Transferability of UNet-Based Downscaling Model for High-Resolution Temperature Data Across Diverse Regions

Marc Benitez[1,2][0009-0006-2072-9124] *, Mirta Rodriguez[2][0009-0008-1767-8711], Javier Panadero[1][0000-0002-3793-3328], Omjyoti Dutta[0009-0005-4333-1883], and Tomas Margalef[1][0000-0001-6384-7389]

[1] Computer Architecture and Operating Systems Department. Universitat Autònoma de Barcelona. C/ de les Sitges, s/n 08193 Cerdanyola del Vallès, Spain
marc.benitezb@autonoma.cat, {javier.panadero, tomas.margalef}@uab.cat
[2] Mitiga Solutions S.L., Carrer de Julià Portet, 3, 08002 Barcelona, Spain
{marc.benitez, mirta.pinilla}@mitigasolutions.com

Abstract. High-resolution atmospheric data is essential for understanding local atmospheric processes, however it is computationally expensive to achieve such high resolutions through physical models. Recently, deep learning techniques, particularly those used in Single Image Super-Resolution, have emerged as a promising approach for statistical downscaling. However, much of the existing research has focused on enhancing model performance within small geographical regions, with limited attention given to the transferability of these models to diverse areas outside of their training domain. This paper introduces a methodology that evaluates the ability of a UNet model to downscale daily 2-meter temperature data outside its training region. The proposed approach uses one-third of the Contiguous United States to train the model, and assesses its performance on unseen areas. Our experimental design deliberately tests both spatial and temporal generalization, demonstrating that relatively compact models can effectively transfer downscaling capabilities to new regions. This results in improvements across key performance metrics including Mean Absolute Error, Root Mean Square Error, and Peak Signal-to-Noise Ratio. Additionally, our approach significantly reduces computational costs while improving downscaling accuracy across diverse climatic and topographic conditions.

Keywords: Downscaling · UNet · Temperature · Transfer learning

1 Introduction

High-resolution meteorological and climate data is essential for advancing our understanding of local atmospheric processes, such as extreme weather events, and future climatic scenarios. However, this local-scale data is often unavailable and

* Corresponding author: marc.benitezb@autonoma.cat

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

must be generated through downscaling, a process that refines coarse-resolution data into finer-scale information.

Downscaling techniques can be categorized into two approaches: dynamical downscaling and statistical downscaling. Dynamical downscaling uses numerical weather prediction models which solve complex physical equations to simulate atmospheric conditions at high resolutions. While these methods are physically consistent, they are computationally expensive, making them unfeasible for large areas or long-term simulations. In contrast, statistical downscaling uses historical relationships between coarse-scale and local-scale data to infer finer details. Although these methods are computationally cheaper than dynamical downscaling, they can produce overly smooth maps and less accurate results, limiting its usage in certain applications.

Recently, deep learning (DL) methods used in the computer vision task of single image super-resolution [15] have gained popularity as statistical downscaling techniques. The use of architectures, such as Convolutional Neural Networks (CNNs), UNets [13], Generative Adversarial Networks (GAN) [1, 5], and more recently, Diffusion Models (DM) [8], has spread across the community, demonstrating promising results in several studies. Generative architectures like GAN and DM even have the ability to produce ensembles, a highly sought-after feature in downscaling and weather forecasting applications. Despite the growing adoption of DL architectures, most efforts in the field have focused on improving model performance using different architectures for small regions [8, 6, 10, 7, 2], validating the models within the same small region in years outside of the training set. There has been limited work on analysing and understanding the ability of these models to transfer knowledge to larger areas outside of their training region [9].

This paper introduces a methodology that evaluates the effects of training a DL model (UNet) to downscale daily 2-meter temperature in a relatively small area and use the resulting model to produce predictions in unseen regions with diverse climatic and topographic features. We choose the Contiguous United States (CONUS) as our training region, using only one third of the available area to train the model, reducing the cost and time needed to train it. The rest of the data is used exclusively for validation and testing. We study both the spatial and temporal generalization capabilities of the downscaling model by also selecting new years for evaluation. The results show that small models are capable of transferring their downscaling skill to unseen areas, reducing metrics such as Mean Absolute Error (MAE), Root Mean Square Error (RMSE) and Peak Signal-To-Noise Ratio (PSNR).

2 Proposed Methodology

This section outlines the main components of the study, illustrated in Fig. 1. We begin by defining the study region, followed by a description of the datasets used and variables selected (Fig. 1a). Next, we define the model architecture and training details. Subsequently, we explain the spatial and temporal data

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

partitioning, shown in Fig. 1b. Finally, we introduce the evaluation metrics used to assess model performance.

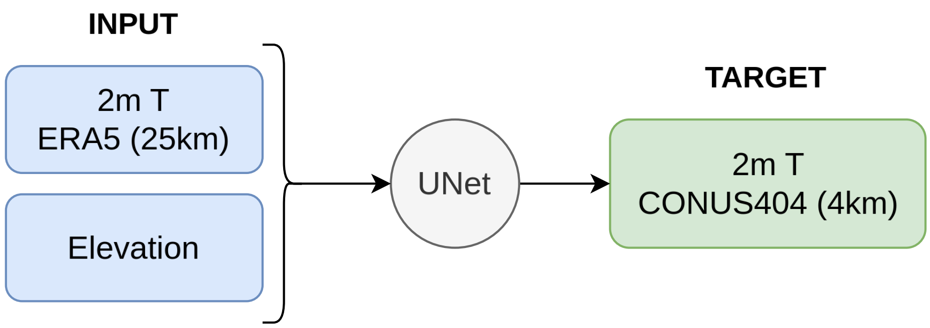

(a) Sketch of the training process. From left to right; input data (low-resolution temperature field from ERA5 and elevation from CONUS404), model architecture and target data (high-resolution temperature field from CONUS404.

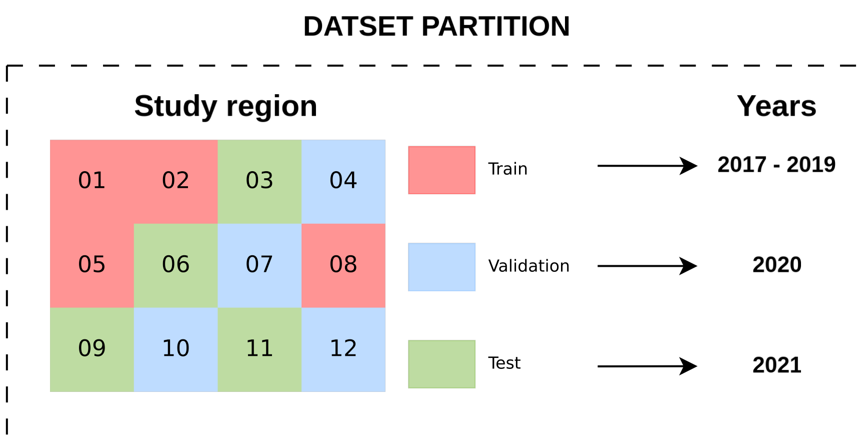

(b) Image showing the spatial and temporal splitting of the dataset. The left side shows in red the training set (splits 01, 02, 05, 08), in blue the validation set (splits 04, 07, 10, 12) and in green the test set (splits 03, 06, 09, 11). On the right side we show the training (2017-2019), validation (2020) and test years (2021).

Fig. 1: Overview of the proposed methodology (a) and dataset temporal and spatial splitting (b).

2.1 Region of Study

Our study focuses on a large domain to evaluate the temporal and spatial generalization capabilities of our downscaling method. The study region is located on

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

the CONUS but it also includes parts of Canada, Mexico and other countries, such as Cuba and the Bahamas (Fig 2). This area has been selected due to the availability of a high-resolution dataset that we use as our target, this dataset is presented in section 2.2. The region covers 5,472 km from west to east and 4,064 km from north to south at 4km resolution, which in terms of pixel size, translates to 1368 by 1016 pixels. Working with such a large domain ensures that we are able to sample from a variety of climatic and topographically diverse regions to study the ability of the models to generalize across diverse environments.

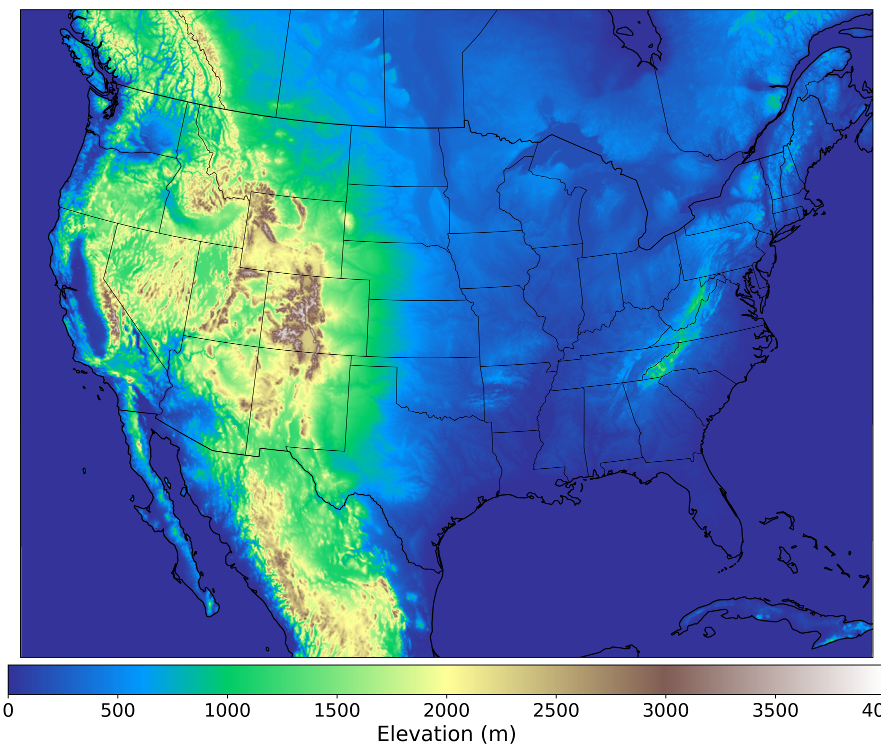

Fig. 2: Spatial boundaries and topographical features of the study region. The image shows the high-resolution elevation map that has been used as input for our model and during the creation of the target dataset (CONUS404).

2.2 Datasets and Selected Variables

The ERA5 dataset [3], produced by the European Center for Medium-Range Weather Forecast (ECMWF) using the Integrated Forecast System (IFS), provides hourly global meteorological data at 0.25° spatial resolution. The dataset includes a wide range of atmospheric variables at single levels (2D) and pressure

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

or model levels (3D) and its temporal coverage starts from 1940 to near real time.

The CONUS404 dataset [12], developed through a collaboration between the National Center for Atmospheric Research (NCAR) and the U.S. Geological Survey (USGS), is a high-resolution hydro-climate dataset covering the CONUS area at a 4km resolution during a 40-year period. The dataset consists of Weather Research and Forecasting model outputs that use ERA5 as its boundary conditions to force the simulation.

Both datasets use different grids systems: ERA5 uses a regular latitude-longitude grid, while CONUS404 uses a Lambert Conformal Conic projection with a 4km grid. To homogenize the projections of both datasets we use the Climate Data Operator (CDO) library [14] to bilinearly interpolate ERA5 onto the target grid. This interpolated data serves both as input and the baseline for our DL model.

This work focuses on downscaling daily 2-meter temperature across the study region in the 2017-2021 period. Additionally, we incorporate elevation data from the target dataset to represent topographical features, helping the model learn how these features interact with daily 2-meter temperature and transfer its results onto unseen regions.

2.3 Model

Our DL model is based on the UNet architecture. We implement this model using the Segmentation Models Python library [4], which allows us to change the UNet encoder for other models, such as the RegNetY [11] family of networks.

The UNet architecture consists of two components: the encoder (contracting path) and the decoder (expanding path). The encoder, in our case a RegNetY, captures contextual information, extracting features through a series of convolutional layers with grouped convolutions and squeeze-and-excitation blocks. The layers progressively reduce spatial dimensions while increasing feature depth at each step, learning hierarchical representations of the input. The decoder reconstructs the spatial resolution through upsampling operations and skip connections, preserving details that would otherwise be lost during downsampling.

We choose RegNetY architecture as our encoder due to its flexibility in size and improved feature extraction capabilities. The Segmentation Models library uses the encoders implemented in the PyTorch Image Models Python library [16] that implements 15 different sizes of the RegNetY encoder, allowing us to evaluate the performance of various encoders without changing the architecture.

The resulting model uses the encoder named "regnety_032" and has 24M trainable parameters. We trained it for 250 epochs with a batch size of 128, which took 2 hours on a single NVIDIA H100 GPU (64GB) at the Barcelona Supercomputing Center. We used the Adam optimizer with an initial learning rate of 0.0005 and a Cosine Annealing Warm Restarts scheduler (T_0=20, T_{mult}=2). For the objective function, we used the Mean Square Error (MSE) loss. All data has been standardized by subtracting the mean and dividing by the standard deviation of the train set.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

2.4 Dataset Partition

A large training domain typically provides the model with a diverse set of meteorological patterns and spatial variability, which can enhance its ability to capture complex relationships and improve generalization to unseen regions. However, this comes at the cost of increased computational resources and hardware requirements. Training with smaller areas lowers the computational needs and allows us to study the generalization capabilities of the model by partitioning the dataset spatially and reserving some areas only for evaluation purposes.

As mentioned in section 2.1, the full region of study has a shape 1368x1016 pixels, where each pixel represents a 4x4km area. The leftmost part of the initial area (Pacific Ocean) was dropped since it only contained ocean pixels. The resulting area is then divided into 12 smaller regions of 310x310 pixels (Fig 3), totalling around 1240x1240km per tile. The inputs of the network are 256x256 random crops of the aforementioned regions.

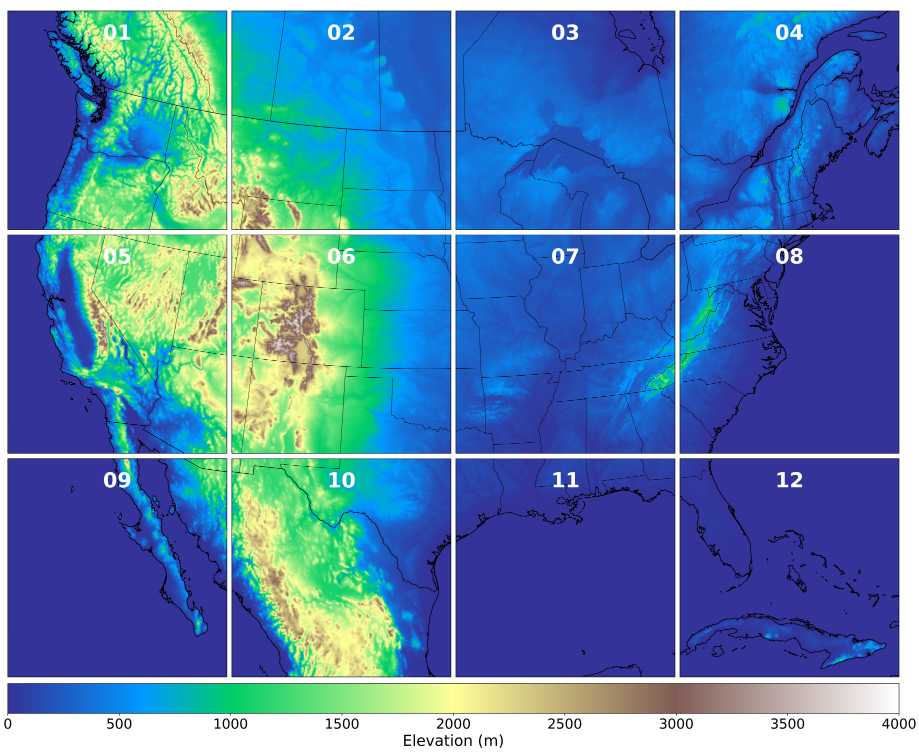

Fig. 3: Patches created from the original study region after eliminating part of the Pacific Ocean. The white number on each tile represents the identifier used to separate the tiles into training, test and validation splits (shown in Fig.1b and Table 1).

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

Each data split consists of non-overlapping years and spatial areas (Table 1). The model is trained on data from years 2017, 2018, and 2019. The validation uses year 2020 and testing year 2021. The dataset is spatially partitioned as mentioned above and each resulting region is exclusively used for training, validation or testing. Regions 01, 02, 05, and 08 are used in the training split, totalling 4,380 samples. Validation uses regions 04, 07, 10, and 12 (1464 samples). Finally, regions 03, 06, 09, and 11 are reserved for testing.

Table 1: Separation of the years and areas (defined in Fig. 1b and 3) used for the train, validation, and test partitions. The last column shows the number of exclusive samples (patches) of that split.

	Years	**Reserved Areas**	**# Exclusive samples**
Train	2017, 2018, 2019	01, 02, 05, 08	4380
Validation	2020	04, 07, 10, 12	1464
Test	2021	03, 06, 09, 11	1460

2.5 Evaluation Metrics

Three metrics are used to evaluate the quality of the downscaled temperature fields: MAE, RMSE and PSNR.
MAE measures the average magnitude of the absolute differences between predicted and target values, without considering their direction.

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{1}$$

where y_i represents the target value (i.e. CONUS404), $\hat{y}_i$ the predicted value and n the number of observations, in this case pixels.

RMSE calculates the square root of the average squared differences between predicted and target values. By squaring the errors before averaging, RMSE gives more weight to larger deviations.

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{2}$$

where y_i represents the target value, $\hat{y}_i$ the predicted value and n the number of observations, in this case pixels.

PSNR is a commonly used metric in the field of image processing and compression to evaluate the quality of a reconstructed or predicted image compared to the original. It quantifies the ratio between the maximum possible power of

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

the signal (target image) and the power of the noise (error introduced of the prediction).

$$\mathrm{PSNR} = 10 \cdot \log_{10}\left(\frac{\mathrm{MAX}_I^2}{\mathrm{MSE}}\right) \tag{3}$$

where MAX_I is the maximum possible pixel value of the image.

Additionally, Power Spectral Density (PSD) is used as a measure of the spatial variability of temperature patterns across different scales. The 2-meter temperature fields are decomposed into different spatial frequency components, which reveal the representation of scale-dependent structure within the data. The power spectra is computed by applying a 1D Fast Fourier Transform to each row (latitude) of the temperature field and squaring the result. Finally the average of the obtained power spectra is calculated over all rows to produce the power spectrum.

3 Experimental Results

Results presented in this section are obtained using year 2021 in the entire region of study. Table 2 shows the results of both methods (bilinear interpolation and UNet) divided by season and data split (i.e. train, test and validation). The UNet outperforms bilinear interpolation in annual metrics across all splits, though the degree of improvement is not uniform for each split and season. For MAE, the performance gain of the UNet is 0.23 (16%), 0.15 (12%) and 0.02 (2%) for train, validation and test splits respectively. Similar behaviour can be observed in RMSE, where the increase is slightly higher at 0.33 (17%), 0.23 (13%) and 0.10 (6%) for each split. The difference between MAE and RMSE gain implies that the UNet reduces large errors across all regions, but fails to correct smaller errors. This pattern is expected with the selected loss function, MSE loss, which penalizes large errors more heavily than smaller ones. The PSNR metric shows a similar trend, with the test split achieving considerably less performance gain (0.22) when compared to train (2.10) and validation (0.85) splits.

Seasonal analysis reveals that the performance of the UNet is better than bilinear interpolation in most seasons. A decrease in performance gain is observed in the summer and autumn seasons, due to the limitations of the UNet model in areas where the overall baseline errors are lower.

Fig. 4 shows the spatial distribution of MAE (top) and RMSE (bottom) across the study region. In the case of MAE maps, there are large discrepancies between bilinear interpolation and CONUS404 on mountainous regions, such as the Rocky Mountains and the Appalachian Mountains. On the other hand, the UNet has learned the relationship between elevation and temperature change, successfully correcting the larger errors, however the correction is overly smooth for orographically complex areas, as it can be observed in the Rocky Mountains where the UNet has increased MAE values in some locations. This effect is caused, as previously explained, by the MSE loss function. The same pattern is present in the RMSE maps, where errors are consistently lower along the East coast compared to the West coast. Although the UNet reduces extreme errors

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

Table 2: Comparison of MAE, RMSE and PSNR metrics between baseline and UNet models for test year 2021 and across seasons (DJF: Winter, MAM: Spring, JJA: Summer, SON: Autumn) for train, validation and test splits.

		Model	**DJF**	**MAM**	**JJA**	**SON**	**Annual**
MAE↓	Train	Interpolation	1.75	1.34	1.26	1.20	1.39
		UNet	1.43	0.99	1.12	1.11	1.16
	Validation	Interpolation	1.61	1.36	0.91	1.06	1.23
		UNet	1.42	0.99	0.90	1.02	1.08
	Test	Interpolation	1.41	1.09	0.93	0.96	1.10
		UNet	1.27	1.01	0.98	1.05	1.08
RMSE↓	Train	Interpolation	2.47	1.91	1.79	1.69	1.98
		UNet	2.08	1.35	1.55	1.51	1.65
	Validation	Interpolation	2.27	1.85	1.25	1.42	1.74
		UNet	2.02	1.35	1.20	1.35	1.51
	Test	Interpolation	2.13	1.62	1.31	1.37	1.64
		UNet	1.90	1.42	1.37	1.43	1.54
PSNR↑	Train	Interpolation	30.02	32.13	33.21	33.34	32.18
		UNet	32.32	35.25	34.77	34.74	34.28
	Validation	Interpolation	29.91	31.17	33.66	32.68	31.86
		UNet	30.53	33.31	33.94	33.03	32.71
	Test	Interpolation	31.97	33.47	34.64	34.50	33.65
		UNet	32.39	34.19	34.75	34.12	33.87

in mountainous regions, the smoothing effect of the predictions introduces the slightly larger errors that we observe in the Western CONUS.

Fig. 5 shows the power spectra for ERA5, CONUS404 and UNet prediction. The image shows that the UNet is more capable of capturing local features than bilinear interpolation. Specifically, bilinear interpolation begins to decline around the 10^{-2} frequency, which represents patterns finer than 100km. In contrast, the UNet model produces a power spectra that closely resembles the target one, only declining right before 10^{-1} frequency.

Fig. 6 compares monthly 2-meter temperature distributions at Mount Harvard (38.93°N, 106.32°W in tile 06 from the test split), the highest point in our elevation map at 3,827m. The bilinear interpolation trend shows that this method does not capture the effect of elevation on 2-meter temperature. In contrast, the UNet does capture the trend, corrects the values and even causes a cold bias by overcorrecting during summer months, specially in July and August.

Fig. 7 presents one example for each season of UNet downscaled 2-meter temperature fields and its ERA5 and CONUS404 counterparts. All ERA5 images are initially blurry, but the UNet model has largely improved the representation

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

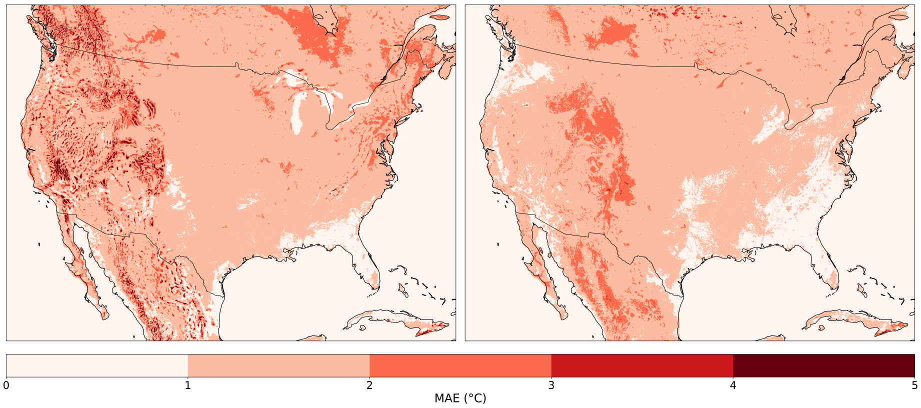

(a) Comparison of MAE for bilinear interpolation (left) and our model (right) during year 2021. The UNet shows consistent results across unseen regions and demonstrates significant skill in mountainous areas.

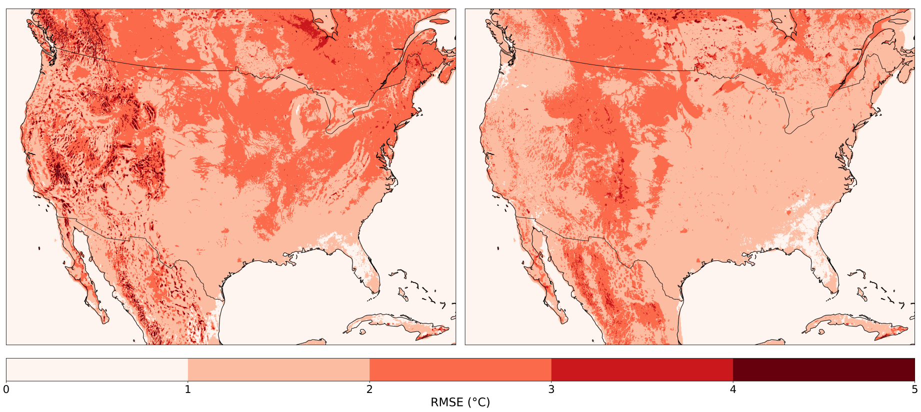

(b) Comparison of RMSE for bilinear interpolation (left) and our model (right) during year 2021. The RMSE maps show that the UNet is capable of mitigating the biases across all the region, specially on the northeastern part of the study region.

Fig. 4: Error maps of the study region for year 2021. Dark reds indicate high mean error, while lighter red tones represent lower error values.

of temperature fields over mountainous regions. While not of the same quality of the CONUS404 groundtruth, the results of our model are promising since it is capable of producing consistent results outside of the train regions.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

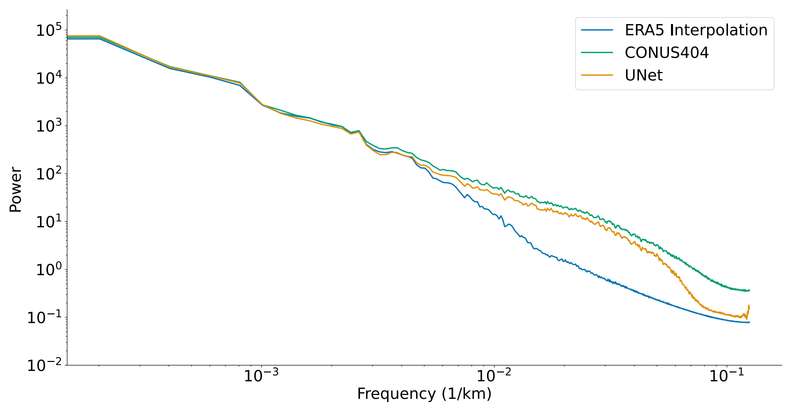

Fig. 5: Power spectra of the interpolated ERA5 input (blue), CONUS404 (green) and UNet (orange), computed over the whole region for year 2021.

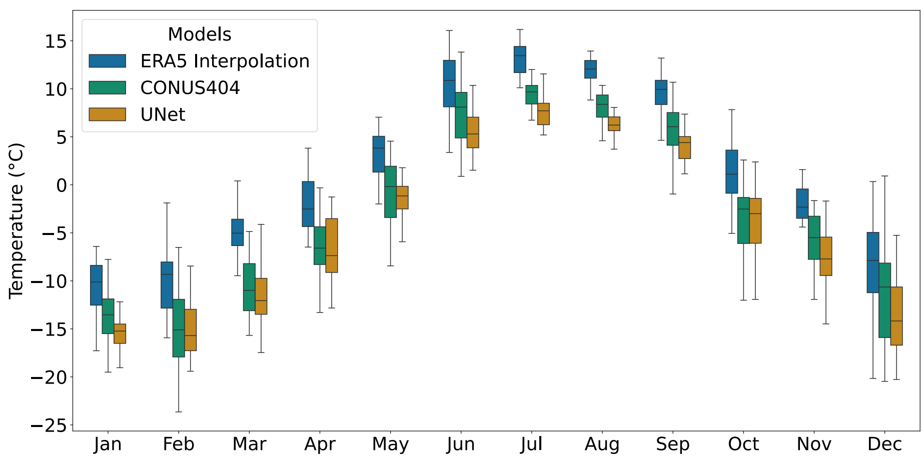

Fig. 6: Distributions of ERA5, CONUS404 and our UNet for daily 2-meter temperature (ºC) values at Mount Harvard grouped by month for year 2021.

4 Conclusion and Future Work

We successfully trained a small UNet (24M parameters) for downscaling daily 2-meter temperature from 25km to 4km in the CONUS region, using only one third of the inference region. Our evaluation shows that the model outperforms bilinear interpolation, this increase in performance is maintained in areas outside of the train region, opening the way to produce large-scale downscaling.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

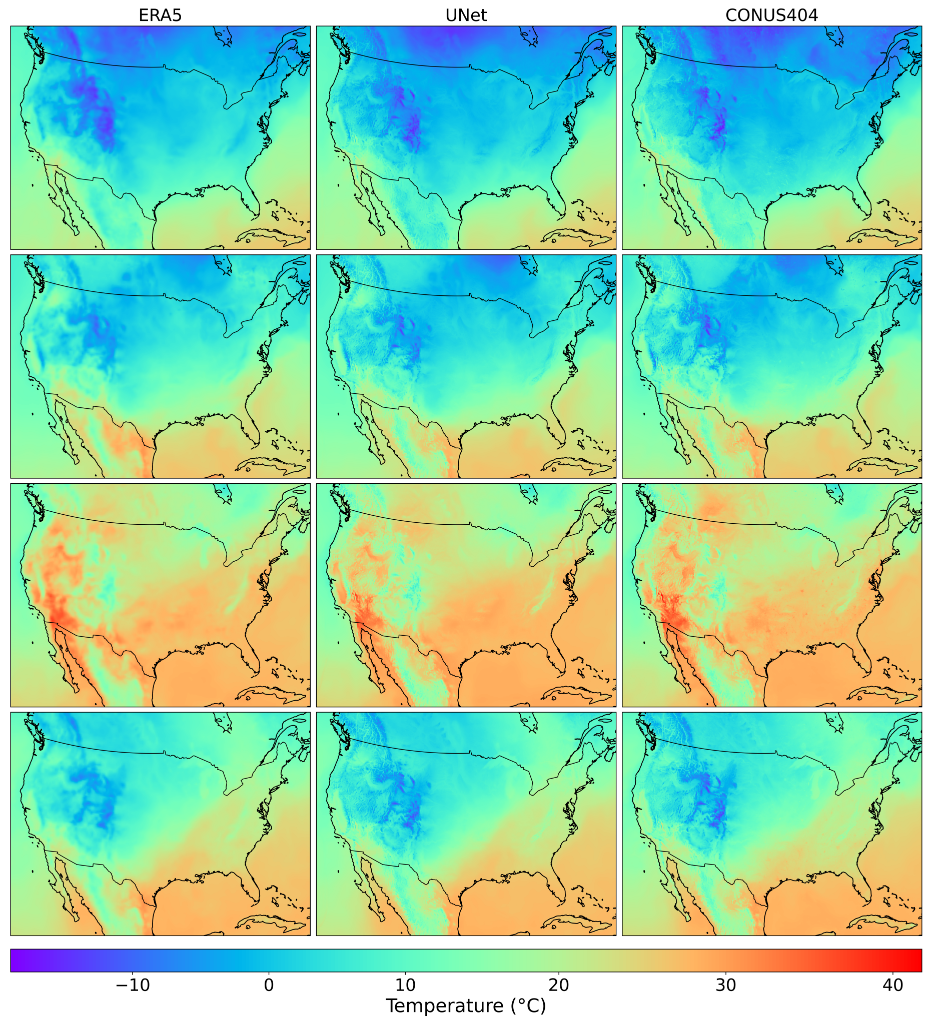

Fig. 7: Four samples extracted from selected days, starting from top to bottom: 2021-01-15, 2021-04-15, 2021-07-15 and 2021-10-15. (Left) ERA5 2-meter temperature fields, (Middle) UNet predicted 2-meter temperature field, (Right) CONUS404 2-meter temperature field

Future work includes an analysis of the temporal requirements by transitioning from daily data to 3-hourly data. This change aims to improve the model's ability to capture finer temporal patterns and to improve its seasonal and overall performance.

We also plan to experiment with the loss function, encouraging the presence of high-frequency details by adding extra penalization to blurry results.

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

Finally, we will benchmark our results to other DL architectures or ensemble methods, which could result in enhanced predictive performance. We also plan to extend this methodology to other variables (precipitation and wind speed) and study regions.

Acknowledgments. This work is supported by the Industrial Doctorate Plan of the Department of Research and Universities of the Generalitat de Catalunya under contract 2023DI00012. It also has been granted by the Spanish Ministry of Science and Innovation MCIN AEI/10.13039/501100011033 under contract PID2023-146193OB-I00 and by the Catalan government under grant 2021 SGR-00574. The model presented in this study were supported by the Barcelona Supercomputing Center (BSC) and Red Española de Supercomputación (RES) under project Inno-2-2024-0003.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. Commun. ACM **63**(11), 139–144 (Oct 2020). https://doi.org/10.1145/3422622, https://doi.org/10.1145/3422622
2. Harder, P., Hernandez-Garcia, A., Ramesh, V., Yang, Q., Sattegeri, P., Szwarcman, D., Watson, C., Rolnick, D.: Hard-constrained deep learning for climate downscaling. Journal of Machine Learning Research **24**(365), 1–40 (2023), http://jmlr.org/papers/v24/23-0158.html
3. Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., Chiara, G.D., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R.J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., Thépaut, J.N.: The era5 global reanalysis. Quarterly Journal of the Royal Meteorological Society **146**, 1999–2049 (7 2020). https://doi.org/10.1002/QJ.3803
4. Iakubovskii, P.: Segmentation models pytorch. https://github.com/qubvel/segmentation_models.pytorch (2019)
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**(9), 7211–7223 (2021). https://doi.org/10.1109/TGRS.2020.3032790
6. Mardani, M., Brenowitz, N., Cohen, Y., Pathak, J., Chen, C.Y., Liu, C.C., Vahdat, A., Nabian, M.A., Ge, T., Subramaniam, A., Kashinath, K., Kautz, J., Pritchard, M.: Residual corrective diffusion modeling for km-scale atmospheric downscaling (2024), https://arxiv.org/abs/2309.15214
7. Baño Medina, J., Manzanas, R., Gutiérrez, J.M.: Configuration and intercomparison of deep learning neural models for statistical downscaling. Geoscientific Model Development **13**(4), 2109–2124 (2020). https://doi.org/10.5194/gmd-13-2109-2020, https://gmd.copernicus.org/articles/13/2109/2020/

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14

8. Merizzi, F., Asperti, A., Colamonaco, S.: Wind speed super-resolution and validation: from era5 to cerra via diffusion models. Neural Computing and Applications **36**(34), 21899–21921 (Dec 2024). https://doi.org/10.1007/s00521-024-10139-9, https://doi.org/10.1007/s00521-024-10139-9
9. Prasad, A., Harder, P., Yang, Q., Sattegeri, P., Szwarcman, D., Watson, C., Rolnick, D.: Evaluating the transferability potential of deep learning models for climate downscaling (2024), https://arxiv.org/abs/2407.12517
10. Pérez, A., Cruz, M.S., Martín, D.S., Gutiérrez, J.M.: Transformer based super-resolution downscaling for regional reanalysis: Full domain vs tiling approaches (2024), https://arxiv.org/abs/2410.12728
11. Radosavovic, I., Kosaraju, R.P., Girshick, R., He, K., Dollár, P.: Designing network design spaces. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10425–10433 (2020). https://doi.org/10.1109/CVPR42600.2020.01044
12. Rasmussen, R.M., Chen, F., Liu, C.H., Ikeda, K., Prein, A., Kim, J., Schneider, T., Dai, A., Gochis, D., Dugger, A., Zhang, Y., Jaye, A., Dudhia, J., He, C., Harrold, M., Xue, L., Chen, S., Newman, A., Dougherty, E., Abolafia-Rosenzweig, R., Lybarger, N.D., Viger, R., Lesmes, D., Skalak, K., Brakebill, J., Cline, D., Dunne, K., Rasmussen, K., Miguez-Macho, G.: Conus404 the ncar–usgs 4-km long-term regional hydroclimate reanalysis over the conus. Bulletin of the American Meteorological Society **104**, E1382–E1408 (2023). https://doi.org/10.1175/BAMS-D-21-0326.1
13. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. pp. 234–241. Springer International Publishing, Cham (2015)
14. Schulzweida, U.: Cdo user guide (Oct 2023). https://doi.org/10.5281/zenodo.10020800, https://doi.org/10.5281/zenodo.10020800
15. Wang, Z., Chen, J., Hoi, S.C.H.: Deep learning for image super-resolution: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence **43**(10), 3365–3387 (2021). https://doi.org/10.1109/TPAMI.2020.2982166
16. Wightman, R.: Pytorch image models. https://github.com/rwightman/pytorch-image-models (2019). https://doi.org/10.5281/zenodo.4414861

ICCS Camera Ready Version 2025
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-97564-6_14