A U-Net Based Approach for Inferring the Structure of Extreme Weather Events from Coarse Model Output

Leroy Bird, Jordis Tradowsky, Greg Bodeker, William Herewini

About EWERAM: The Extreme Weather Event Real-time Attribution Machine

In the immediate aftermath of an extreme weather event (EWE), the public wants to find out to what extent anthropogenic climate change has contributed to the **frequency** and **severity** of the specific event.

To understand how the frequency of an event class has changed, we rely on thousands of years of coarse resolution climate model simulations, of both pre-industrial along with current day simulations.

Currently our best data set for frequency attribution is weather@home which is run over New Zealand and Australian at 50 km resolution. For a particular year, thousands of years worth of simulations are run using conditions from the given year and under pre-industrial conditions.

A straightforward approach for extreme temperature attribution

If a temperature record, typically occurring on a large scale, was broken in Dunedin we can find the grid cell of Dunedin and count the number of times the given threshold was exceeded in anthropogenic and pre-industrial weather@home simulations. While there may be a bias between weather@home and reality, this bias should consistently be in the pre-industrial and the anthropogenic simulations.

Problems with precipitation attribution

Extreme precipitation events are often very localized and can occur over a short period of time. This problem is exacerbated by New Zealand's rough topography which is not resolved in coarse resolution models. Looking at coarse resolution precipitation fields from models does not tell us if an extreme precipitation event occurred within a localized region of the grid cell, or if it was mild precipitation over a large area, nor do we know where this event occurred, i.e. was it over the town or over the mountains 30 kms away?

If we want to know how the frequency of an extreme precipitation event has changed, we first need to find all synoptically similar cases by defining classes of events, see Fig.1. In a next step, we can then evaluate the probability of extreme rain occuring on the sub-grid scale.

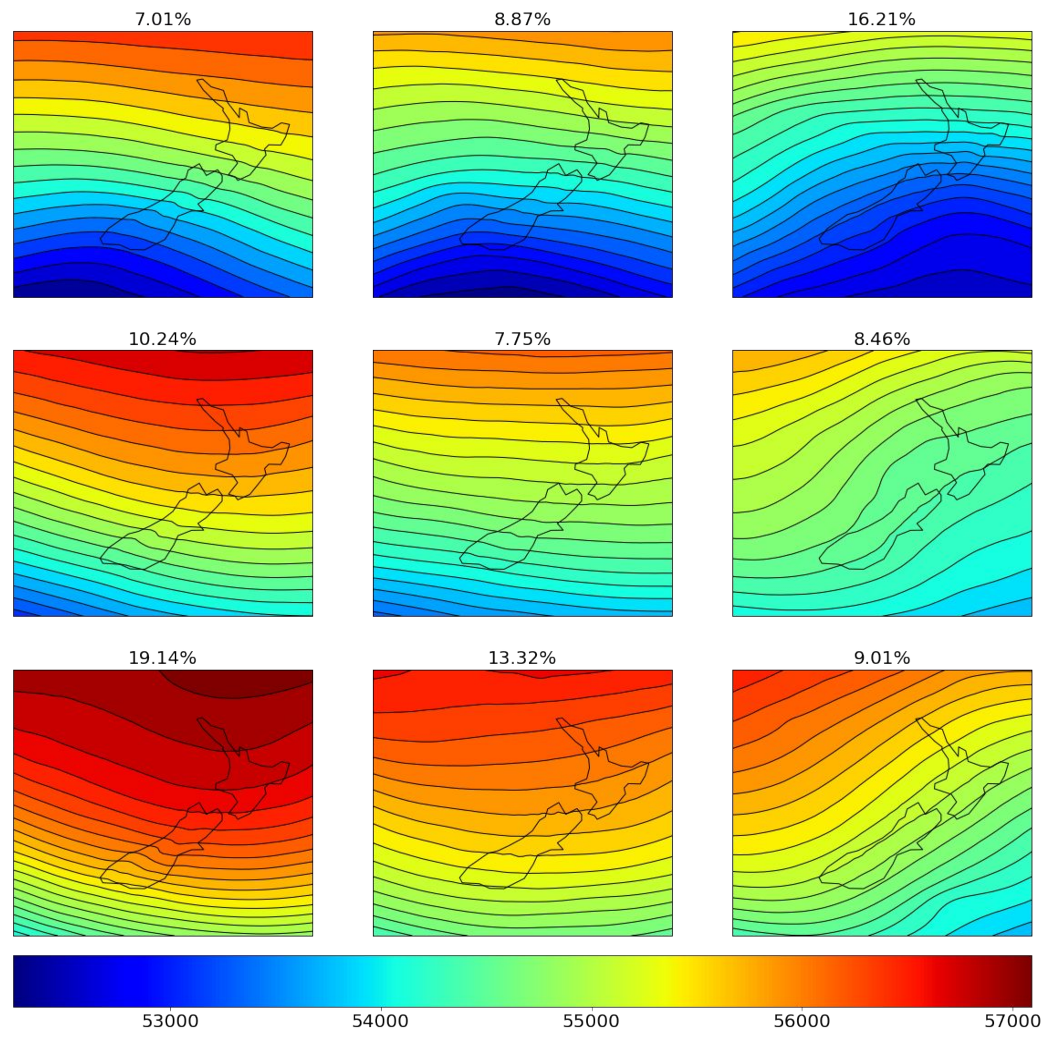

Figure 1: A self organizing map of geopotential at 500 hPa from ERA5. An important tool for grouping synoptic patterns.

An introduction to U-Nets, A machine learning approach

U-Nets are a convolutional neural network based deep learning architecture which are traditionally used for image to image problems, for example image segmentation. It consists of a down scaling branch where higher level but large scale features can be extracted, then an upscaling branch where these features are scaled up to match the target resolution. However, the network contains cross connections which remove the information bottleneck of the downscaling.

The training data we used as a proof of concept was ERA5 (30 km grid), as the coarse resolution input to the U-Net, and ERA5-Land (9 km grid) as the high resolution labels. ERA5-Land is dynamically consistent with ERA5, so together they make great training pairs. The downscaling side of the U-Net was based off a ResNet-34 backbone. The loss function is simply the mean absolute deviation of the U-Nets prediction and ERA5-Land.

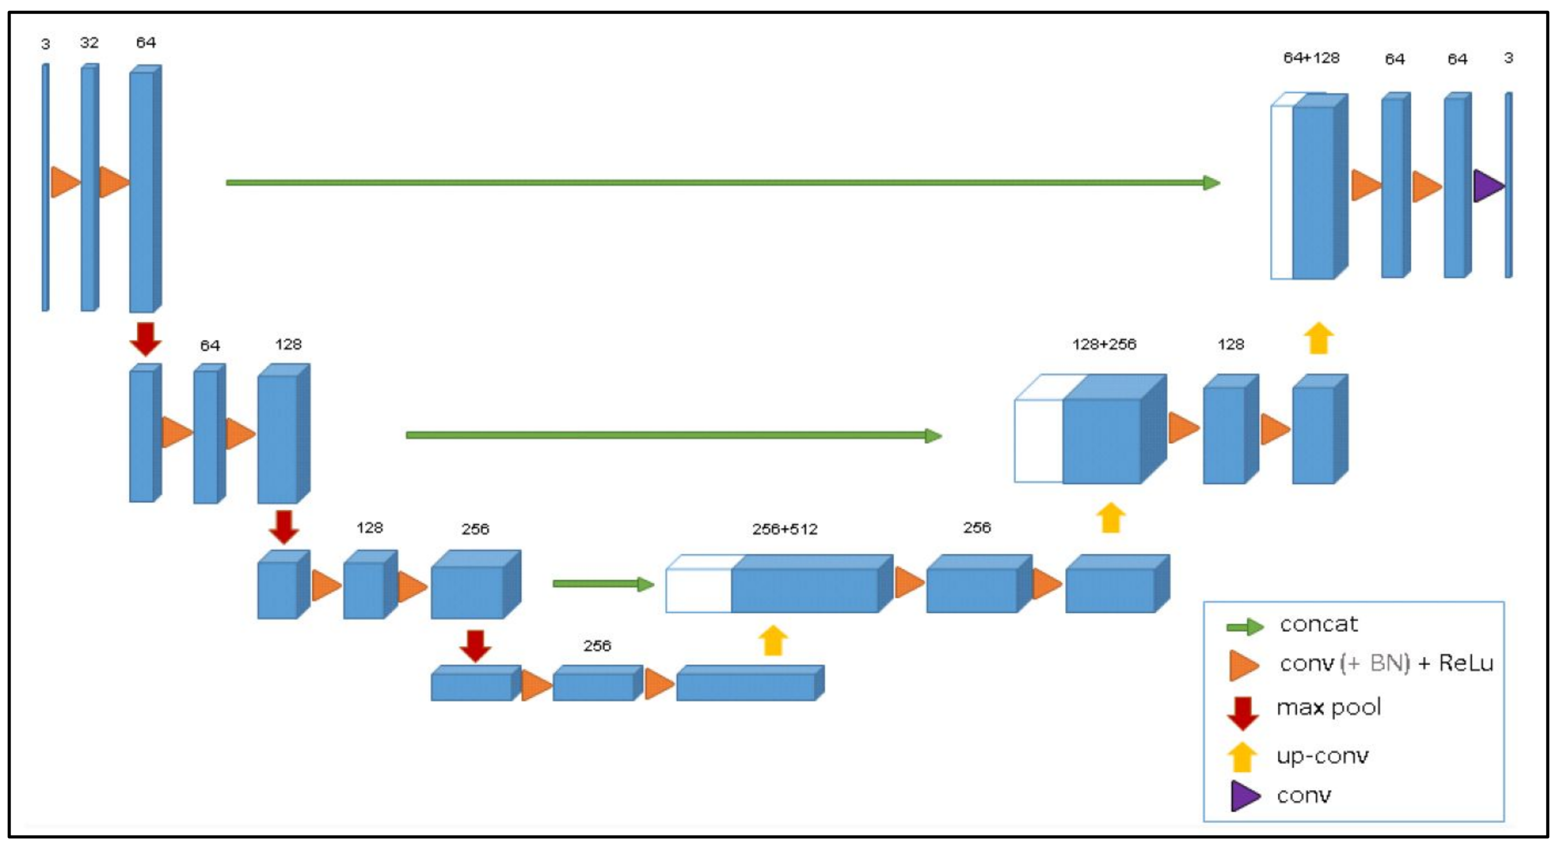

Figure 2: A simplified U-Net-based architecture for illustrative purposes only.

Increasing both temporal and spatial resolution

Unfortunately, increasing the spatial resolution of precipitation maps alone does not account for errors associated with the model, nor does it accurately represent the subgrid scale dynamics. Additionally, many extreme precipitation events occur over just a few hours, so we need some way of generating a more accurate temporal picture of these events.

Instead of feeding precipitation maps into the U-Net we feed in temperature and wind on three pressure levels along with a other 2D variables at t_0, like total column water. The labels are the precipitation from ERA5-Land over each of the next 3 hours, t_1, t_2 and t_3. So the U-Net is providing both temporal and spatial downscaling.

Figure 3

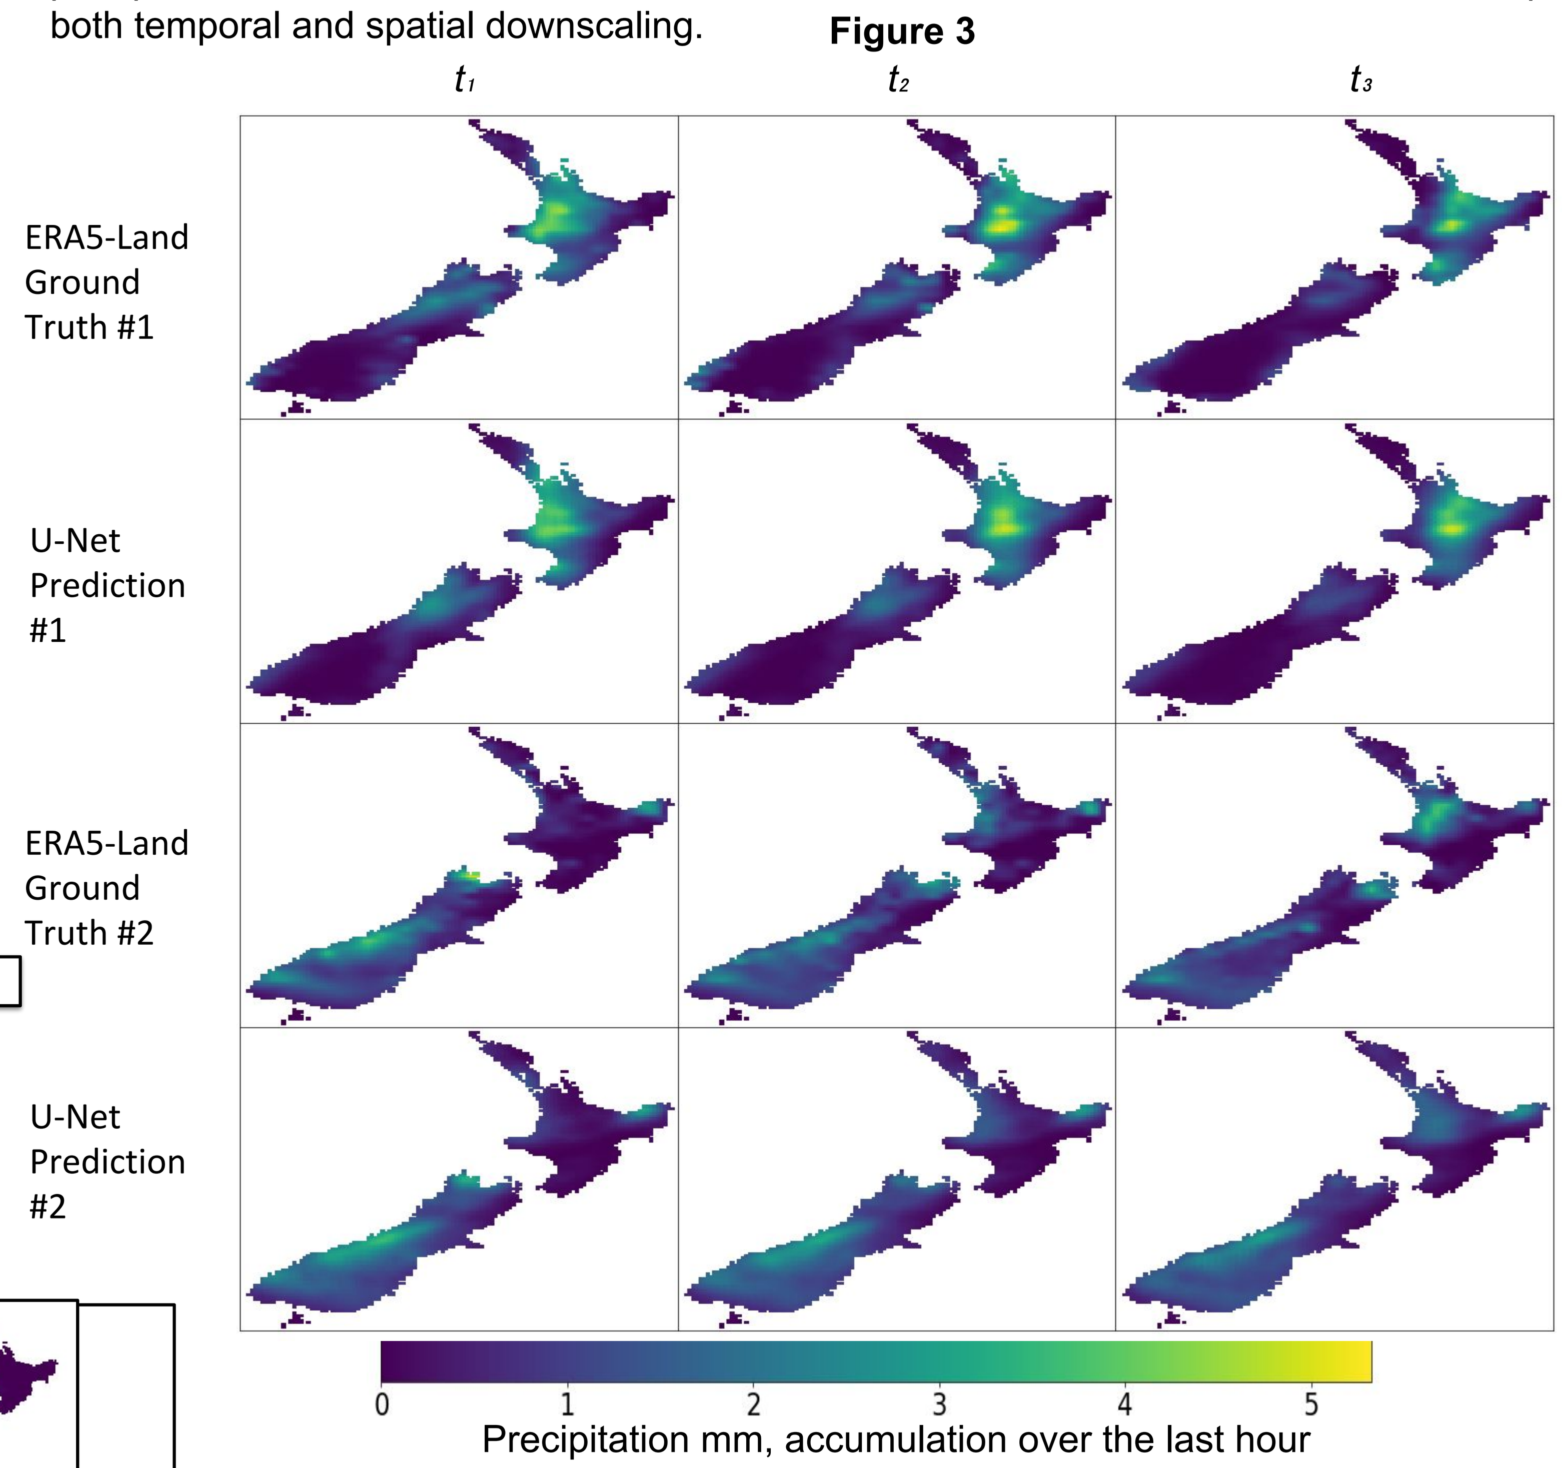

Adding Stochasticity to the U-Net and Conditional Generative Adversarial Networks

While providing both spatial and temporal downscaling gets us closer to understanding the event, it doesn't provide the full distribution of possible outcomes. Notice on Figure 3, the case where the U-net predicts smooth rain, while in reality there were multiple localized convective cells producing more intense rain. Ideally, we want to produce the full distribution of possible events given the coarse resolution model data.

In order to achieve this, we add some normally distributed noise into the upscaling layers of the U-Net. The loss function is then modified to encourage variability.

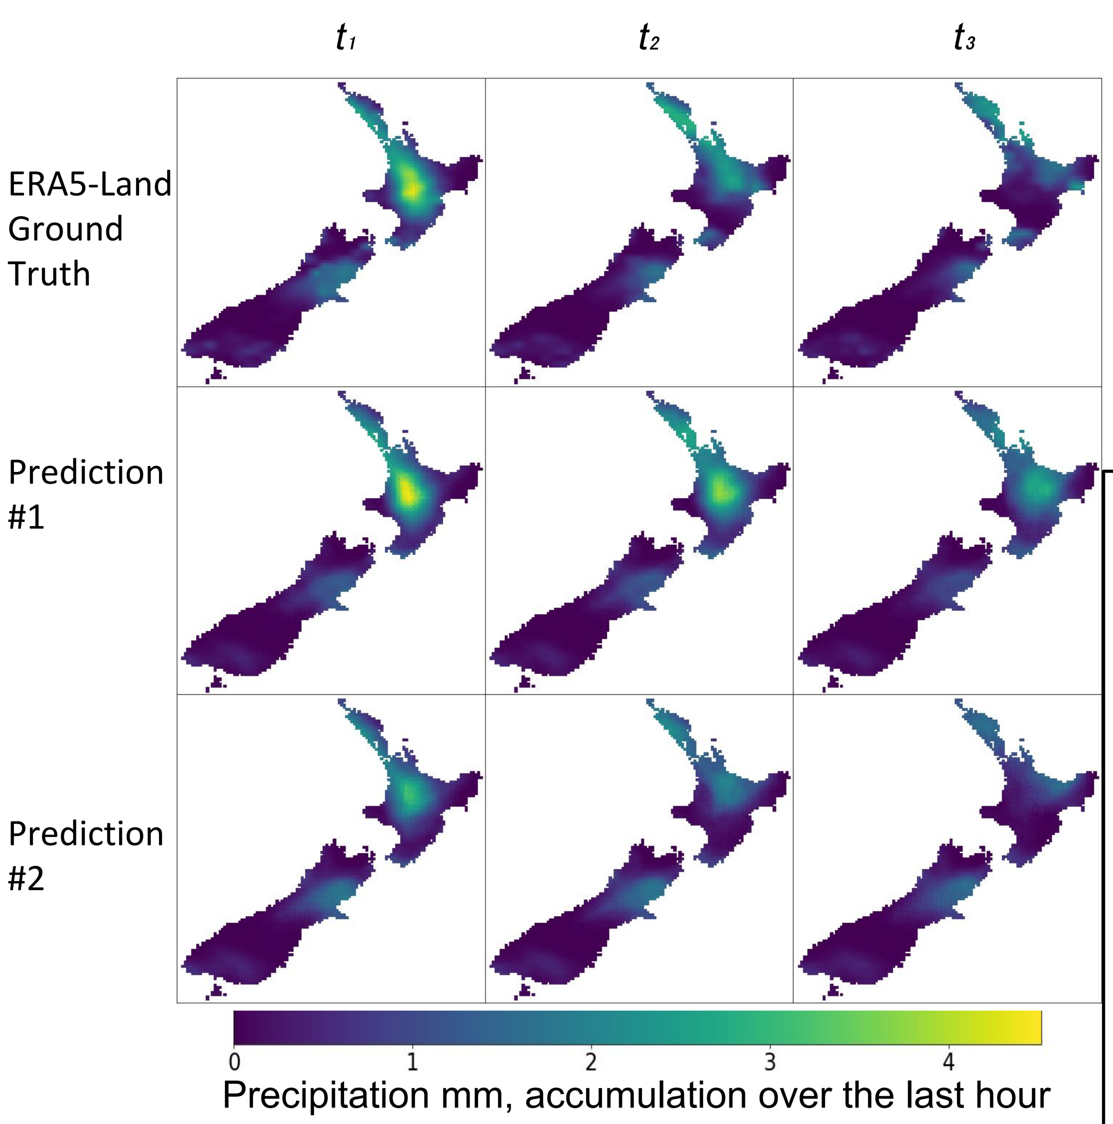

Figure 4: Row One, the ground truth labels over the three hours.

Row two and row three, are two passes of the U-Net, with the same input state, but different random noise applied within the U-Net.

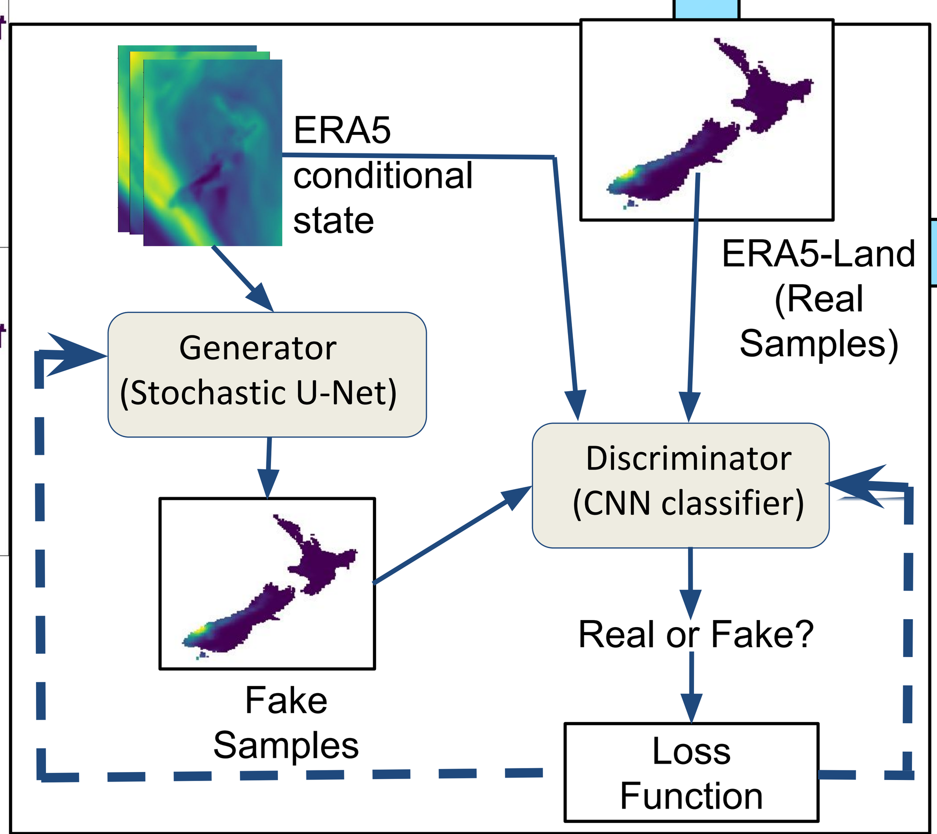

Figure 5: A overview of a conditional generative adversarial network.

Despite the clear variation in the predictions, we are not really getting a realistic series of the events because we are not punishing the U-Net for dynamical consistency and realism. The next step would be to add a conditional generative adversarial network which decides if a given sample is drawn from the same distribution as the real samples.

Putting it all together and next steps

- After finding all events in the synoptic class corresponding to the extreme event under investigation, we can pass the low resolution model output into the U-Net generator to produce many possible high resolution samples.
- Then we can match the generated samples to the high resolution precipitation maps of the real event.
- For every low resolution map of the given class we can now derive the probability that it leads to an extreme event and therefore derive the relative frequency of these events in pre-industrial and anthropogenic simulations.
- However, first we will investigate if the U-Net can accurately approximate the true distribution of events, so we don't bias the results.