Hyperlocal Near-Surface Wind Downscaling as a Building Block for Dynamic Line Rating

David Lajevec[1,3], Andrej Vilhar[2], Aleksandra Rashkovska[1], and Dragi Kocev[1]

[1] Jožef Stefan Institute, Ljubljana, Slovenia
[2] Operato d.o.o., Ljubljana, Slovenia
[3] University of Ljubljana, Faculty of Mathematics and Physics, Ljubljana, Slovenia
dragi.kocev@ijs.si

Abstract. Dynamic line rating (DLR) increases the usable capacity of overhead power lines by estimating, from the local weather, the maximum admissible current of each span. The dominant and least predictable weather input is the near-surface wind, which numerical weather prediction (NWP) resolves only on a grid too coarse for a single span. We adapt a convolutional, terrain-aware statistical downscaling approach to DLR and study it on German DWD observations and ICON-EU forecasts, with evaluation under a stratified, spatially aware station split that measures generalisation to unseen stations. The model is a three-branch multimodal convolutional neural network that fuses the NWP fields with terrain descriptors at two resolutions. It improves wind estimates substantially over the interpolated NWP by exploiting high-resolution terrain, with the high-resolution terrain branch contributing most of the gain; the NWP side needs only a compact low-level wind input. The ultimate goal of this work is to improve these methods and provide a sound basis for more informed and reliable decision making by transmission system operators.

Keywords: Dynamic line rating · Statistical downscaling · Near-surface wind · Convolutional neural networks · Multimodal learning.

1 Introduction

Electric power systems face a steadily growing need to transmit more energy. The expansion of weather-dependent renewable generation, which is spatially more dispersed than conventional generation, together with electrification of demand, increases the load on existing transmission corridors. Because building new overhead lines is expensive, slow, and spatially constrained, it is attractive to use the existing grid closer to its physical limits without compromising safety. Dynamic line rating (DLR), also referred to as dynamic thermal rating, addresses this need by estimating, from a physical conductor model and the surrounding weather, the maximum admissible current, called the ampacity, of each line span [1].

The ampacity is governed by the conductor heat balance and is therefore sensitive to local weather, most strongly to wind speed and wind direction [1, 16].

The limiting span of a line is usually its most exposed or most sheltered one, so a useful DLR estimate requires weather information at the level of an individual span. Numerical weather prediction (NWP) models provide such information only on a relatively coarse spatial grid. However, at the level of a single span, local effects of terrain, vegetation, and line orientation are not resolved.

A physically grounded way to add spatial detail is dynamical downscaling, in which a higher-resolution model is nested inside the NWP output. For operating many spans frequently this is, however, computationally expensive. Statistical downscaling instead learns a mapping from past forecast-observation pairs and is much cheaper, while also correcting location-specific systematic NWP errors. Convolutional approaches that take the local topography as an additional input and learn to correct the NWP wind field according to orographic features have proven particularly effective for near-surface wind in complex terrain [3, 12, 11]. Related work has also demonstrated the full path from NWP to a per-span DLR estimate [13].

We take the convolutional, terrain-aware downscaling approach of [3] as a starting point rather than reproduce it verbatim, and adapt it to our DLR setting in two main respects. First, we drive the downscaling from a substantially coarser NWP model than the prior downscaling work, which operated close to the kilometre scale [3, 11], testing whether the terrain-aware correction still adds value when the input field is far smoother. Second, whereas the DLR-oriented work of [13] reports an end-to-end pipeline without a reproducible protocol, we contribute a systematic approach built on publicly available inputs and a stratified, spatially aware station split that measures generalisation to unseen locations. Concretely, we organise the model as a three-branch multimodal network (one NWP and two terrain branches) and select the input variables suited to our data. Methodologically, the task is one of multimodal fusion of heterogeneous spatio-temporal sources—coarse, time-varying NWP fields and static high-resolution terrain—into a single point prediction.

Our contributions are empirical. First, we show that terrain-aware convolutional downscaling still adds substantial value when driven by a much coarser ($\sim$7 km) NWP model than the $\sim$1 km-scale inputs of prior work, i.e. when the input wind field is far smoother. Second, we provide a reproducible, spatially honest evaluation protocol on fully public data (DWD, ICON-EU, Copernicus DEM), built around a stratified station split whose balance we verify statistically.

2 Background and Related Work

In this section, we review the background and prior work on which our approach builds. We first describe how dynamic line rating depends on the surrounding weather and why wind is the dominant and most uncertain input, and then survey statistical downscaling methods for near-surface wind, leading to the convolutional, terrain-aware approach we adopt.

The temperature of an overhead conductor, and hence its ampacity, is set by the balance between Joule heating and heat exchange with the environment through convective and radiative cooling, solar absorption, and, during rain, evaporative cooling [1, 10]. The ampacity is the largest current for which the steady-state conductor temperature stays at or below a prescribed limit, and is obtained by solving this heat balance as an inverse problem [1]. Depending on how much weather information enters the balance, one distinguishes static, ambient-adjusted, and dynamic ratings, the last using the full set of time-varying weather inputs and thus allowing the highest average utilisation of the line [1].

Among these inputs, the convective cooling, and therefore the ampacity, is dominated by wind speed and wind direction. These are at the same time the hardest weather variables to forecast reliably and are consequently the main source of ampacity uncertainty [16]. This motivates our focus on accurate, location-specific near-surface wind.

Classical statistical methods such as model output statistics [5], analog methods [20], and quantile mapping [6] assume a fixed functional form and treat each location on its own. Machine learning methods instead fit a flexible non-linear relation and can combine heterogeneous inputs, as in neural post-processing of NWP forecasts [17]; for near-surface wind in complex terrain, convolutional architectures that read the surrounding topography and adjust the forecast field to orographic features are particularly effective [3, 12, 11], and related pipelines target per-span DLR estimates [13]. We build on the convolutional downscaling approach of [3].

3 Data

In this section, we present the data used to train and evaluate the model. We describe the DWD wind observations that serve as ground truth, the ICON-EU forecasts that provide the weather inputs, and the terrain features derived from the Copernicus DEM, and we then detail the stratified, spatially aware station split together with its statistical verification.

3.1 Wind Observations

As ground truth, we use near-surface wind observations from the automatic synoptic network of the German Meteorological Service (Deutscher Wetterdienst, DWD) [2], a dense and freely available network that spans diverse terrain, from the North German Plain to the pre-Alpine foothills. Each station reports, at a 10-minute resolution, the mean wind speed (`FF_10`) and the prevailing wind direction (`DD_10`) at the standard measurement height of 10 m above ground, which coincides with the reference height of the near-surface wind in ICON-EU. We use the period 2021–2024, partitioned temporally into a training span (2021–2022), a validation year (2023), and a test year (2024). Because wind direction is a circular quantity, and because NWP models represent wind in the same way, we predict the two horizontal wind components (u, v) rather than speed and direction, and recover speed and direction from them.

After restricting the network to stations that provide observations from 2021 onward and excluding a single high-mountain station at 2956 m, whose flow regime is not representative of the conditions relevant to DLR, 266 stations remain. The partitioning procedure and its verification are described in Sect. 3.4.

3.2 Numerical Weather Prediction (ICON-EU)

As the source of weather forecasts, we use the regional ICON-EU model of DWD [19], the European refinement of the icosahedral non-hydrostatic ICON model, with a native horizontal grid spacing of $\sim$7 km and hourly output. The data is available as a public archive maintained by Open Climate Fix on HuggingFace [14], with four runs per day (00, 06, 12, 18 UTC).

From the available fields, we keep $C = 8$ NWP channels, the near-surface wind components u_{10} and v_{10}, and the horizontal wind components (u, v) at the 1000, 950, and 850 hPa isobaric levels. These channels were chosen through a preliminary input ablation over candidate NWP fields, in which the near-surface and low-level wind dominated the predictive skill, while higher-level, thermodynamic, and stability fields gave no consistent improvement.

3.3 Terrain Features

The NWP fields carry the temporal dynamics of the atmosphere but, at a grid spacing of several kilometres, cannot resolve how the local terrain shapes the near-surface wind through channelling, sheltering, and speed-up over ridges. High-resolution dynamical models reproduce these effects explicitly by refining the represented topography [8], which is precisely the spatial detail missing at the span level. We instead supply this information to the model through static terrain inputs.

For each station, we extract a patch centred on the station coordinate from the Copernicus GLO-30 digital elevation model (DEM), with a native resolution of $\sim$30 m and global coverage [4]. We use the DEM at two resolutions, a coarse one covering a large neighbourhood and a fine one covering the immediate surroundings, which feed the two terrain branches of the model (Sect. 4.2). For this paper, we deliberately restrict the terrain modality to the DEM. Richer modalities such as surface models, optical and radar imagery, and land cover are left for future work.

From the elevation field $z(x, y)$, we derive two standard terrain descriptors, the slope β and the aspect α. Denoting the horizontal derivatives by $z_x = \partial z/\partial x$ and $z_y = \partial z/\partial y$, the slope is the angle the elevation gradient makes with the horizontal plane,

$$\beta(x, y) = \arctan\sqrt{z_x^2 + z_y^2}, \tag{1}$$

and the aspect is the compass direction of steepest descent, measured clockwise from north,

$$\alpha(x, y) = \operatorname{atan2}(-z_x, -z_y) \bmod 2\pi. \tag{2}$$

Table 1. Resulting split (266 stations) and per-set wind/elevation statistics over a common reference year (2021, a training year), so the comparison reflects each set's station-level climatology and the spatial balance of the split independently of the held-out evaluation years (mean ± std across stations).

Set	N	mean speed [m/s]	p90 speed [m/s]	dir. entropy	elevation [m]
Train	171	3.43 ± 1.33	6.15 ± 2.16	0.937 ± 0.032	301 ± 290
Early-stopping	15	3.25 ± 1.79	5.86 ± 2.93	0.942 ± 0.023	277 ± 250
Val	40	3.50 ± 1.38	6.26 ± 2.22	0.939 ± 0.029	308 ± 277
Test	40	3.39 ± 0.97	6.07 ± 1.48	0.938 ± 0.039	301 ± 291

The derivatives are approximated by Horn's weighted 3×3 operator [7]. Intuitively, an exposed ridge and a sheltered valley within the same NWP cell have systematically different slope and aspect, and hence a different actual wind speed, which the terrain branches can exploit.

3.4 Station-Splitting Algorithm

We evaluate the model on both spatial and temporal generalisation, that is its ability to predict at stations and in periods it has never seen. Temporal generalisation is obtained by holding out an entire year for testing (Sect. 5.1), whereas spatial generalisation requires a careful partition of the stations, which is the focus of this section. A naive random split of stations is problematic for two reasons. First, nearby stations are spatially correlated, so placing them in different sets inflates the apparent test performance. Second, terrain types can be distributed unevenly across the sets, leaving them unbalanced.

We therefore use a stratified, spatially aware split. Stations are grouped into a coarse $2^\circ \times 2^\circ$ geographic grid and into three elevation classes (0–300, 300–700, $> 700\,\mathrm{m}$), and the cross product defines the strata. Elevation enters the stratification because near-surface wind in complex terrain depends strongly on it, with higher and more exposed sites being systematically windier [18]; such stations are also comparatively rare, so an unstratified split risks concentrating them in a single set. A first stratified pass holds out 30 % of the stations, which a second pass splits, block by block, into a validation and a test set of 40 stations each. The remaining 171 stations form the training set. To obtain a held-out set for early stopping, we apply the same stratified split once more to the training set alone, carving out an early-stopping set of 15 whole stations that is spatially disjoint from the rest of the training data, so that model selection tracks the same spatial generalisation as the final metric.

The resulting sets are statistically comparable. Their wind and elevation statistics (Table 1) are close, and the directional distributions overlap (Fig. 1). We confirmed this with Kruskal–Wallis and pairwise Kolmogorov–Smirnov tests at the $\alpha = 0.05$ level, which never reject equality of distributions across the sets.

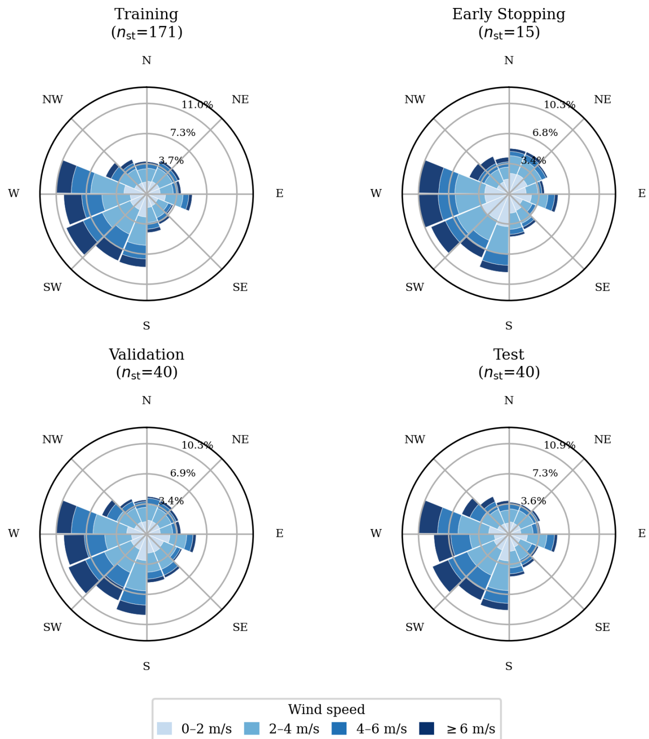

Fig. 1. Wind roses for the four data splits over 2021, showing similar directional distributions across sets.

4 Methods

In this section, we describe the proposed method. We first present the three-branch multimodal downscaling architecture, then give its detailed layer-level configuration, and finally specify the training setup, including the loss, augmentation, and optimisation.

4.1 Multimodal Downscaling Architecture

We follow the convolutional, point-wise downscaling design used in [3]. For a given station the model predicts the horizontal wind components (u, v) at the centre of an input patch, from the surrounding NWP fields and the surrounding high-resolution terrain. The prediction is point-wise and time-independent, which makes the model easy to apply at arbitrary new locations.

In our adaptation, the model has three input branches: one for the NWP fields and two terrain branches for the DEM-derived inputs at a coarse and a fine resolution (Sect. 4.2). Pointwise metadata enter as additional scalars: the

Table 2. Input branches of the proposed model. Both terrain branches use the Copernicus GLO-30 DEM; the NWP patch is bilinearly upsampled from the ICON-EU grid.

Branch	Inputs	Resolution	Patch (extent)
NWP	u_{10}, v_{10}, (u, v) at 1000/950/850 hPa	273 m/px	77×77 (21 km)
Coarse terrain	DEM30, slope, aspect	273 m/px	77×77 (21 km)
Fine terrain	DEM30, slope, aspect	53 m/px	77×77 (4.1 km)
Point / meta	hour, day-of-year, norm. elevation	—	—

cyclic time of day and day of year, each as a $(\sin, \cos)$ pair, and the normalised elevation, five values in total. Each modality is encoded by its own convolutional stem; the branch descriptors are fused by concatenation and mapped by fully connected layers to the output, with u and v produced by two jointly trained heads as in [3]. The layer-level configuration is given in Sect. 4.2.

4.2 Model Configuration

The NWP branch takes a 4×4 patch of ICON-EU cells centred on the station and, following [3], bilinearly interpolates it onto the same 77×77 grid as the terrain branches, spanning $\sim$21 km with the station at the patch centre. Placing all three branches on a common 77×77 grid lets each use an identically structured convolutional stem. Both terrain branches are derived from the same Copernicus GLO-30 DEM at the two resolutions in Table 2, each augmented with the slope and aspect descriptors of Eqs. (1)–(2). Both branches are re-griddings of the same native $\sim$30 m DEM onto a common 77×77 stem; the coarse and fine labels refer to the patch extent and grid spacing, not to a different data source.

All three branch inputs therefore enter the network at the same 77×77 resolution. Each branch then has its own identically structured convolutional stem (separate weights) of four 3×3 convolution blocks with channel widths 64, 128, 192, and 256; every block applies batch normalisation and a SiLU activation, and the first three blocks halve the spatial resolution through 2×2 max-pooling ($77 \rightarrow 38 \rightarrow 19 \rightarrow 9$). Each stem reduces its final 9×9 feature map to a 512-dimensional descriptor by concatenating a central 3×3 average with a global average pool.

The three branch descriptors are concatenated with the five metadata scalars into a single 1541-dimensional fusion vector, which two fully connected layers of width 384 and 192 (SiLU, dropout 0.1) map to the wind components through two linear heads, one for u and one for v. At inference, we apply test-time augmentation: the prediction is averaged with the one obtained from the east–west-mirrored inputs, which reduces variance. The complete model has about 2.9 M parameters and is trained on a fixed 10 % subset of forecast lead times, which amounts to about 3.4 M forecast–observation examples (Sect. 4.3). In the preliminary experiments, the validation error already saturated well below this data size, so the subset keeps training tractable at no meaningful loss in accuracy.

4.3 Training Setup

The model is trained with the Adam optimiser [9] under a cosine-annealed learning rate. Instead of a plain mean-squared or mean-absolute error, which tends to make the predicted wind-component distributions too narrow, we minimise the distribution-aware regression loss of [3]; it rescales the predicted components towards the observed wind speed and applies an asymmetric, pinball-inspired weight to remove the systematic speed bias, both of which preserve the wind-speed distribution that DLR ultimately depends on. Observations flagged as missing are excluded from the loss through a per-sample validity mask, so that no synthetic targets are introduced.

We use rotation augmentation. At each step the terrain patches and the corresponding wind vector are rotated by a common random angle. Because the terrain and the wind vector are rotated jointly, this exploits the approximate rotational equivariance of the terrain–wind relationship; it enlarges the effective set of orographic configurations seen during training, and acts as a strong regulariser, which is particularly important for generalisation to unseen stations. Because the model converges quickly and tends to overfit, we use early stopping [15] and halt training at the epoch that minimises the loss on the held-out early-stopping set (Sect. 3.4). Evaluating the stopping criterion on stations disjoint from the training set, rather than on a random sample of training stations, makes it track the same spatial generalisation as the final evaluation. We use a batch size of 256, a peak learning rate of 2.7×10^{-4}, weight decay 5×10^{-7}, and at most 20 epochs with a patience of 6; these values were chosen by a small, preliminary hyperparameter search.

5 Experiments and Results

In this section, we present and analyse the experimental results. We first define the evaluation setup, metrics, and baselines, then report the deterministic downscaling performance against those baselines, quantify the contribution of each terrain branch through an ablation, and close with a discussion of the findings and the main limitations.

5.1 Experimental Setup and Metrics

We evaluate on the test set (40 unseen stations) over the held-out year 2024, so that the reported numbers reflect spatio-temporal generalisation. For wind speed, we report the mean absolute error (MAE), the root-mean-square error (RMSE), and the mean bias error (MBE; defined as predicted minus observed, so a negative value means the wind speed is underestimated). For wind direction, we report the mean absolute error, computed circularly and restricted to events with wind speed above 1 m/s.

We compare against two baselines that are reproducible on our own data. The first is the interpolated NWP, obtained by bilinearly interpolating the ICON-EU near-surface wind to the station location; it is the unambiguous reference

Table 3. Performance on the held-out test set (year 2024). Proposed model vs. the interpolated-NWP and ANN post-processing baselines. The proposed model is the mean over three seeds (wind-speed MAE 1.058 ± 0.008). Direction MAE is restricted to events with wind speed above 1 m/s.

Model	MAE [m/s]	RMSE [m/s]	Bias [m/s]	MAE dir [°]
NWP (interpolated)	1.271	1.890	−0.71	**25.6**
ANN post-processing	1.157	1.510	+0.36	25.9
Proposed model	**1.058**	**1.437**	**0.00**	25.9

for the value added by downscaling. The second is a classical machine-learning post-processing baseline (ANN): two small multilayer perceptrons that correct the interpolated NWP wind speed and direction from pointwise features (elevation, local terrain descriptors, and the NWP wind at the lowest levels). Because this baseline sees the same information but without the convolutional terrain branches, it isolates the contribution of the terrain-reading architecture from that of non-linear post-processing *per se*. We deliberately do not reimplement other convolutional downscaling methods, because we do not have the exact data used in those works and a like-for-like comparison would not be possible.

5.2 Deterministic Downscaling

The proposed model improves on the interpolated NWP (Table 3) most clearly in wind-speed MAE ($1.271 \rightarrow 1.058$, $-17\,\%$) and RMSE ($-24\,\%$), and it essentially removes the speed bias ($-0.71 \rightarrow 0.00$ m/s). It also clearly beats the ANN post-processing baseline ($1.157 \rightarrow 1.058$), which confirms that the gain comes from the convolutional terrain-reading architecture and not from non-linear post-processing alone. The gain holds over the whole forecast horizon: Fig. 2 shows that the wind-speed MAE of both the baseline and our model grows with lead time, as expected, but the model retains a roughly constant margin (about 0.2 m/s) out to 72 h. Wind direction, by contrast, is essentially unchanged (NWP 25.6°, model 25.9°), which is expected because direction is set mainly by the driving NWP and is little affected by the local terrain correction. The result is stable across random seeds: the test wind-speed MAE over three seeds is 1.058 ± 0.008 m/s, and the per-station MAE has a spread of 0.29 m/s, reflecting genuine site-to-site difficulty rather than seed noise.

The benefit is strongly elevation-dependent (Table 4): at low-lying stations the interpolated NWP is already fairly accurate and the correction is modest ($1.122 \rightarrow 1.067$), whereas at high, exposed sites—precisely the most DLR-relevant spans, and where the coarse NWP is worst—the model cuts the wind-speed MAE by a third ($1.803 \rightarrow 1.163$).

5.3 Terrain Ablation

To quantify where the gain comes from, we ablate the two terrain branches on the validation set (Table 5). With both terrain branches removed, the network sees

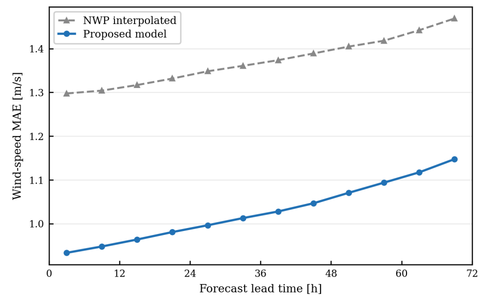

Fig. 2. Wind-speed MAE as a function of forecast lead time on the held-out test set, for the proposed model (mean ± std over three seeds, shaded) and the interpolated-NWP baseline.

Table 4. Wind-speed MAE on the held-out test set (year 2024) by station-elevation band: interpolated NWP vs. the proposed model (mean over three seeds).

Elevation band	samples	NWP MAE [m/s]	Proposed MAE [m/s]
0–200 m	1.25 M	1.122	**1.067**
200–500 m	0.88 M	1.104	**0.971**
500–2000 m	0.62 M	1.803	**1.163**

only the NWP fields and barely improves on the interpolated baseline (val-2023 MAE 1.284 vs. 1.371): a neural network on the NWP input alone is *not* enough, so the gain is not a generic machine-learning effect. Adding the coarse terrain branch alone helps only marginally (1.284 → 1.273), whereas the fine, high-resolution terrain branch accounts for most of the improvement (1.284 → 1.129). The two branches are complementary: combining them is best (1.058), the coarse context contributing once the fine detail is present. This confirms that the value of the method lies specifically in the terrain-aware correction, and predominantly in its high-resolution component. The full configuration in the last row is exactly the model of Table 2 that we evaluate on the test set in Sect. 5.2; its validation MAE of 1.058 coincides with the test MAE of 1.058 reported there, evidence that the spatially honest split generalises consistently across both unseen stations and the unseen year.

5.4 Discussion and Limitations

Our results indicate that convolutional terrain-aware downscaling transfers to a new geography and application, where rotation augmentation and a spatially honest evaluation support generalisation to unseen stations.

Several limitations remain. The model is point-wise and time-independent, ignoring temporal context and imposing no physical constraints, so its skill is bounded by the driving NWP; where the interpolated NWP is already accurate the achievable gain is small. Finally, the present study deliberately stops

Table 5. Terrain ablation on the validation set (year 2023), wind-speed error as mean ± std over three seeds.

Configuration	MAE [m/s]	RMSE [m/s]
NWP only (no terrain)	1.284 ± 0.078	1.766 ± 0.153
NWP + coarse terrain	1.273 ± 0.024	1.772 ± 0.023
NWP + fine terrain	1.129 ± 0.033	1.546 ± 0.043
NWP + coarse + fine (full)	$\mathbf{1.058 \pm 0.012}$	$\mathbf{1.426 \pm 0.020}$

at the deterministic wind estimate. We treat fixing the optimal point-forecast configuration (input variables, terrain resolution, and training regime) as a prerequisite, and the probabilistic forecast as the natural next step: a predictive wind *distribution* is what propagates, through the conductor heat balance [1], to the ampacity distribution that DLR ultimately uses. That probabilistic head and the ampacity computation itself are out of scope here and left for future work; accordingly, the present results should be read as the wind-downscaling building block for DLR rather than an end-to-end ampacity study.

6 Conclusion and Future Work

We adapted a convolutional, terrain-aware statistical downscaling approach to dynamic line rating on German DWD observations and ICON-EU forecasts, and evaluated it under a stratified, spatially aware station split that measures generalisation to unseen locations. The three-branch model improves substantially over the interpolated NWP across the whole forecast horizon, and the ablation shows that the high-resolution terrain branch contributes most of this gain, while the NWP side needs only a compact low-level wind input. All inputs are public (DWD, ICON-EU, Copernicus DEM), and we will release the code and the exact station split to support reproducibility.

Future work includes extending the deterministic forecast with a probabilistic head whose wind distribution propagates through the conductor heat balance to an ampacity distribution, adding temporal context (multiple forecast steps), predicting at the intra-hour (10-minute) scale from hourly inputs, and incorporating richer terrain and land-surface modalities (high-resolution surface models and optical/radar imagery).

Acknowledgments. We acknowledge the support of the Slovenian Research Agency through the grants L2-70125, GC-0001 and GC-0006, and through the Horizon Europe grant 101254461 SLAIF (Slovenian AI Factory), by EuroHPC JU and the Slovenian Ministry of Education, Science and Youth.

References

1. CIGRE Working Group B2.43: Guide for thermal rating calculations of overhead lines. Technical Brochure 601, CIGRE (2014)

2. Deutscher Wetterdienst: Climate data center (cdc): 10-minute station observations. https://opendata.dwd.de, accessed 2026
3. Dujardin, J., Lehning, M.: Wind-Topo: Downscaling near-surface wind fields to high-resolution topography in highly complex terrain with deep learning. Quarterly Journal of the Royal Meteorological Society **148**(744), 1368–1388 (2022)
4. European Space Agency: Copernicus GLO-30 digital elevation model. https://doi.org/10.5270/ESA-c5d3d65 (2022), ∼30 m global DEM
5. Glahn, H.R., Lowry, D.A.: The use of model output statistics (MOS) in objective weather forecasting. Journal of Applied Meteorology **11**(8), 1203–1211 (1972)
6. Gudmundsson, L., Bremnes, J.B., Haugen, J.E., Engen-Skaugen, T.: Technical note: Downscaling RCM precipitation to the station scale using statistical transformations – a comparison of methods. Hydrology and Earth System Sciences **16**(9), 3383–3390 (2012)
7. Horn, B.K.P.: Hill shading and the reflectance map. Proceedings of the IEEE **69**(1), 14–47 (1981)
8. Jiménez, P.A., Dudhia, J.: Improving the representation of resolved and unresolved topographic effects on surface wind in the WRF model. Journal of Applied Meteorology and Climatology **51**(2), 300–316 (2012)
9. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations (ICLR) (2015)
10. Kosec, G., Maksić, M., Djurica, V.: Dynamic thermal rating of power lines – model and measurements in rainy conditions. International Journal of Electrical Power & Energy Systems **91**, 222–229 (2017)
11. Le Toumelin, L., Gouttevin, I., Galiez, C., Helbig, N.: A two-fold deep-learning strategy to correct and downscale winds over mountains. Nonlinear Processes in Geophysics **31**(1), 75–97 (2024)
12. Le Toumelin, L., Gouttevin, I., Helbig, N., Galiez, C., Roux, M., Karbou, F.: Emulating the adaptation of wind fields to complex terrain with deep learning. Artificial Intelligence for the Earth Systems **2**(1), e220034 (2023)
13. Manninen, H., Lippus, M., Rute, G.: Dynamic line rating using hyper-local weather predictions: A machine learning approach. arXiv:2405.12319 (2024), preprint
14. Open Climate Fix: DWD ICON-EU forecast archive. https://huggingface.co/datasets/openclimatefix/dwd-icon-eu, accessed 2026
15. Prechelt, L.: Early stopping — but when? In: Neural Networks: Tricks of the Trade, Lecture Notes in Computer Science, vol. 7700, pp. 53–67. Springer, 2nd edn. (2012)
16. Rashkovska, A., Jančič, M., Depolli, M., Kosmač, J., Kosec, G.: Uncertainty assessment of dynamic thermal line rating for operational use at transmission system operators. IEEE Transactions on Power Systems **37**(6), 4642–4650 (2022)
17. Rasp, S., Lerch, S.: Neural networks for postprocessing ensemble weather forecasts. Monthly Weather Review **146**(11), 3885–3900 (2018)
18. Troen, I., Petersen, E.L.: European Wind Atlas. Risø National Laboratory, Roskilde, Denmark (1989)
19. Zängl, G., Reinert, D., Rípodas, P., Baldauf, M.: The ICON (icosahedral nonhydrostatic) modelling framework. Quarterly Journal of the Royal Meteorological Society **141**(687), 563–579 (2015)
20. Zorita, E., von Storch, H.: The analog method as a simple statistical downscaling technique: Comparison with more complicated methods. Journal of Climate **12**(8), 2474–2489 (1999)