

---

# Combining deep generative models with extreme value theory for synthetic hazard simulation: a multivariate and spatially coherent approach

---

**Alison Peard and Jim W. Hall**  
School of Geography and the Environment  
University of Oxford  
Oxford, UK  
alison.peard@ouce.ox.ac.uk

## Abstract

Climate hazards can cause major disasters when they occur simultaneously as compound hazards. To understand the distribution of climate risk and inform adaptation policies, scientists need to simulate a large number of physically realistic and spatially coherent events. Current methods are limited by computational constraints and the probabilistic spatial distribution of compound events is not given sufficient attention. The bottleneck in current approaches lies in modelling the dependence structure between variables, as inference on parametric models suffers from the curse of dimensionality. Generative adversarial networks (GANs) are well-suited to such a problem due to their ability to implicitly learn the distribution of data in high-dimensional settings. We employ a GAN to model the dependence structure for daily maximum wind speed, significant wave height, and total precipitation over the Bay of Bengal, combining this with traditional extreme value theory for controlled extrapolation of the tails. Once trained, the model can be used to efficiently generate thousands of realistic compound hazard events, which can inform climate risk assessments for climate adaptation and disaster preparedness. The method developed is flexible and transferable to other multivariate and spatial climate datasets.

## 1 Introduction

Compound events, defined as combinations of climate variables, not all of which are necessarily extreme, but which lead to extreme impacts [1], can wreak far greater damages than those composed of a single driver [2]. Most major recent climate catastrophes, such as droughts caused by low precipitation and high temperatures; and coastal flooding caused by co-occurring high tides, heavy precipitation, and strong onshore winds, were compound events. Univariate risk assessments have been shown to significantly underestimate true hazard levels [2]. Additionally, hazard maps which assume spatially-homogeneous return periods are physically unrealistic and lead to inaccurate risk assessments [3]. For successful climate adaptation policies to be developed it is paramount that scientists develop accurate and efficient methods to model the spatial distributions of compound hazard events.

There are two classical approaches for modelling the distribution of multivariate extreme events: copulas [4, 5] and the multivariate conditional exceedance model [6]. Both approaches decompose the problem into modelling the univariate marginals and the dependence structure between them. Here, marginal refers to the probability distribution of a single hazard at a point in space. Copulas are based on Sklar's theorem [7] and accurate modelling relies on choosing the most appropriate copula family. The multivariate conditional exceedance model replaces the copula with a semiparametric

conditional model that learns the conditional distribution of all the marginals, given that one of them is extreme, and introduces flexibility into the asymptotic dependence structure between pairs of variables. Both methods rely on parameter estimation, hence inference becomes computationally inefficient as dimensionality increases.

Machine learning is increasingly being used to address the shortcomings of traditional approaches in multivariate extreme modelling. Letizia and Tonello [8] developed a fully nonparametric model using two separate GANs to capture both the marginal distributions and the dependence structures. Bhatia, Jain, and Hooi [9] developed ExGAN, which generates samples of desired extremeness levels for rainfall data over the United States using a conditional GAN. Boulaguiem et al. [10] combined extreme value theory with deep convolutional GANs (DCGANs) to generate annual maximum precipitation and temperature fields over Europe, training a separate, single-channel model for each of the two variables.

We develop an extension on the work of Boulaguiem et al. [10] to the multi-dimensional case by simultaneously training a DCGAN on three channels of gridded climate data: wind speed, significant wave height, and total precipitation. By training on daily maxima rather than annual maxima, we capture the propensity for extremes to co-occur on a given day. We present preliminary results and demonstrate that the model generates physically realistic samples that capture the dependence structure across space and between variables. Such a model can be used to efficiently generate a large synthetic ensemble of multivariate events, addressing the urgent need for better modelling of compound events in climate risk assessment [2].

## 2 Methodology

**Data** Hourly wind speeds [ $\text{ms}^{-1}$ ], significant height of combined wind waves and swell [m], and total precipitation [m] for the years 2013-2022 are obtained from the *ERA5 hourly data on single levels from 1940 to present* reanalysis product [11]. The data has  $0.25^\circ$  spatial resolution. We obtain the data over the Bay of Bengal (longitude:  $10\text{-}25^\circ$  East, latitude:  $80\text{-}95^\circ$  North) and calculate the daily maxima. We resize the images to  $18 \times 22$  pixels using bilinear interpolation and zero-pad them to  $20 \times 24$  pixels.

**Normalisation** It is conventional to normalise data before training a neural network, to stabilise gradient descent and improve convergence. Boulaguiem et al. [10] propose to replace standard re-scaling methods with the probability integral transform [12, 13], where each observation  $X^j \in \mathbb{R}$  is replaced by its empirical distribution function  $\hat{F}_X(X^j)$ , which follows a uniform distribution over the interval  $[0,1]$ . This effectively separates the dependence structure from the marginal distributions. The neural network is then trained on the uniform-transformed data. For the inverse transform, Boulaguiem et al. [10] fit a generalised extreme value (GEV) to each sequence of annual maxima in its original scale. The GEV parameters are shape  $\xi$ , location  $\mu$ , and scale  $\sigma$ .<sup>1</sup> The shape parameter determines the shape of the tails, which are bounded below for  $\xi > 0$  and bounded above for  $\xi < 0$  [16]. The corresponding probability point function is then used to transform generated uniform samples back to the original scale.

Figure 1 shows the fitted parameters for wind speed over the Bay of Bengal (corresponding plots for significant wave height and precipitation are shown in Appendix B). The fitted parameters indicate that wind is stronger and has higher variance offshore, but is also more likely to have a negative shape parameter  $\xi$ , indicating that it is bounded above. Describing offshore winds using a Weibull distribution is common practice, though more complex models have been proposed [17]. Significant wave height shows a relatively homogeneous distribution, with primarily positive  $\xi$  indicating light-tailed distributions with no upper bound. Precipitation appears heavier and more varied over deep water and along the coastlines. The distribution of the shape parameter for precipitation appears less physical than other results. This may be due to the large number of zeros in the data biasing the fit.

<sup>1</sup>This fitting relies on assumptions of independence and stationarity that underlie the Fisher-Tippett-Gnedenko theorem [14, 15, 16] and the climate variables in this paper exhibit some autocorrelation and seasonality, as is typical for climate data. However, it has been shown that the dependence assumption can be relaxed in cases where the long-range dependence at extreme levels is weak, as the data will still fall into the same family of GEV distributions ([16] Chapter 5). Following a similar argument to Heffernan and Tawn [6], we do not attempt to correct the seasonal component of the data as we are interested in the overall joint distribution of daily maxima, and not their temporal evolution.

For the purposes of this paper, we leave further exploration of suitable marginal distributions as an opportunity for future work.

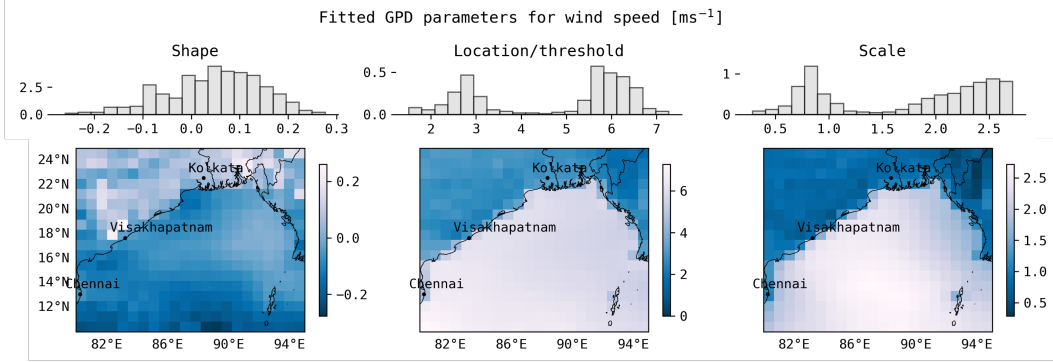


Figure 1: Fitted generalised extreme value (GEV) distribution parameters shape ( $\xi$ ), location ( $\mu$ ), and scale ( $\sigma$ ), over the Bay of Bengal for wind speed in  $\text{ms}^{-1}$ .

**Deep convolutional GAN** GANs [18], are generative deep learning models composed of two competing neural networks which can implicitly learn the distribution of the training data [19] and capture local and global spatial patterns. The generator

$$G : \mathbb{R}^{100 \times M} \rightarrow \mathbb{R}^{20 \times 24 \times M},$$

with  $M \in \mathbb{N}$  representing the number of samples, generates  $M$  fake images of size  $20 \times 24$  from a latent space  $z \in \mathbb{R}^{100 \times M}$  with elements  $z_i^j \sim \text{Norm}(0, 1)$ , and the discriminator

$$D : \mathbb{R}^{20 \times 24 \times M} \rightarrow [0, 1]^M,$$

attempts to distinguish fake data  $G(\mathbf{Z})$  from real data  $\mathbf{U}$ . Elements of both  $G(\mathbf{Z})$  and  $\mathbf{U}$  have uniform distributions with support  $[0, 1]$ . The discriminator seeks to maximise the number of instances of  $D(\mathbf{U}) = 1$  and minimise the instances of  $D(G(\mathbf{Z})) = 1$ . Likewise, the generator seeks to maximise the occurrences of  $D(G(\mathbf{Z})) = 1$ . This setup is captured in the cross-entropy objective function,

$$\min_G \max_D \mathbb{E}_{\mathbf{U}} [\log D(\mathbf{U})] - \mathbb{E}_{G(\mathbf{Z})} [\log D(G(\mathbf{Z}))]. \quad (1)$$

Deep convolutional GANs (DCGANs) [20] employ convolutional and deconvolution layers, which can learn local patterns and features of image data in the generator and the discriminator, respectively. We train a DCGAN on the data transformed to uniform marginals, where each sample is an image consisting of three channels corresponding to wind speed, significant wave height, and total precipitation. We outline the architecture and training procedure in more detail in Appendix A.

### 3 Results and discussion

We present results for data generated by a DCGAN model trained for 1000 epochs over a training set of size  $N = 1000$  with a batch size of 50, then transformed to the original scale using the fitted GEV distributions. Figure 2 shows a sample from the training (left block) and generated (right block) data, where each group of three images represents daily maximum fields for wind speed (left), significant wave height (centre), and total precipitation (right). (Larger figures are shown in Appendix B.) Visually, the generated images appear similar to the training images, with good variation between samples. Wind and waves display higher values in similar locations on the same day in both sets. Significant wave height accounts for wind-sea waves, which are driven by local winds [11], and it's clear the GAN has learned this relationship.

To quantify the extremal dependency between variables, we use the extremal coefficient  $\theta$  [5, 21, 22], which takes values in the range  $[1, D]$  for a  $D$ -dimensional sample, indicating total dependence or independence between variables, respectively. In a three-dimensional setting with  $N$  samples, this can be estimated using,

$$\hat{\theta}_{123} = \frac{N}{\sum_{n=1}^N \min \left( \frac{1}{Y_{n1}}, \frac{1}{Y_{n2}}, \frac{1}{Y_{n3}} \right)}$$




Figure 2: Heatmaps for the training (left) and generated (right) data. Daily maximum wind speed, significant wave height, and precipitation are shown side-by-side for each sample, where darker colours indicate higher values in each case. All heatmaps of the same variable have the same colour scale. (Larger figures for both are shown in Appendix B.)

where the  $Y_{ni}$  have unit Fréchet distributions, which are easily obtained from the uniformly distributed marginals using  $Y_{ni} = -\log(U_{ni})^{-1}$ . The extremal coefficient can be understood as the effective number of independent variables in the sample, while the extremal correlation  $\hat{\chi} = 1 - \hat{\theta}$  forms an extreme analogue for the Pearson correlation coefficient. Comprehensive discussions and derivations of these coefficients are given in Davison, Padoan, and Ribatet [5] and Smith [21]. Figure 3 compares the distribution of the estimator  $\hat{\theta}$  between the train, test, and generated sets across [A] space and [B] the wind, wave, and precipitation dimensions. Additional figures in Appendix B show the model has learned the distribution of the spatial correlations and the spatial extremal correlations.

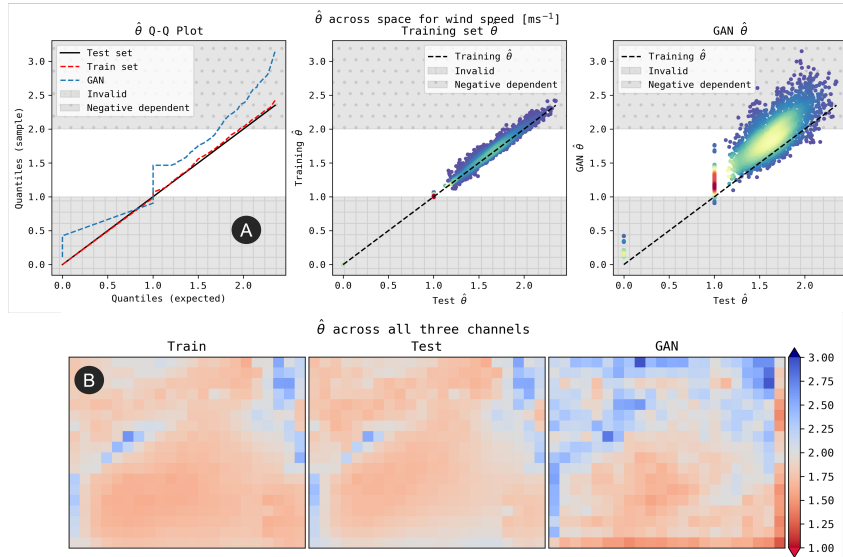


Figure 3: **A:** Distributions of  $\hat{\theta}$  across space for wind data for all 4950 possible pixel pairs: Q-Q plots comparing test, train, and generated sets (left), and scatter plots comparing the train and generated set to the test set (centre and right). **B:** Heatmaps showing three dimensional extremal coefficient estimates  $\hat{\theta}$  for the train, test, and generated data. The GAN-generated data shows more noise, and a bias towards underestimation, particularly near land, as seen by the higher  $\hat{\theta}$  in all plots, but has learned the general spatial structure of the multivariate extremal dependencies.

Overall, the model has learned the extremal dependence structure across space and climate variables but exhibits bias towards larger  $\hat{\theta}$ , i.e., underestimating the dependence between the most extreme events. Analysis of the generated data indicates the DCGAN is generating fewer samples near the highest percentiles for each marginal. GANs generally produce samples from the mean of the training distribution so this is not wholly unexpected [9, 23] and may be improved by replacing the standard normal latent space vectors with a more heavy-tailed distribution [23] or by transfer learning on extreme subsets of the training data.

## 4 Conclusion

In this paper, we have demonstrated the usefulness of deep generative models for learning dependence structures in high dimensional climate data, and how this can be combined with extreme value theory to facilitate interpretability and to control extrapolation to extreme values for spatial, compound events. We have applied this methodology to a dataset of daily maximum wind speed, significant wave height, and precipitation from 2013 to 2022 over the Bay of Bengal and developed a model that can efficiently generate a large number of synthetic compound events.

The model learns the dependence structure well both across space and between different variables, though it displays more overall noise and has a tendency to underestimate the most extreme events. This underestimation is likely due to GAN’s tendency to reproduce samples from the centre of a distribution and could potentially be improved by using heavier-tailed distributions, such as the Student’s  $t$ -distribution, for the latent space variables. To avoid potential side-effects, we stress that when using this model as part of a larger project, care should be taken to ensure the input data is reliable, or that its limitations are explicitly known, as the model will propagate any errors in the input data into generated data.

There are ample opportunities for future development of this work, including a deeper exploration of distributions to be fitted to the univariate marginals; the use of alternative deep generative models such as diffusion models, which have more stable training, and normalising flow models, which explicitly learn the data distribution; or by making improvements to the current model through architecture modifications, further parameter tuning, and transfer learning on extreme subsets of the training data.

The methods in this paper construct a computationally efficient model for generating synthetic datasets of compound events that are coherent in space and time. Such datasets can be used as inputs for further risk modelling and are critically important for understanding the true spatial distribution of the risk of compound extreme events and informing climate adaptation and coastal disaster preparedness policies.

## Acknowledgments and Disclosure of Funding

The authors would like to thank Geoff Nicholls and Steven Reece for their valuable feedback on this work. This work was supported by the Engineering and Physical Sciences Research Council (EPSRC).

## References

- [1] Leonard, M. et al. “A compound event framework for understanding extreme impacts”. In: *Wiley Interdisciplinary Reviews: Climate Change* 5.1 (2014), pp. 113–128.
- [2] Zscheischler, J. et al. “Future climate risk from compound events”. In: *Nature Climate Change* 8.6 (2018), pp. 469–477.
- [3] Vorogushyn, S. et al. “Evolutionary leap in large-scale flood risk assessment needed”. In: *Wiley Interdisciplinary Reviews: Water* 5.2 (2018), e1266.
- [4] Nelson, R. *An Introduction to Copulas*. Springer, New York, 2006.
- [5] Davison, A. C., Padoan, S. A., and Ribatet, M. “Statistical Modeling of Spatial Extremes”. In: *Statistical Science* 27.2 (2012), pp. 161–186. DOI: 10.1214/11-STS376. URL: <https://doi.org/10.1214/11-STS376>.
- [6] Heffernan, J. E. and Tawn, J. A. “A conditional approach for multivariate extreme values (with discussion)”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 66.3 (2004), pp. 497–546.
- [7] Sklar, A. “Fonctions de répartition à n dimensions et leurs marges”. In: *Publications de l’Institut de statistique de l’Université de Paris* 8.1 (1959), pp. 229–231.
- [8] Letizia, N. A. and Tonello, A. M. “Segmented generative networks: Data generation in the uniform probability space”. In: *IEEE Transactions on Neural Networks and Learning Systems* 33.3 (2020), pp. 1338–1347.
- [9] Bhatia, S., Jain, A., and Hooi, B. “ExGAN: Adversarial generation of extreme samples”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 8. 2021, pp. 6750–6758.
- [10] Boulaguiem, Y. et al. “Modeling and simulating spatial extremes by combining extreme value theory with generative adversarial networks”. In: *Environmental Data Science* 1 (2022), e5.
- [11] Hersbach, H. et al. *ERA5 hourly data on single levels from 1940 to present*. (Accessed on 10-09-2023). May 2023. DOI: 10.24381/cds.adbb2d47. URL: <https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels?tab=overview>.
- [12] Fisher, R. A. *Statistical Methods for Research Workers*. Oliver and Boyd, 1932.
- [13] Casella, G. and Berger, R. L. *Statistical inference*. 2nd. Probability Integral Transform: Theorem 2.1.10, p.54. Duxbury Press, 2001.
- [14] Fisher, R. and Tippett, L. “Limiting forms of the frequency distribution of the largest or smallest member of a sample”. In: *Proceedings of the Cambridge Philosophical Society* 24.02 (1928), pp. 180–190.
- [15] Gnedenko, B. “Sur la distribution limite du terme maximum d’une serie aleatoire”. In: *Annals of Mathematics* (1943), pp. 423–453.
- [16] Coles, S. et al. *An Introduction to Statistical Modeling of Extreme Values*. Vol. 208. Springer, 2001.
- [17] Morgan, E. C. et al. “Probability distributions for offshore wind speeds”. In: *Energy Conversion and Management* 52.1 (2011), pp. 15–26.
- [18] Goodfellow, I. et al. “Generative adversarial nets”. In: *Advances in Neural Information Processing Systems* 27 (2014).
- [19] Creswell, A. et al. “Generative adversarial networks: an overview”. In: *IEEE Signal Processing Magazine* 35.1 (2018), pp. 53–65.
- [20] Radford, A., Metz, L., and Chintala, S. “Unsupervised representation learning with deep convolutional generative adversarial networks”. In: *arXiv preprint arXiv:1511.06434* (2015).
- [21] Smith, R. L. “Max-stable processes and spatial extremes”. In: *unpublished manuscript* (1990).
- [22] Strokorb, K. and Schlather, M. “Characterizing extremal coefficient functions and extremal correlation functions”. In: *arXiv preprint arXiv:1205.1315* (2012).
- [23] Huster, T. et al. “Pareto GAN: Extending the representational power of GANs to heavy-tailed distributions”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 4523–4532.
- [24] Arjovsky, M., Chintala, S., and Bottou, L. “Wasserstein generative adversarial networks”. In: *International Conference on Machine Learning*. PMLR. 2017, pp. 214–223.
- [25] Lucic, M. et al. “Are GANs created equal? A large-scale study”. In: *Advances in Neural Information Processing Systems* 31 (2018).
- [26] Biewald, L. *Experiment Tracking with Weights and Biases*. Software available from wandb.com. 2020. URL: <https://www.wandb.com/>.

## A Theory

### A.1 Normalisation using extreme value theory

**Fisher-Tippett-Gnedenko theorem** Given a sequence of  $n$  independent and identically-distributed (i.i.d.) random variables  $X^1, X^2, \dots, X^n$  with a common distribution function  $\mathbb{P}(X^j \leq x) = F(x)$ , the Fisher-Tippett-Gnedenko theorem [14, 15] provides us with results about the behaviour of the sample maximum  $Y^{(n)} = \max(X^1, X^2, \dots, X^n)$  as  $n$  approaches infinity, in a manner analogous to the central limit theorem for the sample mean [16]. It is easily shown that the distribution function of  $Y^{(n)}$  is given by  $[F(x)]^n$ , however this becomes degenerate for large  $n$ . Hence a sequence of shift and scale parameters  $\{b_n\}_{n \in \mathbb{N}}$  and  $\{a_n : a_n > 0\}_{n \in \mathbb{N}}$ , is applied. The limit is instead defined for the distribution of

$$\tilde{Y}^{(n)} = \frac{Y^{(n)} - b_n}{a_n}.$$

The Fisher-Tippett-Gnedenko theorem shows that the limiting distribution must be one of the three distributions which make up the generalised extreme value (GEV) distribution: Gumbel, Fréchet, or Weibull. The GEV is defined three parameters  $(\xi, \mu, \sigma)$ , corresponding to its shape, location, and scale, respectively.

In practice, however, the  $X^j$  are rarely i.i.d. Climate data has trends, seasonality, and dependencies in time. This motivates a *block maxima* approach to extreme value modelling, where the maxima are taken over intervals of size  $k$  of the observations. For sufficiently large  $k$ , e.g.,  $k = 365$  for annual maxima, the maxima can be assumed to be independent and approximately stationary, and a GEV is fitted to these block maxima. Coles et al. [16] shows that, provided dependence becomes negligible for sufficiently far apart samples, the independence assumption can be relaxed as autocorrelation only affects the shape and scale parameters but not what the shape parameter.

### A.2 DCGAN architecture and training

A training set of 2000 days is sampled from the data, and the rest is withheld as a test set. We refer to the sequence of daily maxima for a single spatial location and variable as a marginal. The entire training set is used to fit a GEV distribution to each marginal and subset of size  $N$  of this set is then used for training the DCGAN.

The generator consists of an input layer that takes a 100-dimensional tensor  $z \in \mathbb{R}^{100}$ , followed by a dense layer with 25,600 neurons. The output of the dense layer is then reshaped into a  $5 \times 5 \times 1024$  tensor and passed through batch normalization, a leaky ReLU activation function, and a dropout layer. The tensor is then passed through two transposed convolutional layers, each with 512 filters, followed by batch normalization, a leaky ReLU activation function, and a dropout layer. The final layer is a transposed convolutional layer with three filters and a sigmoid activation function. The total number of trainable parameters in the generator is 10,490,883.




Figure 4: DCGAN generator architecture, adapted to the multivariate case from Boulaguiem et al. [10]. Layers are indicated below the blocks, LReLU stands for a leaky ReLU layer, BN for a batch normalisation layer, and sigmoid for a sigmoid activation layer. Above the blocks, F represents the filter size, and S the stride of the deconvolutional layers. The output has three channels and dimensions of  $20 \times 24$ .

The discriminator has three convolutional layers with 64, 128, and 256 filters respectively, each followed by LeakyReLU activation and dropout. The output of the last convolutional layer is reshaped to  $1 \times 6400$ . The output is then reshaped to a scalar and passed through a sigmoid activation function

to scale predictions to the  $[0, 1]$  interval, where 0 labels the input data as fake and 1 labels it as real. The model has a total of 405,057 parameters, with 404,289 of them being trainable.




Figure 5: DCGAN discriminator architecture, adapted to the multivariate case Boulaguiem et al. [10] Layers are indicated below the blocks, LReLU stands for a leaky RELU layer, BN for a batch normalisation layer, FC for a fully-connected layer, and sigmoid for a sigmoid activation layer. Above the blocks, F represents the filter size, and S the stride of the deconvolutional layers. The output has a scalar in the interval  $[0,1]$ .

GANs are notoriously unstable to train and several modifications have been suggested to improve their stability. We use label smoothing, where real images are assigned a label of 0.9 instead of 1, in the discriminator objective. Further modifications to the DCGAN architecture have been suggested, such as the Wasserstein GAN [24] which replaces the cross-entropy objective with the first Wasserstein distance and places additional constraints on the discriminator. However, motivated by the results of Lucic et al. [25] which find correct hyperparameter selection to be more powerful in improving GAN performance than architecture modifications, we focus on hyperparameter tuning rather than architecture changes.

We use the WandB [26] machine learning experiment tracking tool to perform a comprehensive Bayesian grid search and tune the hyperparameters. A table of all the hyperparameters and configurations explored, along with the final selections, is shown in Appendix B.

## B Tables and figures

### B.1 Tables

Table 1: Parameters included in the WandB Bayesian grid search and final selections.

Parameter	Values	Selected Value	Description
seed	[0, 1, 2, 6, 7, 42]	7	Included in sweep for reproducibility.
learning_rate	0.0001 - 0.0003	0.00013367626823798716	Controls the optimizer step size and rate of convergence.
beta_1	0.1 - 0.5	0.22693882275467836	Controls the exponential decay rate for the first moment estimates of the gradient for the Adam optimizer.
lrelu	0.1 - 0.4	0.2991161912395133	The gradient to assign to negative values in the Leaky ReLU function.
dropout	0.3 - 0.6	0.44053850596844424	Frequency the dropout function sets inputs to zero in training to prevent overfitting.
training_balance	[1, 2]	2	Ratio of training loops for discriminator vs. generator.

### B.2 Figures

#### GEV parameters

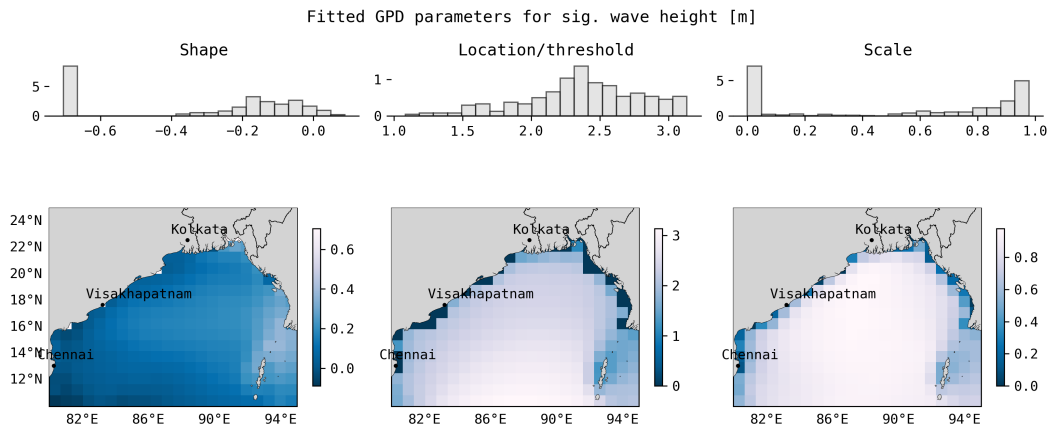


Figure 6: Fitted generalised extreme value (GEV) distribution parameters shape, location, and scale, over the Bay of Bengal for significant wave height from wind waves and swell in m. Here, the value -999 is used on land to represent no data.

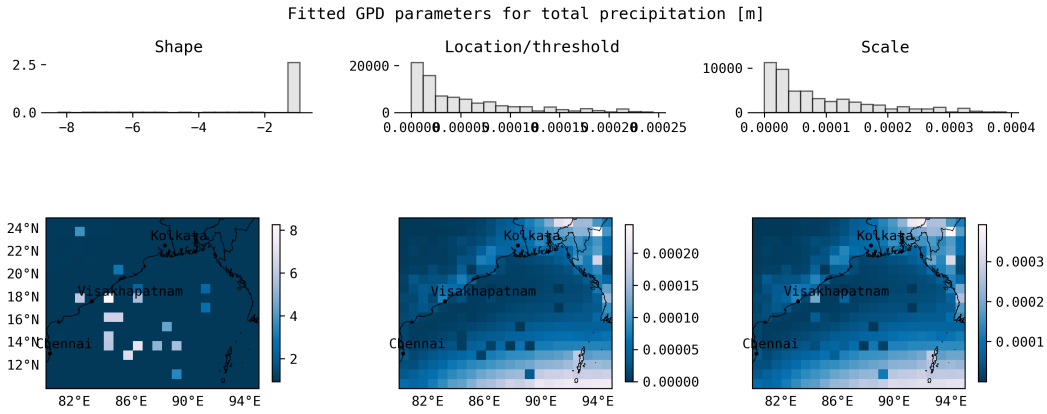


Figure 7: Fitted generalised extreme value (GEV) distribution parameters shape, location, and scale, over the Bay of Bengal for total precipitation m.

### Training and generated data




Figure 8: Sample of training data. Daily maximum wind speed (left), significant wave height (centre), and precipitation (right) are shown side-by-side for each sample.




Figure 9: Sample of generated data. Daily maximum wind speed (left), significant wave height (centre), and precipitation (right) are shown side-by-side for each sample.

### Multivariate extremal coefficients

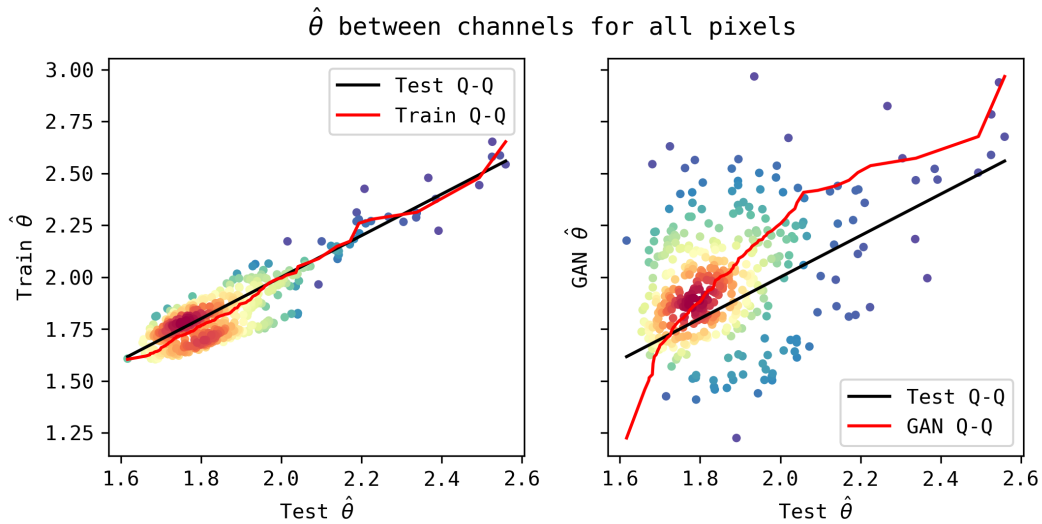


Figure 10: Scatter plots of extremal coefficient  $\hat{\theta}$  across all three channels for all 396 pixels for the test and train sets (left) and test and generated sets (right). Solid red lines show the Q-Q plots for the train and generated sets versus the test set.

## Spatial correlations

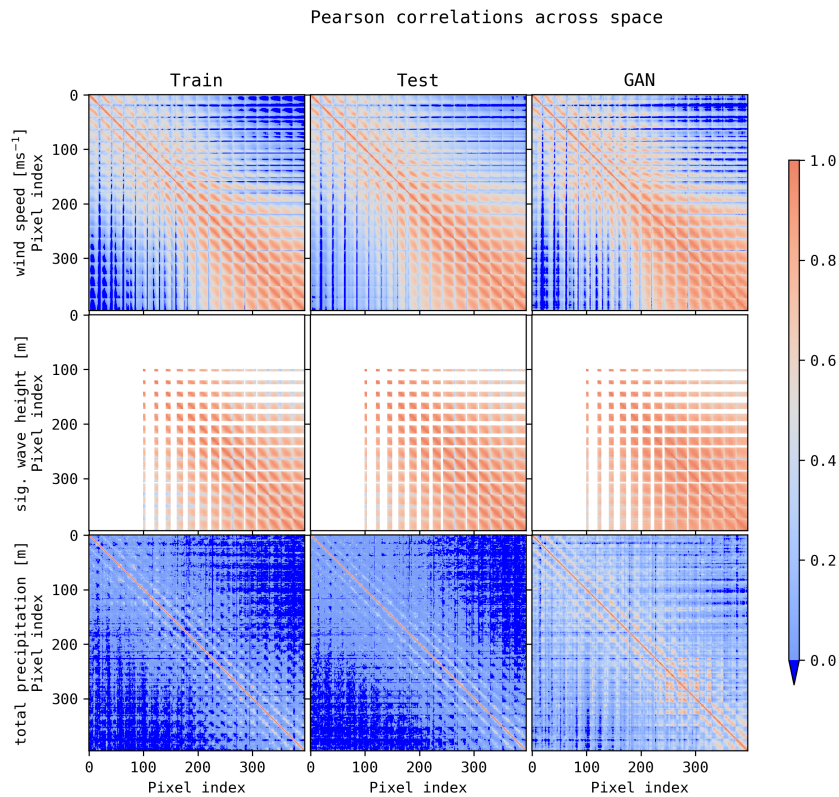


Figure 11: Matrix of pairwise Pearson correlations between all  $18 \times 22 = 396$  pixels for the training, test, and generated data. Clearly, the GAN captures the spatial correlation structure between the variables almost identically for wind speed (top row) and significant wave height (middle row), but underestimates the negative correlations for total precipitation (bottom row). We see a pattern of more positive correlations between nearby pixels for wind and precipitation. A group of positively correlated pixels emerges in the bottom right corner of the wind heatmaps. This corresponds to offshore pixels being more correlated to each other than to land pixels and aligns with observations made about the GEV fit to the wind data.




Figure 12: Matrix of pairwise extremal correlations  $\hat{\chi}$  between all  $18 \times 22 = 396$  pixels for the training, test, and generated data. The GAN captures the overall spatial pattern of the extremal correlations, correctly identifying that extreme winds (top row) offshore are more positively correlated to each other than to onshore extreme winds. The significant wave height (middle row) extremal correlations follow the spatial distribution train and test sets, with a tendency to underestimate extremal correlation as distance increases. Total precipitation (bottom row) appears to capture the overall distribution but appears noisier, with more negative  $\hat{\chi}$ s corresponding to the land pixels in the top left (approx pixels 1 to 150) of the heatmap.