
Sparse Local Implicit Image Function for sub-km Weather Downscaling

Yago del Valle Inclan Redondo

Recursive Inc.

yago.valle@recursiveai.co.jp

Enrique Arriaga-Varela

Recursive Inc.

enrique.arriaga@recursiveai.co.jp

Dmitry Lyamzin

Recursive Inc.

dmitry.lyamzin@recursiveai.co.jp

Pablo Cervantes

Recursive Inc.

pablo.cervantes@recursiveai.co.jp

Tiago Ramalho

Recursive Inc.

tmramalho@recursiveai.co.jp

Abstract

We introduce SpLIIF to generate implicit neural representations and enable arbitrary downscaling of weather variables. We train a model from sparse weather stations and topography over Japan and evaluate in- and out-of-distribution accuracy predicting temperature and wind, comparing it to both an interpolation baseline and CorrDiff. We find the model to be up to 50% better than both CorrDiff and the baseline at downscaling temperature, and around 10-20% better for wind.

1 Introduction

Accurate, high-resolution weather data is a cornerstone for addressing the multifaceted challenges of global climate change. [1, 2] To meet this need, machine learning provides a powerful alternative to computationally intensive numerical models.[3] The field has seen rapid progress, evolving from CNNs,[4] to GANs,[5, 6] to more recent transformer-based architectures[7–10] and state-of-the-art diffusion models,[11–15] that excel at generating the realistic, high-frequency spatial details.

Despite these advances, a critical accuracy gap remains, stemming from a reliance on gridded reanalysis datasets. These datasets, while comprehensive, can diverge from the ground-truth weather captured by sparse, irregularly-spaced weather stations.[16] For effective climate adaptation—which demands precise, site-specific forecasting—this discrepancy can be a major hurdle. The key challenge, therefore, is to develop models that can directly leverage the high-fidelity signal of on-the-ground station data to drive more accurate and reliable predictions.

To bridge this gap, we introduce the Sparse Local Implicit Image Function (SpLIIF). While prior work has individually leveraged sparse station data,[17, 18] topography,[9] or continuous representations,[19, 20] SpLIIF uniquely combines all three, using an Implicit Neural Representation (INR)[21] to generate a continuous model of the weather—an approach well-suited to capturing the inherently continuous and non-linear dynamics of atmospheric systems. We compare this approach to a state-of-the-art model (CorrDiff)[14] that downscales dense weather maps using a U-Net CNN and generates physically consistent high-resolution results via diffusion. Our approach proves highly accurate and outperforms CorrDiff with a 50% improvement in temperature prediction and 20% in wind speed, particularly in complex terrain. Hence, our work contributes to the development of more

accurate weather models and therefore more effective and practical tools for global climate change adaptation and resilience.

2 Methods

Model architecture – SpLIIF can simultaneously process two types of input weather data: dense gridded data $X_d \in \mathbb{R}^{H \times W \times C_d}$ (e.g. from numerical weather prediction models), and/or irregularly spaced sparse point data $X_{sp} \in \mathbb{R}^{N \times C_{sp}}$ (e.g. observations from weather stations).

X_{sp} is first transformed into a dense tensor, $X'_d \in \mathbb{R}^{H' \times W' \times C_{sp}}$, using an Inverse Distance Weighting (IDW) interpolation scheme with learnable weights. If provided, X_d is interpolated to this intermediate resolution ($H' \times W'$) and concatenated along their channel dimension $\in \mathbb{R}^{H' \times W' \times (C_{sp} + C_d)}$. This raw feature representation passes through an MLP to project the features into an initial latent space representation $L_0 \in \mathbb{R}^{H' \times W' \times C_L}$. L_0 is then interpolated to the desired resolution and concatenated with the static, high-resolution topography data, $L_1 \in \mathbb{R}^{H'' \times W'' \times (C_{topo} + C_L)}$. This high-resolution latent space is then passed through an Enhanced Deep Super-Resolution (EDSR) network[22] that learns and sharpens spatial features, whilst retaining the size of the feature space.

Decoding and loss calculation – From the latent space, output weather variables can be decoded by interpolating the higher-dimensional space at N arbitrary query coordinates before passing through a final MLP that produces the output weather variables $\in \mathbb{R}^{N \times C_{out}}$. When calculating losses, the N locations are positions of weather stations not present in X_{sp} and used to calculate the $L1$ loss (during training) or the RMSE (during evaluation). This decoding step, where the higher-dimensional latent space is interpolated to potentially off-grid station locations is the key aspect that allows SpLIIF to learn a continuous representation of the weather.

Datasets – *JMA Weather Station Data (sparse input)*[23] (training and evaluation), hourly observations of temperature, wind speed and wind direction. Stations are distributed across the entirety of Japan. Median distance between adjacent JMA stations is 0.1 degrees. *Topography Data* (training and evaluation)[24], downsampled to 600m resolution. *ERA5 (dense input)* (only evaluation)[25] just 2m_temperature, u10m and v10m variables.

Training – The model was trained using hourly weather data collected in 2018, from 70% of the weather stations covering Japan with the exception of Hokkaido. The weather features used were temperature (normalized from $[-30, 40]^\circ\text{C}$ to $[-1, 1]$), wind strength (normalized from $[0, 30]\text{m/s}$ to $[0, 1]$), and wind direction (decomposed along u and v directions). To manage memory usage and increase batch diversity, this region was partitioned into square patches ($\sim 150 \times 150\text{km} \equiv H'' = W'' = 256$). Within each patch, subsets of up to 30 stations are selected, 80% of which are used as inputs, whilst the rest are used to calculate $L1$ loss on all variables. Batches comprised 10 patches, exposing the model to a range of topographical features and dynamic weather patterns in each training step, thereby promoting robust generalization. Training took 10h on a single T4 GPU (compared to a few thousand hours on A100s for CorrDiff).

3 Results

3.1 SpLIIF vs Baseline

To evaluate the model’s ability to generalize to unseen locations, we measured the Root Mean Square Error (RMSE) between SpLIIF’s predictions and ground-truth observations. The evaluation was performed by taking the same input stations as during the training, but comparing the predictions to the 30% of stations held out from the training set. We used a random 10% of time slices from the 2018 test year, and performed inference using multiple configurations, varying the number of nearby stations used as input. The results, shown in Figure 1, are presented as the percentage improvement of SpLIIF’s RMSE over an Inverse Distance Weighting (IDW) baseline, and plotted as functions of both the evaluation station’s altitude and the number of input stations.

SpLIIF consistently outperformed the baseline across all tested variables, with the magnitude of the improvement depending on the meteorological variable, station altitude, and the number of input stations. For temperature (Figure 1a), improvement was most pronounced at higher altitudes. While gains at sea level were modest (1-5%), they increased significantly with elevation, reaching between




Figure 1: Percentage improvement in RMSE of SpLIIF over the baseline by altitude and number of input stations for (a) temperature, (b) wind speed, and (c) wind angle. Shaded regions indicate the standard deviation on the mean across the different time slices and spatial patches.

25% and 50% at altitudes approaching 1000m. For wind speed (Figure 1b), SpLIIF achieved a 5% to 30% improvement over the baseline. Unlike temperature, this improvement did not show a monotonic trend with altitude, but it did increase as the number of input stations decreased. Finally, the improvement for wind angle (Figure 1c) was more modest ($\lesssim 10\%$), with a slight increase at higher altitudes and no clear dependence on the number of input stations.

These findings can be interpreted by considering the interplay between SpLIIF’s two primary data sources: the sparse station observations and the continuous, high-resolution topography map. In data-dense scenarios with many input stations, both model and baseline have sufficient data to construct an accurate downscaled representation. In data-sparse scenarios, however, the model increasingly leverages the topography to inform its predictions, which the baseline has no access to. Hence explaining why the largest performance gain occurs when observational data is limited. Additionally, the more substantial improvements seen for temperature are physically intuitive; temperature has a strong, direct relationship with elevation (lapse rate), which the model can effectively learn from the topography. In contrast, wind patterns are influenced by more complex, local aerodynamic effects that are not as easily captured by elevation alone.

3.2 SpLIIF vs CorrDiff

Next, we evaluated SpLIIF’s out-of-distribution generalization against a state-of-the-art baseline, CorrDiff. We compared the absolute error of both models on data from Hokkaido, a region unseen during training for either model, using ground-truth data from JMA stations for 10% of hours in 2023. For SpLIIF, we tested inference using both JMA station data and ERA5 reanalysis as inputs. Figure 2 plots the error distributions, where superior performance is indicated by a higher density of small errors and a lower density of large errors.

For temperature (Figure 2a), SpLIIF demonstrates a clear advantage, producing approximately twice as many predictions with an absolute error below 1°C and a tenfold reduction in errors exceeding 10°C , leading to a 50% reduction in the mean error. Notably, SpLIIF’s performance was even stronger when downscaling from ERA5 inputs, despite being trained exclusively on JMA station data. This result highlights the INR architecture’s ability to generalize to different input data sources without re-training. The performance gains for wind speed were more moderate, with SpLIIF showing 15% more low-error predictions and a threefold reduction in high-error predictions (Figure 2b), leading to a 20% reduction in the mean error. For wind direction (Figure 2c), SpLIIF using JMA inputs yielded the best results, with a 10% reduction in the mean error. We hypothesize that direct ground-truth observations provide a crucial signal for local wind dynamics—influenced by microscale topography—that is absent in coarser reanalysis data.

A qualitative comparison further highlights these differences. Figure 2d shows a downscaled cold front from April 2023. SpLIIF’s temperature map contains more high-resolution features that directly correspond to the underlying topography, details that are absent in the CorrDiff output. Similarly, SpLIIF’s wind fields show more complex patterns in mountainous regions, with wind directions aligning more closely with valley structures and contouring mountain ranges, demonstrating a superior ability to capture topographically-induced weather effects.




Figure 2: (a-c) Probability density functions (PDFs) of absolute errors for downscaled predictions CorrDiff, SpLIIF using sparse JMA station data as input (SpLIIF-JMA), and SpLIIF using ERA5 as input (SpLIIF-ERA5). Error distributions are shown for (a) temperature (°C), (b) wind speed (m/s), and (c) wind direction (°). (d) Example inference during a cold front. First column shows topography, and the next three columns respectively show the temperature and wind from ERA5, SpLIIF-ERA5 and CorrDiff. Bolded arrows are a guide-to-the-eye for wind along a valley.

4 Conclusion and future work

In this work, we introduced SpLIIF, a downscaling model that combines a continuous implicit neural representation with sparse station data and high-resolution topography. Our model significantly outperforms traditional baselines, reducing temperature RMSE by up to 50% in complex terrain, and surpasses a state-of-the-art diffusion model in out-of-distribution tests. These results highlight the value of integrating real-world, sparse data and demonstrate that INR-based architectures offer a flexible and powerful alternative to the grid-dependent backbones common in generative models. By generating more physically plausible weather fields from sparse inputs, SpLIIF contributes to the development of more effective tools for climate change adaptation.

Future work will enhance SpLIIF’s applicability to climate action by extending it to predict variables critical for extreme events, like precipitation, and by exploring hybrid architectures that integrate our data-driven approach with physically-constrained models.

References

1. Rampal, N. *et al.* Enhancing Regional Climate Downscaling through Advances in Machine Learning. en. *Artificial Intelligence for the Earth Systems*. <https://journals.ametsoc.org/view/journals/aies/3/2/AIES-D-23-0066.1.xml> (Apr. 2024).
2. Miller, S., Ormaza-Zulueta, N., Koppa, N. & Dancer, A. Statistical downscaling differences strongly alter projected climate damages. en. *Commun Earth Environ* **6**, 145. ISSN: 2662-4435. <https://www.nature.com/articles/s43247-025-02134-2> (Feb. 2025).

3. Bracco, A. *et al.* Machine learning for the physics of climate. en. *Nat Rev Phys* **7**, 6–20. ISSN: 2522-5820. <https://www.nature.com/articles/s42254-024-00776-3> (Jan. 2025).
4. Höhle, K., Kern, M., Hewson, T. & Westermann, R. A comparative study of convolutional neural network models for wind field downscaling. en. *Meteorological Applications* **27**, e1961. ISSN: 1469-8080. <https://onlinelibrary.wiley.com/doi/abs/10.1002/met.1961> (2020).
5. Leinonen, J., Nerini, D. & Berne, A. Stochastic Super-Resolution for Downscaling Time-Evolving Atmospheric Fields With a Generative Adversarial Network. *IEEE Transactions on Geoscience and Remote Sensing* **59**, 7211–7223. ISSN: 1558-0644. <https://ieeexplore.ieee.org/abstract/document/9246532> (Sept. 2021).
6. Oyama, N., Ishizaki, N. N., Koide, S. & Yoshida, H. Deep generative model super-resolves spatially correlated multiregional climate data. en. *Sci Rep* **13**, 5992. ISSN: 2045-2322. <https://www.nature.com/articles/s41598-023-32947-0> (Apr. 2023).
7. Pérez, A., Cruz, M. S., Martín, D. S. & Gutiérrez, J. M. *Transformer based super-resolution downscaling for regional reanalysis: Full domain vs tiling approaches* Oct. 2024. <http://arxiv.org/abs/2410.12728>.
8. Sinha, S., Benton, B. & Emami, P. On the effectiveness of neural operators at zero-shot weather downscaling. en. *Environmental Data Science* **4**, e21. ISSN: 2634-4602. <https://www.cambridge.org/core/journals/environmental-data-science/article/on-the-effectiveness-of-neural-operators-at-zeroshot-weather-downscaling/70E4F46811337C0DF5BC99AEEB696544> (Jan. 2025).
9. Liu, Z. *et al.* MambaDS: Near-Surface Meteorological Field Downscaling With Topography Constrained Selective State-Space Modeling. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–15. ISSN: 1558-0644. <https://ieeexplore.ieee.org/document/10752514> (2024).
10. Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J. K. & Grover, A. *ClimaX: A foundation model for weather and climate* Dec. 2023. <http://arxiv.org/abs/2301.10343>.
11. Huang, L., Gianinazzi, L., Yu, Y., Dueben, P. D. & Hoefler, T. *DiffDA: a diffusion model for weather-scale data assimilation* in *Proceedings of the 41st International Conference on Machine Learning* **235** (JMLR.org, Vienna, Austria, July 2024), 19798–19815.
12. Sundar, R., Hu, Y., Parashar, N., Blanchard, A. & Dodov, B. *TAUDiff: Highly efficient kilometer-scale downscaling using generative diffusion models* Mar. 2025. <http://arxiv.org/abs/2412.13627>.
13. Merizzi, F., Asperti, A. & Colamonaco, S. Wind speed super-resolution and validation: from ERA5 to CERRA via diffusion models. en. *Neural Comput & Applic* **36**, 21899–21921. ISSN: 1433-3058. <https://doi.org/10.1007/s00521-024-10139-9> (Dec. 2024).
14. Mardani, M. *et al.* Residual corrective diffusion modeling for km-scale atmospheric downscaling. en. *Commun Earth Environ* **6**, 124. ISSN: 2662-4435. <https://www.nature.com/articles/s43247-025-02042-5> (Feb. 2025).
15. Brenowitz, N. D. *et al.* *Climate in a Bottle: Towards a Generative Foundation Model for the Kilometer-Scale Global Atmosphere* July 2025. <http://arxiv.org/abs/2505.06474>.
16. Vanella, D. *et al.* Comparing the use of ERA5 reanalysis dataset and ground-based agrometeorological data under different climates and topography in Italy. *Journal of Hydrology: Regional Studies* **42**, 101182. ISSN: 2214-5818. <https://www.sciencedirect.com/science/article/pii/S2214581822001951> (Aug. 2022).
17. Liu, Z. *et al.* *Kolmogorov Arnold Neural Interpolator for Downscaling and Correcting Meteorological Fields from In-Situ Observations* Jan. 2025. <http://arxiv.org/abs/2501.14404>.
18. Manshausen, P. *et al.* *Generative Data Assimilation of Sparse Weather Station Observations at Kilometer Scales* Apr. 2025. <http://arxiv.org/abs/2406.16947>.
19. Luo, Y., Fang, S., Wu, B., Wen, Q. & Sun, L. *Physics-Guided Learning of Meteorological Dynamics for Weather Downscaling and Forecasting* in *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2* (Aug. 2025), 2010–2020. <http://arxiv.org/abs/2505.14555>.
20. Radhakrishnan, G. *et al.* *Continuous latent representations for modeling precipitation with deep learning* en-US. in *Climate Change AI* (Climate Change AI, Dec. 2024). <https://www.climatechange.ai/papers/neurips2024/24>.

21. Chen, Y., Liu, S. & Wang, X. *Learning Continuous Image Representation With Local Implicit Image Function* en. in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021), 8628–8638. https://openaccess.thecvf.com/content/CVPR2021/html/Chen_Learning_Continuous_Image_Representation_With_Local_Implicit_Image_Function_CVPR_2021_paper.html.
22. Lim, B., Son, S., Kim, H., Nah, S. & Mu Lee, K. *Enhanced Deep Residual Networks for Single Image Super-Resolution* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), 136–144. https://openaccess.thecvf.com/content_cvpr_2017_workshops/w12/html/Lim_Enhanced_Deep_Residual_CVPR_2017_paper.html?ref=https://githubhelp.com.
23. Japan Meteorological Agency. *AMeDAS* <https://www.jma.go.jp/jma/en/Activities/amedas/amedas.html>.
24. European Space Agency. *Copernicus Global Digital Elevation Model* 2024. <https://doi.org/10.5270/ESA-c5d3d65>.
25. Hersbach, H. *et al.* ERA5 hourly data on single levels from 1940 to present. *Copernicus Climate Data Store*. 10.24381/cds.adbb2d47 (2023).