
GENERATIVE DIFFUSION-BASED DOWNSCALING FOR CLIMATE

R. A. Watt*, L. A. Mansfield*
Stanford University
Palo Alto, California
{rwatt1, lauraman}@stanford.edu

ABSTRACT

Downscaling, or super-resolution, provides decision-makers with detailed, high-resolution information about the potential risks and impacts of climate change, based on climate model output. Machine learning algorithms are proving themselves to be efficient and accurate approaches to downscaling. Here, we show how a generative, diffusion-based approach to downscaling gives accurate downscaled results. We focus on an idealised setting where we recover ERA5 at 0.25° resolution from coarse grained version at 2° resolution. The diffusion-based method provides superior accuracy compared to a standard U-Net, particularly at the fine scales, as highlighted by a spectral decomposition. Additionally, the generative approach provides users with a probability distribution which can be used for risk assessment. This research highlights the potential of diffusion-based downscaling techniques in providing reliable and detailed climate predictions.

Keywords Diffusion · Generative models · Machine learning · Climate models · Downscaling · Super-resolution

1 Introduction

Climate change poses a threat to humans and ecosystems all over the world. The potential impacts range from increased risk of damage caused by extreme events such as heatwaves, heavy rainfall, and flooding events, to disruptions in biodiversity, agriculture, and threats to food security. This underscores the critical need for accurate predictions of future climate conditions for effective mitigation and adaptation strategies.

Global Climate Models (GCMs) have been pivotal in our efforts to simulate the Earth’s complex system and project future climate scenarios. GCMs are sophisticated computer models that incorporate components from the atmosphere, ocean, land, vegetation, and sea-ice, creating a coupled Earth system. However, the computational cost associated with running GCMs at high resolutions presents a significant challenge. Currently, GCMs operate on a grid with a typical resolution of approximately 1° (roughly 100 kilometers) [Chen et al., 2021b]. For assessing climate change impacts, we often require more localised predictions, crucial for addressing regional impacts. This limitation prompts the exploration of alternative methodologies, such as downscaling, to enhance the spatial precision of climate predictions.

In recent years, machine learning techniques have proven capable of making substantial strides in predicting climate patterns. Many believe it will revolutionise our approach to climate modeling, by offering computationally efficient methods to improve, or even replace, traditional GCMs [Balaji, 2021, Schneider et al., 2023, Mansfield et al., 2023, Rolnick et al., 2023]. The challenges of climate and weather prediction have caught the attention of large technology companies, such as Google, NVIDIA, Microsoft, Huawei, and others, who are leveraging machine learning to facilitate numerical weather prediction [Lam et al., 2023, Pathak et al., 2022, Bi et al., 2022, Nguyen et al., 2023, Chen et al., 2023]. There is also a growing interest in machine learning emulators of GCMs [Mansfield et al., 2020, Watson, 2022, Watt-Meyer et al., 2023] and in embedding machine learning components into GCMs (such as subgrid-scale parameterisations) [Kochkov et al., 2024, O’Gorman and Dwyer, 2018, Gentile et al., 2018], amongst other uses [Molina et al., 2023].

One specific avenue gaining popularity is downscaling, a technique where coarse-resolution climate models are refined using machine learning to generate high-resolution predictions. This could enhance the cost-effectiveness of coarse

*Equal contribution

climate models and their precision on regional scales. This paper uses a generative machine learning approach based on diffusion [Sohl-Dickstein et al., 2015] for downscaling climate data. Using the ERA5 reanalysis dataset, we show how a diffusion models can enhance a coarse resolution map of temperature and winds over the USA [Hersbach et al., 2020]. The benefit is that ensembles can be produced, which are crucial for assessing model uncertainty in climate change studies [Hawkins and Sutton, 2009]. Diffusion-based generative models have been applied in similar downscaling applications [Mardani et al., 2023, Addison et al., 2022, Nath et al., 2024], although they have not yet gained prominence in the climate downscaling literature. Our aim is to contribute to the evolving landscape of machine learning climate downscaling, with a diffusion-based approach applied on continental scales for one of the first times. While previous studies focus on downscaling to km-scale resolution [Mardani et al., 2023, Addison et al., 2022, Nath et al., 2024], we envision how diffusion could be used to downscale conventional climate models with typical CMIP-resolution, i.e., $\mathcal{O}(1^\circ)$ or $\mathcal{O}(100 \text{ km})$, to a higher resolution that would typically require a regional climate model, i.e., $\mathcal{O}(0.1^\circ)$ or $\mathcal{O}(10 \text{ km})$. In agreement with several recent studies, our results show that diffusion-based generative models are a promising approach for climate data, both for downscaling [Mardani et al., 2023, Addison et al., 2022, Nath et al., 2024] and in other applications [Huang et al., 2024, Price et al., 2023, Chan et al., 2024, Bassetti et al., 2023, Cachay et al., 2023].

2 Background

2.1 Downscaling

Climate downscaling, often referred to as super-resolution (based on the machine vision literature), is the process of refining predictions from a low-resolution climate model to a higher resolution. This step is crucial for addressing the limitations of global climate models (GCMs) and tailoring climate predictions to local or regional scales. Traditionally, two main approaches have been employed for downscaling: dynamical downscaling and statistical downscaling, where the latter includes machine learning techniques growing in popularity recently.

Dynamical downscaling involves running a regional climate model (RCM) at a higher resolution, typically 10-50 km, over a specific region of interest [Sunyer et al., 2012, Giorgi, 2019, Tapiador et al., 2020, Giorgi and Jr, 2015]. The low-resolution output from a GCM serves as both the boundary and initial conditions for the RCM. RCMs offer the advantage of enhanced spatial resolution, while guaranteeing the output is dynamically consistent, rendering them popular amongst stakeholders and policymakers for important decisions [Gutowski et al., 2020]. However, RCMs also have inherent biases, like GCMs, and may still require postprocessing techniques to remove biases. Importantly, dynamical downscaling comes with a notable drawback in the high computational cost of running high-resolution RCMs.

Statistical downscaling, in contrast, utilises statistical methods to establish relationships between coarse-resolution model outputs and high-resolution observations in historical data [Maraun and Widmann, 2018b, Vandal et al., 2019]. One approach is to employ a regression model which directly predicts high resolution model variables from low resolution model variables. This assumes that the low resolution model is a “perfect” predictor of the high resolution model (i.e., there are no model biases). This is traditionally known as “perfect prognosis” [Schubert, 1998]. Since low resolution models are often missing processes and feedbacks, this approach alone can lead to inaccurate downscaled predictions [Maraun and Widmann, 2018a]. An alternative approach, known as “model output statistics”, aims to match the predicted statistics to observed statistics, such as the mean, standard deviation, and/or quantiles. This automatically accounts for model biases and is used for the NASA Earth Exchange Global Daily Downscaled Projections [Thrasher et al., 2022]. Another approach lies in “stochastic weather generators”, that simulate stochasticity of weather data based on characteristics of observations [Wilby et al., 1998]

Although faster than dynamical downscaling, these traditional methods often exhibit poor performance in extrapolating to new climates. This limitation arises from the assumption that the same statistical relationships hold under changing conditions, known as stationarity, which we cannot necessarily expect to be true [Fowler et al., 2007, Maraun et al., 2017]. One solution could be hybrid downscaling, such as statistical downscaling applied to RCM outputs [Sunyer et al., 2012]. Alternatively, machine learning/artificial intelligence may offer solutions to this problem.

Machine learning techniques are becoming a popular choice for statistical downscaling that can be framed as perfect prognosis [Baño-Medina et al., 2022], model output statistics [Pour et al., 2018], weather generators [Wilby et al., 1998] or a combination of these [Vandal et al., 2019]. Many of these approaches are based on developments in computer vision for super-resolution of images [Glasner et al., 2009].

2.2 Machine learning

Downscaling can be viewed as a supervised learning task, where the goal is to learn the variables on the high resolution grid (the target), from the variables on the low resolution grid (the input). The majority of machine

learning studies aim to directly learn the mapping between the low resolution input and the high resolution target with a wide variety of machine learning architectures, including random forests [Hasan Karaman and Akyürek, 2023, Medrano et al., 2023, Chen et al., 2021a], support vector machines [Pour et al., 2018], and convolutional neural network (CNN) architectures [Baño-Medina et al., 2022, Jiang et al., 2021, Wang et al., 2019], including U-Net architectures [Agrawal et al., 2019, Kaparakis and Mehrkanon, 2023, Adewoyin et al., 2021, Cho et al., 2023], convolutional autoencoders [Babaousmail et al., 2021] and fourier neural operators [Yang et al., 2023]. These studies show that machine learning can be more accurate than conventional statistical downscaling, with the benefit of low computational cost compared against dynamical downscaling approaches. However, end-users are often concerned when there is a risk of unrealistic predictions from black-box machine learning methods, especially when used in new scenarios such as climate change. Recently, physics-informed machine learning has shown to be a potential solution to this, by embedding known constraints from physics [Harder et al., 2024, Harder et al., 2022].

Generative models, which aim to learn underlying probabilities distributions of the data, are a promising approach for matching downsampled statistics to that of observations (similar to statistical downscaling with “model output statistics” described above). Many high resolution images can be generated through sampling, or *conditioning on*, the low resolution data (also making them a type of “stochastic weather generator”). For example, several recent studies use Generative Adversarial Networks (GANs) to generate high resolution images conditioned on low resolution images [Leinonen et al., 2021, Harris et al., 2022, Oyama et al., 2023, Wang et al., 2021, Price and Rasp, 2022]. Diffusion-based approaches have become established as a state-of-the-art technique for image generation (e.g., Stable Diffusion [Rombach et al., 2022], DALL-E2 [Ramesh et al., 2022], amongst others [Saharia et al., 2022, Nichol et al., 2022]) but have not yet become widely used for climate downscaling. We expect to see more diffusion-based downscaling methods in the near future, following [Bischoff and Deck, 2023, Wan et al., 2023] who used diffusion for downscaling turbulent fluid data and [Mardani et al., 2023, Addison et al., 2022, Nath et al., 2024] who presented diffusion for downscaling climate model output on localised scales. In a similar realm, [Groenke et al., 2020] use unsupervised normalising flows to downscale climate variables. In this study, we demonstrate the performance of a diffusion-based approach to downscaling on a continental scale.

The probabilistic nature of generative machine learning makes them particularly desirable for risk assessment studies, such as those aiming to quantify the likelihood of extreme events. Both weather and climate modelling communities have long used ensembles of simulations to assess model, scenario, and initial condition uncertainty [Hawkins and Sutton, 2009]. There is a growing interest in generative machine learning for generating ensembles from one climate/weather model simulation [Li et al., 2023]. In a downscaling setting, ensembles could be leveraged to determine the trustworthiness of predictions, for example, by highlighting increased model uncertainty when applied to out-of-distribution samples that are likely to occur in a non-stationary climate [Fowler et al., 2007].

3 Methods

3.1 Data

The ERA5 reanalysis dataset, made publicly available by ECMWF [Hersbach et al., 2020], is used in this study. Reanalysis datasets optimally combine observations and models through data assimilation techniques. ERA5 includes hourly estimates for a wide range of atmospheric, land, and oceanic climate variables on a 0.25° resolution longitude-latitude grid with 137 vertical levels, from January 1940 to present day. This dataset is becoming widely used in other machine learning weather/climate prediction studies [Pathak et al., 2022, Bi et al., 2022, Lam et al., 2023, Chen et al., 2023].

We consider downscaling of three variables over the continental USA:

- Air temperature at 2 m
- Zonal wind at 100 m
- Meridional wind at 100 m

We focus only on these three variables over the USA as a demonstration of diffusion for downscaling on continental scales. Although relevant for extreme events, we do not include precipitation here because ERA5 reanalysis at 0.25° resolution is likely too coarse to capture its spatial intermittency [Bihlo, 2021]. Diffusion-based approaches to downscaling precipitation on kilometer scales can be found in [Mardani et al., 2023, Harris et al., 2022, Nath et al., 2024].

ERA5 provides hourly data from 1940 to present day, however, we expect this to be highly correlated in the temporal domain. To reduce the dataset size without significant loss of information, we subsample randomly in time to select

Table 1: Table describing inputs and outputs in this study. By "fixed" inputs we mean the same for all inputs.

Inputs	Outputs
Inputs (coarse grid):	Difference (fine grid - coarse grid):
Temperature at 2m level	Temperature at 2m level
Zonal winds at 100m level	Zonal winds at 100m level
Meridional winds at 100m level	Meridional winds at 100m level
Fixed inputs (fine grid):	
Land-sea mask	
Height of land surface	
Embeddings (scalars):	
Month	
Time of day	

only 30 timesteps per month. This reduces the dataset size by $1/24$ and gives us data sampled approximately once per day, at different times of day. We use years 1950-2022 and separate this into a training dataset (1950-2017) and testing dataset (2018-2022).

We aim to downscale a coarsened version of the ERA5 dataset back onto the original 0.25° resolution grid. For the coarsened dataset, we use bi-linear interpolation onto a 2° resolution grid, to be approximately consistent with typical climate model resolution (e.g., $\mathcal{O}(1^\circ)$ in CMIP6 models [Chen et al., 2021b]). This corresponds to approximately $8\times$ scaling of resolution. Note that this approach assumes the coarse resolution data is not biased, following the perfect prognosis approach to downscaling. To apply this to real climate model data, one may first need to carry out bias correction [Mardani et al., 2023].

3.2 Baseline U-Net for downscaling

Our problem is a supervised learning problem, where we aim to learn the fine resolution variables from the coarse resolution variables. We will compare two approaches for this task: firstly, a U-Net architecture and secondly, a diffusion-based generative model. Since the inputs and outputs are both images with a channel for each variable, this requires an image-to-image model such as a U-Net [Ronneberger et al., 2015]. A U-Net architecture is a neural network comprised of several encoding convolutional layers followed by several decoding up-convolutional layers, making them ideal for computer vision tasks when the input and output images are of the same size. They have shown excellent performance on image processing tasks, such as segmentation [Ronneberger et al., 2015] and super-resolution [Hu et al., 2019]. As we will see in the following section, our diffusion-based approach also uses a U-Net, allowing us to use the same base model, making for a fair comparison between these methods.

Our input variables are the air temperature at 2 m, and zonal and meridional winds at 100 m, all coarsened onto the 2° resolution grid. Note that after coarsening, these variables are stored on the fine resolution (0.25°) grid using bi-linear interpolation, ensuring that the inputs and outputs have the same dimensions and allowing for a simple U-Net architecture. To improve efficiency of the U-Net, rather than directly learning the output on the fine resolution grid, we learn the difference between the output on the fine and coarse grid. As well as these inputs, we also provide the U-Net with fixed inputs for spatial quantities that describe the geography and are constant in time. These are the land-sea mask and the height of the land surface, defined on the fine resolution grid. We expect these to aid learning of details around the coast, the great lakes and the mountains. This means the input image is of size $(128 \times 256 \times 5)$ while the output image is of size $(128 \times 256 \times 3)$. Finally, we also provide scalar values representing the month and the time of day, allowing the U-Net to learn differences in the diurnal and seasonal cycles. These are input to the each block of the U-Net, using a shallow neural network. Table 1 shows these inputs and outputs to the network. All variables are normalised to have standard scaling with zero mean and unit variance. The U-Net is trained by minimising the Mean Squared Error (MSE) between the U-Net predicted image and the samples from the data.

3.3 Diffusion-based generative model

Generative models using diffusion have shown great success as a method for synthesising images [Sohl-Dickstein et al., 2015]. As a generative model, the goal is to learn some probability distribution $p(\mathbf{x})$ given a set of samples $\{\mathbf{x}\}$. This would be an unconditional model, however, we are interested in building a conditional

model, $p(\mathbf{x}|\mathbf{y})$, where $\mathbf{x}$ are the high resolution image samples and $\mathbf{y}$ are the low resolution image samples. In the following discussion, we will neglect the conditioning on $\mathbf{y}$ for clarity, but note that the conditional form is obtained by replacing $\mathbf{x}$ with $\mathbf{x}|\mathbf{y}$.

Diffusion models take inspiration from thermodynamics whereby a diffusion process $\{\mathbf{x}(t)\}_{t=0}^T$ with $t \in [0, T]$ is constructed which transforms samples from the data distribution $\mathbf{x}(0) \sim p(\mathbf{x})$ to that of an isotropic Gaussian $\mathbf{x}(T) \sim \mathcal{N}(\mathbf{0}, \sigma\mathbf{I})$. Following [Song et al., 2021], the stochastic differential equation for such a process is given by

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + g(t)d\mathbf{w} \quad (1)$$

where $\mathbf{f}(\mathbf{x}, t)$ is the drift coefficient $g(t)$ is the diffusion coefficient and $\mathbf{w}$ is a Wiener process. If this process can be reversed, one can use it to generate samples from $p(\mathbf{x}_0)$. This involves first sampling from the terminal distribution $\mathbf{x}(T)$ (the isotropic Gaussian), then transforming the sample through the reverse processes to the data distribution. It can be shown that the reverse of a diffusion process is itself a diffusion process with the following reverse time stochastic differential equation (SDE) [Anderson, 1982]

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})]dt + g(t)d\bar{\mathbf{w}} \quad (2)$$

where $\bar{\mathbf{w}}$ is a Wiener process with time flowing backwards from T to 0 . $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ is the score function, a vector field which points towards higher density. The score does not depend on the intractable normalisation constant making it easier to evaluate. Obtaining samples by solving the reverse SDE requires knowledge of the score function. To do this we parameterise the score using a neural network $\mathbf{s}_\theta(\mathbf{x}, t)$, with the same U-Net architecture as described above in Section 3.2. The weights of $\mathbf{s}_\theta(\mathbf{x}, t)$ are optimised using score matching [Vincent, 2011] by minimising the following loss [Song et al., 2021]

$$\theta^* = \arg \min_{\theta} \mathbb{E}_t \left[\lambda(t) \mathbb{E}_{\mathbf{x}(0)} \mathbb{E}_{\mathbf{x}(t)|\mathbf{x}(0)} \left[\|\mathbf{s}_\theta(\mathbf{x}, t) - \nabla_{\mathbf{x}(t)} \log p_t(\mathbf{x}(t)|\mathbf{x}(0))\|_2^2 \right] \right] \quad (3)$$

where $\lambda(t)$ is a weighting function, t is sampled uniformly between $[0, T]$. This equation is obtained through minimising the evidence lower bound (ELBO) on the negative log likelihood, $\mathbb{E}[-\log p(\mathbf{x}_0)]$, and reweighting by $\lambda(t)$.

Within the framework discussed above, there are a number of design choices and free parameters to set. In this work, we use the implementation given by Karras et al [Karras et al., 2022]. Their suggestions give a noticeable improvement over previous works [Song et al., 2021, Ho et al., 2020, Nichol and Dhariwal, 2021] for image synthesis. One improvement to note is the use of a higher order (second) integration scheme. This significantly reduces the number of timesteps taken when solving the reverse process making inference significantly more efficient. We use 100 timesteps during evaluation, but find that using just 50 is sufficient for a similar accuracy.

4 Results

After training on years 2050-2017 of ERA5, we evaluate the performance of both the U-Net and the diffusion-based model on years 2018-2022.

4.1 Example

Figure 1 shows the coarse resolution maps, the true fine resolution maps, and the downsampled fine resolution maps for the U-Net and diffusion, for one snapshot for all three variables. Both machine learning approaches demonstrate good predictions when compared against the fine resolution. For temperature, the U-Net and diffusion predictions are virtually indistinguishable. For the winds, there are some smaller scale features in the diffusion prediction, while the U-Net predicts smoother features. We expect this is because U-Net is trained to minimise MSE, which naturally favours smoother fields rather than high frequency variations.

4.2 Probabilistic Predictions

Figure 1(j-l), shows one generated downsampled prediction from the diffusion model. However, one of the benefits of the diffusion-based model lies in its ability to generate multiple predictions, creating probability distributions, rather than one deterministic prediction. We generate an ensemble of diffusion predictions consisting of 30 samples, each of which is generated from a different sample from the terminal Gaussian distribution $\mathbf{x}(T)$. This provides us with valuable information when making high-stakes decisions, for example, by highlighting when there is higher uncertainty amongst ensemble members. We find the ensemble members agree reasonably well and show their 1 standard deviation in Figure 1(m-n) for each variable. There is higher uncertainty in regions where the variable is changing quickly, in

Time: 2022-05-06T19:00:00.000000000




Figure 1: Maps comparing (a-c) coarse resolution (the input) to the fine resolution for (d-f) the truth, (g-i) the U-Net downscaled prediction and (j-l) one generated diffusion downscaled prediction, for all three variables at a single timestep. (m-n) show 1 standard deviation across a 30-member ensemble generated by diffusion.




Figure 2: Maps comparing the error metrics for the U-Net downscaled prediction and the diffusion downscaled prediction, for all three variables.

other words, where there are larger spatial gradients. For the surface temperatures, this increased uncertainty tends to be over mountainous regions and for the winds, to the east of mountains and along fronts where the winds are rapidly changing. This makes sense for a downscaling problem where the prediction is constrained by the coarse grid, but the diffusion generates many possible realisations for how the variables could be interpolated between the coarse grid cells. This also results in a gridded pattern where predictions closer to the coarse grid cells have lower uncertainty.

4.3 Metrics

Here, we aim to compare the results of a deterministic U-Net with a probabilistic diffusion-based model. To compare both, we will use the mean absolute error (MAE). For the diffusion-based model, we use the mean across across the ensemble. To consider the probabilistic nature of diffusion, we will also consider the continuous ranked probability score (CRPS), a generalisation of mean absolute error to compare predicted probability distributions to a single ground truth [Hersbach, 2000, Gneiting and Raftery, 2007]. Note, that we present the CRPS metric for diffusion only and cannot fairly compare this with the deterministic U-Net predictions.

Figure 2 shows maps showing these metrics averaged over the entire test dataset. The gridded pattern appears due to reduced error at grid-cells close to the coarse-resolution grid, also present in Figure 1(m-o). The diffusion-based approach shows slightly lower MAE in the high altitude regions along the Rockies, particularly for temperature. These are the regions that exhibit more variability, both in time and space, which could explain why a probabilistic model performs better. The value in the probabilistic approach is further highlighted by the significant reduction in error when considering the CRPS metric for diffusion, which takes into account all ensemble members.

Table 2 shows the mean metrics, averaged across the entire domain and test dataset. Note that when taking spatial averages, we weight each grid-cell by the cosine of latitude to account for the non-uniform size of grid cells.

Table 2: Results

Metric	Model	Temperature ($^{\circ}\text{C}$)	Zonal winds (m/s)	Meridional winds (m/s)
MAE	Linear Interpolation	1.047	0.733	0.746
	U-Net	0.384	0.335	0.348
	Diffusion	0.328	0.308	0.319
CRPS	Diffusion	0.254	0.224	0.232




Figure 3: (a-c) Power spectrum on a log-scale for all three variables for the truth, diffusion, U-Net and linear interpolation, and (d-f) the difference between the truth and the predicted power spectra for diffusion, U-Net and linear interpolation.

4.4 Spectra

In the example snapshots (Figure 1), it appears that the diffusion model produces more accurate high frequency variations for the winds compared to the U-Net. We validate this further by calculating the power spectra across all wavelengths, for all three variables. Here, we treat the data as an image and take a 2D Fourier transform across this image to estimate the power for each wavelength. These results are robust to transforming the data onto a grid equispaced in distance, accounting for different grid spacing in latitude.

Figure 3(a-c) shows the power spectra for each variable for all methods on a log-scale, compared against the ground truth in black. The power spectra for diffusion is indistinguishable from the ground truth, so we also present the differences between these in Figure 3(d-f). The power spectra is significantly more accurate across all scales in the diffusion model in comparison to the U-Net. For the winds, we see this is more evident at the high wavenumbers. This shows the promise of diffusion for approximating geospatial data for a range of uses, potentially addressing issues of oversmoothing seen in other weather and climate studies [Pathak et al., 2022, Bi et al., 2022, Lam et al., 2023].

5 Conclusions

In this paper, we have presented a generative diffusion-based model for downscaling climate data on continental scales. We demonstrated this by recovering 0.25° resolution ERA5 data from a 2° resolution coarse-grained version of ERA5

and found that the diffusion model outperformed a baseline U-Net with the same architecture. The next step would be to apply this diffusion model to downscale the output of a coarse $\mathcal{O}(1^\circ)$ resolution climate model to a higher resolution, e.g., 0.25° ERA5. This presents the additional challenge that the coarse model output may not match up with the high resolution dataset. One approach could be to first apply a bias-correction technique, as done in [Mardani et al., 2023]. The application of downscaling to climate model output based on historical observational datasets also presents the issue of non-stationarity, whereby we cannot assume that the relationship between the coarse resolution and the high resolution data remains constant under a changing climate. Diffusion-based approaches may be appealing for this task, as they predict an ensemble which can highlight the epistemic uncertainty associated with each prediction.

There are a wide range of avenues to expand upon this study, for example, increasing the number of variables predicted, exploring performance at different resolutions, applying the same method to different regions or to full global data, or incorporating temporal, as well as spatial, downscaling. Downscaling of precipitation is of particular interest, due to its intermittent spatial patterns that is typically overly smooth in climate model output compared to observations and is a crucial variable for extreme events such as flooding [Sunyer et al., 2012, Baño-Medina et al., 2022, Vandal et al., 2017, Xu et al., 2020, Akinsanola et al., 2018, Babausmail et al., 2021, Harris et al., 2022].

This study was carried out without access to expensive, high performance computing systems. Although both training and inference was more expensive than the baseline U-Net, diffusion offers significantly improved performance and the ability to generate ensembles, with a computational cost several orders of magnitude lower than full climate model simulations. This could make impact studies accessible and affordable to a wider range of end-users. We conclude this study by noting that, given the rapid improvements in diffusion-based image generation in the last three years [Zhang et al., 2023], we might expect to see even more advances in diffusion-based techniques applied to climate problems. Even during the time it took us to complete this study, we have seen diffusion models gaining popularity in this field [Mardani et al., 2023, Price et al., 2023, Nath et al., 2024, Huang et al., 2024, Chan et al., 2024], a trend that we hope will continue in the years to come.

Open data statement

All code used in this study is available at <https://github.com/robbiewatt1/ClimateDiffuse>. We followed the diffusion implementation of [Karras et al., 2022] available at <https://github.com/NVlabs/edm>. The ERA5 reanalysis dataset used in this study is publicly available at <https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-complete?tab=overview> [Hersbach et al., 2020].

Acknowledgements

This was done in our spare time and not directly funded. We acknowledge our employer, Stanford University, for support and access to journal articles. LAM would also like to acknowledge support from Schmidt Sciences, that funds her research at the intersection of machine learning and climate modeling.

References

- [Addison et al., 2022] Addison, H., Kendon, E., Ravuri, S., Aitchison, L., and Watson, P. A. (2022). Machine learning emulation of a local-scale UK climate model. *arXiv:2211.16116 [physics]*.
- [Adewoyin et al., 2021] Adewoyin, R. A., Dueben, P., Watson, P., He, Y., and Dutta, R. (2021). TRU-NET: a deep learning approach to high resolution prediction of rainfall. *Machine Learning*, 110(8):2035–2062.
- [Agrawal et al., 2019] Agrawal, S., Barrington, L., Bromberg, C., Burge, J., Gazen, C., and Hickey, J. (2019). Machine Learning for Precipitation Nowcasting from Radar Images. *arXiv:1912.12132 [cs, stat]*.
- [Akinsanola et al., 2018] Akinsanola, A., Ajayi, V., Adejare, A., Adeyeri, O., Gbode, I., Ogunjobi, K., Nikulin, G., and Abolude, A. (2018). Evaluation of rainfall simulations over West Africa in dynamically downscaled CMIP5 global circulation models. *Theoretical and Applied Climatology*, 132(1):437–450.
- [Anderson, 1982] Anderson, B. D. O. (1982). Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326.
- [Babausmail et al., 2021] Babausmail, H., Hou, R., Gnitou, G. T., and Ayugi, B. (2021). Novel statistical downscaling emulator for precipitation projections using deep Convolutional Autoencoder over Northern Africa. *Journal of Atmospheric and Solar-Terrestrial Physics*, 218:105614.

- [Balaji, 2021] Balaji, V. (2021). Climbing down Charney’s ladder: machine learning and the post-Dennard era of computational climate science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2194):20200085. Publisher: Royal Society.
- [Bassetti et al., 2023] Bassetti, S., Hutchinson, B., Tebaldi, C., and Kravitz, B. (2023). DiffESM: Conditional Emulation of Earth System Models with Diffusion Models. arXiv:2304.11699 [physics].
- [Baño-Medina et al., 2022] Baño-Medina, J., Manzanar, R., Cimadevilla, E., Fernández, J., González-Abad, J., Cofiño, A. S., and Gutiérrez, J. M. (2022). Downscaling multi-model climate projection ensembles with deep learning (DeepESD): contribution to CORDEX EUR-44. *Geoscientific Model Development*, 15(17):6747–6758. Publisher: Copernicus GmbH.
- [Bi et al., 2022] Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., and Tian, Q. (2022). Pangu-Weather: A 3D High-Resolution Model for Fast and Accurate Global Weather Forecast. arXiv:2211.02556 [physics].
- [Bihlo, 2021] Bihlo, A. (2021). A generative adversarial network approach to (ensemble) weather prediction. *Neural Networks*, 139:1–16.
- [Bischoff and Deck, 2023] Bischoff, T. and Deck, K. (2023). Unpaired Downscaling of Fluid Flows with Diffusion Bridges. arXiv:2305.01822 [physics].
- [Cachay et al., 2023] Cachay, S. R., Zhao, B., Joren, H., and Yu, R. (2023). DYffusion: A Dynamics-informed Diffusion Model for Spatiotemporal Forecasting. arXiv:2306.01984 [cs, stat].
- [Chan et al., 2024] Chan, M. A., Molina, M. J., and Metzler, C. A. (2024). Hyper-Diffusion: Estimating Epistemic and Aleatoric Uncertainty with a Single Model. arXiv:2402.03478 [cs].
- [Chen et al., 2021a] Chen, C., Hu, B., and Li, Y. (2021a). Easy-to-use spatial random-forest-based downscaling-calibration method for producing precipitation data with high resolution and high accuracy. *Hydrology and Earth System Sciences*, 25(11):5667–5682.
- [Chen et al., 2021b] Chen, D., Rojas, M., Samset, B., Cobb, K., Diongue Niang, A., Edwards, P., Emori, S., Faria, S., Hawkins, E., Hope, P., Huybrechts, P., Meinshausen, M., Mustafa, S., Plattner, G.-K., and Tréguier, A.-M. (2021b). Framing, Context, and Methods. In Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J., Maycock, T., Waterfield, T., Yelekçi, O., Yu, R., and Zhou, B., editors, *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 147–286. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA. Type: Book Section.
- [Chen et al., 2023] Chen, L., Zhong, X., Zhang, F., Cheng, Y., Xu, Y., Qi, Y., and Li, H. (2023). FuXi: a cascade machine learning forecasting system for 15-day global weather forecast. *npj Climate and Atmospheric Science*, 6(1):1–11. Publisher: Nature Publishing Group.
- [Cho et al., 2023] Cho, D., Im, J., and Jung, S. (2023). A new statistical downscaling approach for short-term forecasting of summer air temperatures through a fusion of deep learning and spatial interpolation. *Quarterly Journal of the Royal Meteorological Society*, n/a(n/a). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qj.4643>.
- [Fowler et al., 2007] Fowler, H. J., Blenkinsop, S., and Tebaldi, C. (2007). Linking climate change modelling to impacts studies: recent advances in downscaling techniques for hydrological modelling. *International Journal of Climatology*, 27(12):1547–1578. Publisher: John Wiley & Sons, Ltd.
- [Gentine et al., 2018] Gentine, P., Pritchard, M., Rasp, S., Reinaudi, G., and Yacalis, G. (2018). Could Machine Learning Break the Convection Parameterization Deadlock? *Geophysical Research Letters*, 45(11):5742–5751. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2018GL078202>.
- [Giorgi, 2019] Giorgi, F. (2019). Thirty Years of Regional Climate Modeling: Where Are We and Where Are We Going next? *Journal of Geophysical Research: Atmospheres*, 124(11):5696–5723. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2018JD030094>.
- [Giorgi and Jr, 2015] Giorgi, F. and Jr, W. J. G. (2015). Regional Dynamical Downscaling and the CORDEX Initiative. *Annual Review of Environment and Resources*, 40(Volume 40, 2015):467–490. Publisher: Annual Reviews.
- [Glasner et al., 2009] Glasner, D., Bagon, S., and Irani, M. (2009). Super-resolution from a single image. In *2009 IEEE 12th International Conference on Computer Vision*, pages 349–356. ISSN: 2380-7504.
- [Gneiting and Raftery, 2007] Gneiting, T. and Raftery, A. E. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378.
- [Groenke et al., 2020] Groenke, B., Madaus, L., and Monteleoni, C. (2020). ClimAlign: Unsupervised statistical downscaling of climate variables via normalizing flows. In *Proceedings of the 10th International Conference on Climate Informatics*, pages 60–66, virtual United Kingdom. ACM.

- [Gutowski et al., 2020] Gutowski, W. J., Ullrich, P. A., Hall, A., Leung, L. R., O’Brien, T. A., Patricola, C. M., Arritt, R. W., Bukovsky, M. S., Calvin, K. V., Feng, Z., Jones, A. D., Kooperman, G. J., Monier, E., Pritchard, M. S., Pryor, S. C., Qian, Y., Rhoades, A. M., Roberts, A. F., Sakaguchi, K., Urban, N., and Zarzycki, C. (2020). The Ongoing Need for High-Resolution Regional Climate Models: Process Understanding and Stakeholder Information. *Bulletin of the American Meteorological Society*, 101(5):E664–E683. Publisher: American Meteorological Society Section: Bulletin of the American Meteorological Society.
- [Harder et al., 2024] Harder, P., Hernandez-Garcia, A., Ramesh, V., Yang, Q., Sattigeri, P., Szwarcman, D., Watson, C., and Rolnick, D. (2024). Hard-Constrained Deep Learning for Climate Downscaling. arXiv:2208.05424 [physics].
- [Harder et al., 2022] Harder, P., Watson-Parris, D., Stier, P., Strassel, D., Gauger, N. R., and Keuper, J. (2022). Physics-informed learning of aerosol microphysics. *Environmental Data Science*, 1:e20.
- [Harris et al., 2022] Harris, L., McRae, A. T. T., Chantry, M., Dueben, P. D., and Palmer, T. N. (2022). A Generative Deep Learning Approach to Stochastic Downscaling of Precipitation Forecasts. *Journal of Advances in Modeling Earth Systems*, 14(10):e2022MS003120. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2022MS003120>.
- [Hasan Karaman and Akyürek, 2023] Hasan Karaman, C. and Akyürek, Z. (2023). Evaluation of near-surface air temperature reanalysis datasets and downscaling with machine learning based Random Forest method for complex terrain of Turkey. *Advances in Space Research*, 71(12):5256–5281.
- [Hawkins and Sutton, 2009] Hawkins, E. and Sutton, R. (2009). The Potential to Narrow Uncertainty in Regional Climate Predictions. *Bulletin of the American Meteorological Society*, 90(8):1095–1108.
- [Hersbach, 2000] Hersbach, H. (2000). Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems. *Weather and Forecasting*, 15(5):559–570. Publisher: American Meteorological Society Section: Weather and Forecasting.
- [Hersbach et al., 2020] Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J.-N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qj.3803>.
- [Ho et al., 2020] Ho, J., Jain, A., and Abbeel, P. (2020). Denoising Diffusion Probabilistic Models. arXiv:2006.11239 [cs, stat].
- [Hu et al., 2019] Hu, X., Naiel, M. A., Wong, A., Lamm, M., and Fieguth, P. (2019). RUNet: A Robust UNet Architecture for Image Super-Resolution. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 505–507, Long Beach, CA, USA. IEEE.
- [Huang et al., 2024] Huang, L., Gianinazzi, L., Yu, Y., Dueben, P. D., and Hoefler, T. (2024). DiffDA: a Diffusion model for weather-scale Data Assimilation. arXiv:2401.05932 [cs].
- [Jiang et al., 2021] Jiang, Y., Yang, K., Shao, C., Zhou, X., Zhao, L., Chen, Y., and Wu, H. (2021). A downscaling approach for constructing high-resolution precipitation dataset over the Tibetan Plateau from ERA5 reanalysis. *Atmospheric Research*, 256:105574.
- [Karakakis and Mehrkanoon, 2023] Karakakis, C. and Mehrkanoon, S. (2023). WF-UNet: Weather Data Fusion using 3D-UNet for Precipitation Nowcasting. *Procedia Computer Science*, 222:223–232.
- [Karras et al., 2022] Karras, T., Aittala, M., Aila, T., and Laine, S. (2022). Elucidating the Design Space of Diffusion-Based Generative Models. *Advances in Neural Information Processing Systems*, 35:26565–26577.
- [Kochkov et al., 2024] Kochkov, D., Yuval, J., Langmore, I., Norgaard, P., Smith, J., Mooers, G., Klöwer, M., Lottes, J., Rasp, S., Düben, P., Hatfield, S., Battaglia, P., Sanchez-Gonzalez, A., Willson, M., Brenner, M. P., and Hoyer, S. (2024). Neural General Circulation Models for Weather and Climate. arXiv:2311.07222 [physics].
- [Lam et al., 2023] Lam, R., Sanchez-Gonzalez, A., Willson, M., Wyrnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Meroze, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., and Battaglia, P. (2023). Learning skillful medium-range global weather forecasting. *Science*, 0(0):eadi2336. Publisher: American Association for the Advancement of Science.
- [Leinonen et al., 2021] Leinonen, J., Nerini, D., and Berne, A. (2021). Stochastic Super-Resolution for Downscaling Time-Evolving Atmospheric Fields With a Generative Adversarial Network. *IEEE Transactions on Geoscience and Remote Sensing*, 59(9):7211–7223. Conference Name: IEEE Transactions on Geoscience and Remote Sensing.
- [Li et al., 2023] Li, L., Carver, R., Lopez-Gomez, I., Sha, F., and Anderson, J. (2023). SEEDS: Emulation of Weather Forecast Ensembles with Diffusion Models. arXiv:2306.14066 [physics].

- [Mansfield et al., 2023] Mansfield, L. A., Gupta, A., Burnett, A. C., Green, B., Wilka, C., and Sheshadri, A. (2023). Updates on Model Hierarchies for Understanding and Simulating the Climate System: A Focus on Data-Informed Methods and Climate Change Impacts. *Journal of Advances in Modeling Earth Systems*, 15(10):e2023MS003715. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2023MS003715>.
- [Mansfield et al., 2020] Mansfield, L. A., Nowack, P. J., Kasoar, M., Everitt, R. G., Collins, W. J., and Voulgarakis, A. (2020). Predicting global patterns of long-term climate change from short-term simulations using machine learning. *npj Climate and Atmospheric Science*, 3(1):1–9. Number: 1 Publisher: Nature Publishing Group.
- [Maraun et al., 2017] Maraun, D., Shepherd, T. G., Widmann, M., Zappa, G., Walton, D., Gutiérrez, J. M., Hagemann, S., Richter, I., Soares, P. M. M., Hall, A., and Mearns, L. O. (2017). Towards process-informed bias correction of climate change simulations. *Nature Climate Change*, 7(11):764–773. Publisher: Nature Publishing Group.
- [Maraun and Widmann, 2018a] Maraun, D. and Widmann, M. (2018a). Perfect Prognosis. In Maraun, D. and Widmann, M., editors, *Statistical Downscaling and Bias Correction for Climate Research*, pages 141–169. Cambridge University Press, Cambridge.
- [Maraun and Widmann, 2018b] Maraun, D. and Widmann, M. (2018b). *Statistical Downscaling and Bias Correction for Climate Research*. Cambridge University Press. Google-Books-ID: AMhJDwAAQBAJ.
- [Mardani et al., 2023] Mardani, M., Brenowitz, N., Cohen, Y., Pathak, J., Chen, C.-Y., Liu, C.-C., Vahdat, A., Kashinath, K., Kautz, J., and Pritchard, M. (2023). Generative Residual Diffusion Modeling for Km-scale Atmospheric Downscaling. arXiv:2309.15214 [physics].
- [Medrano et al., 2023] Medrano, S. C., Satgé, F., Molina-Carpio, J., Zolá, R. P., and Bonnet, M.-P. (2023). Downscaling Daily Satellite-Based Precipitation Estimates Using MODIS Cloud Optical and Microphysical Properties in Machine-Learning Models. *Atmosphere*, 14(9):1349. Number: 9 Publisher: Multidisciplinary Digital Publishing Institute.
- [Molina et al., 2023] Molina, M. J., O’Brien, T. A., Anderson, G., Ashfaq, M., Bennett, K. E., Collins, W. D., Dagon, K., Restrepo, J. M., and Ullrich, P. A. (2023). A Review of Recent and Emerging Machine Learning Applications for Climate Variability and Weather Phenomena. *Artificial Intelligence for the Earth Systems*, 2(4). Publisher: American Meteorological Society Section: Artificial Intelligence for the Earth Systems.
- [Nath et al., 2024] Nath, P., Shukla, P., Wang, S., and Quilodrán-Casas, C. (2024). Forecasting Tropical Cyclones with Cascaded Diffusion Models. arXiv:2310.01690 [physics].
- [Nguyen et al., 2023] Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J. K., and Grover, A. (2023). ClimaX: A foundation model for weather and climate. arXiv:2301.10343 [cs].
- [Nichol and Dhariwal, 2021] Nichol, A. and Dhariwal, P. (2021). Improved Denoising Diffusion Probabilistic Models. arXiv:2102.09672 [cs, stat].
- [Nichol et al., 2022] Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., and Chen, M. (2022). GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models.
- [O’Gorman and Dwyer, 2018] O’Gorman, P. A. and Dwyer, J. G. (2018). Using Machine Learning to Parameterize Moist Convection: Potential for Modeling of Climate, Climate Change, and Extreme Events. *Journal of Advances in Modeling Earth Systems*, 10(10):2548–2563. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2018MS001351>.
- [Oyama et al., 2023] Oyama, N., Ishizaki, N. N., Koide, S., and Yoshida, H. (2023). Deep generative model super-resolves spatially correlated multiregional climate data. *Scientific Reports*, 13(1):5992. Number: 1 Publisher: Nature Publishing Group.
- [Pathak et al., 2022] Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., Hassanzadeh, P., Kashinath, K., and Anandkumar, A. (2022). FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators. arXiv:2202.11214 [physics].
- [Pour et al., 2018] Pour, S. H., Shahid, S., Chung, E.-S., and Wang, X.-J. (2018). Model output statistics downscaling using support vector machine for the projection of spatial and temporal changes in rainfall of Bangladesh. *Atmospheric Research*, 213:149–162.
- [Price and Rasp, 2022] Price, I. and Rasp, S. (2022). Increasing the accuracy and resolution of precipitation forecasts using deep generative models. arXiv:2203.12297 [cs, stat].
- [Price et al., 2023] Price, I., Sanchez-Gonzalez, A., Alet, F., Ewalds, T., El-Kadi, A., Stott, J., Mohamed, S., Battaglia, P., Lam, R., and Willson, M. (2023). GenCast: Diffusion-based ensemble forecasting for medium-range weather. arXiv:2312.15796 [physics].

- [Ramesh et al., 2022] Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022). Hierarchical Text-Conditional Image Generation with CLIP Latents. *arXiv:2204.06125* [cs].
- [Rolnick et al., 2023] Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A. S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., Luccioni, A. S., Maharaj, T., Sherwin, E. D., Mukkavilli, S. K., Kording, K. P., Gomes, C. P., Ng, A. Y., Hassabis, D., Platt, J. C., Creutzig, F., Chayes, J., and Bengio, Y. (2023). Tackling Climate Change with Machine Learning. *ACM Computing Surveys*, 55(2):1–96.
- [Rombach et al., 2022] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-Resolution Image Synthesis with Latent Diffusion Models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, New Orleans, LA, USA. IEEE.
- [Ronneberger et al., 2015] Ronneberger, O., Fischer, P., and Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. In Navab, N., Hornegger, J., Wells, W. M., and Frangi, A. F., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351, pages 234–241. Springer International Publishing, Cham. Series Title: Lecture Notes in Computer Science.
- [Saharia et al., 2022] Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S. K. S., Ayan, B. K., Mahdavi, S. S., Lopes, R. G., Salimans, T., Ho, J., Fleet, D. J., and Norouzi, M. (2022). Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *arXiv:2205.11487* [cs].
- [Schneider et al., 2023] Schneider, T., Behera, S., Boccaletti, G., Deser, C., Emanuel, K., Ferrari, R., Leung, L. R., Lin, N., Müller, T., Navarra, A., Ndiaye, O., Stuart, A., Tribbia, J., and Yamagata, T. (2023). Harnessing AI and computing to advance climate modelling and prediction. *Nature Climate Change*, 13(9):887–889. Number: 9 Publisher: Nature Publishing Group.
- [Schubert, 1998] Schubert, S. (1998). Downscaling local extreme temperature changes in south-eastern Australia from the CSIRO Mark2 GCM. *International Journal of Climatology*, 18(13):1419–1438. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/%28SICI%291097-0088%2819981115%2918%3A13%3C1419%3A%3AAID-JOC314%3E3.0.CO%3B2-Z](https://onlinelibrary.wiley.com/doi/pdf/10.1002/%28SICI%291097-0088%2819981115%2918%3A13%3C1419%3A%3AAID-JOC314%3E3.0.CO%3B2-Z).
- [Sohl-Dickstein et al., 2015] Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N., and Ganguli, S. (2015). Deep Unsupervised Learning using Nonequilibrium Thermodynamics.
- [Song et al., 2021] Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2021). Score-based Generative Modeling through Stochastic Differential Equations.
- [Sunyer et al., 2012] Sunyer, M., Madsen, H., and Ang, P. (2012). A comparison of different regional climate models and statistical downscaling methods for extreme rainfall estimation under climate change. *Atmospheric Research*, 103:119–128.
- [Tapiador et al., 2020] Tapiador, F. J., Navarro, A., Moreno, R., Sánchez, J. L., and García-Ortega, E. (2020). Regional climate models: 30 years of dynamical downscaling. *Atmospheric Research*, 235:104785.
- [Thrasher et al., 2022] Thrasher, B., Wang, W., Michaelis, A., Melton, F., Lee, T., and Nemani, R. (2022). NASA Global Daily Downscaled Projections, CMIP6. *Scientific Data*, 9(1):262. Number: 1 Publisher: Nature Publishing Group.
- [Vandal et al., 2019] Vandal, T., Kodra, E., and Ganguly, A. R. (2019). Intercomparison of machine learning methods for statistical downscaling: the case of daily and extreme precipitation. *Theoretical and Applied Climatology*, 137(1):557–570.
- [Vandal et al., 2017] Vandal, T., Kodra, E., Ganguly, S., Michaelis, A., Nemani, R., and Ganguly, A. R. (2017). DeepSD: Generating High Resolution Climate Change Projections through Single Image Super-Resolution. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17*, pages 1663–1672, New York, NY, USA. Association for Computing Machinery.
- [Vincent, 2011] Vincent, P. (2011). A Connection Between Score Matching and Denoising Autoencoders. *Neural Computation*, 23(7):1661–1674. Conference Name: Neural Computation.
- [Wan et al., 2023] Wan, Z. Y., Baptista, R., Chen, Y.-f., Anderson, J., Boral, A., Sha, F., and Zepeda-Núñez, L. (2023). Debias Coarsely, Sample Conditionally: Statistical Downscaling through Optimal Transport and Probabilistic Diffusion Models. *arXiv:2305.15618* [physics].
- [Wang et al., 2021] Wang, J., Liu, Z., Foster, I., Chang, W., Kettimuthu, R., and Kotamarthi, V. R. (2021). Fast and accurate learned multiresolution dynamical downscaling for precipitation. *Geoscientific Model Development*, 14(10):6355–6372. Publisher: Copernicus GmbH.
- [Wang et al., 2019] Wang, Y., Wang, L., Wang, H., and Li, P. (2019). End-to-End Image Super-Resolution via Deep and Shallow Convolutional Networks. *IEEE Access*, 7:31959–31970.

- [Watson, 2022] Watson, P. A. G. (2022). Machine learning applications for weather and climate need greater focus on extremes. *Environmental Research Letters*, 17(11):111004. Publisher: IOP Publishing.
- [Watt-Meyer et al., 2023] Watt-Meyer, O., Dresdner, G., McGibbon, J., Clark, S. K., Henn, B., Duncan, J., Brenowitz, N. D., Kashinath, K., Pritchard, M. S., Bonev, B., Peters, M. E., and Bretherton, C. S. (2023). ACE: A fast, skillful learned global atmospheric model for climate prediction. arXiv:2310.02074 [physics].
- [Wilby et al., 1998] Wilby, R. L., Wigley, T. M. L., Conway, D., Jones, P. D., Hewitson, B. C., Main, J., and Wilks, D. S. (1998). Statistical downscaling of general circulation model output: A comparison of methods. *Water Resources Research*, 34(11):2995–3008. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/98WR02577>.
- [Xu et al., 2020] Xu, M., Liu, Q., Sha, D., Yu, M., Duffy, D. Q., Putman, W. M., Carroll, M., Lee, T., and Yang, C. (2020). PreciPatch: A Dictionary-based Precipitation Downscaling Method. *Remote Sensing*, 12(6):1030. Number: 6 Publisher: Multidisciplinary Digital Publishing Institute.
- [Yang et al., 2023] Yang, Q., Hernandez-Garcia, A., Harder, P., Ramesh, V., Sattigeri, P., Szwarcman, D., Watson, C. D., and Rolnick, D. (2023). Fourier Neural Operators for Arbitrary Resolution Climate Data Downscaling. arXiv:2305.14452 [physics].
- [Zhang et al., 2023] Zhang, C., Zhang, C., Zhang, M., and Kweon, I. S. (2023). Text-to-image Diffusion Models in Generative AI: A Survey. arXiv:2303.07909 [cs].