Nature Communications

https://doi.org/10.1038/s41467-026-70562-5

Article in Press

Reconstructing fine-scale 3D wind fields with terrain-informed machine learning

Received: 19 July 2025

Accepted: 2 March 2026

Cite this article as: Lin, C., Tie, R., Yi, S. *et al.* Reconstructing fine-scale 3D wind fields with terrain-informed machine learning. *Nat Commun* (2026). https://doi.org/10.1038/s41467-026-70562-5

Chensen Lin, Ruian Tie, Shihong Yi, Dongqing Liu, Xiaohui Zhong, Zixin Hu & Hao Li

Reconstructing fine-scale 3D wind fields with terrain-informed machine learning

Chensen Lin[1†], Ruian Tie[1,2†], Shihong Yi[1], Dongqing Liu[3], Xiaohui Zhong[1,4], Zixin Hu[1], Hao Li[1,2,4,5*]

[1]Artificial Intelligence Innovation and Incubation Institute, Fudan University, Shanghai, China.
[2]Shanghai Innovation Institute, Shanghai, China.
[3]Meteorological Observatory, Nanjing Meteorological Bureau, Nanjing, China.
[4]Joint Laboratory for AI-Based Earth System Forecasting, Fudan University, Shanghai, China.
[5]Shanghai Academy of Artificial Intelligence for Science, Shanghai, China.

*Corresponding author(s). E-mail(s): lihao_lh@fudan.edu.cn;
Contributing authors: linchensen@fudan.edu.cn; ratie25@m.fudan.edu.cn; shyi25@m.fudan.edu.cn; cma_ldq@sina.com; zhongxiaohui@fudan.edu.cn; huzixin@fudan.edu.cn;
†These authors contributed equally to this work.

Abstract

Fine-scale near-surface wind field prediction is essential for a wide range of applications. However, most operational and AI-based weather models operate at kilometer-scale resolution, where terrain-induced wind features such as slope jets, flow deflection, and recirculation are systematically averaged out. Here we introduce FuXi-CFD, a machine learning-based framework designed to generate detailed three-dimensional (3D) near-surface wind fields at 30-meter horizontal resolution, using only coarse-resolution atmospheric inputs and high-resolution terrain information. The model is trained on a large-scale dataset generated via computational fluid dynamics (CFD), encompassing a wide range of terrain types and inflow conditions. Although relying only on horizontal wind inputs, FuXi-CFD infers the full 3D wind fields—including latent variables such as vertical velocity and turbulence-related features. It achieves CFD-comparable accuracy

while reducing inference time from hours to seconds. Notably, the model also generalizes well to real-world conditions, as demonstrated by consistent performance against independent wind-tower observations. This capability enables real-time wind field reconstruction for terrain-sensitive applications such as wind turbine siting, power forecasting, and wildfire spread modeling.

Keywords: fine-scale atmospheric modeling, complex terrain, computational fluid dynamics, machine learning

1 Introduction

Fine-scale wind field information is essential for a wide range of high-impact applications, including wind turbine siting, power forecasting, wildfire spread prediction, and urban risk assessment [1–4]. These tasks require accurate, site-specific estimates of wind speed, direction, and turbulence—often at spatial resolutions of tens of meters [5, 6].

While coarse-resolution forecasts may suffice in relatively flat regions, such approaches become inadequate in areas with complex topography. Surface features such as mountains, ridges, and valleys can strongly modulate wind behavior, generating localized phenomena including slope jets (from flow channeling along inclined surfaces), vertical wind shear (due to sudden elevation changes), and recirculation zones (in the lee of steep terrain) [7]. These terrain-induced flow patterns are often poorly correlated with the surrounding kilometer-scale average and may even exhibit opposite flow directions.

Most operational weather forecasting models, including ECMWF's IFS [8], NOAA's GFS [9], and China's GRAPES [10], operate at horizontal resolutions of several kilometers. At this scale, terrain-induced flow structures are systematically averaged out, limiting the utility of these models for site-specific assessments. While such models perform well for synoptic-scale weather prediction, they fall short in resolving the fine-scale variability required for localized or site-specific applications.

Recent advances in machine learning have enabled a new class of data-driven weather forecasting models. Notable examples include GraphCast [11], FourCastNet [12], Pangu-Weather [13], FuXi [14, 15], MetNet [16], and GenCast [17]. These models are typically trained on global reanalysis datasets such as ERA5 [18], which offer native resolutions of around 0.25° (approximately 31 km). While they have achieved remarkable improvements in forecast accuracy and lead time, their spatial resolution has seen little advancement, as the outputs remain fundamentally constrained by the resolution of the training data. This limitation restricts their ability to resolve fine-scale wind features, particularly in regions with complex terrain.

To obtain finer-scale information from coarse atmospheric forecasts, various downscaling techniques have been explored. Statistical approaches, based on interpolation or regression, provide computationally efficient corrections [19] but struggle to generalize in regions lacking dense measurements and cannot reliably recover the three-dimensional, turbulent structures of wind flow [20]. Dynamical downscaling

approaches using mesoscale weather models such as WRF offers improved physical realism [21] but incurs high computational costs and suffers from limitations in turbulence parameterization and numerical stability in steep terrain [22]. These limitations have restricted their application in high-resolution operational wind forecasting, especially in mountainous regions [23].

In recent years, machine-learning–based downscaling has attracted growing interest, and a wide variety of approaches—ranging from convolutional super-resolution networks [24–26], GAN models [25, 27], operator-learning architectures [28–30] and more recent diffusion or Transformer variants [31–33]—have demonstrated encouraging progress in enhancing coarse atmospheric fields. These studies collectively show that data-driven approaches can extract spatial patterns from low-resolution inputs and offer a computationally efficient alternative to traditional statistical or dynamical downscaling.

Most existing works, however, operate within a resolution range constrained by the availability of training data: "high-resolution" targets are typically at the kilometer scale, and the learning task often involves recovering a synthetically degraded version of the same fields. As a result, current ML models typically learn scale transformations from 20–30 km down to 1–3 km, rather than the tens-of-meters scales needed to resolve terrain-induced accelerations or flow deformation. This reflects a broader limitation of today's public datasets rather than any specific modeling choice.

In addition, ML downscaling is generally formulated as a two-dimensional super-resolution problem focused on a single horizontal layer. While this formulation is convenient and has yielded promising improvements, fine-scale mountain flows are inherently three-dimensional, involving slope flows, vertical motions, and strongly deformed streamlines that cannot be inferred from 2D inputs alone. Extending ML approaches toward reconstructing full 3D wind structures therefore remains an open challenge for the field.

Finally, most studies train and evaluate models within the same terrain region, and only a few works have begun exploring generalization across distinct orographic settings. Early results [34] suggest that performance tends to degrade when applied to terrains not represented in the training data—a natural consequence of the limited diversity of existing high-resolution wind datasets.

To resolve wind flows at tens-of-meter resolution in complex terrain, computational fluid dynamics (CFD) [35, 36] remains the most physically grounded modeling approach. By explicitly tackling the Navier–Stokes equations with detailed representations of terrain geometry, surface roughness length, and boundary-layer processes, CFD can capture flow separation, channeling, speed-up effects, and turbulence generated by topographic forcing—features that cannot be inferred from kilometer-scale atmospheric models or purely data-driven downscaling. High-resolution CFD has thus become the de facto standard in applications requiring fine-scale wind characterization, including wind-energy assessment, infrastructure design, dispersion studies, and safety evaluations, where spatial resolutions of 30–100 m are commonly employed [37–39].

Despite its accuracy, CFD remains computationally intensive and operationally demanding. Configuring a realistic domain, preparing terrain and boundary conditions,

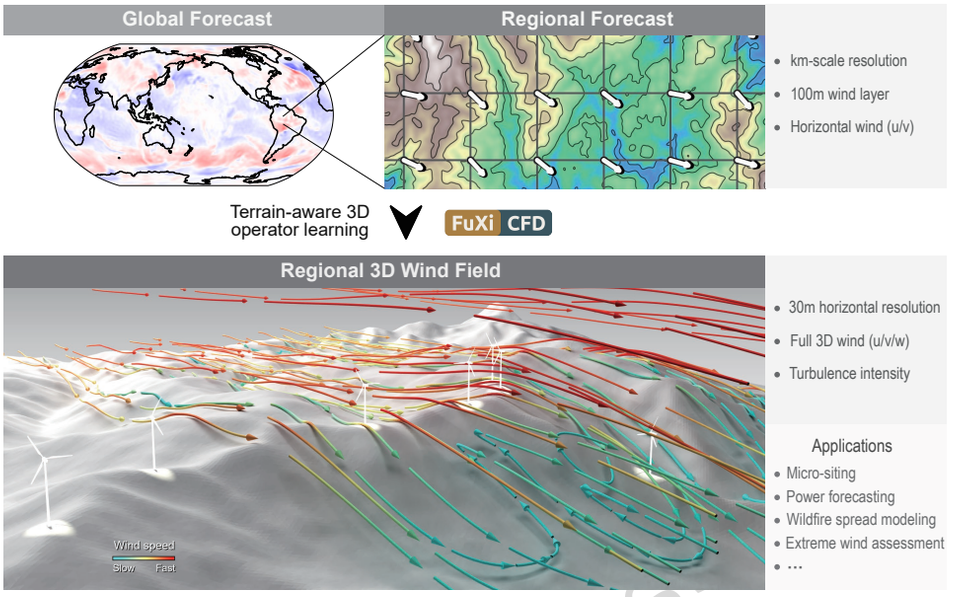

Fig. 1 Terrain-aware 3D wind reconstruction with FuXi-CFD. FuXi-CFD converts kilometer-scale horizontal winds from global or regional forecasts into 30 m-resolution three-dimensional wind fields over complex terrain. The model learns a terrain-response operator from a large computational fluid dynamics (CFD)-generated library, enabling prediction of full velocity vector and turbulence-related quantities from coarse inputs.

and running a high-resolution simulation over complex terrain can require substantial expert effort and runtime, often making it impractical for real-time forecasting.

To address these challenges, this study develops a data-driven framework termed FuXi-CFD (Fig. 1), which combines the physical fidelity of CFD with the efficiency of modern machine learning. We construct a large CFD-informed dataset covering diverse terrain shapes and inflow conditions, explicitly resolving the fine-scale flow responses that are absent from kilometer-scale atmospheric inputs. Building on this dataset, we train a learning model that approximates the terrain–flow operator, enabling the reconstruction of three-dimensional wind fields—including vertical motion and turbulence-related quantities—from coarse horizontal winds alone. The resulting framework offers a physically grounded yet computationally efficient pathway toward high-resolution wind prediction in complex terrain, supporting applications such as micro-siting, power forecasting, wildfire spread modeling, and extreme-wind assessment.

2 Results

2.1 Overall predictive accuracy

Built upon a large CFD-informed dataset and a terrain-aware operator learning framework (see Methods), FuXi-CFD's predictive skill was first assessed by comparing its outputs with the CFD references across unseen terrains. We focus on four key variables—u, v, w, and k—that characterize near-surface and turbine-height winds. Figure 2 summarizes the model–reference agreement at 10 m, 50 m, and 100 m—three heights that span standard meteorological observations to typical rotor-layer elevations.

Across all heights, FuXi-CFD reproduces the horizontal wind components with high fidelity, with predictions tightly following the 1:1 line and showing minimal scatter. The vertical velocity exhibits slightly larger spread but remains well constrained, reflecting the model's ability to recover terrain-induced updraft and downdraft structures from coarse background wind information. In contrast, predictions of turbulent kinetic energy show the largest dispersion. This behavior is expected: although k is included in the training labels, it is not supplied as an explicit input, and thus acts as a latent quantity that must be inferred indirectly from the flow state.

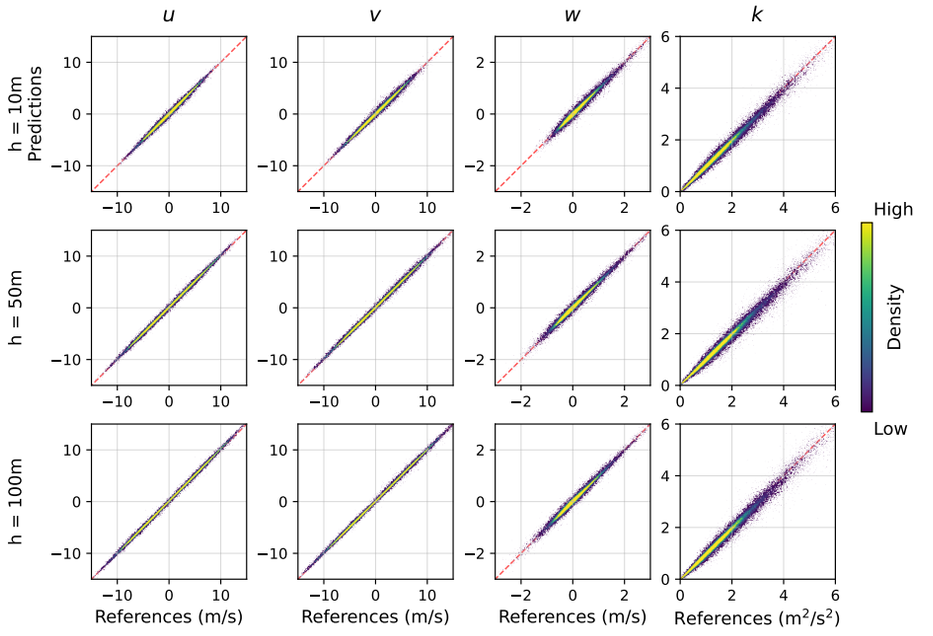

Fig. 2 Agreement between FuXi-CFD predictions and computational fluid dynamics (CFD) references. Parity plots compare predicted and simulated horizontal winds (u, v), vertical velocity (w), and turbulent kinetic energy (k) at 10 m, 50 m, and 100 m height. Predictions of u and v align closely with the CFD truth across all heights, indicating that the model recovers both large-scale inflow directionality and fine-scale terrain-modulated variations. Scatter is largest for k, reflecting its derived nature and dependence on small-scale shear. Overall, the results demonstrate strong predictive skill for primary wind components and reasonable performance for turbulence-related quantities.

Performance improves with height, and the 100 m level shows the strongest agreement across all four variables. This is consistent with the fact that the upper levels are more strongly anchored by the coarse-scale input winds, whereas the near-surface layers rely more heavily on learning terrain-driven adjustments. Overall, the results demonstrate that FuXi-CFD captures both the large-scale momentum distribution and the fine-scale terrain-modulated responses across a wide range of heights.

2.2 Spatial pattern fidelity

To complement the pointwise statistics, we next examine how FuXi-CFD performs at the case level across different terrain conditions (Fig. 3). Panel (a) summarizes all test cases as a function of terrain relief and their mean absolute error. Most cases

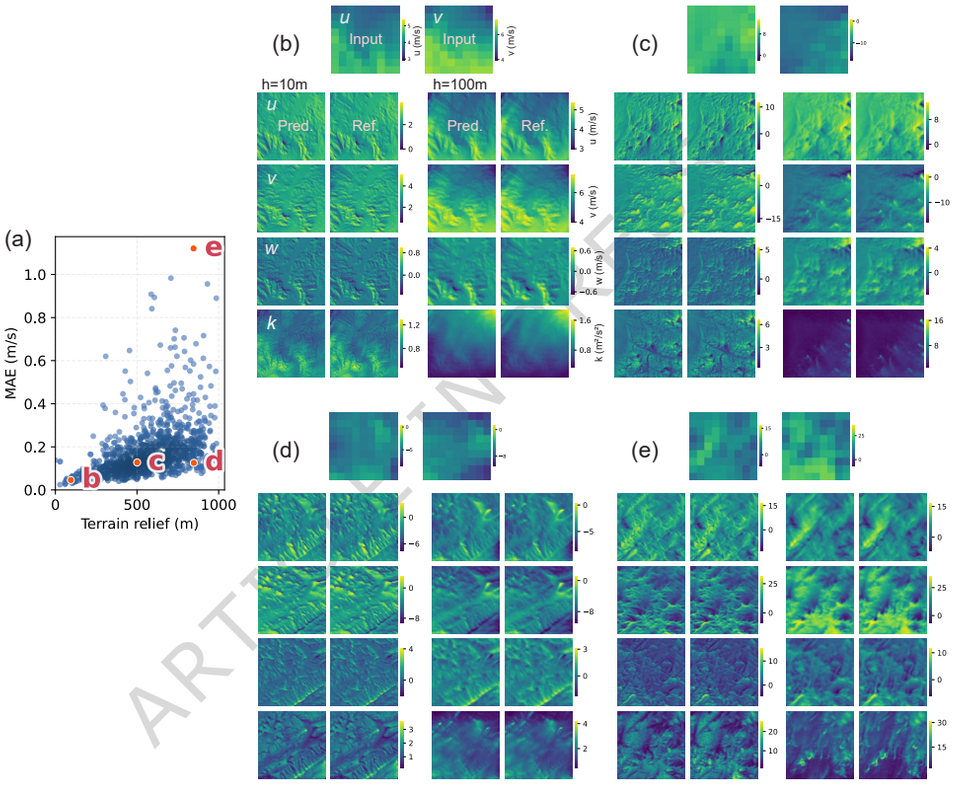

Fig. 3 Examples of multi-height wind-field reconstructions across diverse terrain conditions. (a) Distribution of test cases as a function of terrain relief, with selected examples spanning low to extreme terrain relief. (b–e) For each case, coarse 10 m winds (9×9 grid) are shown alongside high-resolution predictions and computational fluid dynamics (CFD) references at 10 m, 50 m, and 100 m. The reconstructed fields capture speed-up over ridges, flow channeling in valleys, and height-dependent smoothing as terrain influence weakens. Across diverse conditions, FuXi-CFD recovers both broad flow patterns and local gradients with good fidelity, despite a downscaling ratio exceeding 30×. Beyond horizontal winds, the model also infers vertical motion (w) and turbulent kinetic energy (k), which are absent from the inputs yet essential for wind-related applications.

cluster at low errors over a wide range of relief, with a few challenging examples at large terrain relief. From this distribution we select four representative cases (b–e): three lie within the dense cloud of typical performance, covering gentle, moderate, and strongly structured terrain, while case (e) corresponds to an extreme-relief, higher-error example near the upper envelope of the scatter. This design allows us to assess not only average behavior but also how the model behaves in the most demanding settings.

For each selected case, panel (b–e) compares the coarse 100 m input (9×9) with the reconstructed high-resolution fields and the CFD references at 10 m, 50 m, and 100 m for all four variables (u, v, w, k). Across these diverse landscapes, FuXi-CFD preserves the key terrain-forced structures: speed-up over ridges, flow channeling and turning in valleys, and progressive smoothing of small-scale features with height as the direct terrain influence weakens. The predicted patterns closely track the CFD references in both broad gradients and local variations, rather than collapsing to an over-smoothed version of the coarse input. This remains true even for the extreme case (e), where complex orography and larger overall errors make the reconstruction most challenging. Additional three-dimensional streamline visualizations, which provide a direct view of terrain-induced flow structures, are provided in Supplementary Notes 5.

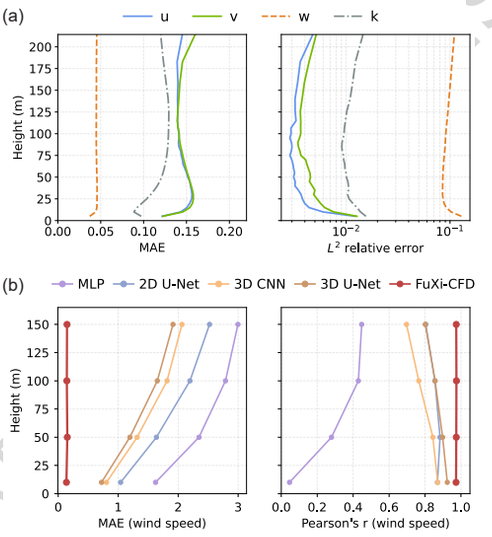

Fig. 4 Height-resolved performance and comparison with conventional models. (a) Vertical profiles of prediction errors for horizontal winds (u, v), vertical velocity (w), and turbulent kinetic energy (k), evaluated using mean absolute error (MAE) and relative L^2 metrics. Errors remain low across the 0–200 m layer, with slightly larger deviations near the surface. (b) Wind-speed prediction performance across models as a function of height. Conventional architectures trained on the same dataset exhibit larger errors and lower correlations than FuXi-CFD.

2.3 Height-resolved performance

FuXi-CFD predicts full three-dimensional wind fields, enabling evaluation of performance across height rather than at a single level. Across the 0–200 m layer (Fig. 4(a)), prediction errors remain consistently low for horizontal winds, with only modest increases near the surface where terrain-induced gradients are strongest. Vertical velocity also shows stable accuracy despite not being provided as input, whereas turbulent kinetic energy exhibits larger deviations consistent with its stronger sensitivity to local shear. Importantly, accuracy does not deteriorate away from the constrained 100 m level, indicating robust vertical generalization.

We further compare FuXi-CFD with several widely used deep-learning architectures trained on the same dataset (Fig. 4(b)). Across all evaluated heights, conventional models exhibit substantially larger errors and lower correlations. Additional comparisons across variables, heights, and representative cases are provided in Supplementary Notes 3.3.

The results show that, even when trained on the same data, conventional models fail to reproduce the height-dependent flow variability resolved by FuXi-CFD. This suggests that reconstructing high-resolution three-dimensional wind fields from coarse inputs is not a trivial mapping but requires models capable of learning terrain-aware three-dimensional flow operators.

2.4 Validation

To assess the applicability and generalization capability of FuXi-CFD under real-world complex terrain conditions, we conduct independent validation using tall-tower observations [40]. As a representative example, the OPE meteorological tower in Europe is considered. Figures 5(a1)-(a3) show a photograph of the tower site, the digital elevation model (DEM) and surface roughness length (z_0) over a surrounding area of $9\,\mathrm{km} \times 9\,\mathrm{km}$. The terrain exhibits moderate elevation variations (on the order of several tens of meters), introducing noticeable but relatively gentle terrain-induced flow variations, making it a practical site for terrain-driven wind-flow validation.

The OPE tower provides wind speed and wind direction measurements at heights of 10 m, 50 m, and 120 m with a temporal resolution of 3 h. We collect three years of observations from 2018 to 2020 and compare them against FuXi-CFD reconstructions. Figure 5(b) presents the mean absolute error (MAE) at 10 m, 50 m, and 120 m as a function of wind speed interval. Two ERA5-based site-level vertical reconstruction baselines are considered. ERA5-SC denotes a single-constraint power-law approach that extrapolates a vertical wind profile using only the ERA5 10 m wind speed, while ERA5-DC employs both the ERA5 10 m and 100 m winds to construct a two-level-constrained profile. The latter reflects the maximum information recoverable from coarse-grid reanalysis at a point location, without explicitly resolving terrain effects or introducing additional expensive dynamical simulations. Detailed formulations and parameter settings of these baselines are provided in the Supplementary Information.

The results indicate that ERA5-DC substantially reduces errors relative to ERA5-SC, confirming its suitability as a strong baseline. Under moderate to strong wind

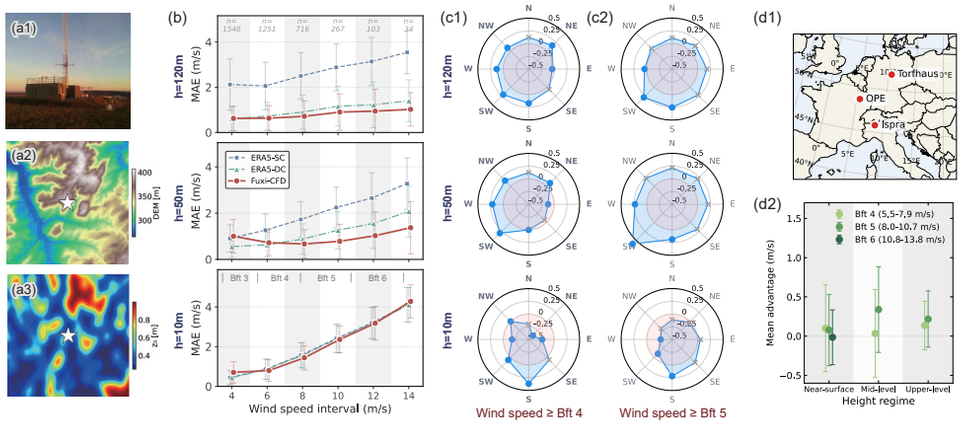

Fig. 5 Validation of FuXi-CFD against tall-tower observations over complex terrain. (a1–a3) Photograph of the European OPE meteorological tower, together with the digital elevation model (DEM) and roughness length distribution over an approximately 9 km × 9 km area centered on the tower. The photograph in (a1) is reproduced from ICOS RI (https://meta.icos-cp.eu/resources/stations/AS_OPE), licensed under CC BY 4.0. (b) Mean absolute error at 120 m, 50 m, and 10 m as a function of wind-speed interval. Baseline comparisons include ERA5-SC (single-constraint extrapolation) and ERA5-DC (dual-constraint extrapolation). ERA5-DC clearly outperforms ERA5-SC, while FuXi-CFD further reduces errors under moderate to strong wind conditions, with the advantage increasing at higher wind speeds. Error bars represent ±1 standard deviation. (c) Directional distribution of the error difference between FuXi-CFD and ERA5-DC. FuXi-CFD exhibits the largest improvement for southwesterly winds, indicating that the model achieves greater performance gains in scenarios where terrain effects are pronounced, with particularly strong advantages in high-wind, high-value conditions. (d1–d2) Statistical validation across three independent European tall-tower sites (OPE, Torfhaus, and Ispra). Consistent performance improvements are observed across sites, height regimes, and wind-speed ranges. The model is trained on complex terrain data from China and validated at European sites, demonstrating stable cross-regional generalization.

conditions, FuXi-CFD further improves reconstruction accuracy relative to ERA5-DC, with the performance advantage increasing systematically with wind speed. This behavior is particularly relevant for applications such as wind power assessment and strong-wind hazard mitigation, where higher wind-speed regimes correspond to rapidly increasing power output or structural loading, and even small improvements in wind-speed accuracy can lead to disproportionately larger impacts on engineering estimates.

To further examine the role of terrain effects, Fig. 5(c) shows the directional dependence of the error difference between FuXi-CFD and ERA5-DC. Based on the DEM, southwesterly inflow corresponds to flow ascending from relatively lower terrain toward the tower location, where terrain-induced effects are expected to be most pronounced. The marked performance advantage of FuXi-CFD in this wind sector indicates its ability to capture direction-dependent terrain modulation of wind-profile structure. Statistics conditioned on Beaufort scale ≥4 and ≥5 further show that this advantage strengthens under increasingly strong mechanically driven flow regimes.

In addition to OPE, two other European tall-tower sites, Torfhaus and Ispra, are included in the validation. As the available measurement heights differ among sites, observations are grouped into three height regimes: near-surface (∼10 m), mid-level

(20–90 m), and upper-level (90–130 m). Figure 5(d2) summarizes the performance differences between FuXi-CFD and ERA5-DC across wind-speed ranges and height regimes. Consistent performance characteristics are observed across sites and conditions, confirming the robustness of the conclusions drawn from the OPE case. Despite being trained solely on simulations over complex terrain in southeastern China, FuXi-CFD exhibits consistent performance at European sites, indicating robust cross-regional generalization.

3 Discussion

This study presents a three-dimensional wind-field reconstruction framework, FuXi-CFD, which integrates large-scale CFD data with deep learning to learn a terrain-response operator and achieve high-resolution (30 m) wind-field reconstruction from kilometer-scale background winds over complex terrain. Results show that FuXi-CFD maintains stable predictive performance on unseen terrains: it not only achieves low overall errors but also accurately reproduces terrain-forced structures such as speed-up over ridges, channeling in valleys, and upwind–downwind slope features, while exhibiting consistent vertical error patterns across the 0–200 m layer. In addition to the high-resolution horizontal winds, the model also infers vertical motion and certain turbulence-related features that are not explicitly provided in the inputs, demonstrating coherent three-dimensional structural consistency.

From the perspective of boundary-layer physics, although the dataset used in this study is constructed under the neutral atmospheric boundary layer assumption, FuXi-CFD remains applicable in real moderate-to-strong wind conditions. Such dynamically dominated regimes are typically characterized by low Richardson numbers (Ri)—indicating that buoyancy contributes far less than shear-generated turbulence—so the flow structure is primarily governed by terrain forcing and shear, with buoyancy acting only as a secondary adjustment. Under these conditions, even without explicitly incorporating thermal stratification, the model maintains robust performance in year-round wind-tower statistics. Under weak-wind conditions, however, the prediction errors are higher, partly because the dataset emphasizes strong-wind cases and contains relatively few weak-wind samples, and partly because thermal effects exert stronger influence on near-surface flow under weak-wind regimes, which are not represented in the current CFD database.

Future work could extend the existing dataset by incorporating CFD cases that include thermal effects, adding background stratification, surface heat flux, and other relevant thermal inputs to improve performance under weak-wind and strongly stable conditions. However, such an extension requires higher-dimensional inputs, broader parameter-space coverage, and significantly increased computational cost per case, and its benefits mainly target the finer characterization of low-wind scenarios. For applications dominated by moderate-to-strong winds—such as wind-energy assessment, hazard forecasting, and low-altitude safety—FuXi-CFD already demonstrates reasonable physical consistency, stable generalization, and clear engineering relevance.

Overall, FuXi-CFD provides an efficient, deployable, and physically grounded pathway for high-resolution three-dimensional wind-field downscaling in complex terrain,

bringing data-driven wind downscaling into the 30-meter regime in a practically meaningful way.

4 Methods

4.1 CFD-informed Dataset

The FuXi-CFD dataset consists of more than 12,000 high-resolution CFD simulations of real mountainous terrain in southeastern China. Each sample covers a 9 km × 9 km region at 30 m horizontal resolution and provides three-dimensional fields of wind velocity components (u, v, w) and turbulent kinetic energy (k). The dataset spans a wide range of terrain geometries, elevation gradients, and roughness patterns, forming a large and physically consistent foundation for learning terrain-induced wind variability.

Two physical inputs define each terrain characteristics: elevation and surface roughness. Elevation is obtained from the 30 m Shuttle Radar Topography Mission (SRTM) dataset [41], while surface roughness is derived from ESA WorldCover land cover maps [42] using established empirical mappings to aerodynamic roughness length [43]. These globally available datasets define the underlying terrain–atmosphere interaction environment.

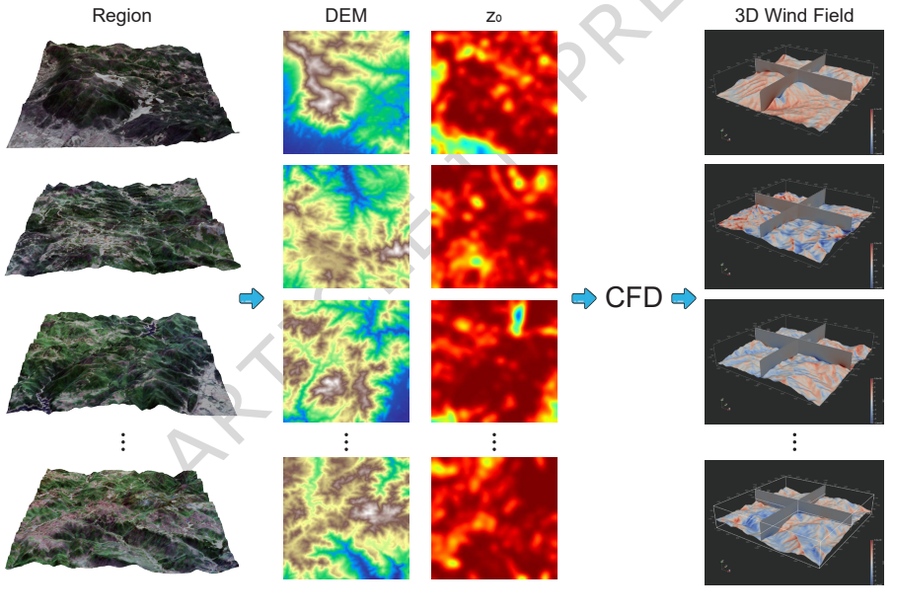

Fig. 6 Computational fluid dynamics (CFD)-informed dataset generation. Each row illustrates a representative terrain block from the dataset, starting from satellite imagery and its corresponding digital elevation model (DEM) and surface roughness length (z_0) fields. For each terrain block and inflow direction, CFD produces a 3D wind field, shown as horizontal slices and two orthogonal vertical sections. These simulations form a comprehensive library of terrain–flow interactions, which FuXi-CFD learns to approximate.

Figure 6 summarizes the dataset-generation workflow. For each selected region, terrain elevation and roughness length are extracted as physical inputs. These inputs, together with a specified inflow condition, drive a high-resolution CFD simulation that produces a three-dimensional wind field over complex terrain.

To ensure physically realistic flow development, each CFD simulation uses a 12 km × 12 km domain, where the central 9 km × 9 km region serves as the output region and the surrounding 1.5 km buffer allows the flow to adjust to terrain forcing before entering the core domain. Simulations are performed using a steady-state Reynolds-Averaged Navier–Stokes (RANS) solver (details in Supplementary Information).

The inflow is initialized with a logarithmic atmospheric boundary layer (ABL) profile. The reference wind speed at 100 m is sampled from a skewed distribution between 3 and 20 m/s, with higher probability around 10 m/s to emphasize typical turbine operating conditions. Wind direction is sampled uniformly from the full 360° range. These choices produce a wide variety of realistic inflow scenarios, helping the machine-learning model learn terrain-induced flow responses across a broad operational envelope.

Each simulation yields full 3D fields of the wind velocity components (u, v, w) and turbulent kinetic energy (k) on a 300 × 300 × 47 grid extending up to 3 km in

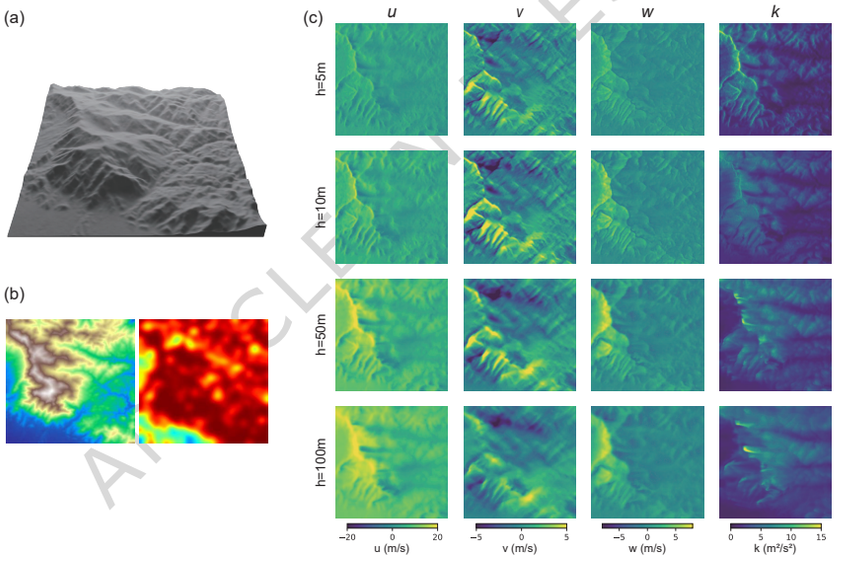

Fig. 7 Terrain-induced variability in computational fluid dynamics (CFD)-generated wind fields. (a–b) Topography and roughness length for an example region. (c) Horizontal winds (u, v), vertical velocity (w), and turbulent kinetic energy (k) at several heights show strong spatial variation driven by terrain forcing. Lower levels (5–10 m) exhibit sharp gradients associated with slope-parallel accelerations and localized separation zones, while flow patterns become smoother with altitude as terrain influence weakens. Turbulence intensity remains elevated over steep slopes and leeward regions, reflecting shear-induced production captured by the CFD solver. These examples highlight the diversity of flow responses that the model must learn to reproduce from coarse inputs.

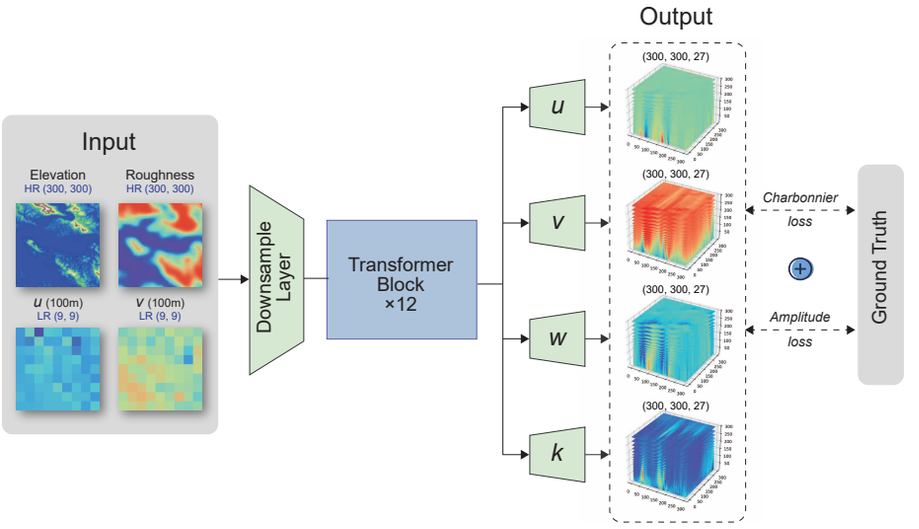

Fig. 8 Architecture of the terrain-informed deep learning model for 3D wind downscaling. The model takes high-resolution terrain (300×300 grid) and coarse 100 m winds (9×9 grid) as inputs. A downsampling layer and twelve transformer encoder blocks extract cross-scale spatial dependencies, which are then decoded into full 3D fields of horizontal winds (u, v), vertical velocity (w), and turbulent kinetic energy (k) at 30 m horizontal resolution and 27 vertical levels. Training is supervised with a combination of Charbonnier and frequency-domain losses to encourage both local accuracy and structural coherence.

altitude. In total, the dataset contains more than 12,000 simulations and represents over 300,000 CPU hours of computation.

The resulting library captures a diverse set of terrain–flow interactions. As illustrated in Figure 7, lower-altitude winds exhibit strong acceleration, separation, and channeling associated with local topographic forcing, while flow patterns become gradually smoother with height. Turbulence intensity remains elevated over steep slopes and leeward regions due to shear-induced production resolved by the RANS model.

Although inflow wind speed and direction determine each CFD solution, they are deliberately excluded from the dataset inputs used for machine-learning training. This design reflects the target application: in practical wind-energy assessment, upstream flow conditions are often uncertain, unavailable, or only coarsely resolved. By training on diverse CFD outcomes without providing inflow metadata, the model learns to infer terrain-sensitive 3D flow patterns from limited inflow information, enabling applications where forward CFD would otherwise be inapplicable. A detailed correlation analysis is provided in Supplementary Notes 1.3, illustrating the non-trivial relationships between dataset inputs and outputs across heights.

Overall, the FuXi-CFD dataset provides a large, physically consistent basis for learning terrain-response operators that reconstruct fine-scale, three-dimensional wind fields.

4.2 Machine learning-based downscaling model

To reconstruct fine-scale three-dimensional wind fields from coarse atmospheric forecasts, we design a deep learning model that maps kilometer-scale wind inputs to high-resolution 3D wind structures. Specifically, the model takes as input 100 m wind components (u, v) sampled on a 9×9 grid at 1 km resolution, covering a 9 km $\times$ 9 km area. In parallel, terrain elevation and roughness length maps are provided at 30 m resolution over the same region. These inputs represent the typical outputs of modern machine learning-based forecasting systems combined with satellite-derived surface data. The model is trained to infer four key physical variables—u, v, w and k—across 27 vertical levels up to 214 m, forming a four-channel volumetric output used for wind energy assessment.

Figure 8 provides an overview of the model architecture. We adopt a shared-encoder architecture with four task-specific decoder branches. The encoder is built upon a Vision Transformer [44] backbone, which efficiently captures terrain-induced spatial correlations and long-range dependencies that are challenging for conventional convolutional models. The four input fields are concatenated into a tensor of shape (batch size, 4, 300, 300), corresponding to a 9 km $\times$ 9 km area at 30 m resolution. The shared encoder extracts latent features from this input, which are then processed by separate decoder branches to reconstruct each target variable.

To enhance spatial fidelity and suppress noise in the predicted fields, we employ a hybrid loss function combining a robust spatial-domain loss with a frequency-domain regularization term. The former promotes local consistency, while the latter encourages the preservation of large-scale structural patterns. The model is trained using the AdamW optimizer with a stepwise learning rate schedule over approximately 128 hours on two A100 GPUs. This joint loss design leads to improved accuracy and more stable convergence across all target variables, particularly for the latent fields w and k that are not directly observable in the input. Additional model details are provided in the Supplementary Notes 3.1.

Once trained, the model generates high-resolution 3D wind and turbulence fields within a fraction of a second per query, representing a speedup of over three orders of magnitude compared to traditional CFD simulations (Supplementary Notes 3.2).

Data availability. The FuXi-CFD dataset generated in this study is publicly available at Zenodo (https://doi.org/10.5281/zenodo.18770845). All data necessary to reproduce the results reported in this paper are publicly available.

Code Availability. The inference code and pre-trained model weights used in this study are publicly available at Zenodo (https://doi.org/10.5281/zenodo.18770845).

References

[1] Barthelmie, R.J., Pryor, S.C.: Automated wind turbine wake characterization in complex terrain. Atmospheric Measurement Techniques **12**(6), 3463–3484 (2019)

[2] Pan, Y., Archer, C.L.: A hybrid wind-farm parametrization for mesoscale and climate models. Boundary-Layer Meteorology **168**(3), 469–495 (2018)

[3] Floors, R., Vincent, C.L., Gryning, S.-E., Peña, A., Batchvarova, E.: The wind profile in the coastal boundary layer: Wind lidar measurements and numerical modelling. Boundary-layer meteorology **147**(3), 469–491 (2013)

[4] Linn, R., Winterkamp, J., Edminster, C., Colman, J.J., Smith, W.S.: Coupled influences of topography and wind on wildland fire behaviour. International Journal of Wildland Fire **16**(2), 183–195 (2007)

[5] Veers, P., Dykes, K., Lantz, E., Barth, S., Bottasso, C.L., Carlson, O., Clifton, A., Green, J., Green, P., Holttinen, H., *et al.*: Grand challenges in the science of wind energy. Science **366**(6464), 2027 (2019)

[6] Clifton, A., Barber, S., Stökl, A., Frank, H., Karlsson, T.: Research challenges and needs for the deployment of wind energy in hilly and mountainous regions. Wind Energy Science **7**(6), 2231–2254 (2022)

[7] He, Y., Chan, P., Li, Q.: Wind characteristics over different terrains. Journal of Wind Engineering and Industrial Aerodynamics **120**, 51–69 (2013)

[8] Haiden, T., Janousek, M., Vitart, F., Bouallegue, Z.B., Ferranti, L., Prates, F., Richardson, D.: Evaluation of ecmwf forecasts, including the 2018 upgrade. Technical Report ECMWF Technical Memorandum 831, European Centre for Medium-Range Weather Forecasts (2018)

[9] Harris, L., Chen, X., Putman, W., Zhou, L., Chen, J.-H.: A scientific description of the gfdl finite-volume cubed-sphere dynamical core (2021)

[10] Chen, D., Xue, J., Yang, X., Zhang, H., Shen, X., Hu, J., Wang, Y., Ji, L., Chen, J.: New generation of multi-scale nwp system (grapes): General scientific design. Chinese Science Bulletin **53**(22), 3433–3445 (2008)

[11] Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., *et al.*: Learning skillful medium-range global weather forecasting. Science **382**(6677), 1416–1421 (2023)

[12] Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., et al.: Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. arXiv preprint arXiv:2202.11214 (2022)

[13] Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., Tian, Q.: Accurate medium-range global weather forecasting with 3d neural networks. Nature **619**(7970), 533–538 (2023)

[14] Chen, L., Zhong, X., Zhang, F., Cheng, Y., Xu, Y., Qi, Y., Li, H.: Fuxi: A cascade machine learning forecasting system for 15-day global weather forecast. npj climate and atmospheric science **6**(1), 190 (2023)

[15] Chen, L., Zhong, X., Li, H., Wu, J., Lu, B., Chen, D., Xie, S.-P., Wu, L., Chao, Q., Lin, C., *et al.*: A machine learning model that outperforms conventional global subseasonal forecast models. Nature Communications **15**(1), 6425 (2024)

[16] Espeholt, L., Agrawal, S., Sønderby, C., Kumar, M., Heek, J., Bromberg, C., Gazen, C., Carver, R., Andrychowicz, M., Hickey, J., *et al.*: Deep learning for twelve hour precipitation forecasts. Nature communications **13**(1), 1–10 (2022)

[17] Price, I., Sanchez-Gonzalez, A., Alet, F., Andersson, T.R., El-Kadi, A., Masters, D., Ewalds, T., Stott, J., Mohamed, S., Battaglia, P., et al.: Gencast: Diffusion-based ensemble forecasting for medium-range weather. arXiv preprint arXiv:2312.15796 (2023)

[18] Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., *et al.*: The era5 global reanalysis. Quarterly journal of the royal meteorological society **146**(730), 1999–2049 (2020)

[19] Schoof, J.T.: Statistical downscaling in climatology. Geography Compass **7**(4), 249–265 (2013)

[20] Martinez-García, F.P., Contreras-de-Villar, A., Muñoz-Perez, J.J.: Review of wind models at a local scale: advantages and disadvantages. Journal of Marine Science and Engineering **9**(3), 318 (2021)

[21] Soares, P.M., Cardoso, R.M., Miranda, P.M., Medeiros, J., Belo-Pereira, M., Espirito-Santo, F.: Wrf high resolution dynamical downscaling of era-interim for portugal. Climate dynamics **39**, 2497–2522 (2012)

[22] Fernández-González, S., Martín, M.L., García-Ortega, E., Merino, A., Lorenzana, J., Sánchez, J.L., Valero, F., Rodrigo, J.S.: Sensitivity analysis of the wrf model: Wind-resource assessment for complex terrain. Journal of Applied Meteorology and Climatology **57**(3), 733–753 (2018)

[23] Goger, B., Rotach, M.W., Gohm, A., Stiperski, I., Fuhrer, O.: Current challenges for numerical weather prediction in complex terrain: Topography representation and parameterizations. In: 2016 International Conference on High Performance Computing & Simulation (HPCS), pp. 890–894 (2016). IEEE

[24] Dupuy, F., Durand, P., Hedde, T.: Downscaling of surface wind forecasts using convolutional neural networks. Nonlinear Processes in Geophysics **30**(4), 553–570 (2023)

[25] Kumar, A., Islam, T., Ma, J., Kashiyama, T., Sekimoto, Y., Mattmann, C.: Windsr: Improving spatial resolution of satellite wind speed through super-resolution. IEEE Access **11**, 69486–69494 (2023)

[26] Lian, J., Huang, S., Shao, J., Chen, P., Tang, S., Lu, Y., Yu, H.: Terrawind: A deep learning-based near-surface winds downscaling model for complex terrain region. Geophysical Research Letters **51**(23), 2024–112124 (2024)

[27] Buster, G., Benton, B.N., Glaws, A., King, R.N.: High-resolution meteorology with climate change impacts from global climate model data using generative machine learning. Nature Energy **9**(7), 894–906 (2024)

[28] Jiang, P., Yang, Z., Wang, J., Huang, C., Xue, P., Chakraborty, T., Chen, X., Qian, Y.: Efficient super-resolution of near-surface climate modeling using the fourier neural operator. Journal of Advances in Modeling Earth Systems **15**(7), 2023–003800 (2023)

[29] Sinha, S., Benton, B., Emami, P.: On the effectiveness of neural operators at zero-shot weather downscaling. Environmental Data Science **4**, 21 (2025)

[30] Ding, J.-W., Hsieh, I.-Y.L.: Super-resolution wind mapping with deep learning for scalable renewable energy planning. Communications Earth & Environment **7**, 51 (2025)

[31] Mardani, M., Brenowitz, N., Cohen, Y., Pathak, J., Chen, C.-Y., Liu, C.-C., Vahdat, A., Nabian, M.A., Ge, T., Subramaniam, A., *et al.*: Residual corrective diffusion modeling for km-scale atmospheric downscaling. Communications Earth & Environment **6**(1), 124 (2025)

[32] Tomasi, E., Franch, G., Cristoforetti, M.: Can ai be enabled to perform dynamical downscaling? a latent diffusion model to mimic kilometer-scale cosmo5. 0_clm9 simulations. Geoscientific Model Development **18**(6), 2051–2078 (2025)

[33] Lopez-Gomez, I., Wan, Z.Y., Zepeda-Núñez, L., Schneider, T., Anderson, J., Sha, F.: Dynamical-generative downscaling of climate model ensembles. Proceedings of the National Academy of Sciences **122**(17), 2420288122 (2025)

[34] 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), 220034 (2023)

[35] Ferziger, J.H., Perić, M.: Computational Methods for Fluid Dynamics vol. 586. Springer, Berlin (2002)

[36] Karniadakis, G.E., Israeli, M., Orszag, S.A.: High-order splitting methods for the incompressible navier-stokes equations. Journal of computational physics **97**(2), 414–443 (1991)

[37] Meteodyn: Meteodyn WT – Wind Resource Assessment Software. https://www.meteodyn.com/

[38] WindSim AS: WindSim – CFD software for wind energy. https://www.windsim.com/

[39] Envision Energy: Greenwich – Wind Farm Design Platform. https://www.envision-group.com/

[40] ICOS Research Infrastructure: Atmospheric tower observations. Carbon Portal (2025). https://data.icos-cp.eu

[41] U.S. Geological Survey: Shuttle Radar Topography Mission (SRTM) 1 Arc-Second Global. U.S. Geological Survey (2014)

[42] Zanaga, D., Van De Kerchove, R., De Keersmaecker, W., Souverijns, N., Brockmann, C., Quast, R., Wevers, J., Grosu, A., Paccini, A., Vergnaud, S., Cartus, O., Santoro, M., Fritz, S., Georgieva, I., Lesiv, M., Carter, S., Herold, M., Li, L., Tsendbazar, N.-E., Ramoino, F., Arino, O.: ESA WorldCover 10 m 2020 v100. Zenodo (2021)

[43] Hasager, C.B., Nielsen, N.W., Jensen, N.O., Boegh, E., Christensen, J.H., Dellwik, E., Soegaard, H.: Effective roughness calculated from satellite-derived land cover maps and hedge-information used in a weather forecasting model. Boundary-Layer Meteorology **109**, 227–254 (2003)

[44] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)

Acknowledgements. This work was supported by the AI for Science Program of Shanghai Municipal Commission of Economy and Informatization (2025-GZL-RGZN-BTBX-02017, 2025-GZL-RGZN-BTBX-02031) and the Smart Grid National Science and Technology Major Project (2024ZD0800400). Part of this research was performed using the CFFF platform of Fudan University. The authors thank Prof. Dongxiao Zhao (Shanghai Jiao Tong University) for helpful discussions and insightful suggestions.

Author contributions. C.L. conceived the study, designed the methodology, performed the analyses and visualization, and wrote the original draft of the manuscript. R.T. contributed to AI model training and writing the original draft of the manuscript. S.Y., D.L., X.Z. and Z.H. contributed to validation and manuscript review and editing. H.L. conceived the study, secured funding, supervised the project, provided project administration, and contributed to manuscript review and editing.

Competing interests. The authors declare no competing interests.

Figure Legends/Captions

Fig. 1. Terrain-aware 3D wind reconstruction with FuXi-CFD. FuXi-CFD converts kilometer-scale horizontal winds from global or regional forecasts into 30 m-resolution three-dimensional wind fields over complex terrain. The model learns a terrain-response operator from a large computational fluid dynamics (CFD)-generated library, enabling prediction of full velocity vector and turbulence-related quantities from coarse inputs.

Fig. 2. Agreement between FuXi-CFD predictions and computational fluid dynamics (CFD) references. Parity plots compare predicted and simulated horizontal winds (u, v), vertical velocity (w), and turbulent kinetic energy (k) at 10 m, 50 m, and 100 m height. Predictions of u and v align closely with the CFD truth across all heights, indicating that the model recovers both large-scale inflow directionality and fine-scale terrain-modulated variations. Scatter is largest for k, reflecting its derived nature and dependence on small-scale shear. Overall, the results demonstrate strong predictive skill for primary wind components and reasonable performance for turbulence-related quantities.

Fig. 3. Examples of multi-height wind-field reconstructions across diverse terrain conditions. (a) Distribution of test cases as a function of terrain relief, with selected examples spanning low to extreme terrain relief. (b–e) For each case, coarse 10 m winds (9×9 grid) are shown alongside high-resolution predictions and computational fluid dynamics (CFD) references at 10 m, 50 m, and 100 m. The reconstructed fields capture speed-up over ridges, flow channeling in valleys, and height-dependent smoothing as terrain influence weakens. Across diverse conditions, FuXi-CFD recovers both broad flow patterns and local gradients with good fidelity, despite a downscaling ratio exceeding 30×. Beyond horizontal winds, the model also infers vertical motion (w) and turbulent kinetic energy (k), which are absent from the inputs yet essential for wind-related applications.

Fig. 4. Height-resolved performance and comparison with conventional models. (a) Vertical profiles of prediction errors for horizontal winds (u, v), vertical velocity (w), and turbulent kinetic energy (k), evaluated using mean absolute error (MAE) and relative L^2 metrics. Errors remain low across the 0–200 m layer, with slightly larger deviations near the surface. (b) Wind-speed prediction performance across models as a function of height. Conventional architectures trained on the same dataset exhibit larger errors and lower correlations than FuXi-CFD.

Fig. 5. Validation of FuXi-CFD against tall-tower observations over complex terrain. (a1–a3) Photograph of the European OPE meteorological tower, together with the digital elevation model (DEM) and roughness length distribution over an approximately 9 km × 9 km area centered on the tower. The photograph in (a1) is reproduced from ICOS RI (https://meta.icos-cp.eu/resources/stations/AS_OPE), licensed under CC BY 4.0. (b) Mean absolute error at 120 m, 50 m, and 10 m as a function of wind-speed interval. Baseline comparisons include ERA5-SC (single-constraint extrapolation) and ERA5-DC (dual-constraint extrapolation). ERA5-DC clearly outperforms ERA5-SC, while FuXi-CFD further reduces errors under moderate to strong

wind conditions, with the advantage increasing at higher wind speeds. Error bars represent ±1 standard deviation. (c) Directional distribution of the error difference between FuXi-CFD and ERA5-DC. FuXi-CFD exhibits the largest improvement for southwesterly winds, indicating that the model achieves greater performance gains in scenarios where terrain effects are pronounced, with particularly strong advantages in high-wind, high-value conditions. (d1–d2) Statistical validation across three independent European tall-tower sites (OPE, Torfhaus, and Ispra). Consistent performance improvements are observed across sites, height regimes, and wind-speed ranges. The model is trained on complex terrain data from China and validated at European sites, demonstrating stable cross-regional generalization.

Fig. 6. Computational fluid dynamics (CFD)-informed dataset generation. Each row illustrates a representative terrain block from the dataset, starting from satellite imagery and its corresponding digital elevation model (DEM) and surface roughness length (z_0) fields. For each terrain block and inflow direction, CFD produces a 3D wind field, shown as horizontal slices and two orthogonal vertical sections. These simulations form a comprehensive library of terrain–flow interactions, which FuXi-CFD learns to approximate.

Fig. 7. Terrain-induced variability in computational fluid dynamics (CFD)-generated wind fields. (a–b) Topography and roughness length for an example region. (c) Horizontal winds (u, v), vertical velocity (w), and turbulent kinetic energy (k) at several heights show strong spatial variation driven by terrain forcing. Lower levels (5–10 m) exhibit sharp gradients associated with slope-parallel accelerations and localized separation zones, while flow patterns become smoother with altitude as terrain influence weakens. Turbulence intensity remains elevated over steep slopes and leeward regions, reflecting shear-induced production captured by the CFD solver. These examples highlight the diversity of flow responses that the model must learn to reproduce from coarse inputs.

Fig. 8. Architecture of the terrain-informed deep learning model for 3D wind downscaling. The model takes high-resolution terrain (300×300 grid) and coarse 100 m winds (9×9 grid) as inputs. A downsampling layer and twelve transformer encoder blocks extract cross-scale spatial dependencies, which are then decoded into full 3D fields of horizontal winds (u, v), vertical velocity (w), and turbulent kinetic energy (k) at 30 m horizontal resolution and 27 vertical levels. Training is supervised with a combination of Charbonnier and frequency-domain losses to encourage both local accuracy and structural coherence.

Global Forecast

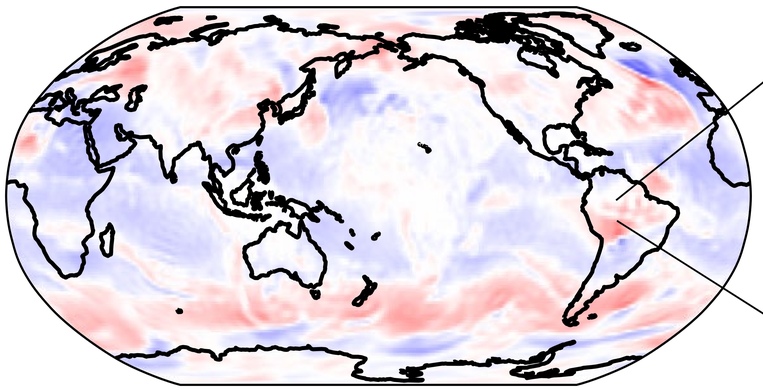

Regional Forecast

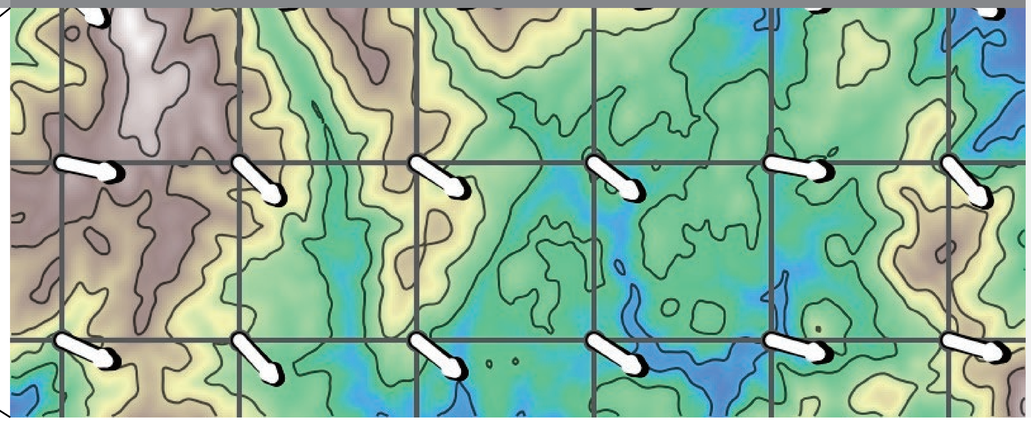

- km-scale resolution
- 100m wind layer
- Horizontal wind (u/v)

Terrain-aware 3D operator learning

FuXi CFD

Regional 3D Wind Field

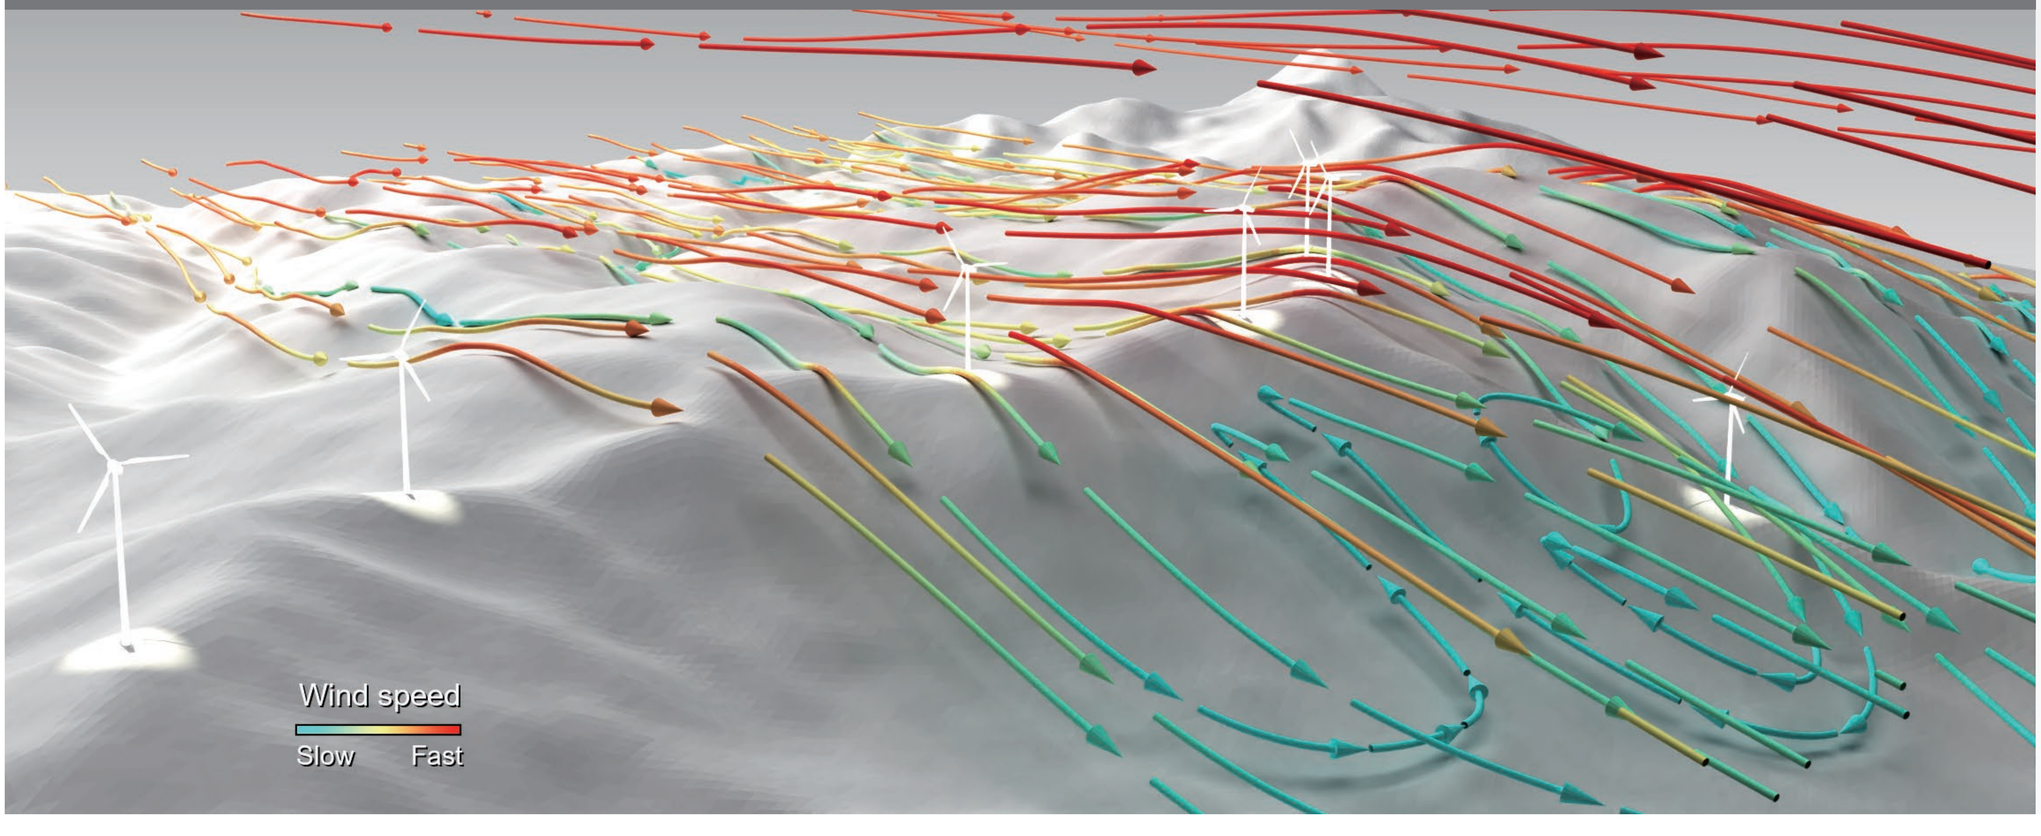

- 30m horizontal resolution
- Full 3D wind (u/v/w)
- Turbulence intensity

Applications

- Micro-siting
- Power forecasting
- Wildfire spread modeling
- Extreme wind assessment
- ...

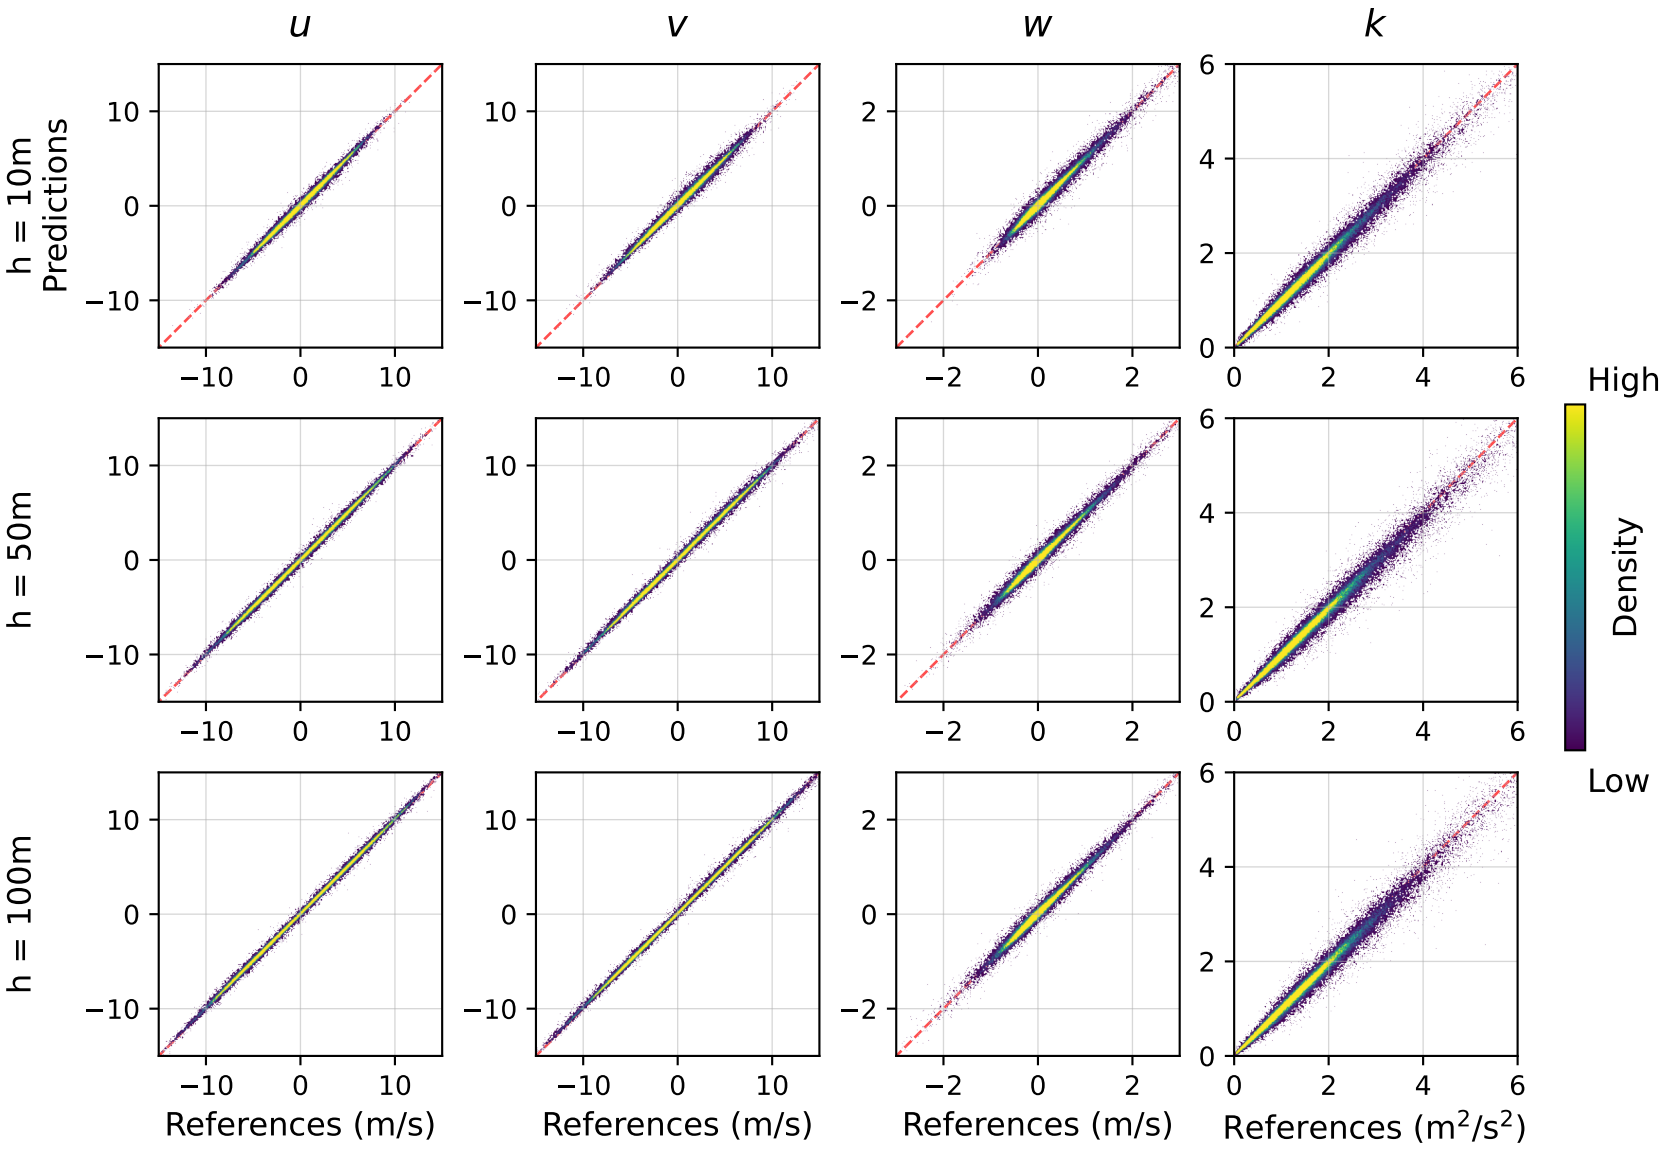
u
v
w
k
h = 10m Predictions
h = 50m
h = 100m
References (m/s)
References (m/s)
References (m/s)
References (m²/s²)
High
Density
Low

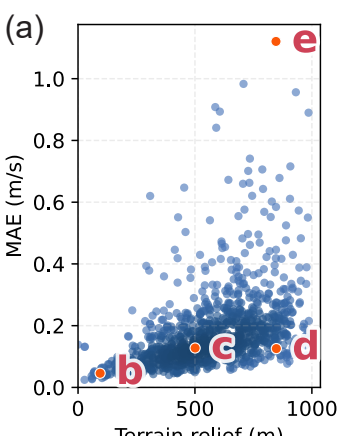

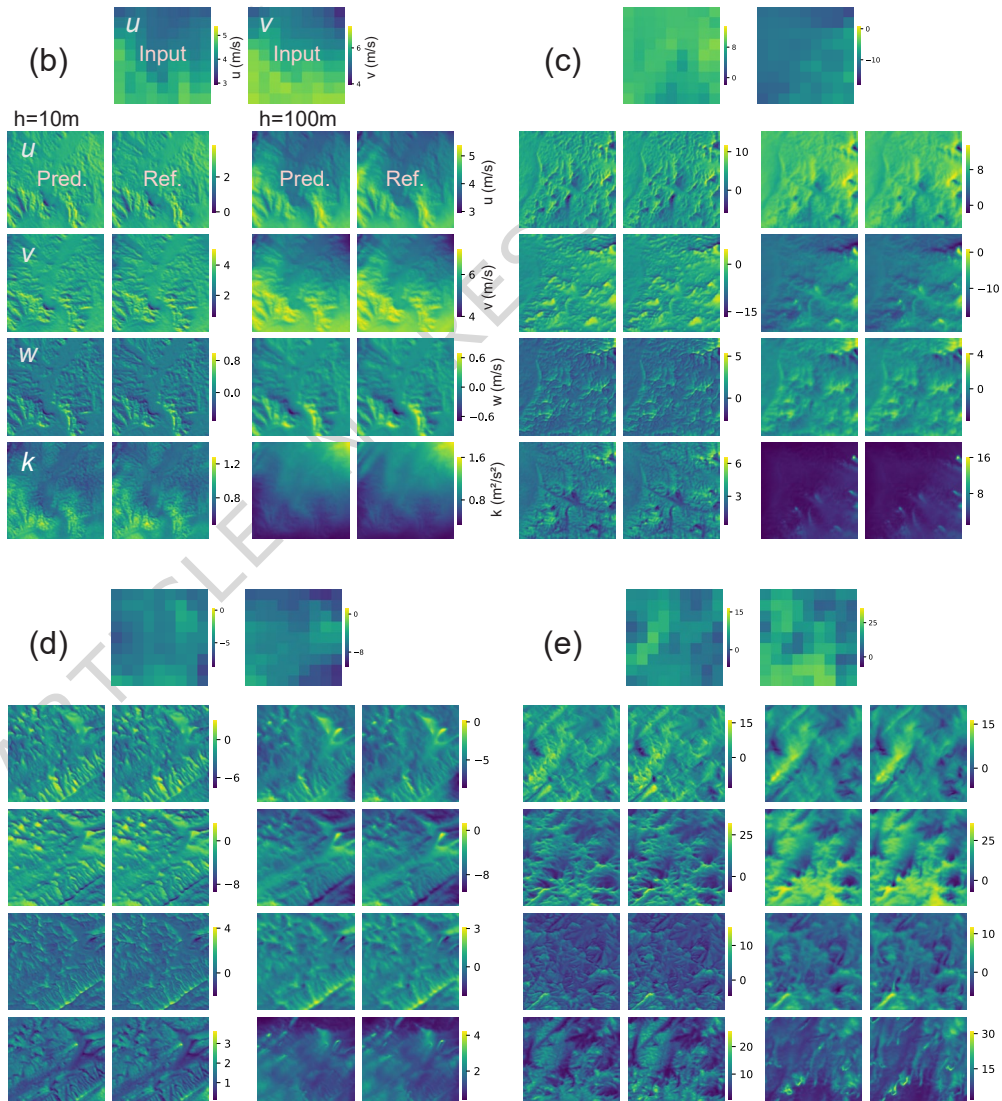

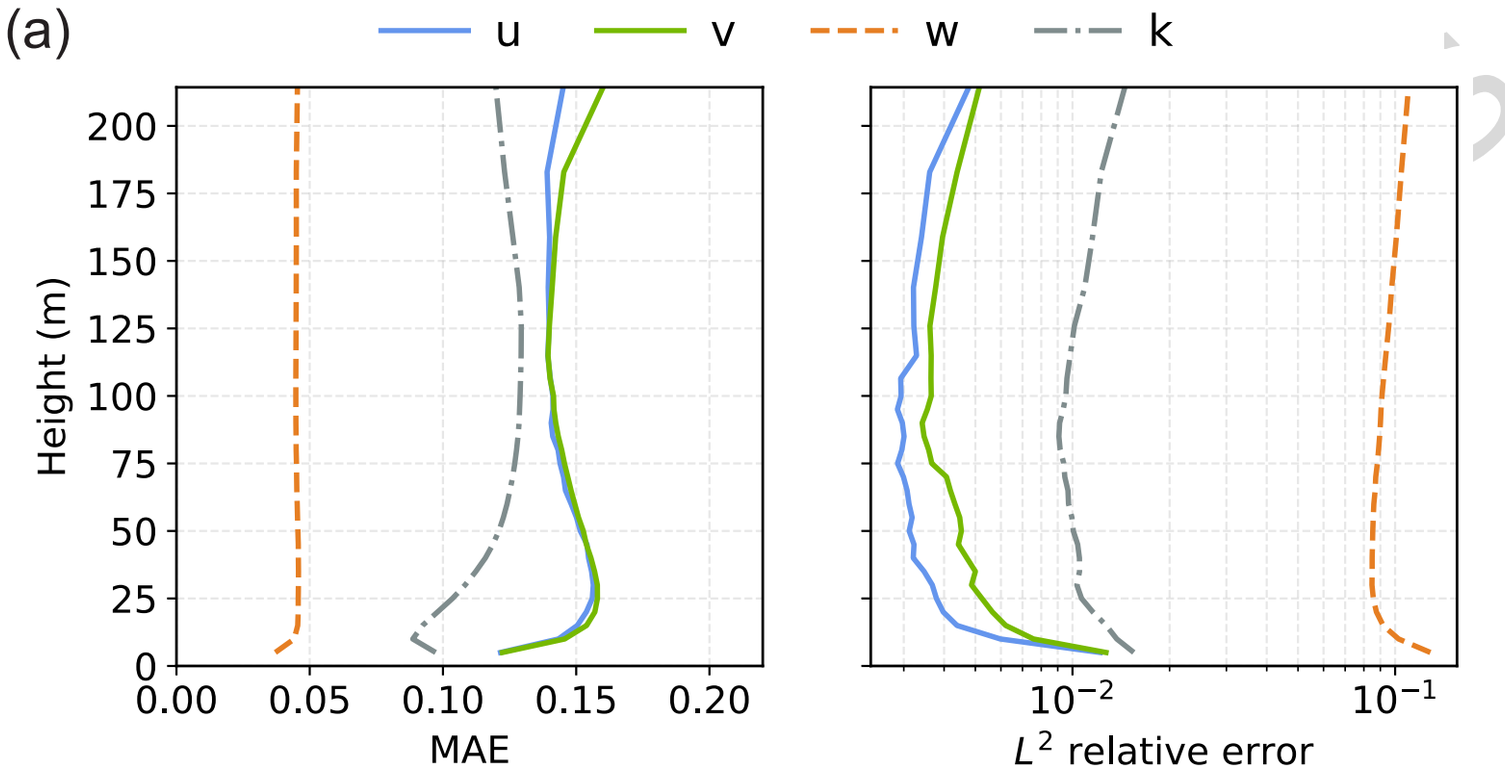

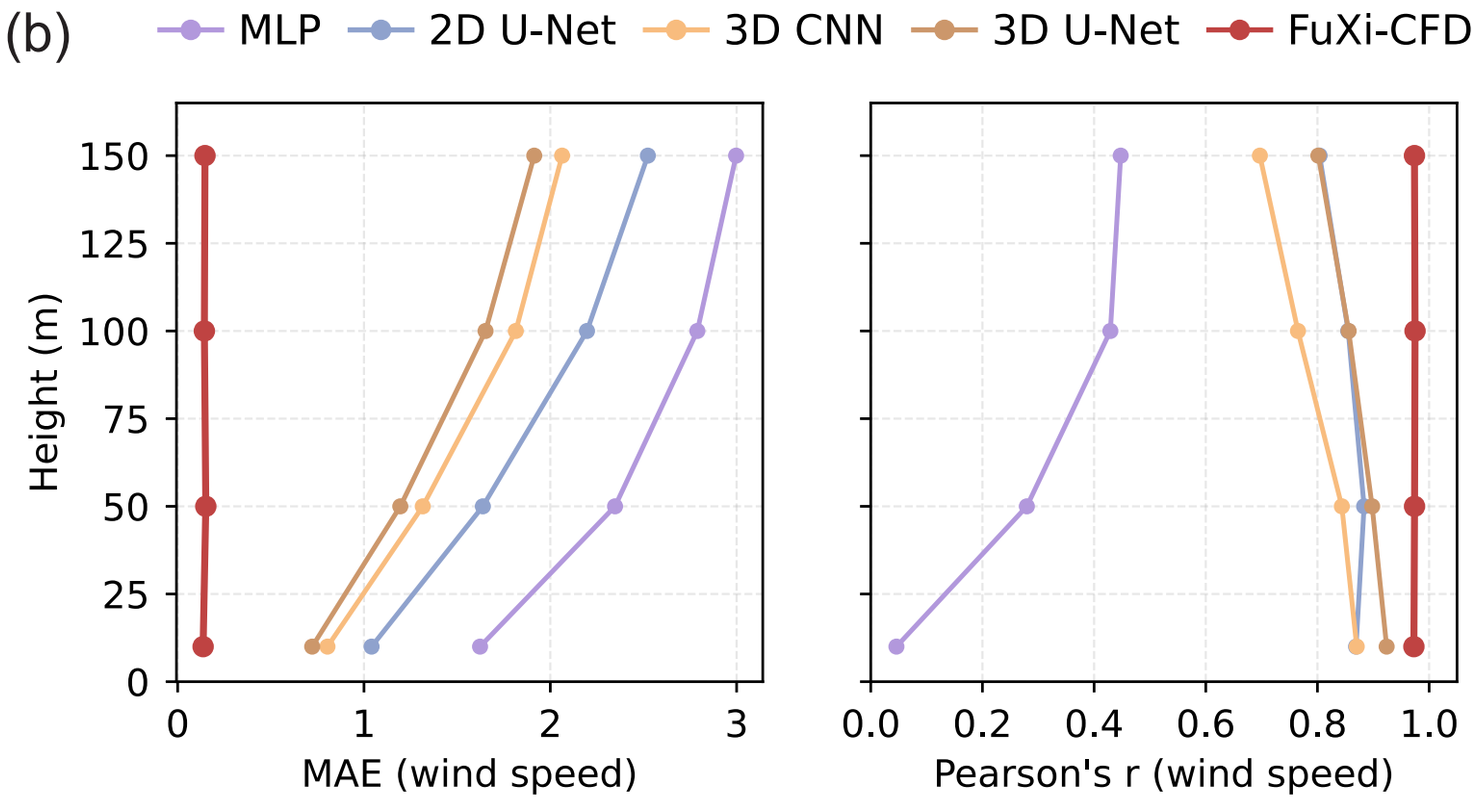


(a1)

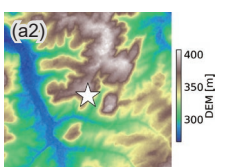
(a2)
400
350 [m]
300
DEM [m]

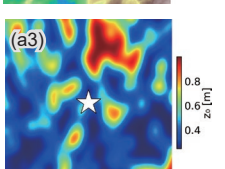
(a3)
0.8
0.6
0.4
z_0 [m]

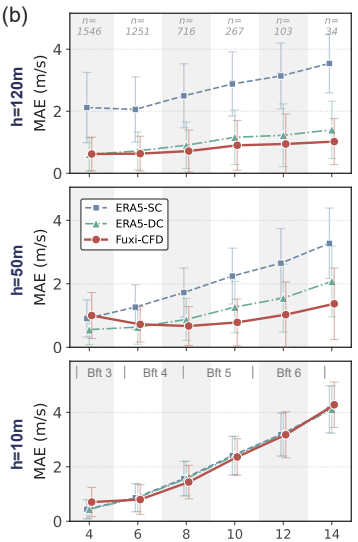
(b)
n= 1546
n= 1251
n= 716
n= 267
n= 103
n= 34
h=120m
h=50m
h=10m
MAE (m/s)
ERA5-SC
ERA5-DC
Fuxi-CFD
Bft 3
Bft 4
Bft 5
Bft 6
Wind speed interval (m/s)

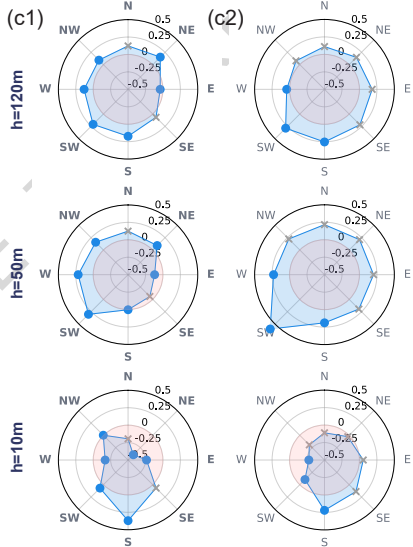
(c1)
(c2)
h=120m
h=50m
h=10m
N
NE
E
SE
S
SW
W
NW
0.5
0.25
0
-0.25
-0.5
Wind speed ≥ Bft 4
Wind speed ≥ Bft 5

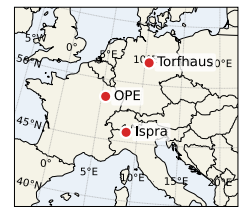
(d1)
Torfhaus
OPE
Ispra

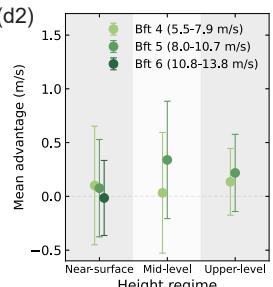
(d2)
Bft 4 (5.5-7.9 m/s)
Bft 5 (8.0-10.7 m/s)
Bft 6 (10.8-13.8 m/s)
Mean advantage (m/s)
Near-surface
Mid-level
Upper-level
Height regime

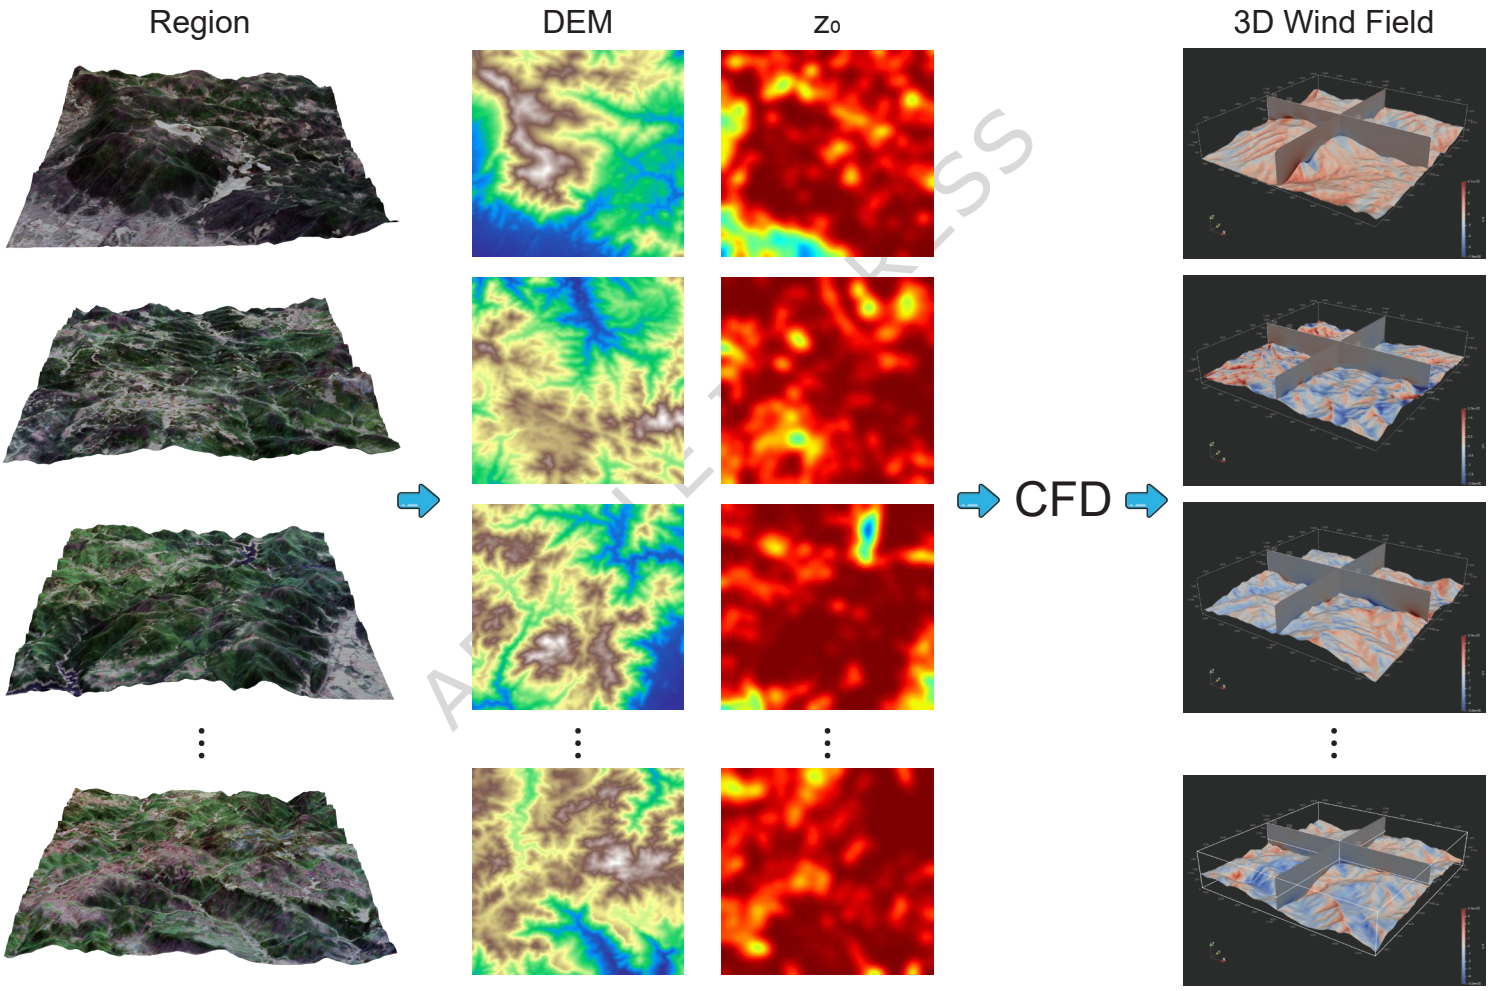

(a)

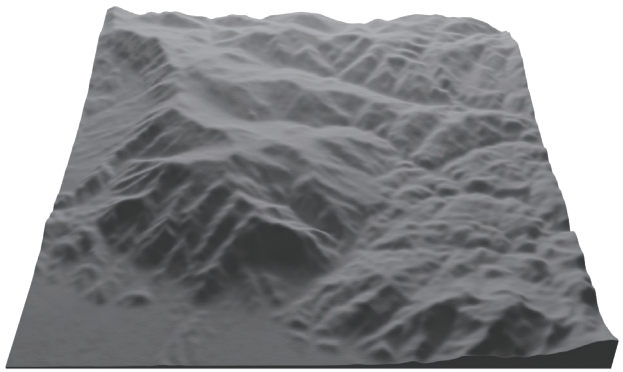

(b)

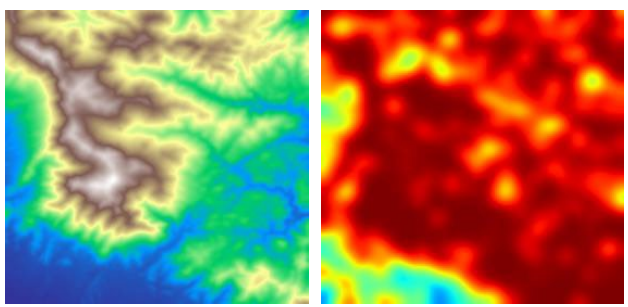

(c)