

Observation-driven correction of numerical weather prediction for marine winds

Matteo Peduto^{1,2}, Qidong Yang¹, Jonathan Giezendanner¹, Devis Tuia², and Sherrie Wang¹

¹Massachusetts Institute of Technology, Cambridge, MA, USA

²École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Key Points:

- We reformulate marine wind forecasting as observation-informed correction of numerical weather prediction using a transformer-based architecture.
- The model reduces forecast errors by 45% at one-hour lead time and 13% at 48 hours across all observation platform types.
- Single pass inference at arbitrary coordinates enables both site-specific and basin scale predictions suitable for operational workflows.

Corresponding author: Matteo Peduto, matteo.peduto@epfl.ch

Corresponding author: Sherrie Wang, sherwang@mit.edu

Abstract

Accurate marine wind forecasts are essential for safe navigation, ship routing, and energy operations, yet they remain challenging because observations over the ocean are sparse, heterogeneous, and temporally variable. We reformulate wind forecasting as *observation-informed correction* of a global numerical weather prediction (NWP) model. Rather than forecasting winds directly, we learn local correction patterns by assimilating the latest in-situ observations to adjust the Global Forecast System (GFS) output. We propose a transformer-based deep learning architecture that (i) handles irregular and time-varying observation sets through masking and set-based attention mechanisms, (ii) conditions predictions on recent observation–forecast pairs via cross-attention, and (iii) employs cyclical time embeddings and coordinate-aware location representations to enable single-pass inference at arbitrary spatial coordinates. We evaluate our model over the Atlantic Ocean using observations from the International Comprehensive Ocean-Atmosphere Data Set (ICOADS) as reference. The model reduces GFS 10-meter wind RMSE at all lead times up to 48 hours, achieving 45% improvement at 1-hour lead time and 13% improvement at 48-hour lead time. Spatial analyses reveal the most persistent improvements along coastlines and shipping routes, where observations are most abundant. The tokenized architecture naturally accommodates heterogeneous observing platforms (ships, buoys, tide gauges, and coastal stations) and produces both site-specific predictions and basin-scale gridded products in a single forward pass. These results demonstrate a practical, low-latency post-processing approach that complements NWP by learning to correct systematic forecast errors.


Plain Language Summary

Accurate wind forecasts over the open ocean are critical for safe maritime navigation, efficient ship routing, offshore energy operations, and storm warnings. However, the ocean has far fewer weather observation stations than land areas, which causes global weather prediction models to miss local conditions. We adopt a practical post-processing approach: rather than replacing an established global forecast model (the Global Forecast System), we correct its predictions using the most recent measurements from ships, buoys, tide gauges, and coastal weather stations. We trained a machine-learning model based on transformer architecture to learn when and how to adjust the forecasted winds based on these fresh observations, their geographic locations, and temporal information including time of day and season. When tested over the Atlantic Ocean, our method reduces forecast errors at all prediction lead times up to 48 hours and performs most reliably in regions where new observations are available. The model also generates predictions rapidly at any location—either for specific sites of interest or across a complete spatial grid—making it well-suited for operational forecasting workflows. This work demonstrates how integrating global weather models with timely observational data and modern artificial intelligence can produce more reliable marine wind forecasts, benefiting maritime operators, offshore energy facilities, and coastal communities.


1 Introduction

Accurate marine wind forecasts are essential for navigation safety, ship routing, offshore energy operations, and hazard early warning, yet they remain challenging because observations over the ocean are sparse, heterogeneous, and unevenly distributed in space and time. Global numerical weather prediction (NWP) systems provide coherent large-scale dynamics but contain biases at many locations (Figure 1).

In this work, we reformulate marine wind forecasting as a *correction* problem: rather than predicting winds from scratch, we learn local patterns using the in-situ observations to adjust NWP such as NOAA’s Global Forecast System (GFS). The central idea is to preserve the large-scale flow provided by NWP while applying localized, observation-informed



(a) Wind vectors over the North Atlantic. Grey arrows indicate the GFS forecast field, black arrows show in-situ observations, and green arrows represent GFS predictions corresponding to observation locations. The inset map outlines the analysis domain.



(b) Comparison between observed and predicted wind components. Green symbols denote GFS forecasts and blue symbols our model outputs. Circles correspond to u and squares to v ; the 1:1 line marks perfect agreement.

Figure 1: Comparison of spatial and statistical wind patterns from GFS forecasts, in-situ observations, and model predictions.

corrections wherever and whenever observations indicate discrepancies. We design a transformer-based neural network architecture that (i) operates on irregular, time-varying sets of observations via masking and attention; (ii) conditions predictions on recent observation-forecast pairs using cross-attention; and (iii) encodes time and location with cyclical and coordinate-aware location representations so that inference can be queried at arbitrary coordinates without interpolation heuristics.

We evaluate our approach over the Atlantic Ocean using observations from the International Comprehensive Ocean-Atmosphere Data Set (ICOADS) as a reference (Section 3.1), with GFS and the European Center for Medium-Range Weather Forecasts Reanalysis (ERA5) as baseline models. Spatial analyses show most consistent improvements along coasts and shipping routes, where observations are more dense. The architecture naturally ingests heterogeneous platforms (ships, buoys, tide gauges, coastal stations) and produces single-pass predictions at arbitrary locations, enabling both gridded and site-specific products suitable for operational workflows.

Our contributions are as follows.

1. **Formulation.** We forecast marine wind by assimilating recent observations at variable locations on the ocean into global numerical model predictions. We predict state-dependent NWP forecast errors to generate corrected forecasts at any arbitrary target locations (Section 4).
2. **Demonstrated performance across the Atlantic.** Applied to 10-m wind forecasting over the entire Atlantic Ocean, the model achieves a **45% reduction in GFS wind-speed error at 1-hour lead time** and **13% at 48 hours**, systematically improving predictions across all observation platforms—including ships, buoys, tide gauges, and coastal stations. These results demonstrate that localized, observation-informed corrections can significantly enhance global model accuracy in operational maritime contexts.
3. **Scalable arbitrary-location inference.** The architecture enables single-pass predictions at any coordinates—from individual platforms to basin-scale grids—without fixed neighborhoods or interpolation. It handles variable observations at each time step through masking and cyclical time and coordinate-aware location representations, providing a flexible framework for real-time, observation-informed inference (Sections 5 and 6).

2 Related Work

Modern weather forecasting couples dynamical models with statistical post-processing to correct systematic errors. Since the 1970s, Model Output Statistics (MOS) and its probabilistic successors have provided principled ways to calibrate NWP outputs against observations (Glahn & Lowry, 1972; Gneiting et al., 2005; Raftery et al., 2005; Wilks, 2011). These approaches address mean and distributional biases but are limited in representing nonlinear, state-dependent, and spatially varying errors. Deep post-processing methods have therefore gained traction: neural networks trained on ensemble and auxiliary predictors improve calibration and sharpness over classical regression (Rasp & Lerch, 2018; Scheuerer & Hamill, 2020), and recent applications show ML-based bias correction enhancing GFS precipitation and near-surface temperature forecasts in real-world settings (Sun et al., 2023; NOAA/NWS/NCEP, 2025). Our work adopts this correction paradigm but targets marine surface winds under sparse, shifting observation geometry, using self- and cross-attention to condition corrections on the most informative recent observations.

End-to-end deep learning has rapidly advanced from nowcasting to global medium-range forecasting with near-instant inference. Data-driven global models such as FourCastNet, GraphCast, and Pangu-Weather demonstrate medium-range skill compet-

itive with leading operational systems while achieving orders-of-magnitude speedups at $\sim 0.25^\circ$ resolution (Pathak et al., 2022; Lam et al., 2023; Bi et al., 2023). For short-range, high-resolution targets, the MetNet family delivers kilometer-scale, minute-level updates for multiple surface variables, including wind (Pathak et al., 2023; Google Research & DeepMind, 2023). These developments make large ensembles and on-demand forecasts at arbitrary coordinates computationally feasible. While our focus is not to replace NWP globally, we leverage these architectural insights (efficient attention, operator-style layers) in a model that can infer at arbitrary offshore locations with low latency.

Marine wind forecasting is uniquely constrained by sparse, heterogeneous observations and marine-specific model errors. Compared to land, the ocean is observationally sparse and relies on a mix of buoys, ships, coastal stations, and satellite surface winds (Atlas et al., 2011; Remote Sensing Systems, 2024). Data gaps degrade initialization and can leave persistent marine wind biases in dynamical models; targeted assimilation has shown measurable improvements, including for near-surface winds over the tropics and high latitudes (Rennie et al., 2021; ECMWF, 2023; Zuo & Hasager, 2023). Physics-guided and hybrid ML systems that fuse NWP with local sensors improve site-specific offshore wind forecasts across heights and lead times (Ye et al., 2023, 2024). Concurrently, deep spatiotemporal models (LSTM/GNN/attention) trained on buoy networks and coastal sensors reduce error relative to raw NWP and classical baselines (Han et al., 2023; Liu et al., 2023; Dong et al., 2024). Our framework fits this line of work by learning observation-informed corrections over the Atlantic domain while remaining agnostic to which platforms are present at any given time.

Learning from irregular, time-varying sets of observations benefits from attention architectures and geospatial encodings. Set-aware transformers and Perceiver-style cross-attention provide permutation-invariant processing of variable-length inputs and flexible conditioning from one set (recent observations) onto another (targets) (Lee et al., 2019; Jaegle et al., 2021). To query forecasts at arbitrary points without gridding artifacts, geospatial encoders that respect spherical geometry yield globally consistent coordinate features for transformers (Rußwurm et al., 2023; Tancik et al., 2020; Sitzmann et al., 2020). Our model follows this recipe: we encode time cyclically, locations with a spherical-harmonic pipeline, and use self-attention to summarize recent observation/forecast pairs before cross-attending them from any set of target points.

Arbitrary-location, low-latency inference is increasingly practical and well aligned with marine applications. Machine learning weather models routinely produce global fields in under a minute and support dense spatiotemporal sampling (Lam et al., 2023; Bi et al., 2023; Pathak et al., 2022). By decoupling targets from a fixed grid and using learned location encoders, one can deliver site-specific, observation-informed corrections anywhere on the ocean. Our work explores this idea for 10-m winds over the Atlantic: a model that scales to arbitrary target sets, integrates the observations, and demonstrably improves GFS across lead times while preserving global context.

3 Data

Our study combines global reanalysis, operational forecasts, and in situ marine observations to evaluate and train the proposed models. Our analysis focuses on the North Atlantic domain (2° – 64° N, 98° W– 11.5° E) over the period April 2015 to September 2024. We used three complementary datasets: **the International Comprehensive Ocean Atmosphere Data Set (ICOADS)** as the observational reference, **the ECMWF Reanalysis v5 (ERA5)** as a physically consistent benchmark, and **the Global Forecast System (GFS)** as the operational forecast baseline.

3.1 International Comprehensive Ocean–Atmosphere Data Set (ICOADS)

The International Comprehensive Ocean–Atmosphere Data Set (ICOADS) provides the most extensive collection of global marine meteorological and ocean surface observations. It aggregates reports from ships, buoys, coastal and Coastal-Marine Automated Network (C-MAN) stations, and other near-surface platforms, harmonized through quality flags and hourly timestamp rounding. After filtering, the dataset includes 6,625 unique platforms with an average of approximately 70 measurements per platform. Ships, moored buoys, and tide gauges record the majority of observations in ICOADS, while drifting buoys contribute by far the least number of observations (Table 1). Figure 2 illustrates the spatial distribution of observations, showing dense coastal coverage and sparser but widespread ship-based sampling across the open ocean.

Table 1: Number of ICOADS observations by platform type (2015–2024).

Platform type	Observations
Ship	7 929 745
Drifting buoy	179 923
Moored buoy	9 677 483
C-MAN station	1 518 359
Coastal station	3 011 731
Tide gauge	8 937 940
Total	34 860 848




Figure 2: Spatial density of ICOADS observations by platform type (2015–2024). Higher densities occur near coastlines, while ship observations dominate the open ocean.

3.2 ECMWF Reanalysis v5 (ERA5)

ERA5 is the fifth-generation global climate reanalysis produced by the European Centre for Medium-Range Weather Forecasts (ECMWF) under the Copernicus Climate Change Service (C3S) (Hersbach et al., 2020; European Centre for Medium-Range Weather Forecasts (ECMWF), 2024; Copernicus Climate Change Service (C3S), 2017). It provides a consistent record of climate variables from January 1940 to the present. ERA5 offers hourly estimates of a wide range of atmospheric, land, and oceanic variables. The dataset covers the entire globe at a spatial resolution of $0.25^\circ \times 0.25^\circ$ and includes 137 vertical atmospheric levels, extending up to about 80 km altitude. In this work, we use

ERA5 as a benchmark against which we compare our machine learning models. Since our focus is on predicting wind velocity and direction, we extracted the horizontal wind components (U and V) at 10 m height. To ensure consistency with in-situ records, ERA5 values were sampled at the exact time and location of each ICOADS observation. In each case, the closest available ERA5 time step is used as the reference value. Tests showed that interpolation among the four nearest ERA5 grid points did not improve ERA5 accuracy relative to ICOADS observations, so we use the nearest-neighbor ERA5 value in this paper.

3.2.1 Global Forecast System (GFS)

The Global Forecast System (GFS) is a global numerical weather prediction model developed by the National Centers for Environmental Prediction (NCEP) (National Centers for Environmental Prediction (NCEP), 2020). It provides gridded forecasts of atmospheric, oceanic, and land-surface variables. Forecasts extend up to 384 hours (16 days) ahead, with hourly lead times available for the first 120 hours and 3-hourly lead times thereafter. New forecast cycles are initialized four times per day (00h, 06h, 12h, 18h UTC) at a base horizontal resolution of 1° . GFS data are accessed through Google Earth Engine (GEE). While convenient, this interface restricts the set of available variables. We downloaded the U and V components of wind velocity at 10 m, air temperature at 2 m, specific humidity at 2 m, and relative humidity at 2 m.

GFS data are sampled at the closest available grid point in space and time for each ICOADS observation. However, unlike ERA5, which provides continuous hourly coverage, GFS is issued only four times daily (every 6 hours), each run producing forecasts with lead times up to 384 hours. To ensure a prediction horizon of 48 hours with the best available prediction from GFS, the 48 forecast values are distributed across eight consecutive cycles. As GFS cycles occur every 6 hours, we make sure to use the appropriate GFS value coming from the right cycle based on the lead time value and the observation measurement time. For an observation at 03 UTC, the 00 UTC GFS cycle provides the 1-, 2-, and 3-hour lead-time forecasts. We obtain the lead times from 4 to 9 hours from the preceding day’s 18 UTC GFS cycle using the appropriate forecast hours. The process is illustrated in Figure A2 in the appendix.

4 Methods and Model

Predicting the future state of the atmosphere is inherently difficult and has been the focus of sustained research across both physics-based numerical weather prediction (NWP) and modern AI weather models. Instead of attempting to generate forecasts directly from observations, Yang et al. (2024) reframed the problem by demonstrating that machine learning can systematically correct errors in existing NWP guidance. This correction-based perspective is attractive because it leverages the strengths of NWP while allowing ML models to focus on residual, data-driven improvements. Our work adopts this same approach.

The original forecasting task, using only past observed values, can be formulated as:

$$\mathbf{w}(t + l\Delta t) = F(\mathbf{w}(t - b\Delta t : t)) \quad (1)$$

where t denotes the current time, $\mathbf{w}$ represents a vector of weather observations, $l\Delta t$ is the lead time (the time into the future for the prediction), $b\Delta t$ is the historical time window used as input, and F is the forecasting function.

The alternative correction-based task is written as:

$$\mathbf{w}(t + l\Delta t) = F(\mathbf{w}(t - b\Delta t : t)) + \mathbf{g}(\mathbf{w}(t - b\Delta t : t + l\Delta t)) \quad (2)$$

where $\mathbf{g}$ represents an external weather forecast, such as NWP.

4.1 General Framework

A major deviation from previous work (Yang et al., 2024) is that, in this work, we do not assume that past ocean observations remain at fixed locations across time steps. Instead, both the number and spatial distribution of available observations vary dynamically. The model must therefore predict at arbitrary positions, with a changing set of reference points between consecutive time steps. A central difficulty is the sparsity of observations over the Atlantic, which raises the question of which past points should be retained for prediction. To assess whether spatial or temporal proximity is more informative, we computed the Pearson correlation coefficients between target points and past observations under varying distance and time thresholds (Fig. A1). The results indicate that correlation is dominated by temporal proximity: the most recent measurements, even if geographically distant, are more predictive than older nearby observations. Consequently, the framework adopts a simpler approach: the most recent observations are used to predict the following time step, regardless of spatial distance. This design enables inference at arbitrary sets of positions for a specific time step across the ocean in a single pass, while maintaining predictive skill.

4.2 Model Implementation

Our objective is to design a correction model for the Global Forecasting System (GFS). The model learns from pairs of past observations and historical forecast values coming from GFS at specific locations. The model must learn how to correct past GFS predictions based on the historical observations. It must also capture which pair is more relevant to correct the GFS forecast for a given target point (Figure 3).

We use Vision Transformers (ViTs) due to their state-of-the-art performance and ability to encode positional information. Following Kazemnejad et al. (2019), positional embeddings $\mathbf{e}$ are added to the input embeddings $\mathbf{X}$, where each $\mathbf{x}_i$ represents the feature vector of the i -th pair element:

$$\mathbf{X} = [x_1 + e_1, x_2 + e_2, \dots, x_n + e_n], \quad (3)$$

This architecture is adopted as the basis of the model, with task-specific modifications.

4.2.1 Self-attention among historical observation-NWP forecast pairs

We process {past observation, GFS forecast} pairs through a self-attention mechanism (Figure 3), enabling the model to assign varying importance to each pair in the correction pattern. Let $\mathbf{X}$ be an input sequence of n historical pairs $[x_1, x_2, \dots, x_n] \in \mathbb{R}^{n \times d}$, where each x_i concatenates the observed state and the corresponding GFS forecast at a given time. Self-attention (Raschka, 2023) produces a new representation as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right) \mathbf{V}, \quad (4)$$

where $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{n \times d_k}$ are the query, key, and value matrices obtained from linear projections of $\mathbf{X}$, and d_k is the dimension of the projected features. The dot product between queries and keys quantifies similarity between the pairs.

In practice, Multi-Head Self-Attention (MHSA) (Vaswani et al., 2017) is applied to this sequence to capture interactions between all pairs across multiple representation subspaces:

$$\text{MHSA}(\mathbf{X}) = [\text{head}_1 \parallel \dots \parallel \text{head}_h] \mathbf{W}^O, \quad (5)$$

where each $\text{head}_i = \text{Attention}(\mathbf{X}\mathbf{W}_i^Q, \mathbf{X}\mathbf{W}_i^K, \mathbf{X}\mathbf{W}_i^V)$ and $\parallel$ denotes concatenation.




Figure 3: Schematic representation of the spatio-temporal learning framework. Each sample consists of in-situ observations at time t (blue circles), corresponding GFS forecasts at t (red squares) and $t + \Delta t$ (magenta squares), and target observations at $t + \Delta t$ (green crosses). The lower panels illustrate how each input is encoded through cyclical time, geographic positional, and sequential embeddings. Past observations and past GFS forecasts are processed through self-attention to capture temporal and spatial dependencies and learn the correction pattern, while cross-attention integrates information from GFS forecasts to predict target wind components at $t + \Delta t$.

In this architecture, self-attention allows the model to learn which historical observation-GFS forecast pairs are most relevant for correcting the GFS forecast at the target points, enabling flexible weighting of past data in generating the corrected prediction.

4.2.2 Cross-attention between past observations and target predictions

The second key component of the model is cross-attention, which enables the network to identify which past observations are most relevant for each target prediction. Whereas self-attention allows tokens within the same sequence to attend to one another, cross-attention relates tokens from two different sequences (Lin et al., 2021).

The attention module receives two inputs:

- a **query** sequence $\mathbf{Q} \in \mathbb{R}^{n_q \times d_k}$ (e.g., target tokens),
- and a **key-value** sequence $\mathbf{K} \in \mathbb{R}^{n_k \times d_k}$, $\mathbf{V} \in \mathbb{R}^{n_k \times d_v}$ (e.g., past observation-forecast pairs).

The cross-attention operation is formally defined as:

$$\text{CrossAttention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right) \mathbf{V}, \quad (6)$$

where n_q and n_k denote the number of query and key/value tokens, respectively, d_k the dimension of queries and keys, and d_v the dimension of the values. This formulation allows each query token to attend selectively to all tokens in the key-value sequence, thereby extracting the most relevant contextual information.

As with self-attention, cross-attention is implemented using multiple heads:

$$MHA_{\text{cross}}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = [\text{head}_1 \parallel \dots \parallel \text{head}_h] \mathbf{W}^O, \quad (7)$$

where each $\text{head}_i = \text{Attention}(\mathbf{QW}_i^Q, \mathbf{KW}_i^K, \mathbf{VW}_i^V)$ and $\parallel$ denotes concatenation.

Cross-attention layers are stacked after self-attention within the decoder blocks. This mechanism is particularly important in our setting, where the available observations vary over time and space, and the model must dynamically determine which inputs are most informative for each prediction.

4.2.3 Masking for Irregular Observations

A distinctive feature of our setting is that both the number and spatial distribution of observations vary at every time step. To ensure the transformer operates consistently under these conditions, we construct binary masks that specify which tokens are valid inputs and which target points should contribute to the loss. These masks are applied during both self-attention and cross-attention layers, preventing the model from attending to nonexistent observations, and during loss computation, ensuring that errors are only accumulated over available targets. This design allows the architecture to gracefully handle changing observation sets dynamically without requiring imputation or fixed spatial grids.

4.2.4 Position and Time Embedding

To fully exploit the capabilities of the ViT architecture, we add both temporal and positional embeddings to the feature representations before entering the transformer layers. We encode each timestamp using four cyclical features capturing daily and annual periodicities: $\sin(\frac{2\pi d}{366})$, $\cos(\frac{2\pi d}{366})$, $\sin(\frac{2\pi h}{24})$, and $\cos(\frac{2\pi h}{24})$ where d is the day of the year and h the hour of the day. This encoding preserves the cyclic nature of time and provides a compact representation for downstream learning, and has been widely adopted in transformer-based forecasting models (Pospíchal et al., 2022; Su et al., 2025).

We adopt the geographic positional encoding scheme of (Rußwurm et al., 2023). This approach combines spherical harmonic basis functions with a Sinusoidal Representation Network (SirenNet). Concretely, each latitude-longitude coordinate is first projected into spherical harmonic coefficients (up to a chosen degree), which are then processed by a neural network to produce a learned embedding vector. This hybrid encoding captures both global spherical geometry (avoiding distortions at high latitudes or across the date line) and fine-grained local structure, making it well-suited for geospatial modeling in our project.

5 Experimental Setup

5.1 Forecasting Task and Baselines

The forecasting task is to predict 10-m wind velocity at arbitrary offshore locations up to 48 h, using as input the most recent set of observations combined with the GFS forecast field. Predictions are evaluated against ICOADS observations (Section 3.1). We use two numerical models as baselines: (i) the Global Forecast System (GFS), which provides the starting point for correction, and (ii) ERA5 reanalysis, used as a reference global dataset. We report comparisons across multiple lead times and stratify by platform type and spatial region in our analysis.

5.2 Data Splitting

To ensure generalization to unseen conditions, we split the dataset chronologically: the first 80% of data temporally are used for training (27.8M samples), the next 10% for validation (3.6M), and the final 10% for testing (3.4M) (Fig. A3). This temporal split prevents data leakage and ensures the model is evaluated on future time periods not seen during training. Within each temporal split, the available observation locations vary dynamically at each time step based on which platform reported measurements at that time. Consequently, the target prediction points—which, recall, are all the observations at some future lead time—also change over time steps. This design ensures the model learns to predict wind corrections at arbitrary ocean locations, promoting spatial generalization and enabling robust performance at previously unseen geographic positions during validation and testing.

5.3 Model Training

We train the models to minimize the mean vector magnitude error between predicted and observed wind velocity components (u , v) concurrently. We monitor performance using RMSE (m s^{-1}). We perform optimization using AdamW with weight decay, an initial learning rate of 1×10^{-4} , and cosine learning-rate scheduling with warm restarts. We trained our models for 100 epochs with early stopping based on validation RMSE and a patience of 25 epochs. Our experiments are run on NVIDIA Tesla V100 GPUs (16 GB). Training a single model requires approximately 24 hours. To ensure reproducibility, all runs use fixed random seeds for data shuffling, parameter initialization, and dropout.

5.4 Hyperparameters and Architecture

The transformer encoder consists of 8 layers with 8 attention heads each, a hidden dimension of 128. Temporal embeddings capture both hourly and annual cycles, while spatial embeddings use spherical harmonics combined with a sinusoidal representation network (Section 4.2.4). Cross-attention layers condition target points on the encoded set of observation-forecast pairs, enabling inference at arbitrary spatial coordinates.


6 Results

We present our results in three parts. First, we evaluate the global performance of our model, understanding its performance throughout the different platform types and in comparison with the global numerical forecasting models. Second, we conduct an extensive spatial analysis to identify potential regional patterns and areas where the model performs well or struggles. Finally, we perform ocean-wide inference, leveraging the model’s capability to generate predictions at arbitrary locations, including at grid points, across the entire ocean domain. We compare the gridded model outputs against the full GFS field. We focus on specific oceanic case studies to examine in detail how the wind field is modified in localized regions, and how the spatial density of nearby observations influences these corrections.

6.1 Global Performance

Our proposed model consistently improves upon GFS forecasts for all lead times up to 48 h. Figure 4a compares the performance against GFS and ERA5 across lead times, while Table 2 quantifies the improvement in terms of RMSE. The model reduces GFS error by 45% at 1 h lead time and by 13% at 48 h. It also outperforms ERA5 up to 12 h.

We tested several additional experiments through architectural variations to assess potential improvements beyond the baseline configuration. These changes are sum-



(a) Mean absolute error of wind predictions as a function of forecast lead time. Our ML model (orange) is compared against the GFS forecasts (blue) and the ERA5 reanalysis (grey).

(b) Mean absolute error difference between the ML model and GFS forecasts, grouped by observation platform type. Negative values indicate an error reduction relative to GFS.

Figure 4: Comparison of model performance across forecast lead times and observation platforms. (a) Overall accuracy relative to GFS and ERA5. (b) Error reduction by platform type.

marized in Table A1. None of the modifications, including the use of longer temporal context, additional GFS predictors, platform metadata, or residual GFS connections, yielded measurable gains.

To further evaluate model behavior, we analyze performance by platform type and by individual observation object. This assessment aims to identify whether certain observation platforms pose greater forecasting challenges or whether specific stations consistently yield higher prediction errors. Figure 4b presents the mean error reduction relative to GFS as a function of lead time and platform category. On average, all platform types exhibit improved skill compared to GFS, though the magnitude of improvement varies. The largest gains are observed for coastal stations, followed by tide gauges, CMAN stations, fixed platforms, moored buoys, and ships, while drifting buoys show the smallest improvement.

6.2 Spatial Performance

To assess how corrections vary across regions and lead times, Figure 5 compares spatial distributions of average wind-speed errors for the proposed model and the GFS baseline at 1, 8, 24, and 48 hour lead times. The left and middle columns show absolute errors for the machine-learning corrector and GFS, respectively, while the right column illustrates their difference (model – GFS), where blue indicates improvement and red degradation. Across all lead times, the model produces lower errors over most of the Atlantic domain, with widespread blue shading in the difference panels. The strongest improvements are found along the eastern seaboard of North America, the Caribbean, and western European coasts, as well as along major shipping corridors. In these regions, local biases in GFS are systematically reduced by the observation-informed corrections. As forecast horizon increases, both models exhibit higher errors. Nevertheless, the correc-

Table 2: Root-mean-square error (RMSE) of 10-m wind speed for several lead times, and relative improvement of the ML model with respect to GFS.

Lead time (h)	Our Model (m s^{-1})	GFS (m s^{-1})	ERA5 (m s^{-1})	Improv. ^a (%)
1	1.51	2.75	2.46	45.1
2	1.82	2.79	2.46	34.6
4	2.09	2.87	2.46	27.3
8	2.33	2.95	2.46	20.9
12	2.43	3.00	2.46	19.1
18	2.54	3.07	2.46	17.5
24	2.62	3.15	2.46	16.8
36	2.80	3.29	2.46	15.1
48	2.99	3.44	2.46	13.1

^a Improvement = $(\text{GFS} - \text{ML})/\text{GFS} \times 100\%$.

tor maintains clear gains up to 48 hours, with the difference maps remaining predominantly blue across the basin. The relative advantage gradually diminishes in magnitude and spatial extent, especially over the central and northern Atlantic, where observational coverage is sparse. Small, isolated regions of degradation appear at high latitudes. In general, our model performs better at locations with more previous observations available nearby (Figure A4). However, we also observe at low latitudes, where data is sparse, that our model manages to improve consistently over GFS.

Model errors inherit the spatial structure of GFS errors, as GFS serves as the baseline for the machine-learning corrector. Error varies geographically across the basin, with higher values at high latitudes and lower values at low latitudes. Mid- and high latitudes typically have more energetic and variable winds that are harder to predict accurately (Zhang et al., 2024). Notably, the model’s performance relative to GFS is spatially heterogeneous. In regions where GFS exhibits systematic biases, the corrector effectively reduces errors. However, for some areas—particularly at higher lead times—the model shows limited improvement or occasional degradation.

6.3 Global Inference

To evaluate the model predictions across the entire Atlantic, we performed global inference on a regular grid of points (0.25° resolution) at multiple time steps. The results reveal clear differences between the corrected fields and the original GFS forecasts (Figure 6). In particular, the ML model introduces local adjustments in the vicinity of past observations, where it modifies the predicted wind velocities for subsequent time steps. This results in fields that are less spatially smooth than GFS, reflecting the localized influence of recent measurements. This is most obviously visualized in the difference field between our forecast and GFS (Figure 6, right column). In some cases, the model also generates isolated adjustments in regions without direct observational input.

We highlight several case studies to illustrate how the model adjusts the wind field over the ocean, using past observations to better align with subsequent measurements (Figure 7). In Figures 7(c) and 7(f), the corrected field northwest of Ireland is adjusted toward a pattern that more closely matches the observed values. Similarly, in Figures 7(a) and 7(d), a distinct circular region of adjusted wind speed is visible around a prior observation northwest of Cuba, marked by the red circle, where the model refines both wind direction and intensity relative to the GFS baseline. In this case, the GFS prediction was




Figure 5: Spatial distribution of wind-speed prediction errors for the machine-learning model and the GFS baseline at lead times of 1 h, 8h, 24 h, and 48 h. Panels (a), (d), (g), and (j) show absolute errors for the model, while panels (b), (e), (h), and (k) show corresponding errors for GFS. Panels (c), (f), (i), and (l) display the difference between the two (model – GFS), where blue indicates lower errors for the model, and red indicates higher errors. Each panel corresponds to the corrected field at the specified lead time.

inaccurate, and the model successfully corrected it. In the second circle, north of the Yucatán, the entire field is adapted based on a previous observation, with changes extending over a large region. The accuracy of these adjustments is confirmed by a subsequent observation. In contrast, Figures 7(b) and 7(c) illustrate a mid-ocean region with sparse prior observations. Here, the model applies only minor adjustments to the GFS baseline due to limited information for guidance. Some changes are still visible near the previous observation. Interestingly, the adapted field captures the future measurement more closely than the past value in the immediate vicinity, but adjustments farther west introduce distortions that result in lower accuracy compared to the GFS forecast.

7 Discussion

Our results show that data-driven methods can meaningfully correct numerical weather prediction and deliver substantial improvements. The transformer-based corrector consistently reduces surface-wind errors across all lead times, outperforming GFS by 45% on average at 1 h and maintaining a 13% average improvement at 48 h. These gains demonstrate that in-situ observations contain exploitable information about short-term wind patterns and that a learned corrector can translate these data into im-




Figure 6: Comparison of our model wind fields, GFS field, and their differences across three consecutive time steps at 1h lead-time. Each row corresponds to a different time step, while columns show (left) our model field, (middle) the GFS baseline field, and (right) the difference between them (ML – GFS). Teal areas in the difference panels indicate lower wind speed for the ML model, and brown areas indicate higher wind speed relative to GFS.

proved predictions. GFS provides a starting point that is physically grounded but biased, while the transformer focuses on learning local patterns and state-dependent errors from the latest observations through its attention mechanisms. Our early experiments over the ocean and prior work have both confirmed that the transformer architecture is better suited to capture observation-conditioned error patterns than simpler models like MLPs and LSTMs (Yang et al., 2024). The model also outperforms ERA5 up to 12 h, indicating that short-term, observation-driven corrections can surpass even retrospectively optimized reanalysis fields.

Model performance improves with higher observational density in space and time. The most consistent error reductions compared to GFS appear where recent observations are available, particularly along coasts and major shipping routes, while gains diminish in mid-ocean regions where measurements are sparse. In data-sparse areas, the model’s corrections naturally weaken, and predictions revert toward the GFS baseline. However, performance can also degrade below GFS skill when the model overfits to isolated measurements without learning consistent local correction patterns.

Performance improvements also vary systematically across platform types. Stationary platforms—including coastal stations, tide gauges, C-MAN stations, and fixed off-shore platforms—exhibit greater improvements than moving platforms such as drifting buoys and ships. This hierarchy likely reflects the model’s ability to learn tailored, location-specific corrections for fixed platforms, exploiting persistent error structures similar to fixed-site correction systems studied in Yang et al. (2024). In contrast, moving platforms present a harder generalization problem: ships follow dynamic routes that may traverse poorly sampled regions, while drifting buoys can enter ocean areas with few training observations. Drifting buoys may represent the hardest correction challenge because, unlike ships that tend to follow established corridors, they drift unpredictably into sparsely observed ocean regions. These platform-dependent differences reflect a broader structural challenge in the marine observing system. Although ICOADS provides the most




Figure 7: Case studies illustrating localized wind-field corrections across three representative regions. The top row shows predictions from the proposed model, and the bottom row shows the corresponding GFS baselines: (a,d) Caribbean, (b,e) Mid-Ocean, and (c,f) East Atlantic. Orange dots indicate past observations used as input, while black arrows represent future observed winds used for evaluation. In observation-dense regions (Caribbean and East Atlantic), the model adjusts wind direction and magnitude to better align with local measurements. In observation-sparse areas (Mid-Ocean), modifications remain limited, retaining the large-scale GFS structure while introducing small, locally consistent refinements.

comprehensive historical archive of marine meteorological observations, its spatial coverage remains biased toward shipping corridors and coastal regions (National Center for Atmospheric Research, 2024), leaving large oceanic areas persistently undersampled. The long-term decline in voluntary observing ships further threatens the continuity of in-situ data streams essential for machine-learning-based corrections (AWE International, 2023). Emerging opportunities, such as integrating meteorological reports transmitted through the global Automatic Identification System (AIS) network, may partially mitigate this gap by supplying near-real-time weather measurements from thousands of vessels, potentially densifying coverage in historically sparse regions and improving correction performance (U.S. Army Engineer Research and Development Center, 2024; Kumpf et al., 2021).

Operational feasibility is demonstrated at basin scale with modest computational requirements. Our global inference experiment indicates that the corrector can be applied across a 0.25° Atlantic grid with modest computational cost, making basin-scale updates feasible on operational timescales. Inference consists of a single forward pass evaluated at arbitrary target coordinates using a single NVIDIA Tesla V100 GPU (16 GB), generating corrected fields in under 5 minutes. This enables rapid regeneration of forecasts as new measurements arrive, integrating into existing pipelines that ingest GFS forecasts and available in-situ observations. While the system is already suitable for operational deployment, there remain many opportunities for future work. Beyond deterministic forecasts, operational products may also benefit from introducing uncertainty quantification through probabilistic or conformal approaches to support ensem-

ble and risk-aware decision-making in operational contexts. Likewise, while this study focuses on the Atlantic basin, extending or retraining the system for other ocean basins will test its robustness to differing climatologies and observational geometries.

8 Conclusion

Our study shows that observation-informed corrections can meaningfully improve global numerical wind forecasts in marine environments. Our transformer-based model merges heterogeneous in-situ measurements with global forecasts across the Atlantic Ocean and can be evaluated at arbitrary points in the study region, avoiding the need for pre-defined grids. Using GFS as a physical prior, the model reduces forecast errors by 45% at 1-hour lead time and 13% at 48 hours, with improvements observed across ships, buoys, and coastal stations. Rather than replacing numerical weather prediction, the model serves as a correction mechanism that adjusts large-scale forecasts using the latest available observations. Its performance gains and computational efficiency highlight the value of combining physical priors with learned, data-driven adjustments, particularly in regions where in-situ data is abundant. The spatial variability of improvements and the dependence on observational density emphasize the continued importance of robust marine observing networks. Ultimately, our work aims to strengthen machine-learning-assisted marine forecasting for applications ranging from offshore energy operations to maritime safety and navigation.

Code and Data Availability

The code and data for this paper are available on GitHub (<https://github.com/Earth-Intelligence-Lab/offshore-wind-forecasting>) and Zenodo (Yang et al., 2024, 2025).

Conflict of Interest declaration

The authors declare there are no conflicts of interest for this manuscript.

Acknowledgments

The MIT Maritime Consortium, the Doherty Professorship, and MIT Sea Grant supported this work.

References

- Atlas, R., et al. (2011). A cross-calibrated, multiplatform ocean surface wind velocity product for meteorological and oceanographic applications. *Bulletin of the American Meteorological Society*, 92(2), 157–174.
- AWE International. (2023). *New insights from old ship data: Challenges in maintaining global marine observations*. <https://www.awe.international/article/1841509/new-insights-old-ship-data>. (Accessed: October 2025)
- Bi, K., et al. (2023). Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619, 533–538.
- Copernicus Climate Change Service (C3S). (2017). *Era5: Fifth generation of ecmwf atmospheric reanalyses of the global climate*. <https://cds.climate.copernicus.eu>. Copernicus Climate Change Service.
- Dong, D., et al. (2024). Short-term marine wind speed forecasting based on dynamic graph embedding and spatiotemporal information. *Journal of Marine Science and Engineering*, 12(3), 502.
- ECMWF. (2023). *The nwp impact of aeolus level-2b winds at ecmwf* (Tech. Rep.). ECMWF Technical Memo. Retrieved from <https://www.ecmwf.int/>

- en/elibrary/81546-nwp-impact-aeolus-level-2b-winds-ecmwf
European Centre for Medium-Range Weather Forecasts (ECMWF). (2024). *Ecmwf reanalysis documentation*. <https://www.ecmwf.int/en/forecasts/datasets/reanalysis-datasets/era5>.
- Glahn, H., & Lowry, D. (1972). The use of model output statistics (mos) in objective weather forecasting. *Journal of Applied Meteorology*, 11(8), 1203–1211.
- Gneiting, T., Raftery, A. E., Westveld, A. H., & Goldman, T. (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum crps estimation. *Monthly Weather Review*, 133(5), 1098–1118.
- Google Research & DeepMind. (2023). *Metnet-3: A state-of-the-art neural weather model available in google products*. Blog post. Retrieved from <https://research.google/blog/metnet-3-a-state-of-the-art-neural-weather-model-available-in-google-products/>
- Han, X., et al. (2023). Enhanced offshore wind speed forecasts along the u.s. east coast: A deep learning framework. *Open Research Europe*, 3, 31.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., ... Thépaut, J.-N. (2020). The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730), 1999–2049. doi: 10.1002/qj.3803
- Jaegle, A., et al. (2021). Perceiver io: A general architecture for structured inputs and outputs. In *ICML*.
- Kazemnejad, A., Dehghani, M., Bahri, D., & Tay, Y. (2019). Positional encoding in transformers. *arXiv preprint arXiv:1908.06982*. Retrieved from <https://arxiv.org/abs/1908.06982>
- Kumpf, A., Hinrichs, M., & Isermann, J. (2021). Improving meteorological models using ships’ weather data communicated via ais. *Journal of Marine Science and Engineering*. Retrieved from https://www.researchgate.net/publication/355261207_Improving_Meteorological_Models_Using_Ships%27_Weather_Data_Communicated_via_AIS (Accessed: October 2025)
- Lam, R. T., et al. (2023). Learning skillful medium-range global weather forecasting. *Science*, 382(6669), 1416–1421.
- Lee, J., et al. (2019). Set transformer: A framework for attention-based permutation-invariant neural networks. In *ICML*.
- Lin, C., Li, S., Zhang, Y., & Ding, Z. (2021). Cat: Cross attention in vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision workshops* (pp. 0–0). Retrieved from <https://arxiv.org/abs/2106.05790>
- Liu, J., et al. (2023). Adaptive graph-learning convolutional network for multi-node offshore wind forecasting. *Journal of Marine Science and Engineering*, 11(4), 879.
- National Center for Atmospheric Research. (2024). *Icoads: Surface marine weather observations*. <https://climatedataguide.ucar.edu/climate-data/icoads-surface-marine-weather-observations>. (Accessed: October 2025)
- National Centers for Environmental Prediction (NCEP). (2020). *Ncep global forecast system (gfs) documentation* (Tech. Rep.). NOAA/National Weather Service. <https://www.nco.ncep.noaa.gov/pmb/products/gfs/>.
- NOAA/NWS/NCEP. (2025). *A machine learning-based bias correction method for 2-m temperature of the gfsv16* (Tech. Rep.). NOAA NCEP Office Note 520. Retrieved from https://library.oarcloud.noaa.gov/noaa_documents.lib/NWS/NCEP/NCEP_office_notes/NCEP_office_note_520.pdf
- Pathak, J., et al. (2022). Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv:2202.11214*.
- Pathak, J., et al. (2023). Deep learning for day forecasts from sparse observations. *arXiv:2306.06079*.
- Pospíchal, J., Kubovčík, M., & Dirgová-Luptáková, I. (2022). Solar irradiance forecasting with transformer model. *Applied Sciences*, 12(17), 8852. Retrieved from

- <https://www.mdpi.com/2076-3417/12/17/8852> doi: 10.3390/app12178852
- Raftery, A. E., Gneiting, T., Balabdaoui, F., & Polakowski, M. (2005). Using bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, 133(5), 1155–1174.
- Raschka, S. (2023). A tour of self-attention. *arXiv preprint arXiv:2308.13086*. Retrieved from <https://arxiv.org/abs/2308.13086>
- Rasp, S., & Lerch, S. (2018). Neural networks for postprocessing ensemble weather forecasts. *Monthly Weather Review*, 146(11), 3885–3900.
- Remote Sensing Systems. (2024). *Ccmp wind vector analysis product*. Product documentation. Retrieved from <https://www.remss.com/measurements/ccmp/>
- Rennie, M. P., et al. (2021). The impact of aeolus wind retrievals on ecmwf global weather forecasts. *Quarterly Journal of the Royal Meteorological Society*, 147(740), 3555–3586.
- Rußwurm, M., Klemmer, K., Rolf, E., Zbinden, R., & Tuia, D. (2023). Geographic location encoding with spherical harmonics and sinusoidal representation networks. *arXiv:2310.06743*.
- Scheuerer, M., & Hamill, T. M. (2020). Postprocessing of ensemble temperature forecasts via quantile regression forests. *Monthly Weather Review*. (Early ML post-processing example)
- Sitzmann, V., et al. (2020). Implicit neural representations with periodic activation functions. In *Neurips*.
- Su, L., Zuo, X., Li, R., Wang, X., Zhao, H., & Huang, B. (2025). A systematic review for transformer-based long-term series forecasting. *Artificial Intelligence Review*, 58(80). doi: 10.1007/s10462-024-11044-2
- Sun, D., et al. (2023). Deep learning improves gfs wintertime precipitation forecasts over the southeastern united states. *Geophysical Research Letters*, 50(22), e2023GL104406.
- Tancik, M., et al. (2020). Fourier features let networks learn high frequency functions in low dimensional domains. In *Neurips*.
- U.S. Army Engineer Research and Development Center. (2024). *Sharing ships’ weather data via ais: Concept and results from multiyear observations*. <https://www.erd.c.usace.army.mil/Media/Publication-Notices/Article/3836011/sharing-ships-weather-data-via-ais-concept-and-results-from-multiyear-observati/>. (Accessed: October 2025)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998–6008).
- Wilks, D. S. (2011). *Statistical methods in the atmospheric sciences* (3rd ed.). Academic Press.
- Yang, Q., Giezendanner, J., Civitarese, D. S., Jakubik, J., Schmitt, E., Chandra, A., ... Wang, S. (2024). Local off-grid weather forecasting with multi-modal earth observation data. *arXiv preprint, arXiv:2410.12938*. Retrieved from <https://arxiv.org/abs/2410.12938> (Last revised 25 August 2025 (v4))
- Ye, F., Brodie, J., Miles, T., & Ezzat, A. A. (2023). AIRU-WRF: A physics-guided spatio-temporal wind forecasting model for u.s. mid-atlantic offshore. In *arxiv:2303.02246*.
- Ye, F., Zhang, X., Stein, M., & Ezzat, A. A. (2024). Deepmide: A multivariate spatio-temporal method for ultra-scale offshore wind energy forecasting. *arXiv:2410.20166*.
- Zhang, W., Zhang, L., & Liu, J. (2024). Satellite-based re-examination of changes in terrestrial near-surface wind speed in the last 30 years. *Frontiers in Earth Science*, 12. Retrieved from <https://doi.org/10.3389/feart.2024.1398638> doi: 10.3389/feart.2024.1398638
- Zuo, H., & Hasager, C. B. (2023). The impact of aeolus winds on near-surface wind

forecasts over tropical ocean and high-latitude regions. *Atmospheric Measurement Techniques*, 16, 3901–3913. doi: 10.5194/amt-16-3901-2023

Appendix A Appendix




Figure A1: Pearson correlation (r) between target wind observations and past measurements as a function of spatial and temporal separation. The x-axis shows the maximum spatial distance threshold (in kilometers) within which past observations are considered, and the colored lines correspond to different temporal windows ($\Delta t_{\max}$) ranging from 5 to 71 h, as indicated in the legend. Each curve therefore represents how correlation decreases with increasing spatial distance for a fixed temporal window. The results are computed using pairs of in-situ observations separated by approximately one hour ($\min \Delta t = -1$ h). Correlations generally decrease both with distance and with increasing temporal window size, indicating that temporal proximity has a stronger influence on correlation than spatial distance.

Modification	Performance gain
Two previous hours of observations	None
Additional GFS variables (Temperature, Humidity, q)	None
Platform/object type encoding	None
Residual GFS connection after cross-attention	None

Table A1: Summary of architectural variations tested. None improved performance relative to the baseline model.






Figure A2: Schematic illustration of the GFS forecast cycle and its alignment with observation times. Each GFS forecast is initialized every 6 hours and provides predictions at multiple lead times. For a given observation or target time (gray circles), the corresponding GFS value is selected from the most recent available forecast cycle based on its valid time. The colored markers indicate which GFS forecast (by initialization time and lead time) is used to match each observation. To construct a continuous sequence of predictions up to a 48-hour horizon, forecasts from eight consecutive GFS cycles are combined, ensuring that each observation uses the most temporally consistent forecast available.



(a) Spatial distribution of the training (blue), validation (orange), and test (green) subsets across the North Atlantic. Each color represents the observation locations assigned to the corresponding subset, ensuring spatial diversity and minimal overlap between regions.



(b) Temporal allocation of samples showing the chronological split between training, validation, and test periods. The temporal segmentation ensures that validation and testing occur strictly after the training period, avoiding temporal leakage.

Figure A3: Spatial and temporal partitioning of the dataset used for model development and evaluation. The top panel illustrates the spatial separation between subsets, while the bottom panel shows the chronological division of samples through time. Together, these splits ensure that the model is trained and evaluated on independent spatial and temporal domains.




Figure A4: Relationship between observation density and spatial error reduction across lead times. Scatter plots showing the average spatial error reduction (model minus GFS) as a function of the fraction of total observations within each $1^\circ \times 1^\circ$ grid cell for lead times of (a) 1 h, (b) 8 h, (c) 24 h, and (d) 48 h. Each point represents a single grid cell, with negative values indicating that the machine learning model outperforms GFS. The dashed horizontal line at $y = 0$ marks the threshold where the model and GFS have equal performance. Linear trend lines (brown) indicate a consistent negative relationship between observation density and error reduction across all lead times. The trend suggests that grid cells with higher observation density tend to show larger improvements over GFS. Conversely, data-sparse regions show more variable performance, with some cells achieving substantial improvements (up to 4 m s^{-1} error reduction) while others show degradation.