Reinstating Continuous Climate Patterns From Small and Discretized Data

Xihaier Luo [1] **Xiaoning Qian** [1 2] **Nathan Urban** [1] **Byung-Jun Yoon** [1 2]

Abstract

Wind energy is a leading renewable energy source. It does not pollute the environment and reduces greenhouse gas emissions that contribute to global warming. However, current wind characterization is performed at a resolution insufficient for assessing renewable energy resources in different climate scenarios. In this paper, we advocate the use of deep generative models for wind field representation learning. In contrast to existing approaches, we formulate the generative model as an explicit function of the spatial coordinate, thereby learning a continuous representation of the wind field, which can extrapolate from discretized data with demonstrated generalizability. We extend the concept of conditional neural fields by encoding the local turbulent wind properties into latent variables. Such resolution enhancement enables essential localized analyses of renewable energy resources' long-term economic sustainability.

1. Introduction

With the increased reliance on wind-based energy generation in modern power systems, it is critical to properly capture the dynamic behaviors of climatological wind when planning power system operation (Pryor et al., 2020; Veers et al., 2019; Manwell et al., 2010). However, due to the low viscosity of air, nearly all winds appear turbulent by nature, particularly in the troposphere, which houses the majority of critical systems such as infrastructure, ecological systems, and power plants. Many studies have been conducted in order to develop global climate simulation models that provide an accurate characterization of wind (Al-Yahyai et al., 2010; Cassola & Burlando, 2012; Zhang et al., 2022). A typical climate model used in wind data generation has a resolution of around 1°, or approximately 100 kilometer (km) near the equator. This resolution is insufficient for accurately assessing renewable energy resources, which typically necessitates a resolution finer than 10 km, preferably 2 km (Cox et al., 2018). As a result, we can only achieve suboptimal utilization and development of renewable energy farms.

Downscaling is a common method for obtaining local scale projections. Widely used downscaling methods can be classified into two classes: (1) *dynamical* methods and (2) *statistical* methods. The first class of methods typically utilizes numerical models to simulate atmospheric processes at a higher resolution within a limited geographical area. Regional climate models, for example, or limited-area models, use large-scale and lateral boundary conditions from climate models to generate higher resolution outputs (Wood et al., 2004; Murphy, 1999; Maraun et al., 2015). These models are commonly resolved at the 0.5° latitude and longitude scale and are used to parameterize physical atmospheric processes. The second class of methods employs statistical approaches to establish empirical relationships between climate-resolution climate variables and local climate, i.e., learning a low-resolution to high-resolution mapping from historical climate observations (Jakob Themeßl et al., 2011; Tang et al., 2016; Maraun et al., 2010). One critical limitation of such downscaling methods is that they typically fail to explicitly account for spatial dependencies in both low- and high-resolution climate data.

Alternatively, machine learning (ML) methods have been actively investigated in many recent studies to quickly generate high-resolution climate data (Stengel et al., 2020; Kurinchi-Vendhan et al., 2021; Yang et al., 2023). The process is known as image super-resolution (SR). The goal is to build a ML model – in most cases, a deep learning model – to accurately generate a high-resolution (HR) image output given low-resolution (LR) input image. If we are going to use the predictions of these neural network-based SR models to make decisions, we must ensure that their SR outputs are accurate and realistic. But what essentially is a climatological wind field? It is, by definition, a *continuous* physical property. However, when a numerical model attempts to process a climate field, it must first store and represent the field as 2D/3D arrays of pixels/voxels, with resolution controlling the trade-off between complexity and precision. Despite the

[1]Computational Science Initiative Brookhaven National Laboratory Upton, NY 11973, USA [2]Texas A&M University, College Station, TX 77843, USA. Correspondence to: <xluo@bnl.gov, byoon@bnl.gov>.

Accepted after peer-review at the 1st workshop on Synergy of Scientific and Machine Learning Modeling, SynS & ML ICML, Honolulu, Hawaii, USA. July, 2023.

need for an accurate continuous representation of climate fields, no general methods are available to date (Liu et al., 2020; Fukami et al., 2019; Güemes et al., 2021; Yousif et al., 2021; Kim et al., 2021; Luo & Kareem, 2021).

In this paper, we concentrate on the problem of learning continuous generative representations of wind data in a SR formulation to address the aforementioned shortcoming of existing methods. To summarize, the contributions of this article are as follows:

- A novel implicit neural network (INN) model for a continuous representation of climate data, such as turbulent wind fields.
- We encode climate data into a latent space and use random Fourier feature mapping for spatial coordinates. This improves characterization of high-frequency signals.
- We experimentally demonstrate that the proposed INN model allows for fast and accurate super-resolution up to $\times 20$ times, without requiring such high-resolution data during the training phase.

2. Method

In the weather and climate data set, a geophysical variable can be thought of as a scalar function of some physical variables, such as latitude, longitude, and time. Often, the underlying field generation process lacks a well-defined analytic form. Thus, functions are described by parameters Θ that climate scientists hand-craft and optimize. We refer to such a field as: $\mathbf{v} = F_\Theta(\mathbf{s})$, where $\mathbf{v}$ denotes the geophysical variable (in this paper, it is the wind speed), $\mathbf{s}$ denotes the spatial coordinate, and $F(\cdot)$ denotes a numerical weather model. Climate data is typically index sampled functions with discrete values, such as camera pixels, or discrete function parameterizations with voxels or discretized level sets. In light of this, the super-resolution task is commonly defined as follows:

$$\mathbf{v}_{high} = \mathcal{M}(\mathbf{v}_{low}) \tag{1}$$

Ideally, we want to access the climate property not just at fixed discrete locations, but at every possible point $\mathbf{s} \in \mathbb{R}^2$. Furthermore, we argue that such a super-resolution formulation (Eq. 1) ignores the inherent physics encapsulated in $F(\cdot)$. As a result, we present an alternative viewpoint: we compress the low-resolution data by training a neural network $\hat{F}(\cdot)$ to capture the desired climate property as a continuous scalar function of spatial coordinates:

$$\mathbf{v} = \hat{F}(\mathbf{s}) \tag{2}$$

Once the model has been trained $\hat{F} \approx F$, we can use it to make inferences for any continuous spatial coordinate in the domain and generate $\mathbf{v}_{high}$ at any resolution.

Aside from deep residual networks as the backbone, we utilized two additional tools to ensure that the neural network model can capture high-frequency signals in wind data and maintain good performance across different samples. The techniques we adopt for such a new continuous representation learning formulation are known as *Projection* and *Conditioning*.

Projection. Because the modeling function $\hat{F}(\cdot)$ is implicitly defined, continuous, and differentiable, current MLP-based network architectures have emerged as a powerful tool for representation learning. However, such implicit neural representations are incapable of modeling fine-grained signals, a phenomenon known as spectral bias in the literature (Basri et al., 2020; Rahaman et al., 2019). Fourier feature mapping has recently been discovered to be effective in overcoming the spectral bias of coordinate-based MLPs towards low frequencies by allowing them to learn at much higher frequencies (Tancik et al., 2020). As a result, we design a projection layer in our proposed network in order to better capture climatological wind details. The goal of this projection layer is to normalize the spatial coordinates and transform them into Fourier features. We specifically test three types of embeddings:

(1) Basic: $\gamma(\mathbf{s}) = [\cos(2\pi\mathbf{s}), \sin(2\pi\mathbf{s})]^{\mathrm{T}}$. (2) Positional: $\gamma(\mathbf{s}) = \left[\ldots, \cos\left(2\pi\sigma^{j/m}\mathbf{s}\right), \sin\left(2\pi\sigma^{j/m}\mathbf{s}\right), \ldots\right]^{\mathrm{T}}$ for $j = 0, \ldots, m-1$. For each dimension, we will use log-linear spaced frequencies, with the scale σ determined by a hyperparameter sweep for each task and dataset. (3) Gaussian: $\gamma(\mathbf{s}) = [\cos(2\pi\mathbf{B}\mathbf{s}), \sin(2\pi\mathbf{B}\mathbf{s})]^{\mathrm{T}}$, where *cos* and *sin* are used elementwise. $\mathbf{B}$ is a blocked diagonal random matrix where the non-zero values are sampled independently and randomly from a Gaussian distribution with a zero mean and standard deviation of σ.

Conditioning. For a climate data set, it is neither feasible nor practical to train a unique neural network for each data instance. Instead, we want a model that can represent a collection of instances, learn their inherent and interconnected properties, and embed them in a low-dimensional latent space (Xie et al., 2022; Chen et al., 2021). To that end, as shown in Fig. 1, we introduce a latent vector $\mathbf{z}$, which can be thought of disentangled input-dependent critical information by encoding the available data instance, as an additional conditioning input to our proposed neural network for continuous super-resolution. Conceptually, the projected spatial coordinates are integrated with this latent vector $\mathbf{z}$ to generate data-dependent continuous wind fields. Formally, $\hat{F}(\cdot)$ stated in Eq. 2 is now a function of a latent code $\mathbf{z}_i$ and a query of 2D locations, which outputs the wind

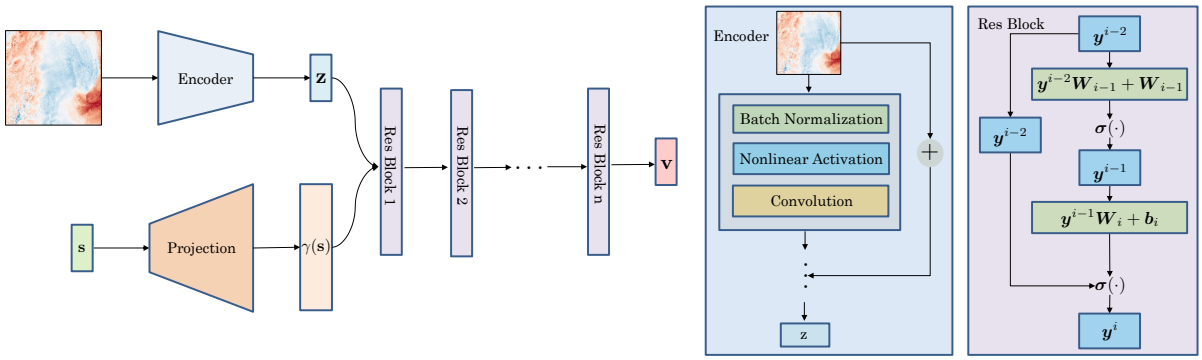

Figure 1. Schematic diagram illustrating the proposed implicit neural network (INN) model. The model is comprised of three key components: (1) a convolutional neural network-based encoder, trained to map an input image to a latent space; (2) a projection network that maps input coordinates to Fourier features, achieved through techniques such as sinusoidal mapping; and (3) a series of fully connected layers incorporating batch normalization and skip connections. By integrating these components, the INN model is capable of representing discretized climate data in the continuous domain, enabling its presentation at an arbitrarily high resolution.

speed at that location for an instance indexed by i:

$$u = \hat{F}(\mathbf{s}, \mathbf{z}_i) = \hat{F}(\gamma(\mathbf{s}), \phi(\mathbf{v}_i)), \tag{3}$$

where $\gamma(\cdot)$ is the projection of the coordinates and $\phi(\cdot)$ is the encoder. This formulation allows $\hat{F}(\cdot)$ to model a set of data instances with a single neural network by conditioning the network output on a latent vector, which is generated from the low-resolution data instance.

3. Experiments

Data set. The wind data was taken from the National Renewable Energy Laboratory Wind Integration National Database (WIND) Toolkit (Draxl et al., 2015a;b). The WIND Toolkit combines a meteorological data set provided by a mesoscale model, a power data set, and a forecast data set. It provides high spatial and temporal resolution wind power, wind power forecasting, and meteorological data for over 126,000 locations across the continental United States during a 7-year span. The simulated forecasts were developed using the Weather Research and Forecasting Model, which operates on a 2-km by 2-km grid with a 20-meter (m) resolution from the ground to 160 m above ground. The spatial resolution of the WIND Toolkit is 2-km $\times$ 1-hour (hr) spatio-temporal resolution. As a result, the wind data set is 1602 (latitude) $\times$ 2976 (longitude) $\times$ 61368 (number of instances), or almost 1.2 TB per wind component. We sample the data using standard baselines to make training and validation more efficient and replicable to others (Stengel et al., 2020). Wind test data is taken at a 4-hr temporal resolution, with the first 1000 snapshots used for experiments. The spatial resolution is reduced, and the final data set is 1500 $\times$ 2000 $\times$ 1000 in size, with each wind component requiring 6 GB of RAM. Specifically, there are 800 training instances, 100 validation instances, and 100 test instances in the data set. We used bicubic interpolation to generate a pair of low-resolution and high-resolution samples for each instance. For example, suppose we set the size of the input low-resolution sample to 75 $\times$ 100. The proposed model's super-resolution ability to scale from $\times 1$ to $\times 20$ can then be tested. We trained the model on common upscaling scales, from $\times 1$ to $\times 5$, then aggressively tested it not only on in-distribution scales ($\times 1 \sim \times 5$) but also on out-of-distribution scales ($\times 10 \sim \times 20$).

Implementation details. We utilize EDSR-baseline (excluding its up-sampling modules) as the encoder in our approach (Lim et al., 2017). The encoder consists of a single main branch with 16 residual blocks, facilitating parameter sharing across various upscaling ratios. The residual blocks are configured with 3×3 kernels. The initial layer of the encoder performs convolution and yields 64 feature maps. To maintain range flexibility within the network, we have excluded batch normalization layers from our network. This is because batch normalization layers normalize the features, thereby removing their range adaptability. In our decoding function, we employ a 5-layer MLP with ReLU activation and hidden dimensions set to 256. An Adam variant called decoupled weight decay regularization is adopted as the optimizer with a weight decay $\lambda = 0.0001$ (Loshchilov & Hutter, 2019). All models are trained for $1,000$ epochs with an initial learning rate of 0.001 that decays by $\gamma = 0.995$ every epoch. In the respective experiments, we used an NVIDIA A100 GPU with 40GB HBM2 memory.

Evaluation Metrics. We have employed three measures to evaluate model performance in addition to the mean squared error (MSE)-based loss between prediction and target. First, peak signal-to-noise ratio (PSNR) is defined as the ratio of a signal's maximum possible value (power) to the power of distorting noise that affects the quality of its representation. Second, the structural similarity index (SSIM) is a perceptual metric that evaluates the degradation of image quality

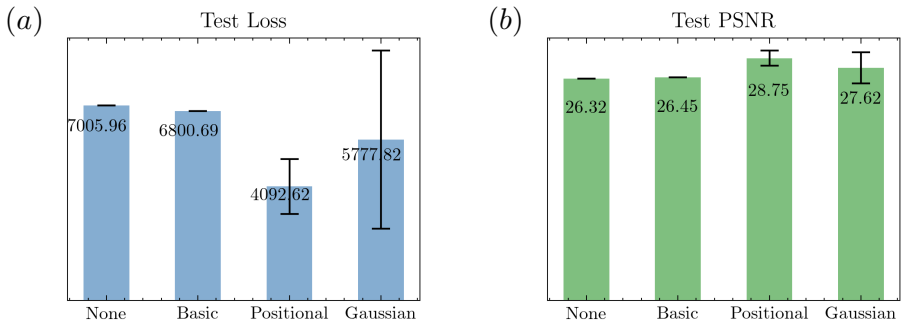

Figure 2. Test results. The MLP model was tested with and without different encoding methods using uniformly sampled hyperparameters.

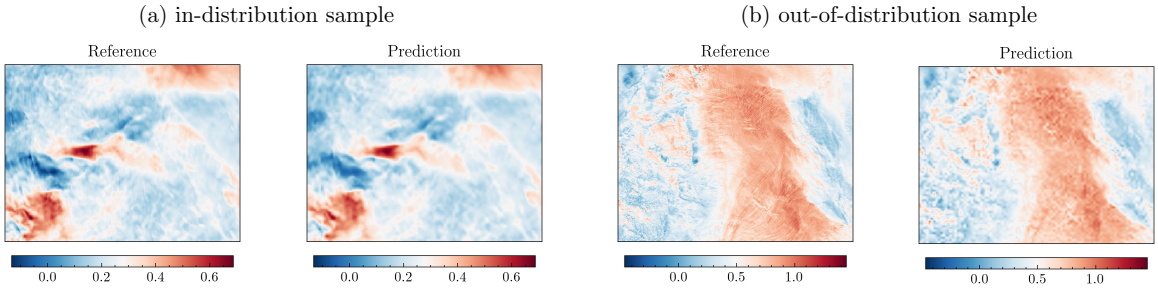

Figure 3. Model prediction performance is evaluated using in-distribution and out-of-distribution upscaling samples. The term *out-of-distribution* encompasses not only the test data samples that were not employed during the training phase but also extends to a more aggressive upscaling ratio (e.g., $\times 15$) that was not utilized in the training process (e.g., $\times 3$).

caused by data compression or transmission losses. Third, the mean absolute error (MAE) is employed.

3.1. Embedding results

The objective of this experiment is to demonstrate the efficacy of embedding methods in learning high-frequency signals. We experiment with four different models: (1) MLP. (2) MLP + Basic Encoding. (3) MLP + Positional Encoding. (4) MLP + Gaussian Encoding. The MLP architecture, which takes the form of $2 \to 256 \to 256 \to 256 \to 256 \to 1$, remains unchanged. It is worth noting that network architecture and hyperparameter search are two of the most important challenges and sources of innovation in deep learning, which are mostly problem-specific and empirical. Because we want to show how to use Fourier features-based embedding methods to address the learning issue of high-frequency signals, we will only focus on hyperparameter tuning for encoding. Grid search is used to evaluate candidates derived from a grid of parameter values specified by the grid parameters. The frequency constant σ is sampled from $\mathcal{U}[1, 50]$ in the case of positional encoding, and m, which denotes the number of frequencies to map to, is sampled from $\mathcal{U}[10, 50]$. Meanwhile, in the case of Gaussian encoding, the standard deviation is sampled from $\mathcal{U}[1, 50]$. Fig. 2 shows the test loss (the lower the better) and the test PSNR (the higher the better). The obtained results show that positional/Gaussian encoding improves performance significantly over basic encoding and the model trained without encoding. The mean performance suggests that positional encoding performs slightly better than Gaussian, at least with the hyperparameters we tested. The Gaussian encoding with hyperparameter $\sigma = 15$ provides the best performance, whether measured by the loss value or the test PNSR.

3.2. Qualitative evaluation results

To quantify the effectiveness of the learned continuous representation, we propose that, in addition to evaluating super-resolution tasks of scales in training distribution, we additionally evaluate extremely large upscaling ratios that are not in training distribution. Specifically, the upscaling ratios are randomly sampled from the uniform distribution $\mathcal{U}[2, 5]$ during the training period. During testing, the models are evaluated on previously unseen instances with considerably higher upscaling ratios sampled from the uniform distribution $\mathcal{U}[10, 20]$. Fig. 3 shows a qualitative comparison. The in-distribution sample is tested with an upscaling ratio of 3 and the out-of-distribution sample is tested with an upscaling ratio of 15. We can see from the visualization that INN captures the macro- and meso-scale climate patterns very well. For continuous learning, the first row of each sub-figure in Fig. 4 shows the predicted high-resolution representation with various upscaling ratios, and the second

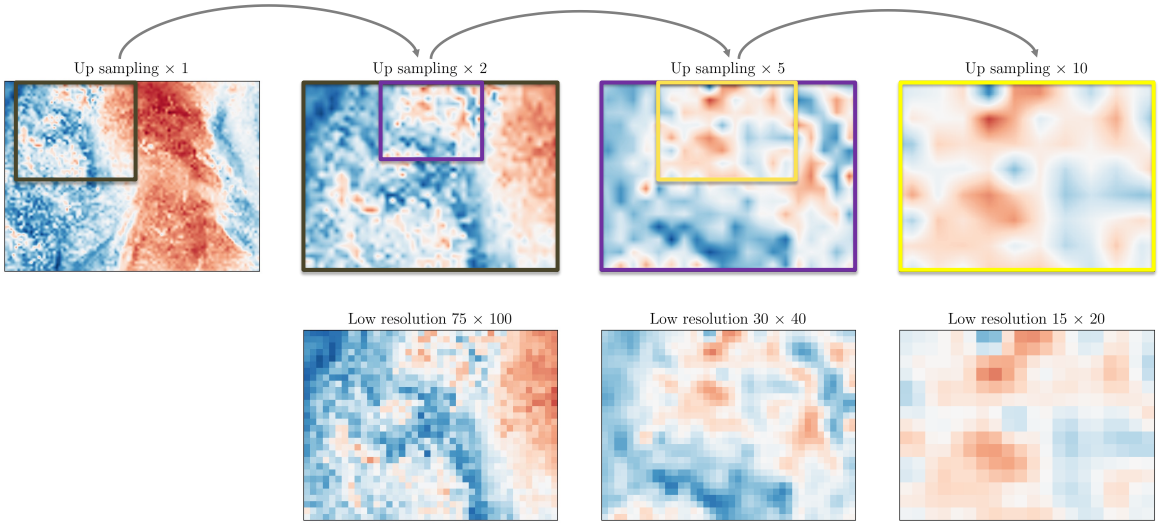

Figure 4. Qualitative illustration of learning continuous representation. The performance of the model is evaluated on a randomly selected instance from the test dataset. Different upscaling ratios were tested, and the resulting prediction outcomes are displayed in the first row. The second row showcases the original resolution data. By zooming in on the visualization of the wind speed data, we can observe the effectiveness of the proposed continuous super-resolution INN model.

	In-Distribution Performance				Out-of-Distribution Performance			
	Linear	Cubic	Nearest	INN	Linear	Cubic	Nearest	INN
PSNR ↑	31.969	32.299	29.507	**33.865**	27.447	27.364	26.154	**33.011**
SSIM ↑	0.8271	0.8347	0.7490	**0.9762**	0.7256	0.7321	0.6618	**0.9744**
MAE ↓	0.0201	0.0193	0.0259	**0.0192**	0.0509	0.0505	0.0585	**0.0197**

Table 1. Performance assessment results in terms of PSNR (peak signal-to-noise ratio), SSIM (structural similarity index measure), and MAE (mean absolute error). The best prediction performance is emphasized using bold font, where INN significantly outperforms all other interpolation-based baselines.

row shows the corresponding low-resolution representation. The results show that INN representation bridges the gap between discrete and continuous representations in 2D climate data by providing a framework for naturally and effectively exploiting information from image ground truths at various resolutions.

3.3. Quantitative evaluation results

Table 3.1 compares our proposed INN to the other three continuous super-resolution methods: bilinear, bicubic, and nearest interpolation. The PSNR, SSIM, and MAE are calculated. We compute prediction performance using different upscaling ratios, as opposed to traditional fixed upscaling ratio-based super-resolution. We tested 100 instances with various upscaling ratios ranging from 2 to 5 (In-Distribution Performance) as well as from 10 to 20 (Out-of-Distribution Performance). It should be pointed out that the proposed INN model outperforms all baseline interpolation methods, particularly when tested on a high upscaling ratio taken from the out-of-distribution range. Overall, the calculated results show that INN can model complex wind signals by learning a realistic approximation from low-resolution data.

4. Conclusion

In this work, we have proposed a coordinate-based deep learning model for continuous super-resolution (SR) of climate data. We have specifically developed an implicit neural network (INN) model for learning continuous, rather than discrete, representations of climate data. Notably, our INN-SR model only requires low-resolution samples as labels rather than its high-resolution counterparts. The INN-SR model, once trained with the WIND Toolkit, can generate data of any size. We have demonstrated that a learned continuous representation can generalize to much higher precision than the training scales while maintaining high fidelity. We have also tested the INN-SR model's performance on unseen test data. The proposed INN-SR model significantly outperforms standard interpolation-based SR approaches in terms of accuracy and robustness.

Broader impact

Broader impact Many physical systems and phenomena occur on a continuous basis in nature, as we know. They are formally characterized by partial differential equations (PDEs), which reflect the continuous nature of the physical systems. However, discretization is unavoidable during numerical simulations. Our study has the potential to improve the representation of current physical data. As a result, we believe this study can benefit both the broader scientific community and the general public in better understanding the physical systems of interest.

Fairness and ethic issues Our work involves significant analysis and experiments on datasets utilizing the WIND Toolkit, with no concerns about personal identity or privacy. As a result, our work is free of ethical and privacy concerns. Furthermore, the evaluations have been performed in a manner that guarantees fairness for all baselines and methods.

Acknowledgements

This research were supported by funding from the Advanced Scientific Computing Research program in the Department of Energy's Office of Science under project $B\&R\#KJ0402010$.

References

Al-Yahyai, S., Charabi, Y., and Gastli, A. Review of the use of numerical weather prediction (nwp) models for wind energy assessment. *Renewable and Sustainable Energy Reviews*, 14(9):3192–3198, 2010.

Basri, R., Galun, M., Geifman, A., Jacobs, D., Kasten, Y., and Kritchman, S. Frequency bias in neural networks for input of non-uniform density. In *International Conference on Machine Learning*, pp. 685–694. PMLR, 2020.

Cassola, F. and Burlando, M. Wind speed and wind energy forecast through kalman filtering of numerical weather prediction model output. *Applied energy*, 99:154–166, 2012.

Chen, Y., Liu, S., and Wang, X. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8628–8638, 2021.

Cox, S. L., Lopez, A. J., Watson, A. C., Grue, N. W., and Leisch, J. E. Renewable energy data, analysis, and decisions: a guide for practitioners. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2018.

Draxl, C., Clifton, A., Hodge, B.-M., and McCaa, J. The wind integration national dataset (wind) toolkit. *Applied Energy*, 151:355–366, 2015a.

Draxl, C., Hodge, B., Clifton, A., and McCaa, J. Overview and meteorological validation of the wind integration national dataset toolkit. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2015b.

Fukami, K., Fukagata, K., and Taira, K. Super-resolution reconstruction of turbulent flows with machine learning. *Journal of Fluid Mechanics*, 870:106–120, 2019.

Güemes, A., Discetti, S., Ianiro, A., Sirmacek, B., Azizpour, H., and Vinuesa, R. From coarse wall measurements to turbulent velocity fields through deep learning. *Physics of Fluids*, 33(7):075121, 2021.

Jakob Themeßl, M., Gobiet, A., and Leuprecht, A. Empirical-statistical downscaling and error correction of daily precipitation from regional climate models. *International Journal of Climatology*, 31(10):1530–1544, 2011.

Kim, H., Kim, J., Won, S., and Lee, C. Unsupervised deep learning for super-resolution reconstruction of turbulence. *Journal of Fluid Mechanics*, 910, 2021.

Kurinchi-Vendhan, R., Lütjens, B., Gupta, R., Werner, L., and Newman, D. Wisosuper: Benchmarking super-resolution methods on wind and solar data. *arXiv preprint arXiv:2109.08770*, 2021.

Lim, B., Son, S., Kim, H., Nah, S., and Mu Lee, K. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 136–144, 2017.

Liu, B., Tang, J., Huang, H., and Lu, X.-Y. Deep learning methods for super-resolution reconstruction of turbulent flows. *Physics of Fluids*, 32(2):025105, 2020.

Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

Luo, X. and Kareem, A. Dynamic mode decomposition of random pressure fields over bluff bodies. *Journal of Engineering Mechanics*, 147(4):04021007, 2021.

Manwell, J. F., McGowan, J. G., and Rogers, A. L. *Wind energy explained: theory, design and application*. John Wiley & Sons, 2010.

Maraun, D., Wetterhall, F., Ireson, A., Chandler, R., Kendon, E., Widmann, M., Brienen, S., Rust, H., Sauter, T., Themeßl, M., et al. Precipitation downscaling under climate

change: Recent developments to bridge the gap between dynamical models and the end user. *Reviews of geophysics*, 48(3), 2010.

Maraun, D., Widmann, M., Gutiérrez, J. M., Kotlarski, S., Chandler, R. E., Hertig, E., Wibig, J., Huth, R., and Wilcke, R. A. Value: A framework to validate downscaling approaches for climate change studies. *Earth's Future*, 3(1):1–14, 2015.

Murphy, J. An evaluation of statistical and dynamical techniques for downscaling local climate. *Journal of Climate*, 12(8):2256–2284, 1999.

Pryor, S. C., Barthelmie, R. J., Bukovsky, M. S., Leung, L. R., and Sakaguchi, K. Climate change impacts on wind power generation. *Nature Reviews Earth & Environment*, 1(12):627–643, 2020.

Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., and Courville, A. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pp. 5301–5310. PMLR, 2019.

Stengel, K., Glaws, A., Hettinger, D., and King, R. N. Adversarial super-resolution of climatological wind and solar data. *Proceedings of the National Academy of Sciences*, 117(29):16805–16815, 2020.

Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33: 7537–7547, 2020.

Tang, J., Niu, X., Wang, S., Gao, H., Wang, X., and Wu, J. Statistical downscaling and dynamical downscaling of regional climate in china: Present climate evaluations and future climate projections. *Journal of Geophysical Research: Atmospheres*, 121(5):2110–2129, 2016.

Veers, P., Dykes, K., Lantz, E., Barth, S., Bottasso, C. L., Carlson, O., Clifton, A., Green, J., Green, P., Holttinen, H., et al. Grand challenges in the science of wind energy. *Science*, 366(6464):eaau2027, 2019.

Wood, A. W., Leung, L. R., Sridhar, V., and Lettenmaier, D. Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs. *Climatic change*, 62(1):189–216, 2004.

Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., and Sridhar, S. Neural fields in visual computing and beyond. In *Computer Graphics Forum*, volume 41, pp. 641–676. Wiley Online Library, 2022.

Yang, Q., Hernandez-Garcia, A., Harder, P., Ramesh, V., Sattegeri, P., Szwarcman, D., Watson, C. D., and Rolnick, D. Fourier neural operators for arbitrary resolution climate data downscaling. *arXiv preprint arXiv:2305.14452*, 2023.

Yousif, M. Z., Yu, L., and Lim, H.-C. High-fidelity reconstruction of turbulent flow from spatially limited data using enhanced super-resolution generative adversarial network. *Physics of Fluids*, 33(12):125119, 2021.

Zhang, Z., Huang, P., Bitsuamlak, G., and Cao, S. Large-eddy simulation of wind-turbine wakes over two-dimensional hills. *Physics of Fluids*, 34(6):065123, 2022.