Climate Change Projections: Methods Report

Horizons Region, New Zealand

Climate Change Projections from CCAM-Downscaled CMIP6 Simulations

---

Prepared by

Bodeker Scientific

www.bodekerscientific.com

For

Horizons Regional Council

Report:	Climate Impacts Assessment
Baseline:	1986–2005
Scenarios:	SSP1-2.6, SSP2-4.5, SSP3-7.0, SSP5-8.5
Future periods:	2021-2040, 2041-2060, 2080-2099

September 23, 2026

Contents

1 Introduction **2**

2 Developing a more comprehensive estimate of uncertainty on projections **2**

3 Generating observational data to validate the CCAM baselines **5**

3.1 Temperature 5

3.2 Precipitation 11

1 Introduction

In producing a climate analysis report for the Horizons Region, Bodeker Scientific developed and applied several advanced methods for data analysis and presentation. To avoid encumbering the other reports generated as part of the collection of reports, these methods are summarised in this methods-specific report.

2 Developing a more comprehensive estimate of uncertainty on projections

The projections synthesised in the variable-specific reports were generated using the Conformal Cubic Atmospheric Model (CCAM) by Earth Sciences New Zealand (ESNZ). We acknowledge ESNZ and thank them for making these simulations freely and easily available. CCAM is a regional climate model that dynamically downscales outputs from global climate models to produce higher-resolution climate projections over New Zealand. Papers describing the generation of these climate projections for New Zealand are available here and here.

Regional climate modelling is necessary because global climate models typically operate at spatial resolutions that are too coarse to adequately represent New Zealand's complex topography. Mountains, coastlines, and local climatic gradients strongly influence regional climate and require higher spatial resolution to be represented realistically.

Six CMIP6 global climate models were used to drive the CCAM simulations:

1. ***ACCESS-CM2*** is the second-generation coupled climate model developed by the Australian research community through the Australian Community Climate and Earth System Simulator (ACCESS) project. It combines atmospheric, oceanic, sea-ice, and land-surface components and was developed jointly by the Australian Bureau of Meteorology and Commonwealth Scientific and Industrial Research Organisation (CSIRO). The model is widely used for climate projections in the Australia/New Zealand region, including studies of temperature, precipitation, and extreme events.
2. ***AWI-CM-1-1-MR*** is a coupled Earth system model developed by the Alfred Wegener Institute in Germany. One of its key strengths is in representing ocean and sea-ice processes. Its ocean component uses an unstructured mesh, allowing much finer spatial resolution in regions of interest such as the polar regions while retaining manageable computational costs elsewhere. The model is frequently used in studies of polar climate change and large-scale ocean circulation.
3. ***CNRM-CM6-1*** is developed by the Centre National de Recherches Météorologiques and Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS) in France. It includes interactive representations of the atmosphere, ocean, land surface, and sea ice, together with substantial improvements over earlier versions in cloud physics and surface processes. The model is widely used in both European and international climate assessments.
4. ***EC-Earth3*** is a European climate model developed by a consortium of research institutions from several countries, building on operational weather-forecasting systems from the European Centre for Medium-Range Weather Forecasts (ECMWF). It incorporates sophisticated atmosphere, ocean, land, and sea-ice components and is designed to support a broad range of climate applications, from seasonal forecasting to long-term climate projections. It performs particularly well in reproducing large-scale patterns of historical climate.
5. ***GFDL-ESM4*** is an Earth System Model developed by the Geophysical Fluid Dynamics Laboratory (GFDL), part of the US National Oceanic and Atmospheric Administration (NOAA). In addition to the physical climate system, it includes detailed representations of biogeochemical cycles, allowing interactions between climate, ecosystems, and atmospheric composition to be simulated. The model is widely used for studies of carbon cycling, ocean processes, and future climate change.
6. ***NorESM2-MM*** is the medium-resolution version of the second-generation Norwegian Earth System Model,

developed by a consortium led by the Norwegian Climate Centre. It is based partly on components from the Community Earth System Model but incorporates substantial Norwegian developments, particularly in aerosol processes and ocean dynamics. The model is commonly used in research on high-latitude climate, atmospheric circulation, and climate feedbacks.

The availability of a set of six ensemble members, rather than relying on a single simulation, is a key strength of the modelling framework used in this assessment. Ensemble approaches allow both the central tendency and uncertainty of projected changes to be assessed. That said, an ensemble of six models does not capture the full range of plausible model uncertainty. To provide a broader estimate of the possible uncertainties in the climate projections, we have considered the equilibrium climate sensitivity (ECS) of each of the six models in the context of the much wider range of 53 CMIP6 models (see Figure 1). The ECS of a climate model is a measure

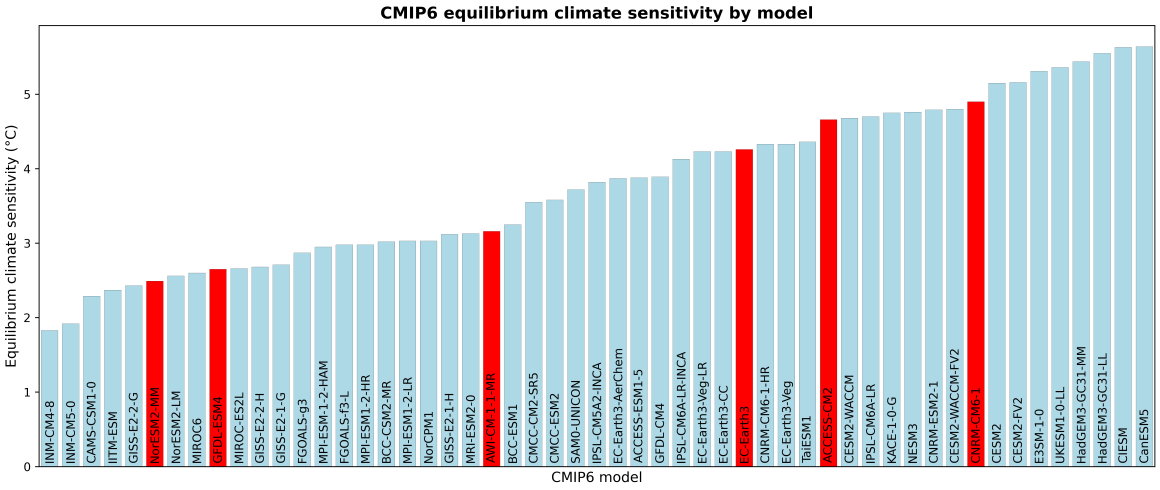


Figure 1: The location of the six guiding CMIP6 models across the full range of equilibrium climate sensitivities (ECSs) of all 53 CMIP6 models. The original of this figure can be downloaded here.

of how strongly the simulated climate in the model responds to an increase in atmospheric greenhouse gas (GHG) concentrations. More specifically, it is the amount by which the global average surface temperature would eventually rise if the concentration of carbon dioxide in the atmosphere were doubled and then held at that level for long enough for the climate system to reach a new equilibrium. For example, if a climate model has an ECS of 3°C, this means that doubling carbon dioxide would eventually make the Earth about 3°C warmer on average than it was before the increase. The word 'equilibrium' is important. The full warming does not happen immediately because the oceans absorb large amounts of heat and take many decades to centuries to adjust. ECS therefore describes the long-term response, rather than the warming expected within the next few decades. Different climate models have different ECS values because they represent climate processes somewhat differently. In particular, they differ in the strength of feedbacks involving water vapour, clouds, sea ice and other parts of the climate system.

We consider here projected changes in the number of growing degree days (GDDs) over a 5°C threshold (hereafter GDD5) as an example of the application of the ECS method in expanding the range of projection uncertainty beyond that which can be derived from the six models. Projected changes in the number of growing degree days, under a selected SSP scenario (the choice of scenario is not material to this methodological demonstration), for a randomly selected location in the Horizons Region, as a function of the equilibrium climate sensitivity for each model, is shown in Figure 2. It is clear that there is some dependence of the projected change in growing degree days on the model ECS; as expected, models with a higher ECS tend to simulate larger increases in the number of growing degree days, although there are some deviations, e.g., the CNRM-CM6-1 model generates a smaller GDD5 change than would be expected from its ECS. By fitting a straight line to these data we can quantify, with some uncertainty (given that the models do not sit perfectly on a straight line), the dependence of the change in growing degree days on the ECS of the six underlying models. This straight line fit, together with its uncertainty (5^{th} to 95^{th} percentile), is shown in Figure 3. As shown on the figure, in this

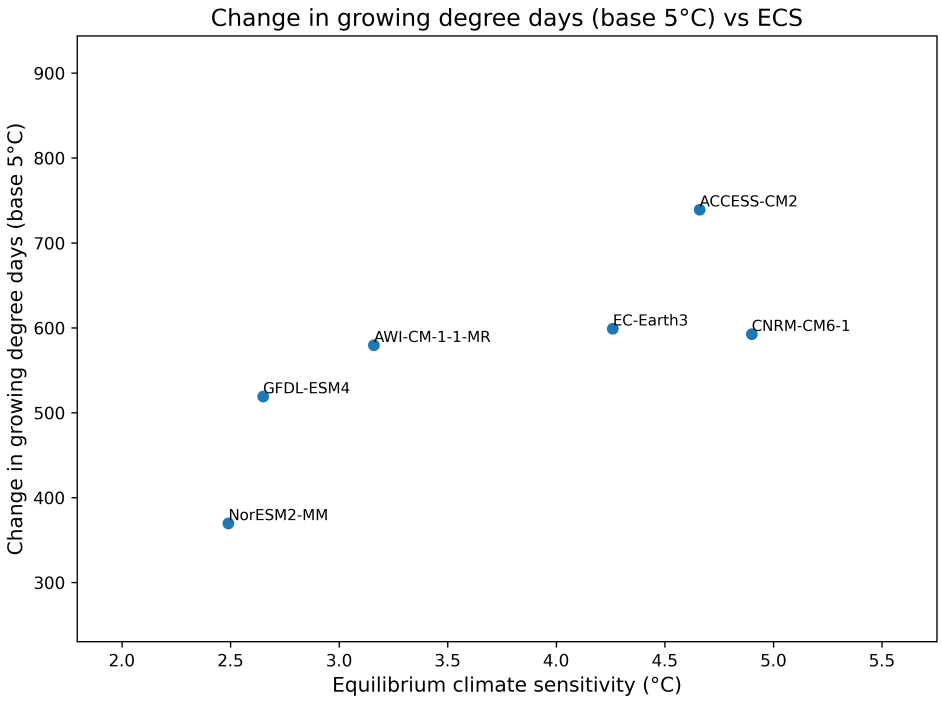


Figure 2: Projected changes in the number of growing degree days for a randomly selected location in the Horizons Region, under a selected SSP scenario, as a function of the equilibrium climate sensitivity (ECS) of each of the underlying six guiding CMIP6 models.

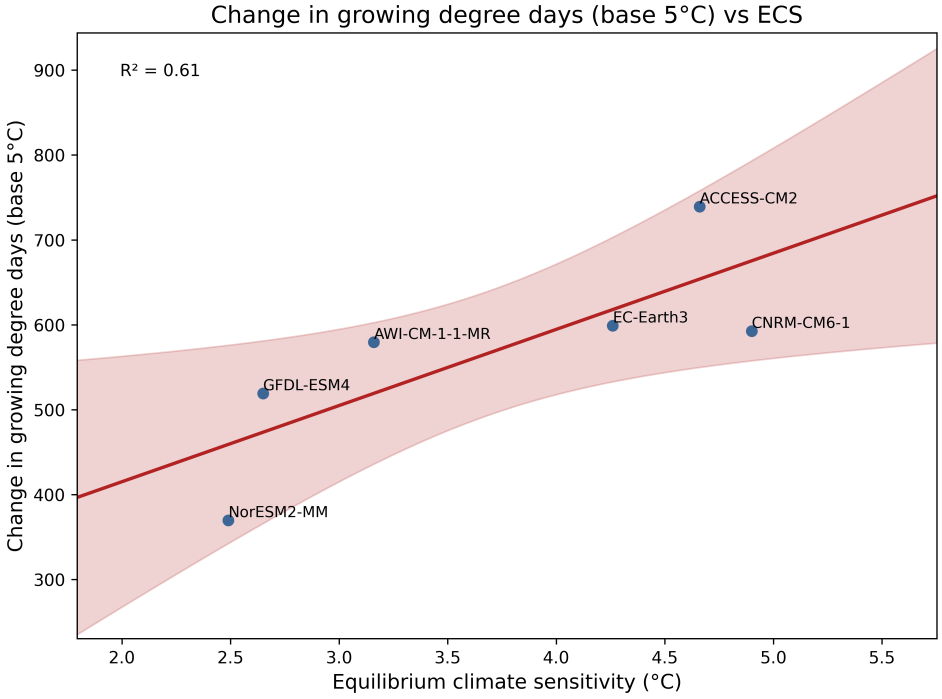


Figure 3: The straight line fit to the projected changes in the number of growing degree days shown in Figure 2, together with its 5^{th} to 95^{th} percentile uncertainty range (shaded region).

particular case the straight line explains 61% of the variance in the projected changes in GDD5 as a function of ECS. By evaluating the fitted regression and its uncertainty distribution at each of the 53 CMIP6 ECS values (see Figure 4), and extracting percentiles from the 1^{st} to the 99^{th} percentile at each ECS value, we generate 5,247 possible GDD5 change data points from which we can calculate more robust uncertainty statistics.

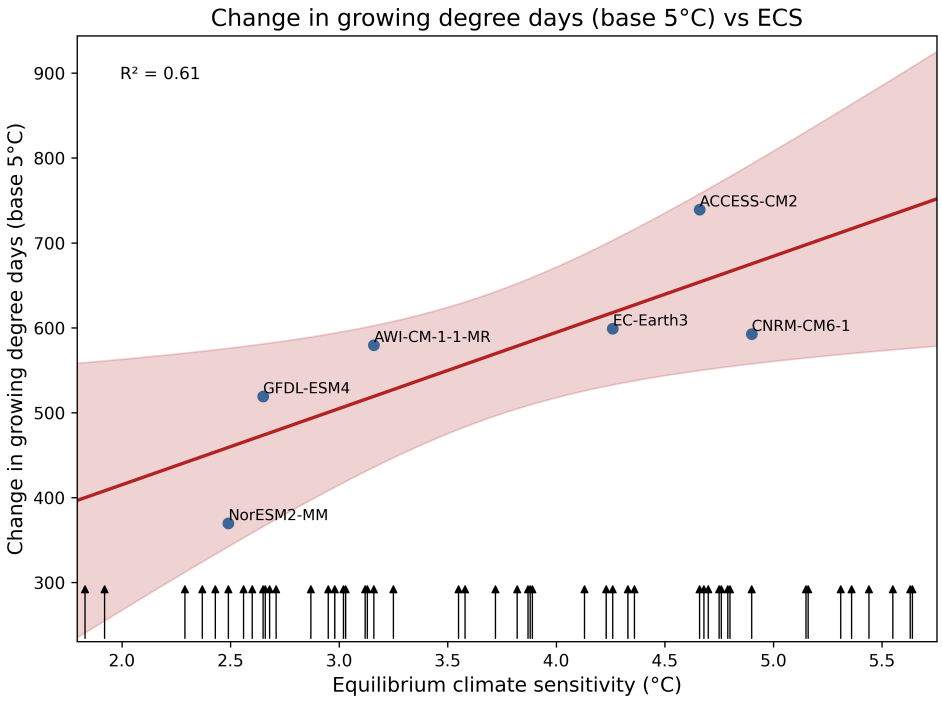


Figure 4: As for Figure 3 but now showing on the x-axis the ECS values of the 53 CMIP6 models, together with its 5^{th} to 95^{th} percentile uncertainty range (shaded region).

The same process is repeated in each model grid cell over the Horizons Region and for each of the 25 climate variables that have been analysed. It makes it possible to calculate the 5^{th} and 95^{th} percentile (P5 and P95) fields that are then compared with the minimum and maximum fields across the six CCAM models to assess the extent to which those minimum and maximum fields can be viewed as valid proxies for the 5^{th} and 95^{th} percentile fields of the more complete distribution of uncertainties. When averaged over sub-regions, it is then these statistics that form the basis for the 'violin plots' that are shown in the climate variable-specific reports. An example of one such plot is shown in Figure 5. The shaded areas in Figure 5 are the violin plots. At each vertical location on the violin plot, the width is indicative of the likelihood of that value (read off the Y axis). As can be seen, the absolute maximum and absolute minimum in the uncertainty of the projected changes often exceed the minimum and/or the maximum across the six models. The utility of this method to explore the uncertainty range in the projections depends on the extent to which ECS explains the scatter in the projections across the six CCAM model simulations (see Figure 6). In this case R^2 values vary between $\sim$0.43 and $\sim$0.65, indicating that ECS explains between 43% and 65% of the variance in GDD5, depending on location. In general, locations with lower R^2 value will have larger estimated uncertainties on their projections using this ECS-based method.

3 Generating observational data to validate the CCAM baselines

3.1 Temperature

An hourly temperature data set extending from 1970 to 2024 was generated to provide a baseline comparison for several of the temperature-based CCAM climate variables. The European Centre for Medium-Range

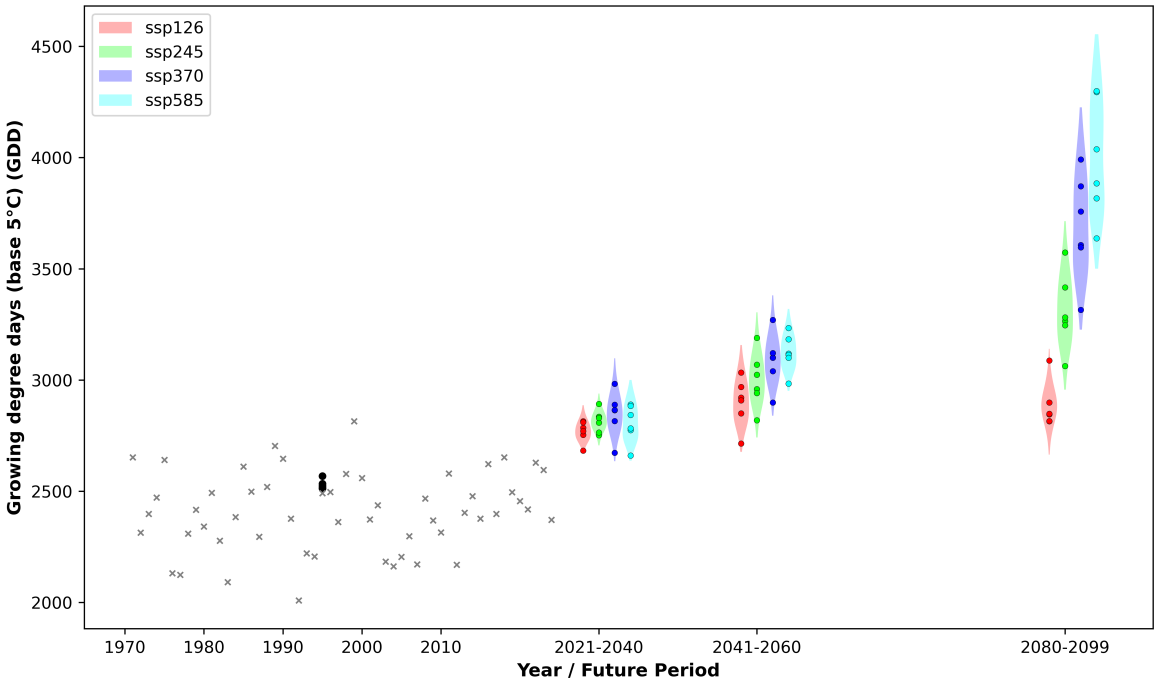


Figure 5: The time evolution of growing degree days (base 5°C) from the 1986-2005 baseline in units of GDD over the NorthernZone-Taumarunui region. Small grey crosses show observational data. Small black circles show the six model-derived means over the 1986-2005 period. Small coloured circles show the individual model values extracted from the data files provided by MfE with different colours representing different SSP scenarios. The violin plots show indicative uncertainties derived using the ECS-based uncertainty analysis. The change values for future periods, extracted from the data files, are offset by the *mean* of the baseline model data. The original of this figure can be downloaded here.

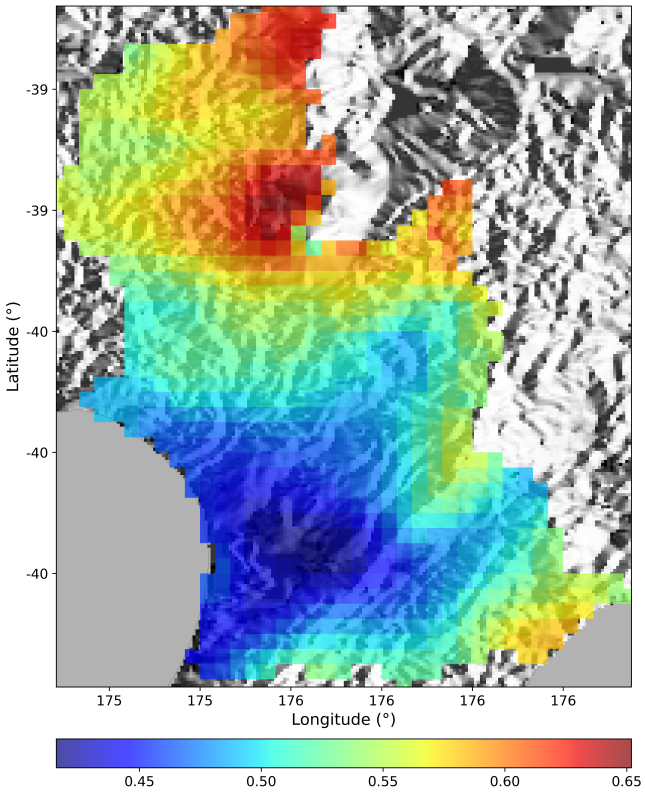


Figure 6: The R^2 values for the fits of GDD5 against the equilibrium climate sensitivity of the six CMIP6 models. In this case it is the change in GDD5 from 1986-2005 to 2041-2060 under the SSP3-7.0 scenario that is analysed. The original of this figure is available here.

Weather Forecasts (ECMWF) generates an hourly reanalysis over all land regions globally at 0.1° resolution (approximately 10 km resolution). One such field for 00:00 UTC on 1 January 2010 is shown in Figure 7. The

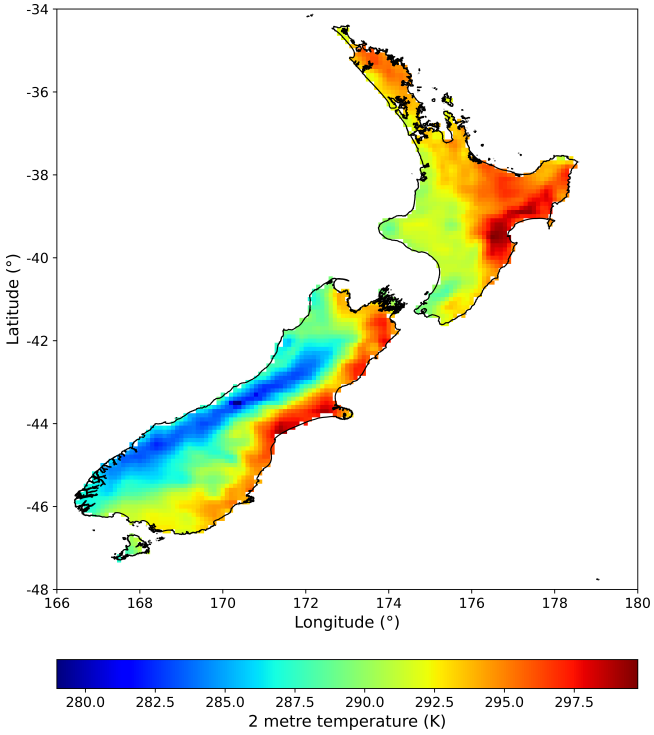


Figure 7: The surface (2 m) temperature field from the ERA5-land reanalysis data set for 00:00 UTC on 1 January 2010. This field is provided at 0.1° x 0.1° resolution.

data range of this field is 278.98K to 299.77K. To create a smoother field that is more amenable to interpolation to higher resolution, the first step is to adjust the field to equivalent sea-level temperature by warming each grid cell using a lapse rate of 6 K km^{-1} and the ERA5 elevation field (not shown). The resultant field is still at 0.1° by 0.1° resolution (Figure 8). The data range for this field is 286.70K to 303.09K. The sea-level temperature field is then spatially interpolated to 0.01° by 0.01° resolution. The expectation is that because the effects of topography would have been removed from this field, the spatial gradients are much smoother and, as a result, the interpolation more justifiable. The 0.01° by 0.01° resolution interpolated field is then transformed (cooled) back to the surface, again using a 6 K km^{-1} lapse rate, but now using a 0.01° by 0.01° resolution elevation field. The original ERA5-land field at 0.1° by 0.1° resolution and the field now interpolated to 0.01° by 0.01° resolution are shown using the same colour scale in Figure 9. The data shown in Figure 7 and in the left-hand panel of Figure 9 are the same - just the colour scale is different.

The relatively low resolution of the ERA5-land data set results in many grid cells around the coastline not having a value since they were classed as ocean cells in the ERA5-land data set. For all cells that don't have an ERA5-derived value, but whose altitude is above sea-level, the nearest value, obtained through a KD-tree search, is used to fill the missing value. The result is a filled data set over the entire land surface of New Zealand.

However, we cannot be sure that these fields are not biased as the ERA5 data assimilation process includes measurements from only a very few meteorological stations over the land surface of New Zealand. To correct any potential biases, we next consider *all* available measurements from meteorological stations around New Zealand that measure 2 metre temperatures. Differences between these temperatures and the ERA5 values at the nearest 0.01° by 0.01° grid cell (ERA5 minus station value) are shown in Figure 10.

The next step is to statistically model these differences and subtract the modelled difference field from the ERA5 field. This is done as two steps:

1. The difference field is modelled as linear functions of latitude and longitude. Because we are quite con-

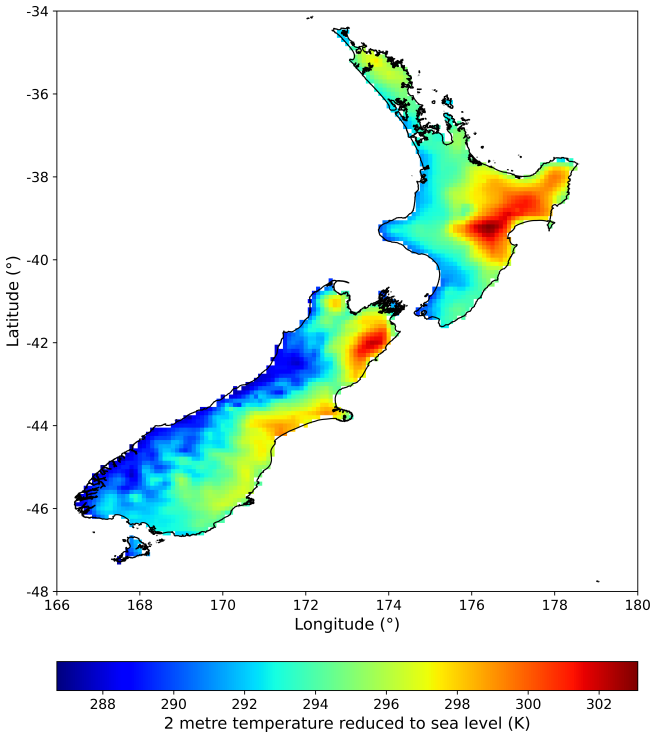


Figure 8: The surface (2 m) temperature field shown in Figure 7 reduced to sea level using a lapse rate of 6 K km^{-1}.

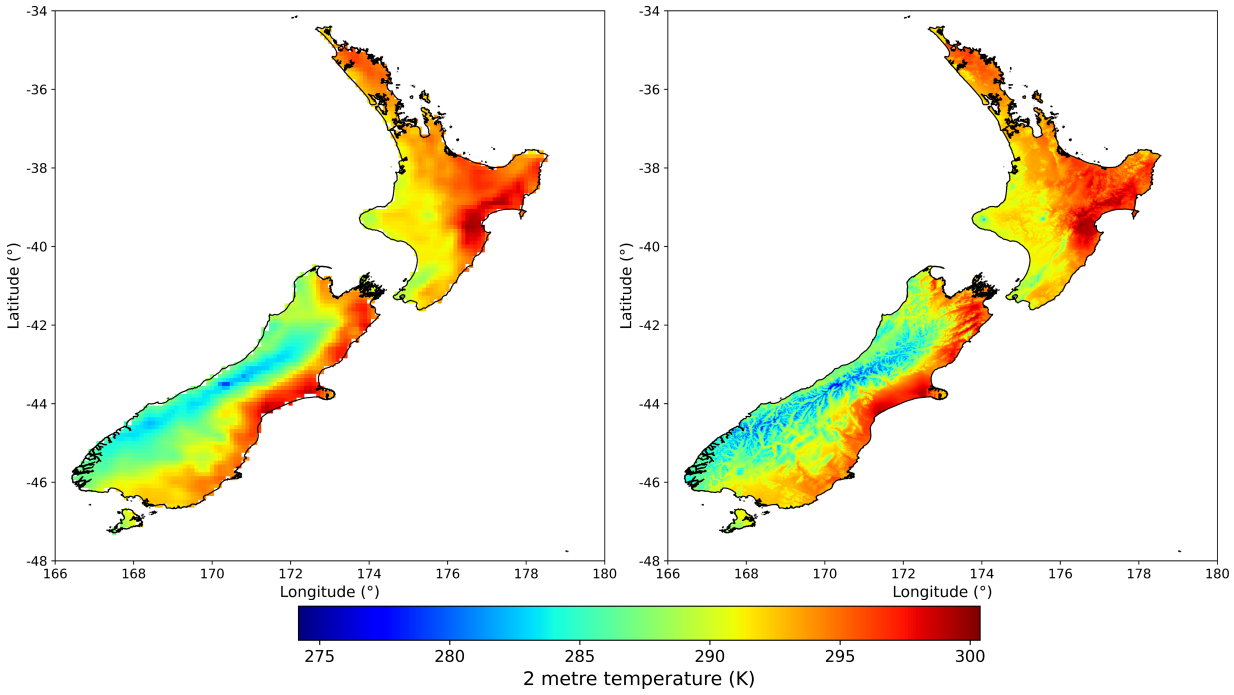


Figure 9: The original ERA5-land surface temperature field at 0.1° by 0.1° resolution (left). The 0.01° by 0.01° resolution field generated through spatial interpolation on the sea level temperature field (right).

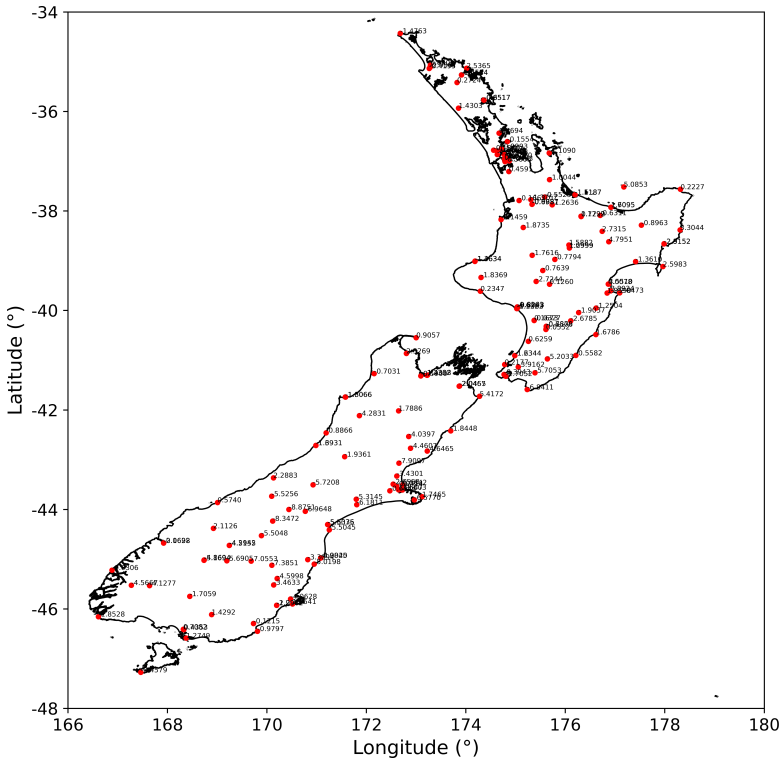


Figure 10: Differences between the 0.01° by 0.01° resolution ERA5-derived temperature field and the measurements from the nearest meteorological stations (ERA5 minus station value).

fident that we can capture a SW-NE gradient, but less confident that we can capture a SE-NW gradient, values of 0.0 are added to the difference field at 34°S, 166°E and 48°S, 180°E (i.e., the top left and bottom right corners of the domain) to anchor the fit. More complicated models can be applied, e.g., where, in addition to linear functions of latitude and longitude, a $ln(z)$ term (where z is the altitude) can be included. However, this can result in spurious model values when the altitude dependence cannot be robustly determined. The goal in this first step is simply to model the most broad-scale morphology of the difference field so that the Kriging (see next step) doesn't need to do all of the work. The difference field modelled in this way is shown in Figure 11.

2. The residuals from that first fit are then interpolated to the 0.01° by 0.01° grid using Kriging interpolation. Those residuals, and their Kriging interpolated field, are shown in Figure 12.

The fields shown in Figure 11 and Figure 12 are then subtracted from the field shown in the right-hand panel of Figure 9 to create the final gauge-corrected 2 metre temperature field which is shown in Figure 13. In addition to the effects of the topography on surface temperature being far more apparent, the temperatures, e.g., around Ashburton are much higher in the high-resolution field than in the original ERA5 field. Hourly temperature measurements at three stations around Ashburton were 307.2K, 306.7K, and 305.7K respectively which agrees very well with the temperatures shown in the high-resolution ERA5 field indicating that the bias correction has worked well. It should be noted, however, that data gaps in the national climate database mean that the differences for some times cannot be calculated. Where insufficient differences for a time interval can be computed (currently 10), differences from the preceding and following hour are also used. But this may still result in periods for which there are insufficient differences to calculate a gauge-corrected field in which case this field will be missing from the data set. These gauge-corrected temperature fields are then spatially sub-sampled to the Horizons Region and are used to calculate the various temperature-specific climate variables needed to validate the CCAM baseline (climatological) values..

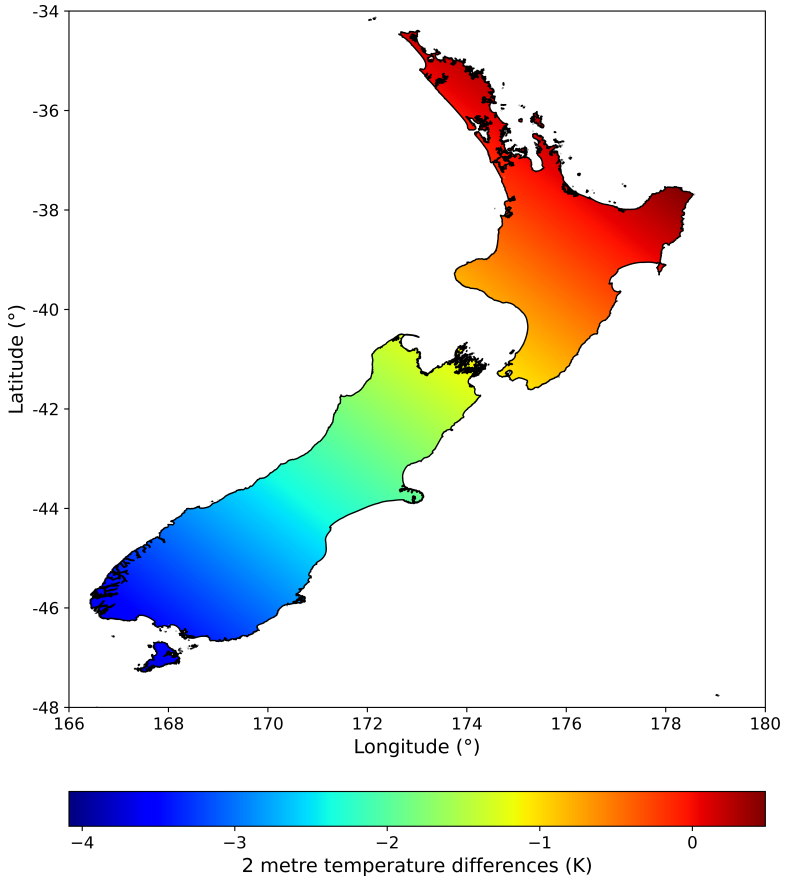


Figure 11: The difference field modelled as a function of latitude and longitude.

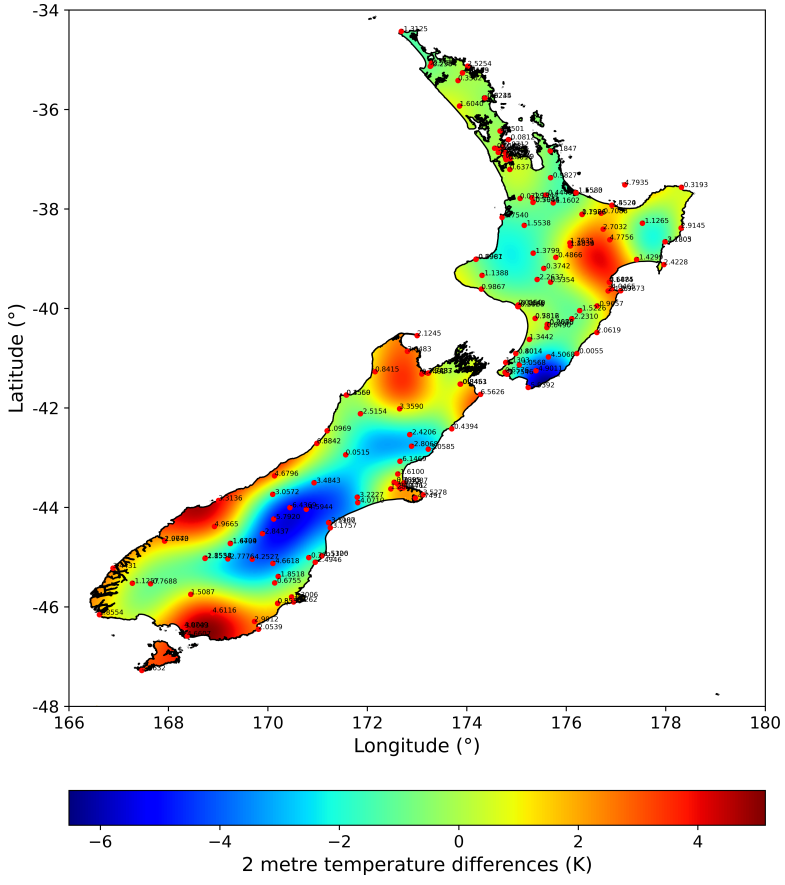


Figure 12: The residuals resulting from the initial modelling of the difference field and their Kriging interpolated field.

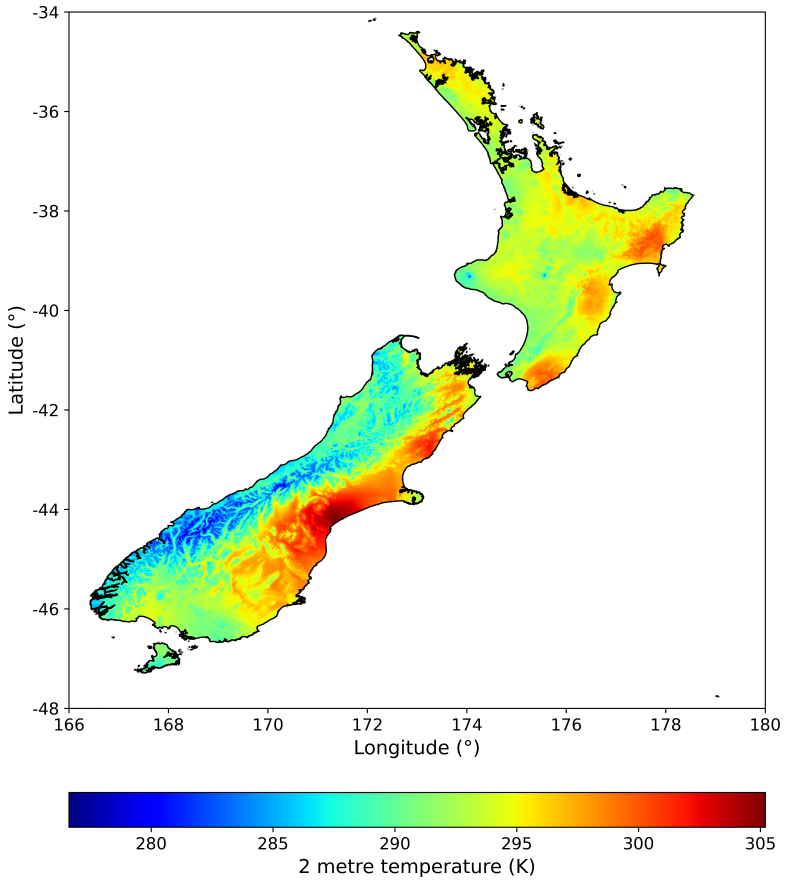


Figure 13: The final gauge-corrected 2 metre temperature field at 0.01° by 0.01° resolution.

3.2 Precipitation

Because data from only a few sites from around New Zealand provided observations to ECMWF for assimilation into their reanalysis product, the ERA5 precipitation data fields, provided at 0.1° by 0.1° resolution, tend to significantly underestimate the precipitation depth. Cognisant of that shortcoming, a gauge-corrected hourly precipitation data set was created as part of a piece of work that Bodeker Scientific led under contract to the New Zealand Dam Safety Hydrology Group (DSHG). That project is described here. The data set was created by Mike Kittridge as part of that project; Mike has kindly permitted that data set to be used here.

To create the data set, a log Pearson Type III distribution was fitted to 14-day precipitation totals from both individual cells in the ERA5 data set and gauges with records longer than 50 years to derive an associated probability distribution function (PDF). There are 1,200 gauges with records of 25 years or more, and 255 gauges with records of 70 years or more. The log Pearson Type III distribution was also applied to gauges with shorter than a 50-year record using the shape parameters from the nearest gauge with a greater than 50-year record. PDFs were then interpolated between gauge locations. ERA5 values in each cell, at each time step, were adjusted using the difference between the ERA5 PDF and the PDF derived from gauge records.

An example of the effect of the gauge correction on the ERA5 precipitation fields is shown in Figure 14. The gauge-correction of the ERA5 precipitation fields results in significant improvements in the accuracy of the fields. These gauge-corrected hourly precipitation fields, from 1950 to 2020, were then used to calculate precipitation-based statistics for comparison against the CCAM baseline values (noting the limitations of the observational data set being at 0.1° by 0.1° resolution whereas the CCAM fields are at 0.05° by 0.05° resolution).

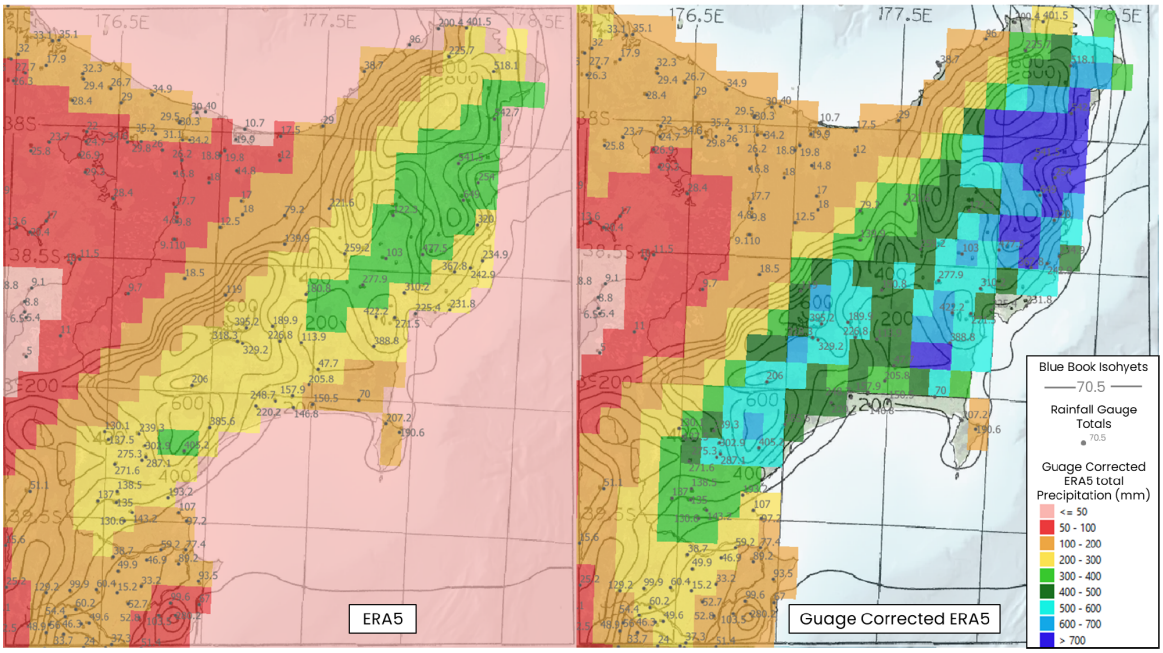


Figure 14: Comparison of total precipitation for Cyclone Bola over the 3-day period from 6 to 8 March 1988 for the East Cape: (left) the coloured cells show the raw ERA5 total precipitation field and (right) the coloured cells show the gauge-corrected ERA5 total precipitation field. Also shown in both panels is the total precipitation map historically used to estimate possible maximum precipitation in New Zealand and the gauge locations with their associated 3-day total precipitation values extracted from the National Climate Database (where available).