more notable.

3.3.2 Rx1day Climatology

For the Rx1Day climatology, the regression baseline and $\lambda_{adv} = 0.0001$ cases again have the highest MAE (Figure 6c) and are generally dry-biased (Figure 8a(i-ii)) for both loss function configurations. The MAE decreases for higher lambda, except at $\lambda_{adv} = 0.1$ with no intensity constraint where there is a sharp increase in MAE (250%) and the RX1Day climatology is significantly overestimated (Figure 8a(ii) (rightmost panel)). On the other hand, the lowest MAE is achieved with this same $\lambda_{adv} = 0.1$ with the intensity constraint. The spatial patterns in the RX1Day climatology for $\lambda_{adv} = 0.1$ also match the ground truth. Overall, the RX1Day climatology performance seems to most robust across λ_{adv} values when the intensity constraint is applied.

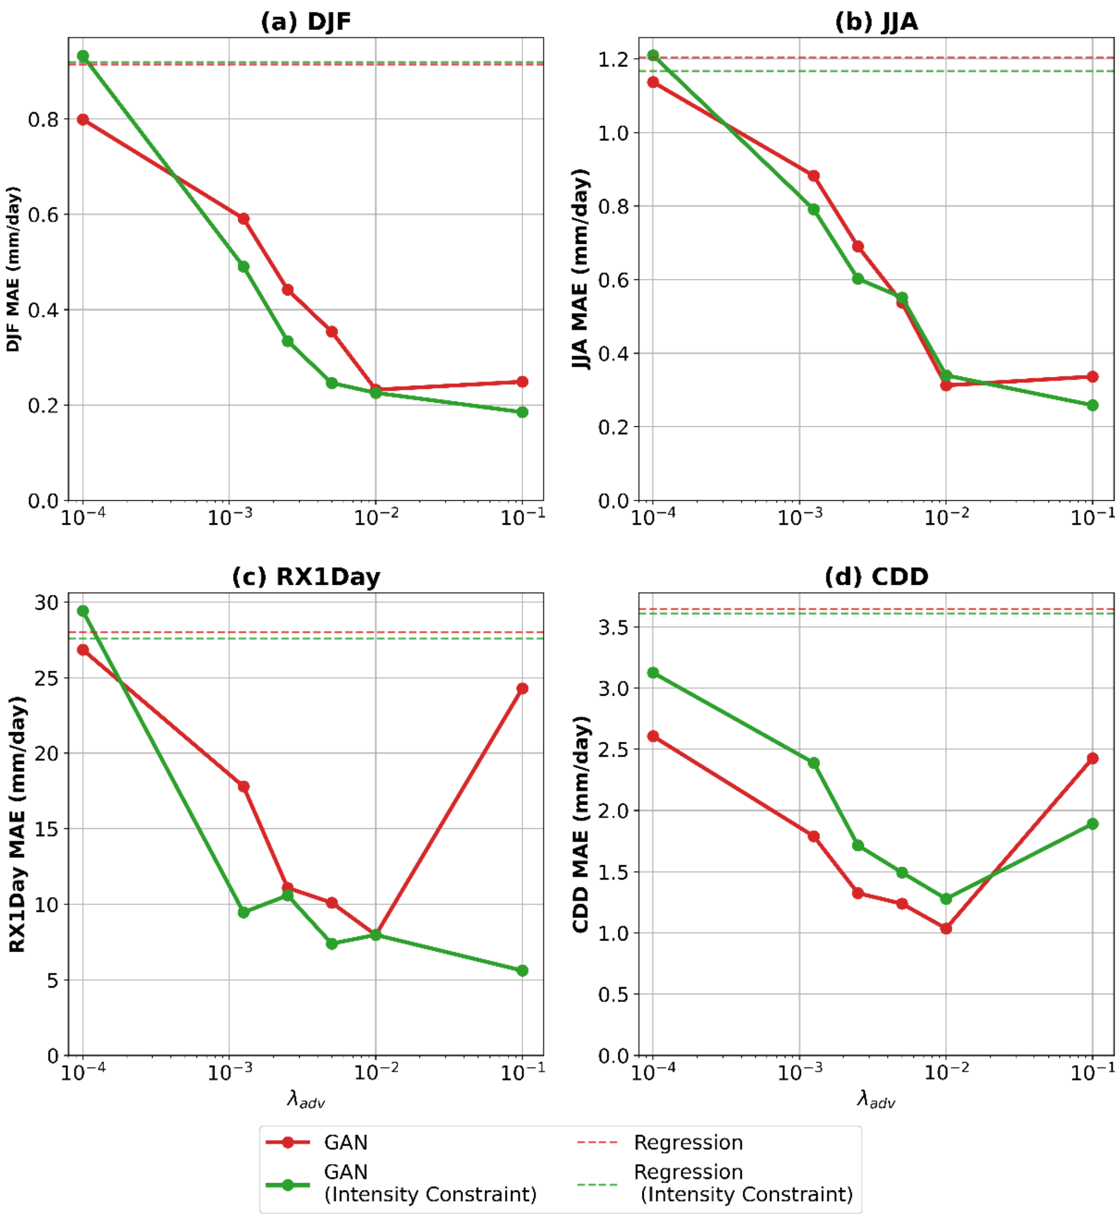

Figure 6: The MAE as a function of λ_{adv} for the GAN trained with (green) and without (red) the intensity constraint across four key statistics — mean DJF (a) and JJA (b) precipitation, RX1Day (c), and CDD (d) — relative to ground truth CCAM RCM simulation from EC-Earth3. The performance of the regression baseline is shown as the dashed line, both with (green) and without (red) the intensity constraint.

3.3.3 Consecutive Dry Days

The results for CDD show the same trends as those for Rx1Day, except at λ_{adv}=0.1 for where the MAE abruptly increases for both loss function configurations. Upon visual inspection in Figure 8b, the MAE increase appears to be due to an overestimation in CDD over the ocean, particularly on the eastern coast of the South Island and the northern region of the North Island of New Zealand. Interestingly, the configuration with the intensity constraint appears to have a larger MAE across all values of λ_{adv} compared to the configuration without it.

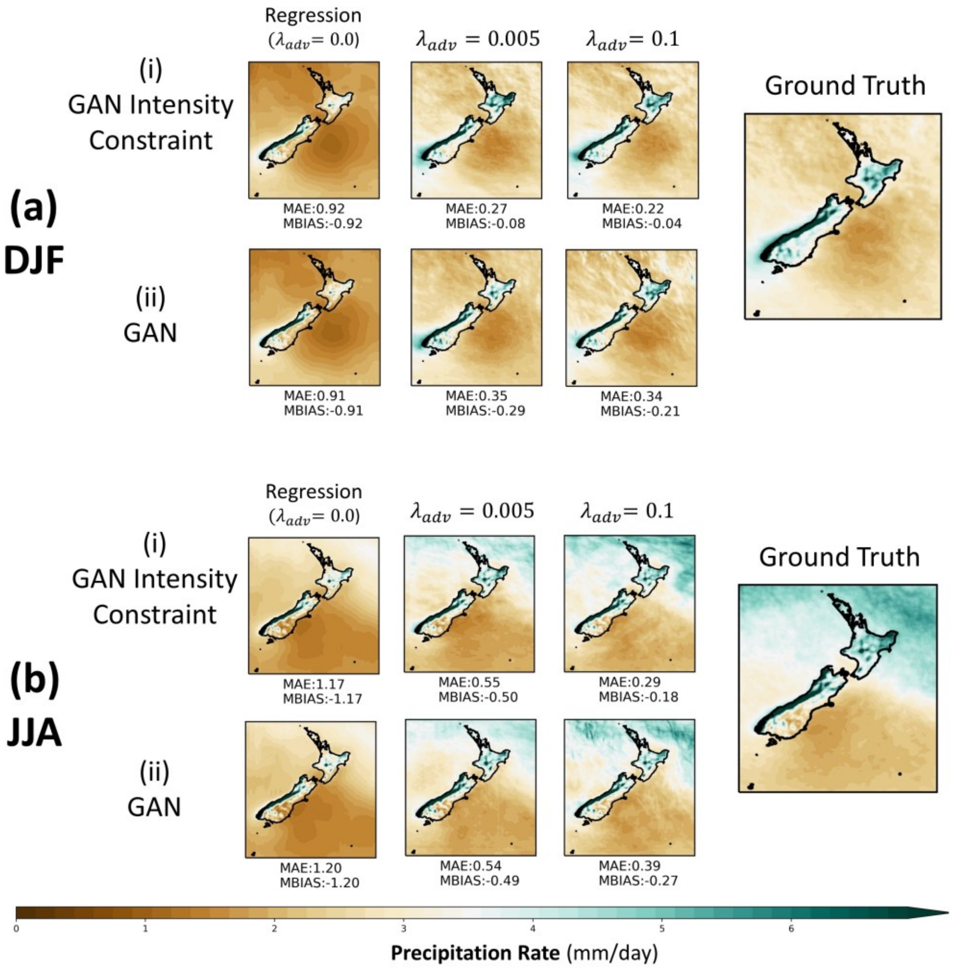

Figure 7: The performance of the two GAN loss function configurations as a function of λ_{adv}; with (i) and without the intensity constraint (ii) in generating DJF and JJA climatological precipitation relative to ground truth CCAM RCM simulations (EC-Earth3) for a single ensemble member. The regression baseline is indicated by λ_{adv}= 0.0. The text for each subplot shows the MAE and the mean bias (MBIAS) relative to ground truth.

3.3.4 In-sample Performance

It is also important to discuss in-sample performance, that is, evaluating the emulator on the same RCM simulation as it was trained on (ACCESS-CM2) between 1986-2005. Differences between in-sample and out-of-sample performance can shed light on the emulator's ability to generalize further (for example, to other GCMs). The in-sample performance across the four metrics is generally significantly better (lower error) than on EC-Earth3, particularly for the regression baseline and lower values of λ_{adv} (≤ 0.01). The higher in-sample performance suggests that the algorithm may have slightly overfitted to the ACCESS-CM2 training distribution despite efforts to prevent it. However, for λ_{adv} =0.1

with intensity constraint, in-sample and out-of-sample performances are similar for all metrics except for CDD, as illustrated in Supplementary Figure S9-S10. One potential explanation is that adversarial training mitigates overfitting, allowing the algorithm to learn more generalizable relationships, though further research would be required to test this. We also assessed the out-of-sample emulator performance on the NorESM2-MM GCM (i.e. the model trained on ACCESS-CM2 is applied to NorESM2-MM predictor fields). The results were nearly identical to the EC-Earth3 evaluation, as summarized in Supplementary Figure S8. This result is important as is implies a GAN emulator trained only on one RCM/GCM simulation pair can be broadly applied to historical climates from other GCMs, a finding that differs from the common view of GANs as being unstable.

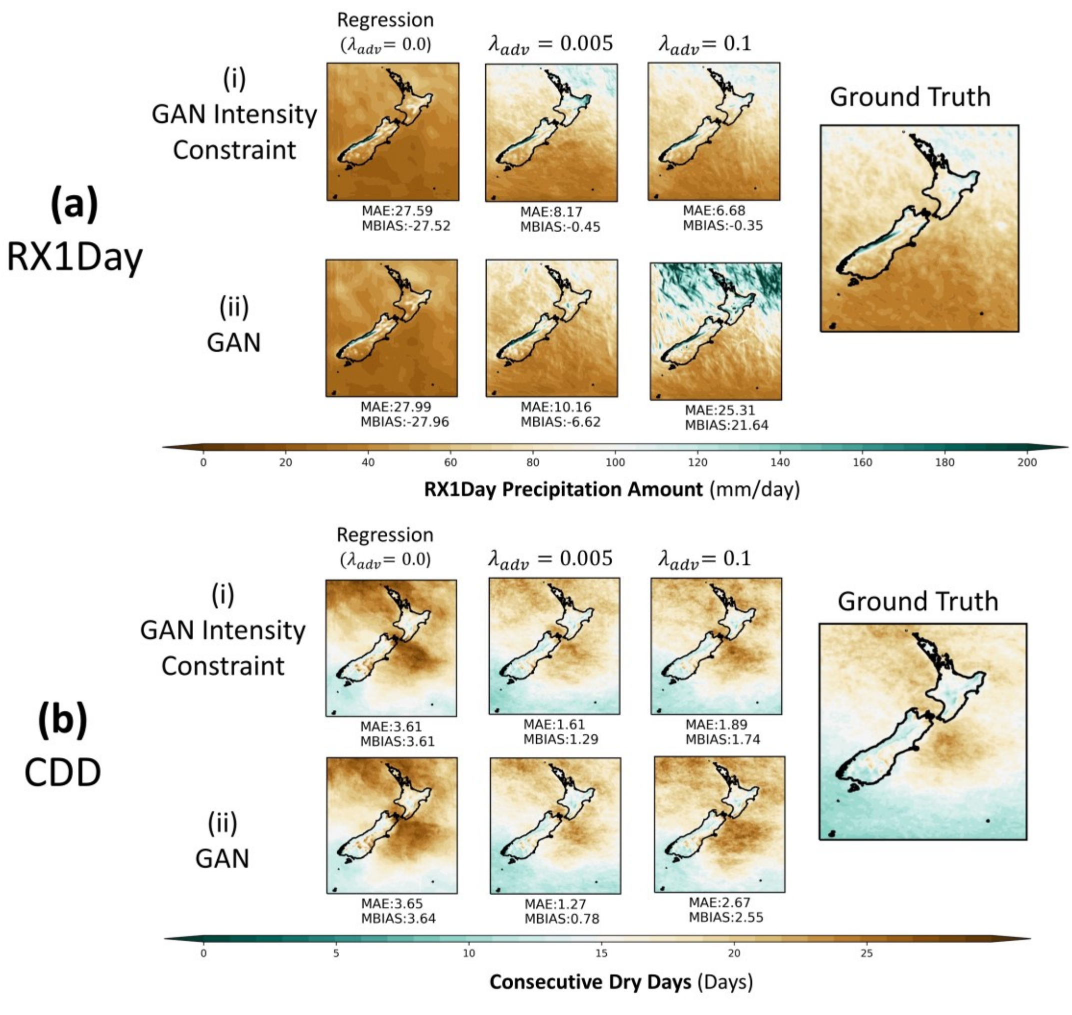

Figure 8: The performance of the two GAN loss function configurations as a function of λ_{adv}; with (i) and without the intensity constraint (ii) in generating climatological RX1Day and CDD relative to ground truth CCAM RCM simulations (EC-Earth3) for a single ensemble member. The regression baseline is indicated by λ_{adv}= 0.0. The text for each subplot shows the MAE and the mean bias (MBIAS) relative to ground truth.

3.3.5 Summary

Overall, when considering all climate statistical metrics, the lowest MAE scores occur when λ_{adv} is set between 0.05 and 0.1 with the intensity constraint. Note, that for this range of λ_{adv} we also see good performance in accurately capturing precipitation distribution. In comparison, without the intensity constraint, the best performance is generally achieved when λ_{adv} is between 0.0025 to 0.01, with the lowest scores notably at $\lambda_{adv} = 0.01$. However, the larger values of λ_{adv} within this range (0.005, 0.01) do not accurately capture precipitation distribution (as detailed in Section 3.2), making $\lambda_{adv} = 0.0025$ the only viable option.

3.4 Ensemble Statistics

Moving beyond standard downscaling metrics, this section assesses the ensemble spread produced by GANs. It aims to determine whether GANs can skillfully generate ensembles that capture the "true" variability of potential outcomes that is essential for uncertainty quantification in a downscaling or weather generation context. This study uses the spread-error relationship (section 3.4.1) and the Continuous Ranked Probability Score (CRPS; section 3.4.2) metrics, which are commonly used for evaluating ensemble weather forecasts (i.e. Doblas-Reyes et al., 2005; Leutbecher & Palmer, 2008; Palmer et al., 2008), and more recently for DL-based weather forecasts (Harris et al., 2022; Kochkov et al., 2024; Price & Rasp, 2022; Ravuri et al., 2021a; Vosper et al., 2023).

3.4.1 Spread-error Relationship

The spread-error relationship evaluates an ensemble's dispersion – also commonly known as the ensemble's calibration. The spread-error relationship describes a relation between spread of the ensemble about its mean (RMSS) and the error in the ensemble mean (hereon referred to as RMSE) (Doblas-Reyes et al., 2005; Fortin et al., 2014; Leutbecher & Palmer, 2008; Palmer et al., 2008). A well-calibrated (statistically perfect) ensemble of infinite size generally has a linear spread-error relationship (black dashed line in Figure 9a-b), meaning that the average distance between the ground truth and the ensemble mean equals the average distance between individual ensemble members and the ensemble mean. The key characteristic of a well-calibrated ensemble is that individual ensemble members are not statistically distinguishable from the ground truth data. This relationship has been widely used in the ensemble weather forecasting (Leutbecher & Palmer, 2008), and has recently examined in a downscaling context (Vosper et al., 2023).

The spread-error relationship is computed for a 10-member ensemble spanning the 20-year evaluation period for each λ_{adv}. Each ensemble member is distinguished by a unique noise vector for the same large-scale predictor variables, as depicted in Figures 2b and c. We also generated a 100-member ensemble spanning one year (i.e. larger ensemble but shorter duration) to understand the impact of ensemble size on the spread-error relationship, which did not alter our findings (not shown). Similar to previous studies (Kochkov et al., 2024; Vosper et al., 2023), to compute spread error curves, we first average RMSS and RMSE values over time, then compute the mean RMSS across all RMSE bins. Due to our smaller ensemble size ($n = 10$), the RMSS and RMSE values are adjusted by the factors 1.11 ($\frac{n}{n-1}$) and 0.9 ($\frac{n}{n+1}$), respectively, as outlined in Vosper et al., (2023) & Leutbecher and Palmer, (2008).

In the regression baseline ($\lambda_{adv} = 0.0$) and $\lambda_{adv} = 0.0001$ cases, there is no spread amongst their ensemble members (the RMSS is zero) and thus the slope of the spread-error relationship is infinite (Figure 9a-b). When the intensity constraint is used, increasing λ_{adv} improves dispersion or calibration, and when λ_{adv} is between 0.005 and 0.1 the spread-error curve and its slope are more closely aligned with the perfect ensemble ($y = x$). Visual inspection of an individual case (Figure 9c) likewise shows that for small values of λ_{adv} no dispersion is evident among ensemble members, while for λ_{adv} greater than 0.005 dispersion becomes more pronounced. Most importantly the dispersion appears perceptually realistic, where each member's precipitation patterns are different but all consistent with large-scale circulation patterns.

Conversely, when no intensity constraint is used, for nearly all values of λ_{adv} the spread-error curves are primarily under-dispersive or poorly calibrated, as illustrated in Figure 9b. The ensemble is well-calibrated when $\lambda_{adv} = 0.01$ as the spread-error curves are close to a perfect ensemble ($y = x$), but rapidly transitions to being over dispersive for larger $\lambda_{adv}(\geq 0.01)$. Overall, in the absence of the intensity constraint, there is some instability or heightened sensitivity to the spread-error relationship as a function of λ_{adv}.

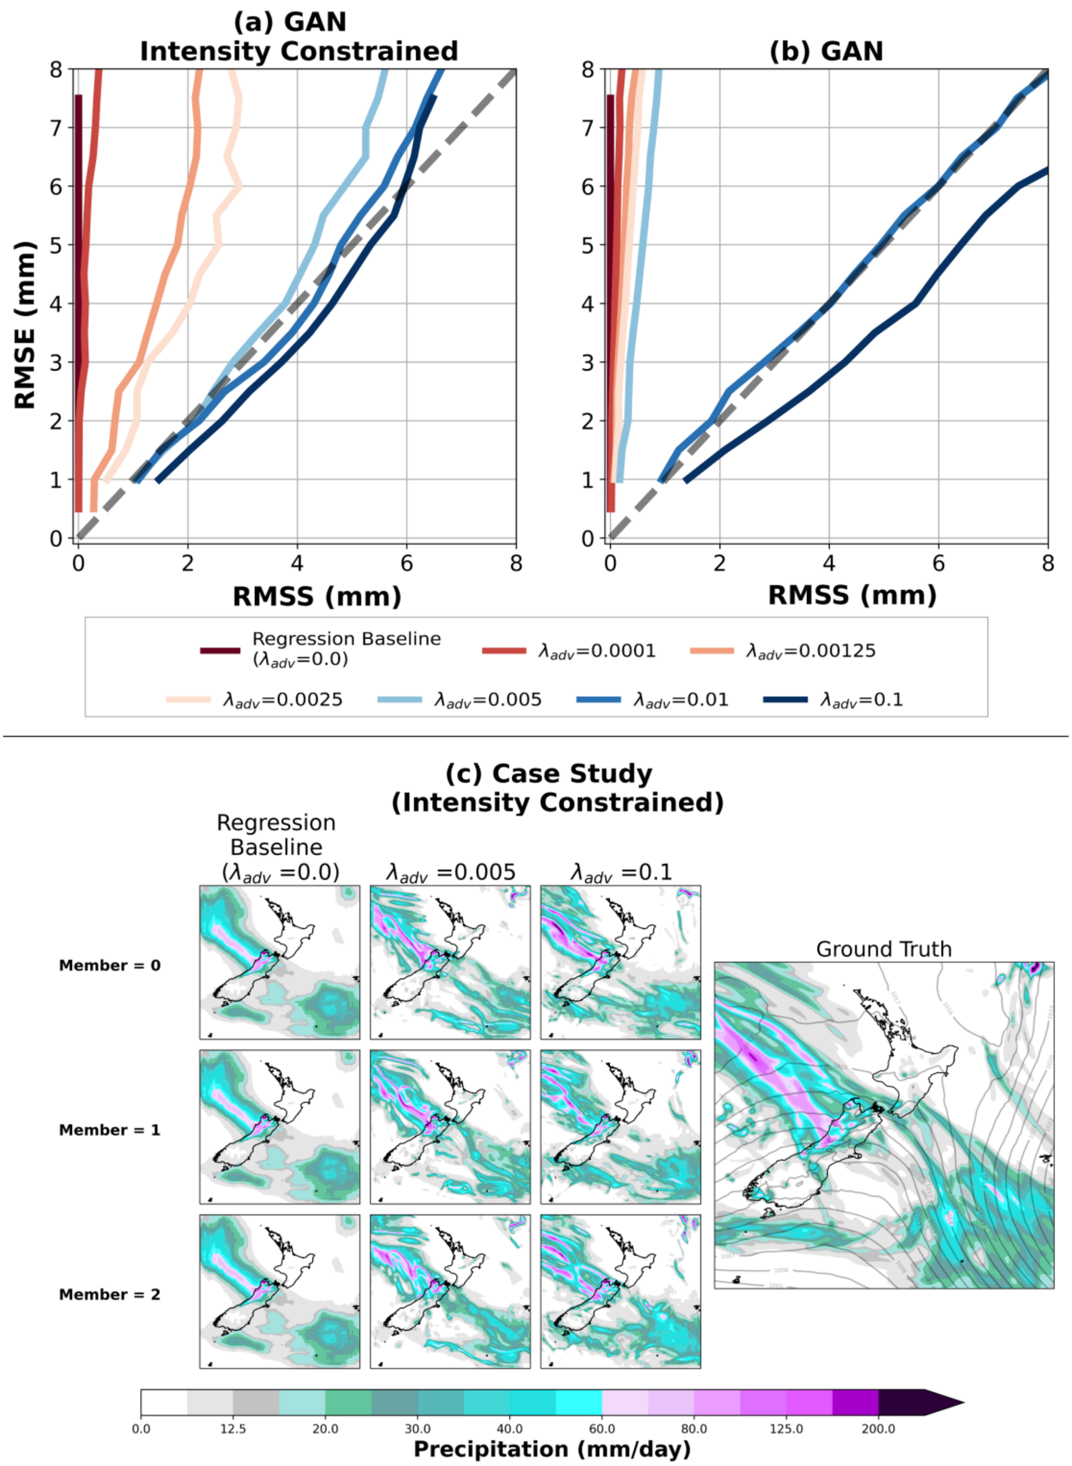

Figure 9: (a) The spread-error relationship as a function of λ_{adv} when the intensity constraint is used. (b) The spread-error relationship as a function of λ_{adv} when intensity constraint is not used. (c) Examples of solutions from three ensemble members across all values of λ_{adv}, and its corresponding ground truth CCAM precipitation (b). (c) shows a case study (2004-01-16) from EC-Earth3 to illustrate dispersion across three different ensemble members.

3.4.2 Continuous Ranked Probability Score (CRPS)

The Continuous Ranked Probability Score (CRPS) is a proper scoring rule that assesses the accuracy of an ensemble in representing the full range of uncertainty within a prediction (Gneiting & Katzfuss, 2014; Gneiting & Raftery, 2007; Hersbach, 2000; Lerch et al., 2017; Matheson & Winkler, 1976). The CRPS measures the distance between the predicted probability distribution and ground truth, but also assesses the ensemble's spread or calibration. The CRPS is often interpreted as a generalization of the MAE (absolute difference between a prediction and ground truth) for probabilistic forecast evaluation. In the case of a deterministic prediction (e.g., regression baselines or a single ensemble member) the CRPS equals the MAE, allowing for a comparison between deterministic and ensemble predictions. One notable advantage of the CRPS is that it is less sensitive to double counting of position (location of where precipitation occurs) and intensity (precipitation amount) errors, which is a widely known limitation of MAE in ensemble forecast evaluation (Hersbach, 2000).

Both the MAE and CRPS are computed on an individual prediction basis (per grid cell) and averaged across all timesteps (7300 timesteps), latitudes, and longitudes. Here, the MAE is calculated for each ensemble member and then averaged across all members. To reduce the effect of outlier precipitation values on the computation of the MAE and CRPS, we exclude grid points for a given timestep (and all corresponding members) when at least one ensemble member has a precipitation value exceeding 2000 mm.

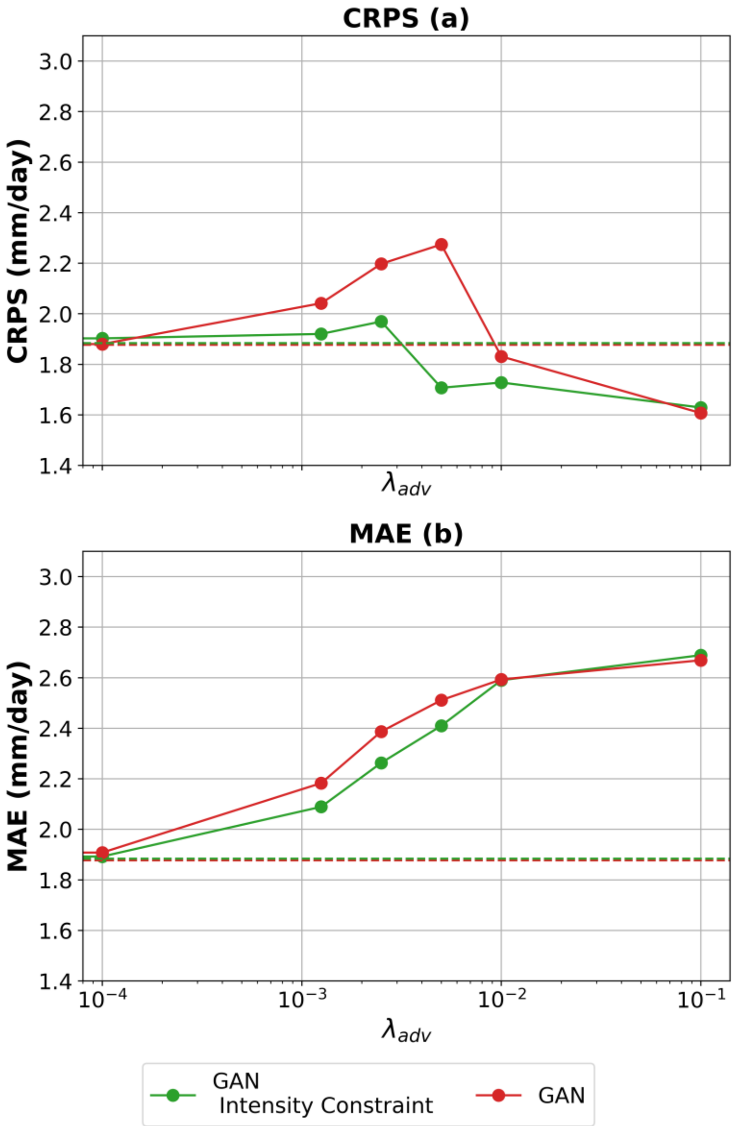

Figure 10: The CRPS (a) and MAE (b) as a function of λ_{adv} for both loss function configurations on EC-Earth3 relative to ground truth, with (green) and without the intensity constraint (red).

The MAE (Figure 10b) is lowest for the regression baseline ($\lambda_{adv} = 0$) and increases rapidly as a function of λ_{adv}, where it is 50% greater when $\lambda_{adv} = 0.1$ for both loss function configurations. The regression baseline's lower MAE is expected, as it directly optimizes for the Mean Squared Error (MSE) during training, which aligns closely with the MAE metric. This relationship between MAE and λ_{adv} is also somewhat expected, as by design larger λ_{adv} values allow for more deviation from the regression baselines (which are optimized for MSE), leading to an increased MAE. Several studies have also noted that GANs have a higher MAE scores than regression-based DL algorithms (i.e. J. Wang et al., 2021).

As for the CRPS metric, GANs only outperform the regression baseline at certain values of λ_{adv}. For instance, the GAN's CRPS is larger than the regression baseline when

λ_{adv} ≤0.0025 with the intensity constraint, and λ_{adv} ≤0.005 without it (Figure 10a). On the other hand, the GAN's CRPS is lower than the regression baseline when λ_{adv} ≥0.005 with the intensity constraint, and λ_{adv} ≥0.01 without it, where in both loss configurations the lowest CRPS is achieved when $\lambda_{adv} = 0.1$. Note, the CRPS scores with the intensity constraint are typically lower than those without across all λ_{adv} values, except at λ_{adv} = 0.1, where the scores are similar.

3.4.3 Summary

In summary of our results from both CRPS and spread-error metrics, we find that smaller values of λ_{adv} (<0.05) tend to generate under-dispersive (poorly calibrated) ensembles with larger or similar CRPS scores to the regression baseline for both loss configurations. When λ_{adv} is between 0.005 and 0.1, GANs trained with an intensity constraint generate dispersive (well-calibrated) ensembles with lower CRPS scores than the regression-baseline. However, GANs trained without an intensity constraint produce well-calibrated ensembles and lower CRPS scores than the regression-baseline only when $\lambda_{adv} = 0.01$ and become over-dispersive for larger λ_{adv}. Additionally, GANs trained with an intensity constraint are more dispersive and have lower CRPS scores across all λ_{adv} than those trained without it, suggesting that the intensity constraint improves robustness beyond its intended design.

4 Discussion

4.1 The Importance of Constraints

In general, smaller values of λ_{adv} (typically below 0.005) are common amongst downscaling (climate and weather) and super-resolution studies when no intensity constraint is used (Harris et al., 2022; Ledig et al., 2017; Leinonen et al., 2021; Vosper et al., 2023; X. Wang et al., 2018). Our optimal range of λ_{adv} ($0.00125 \leq \lambda_{adv} < 0.005$) without the intensity constraint aligns with these values. Our findings show that this range of λ_{adv} balances good performance in capturing rainfall mean variations (Figure 7 & 8) and distribution e.g. for extreme events (Figure 4). While larger values of λ_{adv} (>0.005) perform better on the former, they drastically overestimate extreme precipitation events (>200mm/day). As λ_{adv} becomes too small, GAN performance converges towards that of regression-based DL algorithm, which generally performs poorly across all metrics.

The agreement between our optimal λ_{adv} range (without the intensity constraint) and other studies is promising, but one should be cautious about this range of λ_{adv} as they have not been properly assessed in literature for their performance in climate settings (i.e. how well they capture climate statistics). Our results demonstrate that GANs within this λ_{adv} range produce under-dispersive ensembles (Figure 9b), limiting their usefulness for uncertainty quantification (see also section 4.2). Additionally, their errors on climate statistical metrics are much higher than larger λ_{adv} values (Figure 6-9). More broadly, there are other challenges with training without the intensity constraint, such as the case for large λ_{adv}, where precipitation extremes significantly overestimated. This raises concerns about

GAN robustness (without the intensity constraint) under climate change, due to potential unreliability in simulating extreme events.

Our study also shows that an intensity constraint in the loss function improves the robustness of GANs and allows training with large λ_{adv}. Larger values of λ_{adv} ($\geq$0.005) generate more dispersive ensembles in the results and improve accuracy in climate statistical metrics compared to smaller λ_{adv} values (below 0.005), all while accurately representing the precipitation distribution. Several studies have also incorporated intensity constraints into GAN loss functions. These studies have used significantly larger values of λ_{adv} (e.g., Ravuri et al., 2021: $\lambda_{adv} = 0.05$; Price & Rasp, 2022: $\lambda_{adv} = 0.1$). They reported substantial improvement over regression-based DL algorithms, focusing primarily on metrics such as CRPS and performance on extreme events in a weather forecasting context. However, they did not directly compare their results to those without an intensity constraint.

4.2 Stochastic Weather Generation with GANs

The application of GANs as a stochastic weather generator remains both under-utilized and under-evaluated in climate science. Stochastic weather generators can generate large ensembles (or sequences) of climate fields (i.e. Ailliot et al., 2015; Benoit et al., 2018; Furrer & Katz, 2008; Steinschneider et al., 2019; Verdin et al., 2018), which can be used to estimate the likelihood of a certain extreme event occurring (i.e. average recurrence interval), thereby offering valuable insights for disciplines such as catastrophe modeling. Recently, several studies have used generative DL algorithms (including GANs) in a similar capacity to stochastic weather generators (Boulaguiem et al., 2022; Brochet et al., 2023; Peard & Hall, 2023; Sha et al., 2024). GANs may have certain benefits over stochastic weather generators, such as their ability to learn complex spatio-temporal relationships (Sha et al., 2024). This may help them better simulate extreme phenomena like cyclones and atmospheric rivers, though further comparison with traditional stochastic weather generators is needed. Although GANs show promise in this context, their success is ultimately hinged on their ability to generate sufficiently dispersive ensembles (that capture the true variability of all possible outcomes).

Several studies, which have mainly focused on weather forecasting have assessed the calibration (dispersion) of GAN-generated ensembles (i.e. Harris et al., 2022; Price & Rasp, 2022; Ravuri et al., 2021b; Vosper et al., 2023). Collectively, these studies suggest that GANs can produce well-calibrated ensemble predictions across a large range of λ_{adv} (0.001-0.1), and thus for this purpose we cannot expect a single value of λ_{adv} to work across all problems and regions. However, our study introduces key insights into using GANs for uncertainty quantification not previously detailed in literature. Firstly, our study emphasizes the importance of exploring the λ_{adv} parameter, due to its significant impact on ensemble dispersion (calibration). Secondly, incorporating constraints (i.e. intensity constraints) to the loss function can not only improve ensemble dispersion across λ_{adv}, but also yields more robust performance compared to traditional GAN implementations.

4.3 Limitations

Our study has only focused on historical training and evaluation. Further research should focus on considering how well GANs extrapolating to future scenarios, especially in warmer climates, may require broader training across both historical and future periods, as well as multiple RCM simulations (Bano-Medina et al., 2023; Chadwick et al., 2011; Doury et al., 2022; Holden et al., 2015). The choice of training simulation may impact the algorithm's ability to extrapolate to future climates across multiple GCMs (Bano-Medina et al., 2023; Rampal et al., 2024). For example, warmer RCM simulations (i.e. with a higher equilibrium climate sensitivity), may offer greater diversity in extreme events and climate variability. When assessing emulator performance in future climates, one should consider evaluating the emulator's ability to reproduce the RCM's climate change signals and non-stationary changes like trends in extreme precipitation. Examples of these evaluation strategies are provided in Bano-Medina et al. (2023), Rampal et al. (2024), and Doury et al. (2022).

Further development of statistical constraints incorporated into the loss function should also be considered. In our case, the intensity constraint configuration performs exceptionally well across various evaluation metrics but appears to have a lower skill for CDD. A potential explanation for this lower skill could relate to the concept of metric transitivity (Abramowitz et al., 2019), in which optimizing the algorithm to perform well on specific metrics (i.e. intensity) means it performs slightly worse on other metrics which depend more on the temporal aspects of precipitation (i.e. CDD). In future work, applying additional constraints tailored for CDD could potentially improve the skill for this metric.

Our research has only focused on downscaling within the New Zealand domain, and thus it is unclear how generalizable our intensity constraint modification and optimal λ_{adv} value is across different domains, especially those larger in size (i.e. CORDEX domains) and involving various variables. Although not detailed here, preliminary evidence, which will be explored in a subsequent study, indicate that this optimal λ_{adv} range with the intensity constraint successfully downscales precipitation in different regions and for other variables (i.e. temperature), though further testing is needed to confirm its robustness.

Lastly, it is important to highlight common criticisms of GANs, such as "mode collapse" or a lack of diversity in generated samples (Che et al., 2017; Dubiński et al., 2023; Mao et al., 2019; Salimans et al., 2016; Srivastava et al., 2017). While we acknowledge such criticisms, our study suggests that GANs can be very effective in downscaling with careful training strategies (as detailed in Section 2). While diffusion models are an emerging type of downscaling technique (and stable) (Addison et al., 2022, Leinonen et al., 2023), they are significantly slower that GANs in training and inference time.

5 Conclusion

This study demonstrates that conditional Generative Adversarial Networks (GANs) can improve upon several of the limitations of regression-based deep learning (DL) algorithms for downscaling in a historical climate setting. We also highlighted the broader potential of GANs for stochastic weather generation, noting their skill in generating ensembles that accurately encompass the full spectrum of possible outcomes.

We trained a series of GANs on a single historical RCM simulation (ACCESS-CM2) and tested their performance on two completely unseen GCMs (EC-Earth3 and NorESM2-MM) to assess their generalization potential for downscaling across different GCM/RCM combinations.

The best-performing GANs examined here outperformed regression-based DL algorithms across various metrics relative to ground truth RCM simulations. While previous studies have found promising results using GANs for a few problems with limited metrics, we measure skill across a wide range of metrics, extending beyond conventional error metrics (e.g. mean absolute error) used in many DL studies. Crucially we examine climate statistical metrics (climatology of seasonal precipitation, wettest day of the year and length of the longest dry spell), temporal variability, the precipitation intensity distribution, and ensemble calibration (dispersion), which are much more relevant for climate studies.

We investigated how the hyperparameter λ_{adv} (weighting of the adversarial loss) impacts the skill of the GAN, which has been largely unexplored in literature. GAN performance was strongly dependent on λ_{adv} in standard implementations, such that λ_{adv} cannot be too big or too small (i.e. there is no convergence to good behavior in either limit), and selecting an optimal value requires trade-offs. Larger values of λ_{adv} (≥ 0.01) would perform well across most metrics but can drastically overestimate precipitation intensity which diverges monotonically as λ_{adv} increases. Smaller values would perform well on precipitation intensity but less well for climate statistics and generating well-calibrated (dispersive) ensembles needed to assess uncertainty. In this situation we cannot be confident that a value of λ_{adv} tuned to work well in the historical climate or situation would generalize to an unobserved climate scenario.

However, by incorporating a simple intensity constraint into the loss function of the GAN, we significantly improved the robustness of GAN performance, thereby requiring fewer trade-offs when selecting an optimal λ_{adv}. The intensity constraint allows for the selection of larger λ_{adv} (≥ 0.005), hence a stronger weighting of the adversarial loss, which performs well across all evaluation metrics, including precipitation intensity, but can also generate well-calibrated (dispersive) ensembles required for stochastic weather generation.

While we found an optimal range of λ_{adv} between 0.005 and 0.1, we strongly recommend thoroughly exploring and testing this hyperparameter during training in different contexts (i.e. across different regions and variables). We also emphasize the importance of statistical constraints for tailored GAN design and the use of climate-relevant evaluation metrics. Further work will be required to see whether the range of $\boldsymbol{\lambda_{adv}}$ found to succeed here indeed generalizes to both future climates and across multiple variables.

Acknowledgments.

NR, SS, SH, GA acknowledge the support of the Australian Research Council Centre of Excellence for Climate Extremes (CLEX; CE170100023). NR and PG received funding from the New Zealand MBIE Endeavour Smart Ideas Fund (C01X2202). The authors would also like to acknowledge the New Zealand eScience Infrastructure for providing access to GPUs. The authors declare no conflicts of interest.

Data Availability Statement

The code and datasets supporting this study are accessible to the public. The code can be found on GitHub (https://github.com/nram812/A-Robust-Generative-Adversarial-Network-Approach-for-Climate-Downscaling), and the training and validation datasets are available on Zenodo (https://doi.org/10.5281/zenodo.10889046).

References

Abu-Srhan, A., Abushariah, M. A. M., & Al-Kadi, O. S. (2022). The effect of loss function on conditional generative adversarial networks. *Journal of King Saud University - Computer and Information Sciences*, *34*(9), 6977–6988. https://doi.org/10.1016/j.jksuci.2022.02.018

Ackerley, D., Dean, S., Sood, A., & Mullan, A. B. (2012). Regional climate modelling in New Zealand: Comparison to gridded and satellite observations. *Weather and Climate*, *32*(1), 3–22. https://doi.org/10.2307/26169722

Ailliot, P., Allard, D., Monbet, V., & Naveau, P. (2015). Stochastic weather generators: An overview of weather type models. *Journal de La Société Française de Statistique*, *156*(1), 101–113.

Annau, N. J., Cannon, A. J., & Monahan, A. H. (2023). Algorithmic Hallucinations of Near-Surface Winds: Statistical Downscaling with Generative Adversarial Networks to Convection-Permitting Scales. *Artificial Intelligence for the Earth Systems*, *2*(4). https://doi.org/10.1175/AIES-D-23-0015.1

Arjovsky, M., Chintala, S., & Bottou, L. (2017). *Wasserstein GAN* (arXiv:1701.07875). arXiv. https://doi.org/10.48550/arXiv.1701.07875

Babaousmail, H., Hou, R., Gnitou, G. T., & Ayugi, B. (2021). Novel statistical downscaling emulator for precipitation projections using deep Convolutional Autoencoder over Northern Africa. *Journal of Atmospheric and Solar-Terrestrial Physics*, *218*, 105614. https://doi.org/10.1016/j.jastp.2021.105614

Bailie, T., Koh, Y. S., Rampal, N., & Gibson, P. B. (2024). Quantile-Regression-Ensemble: A Deep Learning Algorithm for Downscaling Extreme Precipitation. *Proceedings of the AAAI Conference on Artificial Intelligence*, *38*(20), Article 20. https://doi.org/10.1609/aaai.v38i20.30193

Bano-Medina, J., Iturbide, M., Fernandez, J., & Gutierrez, J. M. (2023). *Transferability and explainability of deep learning emulators for regional climate model projections: Perspectives for future applications* (arXiv:2311.03378). arXiv. http://arxiv.org/abs/2311.03378

Baño-Medina, J., Manzanas, R., & Gutiérrez, J. M. (2020). Configuration and intercomparison of deep learning neural models for statistical downscaling. *Geoscientific Model Development*, *13*(4), 2109–2124. https://doi.org/10.5194/gmd-13-2109-2020

Bartók, B., Wild, M., Folini, D., Lüthi, D., Kotlarski, S., Schär, C., Vautard, R., Jerez, S., & Imecs, Z. (2017). Projected changes in surface solar radiation in CMIP5 global climate models and in EURO-CORDEX regional climate models for Europe. *Climate Dynamics*, *49*(7), 2665–2683. https://doi.org/10.1007/s00382-016-3471-2

Benestad, R. E. (2004). Empirical-statistical downscaling in climate modeling. *Eos, Transactions American Geophysical Union*, *85*(42), 417–422. https://doi.org/10.1029/2004EO420002

Benestad, R. E. (2010). Downscaling precipitation extremes: Correction of analog models through PDF predictions. *Theoretical and Applied Climatology*, *100*(1–2), 1–21. https://doi.org/10.1007/s00704-009-0158-1

Benoit, L., Allard, D., & Mariethoz, G. (2018). Stochastic Rainfall Modeling at Sub-kilometer Scale. *Water Resources Research*, *54*(6), 4108–4130. https://doi.org/10.1029/2018WR022817

Boé, J., Mass, A., & Deman, J. (2023). A simple hybrid statistical–dynamical downscaling method for emulating regional climate models over Western Europe. Evaluation, application, and role of added value? *Climate Dynamics*, *61*(1), 271–294. https://doi.org/10.1007/s00382-022-06552-2

Boé, J., Somot, S., Corre, L., & Nabat, P. (2020). Large discrepancies in summer climate change over Europe as projected by global and regional climate models: Causes and consequences. *Climate Dynamics*, *54*(5), 2981–3002. https://doi.org/10.1007/s00382-020-05153-1

Boulaguiem, Y., Zscheischler, J., Vignotto, E., Van Der Wiel, K., & Engelke, S. (2022). Modeling and simulating spatial extremes by combining extreme value theory with generative adversarial networks. *Environmental Data Science*, *1*, e5. https://doi.org/10.1017/eds.2022.4

Brochet, C., Raynaud, L., Thome, N., Plu, M., & Rambour, C. (2023). Multivariate Emulation of Kilometer-Scale Numerical Weather Predictions with Generative Adversarial Networks: A Proof of Concept. *Artificial Intelligence for the Earth Systems*, *2*(4). https://doi.org/10.1175/AIES-D-23-0006.1

Chadwick, R., Coppola, E., & Giorgi, F. (2011). An artificial neural network technique for downscaling GCM outputs to RCM spatial scale. *Nonlinear Processes in Geophysics*, *18*(6), 1013–1028. https://doi.org/10.5194/npg-18-1013-2011

Chapman, S., Syktus, J., Trancoso, R., Thatcher, M., Toombs, N., Wong, K. K.-H., & Takbash, A. (2023). Evaluation of Dynamically Downscaled CMIP6-CCAM Models Over Australia. *Earth's Future*, *11*(11), e2023EF003548. https://doi.org/10.1029/2023EF003548

Che, T., Li, Y., Jacob, A. P., Bengio, Y., & Li, W. (2017). *Mode Regularized Generative Adversarial Networks* (arXiv:1612.02136). arXiv. https://doi.org/10.48550/arXiv.1612.02136

Deser, C., Phillips, A., Bourdette, V., & Teng, H. (2012). Uncertainty in climate change projections: The role of internal variability. *Climate Dynamics*, *38*(3), 527–546. https://doi.org/10.1007/s00382-010-0977-x

Deser, C., & Phillips, A. S. (2023). A range of outcomes: The combined effects of internal variability and anthropogenic forcing on regional climate trends over Europe. *Nonlinear Processes in Geophysics*, *30*(1), 63–84. https://doi.org/10.5194/npg-30-63-2023

Di Virgilio, G., Evans, J. P., Di Luca, A., Grose, M. R., Round, V., & Thatcher, M. (2020). Realised added value in dynamical downscaling of Australian climate change.

Climate Dynamics, *54*(11–12), 4675–4692. https://doi.org/10.1007/s00382-020-05250-1

Di Virgilio, G., Evans, J. P., Di Luca, A., Olson, R., Argüeso, D., Kala, J., Andrys, J., Hoffmann, P., Katzfey, J. J., & Rockel, B. (2019). Evaluating reanalysis-driven CORDEX regional climate models over Australia: Model performance and errors. *Climate Dynamics*, *53*(5–6), 2985–3005. https://doi.org/10.1007/s00382-019-04672-w

Doblas-Reyes, F. J., Hagedorn, R., & Palmer, T. N. (2005). The rationale behind the success of multi-model ensembles in seasonal forecasting – II. Calibration and combination. *Tellus A: Dynamic Meteorology and Oceanography*, *57*(3), 234. https://doi.org/10.3402/tellusa.v57i3.14658

Doury, A., Somot, S., Gadat, S., Ribes, A., & Corre, L. (2022). Regional climate model emulator based on deep learning: Concept and first evaluation of a novel hybrid downscaling approach. *Climate Dynamics*, *60*(5), 1751–1779. https://doi.org/10.1007/s00382-022-06343-9

Dubiński, J., Deja, K., Wenzel, S., Rokita, P., & Trzciński, T. (2023). *Selectively increasing the diversity of GAN-generated samples* (arXiv:2207.01561). arXiv. http://arxiv.org/abs/2207.01561

Feser, F., Rockel, B., von Storch, H., Winterfeldt, J., & Zahn, M. (2011). Regional Climate Models Add Value to Global Model Data: A Review and Selected Examples. *Bulletin of the American Meteorological Society*, *92*(9), 1181–1192. https://doi.org/10.1175/2011BAMS3061.1

Fortin, V., Abaza, M., Anctil, F., & Turcotte, R. (2014). Why Should Ensemble Spread Match the RMSE of the Ensemble Mean? *Journal of Hydrometeorology*, *15*(4), 1708–1713. https://doi.org/10.1175/JHM-D-14-0008.1

Fowler, H. J., Blenkinsop, S., & Tebaldi, C. (2007). Linking climate change modelling to impacts studies: Recent advances in downscaling techniques for hydrological modelling. *International Journal of Climatology*, *27*(12), 1547–1578. https://doi.org/10.1002/joc.1556

Furrer, E. M., & Katz, R. W. (2008). Improving the simulation of extreme precipitation events by stochastic weather generators. *Water Resources Research*, *44*(12). https://doi.org/10.1029/2008WR007316

Gensini, V. A., Haberlie, A. M., & Ashley, W. S. (2023). Convection-permitting simulations of historical and possible future climate over the contiguous United States. *Climate Dynamics*, *60*(1), 109–126. https://doi.org/10.1007/s00382-022-06306-0

Gibson, P. B., Rampal, N., Dean, S. M., & Morgenstern, O. (2024). Storylines for Future Projections of Precipitation Over New Zealand in CMIP6 Models. *Journal of Geophysical Research: Atmospheres*, *129*(5), e2023JD039664. https://doi.org/10.1029/2023JD039664

Gibson, P. B., Stone, D., Thatcher, M., Broadbent, A., Dean, S., Rosier, S. M., Stuart, S., & Sood, A. (2023). High-Resolution CCAM Simulations Over New Zealand and the South Pacific for the Detection and Attribution of Weather Extremes. *Journal of Geophysical Research: Atmospheres*, *128*(14), e2023JD038530. https://doi.org/10.1029/2023JD038530

Giorgi, F., Brodeur, C. S., & Bates, G. T. (1994). Regional Climate Change Scenarios over the United States Produced with a Nested Regional Climate Model. *Journal of*

Climate, *7*(3), 375–399. https://doi.org/10.1175/1520-0442(1994)007<0375:RCCSOT>2.0.CO;2

Giorgi, F., Jones, C., & Asrar, G. R. (2009). *Addressing climate information needs at the regional level: The CORDEX framework*.

Gneiting, T., & Katzfuss, M. (2014). Probabilistic Forecasting. *Annual Review of Statistics and Its Application*, *1*(Volume 1, 2014), 125–151. https://doi.org/10.1146/annurev-statistics-062713-085831

Gneiting, T., & Raftery, A. E. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, *102*(477), 359–378. https://doi.org/10.1198/016214506000001437

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative Adversarial Nets. *Advances in Neural Information Processing Systems*, *27*. https://proceedings.neurips.cc/paper_files/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html

Groenke, B., Madaus, L., & Monteleoni, C. (2020). ClimAlign: Unsupervised statistical downscaling of climate variables via normalizing flows. *Proceedings of the 10th International Conference on Climate Informatics*, 60–66. https://doi.org/10.1145/3429309.3429318

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. C. (2017). Improved Training of Wasserstein GANs. *Advances in Neural Information Processing Systems*, *30*. https://proceedings.neurips.cc/paper_files/paper/2017/hash/892c3b1c6dccd52936e27cbd0ff683d6-Abstract.html

Harris, L., McRae, A. T. T., Chantry, M., Dueben, P. D., & Palmer, T. N. (2022). A Generative Deep Learning Approach to Stochastic Downscaling of Precipitation Forecasts. *Journal of Advances in Modeling Earth Systems*, *14*(10), e2022MS003120. https://doi.org/10.1029/2022MS003120

Hawkins, E., & Sutton, R. (2009). The Potential to Narrow Uncertainty in Regional Climate Predictions. *Bulletin of the American Meteorological Society*, *90*(8), 1095–1108. https://doi.org/10.1175/2009BAMS2607.1

Hawkins, E., & Sutton, R. (2011). The potential to narrow uncertainty in projections of regional precipitation change. *Climate Dynamics*, *37*(1–2), 407–418. https://doi.org/10.1007/s00382-010-0810-6

Hersbach, H. (2000). Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems. *Weather and Forecasting*, *15*(5), 559–570. https://doi.org/10.1175/1520-0434(2000)015<0559:DOTCRP>2.0.CO;2

Hobeichi, S., Nishant, N., Shao, Y., Abramowitz, G., Pitman, A., Sherwood, S., Bishop, C., & Green, S. (2023). Using Machine Learning to Cut the Cost of Dynamical Downscaling. *Earth's Future*, *11*(3), e2022EF003291. https://doi.org/10.1029/2022EF003291

Holden, P. B., Edwards, N. R., Garthwaite, P. H., & Wilkinson, R. D. (2015). Emulation and interpretation of high-dimensional climate model outputs. *Journal of Applied Statistics*, *42*(9), 2038–2055. https://doi.org/10.1080/02664763.2015.1016412

Hoogewind, K. A., Baldwin, M. E., & Trapp, R. J. (2017). The Impact of Climate Change on Hazardous Convective Weather in the United States: Insight from High-Resolution Dynamical Downscaling. *Journal of Climate*, *30*(24), 10081–10100. https://doi.org/10.1175/JCLI-D-16-0885.1

Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2018). *Image-to-Image Translation with Conditional Adversarial Networks* (arXiv:1611.07004). arXiv. https://doi.org/10.48550/arXiv.1611.07004

Isphording, R. N., Alexander, L. V., Bador, M., Green, D., Evans, J. P., & Wales, S. (2023). A Standardized Benchmarking Framework to Assess Downscaled Precipitation Simulations. *Journal of Climate*, *1*(aop). https://doi.org/10.1175/JCLI-D-23-0317.1

Izumi, T., Amagasaki, M., Ishida, K., & Kiyama, M. (2022). Super-resolution of sea surface temperature with convolutional neural network- and generative adversarial network-based methods. *Journal of Water and Climate Change*, *13*(4), 1673–1683. https://doi.org/10.2166/wcc.2022.291

Jones, R. G., Murphy, J. M., & Noguer, M. (1995). Simulation of climate change over europe using a nested regional-climate model. I: Assessment of control climate, including sensitivity to location of lateral boundaries. *Quarterly Journal of the Royal Meteorological Society*, *121*(526), 1413–1449. https://doi.org/10.1002/qj.49712152610

Kochkov, D., Yuval, J., Langmore, I., Norgaard, P., Smith, J., Mooers, G., Klöwer, M., Lottes, J., Rasp, S., Düben, P., Hatfield, S., Battaglia, P., Sanchez-Gonzalez, A., Willson, M., Brenner, M. P., & Hoyer, S. (2024). *Neural General Circulation Models for Weather and Climate* (arXiv:2311.07222). arXiv. http://arxiv.org/abs/2311.07222

Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., & Shi, W. (2017). *Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network* (arXiv:1609.04802). arXiv. http://arxiv.org/abs/1609.04802

Leinonen, J., Nerini, D., & Berne, A. (2021). Stochastic Super-Resolution for Downscaling Time-Evolving Atmospheric Fields With a Generative Adversarial Network. *IEEE Transactions on Geoscience and Remote Sensing*, *59*(9), 7211–7223. https://doi.org/10.1109/TGRS.2020.3032790

Lerch, S., Thorarinsdottir, T. L., Ravazzolo, F., & Gneiting, T. (2017). Forecaster's Dilemma: Extreme Events and Forecast Evaluation. *Statistical Science*, *32*(1), 106–127. https://doi.org/10.1214/16-STS588

Leutbecher, M., & Palmer, T. N. (2008). Ensemble forecasting. *Journal of Computational Physics*, *227*(7), 3515–3539. https://doi.org/10.1016/j.jcp.2007.02.014

Liu, C., Ikeda, K., Rasmussen, R., Barlage, M., Newman, A. J., Prein, A. F., Chen, F., Chen, L., Clark, M., Dai, A., Dudhia, J., Eidhammer, T., Gochis, D., Gutmann, E., Kurkute, S., Li, Y., Thompson, G., & Yates, D. (2017). Continental-scale convection-permitting modeling of the current and future climate of North America. *Climate Dynamics*, *49*(1), 71–95. https://doi.org/10.1007/s00382-016-3327-9

Mao, Q., Lee, H.-Y., Tseng, H.-Y., Ma, S., & Yang, M.-H. (2019). *Mode Seeking Generative Adversarial Networks for Diverse Image Synthesis* (arXiv:1903.05628). arXiv. http://arxiv.org/abs/1903.05628

Maraun, D. (2016). Bias Correcting Climate Change Simulations—A Critical Review. *Current Climate Change Reports*, *2*(4), 211–220. https://doi.org/10.1007/s40641-016-0050-x

Maraun, D., Wetterhall, F., Ireson, A. M., Chandler, R. E., Kendon, E. J., Widmann, M., Brienen, S., Rust, H. W., Sauter, T., Themeßl, M., Venema, V. K. C., Chun, K. P., Goodess, C. M., Jones, R. G., Onof, C., Vrac, M., & Thiele-Eich, I. (2010). Precipitation downscaling under climate change: Recent developments to bridge the gap between dynamical models and the end user. *Reviews of Geophysics*, *48*(3). https://doi.org/10.1029/2009RG000314

Mardani, M., Brenowitz, N., Cohen, Y., Pathak, J., Chen, C.-Y., Liu, C.-C., Vahdat, A., Kashinath, K., Kautz, J., & Pritchard, M. (2023). *Generative Residual Diffusion Modeling for Km-scale Atmospheric Downscaling* (arXiv:2309.15214). arXiv. http://arxiv.org/abs/2309.15214

Matheson, J. E., & Winkler, R. L. (1976). Scoring Rules for Continuous Probability Distributions. *Management Science*, *22*(10), 1087–1096. https://doi.org/10.1287/mnsc.22.10.1087

McGregor, J. L., & Dix, M. R. (2008). An Updated Description of the Conformal-Cubic Atmospheric Model. In K. Hamilton & W. Ohfuchi (Eds.), *High Resolution Numerical Modelling of the Atmosphere and Ocean* (pp. 51–75). Springer. https://doi.org/10.1007/978-0-387-49791-4_4

Miralles, O., Steinfeld, D., Martius, O., & Davison, A. C. (2022). Downscaling of Historical Wind Fields over Switzerland Using Generative Adversarial Networks. *Artificial Intelligence for the Earth Systems*, *1*(4). https://doi.org/10.1175/AIES-D-22-0018.1

Mirza, M., & Osindero, S. (2014). *Conditional Generative Adversarial Nets* (arXiv:1411.1784). arXiv. https://doi.org/10.48550/arXiv.1411.1784

Nishant, N., Hobeichi, S., Sherwood, S. C., Abramowitz, G., Shao, Y., Bishop, C., & Pitman, A. J. (2023). Comparison of a novel machine learning approach with dynamical downscaling for Australian precipitation. *Environmental Research Letters*. https://doi.org/10.1088/1748-9326/ace463

Oyama, N., Ishizaki, N. N., Koide, S., & Yoshida, H. (2023). Deep generative model super-resolves spatially correlated multiregional climate data. *Scientific Reports*, *13*(1), 5992. https://doi.org/10.1038/s41598-023-32947-0

Palmer, T. N., Doblas-Reyes, F. J., Weisheimer, A., & Rodwell, M. J. (2008). Toward Seamless Prediction: Calibration of Climate Change Projections Using Seasonal Forecasts. *Bulletin of the American Meteorological Society*, *89*(4), 459–470. https://doi.org/10.1175/BAMS-89-4-459

Peard, A., & Hall, J. (2023). *Combining deep generative models with extreme value theory for synthetic hazard simulation: A multivariate and spatially coherent approach* (arXiv:2311.18521). arXiv. https://doi.org/10.48550/arXiv.2311.18521

Prein, A. F., Langhans, W., Fosser, G., Ferrone, A., Ban, N., Goergen, K., Keller, M., Tölle, M., Gutjahr, O., Feser, F., Brisson, E., Kollet, S., Schmidli, J., van Lipzig, N. P. M., & Leung, R. (2015). A review on regional convection-permitting climate modeling: Demonstrations, prospects, and challenges. *Reviews of Geophysics*, *53*(2), 323–361. https://doi.org/10.1002/2014RG000475

Price, I., & Rasp, S. (2022). Increasing the accuracy and resolution of precipitation forecasts using deep generative models. *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, 10555–10571. https://proceedings.mlr.press/v151/price22a.html

Rampal, N. (2024). Enhancing Regional Climate Downscaling Through Advances in Machine Learning in: Artificial Intelligence for the Earth Systems. *Artificial Intelligence for the Earth Systems*. https://doi.org/10.1175/AIES-D-23-0066.1

Rampal, N., Gibson, P. B., Sood, A., Stuart, S., Fauchereau, N. C., Brandolino, C., Noll, B., & Meyers, T. (2022). High-resolution downscaling with interpretable deep learning:

Rainfall extremes over New Zealand. *Weather and Climate Extremes*, *38*, 100525. https://doi.org/10.1016/j.wace.2022.100525

Rampal, N., Hobeichi, S., Gibson, P. B., Baño-Medina, J., Abramowitz, G., Beucler, T., González-Abad, J., Chapman, W., Harder, P., & Gutiérrez, J. M. (2024). Enhancing Regional Climate Downscaling through Advances in Machine Learning. *Artificial Intelligence for the Earth Systems*, *3*(2). https://doi.org/10.1175/AIES-D-23-0066.1

Rasp, S., Dueben, P. D., Scher, S., Weyn, J. A., Mouatadid, S., & Thuerey, N. (2020). WeatherBench: A Benchmark Data Set for Data-Driven Weather Forecasting. *Journal of Advances in Modeling Earth Systems*, *12*(11). https://doi.org/10.1029/2020MS002203

Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., Fitzsimons, M., Athanassiadou, M., Kashem, S., Madge, S., Prudden, R., Mandhane, A., Clark, A., Brock, A., Simonyan, K., Hadsell, R., Robinson, N., Clancy, E., Arribas, A., & Mohamed, S. (2021a). Skilful precipitation nowcasting using deep generative models of radar. *Nature*, *597*(7878), 672–677. https://doi.org/10.1038/s41586-021-03854-z

Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., Fitzsimons, M., Athanassiadou, M., Kashem, S., Madge, S., Prudden, R., Mandhane, A., Clark, A., Brock, A., Simonyan, K., Hadsell, R., Robinson, N., Clancy, E., Arribas, A., & Mohamed, S. (2021b). Skilful precipitation nowcasting using deep generative models of radar. *Nature*, *597*(7878), Article 7878. https://doi.org/10.1038/s41586-021-03854-z

Reddy, P. J., Matear, R., Taylor, J., Thatcher, M., & Grose, M. (2023). A precipitation downscaling method using a super-resolution deconvolution neural network with step orography. *Environmental Data Science*, *2*, e17. https://doi.org/10.1017/eds.2023.18

Renwick, J. A., Mullan, A. B., & Porteous, A. (2009). Statistical Downscaling of New Zealand Climate. *Weather and Climate*, *29*, 24–44. https://doi.org/10.2307/26169704

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. In N. Navab, J. Hornegger, W. M. Wells, & A. F. Frangi (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (pp. 234–241). Springer International Publishing. https://doi.org/10.1007/978-3-319-24574-4_28

Saha, A., & Ravela, S. (2022). *Downscaling Extreme Rainfall Using Physical-Statistical Generative Adversarial Learning* (arXiv:2212.01446). arXiv. http://arxiv.org/abs/2212.01446

Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016). *Improved Techniques for Training GANs* (arXiv:1606.03498). arXiv. https://doi.org/10.48550/arXiv.1606.03498

Sha, J., Chen, X., Chang, Y., Zhang, M., & Li, X. (2024). A spatial weather generator based on conditional deep convolution generative adversarial nets (cDCGAN). *Climate Dynamics*, *62*(2), 1275–1290. https://doi.org/10.1007/s00382-023-06971-9

Sørland, S. L., Schär, C., Lüthi, D., & Kjellström, E. (2018). Bias patterns and climate change signals in GCM-RCM model chains. *Environmental Research Letters*, *13*(7), 074017. https://doi.org/10.1088/1748-9326/aacc77

Srivastava, A., Valkov, L., Russell, C., Gutmann, M. U., & Sutton, C. (2017). VEEGAN: Reducing Mode Collapse in GANs using Implicit Variational Learning. *Advances in Neural Information Processing Systems*, *30*. https://proceedings.neurips.cc/paper/2017/hash/44a2e0804995faf8d2e3b084a1e2db1d-Abstract.html

Steinschneider, S., Ray, P., Rahat, S. H., & Kucharski, J. (2019). A Weather-Regime-Based Stochastic Weather Generator for Climate Vulnerability Assessments of Water Systems in the Western United States. *Water Resources Research*, *55*(8), 6923–6945. https://doi.org/10.1029/2018WR024446

Sun, Y., Deng, K., Ren, K., Liu, J., Deng, C., & Jin, Y. (2024). Deep learning in statistical downscaling for deriving high spatial resolution gridded meteorological data: A systematic review. *ISPRS Journal of Photogrammetry and Remote Sensing*, *208*, 14–38. https://doi.org/10.1016/j.isprsjprs.2023.12.011

Thatcher, M., & McGregor, J. L. (2009). Using a Scale-Selective Filter for Dynamical Downscaling with the Conformal Cubic Atmospheric Model. *Monthly Weather Review*, *137*(6), 1742–1752. https://doi.org/10.1175/2008MWR2599.1

van der Meer, M., de Roda Husman, S., & Lhermitte, S. (2023). Deep Learning Regional Climate Model Emulators: A Comparison of Two Downscaling Training Frameworks. *Journal of Advances in Modeling Earth Systems*, *15*(6), e2022MS003593. https://doi.org/10.1029/2022MS003593

Verdin, A., Rajagopalan, B., Kleiber, W., Podestá, G., & Bert, F. (2018). A conditional stochastic weather generator for seasonal to multi-decadal simulations. *Journal of Hydrology*, *556*, 835–846. https://doi.org/10.1016/j.jhydrol.2015.12.036

Vosper, E., Watson, P., Harris, L., McRae, A., Santos-Rodriguez, R., Aitchison, L., & Mitchell, D. (2023). Deep Learning for Downscaling Tropical Cyclone Rainfall to Hazard-Relevant Spatial Scales. *Journal of Geophysical Research: Atmospheres*, *128*(10), e2022JD038163. https://doi.org/10.1029/2022JD038163

Wang, J., Liu, Z., Foster, I., Chang, W., Kettimuthu, R., & Kotamarthi, V. R. (2021). Fast and accurate learned multiresolution dynamical downscaling for precipitation.

Geoscientific Model Development, *14*(10), 6355–6372. https://doi.org/10.5194/gmd-14-6355-2021

Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Loy, C. C., Qiao, Y., & Tang, X. (2018). *ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks* (arXiv:1809.00219). arXiv. http://arxiv.org/abs/1809.00219

WeatherBench: A Benchmark Data Set for Data-Driven Weather Forecasting—Rasp—2020—Journal of Advances in Modeling Earth Systems—Wiley Online Library. (n.d.). Retrieved October 22, 2021, from https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2020MS002203

Xu, Z., Han, Y., & Yang, Z. (2019). Dynamical downscaling of regional climate: A review of methods and limitations. *Science China Earth Sciences*, *62*(2), 365–375. https://doi.org/10.1007/s11430-018-9261-5

Zhang, X., Alexander, L., Hegerl, G. C., Jones, P., Tank, A. K., Peterson, T. C., Trewin, B., & Zwiers, F. W. (2011). Indices for monitoring changes in extremes based on daily temperature and precipitation data. *WIREs Climate Change*, *2*(6), 851–870. https://doi.org/10.1002/wcc.147