Physics-Informed Machine Learning for Wind Forecasting using WRF

1st Amadej Krepek
Research and Development
Operato d.o.o.
Maribor, Slovenia
amadej.krepek@operato.eu

2nd Iztok Fister, Jr.
Faculty of Electrical Engineering and Computer Science
University of Maribor
Maribor, Slovenia
iztok.fister1@um.si

3rd Andrej Vilhar
Research and Development
Operato d.o.o.
Ljubljana, Slovenia
andrej.vilhar@operato.eu

***Abstract*—Physics-informed machine learning (PIML) integrates physical knowledge into data-driven models to prevent them from producing unrealistic results. Various techniques exist to perform such integration of physics laws. We applied these methods to the domain of wind forecasting, incorporating physical constraints, tuning wind components, and deriving important features. The results improved when these elements were tuned together carefully. Our baseline model was enriched with physical information, with the goal of achieving better mean absolute wind direction metrics compared to the baseline. At a specific location and for a 6-hour forecast horizon, the average wind direction error was reduced by up to 65° in particular cases of high wind speeds, and by 11.9° on average compared to the weather model inputs alone. Compared to the baseline data-driven solution, the average improvement was 9.7°.**

***Index Terms*—physics-informed machine learning, wind forecasting, WRF**

I. Introduction

Accurate wind forecasting is essential for numerous applications, including the calculation of Dynamic Line Rating (DLR) in power transmission systems. In standard DLR heat balance models (e.g. [1]), convective cooling due to wind represents the dominant mechanism of line heat dissipation (see e.g. [2], [3]). Consequently, reliable wind forecasts that detail wind speed and direction are essential for estimating power-line capacity accurately. This paper focuses on improving the accuracy of wind direction forecasts, with the specific objective to:

- improve the **accuracy** of the LSTM (Long Short Term Memory) model [4],
- focus on **short-term** wind forecast [5],
- evaluate using **MAE** (Mean Absolute Error) metric.

PIML improves the LSTM accuracy by combining physical modeling and data-driven methods. The main benefits of the proposed method are:

- faster training,
- better accuracy,
- improved generalization.

This is achieved by injecting physical knowledge into the machine learning pipeline [7].

From a spatial perspective, we focus on local near-surface wind at a single location, but driven by a mesoscale numerical weather prediction (NWP) model (WRF - Weather Research and Forecasting). This means that our predictions operate at the grid scale of WRF (order of kilometers), where topography, valleys and larger terrain features are resolved, but individual buildings, obstacles and street-canyon effects are not explicitly modeled. In other words, our approach does not aim to compete with building-resolving CFD models; instead, it targets the intermediate scale between global models and microscale urban flow, where direction errors from NWP can still be substantial for applications such as dynamic line rating.

In the paper [8], WRF (Weather Research and Forecasting Model), was studied, to predict wind speed on high-resolution terrain. However, it is computationally expensive and requires enormous amounts of data for producing forecasts. We deployed a framework on our machines to generate forecasts used as input to the proposed model.

Based on the highlighted physics-informed solutions, we propose a novel loss function based on the following:

1) the main idea and structure of physics terms was inspired by one of the key approaches to incorporate physics to the machine learning pipeline with custom designed loss functions [9],
2) inspired by creating custom physics-based loss function 3 used for correction and evaluation [10],
3) with our geostrophic wind alignment penalty, defined by eq. 6 as soft constraint predicting wind vectors (u, v) to stabilize flow [11], and
4) by weighted penalties (so-called "soft constraints").

In section two, we cover related work from data-driven solutions to WRF-based and PIML, with description of each category. In section 3, we provide a brief background, with arguments for physics-informed solutions. Section 4 describes the materials and methods, including the variables summarized in the Tables and the visualization of the model pipeline. Section 5 presents and analyzes the experimental results, while Section 6 provides a discussion.

II. Related Work

In this Section we present related work on weather forecasting, with the main devotion to PIML-based methods. The following studies have been identified during the literature review:

1) **Data-Driven**:

979-8-3315-9163-2/26/$31.00 ©2026 IEEE

- FourCastNet [12],
- Aardwark weather [13],
- Pangu-Weather [14],
- LSTM pressure levels [15].

2) **WRF conditioned, PIML:**

- WRF in the Amazon Forest [16],
- AIRU-WRF [17],
- WRF-snowfall prediction [18],
- WRF-tropical cyclone wind fields [19].

A. Data-Driven Weather Forecasting

Data-driven machine learning methods have become increasingly popular. The LSTM predicts the pressure levels used to predict 500 hPa height with approaching the NWP performance accuracy [15]. WeatherBench compared the deep learning method with global weather prediction systems based on data-driven methods [20]. These models have achieved huge gain in performance, and stood by NWP models on forecasting a few hours or days ahead. The biggest requirement is the massive training data - 40 years of historical reanalysis data in FourCastNet's case [12]. For example, in [21], a WRF-BPNN model was developed with a neural network behind to post process WRF forecasts. It demonstrated higher accuracy in temperature and humidity compared to the reference WRF model, while fixing systematic errors in WRF's output. Training on historical winter storms in Northeast U.S., their hybrid approach improved snowfall forecast compared to WRF output [18]. Similar work using WRF-Hydro enhancing runoff prediction with physics-based fusion has gained better results [22]. In summary, the highlighted models are: FourCastNet, producing week-long forecasts, Aardwark weather, replacing entire operational NWP pipeline, Pangu-Weather, with strong deterministic forecast system, deep learning techniques predicting pressure levels experimenting with LSTM.

B. Physics-Guided Wind Forecasting

Predicting hurricane winds with a physics-informed neural network demonstrated that an ML model equipped with physical equations can provide realistic results and stability. This further provides evidence that PIML can generalize from observational data by leveraging the governing equations incorporated into the model [19].

Diving deeper into weather forecasting with physical laws, PINNs (Physics-informed Neural Networks) have become a popular network to solve or constrain applications with PDEs (Partial Differential Equations) using neural networks [6]. Since then, PINNs and hybrid methods have been applied to many problems with fluid dynamics, as well as climate itself [7]. In summary, the approach is to start with purely data-driven weather models, and extend them toward physics-informed methods. Fully data-driven models (e.g. [12], [13], [14], [15], [20]) produce fast and reliable forecasts but require, big data and risk physical inconsistency. At the same time, injecting physics into ML whether through constraints, loss functions or coupling ML with a physics model - has proven effective in multiple contexts ([16], [17], [19], [9]). These additions gather important information together for our approach, which will merge WRF predictions with deep learning and enforce physical information in the forecast. Pure ML often leads to unrealistic results such as negative rain or wind and is optimized for generic error metrics. For that case including soft/hard constraints [23] or penalties for violating known laws avoids out of scope forecasts [16]. AIRU-WRF was proposed, a spatio-temporal model for offshore wind forecasting joining WRF with ML. In practical terms, their physics guided short-term improves in wind forecast accuracy, illustrating the benefit of embedding domain knowledge [17].

III. Background

Traditional ML models can learn complex mappings from big data; however, (1) they often lack sufficiently long periods of quality datasets, and (2) the developed solutions may violate physical laws, leading to poor predictions [6] [17]. PIML addresses these challenges by injecting physical knowledge into the learning phase. By adding equations or a theoretical background, a PIML model can be trained on smaller datasets while still improving the forecast accuracy, and enabling more reliable extrapolation [7]. Another motivation is the computational costs of physics-based models. Complex simulations like WRF are resource-intensive, whereas ML models can produce forecasts more efficient. However, with only data-driven models, we risk forecasting out known physical constraints to the world [17].

In terms of scale, many recent data-driven models (e.g. FourCastNet, Aardvark, Pangu-Weather) already provide skilful large-scale wind direction forecasts on global grids. However, these models do not explicitly resolve local flow-channeling over complex terrain at the kilometer scale, nor the very fine details of flow around buildings and obstacles. Our work targets this intermediate mesoscale-to-local regime: we use WRF to represent terrain-induced flow patterns, and then apply a physics-informed correction at a specific station. The physics we inject through the loss function is therefore primarily related to large-scale balance (via geostrophic wind) and directional consistency, not to building-scale turbulence modeling.

IV. Materials & Methods

The simulation is illustrated in Figure 1. The ICON-EU model [24] supplies the lateral boundary and initial conditions for WRF at a 1-hour interval. After completing the simulation, the resulting WRF outputs were converted into analysis-ready variables using preprocessing tools from `wrf-python` [25].

- **2D (horizontal) variables** – produce a single horizontal field (map) for each time step.
- **3D (vertical) variables** – consist of a stack of horizontal fields across multiple height levels for each time step.

Both groups were transformed into feature vectors at the target station location (46.600282° N, 15.876867° E), using data from August 2025 covering the area of Slovenia. The list of baseline variables is presented in Table I.

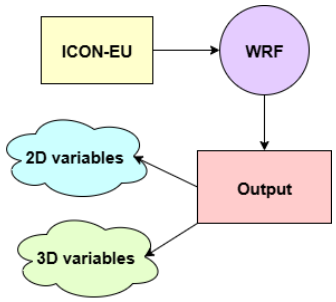

Fig. 1. WRF pipeline from input weather model to output variables

TABLE I
BASE VARIABLES

Name	Units	Description
U10	m/s	10-m eastward wind component
V10	m/s	10-m northward wind component
T2	°C	2-m temperature
Q2	$\mathrm{kg\,kg^{-1}}$	2-m specific humidity
PSFC	Pa	Surface pressure
HGT	m	Terrain height

Additionally, we provide a list of the supplementary variables (Table II), which contribute further to the training of the LSTM model in physics-informed mode.

TABLE II
ADDITIONAL VARIABLES

Name	Units	Description
CWV	$\mathrm{kg\,m^{-2}}$	Column water vapor
RH2	%	2-m relative humidity
MSE_sfc	$\mathrm{J\,kg^{-1}}$	Surface moist static energy
RNET_sfc	$\mathrm{W\,m^{-2}}$	Net surface radiation
SHEAR_0_1KM	$\mathrm{m\,s^{-1}}$	Bulk wind shear (0-1 km)
SHEAR_0_6KM	$\mathrm{m\,s^{-1}}$	Bulk wind shear (0-6 km)
UGEO	$\mathrm{m\,s^{-1}}$	U-component geostrophic wind
VGEO	$\mathrm{m\,s^{-1}}$	V-component geostrophic wind

All the features were generated using the WRF forecasting system and post-processing tools.

A visual summary is shown in Figure 2. Short-Term wind forecasting was built up with supervised sequence learning with a sliding-window approach. This is a time-series problem. Given the last 36 hours of samples, the model is able to predict the 10 m wind vector (u,v) exactly six hours in advance. The blocked k-fold strategy with temporal gap was performed to prevent data leakage. For stability and better metrics, each fold was repeated twice, and the metrics presented were averaged across the folds. The Cross-validation configuration is presented in Table III. After that, the features were standardized using a statistics fit on the training part of each fold, and applied to its validation/test slices.

Two formulas are needed to convert wind vectors (u,v) to wind speed (eq. 1) and wind direction (eq. 2).

$$s = \sqrt{u^2 + v^2}\,. \tag{1}$$

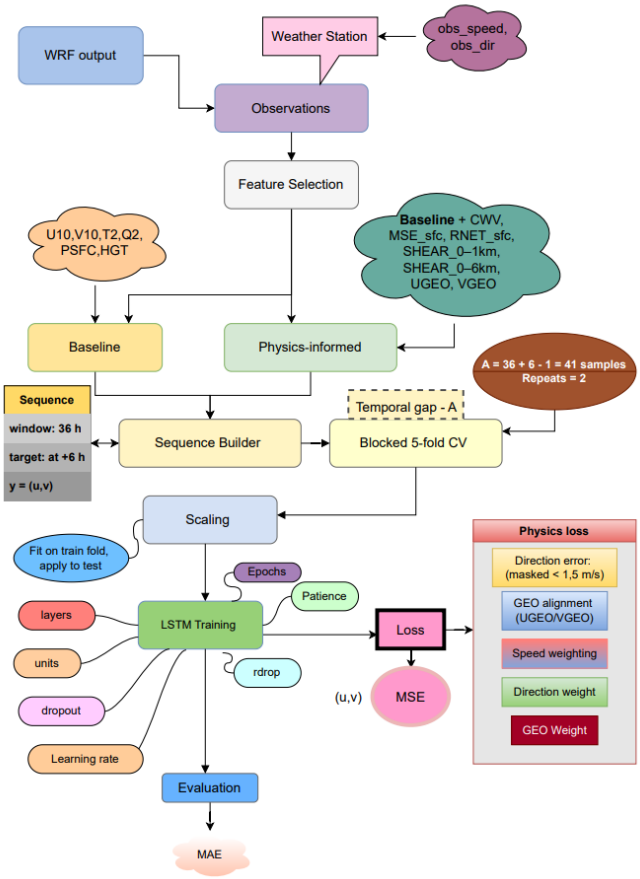

Fig. 2. End-to-end workflow of the proposed physics-informed forecasting system.

TABLE III
CROSS-VALIDATION CONFIGURATION FOR SHORT-TERM WIND FORECASTING

Setting	Description
Cross-validation type	Blocked 5-fold temporal CV
Input window	36 h
Forecast horizon	6 h
Overlap offset	1 h
Gap between folds	$36 + 6 - 1 = 41$ h
Repetitions	Each fold repeated twice

$$\theta = \left(\operatorname{atan2}(u,\, v)\,\frac{180}{\pi} + 180\right) \bmod 360\,. \tag{2}$$

LSTM, with its ability to solve complex problems, was implemented as presented in Table IV. The model predicts the absolute wind vector components (u, v).

The network comprises two stacked, unidirectional LSTM layers with 128 hidden units each, followed by a dense output layer predicting the absolute wind vector components (u, v). A dropout rate of 0.15 is applied after each LSTM layer, and a recurrent dropout of 0.05 is used inside the LSTM units. Training is performed with the Adam optimizer (learning rate

5×10^{-4}), a batch size of 128, and a maximum of 120 epochs. The loss function is a physics-aware composite in `weighted` mode: it combines the mean squared error in (u, v) space with a direction-aware cosine term whose contribution grows with wind speed.

TABLE IV
CONFIGURATION OF THE LSTM PREDICTOR USED IN THIS WORK.

Layer	Output Shape	Details
Input	$(36, F)$	36 h sequence, F features
LSTM 1	$(36, 128)$	`return_sequences=True`
Dropout	$(36, 128)$	rate = 0.15
LSTM 2	(128)	128 units
Recurrent Dropout	(128)	rate = 0.05
Dense (output)	(2)	Predicts (u, v)

A. Physics loss

To incorporate physical knowledge into the learning procedure, we follow one of the suggested strategies in physics-informed machine learning: the use of custom-designed loss functions with "soft" physical constraints [9].

Our loss function approach is described below. It was inspired by [6] [23] [26] [27] (eq. 3):

$$L = L_{\mathrm{uv}} + \lambda_{\mathrm{dir}}\, w_s^{(\alpha)}\, L_{\mathrm{angle}} + \lambda_{\mathrm{geo}}\, L_{\mathrm{geo}}, \tag{3}$$

a) Wind-vector MSE.: The baseline term penalizes errors in the horizontal wind components:

$$L_{\mathrm{uv}} = \mathrm{MSE}(u_{\mathrm{obs}}, u_{\mathrm{pred}}) + \mathrm{MSE}(v_{\mathrm{obs}}, v_{\mathrm{pred}}). \tag{4}$$

and

b) Directional (angular) error.: To improve the prediction of wind direction, we use a cosine-based penalty,

$$L_{\mathrm{angle}} = 1 - \cos(\Delta\theta), \tag{5}$$

where $\Delta\theta$ is the difference between the observed and predicted wind direction [28]. This term is modulated by the weighting factor

$$w_s^{(\alpha)} = \|\mathbf{v}_{\mathrm{obs}}\|^{\alpha},$$

which increases the influence of high-wind events (i.e., cases where directional accuracy is most important).

c) Geostrophic alignment.: To encourage physical coherence with large-scale forecasting, we add a weak regularization term that aligns the predicted near-surface wind $\hat{\mathbf{v}}$ with the diagnostic geostrophic wind $\mathbf{v}_g$:

$$L_{\mathrm{geo}} = 1 - \cos\big(\hat{\mathbf{v}}, \mathbf{v}_g\big). \tag{6}$$

This provides a "soft" dynamical constraint without overriding local terrain-driven wind behavior.

TABLE V
DISTRIBUTION OF SAMPLES BY OBSERVED SPEED BINS FROM PREDICTIONS

	Global	0–1	1–2	2–3	3–4	4–5	5–6
Samples	678	400	194	50	20	9	4

Note: Only the first six bins are reported, as higher bins contain too few samples. Units are in m/s.

TABLE VI
MAE WIND DIRECTION ERROR

Model	Global	0–1	1–2	2–3	3–4	4–5	5–6
WRF	75.8	76.3	68.7	79.2	99.6	112.2	118.5
Baseline	72.9	76.2	77.1	60.7	58.3	43.3	36.4
PIML	**61.7**	**67.3**	**61.7**	**49.5**	**44.3**	**40.2**	**25.4**

d) Hyperparameter selection.: The weights λ_{dir} and λ_{geo} control the contribution of directional and geostrophic terms. Previous studies [11], [29] show that geostrophic information can assist short-term wind forecasts, but excessive weighting may degrade accuracy. Extensive testing in our setup showed that

$$\lambda_{\mathrm{dir}} = 0.25, \qquad \lambda_{\mathrm{geo}} = 0.15$$

provided the best trade-off between physical realism and predictive skill. All experiments used neutral bin weights, meaning that no wind-speed category was artificially emphasized.

V. EXPERIMENTAL RESULTS

Predicting wind speed and especially wind direction requires careful tuning of the physics-based loss function. We were able to achieve much better results using physics-informed philosophy.

Both the global MAE and the MAE computed over binned data are introduced for evaluation. Each bin covers different wind speed intervals. Table V presents the number of samples for each bin and their total number in hours. With increasing wind speed, the number of samples per bin decreases. Bins over 6 m/s are not included as there were only 4 samples altogether beyond this category. The most representative results are presented in Table VI, where we compare the WRF forecast, the baseline LSTM, and the PIML model.

On average, i.e., across the global set of samples, the best performance was achieved by PIML, with an error of 61.7°. Next was the baseline LSTM model, which yielded 72.9°, while the original WRF model showed the worst performance at 75.8°. The wind direction was divided into bins to evaluate performance over specific intervals. As shown in Table IV, PIML achieved the best results across all bins.

The results for wind speed are not presented here, as the improvements were not significant at this stage of the study; notable gains were observed only for some low-speed bins relative to the basic WRF model. Improvements in wind speed remain an open area for future work.

VI. Conclusion

This study shows that physics-informed machine learning can improve local short-term wind direction forecasts when used as a post-processing layer on top of WRF. The combination of WRF-derived features and the proposed physics-based loss achieved consistent reductions in wind-direction error, especially in higher wind-speed bins where accurate direction matters most for applications such as Dynamic Line Rating (DLR).

Methodologically, the work contributes a practical recipe for improving site-specific wind forecasts: extract physically meaningful diagnostics from standard WRF outputs, train a lightweight sequence model on local measurements, and guide learning using a hybrid loss that combines vector MSE, directional consistency and geostrophic alignment. This approach requires only a single on-site anemometer and integrates seamlessly with existing WRF workflows.

In practical terms, the method offers a low-cost way to correct systematic directional biases in complex terrain without modifying the NWP model itself. Future work includes extending the evaluation across seasons and locations, enhancing wind-speed prediction, and assessing the operational impact on DLR and other engineering decisions dependent on short-term wind forecasts.

References

[1] CIGRÉ B2-42 Brochure, Ref. No. 601. (2014). Guide for thermal rating calculations of overhead lines.

[2] Wallnerström, C. J., Huang, Y., and Söder, L. (2015). Impact from dynamic line rating on wind power integration. IEEE Transactions on Smart Grid, 6(1), 343–350. https://doi.org/10.1109/TSG.2014.2341353

[3] Dawson, L., and Knight, A. (2023). Investigating the impact of a dynamic thermal rating on wind farm integration. IET Generation, Transmission and Distribution, 17(10), e12821. https://doi.org/10.1049/gtd2.12821

[4] Van Houdt, G., Mosquera, C., and Nápoles, G. (2020). A review on the long short-term memory model. Artificial Intelligence Review, 53(8), 5929–5955. https://doi.org/10.1007/s10462-020-09838-1

[5] Zhang, Z., Lin, L., Gao, S., Wang, J., Zhao, H., and Yu, H. (2025). A machine learning model for hub-height short-term wind speed prediction. Nature Communications, 16(1), 3195. https://doi.org/10.1038/s41467-025-58456-4

[6] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., and Yang, L. (2021). Physics-informed machine learning. Nature Reviews Physics, 3(6), 422–440. https://doi.org/10.1038/s42254-021-00314-5

[7] Meng, C., Griesemer, S., Cao, D., Seo, S., and Liu, Y. (2025). When physics meets machine learning: a survey of physics-informed machine learning. Machine Learning for Computational Science and Engineering, 1(1), 20. https://doi.org/10.1007/s44379-025-00016-0

[8] Powers, J. G., Klemp, J. B., Skamarock, W. C., Davis, C. A., Dudhia, J., Gill, D. O., Coen, J. L., Gochis, D. J., Ahmadov, R., Peckham, S. E., Grell, G. A., Michalakes, J., Trahan, S., Benjamin, S. G., Alexander, C. R., Dimego, G. J., Wang, W., Schwartz, C. S., Romine, G. S., … Duda, M. G. (2017). The Weather Research and Forecasting Model: Overview, System Efforts, and Future Directions. Bulletin of the American Meteorological Society, 98(8), 1717–1737. https://doi.org/10.1175/BAMS-D-15-00308.1

[9] Kashinath, K., Mustafa, M., Albert, A., Wu, J.-L., Jiang, C., Esmaeilzadeh, S., Azizzadenesheli, K., Wang, R., Chattopadhyay, A., Singh, A., Manepalli, A., Chirila, D., Yu, R., Walters, R., White, B., Xiao, H., Tchelepi, H. A., Marcus, P., Anandkumar, A., … Prabhat. (2021). Physics-informed machine learning: case studies for weather and climate modelling. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 379(2194), 20200093. https://doi.org/10.1098/rsta.2020.0093

[10] Karpatne A, Watkins W, Read J, Kumar V. 2018 Physics-guided neural networks (PGNN): anapplication in lake temperature modeling. In Proc. of the 2018 SIAM Int. Conf. on Data Mining. https://doi.org/10.48550/arXiv.1710.11431

[11] Huang, L., Shu, Y., Yao, J., and Liu, D. (2025). A Physics-Informed Neural Network for Sea Surface Height Prediction in the South China Sea. https://doi.org/10.20944/preprints202509.1707.v1

[12] Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., Hassanzadeh, P., Kashinath, K., and Anandkumar, A. (2022). FourCastNet: A global data-driven high-resolution weather model using adaptive Fourier neural operators. arXiv preprint arXiv:2202.11214. https://arxiv.org/abs/2202.11214

[13] Vaughan, A., Markou, S., Tebbutt, W., Requeima, J., Bruinsma, W. P., Andersson, T. R., Herzog, M., Lane, N. D., Chantry, M., Hosking, J. S., and Turner, R. E. (2024). Aardvark weather: End-to-end data-driven weather forecasting. arXiv preprint arXiv:2404.00411. https://arxiv.org/abs/2404.00411

[14] Bi, K., Xie, L., Zhang, H. et al. Accurate medium-range global weather forecasting with 3D neural networks. Nature 619, 533–538 (2023). https://doi.org/10.1038/s41586-023-06185-3

[15] Weyn, J. A., Durran, D. R., and Caruana, R. (2019). Can Machines Learn to Predict Weather? Using Deep Learning to Predict Gridded 500-hPa Geopotential Height From Historical Weather Data. Journal of Advances in Modeling Earth Systems, 11(8), 2680–2693. https://doi.org/10.1029/2019MS001705

[16] Sharma, H., Shrivastava, M., and Singh, B. (2023). Physics informed deep neural network embedded in a chemical transport model for the Amazon rainforest. Npj Climate and Atmospheric Science, 6(1), 28. https://doi.org/10.1038/s41612-023-00353-y

[17] Ye, F., Brodie, J., Miles, T., and Ezzat, A. A. (2023). AIRU-WRF: A physics-guided spatio-temporal wind forecasting model and its application to the U.S. Mid Atlantic offshore wind energy areas. arXiv preprint arXiv:2303.02246. https://arxiv.org/abs/2303.02246

[18] Khaira, U., Cerrai, D., Thompson, G., and Astitha, M. (2024). Integrating physics-based WRF atmospheric variables and machine learning algorithms to predict snowfall accumulation in Northeast United States. Journal of Hydrology, 644, 132113. https://doi.org/10.1016/j.jhydrol.2024.132113

[19] Hu, F., and Li, Q. (2024). Reconstruction of tropical cyclone boundary layer wind field using physics-informed machine learning. Physics of Fluids, 36(11). https://doi.org/10.1063/5.0234728

[20] Rasp, S., Dueben, P. D., Scher, S., Weyn, J. A., Mouatadid, S., and Thuerey, N. (2020). WeatherBench: A Benchmark Data Set for Data-Driven Weather Forecasting. Journal of Advances in Modeling Earth Systems, 12(11). https://doi.org/10.1029/2020MS002203

[21] Liu, Z., Zhang, J., Yang, Y., Wang, Y., Luo, W., and Zhou, X. (2024). Enhancing Weather Forecast Accuracy Through the Integration of WRF and BP Neural Networks: A Novel Approach. Earth and Space Science, 11(10). https://doi.org/10.1029/2024EA003613

[22] Chen, S., Feng, Y., Li, H., Ma, D., Mao, Q., Zhao, Y., and Liu, J. (2024). Enhancing runoff predictions in data-sparse regions through hybrid deep learning and hydrologic modeling. Scientific Reports, 14(1), 26450. https://doi.org/10.1038/s41598-024-77678-y

[23] Hao, Z., Liu, S., Zhang, Y., Ying, C., Feng, Y., Su, H., and Zhu, J. (2023). Physics-informed machine learning: A survey on problems, methods and applications. arXiv preprint arXiv:2211.08064. https://arxiv.org/abs/2211.08064

[24] F. Prill, D. Reinert, D. Rieger, and G. Zängl, "Working with the ICON Model," 2024.

[25] Visualization & Analysis Systems Technologies. (2017). Geoscience Community Analysis Toolkit: WRF-Python [Software]. Boulder, CO: UCAR/NCAR – Computational and Information Systems Laboratory. https://doi.org/10.5065/D6W094P1

[26] de Bézenac, E., Pajot, A., and Gallinari, P. (2018). Deep learning for physical processes: Incorporating prior scientific knowledge. arXiv preprint arXiv:1711.07970. https://arxiv.org/abs/1711.07970

[27] J. Willard, J., Jia, X., Xu, S., Steinbach, M., and Kumar, V. (2023). Integrating Scientific Knowledge with Machine Learning for Engineering and Environmental Systems. ACM Computing Surveys, 55(4), 1–37. https://doi.org/10.1145/3514228
https://doi.org/10.5676/DWD_pub/nwv/icon_tutorial2024

[28] le Toumelin, L., Gouttevin, I., Galiez, C., and Helbig, N. (2024). A two-fold deep-learning strategy to correct and downscale winds

over mountains. Nonlinear Processes in Geophysics, 31(1), 75–97. https://doi.org/10.5194/npg-31-75-2024

[29] Ezzat, A. A., Jun, M., and Ding, Y. (2018). Spatio-Temporal Asymmetry of Local Wind Fields and Its Impact on Short-Term Wind Forecasting. IEEE Transactions on Sustainable Energy, 9(3), 1437–1447. https://doi.org/10.1109/TSTE.2018.2789685