

Downscaling of precipitation over the Taiwan region by a conditional Generative Adversarial Network

Original

Downscaling of precipitation over the Taiwan region by a conditional Generative Adversarial Network / Iotti, M., Davini, P., von Hardenberg, J., Zappa, G.. - In: POS PROCEEDINGS OF SCIENCE. - ISSN 1824-8039. - 415:(2022). (International Symposium on Grids & Clouds 2022 (ISGC2022)) [10.22323/1.415.0004].

Availability:

This version is available at: 11583/2972646 since: 2022-10-27T13:45:50Z

Publisher:

SISSA

Published

DOI:10.22323/1.415.0004

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Downscaling of precipitation over the Taiwan region by a conditional Generative Adversarial Network

M. Iotti,^{a,*} P. Davini,^b J. von Hardenberg^{c,b} and G. Zappa^d

^a*Dept. of Physics and Astronomy, University of Bologna,
Bologna, Italy*

^b*Institute of Atmospheric Sciences and Climate, CNR-ISAC,
Turin, Italy*

^c*Dept. of Environment, Land and Infrastructure Engineering, Polytechnic University of Turin,
Turin, Italy*

^d*Institute of Atmospheric Sciences and Climate, CNR-ISAC,
Bologna, Italy*

*E-mail: marcello.iotti3@unibo.it, p.davini@isac.cnr.it,
jost.hardenberg@polito.it, g.zappa@isac.cnr.it*

Predicting extreme precipitation events is one of the main challenges of climate science in this decade. Despite the computational power available nowadays, current state-of-the-art Global Climate Models' (GCMs) spatial resolution is still too coarse to correctly represent and predict small-scale phenomena such as convection, therefore precipitation prediction is still imprecise. For this reason, downscaling techniques play a crucial role, both for the understanding of the physical mechanisms behind precipitation onset and development, and for applications like hydrologic studies, risk prediction and emergency management.

Taking advantage of Deep Learning techniques, we exploit a conditional Generative Adversarial Network (cGAN) to train a generative model able to perform precipitation downscaling. This model, a deep Convolutional Neural Network (CNN), takes as input the precipitation field at the scale resolved by GCMs, adds random noise, and outputs a possible realization of the precipitation field at higher resolution, preserving the statistical properties of the input field. Also, being conditioned by the coarse-scale precipitation, the spatial structure of the produced small-scale field is consistent with the one prescribed by the input GCM prediction. To assess the skill of our model, we try to reconstruct the daily total precipitation field over the Taiwan region, starting from a coarsened version of the ERA5 reanalysis dataset. Results show the good ability of the model in capturing the features of the fine-scale precipitation field. In addition, compared to other downscaling techniques, our model has the advantage of being computationally inexpensive at run time and easily generalizable to any geographical domain.

International Symposium on Grids & Clouds 2022 (ISGC 2022)

21 - 25 March, 2022

*Online, Academia Sinica Computing Centre (ASGC), Taipei, Taiwan****

*Speaker

1. Introduction

Extreme precipitation events frequency has increased since the second half of the 20th century in many regions of the world [1]. The trend is expected to not diminish, as the Earth continues to warm due to climate change [2–6]. The social and economic impact of this is considerable: remarkable examples are floodings and landslides, which are the most destructive and therefore the most evident consequences of these events. Also the management of the increased precipitation amount is a challenging task, particularly in densely populated and highly urbanized areas. Indeed the development of human settlements is associated with a modification of the territory, where the natural soil surface is covered by buildings and impervious materials. In such an environment, storm runoff carries a large amount of pollutants, which ends up in streams and lakes, leading to a general degradation of water quality and aquatic ecosystem health [7]. Therefore predicting the frequency of extreme precipitation events is among the major challenges of climate sciences these days. Also the estimation of their occurrence in future scenarios play a central role.

Global Climate Models (GCMs), the main modeling tool of the climatological community, allow the integration of the state of the Earth's atmosphere for the long period of time (from years up to decades) needed for the analysis of the past and present climate, and for future projections and scenarios. To implement the equations governing the evolution of the flow of the Earth's atmosphere, they use a grid covering the entire globe, with a typical horizontal spatial resolution ranging from 50 to 200 km. The available computing power is one of the main factors conditioning this spatial resolution. Despite the continuous development in hardware technologies and code optimization, the trade-off between a higher spatial resolution on one hand, and the availability of hardware resources and the amount of time in which they are employed on the other hand, must allow to obtain analysis and projections at a reasonable cost. This implies that the spatial resolution of current state-of-the-art GCMs is too coarse to properly represent phenomena related to the onset and development of the precipitation, such as convection. In particular GCMs are not able to resolve circulation patterns related to extreme events [8]. In addition, since the precipitation field varies on spatial and temporal scales smaller than those explicitly resolved, their predictions are not adequate for impact studies and for the needs of applied disciplines such as hydrology.

Downscaling refers to any method to infer high-resolution variables from the corresponding low-resolution fields (see e.g. [9, 10]). Many downscaling techniques exist, and they can be grouped into two main families:

1. *Dynamical Downscaling* uses *Regional Climate Models* (RCMs), full atmospheric models with resolution higher than GCMs, which are nested (embedded) in them, but cover only a limited area of interest, to keep the overall computational load manageable.
2. *Statistical Downscaling* looks for statistical relationships between coarse-scale predictors and fine-scale predictands.

Among statistical downscaling methods, *Stochastic Downscaling* techniques deserve a special mention. They implement stochastic disaggregation algorithms, which, starting from the smooth distribution of the precipitation at the large scale, aim to generate a random small-scale precipitation field, satisfying both the large-scale constraints (e.g. the average precipitation intensity) and

	Dynamical Downscaling	Statistical Downscaling
PROs	Physics-based approach	Computationally efficient
CONs	Computationally expensive	Not physics-based Requires many observations Assumes stationarity

Table 1: Main pros and cons of Dynamical and Statistical downscaling techniques.

the consistency with the observed statistical properties of the small-scale precipitation [11]. Table 1 lists the main advantages and drawbacks of the two families of downscaling approaches. RCMs improve GCMs predictions and have a strong theoretical foundation, being based on the integration of the same set of prognostic equations. However many state-of-the-art codes are highly complex modelling systems and thus show scalability issues. Since they add a lot of computational burden to the already huge demands of GCMs, they should be preferred when this cost increase is manageable. Statistical downscaling techniques are computationally very efficient, since they result in relationships easy to implement and compute. Despite this, they often lack a physical foundation, because the relationship between large-scale predictors and small-scale predictands is empirical. Also, to tune this relationship many local-scale observations are required, and these are not always easily available. Finally, the application of statistical downscaling relies on the stationarity of the predictors-predictands relationship. This is not always verified, particularly when considering a future scenario, for which climate change should be properly taken into account.

Seen in the context of image processing, a downscaling task shares many similarities with super-resolution [12], the process of increasing the spatial resolution of an image beyond its original value. In recent years this field of study took advantage from advancements in deep learning techniques. Particularly, it benefited from the introduction of Convolutional Neural Networks (CNNs) and generative models like Generative Adversarial Networks (GANs). In this paper we describe the application of a deep convolutional conditional GAN to the downscaling of the precipitation field produced by a GCM.

The rest of the paper is structured as follows. Section 2 gives an overview about GANs, focusing on the principle behind their training. We also describe the experiment we performed, the dataset we used to train the GAN, and we define the metrics used to evaluate the quality of the generated precipitation field. In Section 3 we evaluate and show the performance of our GAN in reconstructing the daily total precipitation field over the region surrounding the island of Taiwan. Section 4 summarizes and concludes the paper.

2. Methods

2.1 GANs and conditional GANs

A GAN [13] consists of a pair of neural networks trained simultaneously and *adversarially*:

- A **generator**, learning a map from a noise probability distribution to the distribution of the data to be reconstructed.

- A **discriminator**, learning the probability that a sample is real or comes from the generator.

Formally, to capture the distribution of the original data, the generator builds a map from a given noise distribution $p_z(\mathbf{z})$ to a data space p_g . The mapping function modelled by the generator is a differentiable function $G(\mathbf{z}; \theta_g)$ depending on the parameters θ_g . The discriminator represents the function $D(\mathbf{x}; \theta_D)$, which is the probability of \mathbf{x} being from the actual data distribution rather than from p_g . The discriminator is trained like a classifier, to maximize the probability of a correct labelling (samples from the actual distribution labelled as “real” and samples from the generator labelled as “fake”). The generator is trained to cheat the discriminator, maximizing the probability that the discriminator assigns the wrong label (“real”) to a samples generated by it. Using the same notation as above, the generator and the discriminator play a min-max game with the following value function $V(D, G)$ (cf. [13]):

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [\log (1 - D(G(\mathbf{z})))] \quad (1)$$

In the case of a successful training, the data produced by the generator seem realistic, however there is no control over their modes. That is to say, the generated samples seem to come from the distribution of the actual data but it is not possible to specify which subset these samples have to belong to (e.g. if the generator is meant to reproduce the digits from 0 to 9, it is not possible to induce it to reproduce a specific digit). However, it is possible to provide the GAN with additional information in order to condition the generation process [15]. For example, the class labels of the data to be generated can be provided, or data from other modes. Let \mathbf{y} be this additional information; the conditioning of the GAN is performed providing both the generator and the discriminator with \mathbf{y} , in combination with the corresponding sample (a noise sample for the generator and a real/generated sample for the discriminator). Equation 1 becomes (cf. [15]):

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D(\mathbf{x}|\mathbf{y})] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [\log (1 - D(G(\mathbf{z}|\mathbf{y})))] \quad (2)$$

The straightforward application of this idea to an image super-resolution task is to provide both the generator and the discriminator with the coarse-resolution image, to properly direct the generation of the corresponding fine-resolution image.

2.2 Case study and data

To demonstrate the ability of our GAN in performing the downscaling of the precipitation field, we tried to reconstruct the daily total precipitation over the Taiwan region. We considered a squared domain, centered on the city of Taipei (lat 25.10 N, lon 121.59 E), and extended between latitudes 17 N and 32.75 N, and longitudes 113.25 E and 129 E (Figure 1). In order to asses our model’s abilities we carried out a *perfect model experiment*: considering high-resolution precipitation data, we degraded (*upscaled*) them by performing a coarsening operation, and used both the original and the corresponding low-resolution samples to train the GAN. At the end of the training we evaluated the skill of the generator in reproducing realistic high-resolution precipitation fields at the fine scale, using the original dataset as reference. As a data source we used the ERA5 reanalysis dataset [16, 17], which spans the time period from 1950 to present, providing hourly estimates for a large number of atmospheric, ocean and land-surface variables. This dataset covers the entire



Figure 1: The domain under investigation (red box) is centered on the city of Taipei (lat 25.10 N, lon 121.59 E) and includes the Taiwan island and the surrounding ocean (Made with Natural Earth. Free vector and raster map data @ naturalearthdata.com).

globe and has an horizontal resolution of $0.25^\circ \times 0.25^\circ$. We considered a subset of the entire ERA5 dataset, spanning the years from 1950 to 2020, corresponding to 25871 training examples (days). The operational steps of our experiment were as follows:

1. We performed a coarsening operation on the training dataset, reducing its horizontal spatial resolution to $1^\circ \times 1^\circ$. This operation was performed by aggregating every group of 4 adjacent grid cells, and taking the average of the daily total precipitation value for them.
2. The GAN was fed with both the original fine-scale and the coarsened samples. In particular, the generator used only the coarse-scale precipitation fields, while the discriminator was trained with couples of coarse-scale and the corresponding fine-scale examples, drawn from both the original dataset and from the generated samples. The generator performed the downscaling operation, producing precipitation fields with spatial resolution of $0.25^\circ \times 0.25^\circ$.

3. The training progress and the skill of the trained generator were assessed computing some climatologically meaningful metrics, comparing the generated dataset with the reference one.

The GAN was trained for 200 epochs on an NVIDIA Volta V100 GPU, provided by the CINECA consortium.¹ For each epoch the training algorithm processed the whole training dataset, performing the optimization of a loss function defined by Equation 2 for all the 25871 training examples, and updating the generator and the discriminator parameters accordingly.

2.3 Validation

Unlike other deep learning models, the loss function used to train a GAN (Eq. 2, Section 2.1) is not related to the quality of the generated images in a straightforward way. Many different metrics are designed to assess the quality of generated images. However, since our GAN is meant to generate the field of an atmospheric variable, we decided to evaluate the quality of samples produced by the generator using climatologically meaningful metrics. As the first quality metrics, we computed the climatology of the daily total precipitation field, over the whole set of generated samples. For each domain grid point $x^{(i,j)}$ the climatology is defined as:

$$x_{\text{clim}}^{(i,j)} = \frac{\sum_t x_t^{(i,j)}}{T} \quad (3)$$

where $x_t^{(i,j)}$ is the value of the daily total precipitation for that grid point, for the day t , and T is the total number of days in the dataset. In other words the climatology is simply the temporal average of the daily accumulated precipitation for each domain grid point. Then we computed the standard deviation of the daily total precipitation, which for each grid point is defined as:

$$x_{\text{SD}}^{(i,j)} = \sqrt{\frac{\sum_t \left(x_t^{(i,j)} - x_{\text{clim}}^{(i,j)} \right)^2}{T}} \quad (4)$$

with the same meaning of the symbols as above. The standard deviation expresses the dispersion of each grid point's daily total precipitation around the climatology. Lastly we computed the 95th percentile of the daily total precipitation field, which is an important metrics to assess how well the GAN is able to capture the magnitude and the location of extreme precipitation events.

This set of metrics allows to evaluate the quality of the fields produced by the GAN at the end of the training. We also computed them throughout the training process, using the set of precipitation fields generated at each epoch, to monitor the convergence of the networks.

3. Results

Figure 2 shows the climatology of the daily total precipitation field, for both the dataset generated by our GAN at the end of the training (epoch 200) and for the reference (ERA5) dataset. It is possible to notice that the trained GAN is able to reproduce the features of the mean daily total precipitation field. In particular both the position and the value of the maxima of the climatology

¹<https://www.hpc.cineca.it/>.

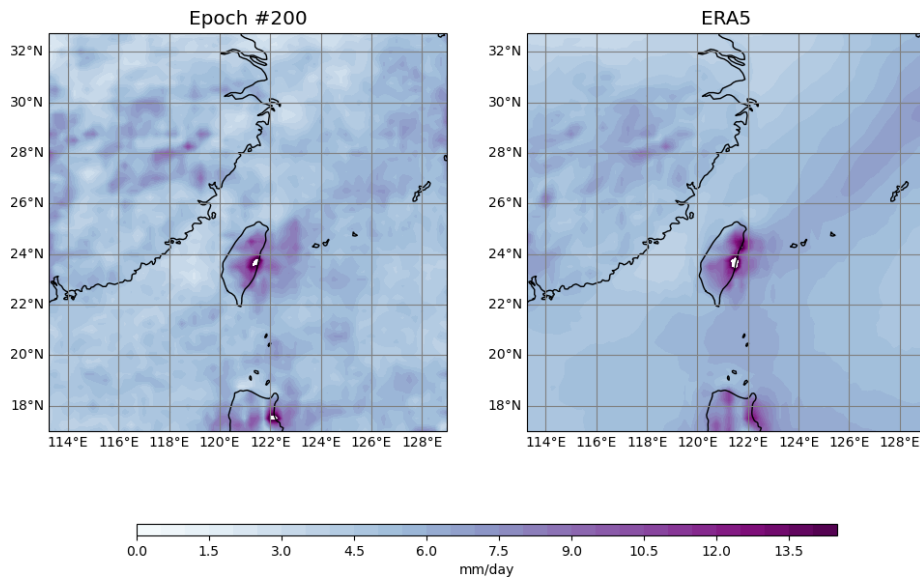


Figure 2: Climatology of the daily total precipitation for the generated dataset (left) and for the reference dataset (right).

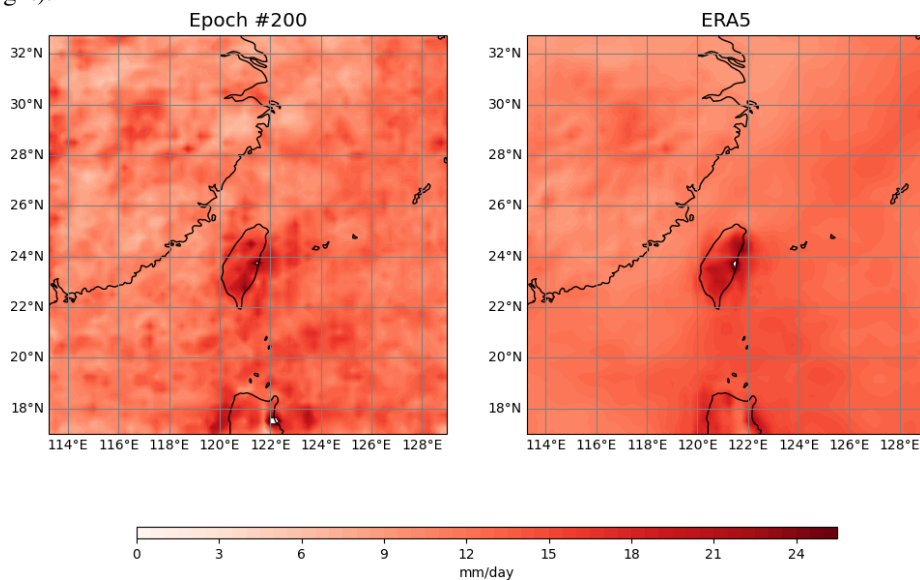


Figure 3: Standard deviation of the daily total precipitation for the generated dataset (left) and for the reference dataset (right).

field (over the island of Taiwan and above the ocean to the east of it; above the northern part of the island of Luzon) are well reproduced. Note also that the GAN manages to capture the excess of average daily total precipitation along the Okinawa Trough (which extends from the center of the domain towards the northeastern part of it), and over the oceanic region surrounding the northern part of the island of Luzon. However the climatology generated by the GAN tends to be noisier than the reference climatology, which appears much smoother, in particular over oceanic regions. Probably this noisy appearance is the residual of the striping effect related to the action of convolution kernels, which is very evident for the first training epochs. Similar considerations apply

to the standard deviation (Figure 3) and the 95th percentile (Figure 4) of the daily total precipitation. For both these metrics the noisy aspect is more pronounced with respect to the climatology, however the salient features and the spatial patterns of the reference dataset are reproduced by the GAN.




Figure 4: 95th percentile of the daily total precipitation for the generated dataset (left) and for the reference dataset (right).






Figure 5: Root-mean-squared errors of climatology, standard deviation and 95th percentile of the generated dataset with respect to the reference (ERA5) dataset, as the training of the GAN proceeds.


The plot in Figure 5 allows to monitor the training process of the GAN, giving a synthesis of the values of the three metrics analyzed above at each training epoch. The root-mean-squared



(a)



(b)



(c)

Figure 6: Average power spectrum of the daily total precipitation for the datasets generated at epoch 1 (a), 100 (b) and 200 (c) (red lines). Each plot also shows the reference (ERA5) average power spectrum (blue lines).

errors for climatology, standard deviation and 95th percentile of the daily total precipitation were computed for each training epoch according to:

$$\text{RMSE}(a) = \sqrt{\frac{\sum_{(i,j)} \left(x_{a, \text{generated}}^{(i,j)} - x_{a, \text{ERA5}}^{(i,j)} \right)^2}{N}} \quad (5)$$

where N is the total number of grid points within the domain and $a = \{\text{climatology, standard deviation, 95th percentile}\}$. The trend of the three metrics suggests that the GAN converges around epoch 150. The following fluctuations, particularly of the 95th percentile and of the standard deviation could probably be related to the local variability in the generated precipitation field. Indeed, in the absence of a physical constraint, such as the dynamics or a parameterization scheme of a Regional Climate Model, the spatial structure of the small-scale precipitation field is difficult to control. This is not necessarily to be considered a drawback of our model, as long as the total amount of precipitation over the domain is preserved and the large-scale structure of the precipitation is respected. In fact, a downscaling procedure aims at generating fine-scale precipitation fields with statistical properties similar to those of the precipitation observed over the corresponding area. Such fields should not be regarded in a deterministic way, but as a possible realization of the small-scale process, consistent with the structure of the large-scale precipitation [18].

Finally, Figure 6 shows the average power spectrum of the daily total precipitation, at the beginning, middle and at the end of the training. The power spectrum gives a quantitative indication of the spatial scales of variability of the precipitation field. We computed it, for both the generated and the reference dataset, on each individual day and then we took the time average. This calculation, carried out throughout the GAN training process, allowed to monitor the development of the GAN downscaling skills from a point of view different from the one shown above. After the first training epochs, the GAN is able to reproduce the reference spectrum for the lowest value of k . This indicates that the conditioning effect of the coarse-scale precipitation field has the desired effect, properly driving the generation of a field with the correct large-scale spatial structure. For higher value of k ($k > 15$) the generated spectrum diverges from the reference one, indicating the higher variability of the generated precipitation at the small scale. It should be emphasized that the fact that our GAN is able to reproduce the precipitation power spectrum with a good accuracy, without being based on physical/meteorological assumptions, is a remarkable achievement.

4. Conclusion

In this paper we presented the application of a deep convolutional Generative Adversarial Network to the task of precipitation downscaling. Such a model is an evolution of plain GANs, constructed by providing both the generator and the discriminator with the coarse-resolution precipitation field, in order to direct the generation process. To demonstrate the ability of the model, we tried to reconstruct the daily total precipitation over the island of Taiwan, in a perfect model setup: starting from a reference dataset, we upscaled it, and we trained the GAN using both the coarse- and the fine-scale datasets. Then we assessed the GAN's downscaling skills by comparing the fine-scale precipitation field from the generator with the reference dataset. The results show the good agreement between the generated and the ERA5 dataset, although the former shows more

variability at small scales, as pointed out both by analyzing the average power spectra of the daily total precipitation, and by commenting the “oscillatory” artifacts in the maps of Section 3. The improvement of model’s performance, with focus on the resolution of this shortcoming will be the subject of future work. The downscaling method we developed offers the advantages of a deep learning-based model, being inexpensive at run time (the computational load is mostly concentrated in the training phase) and easily generalizable to any geographic domain, for which only a re-training with an appropriate dataset is required.

Acknowledgments

ERA5 data [16, 17] were downloaded from the Copernicus Climate Change Service (C3S) Climate Data Store. The results presented in this paper contain modified Copernicus Climate Change Service information 2020. Neither the European Commission nor ECMWF is responsible for any use that may be made of the Copernicus information or data it contains.

We acknowledge the CINECA award under the ISCRA initiative, for the availability of high performance computing resources and support.

References

- [1] P. A. Arias et al., *Technical Summary*. In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, pp. 33–144, (Cambridge University Press, 2021).
- [2] B. Asadieh and N. Y. Krakauer, *Global trends in extreme precipitation: climate models versus observations*. *Hydrology and Earth System Sciences* **19** (2), 877–891 (2015), doi: [10.5194/hess-19-877-2015](https://doi.org/10.5194/hess-19-877-2015).
- [3] H. Madsen et al., *Review of trend analysis and climate change projections of extreme precipitation and floods in Europe*. *Journal of Hydrology* **519**, 3634–3650 (2014), doi: [10.1016/j.jhydrol.2014.11.003](https://doi.org/10.1016/j.jhydrol.2014.11.003).
- [4] M. G. Donat et al., *More extreme precipitation in the world’s dry and wet regions*. *Nature Clim Change* **6** (5), 508–513 (2016), doi: [10.1038/nclimate2941](https://doi.org/10.1038/nclimate2941).
- [5] G. Myhre et al., *Frequency of extreme precipitation increases extensively with event rareness under global warming*. *Sci Rep* **9**, 16063 (2019), doi: [10.1038/s41598-019-52277-4](https://doi.org/10.1038/s41598-019-52277-4).
- [6] S. M. Papalexiou and A. Montanari, *Global and Regional Increase of Precipitation Extremes Under Global Warming*. *Water Resources Research* **55** (6), 4901–4914 (2019), doi: [10.1029/2018WR024067](https://doi.org/10.1029/2018WR024067).
- [7] C. J. Walsh et al., *Urban Stormwater Runoff: A New Class of Environmental Flow Problem*. *PLoS ONE* **7** (9), e45814 (2012), doi: [10.1371/journal.pone.0045814](https://doi.org/10.1371/journal.pone.0045814).
- [8] D. Maraun et al., *Precipitation downscaling under climate change: Recent developments to bridge the gap between dynamical models and the end user*. *Reviews of Geophysics* **48** (3), (2010), doi: [10.1029/2009RG000314](https://doi.org/10.1029/2009RG000314).

- [9] B. C. Hewitson and R. G. Crane, *Climate downscaling: techniques and application*. Climate Research **07** (2), 85–95 (1996), doi: [10.3354/cr007085](https://doi.org/10.3354/cr007085).
- [10] R. L. Wilby and T. M. L. Wigley, *Downscaling general circulation model output: a review of methods and limitations*. Progress in Physical Geography: Earth and Environment **21** (4), 530–548 (1997), doi: [10.1177/030913339702100403](https://doi.org/10.1177/030913339702100403).
- [11] L. Ferraris et al., *A comparison of stochastic models for spatial rainfall downscaling*. Water Resources Research **39** (12), (2003), doi: [10.1029/2003WR002504](https://doi.org/10.1029/2003WR002504).
- [12] M. Reichstein et al., *Deep learning and process understanding for data-driven Earth system science*. Nature **566** (7743), 195–204 (2019), doi: [10.1038/s41586-019-0912-1](https://doi.org/10.1038/s41586-019-0912-1).
- [13] I. Goodfellow et al., *Generative Adversarial Nets*. In *Advances in Neural Information Processing Systems* vol. 27, (Curran Associates Inc., 2014), <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf>.
- [14] I. Goodfellow et al., *Generative Adversarial Networks*. Commun. ACM **63** (11), 139–144 (2020), doi: [10.1145/3422622](https://doi.org/10.1145/3422622).
- [15] M. Mirza and S. Osindero, *Conditional Generative Adversarial Nets*. arXiv:1411.1784 [cs, stat] (2014), <http://arxiv.org/abs/1411.1784> (Accessed: Apr. 01, 2022).
- [16] B. Bell et al., *ERA5 hourly data on single levels from 1950 to 1978 (preliminary version)*, Copernicus Climate Change Service (C3S) Climate Data Store (CDS), (2020), <https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels-preliminary-back-extension?tab=overview> (Accessed: Mar. 01, 2022).
- [17] H. Hersbach et al., *ERA5 hourly data on single levels from 1979 to present*, Copernicus Climate Change Service (C3S) Climate Data Store (CDS), (2018). doi: [10.24381/cds.adbb2d47](https://doi.org/10.24381/cds.adbb2d47) (Accessed: Mar. 01, 2022).
- [18] N. Reborá et al., *RainFARM: Rainfall Downscaling by a Filtered Autoregressive Model*. Journal of Hydrometeorology **7** (4), 724–738 (2006), doi: [10.1175/JHM517.1](https://doi.org/10.1175/JHM517.1).