https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

EGGS-WG: An open source global gridded stochastic weather generator derived from ERA5-Land

Alex Schuddeboom[1], Christian Zammit[2], David Plew[2], Piet Verburg[3], and Aidin Jabbari[2]

[1]School of Physical and Chemical Sciences, University of Canterbury, Christchurch, New Zealand
[2]National Institute of Water and Atmospheric Research, Wellington, New Zealand
[3]School of Geography, Environment and Earth Sciences, Victoria University of Wellington

Correspondence: Alex Schuddeboom (Alex.schuddeboom@canterbury.ac.nz)

Abstract. A new open source and freely available stochastic weather generator, the **E**RA5-land **G**lobal **G**ridded **S**tochastic **W**eather **G**enerator (EGGS-WG), is publicly available. This model offers several advancements over existing alternative stochastic weather generators, including the ability to simulate any terrestrial region of the planet, moving from simulating a single site to an entire gridded domain and increasing the temporal resolution of temperature simulation from daily to hourly. EGGS-WG is designed to be as low resource and portable as possible, ensuring ease of use for any potential end user. This paper describes in detail all of the principles that are used in the model design and code implementation. Preliminary tests are performed over five climatically distinct regions to confirm the ability of the model to reproduce the behaviour of the ERA5-Land dataset. Precipitation occurrence rates and daily rainfall amounts are shown to be reproduced accurately by the model with a range of attributes investigated, including seasonality, spatial correlations and rainfall spells. Although, there are some clear issues in the simulation of precipitation at the most extreme values. Similar analysis is also performed for the air temperature and dew point temperature, showing stronger agreement than the precipitation data. EGGS-WG is designed with a focus on potential future developments, including support for CMIP6 driven future warming and expansion into simulating other variables provided by the ERA5-Land dataset.

1 Introduction

Stochastic Weather Generators (SWGs) are a commonly used type of model, particularly for the simulation of realistic long term sequences of weather and to generate large ensemble probabilistic forecasts (Wilks and Wilby, 1999; Ailliot et al., 2015; Peleg et al., 2017; Sohrabi and Brissette, 2021). The fundamental principle behind SWGs is that by representing weather with a series of statistical relations, we can develop relatively simple and low computational cost models which are able to accurately simulate weather. As these models neglect the physics and chemistry of the earth system, they are fundamentally limited in their capabilities; however, due to their advantages they have been found to have a wide range of applications. This includes, but is not limited to, simulating precipitation for driving complex ecological models (Friend et al., 1997), as a source of multiple variables for crop models (Richardson, 1985), driving hydrological models including stream flow forecasts (Sohrabi and Brissette, 2021) and acting as an intermediary model to downscale results from global climate models (Wilks, 2012).

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

The majority, but not all, of SWGs follow the design principles that guided the initial development of the Weather Generator (WGEN) model (Richardson, 1981; Richardson and Wright, 1984; Wilks and Wilby, 1999). SWGs of this kind, often known as Richardson style models, are used to simulate precipitation, temperature, winds and radiation. This simulation process can be summarised as follows. First, a Markov chain is used to determine if an individual day is going to be wet or dry. Second, if the day is determined to be wet, then the precipitation amount is simulated by using a random number generator in conjunction with an assumed precipitation distribution (although sometimes an empirical distribution is used). Finally, the other required variables are calculated using a first-order auto-regressive model with lag-1 auto-correlation and cross-correlation between all variables (Parlange and Katz, 2000). The values used in this last step are conditioned by the precipitation state to ensure that any relationships between precipitation and other variables are appropriately represented in the model. Prominent models based on this general approach include WGEN (Richardson and Wright, 1984), CLIMGEN (Stockle et al., 2001), CLIGEN (Nicks et al., 1995), WeaGETS (Chen et al., 2012a) and LARSWG (Semenov et al., 1998). There are also more modern models which have a similar design basis but include more sophisticated statistical approaches, such as weatherGEN (Steinschneider and Brown, 2013), RMAWGEN (Cordano and Eccel, 2016) and GWGEN (Sommer and Kaplan, 2017). An interesting chain of models were also developed over California combining weather states, resampling and Gaussian noise (Steinschneider et al., 2019; Rahat et al., 2022; Najibi et al., 2024b, a). There are other kinds of weather generators that fill a similar role but use completely different simulation processes, such as physical modeling (Peleg et al., 2017), alternative statistical approaches (Dawkins et al., 2022; Whitehead and Bebbington, 2024) or machine learning (Ji et al., 2024; Rampal et al., 2025), however, these alternative approaches remain either more computationally expensive, complex or otherwise limited.

One of the major limitations of many of these SWGs is that they are restricted to the simulation of a single location at a given time. This is because these initial designs were based around extending limited time series corresponding to weather stations. There has been a long research history working on the development of Multi-site generators (Wilks, 1998; Khalili et al., 2009; Wilks, 2009; Baigorria and Jones, 2010; Chen et al., 2014; Abbasnezhadi et al., 2019). Changing from a single-site to a multi-site generator is a complex problem, as simulations additionally need to account for spatial correlations. For instance, if an anomalously hot temperature is simulated at one location, it would be expected that a second location within a few kilometers would also be hot. The most popular approach for handling this problem is the usage of spatially correlated random numbers. In particular, the technique described in Wilks (1998) and refined in Brissette et al. (2007) has proven extremely effective at generating the matrices needed to enforce appropriate correlations between any random numbers. There are alternative approaches such as those used in Khalili et al. (2009); Baigorria and Jones (2010), however, they are not considered for this project as benefits over traditional methods are relatively minor.

The motivation for this project was to further these developments with the release of an open source and easy to use gridded SWG capable of simulating any region of the planet, the **E**RA5-land **G**lobal **G**ridded **S**tochatic **W**eather **G**enerator (EGGS-WG). The aim is to produce the two most important fields for end-users, precipitation and temperature. A particular focus has been ensuring that the resource requirements are minimized, with the aim of reasonably sized simulations taking less than a day on personal computers. The approaches used in Abbasnezhadi et al. (2019) were shown to be highly effective at generating a gridded SWG over Canada using high resolution data from the Canadian Precipitation Analysis system. By refining this

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

approach and using the ERA-5 Land dataset (Muñoz-Sabater et al., 2021; Copernicus Climate Change Service, 2022) as the basis for our model parameters, we have been able to create a versatile gridded model that can automatically calculate the parameters needed to run the simulation. Moving from the regionally constrained approach in Abbasnezhadi et al. (2019) to a global approach substantially increases the usability of the model. The addition of dew point temperature and relative humidity simulations will also be useful for many end users.

It is important to note that while many of the developments of this model were inspired by Abbasnezhadi et al. (2019), the other freely available global models are generally limited. Compared to currently available global models, moving to a gridded approach from a site based approach is a notable improvement. This could be particularly useful for hydrological analysis in which an entire catchment will need to be considered. Additionally, ERA5-Land provides global land coverage, meaning that parameters should always be solvable without complex estimation strategies such as those used to increase the coverage of CLIGEN (Fullhart et al., 2021, 2023) and will not be constrained to a single region (Abbasnezhadi et al., 2019; Steinschneider et al., 2019; Najibi et al., 2024b). There are some limitations for ERA-5 data, particularly in the representation of extremes, however there are no viable alternatives to reanalysis data in terms of coverage and resolution. This does mean if our model fails to accurately handle extremes, these issues could compound with biases in ERA5 leading to significant differences between simulation and reality. Additionally, the hourly temporal resolution and high spatial resolution provided by ERA5-Land dataset can produce a level of detail missing from other publicly available SWG with lower resolutions. Despite the clear benefits to doing so, there are very few SWGs that simulate temperature hourly, as is done in Abbasnezhadi et al. (2019), with many of the older models simulation daily minimum and maximum temperatures instead due to input data limitations. There are other approaches like that used in the Steinschneider et al. (2019); Najibi et al. (2024b, a) that might be more powerful but their dependence on local weather states makes it infeasible to apply these approaches globally.

The remainder of this paper describes the underlying mathematical details of EGGS-WG, discusses some of the particulars of the implementation, and performs some simple validation tests on output generated. The full code base for EGGS-WG are available on a GPL-3.0 license. This is accessible through a public GitHub, with a version of record tied to this paper released on Zenodo, (https://doi.org/10.5281/zenodo.15086533, Schuddeboom (2025)). Also included on this repository are instructions for operating the SWG and downloading the data needed to run the model. The only files required to run the model, in addition to the code base, are ERA5-Land data. Initially the EGGS-WG is released with the ability to simulate precipitation, temperature and relative humidity, however other variables, including winds and radiation, are planned to be introduced with further development.

2 Model Design

As stated in section 1, the design of EGGS-WG is heavily based on the work of Abbasnezhadi et al. (2019). This general model flow is a slight modification of the Richardson style stochastic weather generators. It has the following three-step structure, with each step explored in detail in a following subsection:

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

1. A first order two state Markov chain is used to determine the precipitation state for each of the grid cells in the simulation domain. To ensure grid cohesion, the random numbers driving these Markov chains are all required to be spatially correlated.

2. For every cell that is determined to be precipitating, a rainfall amount is calculated using a precipitation equation determined specifically for each location. As with the occurrence, the random numbers used are required to be spatially correlated.

3. Hourly air temperature and dew point temperature are simulated using a modified residual approach. This is different from traditional Richardson generators, as the residuals are generated through the usage of conditioned Fourier series.

2.1 Precipitation Occurrence

The simulation of precipitation occurrence using a Markov chain is one of the longest standing and most consistent component of SWGs (Richardson, 1981; Wilks and Wilby, 1999). There has been significant experimentation with different kinds of Markov chains, including using second and higher order chains (Stockle et al., 2001; Chen et al., 2012a), using more than two precipitation states (Stoner and Economou, 2020; Dawkins et al., 2022) and using more sophisticated Markov models to account for spatial correlations (Baigorria and Jones, 2010). A simple two state first order Markov chain is used in this work, as they generally provide adequate performance (Wilks and Wilby, 1999) and more importantly should be more robust for simulating low precipitation regions. However, in certain climates a single order chain has been show to fail to produce sufficient dry spells (Wilks and Wilby, 1999; Chen et al., 2012b, a), as such an exploration of possible improved performance through a more complex Markov chain is intended for the future.

The exact simulation process for the vector of precipitation states, Q, is as follows. Every grid cell, i, is considered either to be wet or dry. A dry day is defined as a day in which less than 1 mm of precipitation is recorded, with all other days considered wet. The vector is filled so that $Q_{i,t} = 1$ indicates that a grid cell, i, at time t is wet, and a value of 0 is used if i is dry at time t. The simulation is driven using two different probabilities, the probability that a dry day transitions to a wet day, $P_{[Q_{i,t-1}=0,Q_{i,t}=1]}$ and the probability of a wet day remaining wet for the next day, $P_{[Q_{i,t-1}=1,Q_{i,t}=1]}$. For each day a uniformly distributed random number, labelled w_i^t, is generated and if it is greater than the probability of a wet day corresponding to the current state, $P_{c,i}$, the day is considered dry. If w_i^t is less than, $P_{c,i}$ the day is considered wet. In its entirety, this process works by initialising Q to a random precipitation state (more details in section 3) and then simulating the next state day by day based on the transition probability of the grid cell. To ensure that the resulting field is spatially coherent, the random numbers, w_i^t, used with the transition probabilities must be spatially correlated, as described in further detail in section 3.4.

2.2 Precipitation Amount

Due to the statistical complexities of the precipitation distribution, many different models for precipitation amount simulation have been used in past research, none of which manage to flawlessly represent real world precipitation. Most prominent ap-

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

proaches amongst existing models assume precipitation follows an exponential (Richardson, 1981), mixed exponential (Wilks, 1998, 1999a) or gamma (Wilks, 1992) distribution. There are also approaches that use an empirical or semi-empirical distribution (Semenov et al., 1998); however, we avoided an empirical approach due to concerns it would limit future development of EGGS-WG, particularly if future climate changes are considered. For EGGS-WG, we choose to use a three parameter mixed exponential distribution:

$$f(\alpha,\beta_1,\beta_2) = \frac{\alpha}{\beta_1}e^{(\frac{-r}{\beta_1})} + \frac{1-\alpha}{\beta_2}e^{(\frac{-r}{\beta_2})} \tag{1}$$

Where r is the observed precipitation amount, α effectively represents the probability that the exponential associated with β_1 represents the precipitation distribution. Therefore, $1-\alpha$ represents the probability of the β_2 distribution. This also requires the enforcement of the bounds $0 < \alpha \leq 1$ and $\beta_1 \geq \beta_2 > 0$. The parameters, α, β_1 and β_2, are calculated independently for every grid cell in the simulation region. Once we have fit these parameters, an inverse CDF (cumulative distribution function) procedure can be used to simulate the precipitation amounts. Following Abbasnezhadi et al. (2019), this can be done with the equation:

$$r_i^t = DLT - \beta_i^t \times ln[1-\Phi(z_i^t)] \tag{2}$$

Where r_i^t is the rainfall that occurs at grid cell i and time t. DLT is the minimum viable precipitation required for a day to be considered wet (1 mm/day for EGGS-WG), and is required because imposing this cutoff on our simulation means that values below 1 mm should not exist in simulation output. z_i^t is the random number (generated from a standard normal distribution) used to drive the simulation. $\Phi(\cdot)$ is the function transforming a Gaussian random number to a uniformly distributed equivalent. β_i^t is a combination of β_1 and β_2:

$$\beta_i^t = \beta_{2,i} + 2(\beta_{1,i} - \beta_{2,i})[1-\Phi(w_l^t)/\alpha_i^t P_{c,i}] \tag{3}$$

This is following the approach introduced in Wilks (1998), which uses the occurrence model data as part of determining β_i^t. This includes both the transition probability, $P_{c,i}$, and driving random number, w_i^t from the occurrence simulation. A combined model approach like this is essential for avoiding unrealistic precipitation fields in the synchronized fields. Both w_i^t and z_i^t must be generated in a spatially correlated way to ensure that the output the model produces is spatially coherent output.

2.3 Air Temperature and Dew Point Temperature

By opting to simulate temperature data at an hourly temporal resolution, EGGS-WG follows a different approach than the majority of publicly available stochastic weather generators. Specifically, the approach used is to generate both a daily and an hourly Fourier series and condition the hourly series based on the relation between the daily series and precipitation state. These Fourier series are used to reduce the underlying data to normally distributed residuals, which in turn are used to determine

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

the parameters necessary for simulation. Fourier series are particularly useful for representing temperature due to the large periodic signals due to seasonal and diurnal variations. As far as the authors are aware this exact approach has only been used by Abbasnezhadi et al. (2019) to date; however, Abbasnezhadi et al. (2019) recognize that their approach is partially based on other existing work (Hansen and Driscoll, 1977; Wilks, 2008). First, the hourly Fourier series is calculated using the entirety of the input data. This is done for every grid point i using the equation:

$$\tilde{P}_H(i,t) = \bar{T}(i) + \sum_n \{\alpha_H(i,n) cos[\frac{2\pi nt}{8760}] + \beta_H(i,n) sin[\frac{2\pi nt}{8760}]\} \tag{4}$$

where $\tilde{P}_H(i,t)$ is the Fourier series fit for either the mean or standard deviation of the hourly temperature data at time t. The series is calculated for both mean and standard deviation independently, as both are required to calculate residuals. $\bar{T}(i)$ is the mean temperature or mean standard deviation over the entire record and α_H and β_H are Fourier coefficients corresponding to a particular harmonic, n. For EGGS-WG, we use the values $n \in [1, 365, 730, 1095, 1460, 2190]$ corresponding to yearly, daily, 12-hourly, 8-hourly, 6-hourly and 4-hourly cycles. This is a possible point of improvement for EGGS-WG, as different selections of n may lead to a more realistic temperature cycle, although initial work has only identified minor impacts of changing these values. α_H and β_H are calculated with the following formulations:

$$\alpha_H(n,i) = \frac{2}{8760} \sum_{t=1}^{8760} \bar{M}_H(t,i) cos(\frac{2\pi nt}{8760}) \tag{5}$$

$$\beta_H(n,i) = \frac{2}{8760} \sum_{t=1}^{8760} \bar{M}_H(t,i) sin(\frac{2\pi nt}{8760}) \tag{6}$$

where $\bar{M}_H(t,i)$ is either the mean or standard deviation over the entire observational record at time t and location i. With both the hourly mean, $\tilde{\mu}_H(t,i)$, and standard deviation, $\tilde{\sigma}_H(t,i)$, calculated using equation 4, a full set of residuals can be calculated by subtracting $\tilde{\mu}_H(t,i)$ and dividing by $\tilde{\sigma}_H(t,i)$ from the original temperature data. Once the data is processed into hourly residuals, they can be used to determine two matrices, $\mathbf{X}$ and $\mathbf{Y}$, needed to run an auto regressive model (Wilks and Wilby, 1999). In particular, this model takes the form:

$$\hat{R}_H(t) = \mathbf{X}\hat{R}_H(t-1) + \mathbf{Y}E(t) \tag{7}$$

where $R_H(t)$ is the series of the residuals that can be converted back to temperatures using the Fourier series. As is clear from equation 7, there are two components to the model: one part dependent on the last hour residuals, $\hat{R}_H(t-1)$, and one part dependent on $E(t)$, which is a vector of randomly generated normal variates with a mean of 0 and standard deviation of 1. Calculation of $\mathbf{X}$ and $\mathbf{Y}$ is a well established process (Wilks, 1998, 1999b; Wilks and Wilby, 1999), using the unlagged, C_0 and lag-1, C_1, correlation matrices between residuals. This lets us calculate $\mathbf{X}$ and $\mathbf{Y}$, with:

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

$$\mathbf{X} = C_1 C_0^{-1} \tag{8}$$

$$\mathbf{G} = C_0 - \mathbf{X} C_1^T \tag{9}$$

where **G** is the Gramian of **Y**, so **Y** can be calculated with the Cholesky decomposition of **G**. With **X** and **Y** calculated, we have everything needed to run our hourly temperature model in equation 7. However, this process would leave the generated temperature data independent of the precipitation state. To condition these fields, we generate two different daily Fourier series, one using data from only wet days and one with only dry days. Again, we follow the approach detailed in, Abbasnezhadi et al. (2019) which uses the following modified versions of equations 4, 5 and 6:

$$\tilde{P}_D^s(i,t) = \bar{T}^s(i) + \sum_n \{\alpha_D^s(i,n) cos[\frac{2\pi nt}{365}] + \beta_D^s(i,n) sin[\frac{2\pi nt}{365}]\} \tag{10}$$

$$\alpha_D^s(n,i) = \frac{2}{365} \sum_{t=1}^{365} \bar{M}_D^s(t,i) cos(\frac{2\pi nt}{365}) \tag{11}$$

$$\beta_D^s(n,i) = \frac{2}{365} \sum_{t=1}^{365} \bar{M}_D^s(t,i) sin(\frac{2\pi nt}{365}) \tag{12}$$

where S is the precipitation state and is used to indicate the variable is only using data specific to that precipitation state. An alternative set of n values are used for this calculation as well, $n \in [1,2,3,4,6]$, to account for the shift from hourly to daily. The difference between the unconditioned series from equation 4 and the conditioned series from equation 10 is calculated and used to adjust the hourly series, resulting in two conditioned hourly series which can be used to transform simulated residuals to conditioned temperature values.

3 Implementation

While section 2 describes the design principles behind EGGS-WG, there are several implementation level issues that need to be further explained to understand how the model functions. Several of these decisions were made primarily from practical concerns with a view towards limiting the resources required to run EGGS-WG, and building EGGS-WG in a way that does not restrict future development. The most substantial of these issues are discussed in the following subsections.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

3.1 ERA-5 Land Data

All the parameters required to run the model are derived from ERA 5 Land output (Muñoz-Sabater et al., 2021). This is a $0.1° \times 0.1°$ spatial resolution dataset with hourly values and a global coverage over land surface. Specifically, the variables used for this model include the 2 m air temperature, 2 m dew point temperature, precipitation and geopotential (used to determine altitude which is output but not used for simulation). The ERA5 download process allows for the data to be subset before being downloaded to disk. As such, disk space and memory costs can be reduced by using a subset covering the simulation domain. Code to download this data is included with the model. Expansion to other ERA5-Land variables including winds and radiation are planned but not included in this version of EGGS-WG. For model validation purposes, the ERA5-Land data is referred to as observational data in this paper, although as ERA5 is a reanalysis product this is not strictly true. In general, ERA5 has been shown to represent precipitation well in comparison to gauge data with a slight wet bias, however the most extreme precipitation values appear to be underestimated (Lavers et al., 2022) which will impact our simulation of the these extreme values.

3.2 Stationarity

Most stochastic weather generators are inherently stationary (Wilks and Wilby, 1999), meaning that the climate simulated is fixed to one point in time and effects such as past climatic changes due to greenhouse gases are unaccounted for. There is, however, an established approach adding climate change support to SWGs and using them for down scaling larger resolution climate models (Wilks, 1992; Kilsby et al., 2007; Wilks, 2012; Keller et al., 2017; Vesely et al., 2019). While we are interested in introducing this capability to EGGS-WG at a later date, currently EGGS-WG is designed to be stationary. To account for the historical changes that have occurred over the observational period, we adjust all the ERA5-Land data so that it is equivalent to data from the year 2022. This means that our SWG is only effectively simulating conditions equivalent to 2022. To adjust our data, we fit a Fourier series to the underlying data. The linear terms of this Fourier series are then used to generate a trend line for the entire observational record. This trend lines can then be inverted and added to the observational data, adjusting the data so that it matches 2022 climate. For example if June showed 1 degree warming between 1950 and 2022, then 1 degree would be added to the 1950 values. Adjustments made using this approach are limited and will not account for some potentially significant changes in variability or precipitation occurrence.

Additionally, the model has been released with a non-stationary mode that is not examined in this paper. Effectively, this non-stationary mode adds the trends observed in the Fourier series back into the model output after the simulation process. This results in the model recreating the observed trends in ERA5-land without any fundamental changes in the simulation calculations. Periods from before the ERA5-land record are all simulated at 1950 conditions, as there is no consistent approach for estimating earlier changes with the ERA5-Land data. Periods from beyond ERA5-land coverage are simulated with the most recently available data, assuming no change into the future. As stated in the above paragraph, climate change forcings will be considered in the future, but this will go beyond the data provided by ERA5-land and utilize the CMIP model output.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

An example figure showing the effects of these adjustments on temperature is included in the supporting information as figure S1.

3.3 Precipitation Parameter Calculation

There are several distinct approaches that could be used to solve for α, β_1 and β_2 in equation 1. In EGGS-WG these variables are estimated by using the SciPy optimization package (Virtanen et al., 2020) to minimize the difference between the theoretical distribution in equation 1 and the ERA5-Land data. These values need to be calculable for every grid point in the ERA5-Land dataset, as any grid point could be a potential part of a simulation. This strongly limits the ability to make any simplifying assumptions about the fit of the precipitation function. Experimentation with different algorithms has shown the most consistent results from the Powell (Powell, 1964) and Nelder-Mead (Nelder and Mead, 1965) algorithms. The Nelder-Mead algorithm is tried first, and when that fails to reach an appropriate solution, the Powell algorithm used. Restricting α to a value between 0.5 and 0.9 has also shown to improve the results in our internal testing.

3.4 Spatial Coherence of Random Numbers

The random numbers used for the generation of precipitation occurrence and amounts simulations are calculated with the exact same process. In particular, the approach used by Brissette et al. (2007) (in turn based on Wilks (1998); Rebonato and Jaeckel (2000)) is the one that is implemented in EGGS-WG. One of the difficulties of this process is that it requires swapping between the underlying uniformly distributed values to normally distributed values to make the simulation process efficient (Brissette et al., 2007; Abbasnezhadi et al., 2019).

First, we focus on the precipitation occurrence simulation. This process starts by the calculation of the Pearson correlation coefficients between all station pairs of observed precipitation sequences, $\rho_{i,j} = corr(Q_{i,:}, Q_{j,:})$, where $Q_{i,:}$ is the precipitation occurrence vector (as defined in section 2) at location i over the entire observational record. This lets us define the matrix C which is populated by $\rho_{i,j}$ for every pair of i and j. Provided that C is positive definite, applying a Cholesky factorization will result in the upper triangular matrix R such that $C = R'R$ (Brissette et al., 2007). Multiplying the lower triangular matrix R' by a matrix of random normally distributed numbers results in a matrix with the same covariance as C, therefore recreating the expected spatial correlations.

There are, however, some issues that make the generation of R' more complex (discussed in detail in Brissette et al. (2007)). Most importantly, the generated correlations will be smaller than those observed. To account for this, larger correlation values than observed will need to be used to generate R'. Calculating the optimal correlation values to use is done iteratively, with the equation:

$$C_{R_{i+1}} = C_{R_i} + \eta(C_{x,obs} - C_{x,syn}) \tag{13}$$

where $C_{x,obs}$ is the observed correlations defined as C above, $C_{x,syn}$ is the correlations resulting from when a model is run with the correlations C_{R_i}. $C_{R_{i+1}}$ is calculated, and the process is repeated until $C_{x,obs} - C_{x,syn}$ reaches a minimum. In

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

EGGS-WG, several different convergence factors η are used based on internal testing and $C_{x,obs}$ is used as C_{R_0}. This approach can result in $C_{R_{i+1}}$ matrices which are not positive definite. To ensure we have a positive definite matrix, we adjust this matrix by finding its eigenvalues, replacing any negative eigenvalue with a small positive value and then recreating the matrix by using the new eigenvalues and old eigenvectors (Rebonato and Jaeckel, 2000). The precipitation occurrence simulation can be handled the same way with one key difference, the correlations between points, ρ, can only be calculated when both points are considered to be wet.

3.5 Temperature Cross Correlation

One of the differences between EGGS-WG and the SWG in Abbasnezhadi et al. (2019) is that EGGS-WG simulates dew point temperature in addition to air temperature. These two variables are obviously highly correlated, so a mechanism for enforcing this correlation in the model is required. This can be done with the well established approach defined in Richardson and Wright (1984); Wilks (1999b). Effectively, this is a modification of the matrices **X** and **Y** in equation 7 to include both air temperature and dew point temperature, and changing $\hat{R}_H$ from a vector of air temperature residuals to a vector that includes residuals for both air and dew point temperatures. This requires calculating both the lag 0 and lag 1 cross correlations between air temperature and dew point temperature, with the redefinition of C_0 and C_1 from section 2.3 as:

$$C_0 = \begin{bmatrix} C_0(T,T) & C_0(T,D) \\ C_0(D,T) & C_0(D,D) \end{bmatrix} \qquad C_1 = \begin{bmatrix} C_1(T,T) & C_1(T,D) \\ C_1(D,T) & C_1(D,D) \end{bmatrix} \tag{14}$$

Where $C_0(T,D)$ indicates the lag 0 correlation matrix between air temperature, T, and dew point temperature, D. This approach can be scaled to as many variables as required, with past work commonly including radiation and winds. It is our intention to use this same approach when investigating the addition of new variables in later model releases.

3.6 Vector Initialisations

There are several different approaches for initializing the vectors for simulation. EGGS-WG uses a simplistic approach where a slightly simplified version of the model is run for a spin up period of 1000 days. The first day of this spin up period is considered entirely dry, with the temperatures starting with a residual value of 0. Internal testing showed that this spin up period was sufficient to start the simulation in a random state.

3.7 Post simulation adjustments

A relatively unique innovation of this model is the ability to provide post simulation adjustments to the simulation output. There are many reasons why a user would not want these adjustments on simulation output, so this is designed to be an easily togglable model feature. Specifically, two series of adjustments are available, a distribution wide precipitation adjustment and a diurnal variation adjustment to temperatures. For the precipitation adjustment, a cubic spline is fit to the differences between the observed and simulated precipitation distributions at specific percentiles. The generated spline is used to calculate

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

an adjustment factor for every simulated precipitation value. This is done using all the precipitation data, so that there is no additional adjustment to account for seasonality or spatial differences. This is also done at a course scale, around 10 distinct bins, to avoid overfitting to the underlying data. The temperature adjustment works by examining the difference between the month specific daily mean and the diurnal cycle relative to the mean. These differences are added to the simulated temperature data, adjusting it so that the diurnal cycle in EGGS-WG better matches simulations. The main reason these are included is that the global nature of the model may lead to cases where some model parameters are poorly fit and adjustment is necessary for them to be adjusted.

3.8 Relative Humidity

EGGS-WG does not directly produce the relative humidity through the simulation process. Instead, relative humidity is calculated by taking the ratio of the saturation vapour pressure, p_s, to observed vapour pressure, p. In particular, we use the Tetens equation for estimating vapour pressure with the air temperature T and dew point temperature D. Note that for this equation to be accurate, temperatures must be given in Celsius.

$$RH = \frac{p}{p_s} \qquad p = 0.61078e^{\frac{17.27D}{D+237.3}} \qquad p_s = 0.61078e^{\frac{17.27T}{T+237.3}} \tag{15}$$

4 Results and Discussion

4.1 Simulation Testing Regions

To evaluate the performance of the EGGS-WG, we examine simulations covering three distinct domains over New Zealand, one over Australia and one over New Guinea, as shown in figure 1. While almost all of New Zealand is the same Köppen type, Cfb, there is substantial variation in the frequency and magnitude of precipitation, particularly when it comes to extreme events (Vishwanathan et al., 2023b, a). The three New Zealand based regions are chosen to explore a range of precipitation conditions. The West Coast region of the South Island is a rainforest known for its relatively frequent and high intensity precipitation due to the orographic effects of the nearby Southern alps. On the other side of the Southern alps is the Canterbury region, dominated by flat and dry plains. Finally, we include a region on the North Island around New Zealand's largest freshwater lake, Lake Taupō. This is in-between the other two regions in precipitation frequency and annual rainfall. In addition, we include an Australian region which is amongst the driest environments on the planet and a New Guinea region amongst the wettest. The New Guinea region is also much larger than the other regions, allowing us to test model performance over a large domain. For the observational data, we use trend adjusted ERA5-Land data from the years 1950 through 2022, whereas for the model data simulations of 110 year duration are used.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

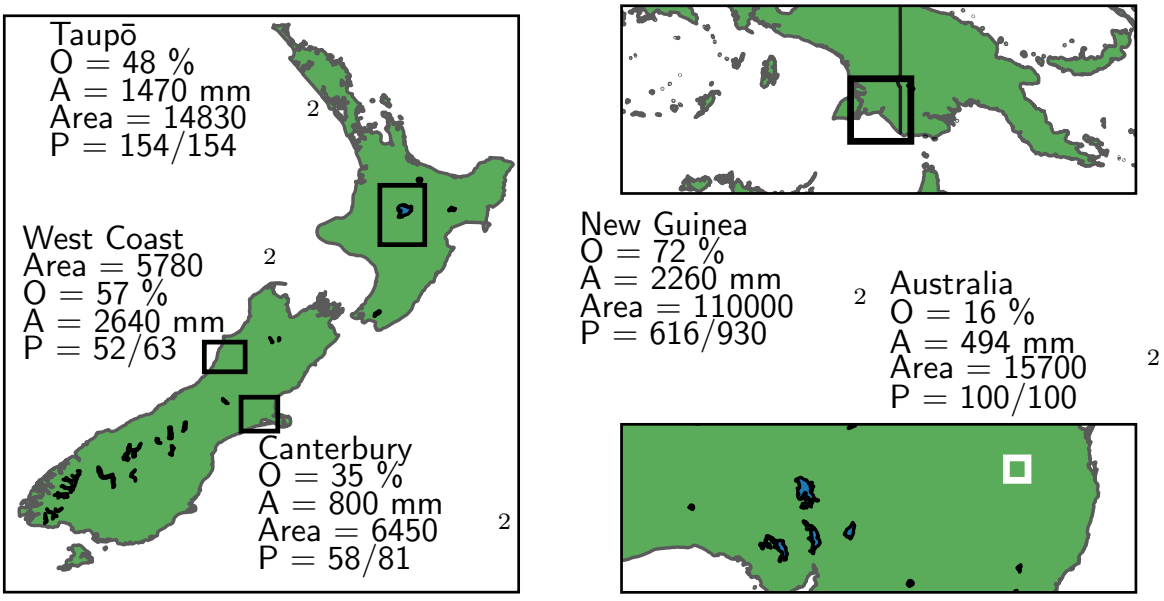

Figure 1. A map showing the five simulation regions in our domain: the West Coast, Canterbury, Taupō, New Guinea and Australia. Included is the ERA5-Land average precipitation occurrence rate over the region, *O*, the annual average precipitation over the region, *A*, the area covered by the region and the total number of grid points in the region, *P*. These values are calculated as averages between 1950 and 2022.

4.2 Precipitation Testing

The first check for model accuracy is to ensure that the overall precipitation occurrence rates and accumulation values match those generated from observations. Figure 2 shows the observational and simulation values for monthly average occurrence rate and precipitation amount for each of the five regions. Clear differences are seen between the regions with occurrence rates ranging from lows around 10 % over Australia and highs around 100 % in New Guinea and monthly amounts ranging from around 10 mm to 300 mm. Generally, the simulation appears to be able to capture to the mean regional properties with some minor differences. All the regions show some level of seasonality, with the monsoon seen in the New Guinea resulting in much larger seasonality than the other regions and the relatively dry Australia and Canterbury regions showing only minor seasonal variation. In general the precipitation occurrence and amount are well captured over all five regions, with the occurrence simulation slightly better matching the observations than the amount simulation.

While figure 2 shows that the domain wide simulation is accurate, it does not give any insight into the spatial structure within these regions. Following an approach established in past research (Wilks, 1999b; Brissette et al., 2007; Abbasnezhadi et al., 2019), the ability of the simulation to capture spatial dynamics was tested using Pearson correlation coefficients. This testing is done by calculating the correlation coefficient between every possible pairing of grid points in the observations and then compared to simulations. All of these calculations are done for each month of the year independently, ensuring that seasonal

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

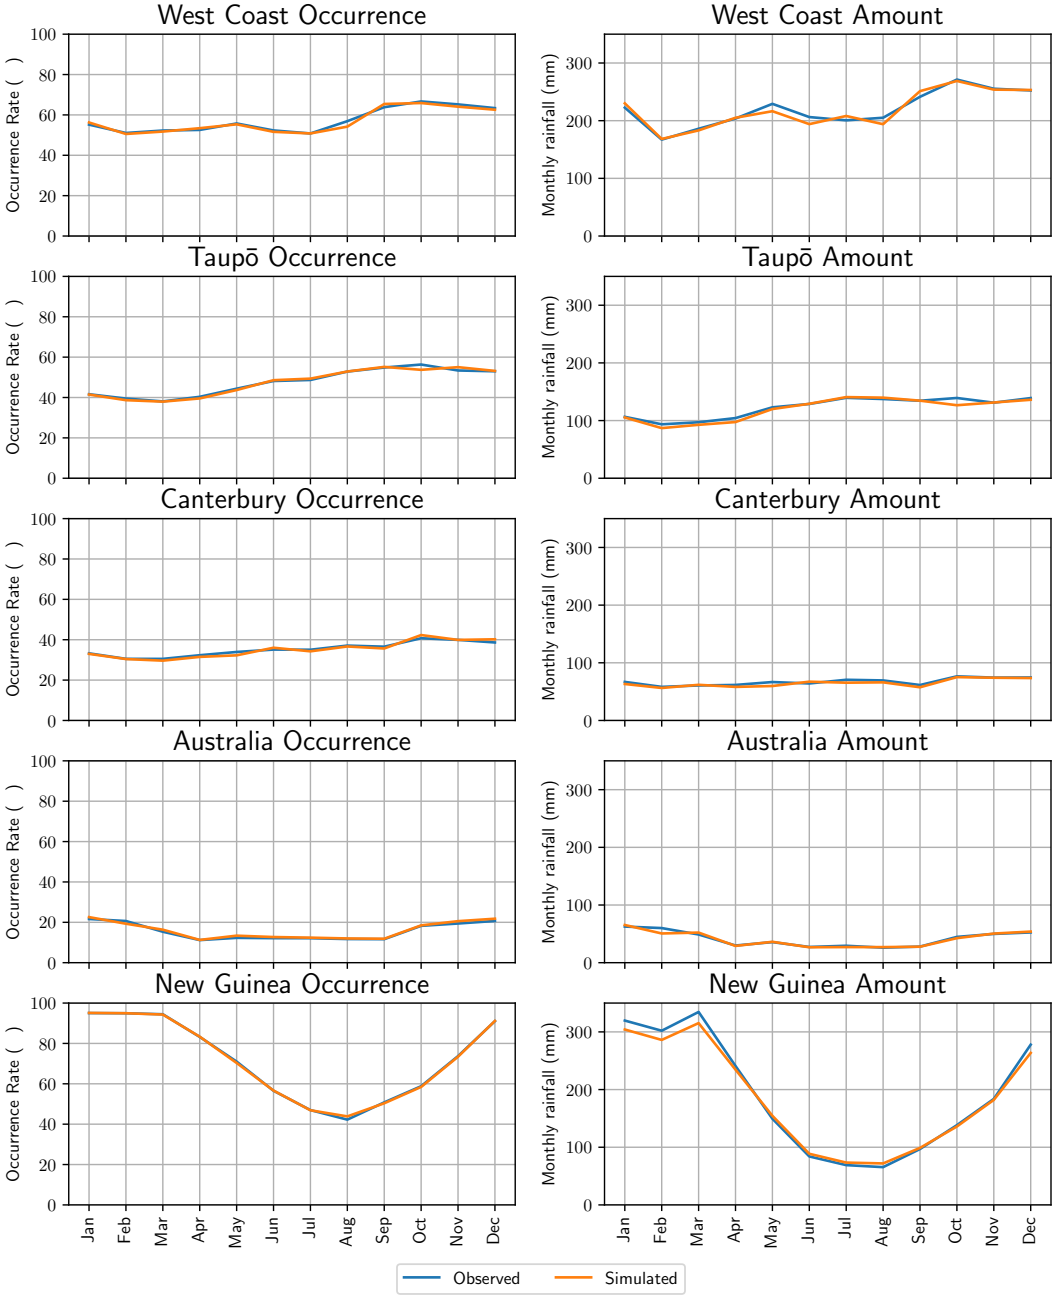

Figure 2. Seasonal average values for the precipitation occurrence and amounts over the five simulation domains. Precipitation amounts are calculated as a monthly total values and then averaged over all of the years in the observations and simulations.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

variation isn't inflating the correlation values. The observed correlation coefficient can then be directly plotted against the simulations to determine if the same spatial patterns of correlation are present.

The resultant correlation coefficient plot is included in figure 3. Through the simulated correlation coefficients, the simulation is shown to be very successful at replicating spatial precipitation occurrence dynamics over all five of our test regions. Interestingly, the observed correlations are substantially different over the different New Zealand domains with strong correlations with tight clustering over the West Coast region, much weaker correlations present over the Taupō region and an intermediary state in the Canterbury region. The simulation is able to capture these dynamics well with a very high level of accuracy in the West Coast and Australia regions, a few outlier grid point pairs that underestimate the observed correlation rate over Canterbury and a wider range but generally accurate values over Taupō. The New Guinea region shows a much different set of correlations than are seen in the other regions, with the main values offset from the 1 to 1 line and points ranging from a correlation of 0 to 1. The spatial dynamics are the very similar for the amount values, with the biggest differences seen in the Canterbury region where a relatively small, but clearly visible subset of grid point pairs show much lower simulated correlations than observed. It is worth noting that this a relatively small portion of the simulation pairs over the Canterbury region, and that the overwhelming majority of grid pairs over this region are well captured. The Australia, West Coast and Taupō regions show stronger agreement than the data over Canterbury. The New Guinea amount data shows a better correlation than it does for occurrence, with the majority of the points landing near the 1 to 1 line despite still covering a large range of correlations.

Another common validation approach in past research has been ensuring that precipitation occurrence spells are well represented (Wilks and Wilby, 1999), where precipitation spells are defined as consecutive days that are either wet or dry days. These results are included in figure 4 with an alternative version included in the supplementary material which uses a logarithmic y-axis to emphasize differences in extremes values. Over all of the regions, the agreement between model and observations is strong, although there is minor disagreement, particularly over the New Guinea region. The biggest differences are seen in the short-lived spells of less than 5 days, although typically the biases cancel each other out over the first 5 days. Interestingly, while there are clear distinct behaviours associated with the different simulation regions, such as the relatively short maximum dry day spells on the West Coast or wet day spells in Canterbury, the model is capable of mimicking these distinct regional behaviours. There appears to be some issues that with the New Guinea simulation for long wet streaks, which are more visible in the supplementary information. This warrants further investigation, but given the relative rarity of these events should not have a significant impact on the simulation. These results lack some of the biases that are expected of using a first order Markov chain, in particular the inability of these models to accurately capture long dry spells (Wilks and Wilby, 1999). This may be due to the climates of the regions analyzed but does suggest that, at least for these simulation domains, the benefits of changing to a different chain depth or design is unlikely to significantly improve EGGS-WG.

For a closer examination of the precipitation amount distribution, a quantile-quantile plot of the precipitation produced by the simulation and seen in the observations is shown as figure 5. All the dry day data is removed from this comparison, as this would effectively be testing the occurrence rate of the simulation, not the precipitation amount distribution. Generally, there is very good agreement between the simulated and observed percentiles over all five of the domains with very little difference visible below the 99^{th} percentile. However, the highest percentile values show notable deviation from the rest of the data,

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

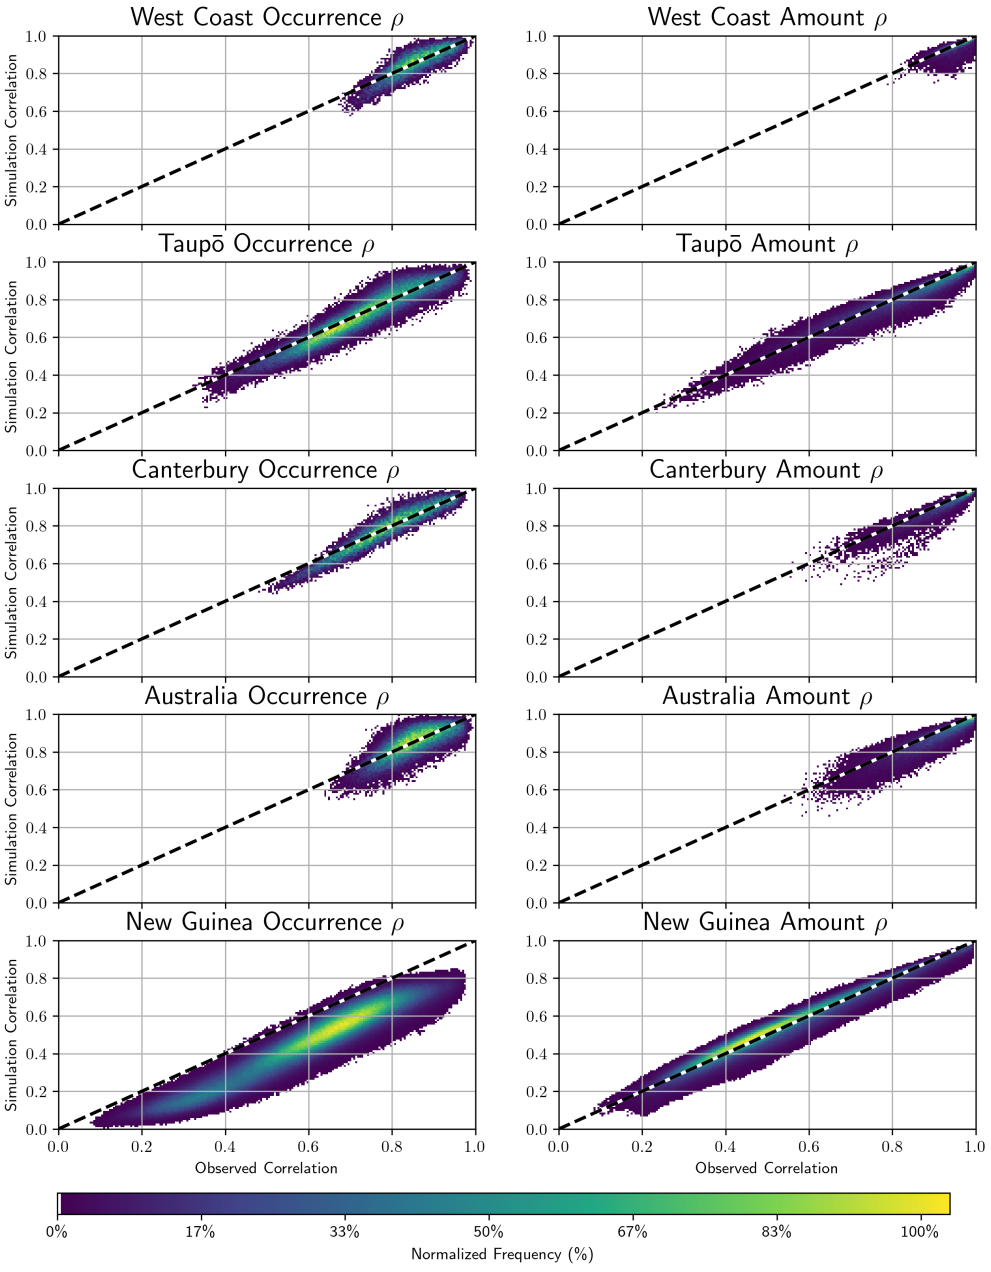

Figure 3. The Pearson correlation coefficient between every set of station pairs for both the simulation and observational data. Values are calculated for both the precipitation occurrence and amount data. The colour of the figure indicates the density of station pairs at a given correlation, ranging from purple for a low number of points to green for a high number of points. The included dashed line indicates the 1:1 correlation line that would be the theoretical ideal for a simulation.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

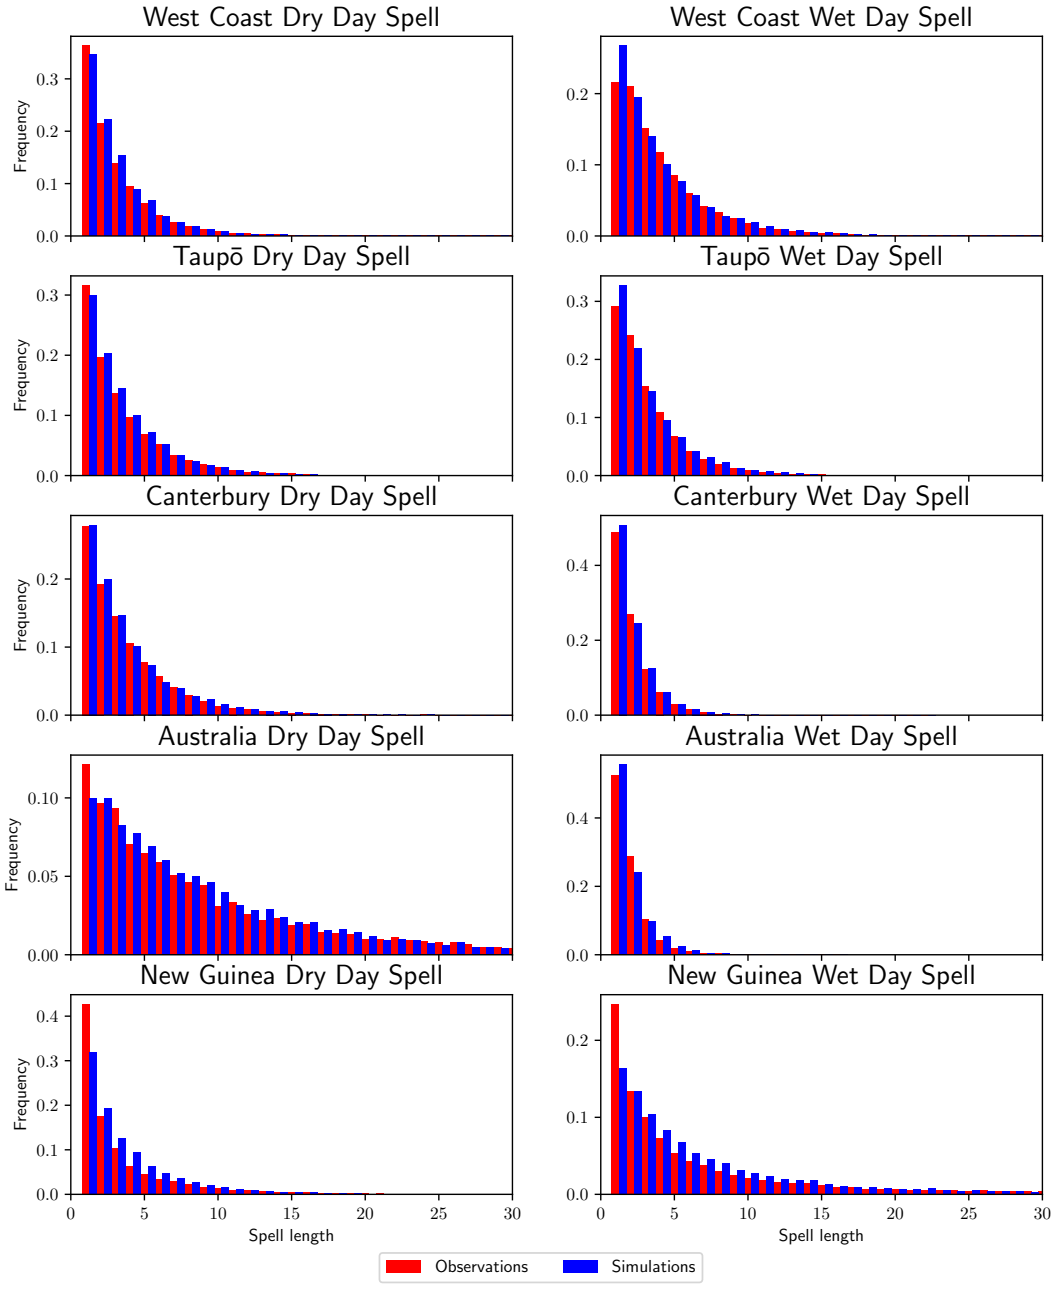

Figure 4. The wet and dry day precipitation spells for the both the simulations and observations over all five of the analyzed domains.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

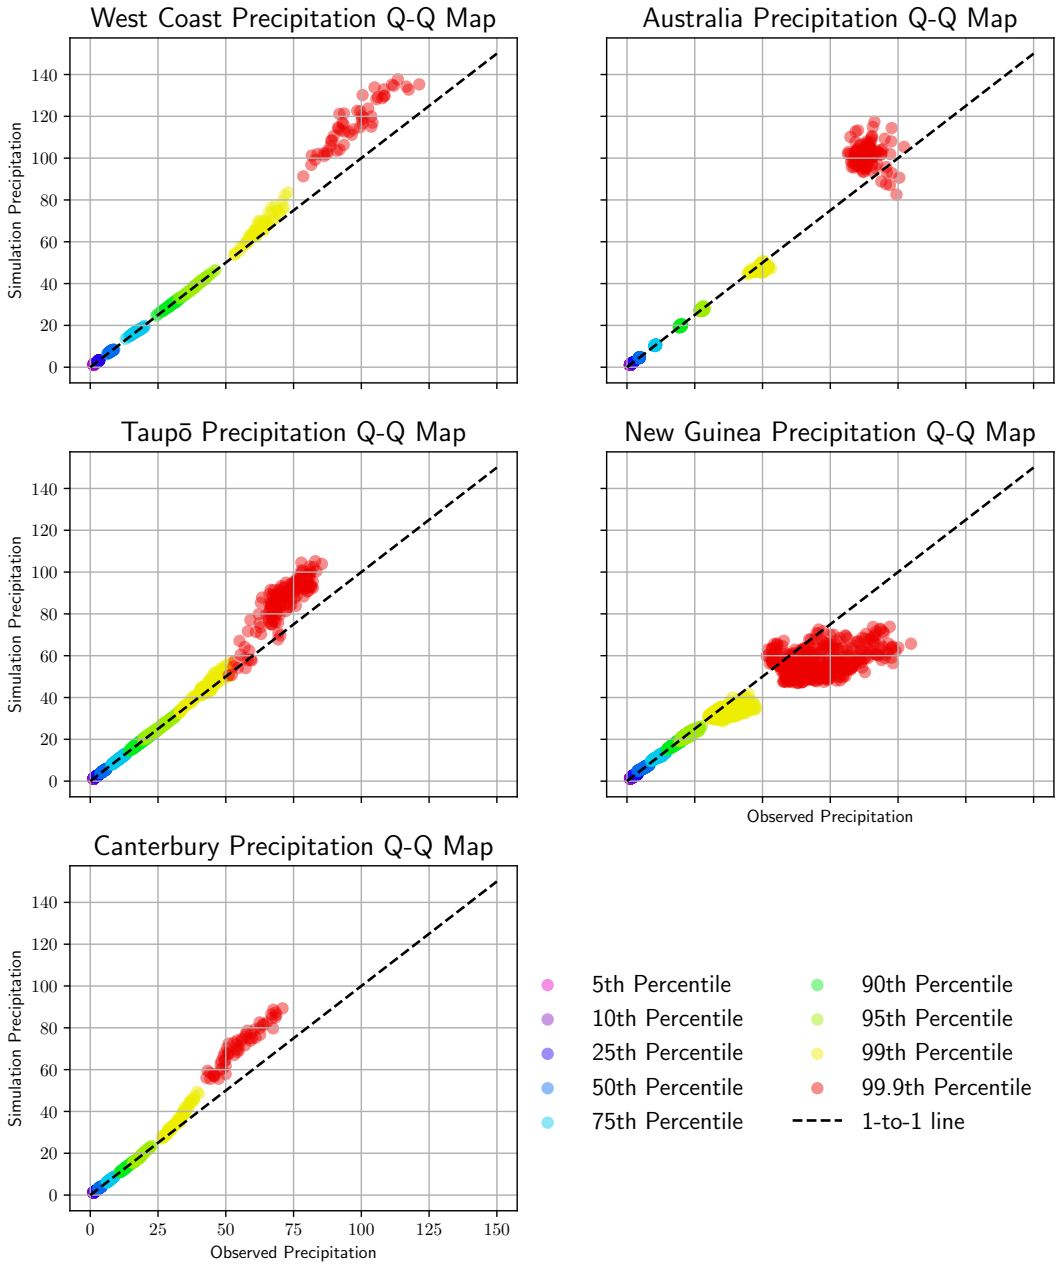

Figure 5. Quantile-Quantile plot of the distributions of observed and simulated precipitation data for every individual grid point in each of the domains. The precipitation data is filtered so that only the wet days are considered. The percentiles included in this plot are the 5^{th}, 10^{th}, 25^{th}, 50^{th}, 75^{th}, 90^{th}, 95^{th}, 99^{th} and 99.9^{th} percentiles. The dashed line indicates the 1:1 correlation line.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

with simulation values consistently larger than observed values. This is deviation is much more pronounced in the Canterbury region than in the Taupō and West Coast regions and much weaker in the Australia region. The only exception to this is the New Guinea region where the extreme precipitation is underestimated by the simulation. This discrepancy could have major implications if the model is used to directly study extreme events, but for general simulations the impact on the simulations is relatively minor due to the relative infrequency of these events. A more aggressive precipitation adjustment than the one described in section 3.7 could be used to improve the simulation of these extreme values. However, to retain the simplicity and to avoid overfitting the underlying data, we choose to keep the post simulation adjustments to precipitation relatively small.

To more comprehensively explore the biases in precipitation distribution, an additional set of figures is included in the supporting information in figures S3-8. These include the general distribution, time aggregated overdispersion plots to align with figures 9 and 10 from Abbasnezhadi et al. (2019) and time period aggregated and full region wide distributions. The results shown in the full precipitation distribution show very strong agreement with figure 5, consistent quality simulation with some minor issues appearing in the representation of the most extreme values. The monthly and yearly standard deviations show results that are comperable to Abbasnezhadi et al. (2019) for the majority of the regions. There is however a clear distinct behaviour over the New Guinea region where simulated variance is lower than is observed, particularly in the annual cumulative values. The monthly and yearly distributions of the mean values show good agreement but some evidence of underestimated variability is seen in the West Coast and New Guinea regions. Moving to the regionwide mean distribution, there is the emergence of significant New Guinea bias which suggests that the region is not varying as coherently in the model as in the observations, however all of the other regions show good agreement between model and observation. Additionally, the distribution of maximum observed values shows reasonable agreement across the distribution. The regionwide maximum values do however show evidence of an underestimate of the most extreme components of the distribution over the New Guinea region. It is unclear how these regional biases are impacted by the larger size of the New Guinea region, which warrants further investigation but is considered out of scope for this work.

4.3 Temperature Testing

Next, we evaluate the quality of the 2 metre air temperature and 2 metre dew point temperature simulation over all five of the simulation domains. Like precipitation, our analysis begins with the examining the seasonal variation in the mean temperatures, as shown in figure 6. These values are calculated by finding every day in the corresponding dataset that belongs to the given month, finding the mean, maximum and minimum of the corresponding day and then taking the mean of all the days in the dataset. All five of these domains show very strong agreement between the simulations and the observations, with a nearly indistinguishable average in mean, minimum and maximum for the air temperature. The results for the dew point temperature values are weaker than air temperature simulations, with clear evidence of reduced simulated variability compared to observations. Still the average values are consistently well represented. Unsurprisingly, the spread between minimum and maximum values is considerably smaller for the dew point temperature than it is for the air temperature. The dew point values are also consistently smaller than the air temperature values, which is expected given the definition of dew point temperature,

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

but important to validate as it is not a requirement directly imposed by EGGS-WG. Interestingly, the dew point values seen over the New Guinea is very tightly grouped showing only a small difference between minimum and maximum values.

Following the approach used to compare the spatial correlations for precipitation values in figure 3, the set of correlation coefficients for every station pair of temperatures is shown in figure 7. Both the 2 m air temperature and the 2 m dew point temperature are analysed, as well as the cross correlation between these two variables. The cross correlation between air temperature and dew point temperature is important to examine, as this capability should be indirectly built into EGGS-WG. As the temporal resolution of the simulated data is hourly, instead of daily, a slightly different approach is needed than for precipitation. The most simplistic approach to account for these cycles is to sperate the data by both the month of the year and the time of day. This does result in 24 times as many grid point correlations as in figure 3, but should not have any other significant impacts on our analysis.

The temperature correlation results in figure 7 show a relatively large spread of values compared to the precipitation equivalent in figure 3. Focusing on the air temperature, we see the highest point density is tightly clustered around the 1:1 line, with minor variance either side of the line. This indicates that the spatial dynamics of air temperature are being well captured by the model in general. While there are some outlying values, the difference in magnitude of the correlation coefficients is relatively small. These are all relatively tightly clustered with only the New Guinea region showing some outliers although this may be at least partially due to having a much larger range of observed correlation values than the other regions. Similar, but weaker, correlations are seen in the dew point temperature difference. One key difference with the dew point temperature data is that the regions of highest density are no longer along the 1:1 line, instead they show a slight deviation indicating a weaker correlation in the simulation than the observations. These differences are minor enough that they should not have a large impact on model accuracy. Interestingly the New Guinea region shows much weaker simulation than all the other cases with the majority of the observed correlations very close to 1 while the simulated correlations are closer to 0.8. We suspect that this is due to statistical issues associated with the minimal range of dew point temperature values over this region as evidenced in figure 6.

The cross correlation between air temperature and dew point temperature shows more substantial differences between simulations and observations. This is somewhat to be expected as it is combining the variance associated with both of the temperatures. The cross correlations are well represented over the Canterbury region but show clear issues over the other regions. West Coast and Taupo show the bulk of simulations underestimate the correlations, but the effect is relatively minor. The Australia region shows a very wide range of correlations both observed and simulated that generally follows the expected 1 to 1 line. The New Guinea region is very complex with two clear peaks, one well represented and one not. Issues over this region could be related to the earlier identified dew point temperature issues, although further investigation would be needed for certainty. The cause of this regional difference is hard to determine, but may be related to the lower rainfall occurrence over the relatively well represented Canterbury and Australia regions. The highest density of grid point correlations fall relatively close to the 1:1 line, however, with a larger separation than seen in either the air temperature or dew point temperature data. This will have implications for the evaluating some simulation outputs, but the correlations still appear close enough to the observed correlations for the majority of purposes.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

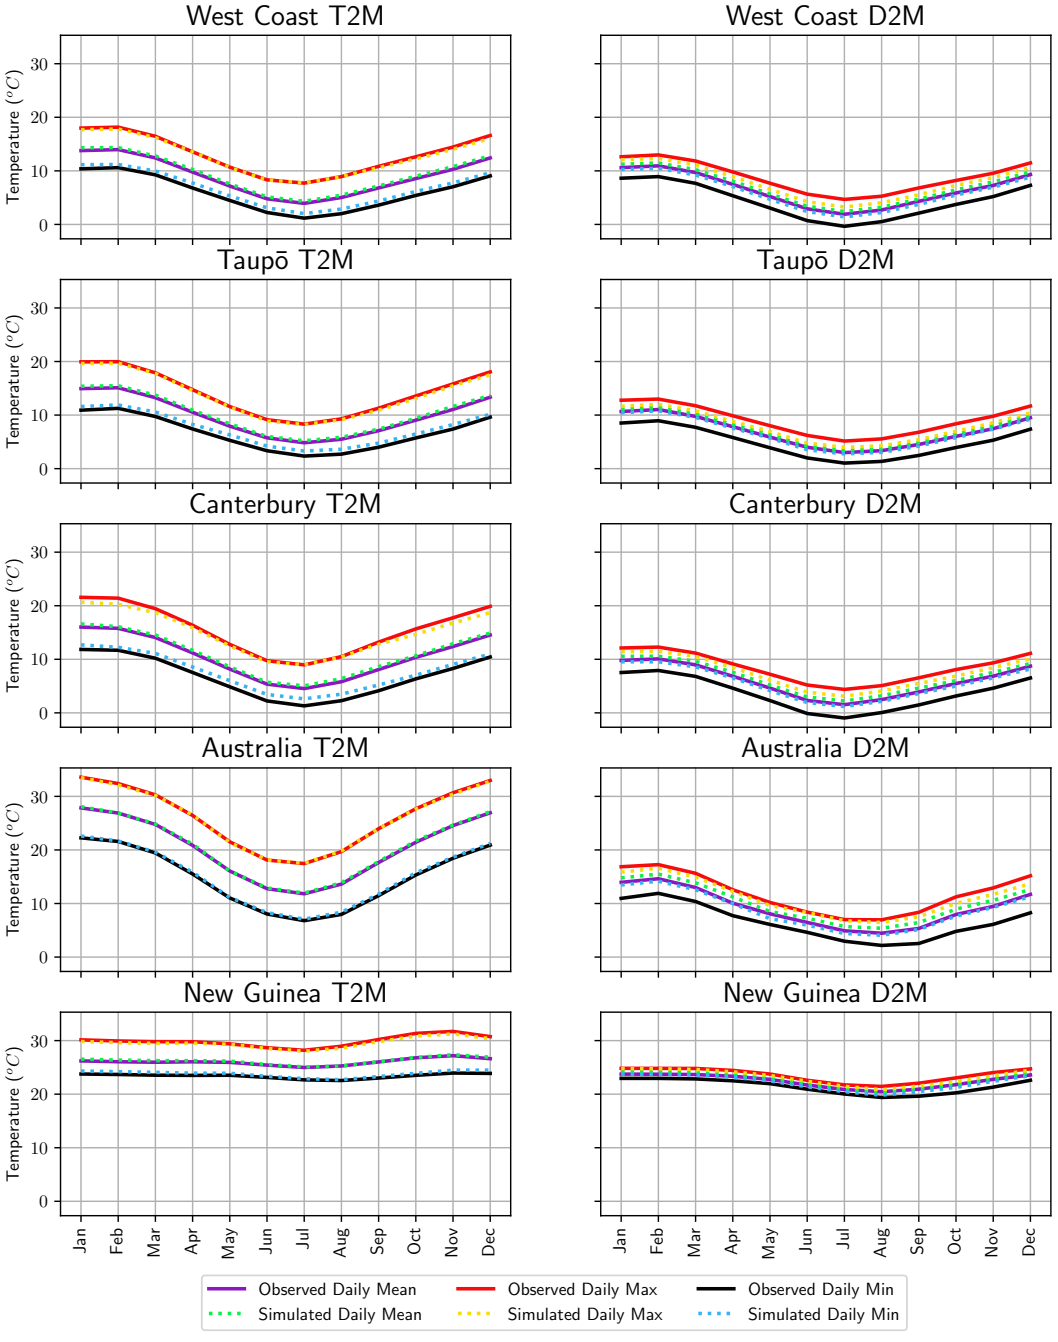

Figure 6. The seasonal cycle of the 2 metre air temperature (T2M) and the 2 metre dew point temperature (D2M) for both the observations and simulations over all five of the domains. Included for each plot is the daily mean, minimum and maximum value for each month of the year.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

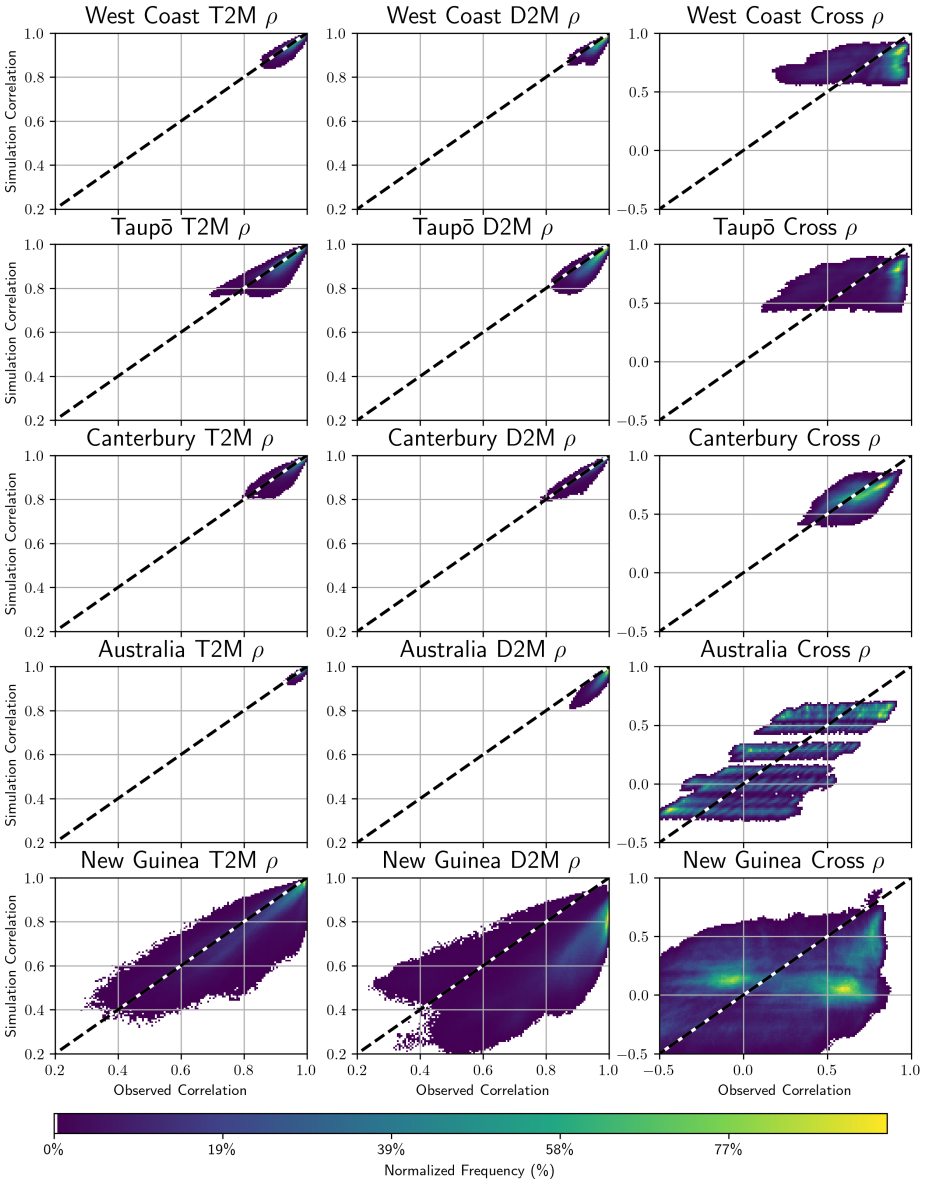

Figure 7. The Pearson correlation coefficient between every set of station pairs for both the simulation and observational data. Values are calculated for both the 2 m air temperature and 2 m dew point temperature data, as well as the correlation between these two variables. The colour of the points indicates the density of station pairs at a given correlation, ranging from purple for a low number of points to green for a high number of points. The included dashed line indicates the 1:1 correlation line that would be the theoretical ideal for a simulation.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

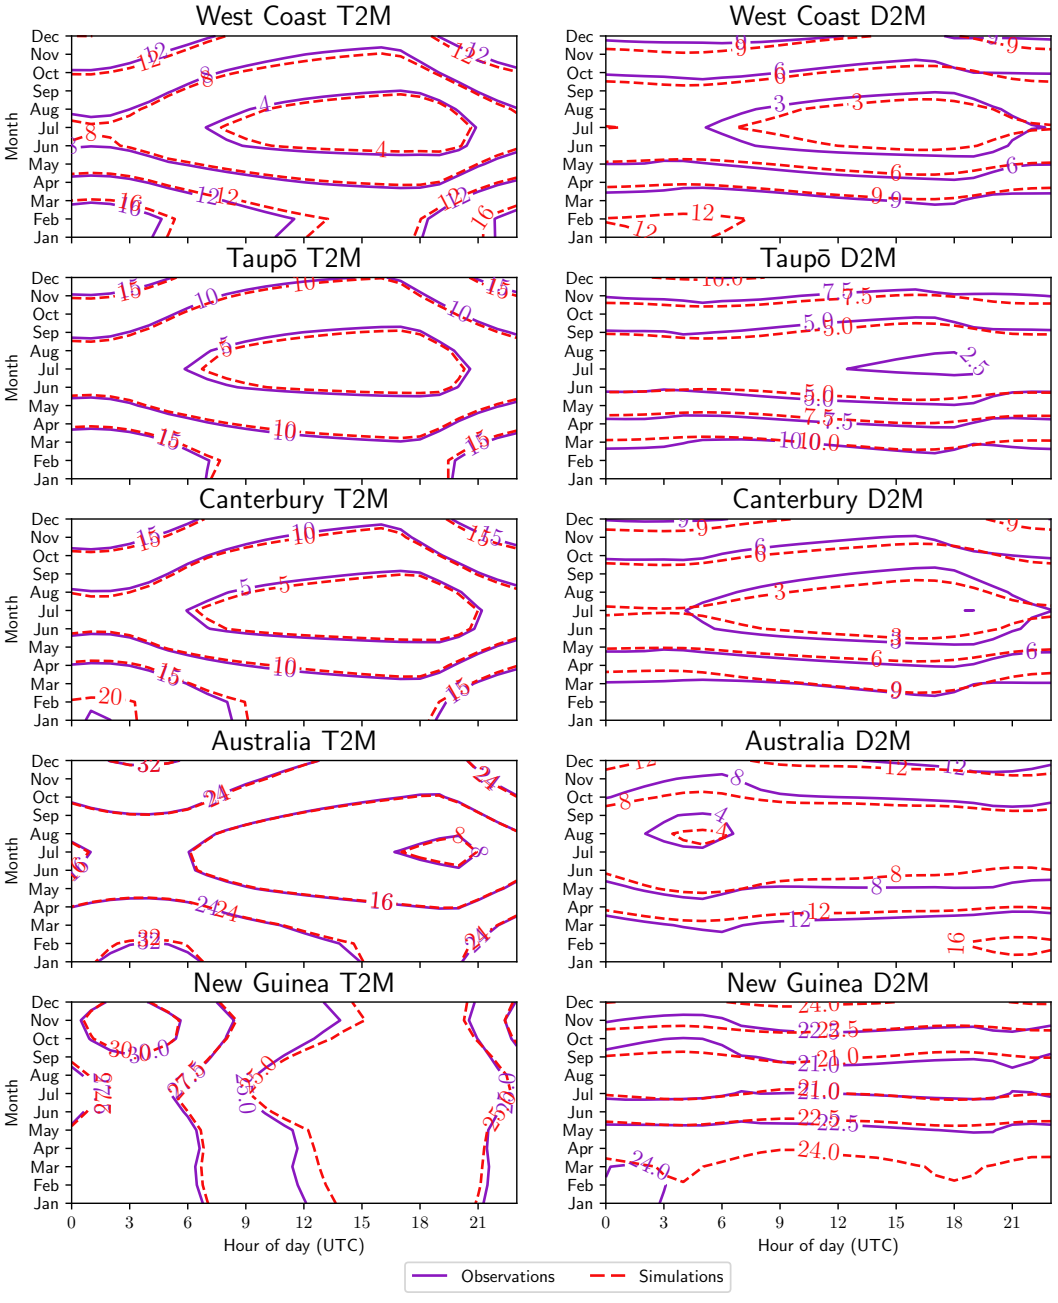

Figure 8. Contour plot of the diurnal variation of 2 m air temperature and dew point temperature for each month in both the simulations and observational data. Note that time on the x-axis is given in UTC and NZST is +12:00.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

As the temperature simulations are hourly, it also needs to be confirmed that the simulation is accurately capturing diurnal variability. The representation of the diurnal cycles of both air temperature and dew point temperature is examined in figure 8. Over all five of the regions, the air temperature values show very strong agreement between observations and simulations. The seasonality is very well captured, which is expected given the earlier results in figure 6. The dew point temperature is also well represented, but does show some minor differences between the observations and simulations. Generally, the representation is relatively consistent over all five of the domains. While the New Guinea region shows a different annual variability pattern than the other regions, this is unsurprising due to its location in the tropics.

While figures 6 and 8 establish that the average temperature values are well simulated by the model, they do not focus on the representation of the wider distribution. To examine the behavior over the entire distribution, a quantile-quantile plot is included in figure 9. Note, this value is calculated separately for each grid-point, month of data and time of day. This means that for example, a single dot will show the 50^{th} percentile of all values at 1 pm for a single month, effective removing biases associated with seasonal cycles from the analysis. Figure 9 only shows values for April specifically as seasonal cycles make the plot very difficult to interpret. The patterns seen here are highly representative of all months, with two extra months (January and August) included separately as examples in the supplementary information. For both the air temperature and dew point temperature, the simulation very accurately recreates the observational temperature values over all five of the investigated domains. There is a slightly better clustering to the 1:1 line in the air temperature than the dew point temperature, however the differences are relatively minor. Unlike the precipitation results shown in figure 5, the extreme values of the distribution do not appear to show any significant deviation from the 1:1 line. This suggests EGGS-WG would be more suitable for studying temperature extremes than it is for studying precipitation extremes. The relatively small range of dew point values over the New Guinea region also suggest the issues identified for this region in figure 7, are more statistical artifact than significant bias.

To further explore any potential temperature biases, we include figures showing the distribution over different timescales, spatial sizes and separately for temperature extremes in the supplementary information (figures S 11 - 40). The mean T2M and D2M values appear very well represented by the model over all time scales for all five of the different regions, both over all individual gridcells and in the regionwide mean values. Interestingly, the maximum values t2m are well generally represented, but there appear to be some issues with replicating the coldest temperatures produced by ERA5 and with the yearly d2m maximums, particularly over the Australian region. Given the climates of the specific regions we have examined, the ERA5 values appear to be unrealistically low and may warrant further investigation. However, since the EGGS-WG is purely based on the ERA5 statistical profile, it does highlight a minor failure in the simulation process. The majority of these results do reinforce that the model is in good alignment with the ERA5 output for all but the coldest temperatures.

4.4 Model Ensemble Variability

The earlier analysis, while comprehensive, is focused on a single model realization. One of the major benefits of stochastic weather generators is their ability to produce a range of simulations with near identical long-term statistical properties but generated from a wide range of realistic time series. This has significant benefits for understanding variability and presenting information to end users that can be interpreted with a storyline approach. To create this ensemble, each of the different

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

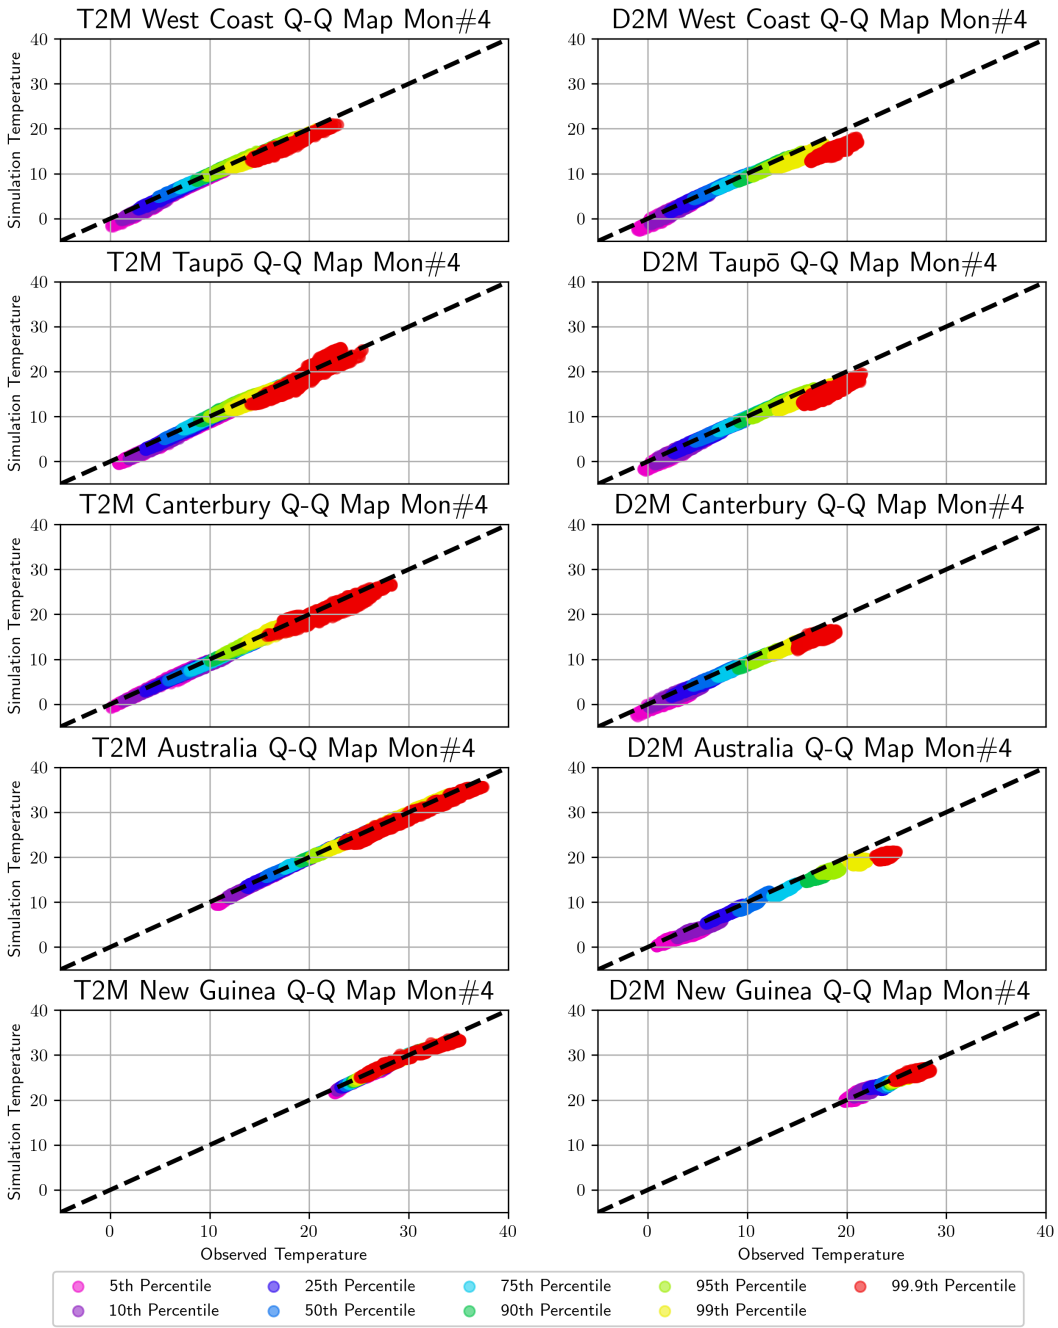

Figure 9. Quantile-quantile comparison of the 2 m air temperature and dew point temperatures for April. Values are calculated independently for each grid point in the domain and for each hour of the day. The percentiles included in this plot are the 5^{th}, 10^{th}, 25^{th}, 50^{th}, 75^{th}, 90^{th}, 95^{th} and 99^{th} percentiles. The dashed line indicates the 1:1 correlation line.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

simulations are separately initiated with the process described in section 3.6, which results in a completely different starting point for each simulation.

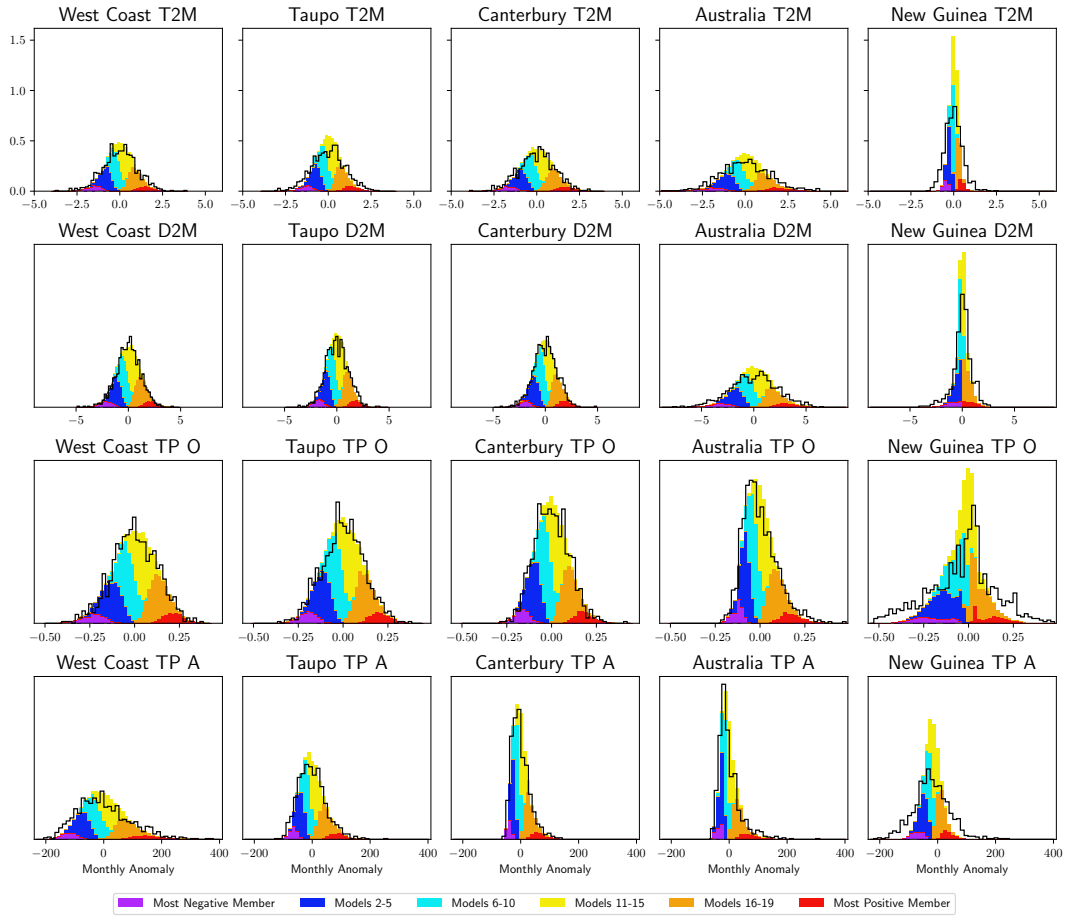

Figure 10. Monthly anomalies of rainfall occurrence, rainfall amount, 2 m air temperature and 2 m dew point temperature for 20 different simulation realizations. The coloring of the histogram indicates the ranking of ensemble members while the black line indicates the equivalent distribution from the ERA5 Land data.

Figure 10 shows is the ranked anomalies of our 20 run ensemble for 2m air temperature, 2m dewpoint temperature, rainfall occurence rate and accumulated rainfall. Effectively each month of our 20 model simulation is considered sperately and the models are split and ranked based on the size of their monthly anomaly. The normalized 20 model ensemble distribution is then compared directly to the equivilent values calculated from the ERA 5 Land dataset. This figure is included to test a few different features of the simulation; the quality of the data when agregated at a monthly level, the range of outcomes produced across an ensemble for a given month and how any biases indentified in these distributions differs over the different regions.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

Examining the 2m air temperature in figure 10, we see a consistent gaussian of similar width produced over all of the regions, strongly inline with the observed distribution, with the exception of the New Guinea region which shows a much more narrow gaussian. A good spread between emsemble members is seen in the other four regions with evidence of notable variation even in the 1^{st} and 20^{th} ranked ensemble members. The results for New Guinea are substantially different from the other regions with a significantly narrower distribution shown in the simulations relative to the the observations. Amongst the ensemble over New Guinea there is also evidence of the most extreme ensemble members showing anomalies close to zero. Very similar biases are seen for the dew point temperature, but with a slightly improved quality of simulation of New Guinea. For the precipitation occurence and accumulation, we see strong agreement between simulations and observations in the first four regions including for some distributions that are clearly non-gaussian in shape. For the New Guinea region we see some distinct behaviour that disagrees with the observations, similar to the temperatures this is mostly in the form of a reduced variability. Interestingly, in the precipitation occurrence there are cases with the most extreme ensemble members showing the monthly mean value. These are most likely due to the large variations due to the monsoon signals, especially during times of close to 100 percent precipitation occurence. Overall, the results presented show that the model does a good job replicating the observations and the ensemble is showing a reasonable degree of variability. The only exception is the New Guinea region which still shows a reasonable simulation but clear biases.

In addition to figure 10, we include two figures in the supplementary material one which shows the monthly anomalies in the ensemble as time series (figure S41) and one that shows the distribution of Pearson correlation coefficients between each of these different ensemble members (figure S42). Specifically, this is the correlation coefficient between time series of monthly mean 2 m air temperature, 2 m dew point temperature, precipitation occurrence and cumulative precipitation between all ensemble members. From these correlation coefficients we can see that the different ensemble members are effectively uncorrelated from one another (with the exception of a minor correlation in the New Guinea data). Combining this information with the results shown in figure 5 and 9, which demonstrate that the individual model runs successfully reproducing the variability within the observations, we can conclude that the individual ensemble runs are able to provide sufficiently variable simulations that are unrelated o each other. As discussed earlier there are some issues at the most extreme portions of the distribution but in general the variability seen between the ensemble members is good.

5 Conclusions

This paper covers the development of a new open source freely available stochastic weather generator, EGGS-WG. The structure and design of this model is based on a well established research framework covering Richardson style generators (Richardson, 1981; Wilks and Wilby, 1999) which has guided the design of many of the most prominent existing stochastic weather generators (Nicks et al., 1995; Semenov et al., 1998; Stockle et al., 2001; Chen et al., 2012a). However, many of these weather generators are lacking features that have been established in research following these initial models (Wilks, 2009; Abbasnezhadi et al., 2019). Modern features of EGGS-WG that are missing from older models include a gridded multisite approach, hourly temperature simulation, simulation of dew point temperature and relative humidity and a global terrestrial simulation

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

domain. By designing this model for portability, low resource requirements and ease of use, it is our aim to provide software that can be as widely adopted as possible.

Comprehensive validation of the EGGS-WG output is performed throughout this paper, focusing on three distinct regions over New Zealand, one in Australia and one in New Guinea. There are four key variables produced by EGGS-WG that are examined to evaluate performance; precipitation occurrence, precipitation amount, 2 m air temperature and 2 m dew point temperature. The mean values for all of these variables are investigated and show a strong agreement between the reanalysis and simulations. Further tests show that the seasonal cycle, diurnal cycles (for temperature only) and spread of the distribution of these variables is also well captured by the simulations. There is an identified reduction in quality with the high-end precipitation extremes (above 99^{th} percentile), although the rarity of these points limits the impact on the overall simulation. The ability of EGGS-WG to recreate observed spatial relationships is also tested and shows generally strong results, except for the cross correlation between air temperature and dew point temperature. The simulation of this relationship is still reasonable but does show consistent underestimation of the observed relationship. Additionally, the model is shown to produce variable time series that can produce a wide range of synthetic story lines that accurately capture the aggregate statistical properties of the observations.

Currently, there are a few major limitations to this model. First and most important is that the model is based entirely on the ERA5-Land data and as such will share any biases present in the ERA5 data. The most commonly discussed of these is a limited capability to represent the most extreme weather events particularly underestimating precipitation extremes. These issues may be improved with the upcoming release of ERA6 which we hope to integrate into our model quickly. The model simulation has also been shown to have a flawed representation of both precipitation extremes and minimum temperatures. Although the model bias produces too much extreme precipitation and the ERA5 produces too little so the exact impact is difficult to determine. There are also other potential limitations including accurate representation of domain wide weather events, the impact of a low depth markov chain and lack of relationship to global teleconnections such as ENSO, which are not evaluated here but may be a focus of future work.

There are also many different future updates being considered for the EGGS-WG model. Most prominent amongst these is adding a climate change forcing component to the model that would allow for direct integration of CMIP6 model data to impact future climate. The second priority is an expansion of the model to include support for other climate variables, most importantly wind vectors, incoming radiation and surface pressure. We have developed an internal version of the model that supports both of these feature, but substantial work on data quality, portability, data accessibility and ease of use is required for them to be added to a public release. There are more speculative features being considered, including integration with existing crop growth models (Nandan et al., 2024; Dumont et al., 2025) or streamflow models (Chen et al., 2012b; Sohrabi and Brissette, 2021), comparison with up and coming GAN based stochastic weather generators (Rampal et al., 2025) and exploration of alternative precipitation schemes to drive an hourly precipitation model. There is also the possibility of future research based on the direct output from the EGGS-WG model, particularly for the investigation of climate extremes. Future work may also move beyond ERA5-Land and rework the model to also work with other modern datasets.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

Code and data availability. The model produced is open source and publically accessible on GitHub under a GPL-3.0 license (https://github.com/alexschuddeboom/EGGS-WG). There is a specific version of record for this paper, V 1.1.0, that is also accessible on Zenodo at https://doi.org/10.5281/zenodo.15307044 (Schuddeboom, 2025). To simplify the ERA5-Land data access process, included with our model is a python script which can be used to download ERA5-Land data directly through their API. Also included on GitHub is a technical manual that briefly describes the operational procedure for the model.

Author contributions. Alex Schuddeboom: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing - Original Draft, Writing - Review & Editing; Christian Zammit: Conceptualization, Funding acquisition, Project administration, Supervision, Writing - Review & Editing; David Plew: Conceptualization, Funding acquisition, Supervision, Writing - Review & Editing; Piet Verburg: Conceptualization, Funding acquisition, Supervision, Writing - Review & Editing; Aidin Jabbari: Supervision, Writing - Review & Editing

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The authors would like to thank the European centre for Medium-range Weather Forecasting for the production of the ERA5-Land dataset, without which this model would not be possible. The ERA5-Land data is accessible through the Copernicus Climate Change Service Climate Data Store Copernicus Climate Change Service (2022). Additionally, we would like to thank the Copernicus Climate Change Service Climate Data Store for making the data widely available and publicly accessible. Funding for this project was provided through both the New Zealand Ministry of Business, Innovation and Employment (MBIE) via the 2022 Endeavour Fund - Smart Ideas (contract C01X2205) and by the Aotearoa New Zealand Tāwhia te Mana Research Fellowships, administered by the Royal Society Te Apārangi (contract MTP-UOC2401).

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

References

Abbasnezhadi, K., Rousseau, A. N., Wruth, A. M., and Zahmatkesh, Z.: Synchronized generation of high-resolution gridded precipitation and temperature fields, Journal of Hydrology, 573, 631–647, https://doi.org/10.1016/j.jhydrol.2019.03.096, 2019.

Ailliot, P., Allard, D., Monbet, V., and Naveau, P.: Stochastic weather generators: an overview of weather type models, Journal de la Société Française de Statistique, 156, 101–113, 2015.

Baigorria, G. A. and Jones, J. W.: GiST: A Stochastic Model for Generating Spatially and Temporally Correlated Daily Rainfall Data, Journal of Climate, 23, 5990–6008, https://doi.org/10.1175/2010JCLI3537.1, 2010.

Brissette, F. P., Khalili, M., and Leconte, R.: Efficient stochastic generation of multi-site synthetic precipitation data, Journal of Hydrology, 345, 121–133, https://doi.org/10.1016/j.jhydrol.2007.06.035, 2007.

Chen, J., Brissette, F., and Leconte, R.: WeaGETS – a Matlab-based daily scale weather generator for generating precipitation and temperature, Procedia Environmental Sciences, 13, 2222–2235, https://doi.org/10.1016/j.proenv.2012.01.211, 2012a.

Chen, J., Brissette, F. P., and Leconte, R.: Downscaling of weather generator parameters to quantify hydrological impacts of climate change, Climate Research, 51, 185–200, https://doi.org/10.3354/cr01062, 2012b.

Chen, J., Brissette, F., and Zhang, X. J.: Multi-Site Stochastic Weather Generator for Daily Precipitation and Temperature, Transactions of the ASABE, 57, 1375–1391, https://doi.org/10.13031/trans.57.10685, 2014.

Copernicus Climate Change Service: Copernicus Climate Change Service (C3S) Climate Data Store (CDS), https://doi.org/10.24381/cds.e2161bac, 2022.

Cordano, E. and Eccel, E.: Tools for stochastic weather series generation in R environment, Italian Journal of Agrometeorology, 21, 31–42, https://doi.org/10.19199/2016.3.2038-5625.031, 2016.

Dawkins, L. C., Osborne, J. M., Economou, T., Darch, G. J., and Stoner, O. R.: The Advanced Meteorology Explorer: a novel stochastic, gridded daily rainfall generator, Journal of Hydrology, 607, 127 478, https://doi.org/10.1016/j.jhydrol.2022.127478, 2022.

Dumont, M., Etheridge, Z., Curtis, A., Beukes, P., and Schuddeboom, A.: Current climate variability of pastoral yield: a case study in Canterbury, New Zealand, Climatic Change, 178, 1–20, https://doi.org/10.1007/s10584-025-03946-z, 2025.

Friend, A., Stevens, A., Knox, R., and Cannell, M.: A process-based, terrestrial biosphere model of ecosystem dynamics (Hybrid v3.0), Ecological Modelling, 95, 249–287, https://doi.org/10.1016/S0304-3800(96)00034-8, 1997.

Fullhart, A., Ponce-Campos, G. E., Meles, M. B., McGehee, R. P., Armendariz, G., Oliveira, P. T. S., Das Neves Almeida, C., de Araújo, J. C., Nel, W., and Goodrich, D. C.: Gridded 20-year climate parameterization of Africa and South America for a stochastic weather generator (CLIGEN), Big Earth Data, 7, 349–374, https://doi.org/10.1080/20964471.2022.2136610, 2023.

Fullhart, A. T., Nearing, M. A., Armendariz, G., and Weltz, M. A.: Climate benchmarks and input parameters representing locations in 68 countries for a stochastic weather generator, CLIGEN, Earth System Science Data, 13, 435–446, https://doi.org/10.5194/essd-13-435-2021, 2021.

Hansen, J. E. and Driscoll, D. M.: A Mathematical Model for the Generation of Hourly Temperatures, Journal of Applied Meteorology, 16, 935–948, https://doi.org/10.1175/1520-0450(1977)016<0935:AMMFTG>2.0.CO;2, 1977.

Ji, H. K., Mirzaei, M., Lai, S. H., Dehghani, A., and Dehghani, A.: Implementing generative adversarial network (GAN) as a data-driven multi-site stochastic weather generator for flood frequency estimation, Environmental Modelling & Software, 172, 105 896, https://doi.org/10.1016/j.envsoft.2023.105896, 2024.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

Keller, D. E., Fischer, A. M., Liniger, M. A., Appenzeller, C., and Knutti, R.: Testing a weather generator for downscaling climate change projections over Switzerland, International Journal of Climatology, 37, 928–942, https://doi.org/10.1002/joc.4750, 2017.

Khalili, M., Brissette, F., and Leconte, R.: Stochastic multi-site generation of daily weather data, Stochastic Environmental Research and Risk Assessment, 23, 837–849, https://doi.org/10.1007/s00477-008-0275-x, 2009.

Kilsby, C., Jones, P., Burton, A., Ford, A., Fowler, H., Harpham, C., James, P., Smith, A., and Wilby, R.: A daily weather generator for use in climate change studies, Environmental Modelling & Software, 22, 1705–1719, https://doi.org/10.1016/j.envsoft.2007.02.005, 2007.

Lavers, D. A., Simmons, A., Vamborg, F., and Rodwell, M. J.: An evaluation of ERA5 precipitation for climate monitoring, Quarterly Journal of the Royal Meteorological Society, 148, 3152–3165, https://doi.org/10.1002/qj.4351, 2022.

Muñoz-Sabater, J., Dutra, E., Agustí-Panareda, A., Albergel, C., Arduini, G., Balsamo, G., Boussetta, S., Choulga, M., Harrigan, S., Hersbach, H., Martens, B., Miralles, D. G., Piles, M., Rodríguez-Fernández, N. J., Zsoter, E., Buontempo, C., and Thépaut, J.-N.: ERA5-Land: a state-of-the-art global reanalysis dataset for land applications, Earth System Science Data, 13, 4349–4383, https://doi.org/10.5194/essd-13-4349-2021, 2021.

Najibi, N., Perez, A. J., Arnold, W., Schwarz, A., Maendly, R., and Steinschneider, S.: A statewide, weather-regime based stochastic weather generator for process-based bottom-up climate risk assessments in California – Part I: Model evaluation, Climate Services, 34, https://doi.org/10.1016/j.cliser.2024.100489, 2024a.

Najibi, N., Perez, A. J., Arnold, W., Schwarz, A., Maendly, R., and Steinschneider, S.: A statewide, weather-regime based stochastic weather generator for process-based bottom-up climate risk assessments in California – Part II: Thermodynamic and dynamic climate change scenarios, Climate Services, 34, https://doi.org/10.1016/j.cliser.2024.100485, 2024b.

Nandan, R., Bandaru, V., Meduri, P., Jones, C., and Lollato, R.: Evaluating the utility of weather generators in crop simulation models for in-season yield forecasting, Agricultural Systems, 220, 104 082, https://doi.org/10.1016/j.agsy.2024.104082, 2024.

Nelder, J. A. and Mead, R.: A Simplex Method for Function Minimization, The Computer Journal, 7, 308–313, https://doi.org/10.1093/comjnl/7.4.308, 1965.

Nicks, A. D., Lane, L. J., and Gander, G. A.: Chapter 2: Weather Generator, in: USDA–Water Erosion Prediction Project hillslope profile and watershed model documentation, July, p. 2.1–2.22, ISBN 10, 1995.

Parlange, M. B. and Katz, R. W.: An Extended Version of the Richardson Model for Simulating Daily Weather Variables, Journal of Applied Meteorology, 39, 610–622, https://doi.org/10.1175/1520-0450-39.5.610, 2000.

Peleg, N., Fatichi, S., Paschalis, A., Molnar, P., and Burlando, P.: An advanced stochastic weather generator for simulating 2-D high-resolution climate variables, Journal of Advances in Modeling Earth Systems, 9, 1595–1627, https://doi.org/10.1002/2016MS000854, 2017.

Powell, M. J. D.: An efficient method for finding the minimum of a function of several variables without calculating derivatives, The Computer Journal, 7, 155–162, https://doi.org/10.1093/comjnl/7.2.155, 1964.

Rahat, S. H., Steinschneider, S., Kucharski, J., Arnold, W., Olzewski, J., Walker, W., Maendly, R., Wasti, A., and Ray, P.: Characterizing Hydrologic Vulnerability under Nonstationary Climate and Antecedent Conditions Using a Process-Informed Stochastic Weather Generator, Journal of Water Resources Planning and Management, 148, https://doi.org/10.1061/(ASCE)WR.1943-5452.0001557, 2022.

Rampal, N., Gibson, P. B., Sherwood, S., Abramowitz, G., and Hobeichi, S.: A Reliable Generative Adversarial Network Approach for Climate Downscaling and Weather Generation, Journal of Advances in Modeling Earth Systems, 17, 1–28, https://doi.org/10.1029/2024MS004668, 2025.

Rebonato, R. and Jaeckel, P.: The Most General Methodology to Create a Valid Correlation Matrix for Risk Management and Option Pricing Purposes, SSRN Electronic Journal, 17, 1767–1768, https://doi.org/10.2139/ssrn.1969689, 2000.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

Richardson, C. W.: Stochastic simulation of daily precipitation, temperature, and solar radiation, Water Resources Research, 17, 182–190, https://doi.org/10.1029/WR017i001p00182, 1981.

Richardson, C. W.: Weather Simulation for Crop Management Models, Transactions of the ASAE, 28, 1602–1606, https://doi.org/10.13031/2013.32484, 1985.

Richardson, C. W. and Wright, D. A.: WGEN: A Model for Generating Daily Weather Variables, United States Department of Agriculture, Agriculture Research Service ARS-8, p. 83, ftp://ftp.biosfera.dea.ufv.br/users/francisca/Franciz/papers/Richardson&Wright.pdf, 1984.

Schuddeboom, A.: EGGS-WG V1.1.0, https://doi.org/10.5281/zenodo.15307044, 2025.

Semenov, M., Brooks, R., Barrow, E., and Richardson, C.: Comparison of the WGEN and LARS-WG stochastic weather generators for diverse climates, Climate Research, 10, 95–107, https://doi.org/10.3354/cr010095, 1998.

Sohrabi, S. and Brissette, F. P.: Evaluation of a stochastic weather generator for long-term ensemble streamflow forecasts, Hydrological Sciences Journal, 66, 474–487, https://doi.org/10.1080/02626667.2021.1873343, 2021.

Sommer, P. S. and Kaplan, J. O.: A globally calibrated scheme for generating daily meteorology from monthly statistics: Global-WGEN (GWGEN) v1.0, Geoscientific Model Development, 10, 3771–3791, https://doi.org/10.5194/gmd-10-3771-2017, 2017.

Steinschneider, S. and Brown, C.: A semiparametric multivariate, multisite weather generator with low-frequency variability for use in climate risk assessments, Water Resources Research, 49, 7205–7220, https://doi.org/10.1002/wrcr.20528, 2013.

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

Stockle, C., Nelson, R. L., Donatelli, M., and Castellví, F.: ClimGen: A flexible weather generation program, Proceedings 2nd International Symposium Modelling Cropping Systems, pp. 229–230, 2001.

Stoner, O. and Economou, T.: An advanced hidden Markov model for hourly rainfall time series, Computational Statistics and Data Analysis, 152, 107 045, https://doi.org/10.1016/j.csda.2020.107045, 2020.

Vesely, F. M., Paleari, L., Movedi, E., Bellocchi, G., and Confalonieri, R.: Quantifying Uncertainty Due to Stochastic Weather Generators in Climate Change Impact Studies, Scientific Reports, 9, 9258, https://doi.org/10.1038/s41598-019-45745-4, 2019.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, I., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., Vijaykumar, A., Bardelli, A. P., Rothberg, A., Hilboll, A., Kloeckner, A., Scopatz, A., Lee, A., Rokem, A., Woods, C. N., Fulton, C., Masson, C., Häggström, C., Fitzgerald, C., Nicholson, D. A., Hagen, D. R., Pasechnik, D. V., Olivetti, E., Martin, E., Wieser, E., Silva, F., Lenders, F., Wilhelm, F., Young, G., Price, G. A., Ingold, G.-L., Allen, G. E., Lee, G. R., Audren, H., Probst, I., Dietrich, J. P., Silterra, J., Webber, J. T., Slavič, J., Nothman, J., Buchner, J., Kulick, J., Schönberger, J. L., de Miranda Cardoso, J. V., Reimer, J., Harrington, J., Rodríguez, J. L. C., Nunez-Iglesias, J., Kuczynski, J., Tritz, K., Thoma, M., Newville, M., Kümmerer, M., Bolingbroke, M., Tartre, M., Pak, M., Smith, N. J., Nowaczyk, N., Shebanov, N., Pavlyk, O., Brodtkorb, P. A., Lee, P., McGibbon, R. T., Feldbauer, R., Lewis, S., Tygier, S., Sievert, S., Vigna, S., Peterson, S., More, S., Pudlik, T., Oshima, T., Pingel, T. J., Robitaille, T. P., Spura, T., Jones, T. R., Cera, T., Leslie, T., Zito, T., Krauss, T., Upadhyay, U., Halchenko, Y. O., and Vázquez-Baeza, Y.: SciPy 1.0: fundamental algorithms for scientific computing in Python, Nature Methods, 17, 261–272, https://doi.org/10.1038/s41592-019-0686-2, 2020.

https://doi.org/10.5194/egusphere-2026-1651
Preprint. Discussion started: 8 May 2026

Vishwanathan, G., McDonald, A., Noble, C., A. Stone, D., Rosier, S., Schuddeboom, A., Kreft, P., Macara, G., Carey-Smith, T., and Bodeker, G. E.: Regional Characteristics of Extreme Precipitation Events Over Aotearoa New Zealand, SSRN Electronic Journal, 44, https://doi.org/10.2139/ssrn.4621746, 2023a.

Vishwanathan, G., McDonald, A., Stone, D. A., Rosier, S., Rana, S., and Noble, C.: Mean and extreme precipitation over Aotearoa New Zealand: A comparison across multiple different estimation techniques, International Journal of Climatology, 43, 3072–3093, https://doi.org/10.1002/joc.8017, 2023b.

Whitehead, M. G. and Bebbington, M. S.: Brief communication: SWM – stochastic weather model for precipitation-related hazard assessments using ERA5-Land data, Natural Hazards and Earth System Sciences, 24, 1929–1935, https://doi.org/10.5194/nhess-24-1929-2024, 2024.

Wilks, D. S.: Adapting stochastic weather generation algorithms for climate change studies, Climatic Change, 22, 67–84, https://doi.org/10.1007/BF00143344, 1992.

Wilks, D. S.: Multisite generalization of a daily stochastic precipitation generation model, Journal of Hydrology, 210, 178–191, https://doi.org/10.1016/S0022-1694(98)00186-3, 1998.

Wilks, D. S.: Multisite downscaling of daily precipitation with a stochastic weather generator, Climate Research, 11, 125–136, https://doi.org/10.3354/cr011125, 1999a.

Wilks, D. S.: Simultaneous stochastic simulation of daily precipitation, temperature and solar radiation at multiple sites in complex terrain, Agricultural and Forest Meteorology, 96, 85–101, https://doi.org/10.1016/S0168-1923(99)00037-4, 1999b.

Wilks, D. S.: High-resolution spatial interpolation of weather generator parameters using local weighted regressions, Agricultural and Forest Meteorology, 148, 111–120, https://doi.org/10.1016/j.agrformet.2007.09.005, 2008.

Wilks, D. S.: A gridded multisite weather generator and synchronization to observed weather data, Water Resources Research, 45, 1–11, https://doi.org/10.1029/2009WR007902, 2009.

Wilks, D. S.: Stochastic weather generators for climate-change downscaling, part II: Multivariable and spatially coherent multisite downscaling, Wiley Interdisciplinary Reviews: Climate Change, 3, 267–278, https://doi.org/10.1002/wcc.167, 2012.

Wilks, D. S. and Wilby, R. L.: The weather generation game: a review of stochastic weather models, Progress in Physical Geography: Earth and Environment, 23, 329–357, https://doi.org/10.1177/030913339902300302, 1999.