

# On the Effectiveness of Neural Operators at Zero-Shot Weather Downscaling

Saumya Sinha\* Brandon Benton, Patrick Emami

National Renewable Energy Laboratory  
Golden, CO, 80401

## Abstract

Machine learning (ML) methods have shown great potential for weather downscaling. These data-driven approaches provide a more efficient alternative for producing high-resolution weather datasets and forecasts compared to physics-based numerical simulations. Neural operators, which learn solution operators for a family of partial differential equations (PDEs), have shown great success in scientific ML applications involving physics-driven datasets. Neural operators are grid-resolution-invariant and are often evaluated on higher grid resolutions than they are trained on, i.e., zero-shot super-resolution. Given their promising zero-shot super-resolution performance on dynamical systems emulation, we present a critical investigation of their zero-shot weather downscaling capabilities, which is when models are tasked with producing high-resolution outputs using higher upsampling factors than are seen during training. To this end, we create two realistic downscaling experiments with challenging upsampling factors (e.g., 8x and 15x) across data from different simulations: the European Centre for Medium-Range Weather Forecasts Reanalysis version 5 (ERA5) and the Wind Integration National Dataset Toolkit (WTK). While neural operator-based downscaling models perform better than interpolation and a simple convolutional baseline, we show the surprising performance of an approach that combines a powerful transformer-based model with parameter-free interpolation at zero-shot weather downscaling. We find that this Swin-Transformer-based approach mostly outperforms models with neural operator layers in terms of average error metrics, whereas an Enhanced Super-Resolution Generative Adversarial Network (ESRGAN)-based approach is better than most models in terms of capturing the physics of the ground truth data. We suggest their use in future work as strong baselines.

## 1 Introduction

Downscaling techniques are used to obtain high-resolution (HR) data from their coarse low-resolution (LR) counterparts. The HR data often includes finer details of physical phenomena than the LR data in complex earth systems such as weather. Downscaling provides insights into climate change and its effects, e.g., the small-scale features and detailed information are crucial for analyzing extreme weather

events that can only be observed at high resolutions. Downscaling can also help upsample medium-range weather forecasts (Jiang et al. 2023) and is useful for optimal grid planning and management of renewable resources such as wind energy (Buster et al. 2024; Stengel et al. 2020; Benton et al. 2024; Kurinchi-Vendhan et al. 2021; Ren et al. 2023).

Although earth system processes such as weather and climate can be approximately expressed as systems of partial differential equations (PDEs), solving these models numerically at sufficiently high resolutions for many practical applications is computationally infeasible. Data-driven downscaling approaches, which promise better efficiency than numerical physics-based solvers, have shown great potential (Ren et al. 2023; Kurinchi-Vendhan et al. 2021; Yang et al. 2023; Jiang et al. 2023; Mikhaylov et al. 2024; Buster et al. 2024). While statistical downscaling methods (Pierce, Cayan, and Thrasher 2014; Wood et al. 2004; Kaczmarek, Isham, and Onof 2014) have been used traditionally, deep learning techniques, in particular, have gained attention due to their ability to efficiently learn complex relationships from large amounts of data. Moreover, the rapid advancement of deep learning in the computer vision field of super-resolution has been adapted with success for downscaling in the atmospheric sciences (Ren et al. 2023; Chen et al. 2022; Kurinchi-Vendhan et al. 2021).

Neural operators (Kovachki et al. 2023) have recently been applied to many scientific machine learning (ML) tasks involving the emulation of physical systems. Unlike traditional neural networks, neural operators approximate a mapping between infinite-dimensional function spaces. For example, neural operators can be used to learn the solution operator for an entire family of PDEs, such as Navier Stokes and Darcy flow (Li et al. 2021). For this application, neural operators are much more efficient than traditional numerical solvers which run on finely discretized grids. Once trained, neural operators are fast to solve any new instance of the PDE (Kovachki et al. 2023; Li et al. 2021). Neural operators have demonstrated the ability to perform *zero-shot super-resolution* (Li et al. 2021; Rahman, Ross, and Azizzadehsheli 2023; Raonic et al. 2023) when emulating physical systems. That is, they can be trained on coarse resolution data and then tested “zero-shot” on a previously unseen fine discretization of a grid.

Neural operator’s ability to perform zero-shot super-

\*Corresponding author: saumya.sinha@nrel.gov.  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

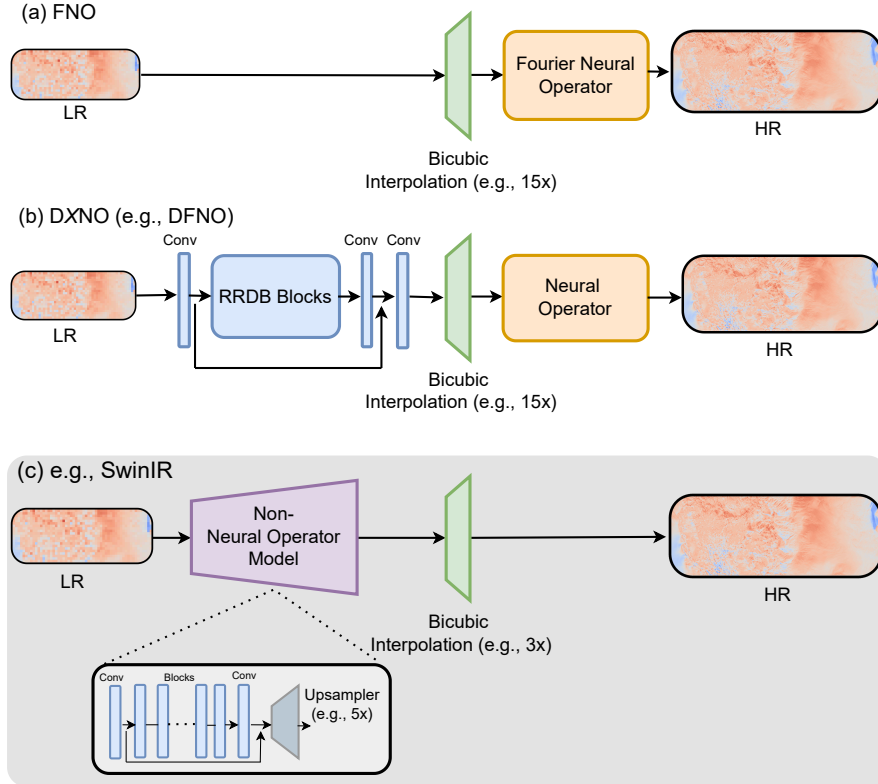


Figure 1: Overview of the neural operator and non-neural-operator zero-shot weather downscaling approaches. We show 5x to 15x zero-shot downscaling as an example. (a,b) For neural operators, the interpolation scale factor is the same as the upsampling factor, e.g., the bicubic interpolation layer upsamples to 5x during training and 15x during evaluation. (c) For regular neural networks (e.g., SwinIR), the model is trained to output at 5x (e.g., using a learnable upsampler such as sub-pixel convolution). At test time, the model generates a 5x output which is then interpolated 3x more to produce the final 15x HR output.

resolution raises the question of whether they can be applied to perform *zero-shot weather downscaling*. Currently, downscaling pipelines train models to map an LR input to an HR output at an upsampling factor (the ratio of the size of the HR grid to the LR grid), and they are evaluated on generating downsampled outputs with the upsampling factor seen during training. In zero-shot weather downscaling, a model is trained with a (small) upsampling factor and then the same model is tasked with producing an HR output at an unseen and higher upsampling factor at test time. The success of neural operators at zero-shot super-resolution when emulating dynamical systems suggests they hold promise for this task as well.

We design challenging experiments to investigate whether neural-operator-based models have an enhanced ability to perform zero-shot weather downscaling. We adapt and expand the learning framework for applying neural operators to this setting as proposed in Yang et al. (2023). One of the key difficulties of zero-shot weather downscaling is generalizing to an upsampling factor where the data at the finest spatial scales contains physical phenomena unseen at the highest resolutions seen during training. We design experiments

that aim to test this setting by using large upsampling factors (e.g., 8x and 15x) and high target resolutions (e.g., 2km x 2km wind speed data).

Overall, our work investigates the zero-shot downscaling potential of neural operators. To summarize, our contributions are:

1. We provide a comparative analysis based on two challenging weather downscaling problems, between various neural operator and non-neural-operator methods with large upsampling factors (e.g., 8x and 15x) and fine grid resolutions (e.g., 2km x 2km wind speed).
2. We examine whether neural operator layers provide unique advantages when testing downscaling models on upsampling factors higher than those seen during training, i.e., *zero-shot downscaling*. Our results instead show the surprising success of an approach that combines a powerful transformer-based model with a parameter-free interpolation step at zero-shot weather downscaling.
3. We find that this Swin-Transformer-based approach mostly outperforms all neural operator models in terms of average error metrics, whereas an Enhanced Super-

Resolution Generative Adversarial Network (ESRGAN)-based approach is better than most models in capturing the physics of the system, and suggests their use in future work as strong baselines. However, these approaches still don't capture variations at smaller spatial scales well, including the physical characteristics of turbulence in the HR data. This suggests a potential for improvement in transformer or GAN-based methods and neural-operator-based methods for zero-shot weather downscaling.

## 2 Related Work

**Weather downscaling with deep learning** Deep learning models have recently shown promise at weather downscaling tasks such as precipitation downscaling (Watson et al. 2020; Chaudhuri and Robertson 2020; Harris et al. 2022). These end-to-end differentiable approaches directly learn to map low-resolution inputs to high-resolution outputs. The most popular approaches such as the Super-Resolution Convolutional Neural Network (SRCNN) (Dong et al. 2015) are based on architectures introduced by the computer vision community for super-resolution (Wang, Chen, and Hoi 2020). These models are used as key baselines in our experiments, thus, we describe them in more detail in Section 4. Works (Buster et al. 2024; Stengel et al. 2020; Tran et al. 2020; Harilal et al. 2022) downscale renewable energy datasets such as wind and solar, as energy system planning depends on high-resolution estimates of these resources. Buster et al. (2024) use a custom GAN by Stengel et al. (2020) (which is similar to ESRGAN (Wang et al. 2018)) trained on Global Climate Models (GCMs) projections to generate high-resolution spatial and temporal features capturing small-scale details otherwise lost in the coarse GCM models. Benton et al. (2024) use the custom GAN (Stengel et al. 2020) for spatiotemporal downscaling of wind data to learn a mapping from low-resolution European Centre for Medium-Range Weather Forecasts Reanalysis version 5 (ERA5) (Hersbach et al. 2020) to high-resolution Wind Integration National Dataset Toolkit (WTK) (Draxl et al. 2015). Through this work, we also perform downscaling of wind data using ERA5 and WTK datasets, but to investigate the utility of neural operator models for the weather downscaling task with a specific emphasis on zero-shot downscaling.

**Related benchmarks for weather downscaling** SuperBench (Ren et al. 2023) introduces super-resolution datasets and benchmarks for scientific applications such as fluid flow, cosmology, and weather downscaling. Their work compares various deep learning methods and analyzes the physics-preserving properties of these models. RainNet (Chen et al. 2022) is one of the first large-scale datasets for spatial precipitation downscaling, spanning over 17 years and covering important meteorological phenomena. Their work also presents an extensive evaluation of many deep learning models. WiSoSuper (Kurinchi-Vendhan et al. 2021) is a benchmark for wind and solar super-resolution. The dataset released by WiSoSuper is based on the National Renewable Energy Laboratory's (NREL's) WTK and National Solar Radiation Database (NSRDB) (Sengupta

et al. 2018) datasets. They compare generative models introduced in Stengel et al. (2020) with other GAN and convolutional neural network (CNN)-based models. In contrast to these benchmarking efforts, our work benchmarks models for weather downscaling but with a focus on neural operator models and zero-shot weather downscaling. In a concurrent work focusing on climate downscaling (Prasad et al. 2024), CNNs, transformers (Alexey 2020), and neural-operator-based models are compared in terms of their ability to pretrain on diverse climate datasets so as to learn transferrable representations across multiple climate variables and spatial regions.

## 3 Background

**Neural operators** Operator learning models such as neural operators (Kovachki et al. 2023) are composed of layers that learn mappings between infinite-dimensional function spaces. In doing so, they approximately learn a continuous operator, which can be realized at any arbitrary grid discretization of the input and output. Thus, neural operators do not depend on the discretization of the grid they are trained on, and we expect them to generalize to grid resolutions different than the ones they are trained on. Li et al. (2021) introduced Fourier neural operators (FNOs), expressing neural operators as a combination of linear integral operators to incorporate the non-local properties of the solution operator with Fourier Transform and non-linear local activation functions (Kovachki et al. 2023), which helps to model non-linear systems and their high-frequency modes. They show improved performance over convolution-based models for complex non-linear PDEs such as the Navier-Stokes equation. With the Fourier layers, the parameters are learned in the Fourier domain, which enables FNOs to be invariant to the grid discretization or resolution. Since neural operators such as FNOs learn resolution-invariant approximations of continuous operators, we aim to understand whether this provides advantages for *zero-shot* weather downscaling.

Yang et al. (2023) adapt Fourier neural operators (FNOs) (Li et al. 2021) to perform downscaling on ERA5 and PDE data. Their proposed model, which they refer to as DFNO, outperforms CNN and GAN-based models. They also evaluate zero-shot downscaling on unseen upsampling factors to observe the model's zero-shot generalization potential. We adapt and expand this downscaling pipeline in our benchmarking study. Our work differs from this paper as we investigate higher upsampling factors (8x and 15x) for training and zero-shot evaluations as opposed to 2x and 4x in their work. We also create a realistic set of experiments that includes LR and HR data sourced from different simulations (ERA5 to WTK downscaling, as described in Section 4.2) and compare various neural operator approaches against strong baselines including powerful transformers.

**Weather downscaling** In weather downscaling, we are given a snapshot of LR weather data (e.g., an image) with a goal of upsampling this data to a higher target resolution. Mathematically, in the standard downscaling problem we have the LR input grid  $\mathbf{x} \in \mathbb{R}^{h \times w \times c}$ , and a target, HR

output,  $\mathbf{y} \in \mathbb{R}^{h' \times w' \times c}$ , where  $h, w \in \mathbb{N}$ ,  $c$  is the number of atmospheric variables, and  $h \times w$  is a lower resolution than  $h' \times w'$ . Deep-learning-based downscaling techniques introduced in Section 2 learn an approximation  $f$  between two finite-dimensional vector spaces  $f : \mathbb{R}^{h \times w \times c} \rightarrow \mathbb{R}^{h' \times w' \times c}$ . We refer to the setting where models are trained and tested on the same upsampling factor as *standard* weather downscaling. In this work, we restrict our focus to only static downscaling problems, i.e., each snapshot represents a single instant in time.

**Zero-shot weather downscaling** In our work, we wish to evaluate the extent to which downscaling models built with resolution-invariant neural operator layers generalize when tested on previously unseen, higher upsampling factors compared to approaches without such layers. The simplest way to obtain an HR image at any arbitrarily fine discretization is a non-learned interpolation scheme such as bicubic interpolation.

We are looking into neural-operator-based downscaling models that learn a mapping  $\mathcal{G}^\dagger : \mathbb{R}^{h \times w \times c} \rightarrow \mathcal{U}$  from  $\mathbf{x} \in \mathbb{R}^{h \times w \times c}$  to a function  $u \in \mathcal{U}$ . We aim to obtain HR outputs  $\mathbf{y} \in \mathbb{R}^{h' \times w' \times c}$  from a discretization of  $u$ , where  $\mathcal{U}$  is an infinite-dimensional function space (Li et al. 2021; Yang et al. 2023). A neural-operator-based downscaling framework (Yang et al. 2023) (Figure 1b) learns a parametric approximation of a mapping from the finite-dimensional LR input space to the infinite-dimensional space,  $G_\theta(\mathbf{x}) : \mathbb{R}^{h \times w \times c} \rightarrow \mathcal{U}$ , as an approximation of  $\mathcal{G}^\dagger$  such that  $G_\theta(\mathbf{x}) := \mathcal{F}_\theta(T^{-1}(f_\theta(\mathbf{x})))$ , with  $\theta$  as the parameters of the model. It is comprised of (a) neural network layers that first learn to map LR inputs to an embedding vector,  $f_\theta : \mathbb{R}^{h \times w \times c} \rightarrow \mathbb{R}^d$ , (b) a discretization inversion operator that converts the vector to a function ( $e \in \mathcal{E}$ ) with  $T^{-1} : \mathbb{R}^d \rightarrow \mathcal{E}(D; \mathbb{R}^{d_e})$ , and (c) neural operator layers  $\mathcal{F}_\theta : \mathcal{E} \rightarrow \mathcal{U}$  that learn to map the function to another function, which can be discretized to produce the HR output  $\mathbf{y} \in \mathbb{R}^{h' \times w' \times c}$ . We refer to these approaches as Downscaling NO models (e.g. DFNO (Yang et al. 2023)). In order to use vanilla FNOs without resolution-dependent neural network layers (a) (as seen in Figure 1(a)), we learn  $G_\theta(\mathbf{x}) := \mathcal{F}_\theta(T^{-1}(\mathbf{x}))$ . Several improvements have since been proposed over FNOs (Guibas et al. 2021; Rahman, Ross, and Aizzadenesheli 2023; Raonic et al. 2023) which we include in our downscaling study and describe in further detail later (Section 4).

## 4 Methodology

We use two experimental setups to compare the performance of neural operators and non-neural-operator-based methods at both standard and zero-shot weather downscaling. First, we downscale ERA5 data, where we learn a mapping from coarsened LR ERA5 to HR ERA5. In our second experiment, we downscale from LR ERA5 to HR WTK. We expect the second task to be more challenging as it presents a more realistic downscaling scenario where the LR inputs belong to a different simulation than the HR data. Thus, we do not assume the LR is a coarsened version of the HR (Ren et al. 2023; Benton et al. 2024).

**Downscaling Neural Operator models** We compare (vanilla) FNO with Downscaling FNO (DFNO), Downscaling U-shaped Neural Operator (Rahman, Ross, and Aizzadenesheli 2023) (DUNO), Downscaling Convolutional Neural Operator (Raonic et al. 2023) (DCNO), and Downscaling Adaptive Fourier Neural Operator (Guibas et al. 2021) (DAFNO). The Downscaling (D) models are based on Yang et al. (2023) (as described in Section 3) with FNO, UNO, CNO, and AFNO as the neural operator layers in the modeling framework. We show details of this model in Figure 1(b). The low-resolution (LR) image first passes through a set of Residual-in-Residual Dense Block (RRDB) blocks, where an RRDB block is composed of multiple levels of residual and dense networks as introduced in ESRGAN (Wang et al. 2018). Then, the embedding is interpolated corresponding to the upsampling factor using bicubic interpolation to obtain a high-resolution output. Finally, this goes through neural operator layers to produce the final downsampled HR image. We can think of this last stage as a post-processing step over the features extracted by the RRDB layers followed by the interpolation. All the Downscaling (D) models are trained with the mean-squared-error (MSE) loss. To perform zero-shot downscaling with either the FNO or DXNO (e.g. DFNO) models at the test time on higher upsampling factors than the ones seen during the training, we use the interpolation layer to increase the resolution by the corresponding higher upsampling factor. Figure 1(a) shows the vanilla FNO model as used in our downscaling pipeline where FNO is post-processing the interpolation output. Further details on the four neural operator models are provided in Appendix A, and hyperparameter and model training details are presented in Appendix B.2.

**Baseline models** We compare all the neural operator models with four baselines: (1) bicubic interpolation, (2) SRCNN, (3) ESRGAN, (4) EDSR, and (5) SwinIR. Super-Resolution Convolutional Neural Network or SRCNN (Dong et al. 2015) is the first CNN-based model to perform single image or spatial super-resolution. SRCNN first upsamples the LR input with bicubic interpolation followed by lightweight CNN layers to obtain the HR image. Enhanced Deep Super-Resolution Network (EDSR) (Lim et al. 2017) introduces deep residual CNN networks to do super-resolution where the CNN layers are followed by an upsampling block performing sub-pixel convolution (Shi et al. 2016). Enhanced Super-Resolution Generative Adversarial Network (ESRGAN) (Wang et al. 2018) is a GAN-based architecture where the generator is composed of many Residual-in-Residual Dense Block (RRDB) blocks. The model is trained with pixel and perceptual loss along with adversarial loss, the perceptual loss minimizes the errors in the feature space and helps improve the visual quality of the generated super-resolved image (Ledig et al. 2017). Swin Transformer for Image Restoration (SwinIR) (Liang et al. 2021) has the advantages of both the CNN and Swin transformer (Liu et al. 2021) layers. It captures long-range dependencies and learns robust features to improve super-resolution performance with the residual Swin Transformer



blocks (RSTB) which is composed of many Swin Transformer layers stacked together with residual connections.

We refer to the implementation of these models as released in the SuperBench (Ren et al. 2023) work, hyperparameter details are added in Appendix B.2. Unlike neural operators, these models have architectures that expect their inputs and outputs to have the same grid resolution at both training and test time. Thus, we add a bicubic interpolation module on the output obtained from these models to produce outputs at higher upsampling factors than seen during training for our zero-shot downscaling experiments (Figure 1(c)).

See Appendix B.3 comparing the parameter count of all the neural operator and non-neural-operator-based models. Notably, we tested the vanilla FNO downscaling framework (Figure 1(a)), by moving the bicubic interpolation module after the FNO layers (as in baseline models, Figure 1(c)), which led to worse performance compared to our proposed pipeline with interpolation before the FNO layers.

#### 4.1 ERA5 to ERA5 downscaling

For our first experiment, we downscale coarsened LR ERA5 to HR ERA5. We use the ERA5 dataset for the entire globe at 25-km spatial resolution. We compare all models using two downscaling paradigms:

1. Standard downscaling: We train and test all the neural operators and baseline models with *the same* upsampling factor of 8x. An upsampling factor of 8x maps LR images of size 90x180 to HR outputs of size 720x1440.
2. Zero-shot downscaling: We first train all the models with an upsampling factor of 4x. Then, during testing, we observe their ability to produce outputs at a *higher* 8x upsampling factor.

**Dataset details** We use the European Centre for Medium-Range Weather Forecasts Reanalysis version (ERA5) dataset released as a part of the SuperBench (Ren et al. 2023) paper. The data is at a 0.25-degree (25km) grid resolution over the globe, i.e., each image has size 720x1440. We have three atmospheric variables (three channels): (1) wind speed,  $\sqrt{u^2 + v^2}$ ,  $u$  and  $v$  being the two components of wind velocity at 10m from the surface, (2) temperature at 2m from the surface, and (3) total column water vapor. This ERA5 dataset consists of image snapshots sampled at 24-hour intervals over an eight year period. Years 2008, 2010, 2011, and 2013 are used for training while 2012 and 2007 are reserved as a validation set for tuning hyperparameters. The years 2014 and 2015 are set aside for testing. Following Ren et al. (2023), we extract eight patches of size 64x64 (for the zero-shot downscaling) from each image to obtain HR images for training. The LR images are created by coarsening the HR patches with bicubic interpolation. We normalize each channel separately with a mean and standard deviation before training. Appendix B.1 has more details.

#### 4.2 ERA5 to WTK downscaling

For this second experiment, we focus on downscaling from LR ERA5 to HR WTK. It should be noted that the LR data in this setup is not obtained by coarsening the HR data but comes from another simulation. We include this experiment

to observe the performance of neural operators in a challenging and more realistic setup; for e.g., where ERA5 serves as boundary conditions for dynamical downscaling with a Numerical Weather Prediction model (Benton et al. 2024). The HR WTK dataset is available over two regions in the US at a 2-km resolution (Benton et al. 2024) and we use two variables: the  $u$  and  $v$  components of the wind velocity at 10m from the surface for this task. We perform the following experiments with the ERA5 to WTK downscaling setup:

1. Standard downscaling: We train and test all the neural operators and the baseline models with an upsampling factor of 5x, to go from 30-km to 6-km. The LR and the HR sizes are 53x53 and 265x265 for one region and 40x106 and 200x530 for the other region.
2. Zero-shot downscaling: For the zero-shot setup, we use the models trained with the 5x upsampling and evaluate them on an upsampling factor of 15, going from 30-km to 2-km. While the LR sizes are the same as above, the HR sizes for the zero-shot case are 795x795 for one region and 600x1590 for the second region.

**Dataset details** We use the National Renewable Energy Laboratory’s (NREL’s) Wind Integration National Dataset Toolkit (Draxl et al. 2015) (WTK) as the ground truth dataset. WTK has a spatial resolution of 2-km and a temporal resolution of 1-hour. We get LR images from the ERA5 dataset (Hersbach et al. 2020), available at 30-km ( $\sim 0.28$  degree) spatial resolution and 1-hour temporal resolution. This data has two channels, the  $u$  and  $v$  components of the wind velocity at 10m from the surface. We have paired ERA5 and WTK datasets over two regions in the US, with image sizes  $\sim 800 \times 800$  and  $\sim 600 \times 1600$  respectively (see (Benton et al. 2024) for more details).

Models are trained to map from coarse 30-km ERA5 to fine 6-km WTK with a 5x upsampling factor. This HR data is created by coarsening the WTK grid from 2-km to 6-km resolution. We realign, i.e. regrid 30-km ERA5 to the 6-km WTK coarsened grid using inverse distance weighted interpolation. For zero-shot experiments, we map the 30-km low-resolution ERA5 to the original 2-km resolution WTK (a 15x upsampling) for evaluation. The year 2007 is split 80/20 between the training and validation. We keep the year 2010 for testing. While training, we ensure that every batch has an equal number of LR and HR pairs from both regions. We extract patches of size 160x160 from the WTK image tiles to obtain the HR images and the corresponding coarse patches (32x32) from ERA5 image tiles as the LR images for training (Benton et al. 2024). This patch size is a hyperparameter tuned on the validation set. We normalize each channel separately with a mean and standard deviation before training. See Appendix B.1 for more details.

#### 4.3 Evaluation metrics

Our study quantitatively analyzes model performance using the following: **(1) Error metrics:** We use four pixel-level error measures: mean-squared-error (MSE), mean-absolute-error (MAE),  $L_\infty$  norm (IN), and the Peak signal-to-noise ratio (PSNR (Ren et al. 2023)). IN is the maximum pixel error between two images and it informs us about the tails

of the pixel error distribution (Ren et al. 2023). **(2) Energy spectrum:** We plot the kinetic energy spectrum (Kolmogorov 1991) for each model which shows a distribution of energy across various wavenumbers (Buster et al. 2024; Stengel et al. 2020; Kurinchi-Vendhan et al. 2021; Benton et al. 2024). These are normalized kinetic energy plots, with wavenumbers measured relative to the domain (or the spatial region) size. We compare each of the models with the energy curve for the ground truth HR. These plots describe how well the models capture physically realistic variations at smaller spatial scales in their downsampled outputs, for example, providing information about the physical characteristics of the turbulence of wind flow captured in model outputs.

## 5 Evaluation

### 5.1 ERA5 to ERA5 downscaling

**Error Metrics** We show the results for the ERA5 to ERA5 standard downscaling experiments and zero-shot experiments in Table 1. Standard downscaling compares models trained and evaluated with an upsampling factor of 8x. We observe that SwinIR outperforms every other model in terms of MSE, MAE, IN, and PSNR. DCNO is a close second and the best-performing neural operator model. The DFNO model shows improved results over the vanilla FNO indicating the advantage of adding convolutional RRDB layers that learn spatial domain features useful for downscaling (as shown in ESRGAN (Wang et al. 2018)). Table 1 zero-shot results show a performance comparison between models trained on a 4x upsampling factor but evaluated zero-shot on generating 8x upsampled HR outputs. The zero-shot experiments show that the SwinIR is still the best-performing model. While DUNO is best among the neural operator models at zero-shot downscaling, DCNO performs much worse than it did on standard downscaling. All the neural operator models are better than bicubic and SRCNN.

**Energy Spectrum** Figure 2 shows zero-shot downsampled wind speed for the SwinIR, ESRGAN, DFNO, DUNO, DCNO, and DAFNO models, alongside the LR, HR, and bicubic interpolated images. We refer to the energy spectrum plot in Figure 3a to show the kinetic energy distributions as functions of wavenumber, across all the downscaling models for the task of zero-shot downscaling. SwinIR best captures the physical properties of the ground truth at low-medium wavenumbers, ESRGAN is better at medium-high wavenumbers but DAFNO is the best at the highest wavenumbers, even though they underestimate the energy content at higher wavenumbers. For the highest wavenumbers or the dissipation range, this underestimation is most significant (for all the models except DAFNO). This is not surprising, as, for example, it is challenging for the models to fill in smaller-scale physical features if they do not see this level of detail when training on smaller upsampling factors. DCNO matches the HR curve for lower wavenumbers but falls behind ESRGAN at higher wavenumbers. DUNO and EDSR are close to but fall behind SwinIR. We observe that SwinIR, ESRGAN, and EDSR produce peaks at the very high-end wavenumbers but the neural operator models except DAFNO do not introduce this high-end noise. See

Appendix C, Figure 5a for the standard downscaling energy spectrum plot.

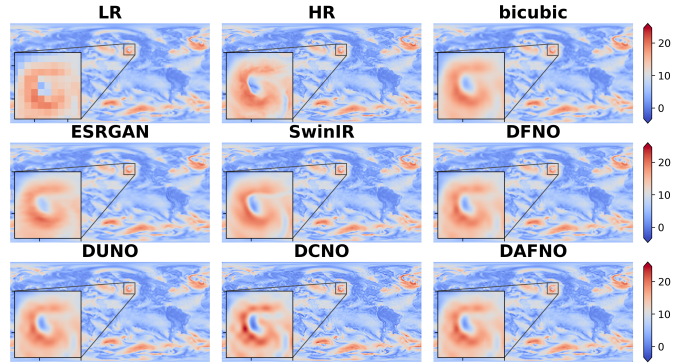


Figure 2: ERA5 wind speed visualizations in  $m/s$  generated from the zero-shot downscaling. We zoom in on a small region for better comparison. SwinIR captures better and finer details of the HR image, over neural operator models, especially in the zoomed-in region. It is also better over regions with complex terrain (e.g. the mountain ranges in North and South America).

### 5.2 ERA5 to WTK downscaling

**Error Metrics** We show the results for the ERA5 to WTK downscaling in Table 2. As discussed in Section 4.2, using ERA5 as the LR for this setup makes it more challenging as the LR data is obtained from a different simulation rather than using a coarsened version of the HR data. Table 2 shows the (1) standard downscaling results obtained from evaluating the models mapping ERA5 to WTK for a 5x upsampling factor and (2) zero-shot downscaling results where we evaluate the models trained on 5x upsampling to generate HR outputs at a 15x upsampling factor. SwinIR remains the best-performing model in both setups. DCNO achieves the best standard downscaling scores among the neural operators but is poor at zero-shot downscaling, as can be seen in Table 2 zero-shot results, where it performs worse than bicubic interpolation. We also observe DAFNO to be performing worse than bicubic interpolation at zero-shot downscaling. EDSR is a close second to SwinIR in both experiments. DUNO performs better than the other neural operators (and bicubic as well as SRCNN) at zero-shot downscaling.

**Energy Spectrum** The energy spectrum plot for the ERA5 to WTK zero-shot downscaling experiment is presented in Figure 3b. ESRGAN comes closest to matching the HR energy spectrum for the zero-shot downscaling. However, as seen in Section 5.1, these models still underestimate the energy content in the high wavenumber range. We observe that SwinIR and EDSR follow behind ESRGAN, but they outperform all the neural operator-based models. DAFNO no longer shows good performance for this experimental setup. FNO performs the worst, consistent with their performance in terms of the average error metrics (as seen in Table 2). Figure 4 compares the zero-shot downsampled wind speed outputs for the bicubic interpolation, ESRGAN,

Table 1: **ERA5 to ERA5 wind speed downscaling results.** MSE has units  $(m/s)^2$ , MAE  $m/s$  and IN  $m/s$ . We bold the best-performing model among all the models and underline the best-performing neural operator model. Results for the other channels are added to the Appendix C.

	Standard Downscaling					Zero-shot Downscaling			
	is NO?	MSE ↓	MAE ↓	IN ↓	PSNR ↑	MSE ↓	MAE ↓	IN ↓	PSNR ↑
bicubic	✗	1.23	0.73	14.82	27.53	1.23	0.73	14.82	27.53
SRCNN	✗	1.14	0.7	14.75	27.83	1.06	0.67	14.56	28.18
ESRGAN	✗	1.29	0.75	15.43	27.3	0.85	0.6	14.51	29.1
EDSR	✗	0.51	0.44	13.6	31.33	0.54	0.45	13.66	31.1
SwinIR	✗	<b>0.36</b>	<b>0.38</b>	<b>12.17</b>	<b>32.84</b>	<b>0.51</b>	<b>0.44</b>	<b>13.2</b>	<b>31.34</b>
FNO	✓	1.36	0.85	14.58	27.09	0.95	0.68	14.17	28.64
DFNO	✓	0.75	0.58	12.53	29.67	0.66	0.51	13.28	30.26
DUNO	✓	0.69	0.53	13.36	30.04	<u>0.63</u>	<u>0.5</u>	<u>13.56</u>	<u>30.44</u>
DAFNO	✓	0.65	0.51	13.77	30.29	<u>0.66</u>	<u>0.52</u>	<u>14.01</u>	<u>30.23</u>
DCNO	✓	<u>0.45</u>	<u>0.43</u>	<u>12.98</u>	<u>31.93</u>	0.92	0.65	14.89	28.76

Table 2: **ERA5 to WTK wind speed downscaling results.** MSE has units  $(m/s)^2$ , MAE  $m/s$  and IN  $m/s$ . We aggregate the error metrics over u and v wind velocity channels. We bold the best-performing model among all the models and underline the best-performing neural operator model.

	Standard Downscaling					Zero-shot Downscaling			
	is NO?	MSE ↓	MAE ↓	IN ↓	PSNR ↑	MSE ↓	MAE ↓	IN ↓	PSNR ↑
bicubic	✗	3.56	1.18	12.87	18.4	4.07	1.25	16.59	19.91
SRCNN	✗	3.16	1.11	12.62	18.83	3.65	1.18	16.39	20.31
ESRGAN	✗	2.75	1.05	13.06	19.27	3.12	1.11	15.96	20.8
EDSR	✗	2.46	0.98	11.89	19.85	2.92	1.05	15.61	21.2
SwinIR	✗	<b>2.29</b>	<b>0.95</b>	<b>11.69</b>	<b>20.11</b>	<b>2.73</b>	<b>1.02</b>	<b>15.33</b>	<b>21.43</b>
FNO	✓	5.23	1.76	14.36	15.24	5.2	1.75	17.82	17.35
DFNO	✓	3.04	1.26	12.56	17.86	3.53	1.33	16.14	19.4
DUNO	✓	2.81	1.09	12.11	18.97	<u>3.3</u>	<u>1.16</u>	<u>15.85</u>	<u>20.43</u>
DAFNO	✓	2.71	1.02	12.12	19.47	4.17	<u>1.19</u>	<u>17.5</u>	<u>19.51</u>
DCNO	✓	<u>2.47</u>	<u>0.99</u>	<u>11.77</u>	<u>19.79</u>	4.66	1.32	17.32	19.51


SwinIR, DFNO, DUNO, DCNO, and DAFNO models. See Appendix C, Figure 5b for the standard downscaling energy spectrum plot.

## 6 Discussion


In the literature, neural operators have performed well at zero-shot super-resolution (Kovachki et al. 2023; Li et al. 2021; Rahman, Ross, and Azzadenesheli 2023; Raonic et al. 2023) when trained to predict the solution of a PDE or when trained to act as an emulator of a time-dependent dynamical system. Their resolution invariance property has also been utilized in Jiang et al. (2023) to train an FNO to act as an emulator for zero-shot super-resolution of weather forecasts. However, our results show that the neural operators under-perform the non-neural operators at zero-shot weather downscaling. Importantly, physical system emulation differs from our static downscaling setting, where we train on pairs of low and high-resolution images. In our case, the neural operators are trained to learn a mapping between resolutions, and are tested on their ability to generalize zero-

shot to higher upsampling factors. We believe this distinction is important to help contextualize our results, which show that neural operators have difficulty with this task.

Our results show that bicubic interpolation followed by a vanilla FNO performs poorly at weather downscaling. In most cases, this performs worse than bicubic interpolation alone. FNO learns spectral features in Fourier space which makes it resolution-invariant—these features are not inherently tied to the resolution of the training dataset. Weather downscaling may benefit from learning spatial features tied to the specific input *and* output grid resolutions. This could be limiting the vanilla FNO’s ability to downscale well. We also evaluate neural operator models with convolutional RRDB blocks before the neural operator layers (the DXNO models (Section 4)), which improves the downscaling performance significantly. The DCNO models, based on CNOs (Raonic et al. 2023), adapt U-Net (Ronneberger, Fischer, and Brox 2015) style convolutions to approximately learn an operator mapping. They perform close to the best model SwinIR in the standard downscaling experiments, but,



(a)



(b)

Figure 3: Figures (a) and (b) show kinetic energy spectrum plots for the ERA5 to ERA5 and ERA5 to WTK zero-shot downscaling, respectively. *Kinetic Energy is normalized and wavenumber is measured relative to the domain size.*

their performance drops significantly in the zero-shot setup, as CNO uses explicit up/down-sampling and thus cannot be applied to different resolutions without some degraded performance (Liu-Schiaffini et al. 2024).

ESRGAN proves to be the best model for capturing the physical properties of the data at medium-high wavenumbers for ERA5 to ERA5 and all wavenumbers for ERA5 to WTK experiments, as measured in our work by kinetic energy plots, for zero-shot downscaling. It is important to note that zero-shot downscaling is a challenging task as we expect the models to produce outputs that have super-resolved physics at the finer scale without training on them. It is possible that ESRGAN learns to generate downscaled outputs with better visual quality because of its architectural design and use of perceptual loss, which may help in capturing the HR physics across spatial scales, yet, we observe that all models underestimate the energy content in the high-




Figure 4: WTK wind speed visualization in  $m/s$  generated from the zero-shot downscaling (figure shows results on one of the two regions). We observe ESRGAN’s downscaled outputs (followed by SwinIR’s) to be sharper with better details than the neural operator models.

wavenumber range for zero-shot downscaling. It seems that SwinIR learns superior-quality features at the smaller up-sampling factor during training, enabling an interpolation on top of SwinIR to generate downscaled outputs better than other models as shown by the average error metrics. Our results suggest that residual Swin-Transformer blocks (RSTB) as adopted in SwinIR are better at extracting high-quality features than RRDB blocks, which can potentially help inform future architectures for downscaling. We recommend that researchers benchmark against powerful non-operator-learning methods with interpolation as strong baselines. However, given that SwinIR and ESRGAN need to use bicubic interpolation (which has no learnable parameters) to do zero-shot downscaling, it could be fundamentally limited in its ability to downscale small-scale physics unseen during training. It is also possible that the set of neural operator models we explored can be improved. Overall, all models appear to be quite far from solving our downscaling tasks.

## 7 Conclusion

This work comprehensively benchmarks neural operators on the task of weather downscaling, with a particular emphasis on critically investigating the zero-shot downscaling capabilities of neural operators. Our analyses involve two studies over (1) learning a mapping from coarsened ERA5 to high-resolution ERA5, and (2) learning a mapping from low-resolution ERA5 wind data to a high-resolution wind data (WTK, 2km x 2km). Our zero-shot downscaling experiments involve challenging upsampling factors: 8x and 15x over the two studies respectively.

With an extensive evaluation using various error metrics and kinetic energy spectrum plots, we show that resolution-invariant neural operators are outperformed by the Swin-Transformer and ESRGAN-based models, even at zero-shot downscaling. This was surprising, as resolution-invariant neural operators were previously shown to be good at zero-shot super-resolution for emulating dynamical systems. While our current study presents limitations of neural operators at weather downscaling, future research may consider

improving the neural operator downscaling frameworks with better feature encoders (Wei and Zhang 2023) or advanced hybrid neural-operator-transformer models (Luo, Qian, and Yoon 2024).

**Acknowledgments** This work was authored by the National Renewable Energy Laboratory (NREL), operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. This work was supported by the Laboratory Directed Research and Development (LDRD) Program at NREL. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes. The research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory.

## References

- Alexey, D. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Bartolucci, F.; de Bézenac, E.; Raonic, B.; Molinaro, R.; Mishra, S.; and Alaifari, R. 2024. Representation equivalent neural operators: a framework for alias-free operator learning. *Advances in Neural Information Processing Systems*, 36.
- Benton, B. N.; Buster, G.; Pinchuk, P.; Glaws, A.; King, R. N.; Maclaurin, G.; and Chernyakhovskiy, I. 2024. Super Resolution for Renewable Energy Resource Data With Wind From Reanalysis Data (Sup3rWind) and Application to Ukraine. *arXiv preprint arXiv:2407.19086*.
- Buster, G.; Benton, B. N.; Glaws, A.; and King, R. N. 2024. High-resolution meteorology with climate change impacts from global climate model data using generative machine learning. *Nature Energy*, 1–13.
- Chaudhuri, C.; and Robertson, C. 2020. CliGAN: A structurally sensitive convolutional neural network model for statistical downscaling of precipitation from multi-model ensembles. *Water*, 12(12): 3353.
- Chen, X.; Feng, K.; Liu, N.; Ni, B.; Lu, Y.; Tong, Z.; and Liu, Z. 2022. RainNet: A large-scale imagery dataset and benchmark for spatial precipitation downscaling. *Advances in Neural Information Processing Systems*, 35: 9797–9812.
- Dong, C.; Loy, C. C.; He, K.; and Tang, X. 2015. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2): 295–307.
- Draxl, C.; Clifton, A.; Hodge, B.-M.; and McCaa, J. 2015. The wind integration national dataset (wind) toolkit. *Applied Energy*, 151: 355–366.
- Guibas, J.; Mardani, M.; Li, Z.; Tao, A.; Anandkumar, A.; and Catanzaro, B. 2021. Adaptive fourier neural operators: Efficient token mixers for transformers. *arXiv preprint arXiv:2111.13587*.
- Harilal, N.; Hodge, B.-M. S.; Monteleoni, C.; and Subramanian, A. 2022. EnhancedSD: Downscaling Solar Irradiance from Climate Model Projections. In *NeurIPS 2022 Workshop on Tackling Climate Change with Machine Learning*.
- Harris, L.; McRae, A. 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.
- Hersbach, H.; Bell, B.; Berrisford, P.; Hirahara, S.; Horányi, A.; Muñoz-Sabater, J.; Nicolas, J.; Peubey, C.; Radu, R.; Schepers, D.; et al. 2020. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730): 1999–2049.
- Jiang, P.; Yang, Z.; Wang, J.; Huang, C.; Xue, P.; Chakraborty, T.; Chen, X.; and Qian, Y. 2023. Efficient Super-Resolution of Near-Surface Climate Modeling Using the Fourier Neural Operator. *Journal of Advances in Modeling Earth Systems*, 15(7): e2023MS003800.
- Kaczmarek, J.; Isham, V.; and Onof, C. 2014. Point process models for fine-resolution rainfall. *Hydrological Sciences Journal*, 59(11): 1972–1991.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kolmogorov, A. N. 1991. Dissipation of energy in the locally isotropic turbulence. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 434(1890): 15–17.
- Kovachki, N.; Li, Z.; Liu, B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.; and Anandkumar, A. 2023. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89): 1–97.
- Kurinchi-Vendhan, R.; Lütjens, B.; Gupta, R.; Werner, L.; and Newman, D. 2021. Wisosuper: Benchmarking super-resolution methods on wind and solar data. *arXiv preprint arXiv:2109.08770*.
- Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4681–4690.
- Li, Z.; Kovachki, N. B.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A. M.; and Anandkumar, A. 2021. Fourier Neural Operator for Parametric Partial Differential Equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.



- Li, Z. L. 2023. ESRGAN-PyTorch. <https://github.com/lizhuoq/ESRGAN-pytorch/>.
- Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Van Gool, L.; and Timofte, R. 2021. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1833–1844.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Mu Lee, K. 2017. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 136–144.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Liu-Schiaffini, M.; Berner, J.; Bonev, B.; Kurth, T.; Aziz-zadenesheli, K.; and Anandkumar, A. 2024. Neural Operators with Localized Integral and Differential Kernels. In *Forty-first International Conference on Machine Learning*.
- Luo, X.; Qian, X.; and Yoon, B.-J. 2024. Hierarchical Neural Operator Transformer with Learnable Frequency-aware Loss Prior for Arbitrary-scale Super-resolution. *arXiv preprint arXiv:2405.12202*.
- Mikhaylov, A.; Meshchaninov, F.; Ivanov, V.; Labutin, I.; Stulov, N.; Burnaev, E.; and Vanovskiy, V. 2024. Accelerating regional weather forecasting by super-resolution and data-driven methods. *Journal of Inverse and Ill-posed Problems*, (0).
- Pathak, J.; Subramanian, S.; Harrington, P.; Raja, S.; Chat-topadhyay, A.; Mardani, M.; Kurth, T.; Hall, D.; Li, Z.; Azizzadenesheli, K.; et al. 2022. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*.
- Pierce, D. W.; Cayan, D. R.; and Thrasher, B. L. 2014. Statistical downscaling using localized constructed analogs (LOCA). *Journal of hydrometeorology*, 15(6): 2558–2585.
- Prasad, A.; Harder, P.; Yang, Q.; Sattegeri, P.; Swarcman, D.; Watson, C.; and Rolnick, D. 2024. Evaluating the transferability potential of deep learning models for climate downscaling. *arXiv preprint arXiv:2407.12517*.
- Rahman, M. A.; Ross, Z. E.; and Azizzadenesheli, K. 2023. U-NO: U-shaped Neural Operators. *Transactions on Machine Learning Research*.
- Raonic, B.; Molinaro, R.; De Ryck, T.; Rohner, T.; Bartolucci, F.; Alaifari, R.; Mishra, S.; and de Bézenac, E. 2023. Convolutional Neural Operators for robust and accurate learning of PDEs. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 77187–77200. Curran Associates, Inc.
- Ren, P.; Erichson, N. B.; Subramanian, S.; San, O.; Lukic, Z.; and Mahoney, M. W. 2023. Superbench: A super-resolution benchmark dataset for scientific machine learning. *arXiv preprint arXiv:2306.14070*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Sengupta, M.; Xie, Y.; Lopez, A.; Habte, A.; Maclaurin, G.; and Shelby, J. 2018. The national solar radiation data base (NSRDB). *Renewable and sustainable energy reviews*, 89: 51–60.
- Shi, W.; Caballero, J.; Huszár, F.; Totz, J.; Aitken, A. P.; Bishop, R.; Rueckert, D.; and Wang, Z. 2016. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1874–1883.
- Stengel, K.; Glaws, A.; Hettinger, D.; and King, R. N. 2020. Adversarial super-resolution of climatological wind and solar data. *Proceedings of the National Academy of Sciences*, 117(29): 16805–16815.
- Tran, D. T.; Robinson, H.; Rasheed, A.; San, O.; Tabib, M.; and Kvamsdal, T. 2020. GANs enabled super-resolution reconstruction of wind field. In *Journal of Physics: Conference Series*, volume 1669, 012029. IOP Publishing.
- Vetterli, M.; Kovačević, J.; and Goyal, V. K. 2014. *Foundations of signal processing*. Cambridge University Press.
- Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; and Change Loy, C. 2018. Esgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, 0–0.
- Wang, Z.; Chen, J.; and Hoi, S. C. 2020. Deep learning for image super-resolution: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(10): 3365–3387.
- Watson, C. D.; Wang, C.; Lynar, T.; and Weldemariam, K. 2020. Investigating two super-resolution methods for downscaling precipitation: ESRGAN and CAR. *arXiv preprint arXiv:2012.01233*.
- Wei, M.; and Zhang, X. 2023. Super-resolution neural operator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18247–18256.
- Wood, A. W.; Leung, L. R.; Sridhar, V.; and Lettenmaier, D. 2004. Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs. *Climatic change*, 62: 189–216.
- Yang, Q.; Hernandez-Garcia, A.; Harder, P.; Ramesh, V.; Sattegeri, P.; Swarcman, D.; Watson, C. D.; and Rolnick, D. 2023. Fourier neural operators for arbitrary resolution climate data downscaling. *arXiv preprint arXiv:2305.14452*.

## A Neural-operator-based model details

The four neural operator models we compare are:

1. **FNO** (Li et al. 2021) FNO uses a combination of linear integral operators and non-linear local activation functions to learn the operator mapping on complex PDEs such as the Navier-Stokes equation. We incorporate localized operator layers (Liu-Schiaffini et al. 2024) into

the FNO model as they help in learning better local features and showed improved performance in our downscaling task compared to FNO without the local layers. We tested adding these localized operators into the other FNO-based models: DFNO and DUNO, but it did not improve their performance.

2. **UNO** (Rahman, Ross, and Azizzadenesheli 2023) is introduced as a deep and memory-efficient architecture that allows for faster training than FNOs. This neural operator has a U-shaped architecture and learns the mapping between function spaces at different domains based on its encoder-decoder framework.
3. **AFNO** (Guibas et al. 2021), a transformer model that uses FNOs for efficient token-mixing instead of the traditional self-attention layers. AFNOs have been used as an integral part of FourCastNet (Pathak et al. 2022), a data-driven weather forecasting model.
4. **CNO** (Raonic et al. 2023) adapts CNNs for learning operators and processing the inputs and outputs as function spaces. They adopt a UNet-based modeling framework to learn a mapping between bandlimited functions (Vetterli, Kovačević, and Goyal 2014), this can help them to learn operators that reduce aliasing errors (Bartolucci et al. 2024) which occur when the neural operator models try to learn a continuous operator on a finite, discretized grid.

## B Training details

### B.1 Additional data details

1. **ERA5 to ERA5 downscaling:** The number of samples (image snapshots) in training, validation, and test are 1460, 730, and 730 respectively. Refer to the main text for the size details (height, width, channels) of each image.

We train all the models two times in this setup.

*Standard downscaling:* This involves training all the models with an upsampling factor of 8x. Following Ren et al. (2023), we do not use the entire image snapshots but random crops of the images for training. We extract eight patches of size 128x128 from each image snapshot to obtain HR image data for training. The corresponding LR image patches of size 16x16 are created by coarsening the HR patches with bicubic interpolation.

*Zero-shot downscaling:* This involves training all the models with an upsampling factor of 4x. For this case, we extract eight patches of size 64x64 from each image snapshot to obtain HR image data for training. The corresponding LR image patches of size 16x16 are created by coarsening the HR patches with bicubic interpolation.

2. **ERA5 to WTK downscaling:** We work with two regions in the US. For each of the regions: the number of samples (image snapshots) in training, validation, and test are 7008, 1752, and 8760 respectively. Refer to the main text for the size details (height, width, channels) of each image.

All the models are trained just once in this setup, with an upsampling factor of 5x. We again use random crops of the images for training. We extract eight patches of size 160x160 from each image snapshot to obtain HR (WTK) image data for training. The LR image patches of size 32x32 are obtained from the corresponding ERA5 image snapshots. We train a single model for both regions, ensuring that every training batch has an equal number of LR and HR pairs from both regions.

The crop size for the ERA5 to ERA5 experiments is the same as the one used in Ren et al. (2023). We tune the crop size for the ERA5 to WTK downscaling experiments and obtain the optimal crop size reported above.

### B.2 Hyperparameters

All the models are trained for 400 epochs.

#### Neural-operator-based models :

All the hyperparameters are tuned over the validation dataset. We perform a sweep over learning rates:  $\{0.005, 0.0001, 0.00001\}$  for all models. The models are trained using the ADAM (Kingma and Ba 2015) optimizer with a batch size of 32, weight decay of  $1e-4$ , and a step learning rate scheduler with a step size of 60. For the Downscaling (D) models, we add the RRDB module, a part of the ESRGAN framework (Wang et al. 2018) before the neural operator layers implementation. We perform a sweep over the number of RRDB blocks:  $\{6, 12, 24\}$  for all the downscaling models.

1. **FNO** We follow the original implementation of FNO from (Li et al. 2021; Kovachki et al. 2023) using most of the default model hyperparameters. We perform a hyperparameter sweep over the number of hidden channels in the lifting and projection blocks, selecting 256 for them, and the number of modes to keep in the Fourier layers, 16 being the best. The best learning rate for the FNO model is found to be 0.005. We use the  $l_p$  loss with  $p = 2$ , reduced over  $dim = 0$  as defined in the original implementation.
2. **DFNO** We keep the selected values for the lifting, projection channels, and the number of models, obtained from tuning the above FNO model. The best learning rate for training DFNO is found to be 0.0001, and the optimal number of RRDB blocks is selected as 12. MSE is used as the loss function.
3. **DUNO** We follow the UNO model implementation released as a part of FNO’s original implementation (Li et al. 2021; Kovachki et al. 2023). While we use most of the default model hyperparameters, we do hyperparameter tuning on the hidden channels (initial width of UNO), selecting 64 as optimal, and the number of output channels of each Fourier layer, selecting them as 64. We found the best learning rate to be 0.0001, and the optimal number of RRDB blocks is 12. MSE is used as the loss function.
4. **DAFNO** We follow the implementation of the AFNO network from the FourCastNet (Pathak et al. 2022) im-

plementation. With most model hyperparameters as default, we perform a hyperparameter sweep over the patch size, choosing 8 as optimal, and the number of blocks (block as defined in (Guibas et al. 2021)), selecting 8. It should be noted that the optimal patch size is found to be 4 when training for the ERA5 to ERA5 zero-shot down-scaling setup. For the DAFNO training, we find the best values for the learning rate to be 0.0001, and the number of RRDB blocks to be 12. MSE is used as the loss function.

5. **DCNO** We follow the original CNO implementation from Raonic et al. (2023) with most of the the default model hyperparameters. We tune the number of layers (upsampling/downsampling blocks) and find the optimal to be 3. For the DCNO training, we find the best values for the learning rate to be 0.0001, and the number of RRDB blocks to be 12. MSE is used as the loss function.

**Baseline models :**

The implementations of SRCNN, EDSR, and SwinIR model pipelines follow the implementations provided by Ren et al. (2023). We follow an open-source implementation of ESRGAN from Li (2023) for our ESRGAN downscaling framework. We keep most of the hyperparameters and training setups from these implementations, but we train each of the baseline models for a fixed 400 epochs (consistent with the neural-operator-based models). We also do a hyperparameter sweep over the learning rates: {0.001,0.0001,0.00001} for all the baseline models, using the validation dataset to tune this hyperparameter. For the ERA5 to ERA5 experiments: we find the optimal learning rate as 0.0001 for EDSR, SwinIR, and ESRGAN, and 0.001 for SRCNN. For the ERA5 to WTK experiments: we find the optimal learning rate as 0.0001 for SRCNN, SwinIR, and ESRGAN, and 0.001 for EDSR.

**B.3 Model parameters**

**C Additional results**

Table 3: Model parameters for all the Baseline and Neural-operator-based models used in the ERA5 to WTK down-scaling setup. SwinIR achieves superior average error metrics (e.g., MSE), as shown in Table 2, while having only marginally higher model parameters than the downscaling neural-operator-based models (except DAFNO). ESRGAN is superior in matching the ground truth energy spectrum (Figure 3b) and is second to DAFNO in parameter count.

<b>Model</b>	<b>#Parameters</b>
SRCNN	0.063M
ESRGAN	39.18M
EDSR	2.144M
SwinIR	12.526M
FNO	1.24M
DFNO	9.89M
DUNO	9.36M
DAFNO	69.15M
DCNO	11.33M



Table 4: **ERA5 to ERA5 temperature downscaling results.** MSE has units  $(K)^2$ , MAE  $K$  and IN  $K$ . We bold the best-performing model among all the models and underline the best-performing neural operator model.

	Standard Downscaling					Zero-shot Downscaling			
	is NO?	MSE ↓	MAE ↓	IN ↓	PSNR ↑	MSE ↓	MAE ↓	IN ↓	PSNR ↑
bicubic	✗	2.47	0.89	20.41	45.99	2.47	0.89	20.41	45.99
SRCNN	✗	2.09	0.83	20.94	46.74	1.95	0.78	20.35	47.02
ESRGAN	✗	3.75	1.27	45.5	44.18	1.54	0.82	16.88	48.04
EDSR	✗	0.68	0.44	15.5	51.6	0.81	0.47	15.91	50.84
SwinIR	✗	<b>0.39</b>	<b>0.35</b>	<b>11.17</b>	<b>54.04</b>	<b>0.81</b>	<b>0.46</b>	<b>16.27</b>	<b>50.84</b>
FNO	✓	1.93	0.9	18.42	47.06	1.34	0.73	17.82	48.64
DFNO	✓	1.46	0.74	18.44	48.3	1.06	0.61	16.93	49.65
DUNO	✓	0.88	0.52	15.61	50.47	0.94	<u>0.53</u>	<u>17.18</u>	<u>50.2</u>
DAFNO	✓	0.9	0.56	15.45	50.39	<u>1.21</u>	0.68	<u>32.75</u>	49.1
DCNO	✓	<u>0.45</u>	<u>0.4</u>	<u>11.11</u>	<u>53.38</u>	1.85	0.81	16.72	47.26

Table 5: **ERA5 to ERA5 total column water vapor downscaling results.** MSE has units  $(kg/m^2)^2$ , MAE  $kg/m^2$  and IN  $kg/m^2$ . We bold the best-performing model among all the models and underline the best-performing neural operator model.

	Standard Downscaling					Zero-shot Downscaling			
	is NO?	MSE ↓	MAE ↓	IN ↓	PSNR ↑	MSE ↓	MAE ↓	IN ↓	PSNR ↑
bicubic	✗	2.72	0.98	29.44	33.18	2.72	0.98	29.44	33.18
SRCNN	✗	2.59	0.96	29.24	33.4	2.27	0.89	29.01	33.97
ESRGAN	✗	3.1	0.94	48.3	32.61	1.71	0.79	26.17	35.19
EDSR	✗	0.94	0.55	23.76	37.81	2.92	1.05	15.61	21.2
SwinIR	✗	<b>0.61</b>	<b>0.47</b>	<b>15.83</b>	<b>39.69</b>	<b>1.14</b>	<b>0.57</b>	<b>25.45</b>	<b>36.97</b>
FNO	✓	3.18	1.12	29.22	32.49	1.83	0.78	28.38	34.9
DFNO	✓	1.95	0.83	29.74	34.63	1.64	0.75	27.77	35.38
DUNO	✓	1.18	0.64	25.47	36.81	<u>1.38</u>	<u>0.66</u>	<u>26.14</u>	<u>36.13</u>
DAFNO	✓	1.16	0.65	25.28	36.87	<u>1.85</u>	<u>0.84</u>	<u>26.62</u>	<u>34.84</u>
DCNO	✓	<u>0.67</u>	<u>0.51</u>	<u>13.53</u>	<u>39.23</u>	2.53	1.01	28.45	33.48

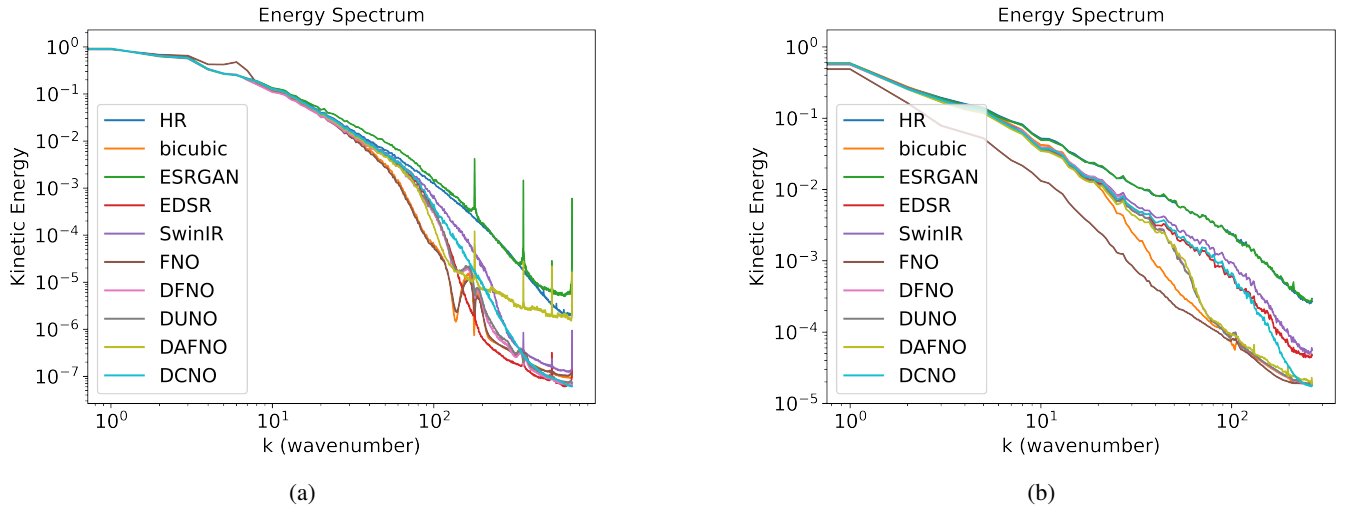


Figure 5: Figures (a) and (b) show kinetic energy spectrum plots for the ERA5 to ERA5 and ERA5 to WTK standard downscaling, respectively. Kinetic Energy is normalized and wavenumber is measured relative to the domain size. ESRGAN matches the HR spectrum for both setups, even at higher wavenumbers. DAFNO is second to ESRGAN at high wavenumbers for the ERA5 to ERA5 standard downscaling. SwinIR and EDSR rank second to ESRGAN for the ERA5 to WTK setup.