

WxC-Bench: A Novel Dataset for Weather and Climate Downstream Tasks

Received: 24 December 2024

Accepted: 5 February 2026

Cite this article as: Shinde, R., Ankur, K., Phillips, C.E. *et al.* WxC-Bench: A Novel Dataset for Weather and Climate Downstream Tasks. *Sci Data* (2026). <https://doi.org/10.1038/s41597-026-06839-7>

Rajat Shinde, Kumar Ankur, Christopher E. Phillips, Aman Gupta, Simon Pfreundschuh, Sujit Roy, Sheyenne Kirkland, Vishal Gaur, Venkatesh Kolluru, Amy Lin, Prajun Trital, Aditi Sheshadri, Udaysankar Nair, Manil Maskey & Rahul Ramachandran

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

SCIENTIFIC DATA

CONFIDENTIAL

COPY OF SUBMISSION FOR PEER REVIEW ONLY

Tracking no: SDATA-24-04381B

WxC-Bench: A Novel Dataset for Weather and Climate Downstream Tasks

Authors: Rajat Shinde (The University of Alabama in Huntsville), Ankur Kumar (The University of Alabama in Huntsville), Christopher Phillips (The University of Alabama in Huntsville), Aman Gupta (Stanford University), Simon Pfreundschuh (Colorado State University), Sujit Roy (The University of Alabama in Huntsville), Sheyenne Kirkland (The University of Alabama in Huntsville), Vishal Gaur (The University of Alabama in Huntsville), Venkatesh Kolluru (The University of Alabama in Huntsville), Amy Lin (The University of Alabama in Huntsville), Prajun Trital (The University of Alabama in Huntsville), Aditi Sheshadri (Stanford University), Udaysankar Nair (The University of Alabama in Huntsville), Manil Maskey, and Rahul Ramachandran (NASA Marshall Space Flight Center)

Abstract:

High-quality openly-accessible machine learning (ML)-ready datasets play a foundational role in developing new artificial intelligence (AI) models or fine-tuning existing models for scientific applications such as weather and climate analysis. However, despite the growing development of new deep learning models for weather and climate, there is a scarcity of curated, pre-processed ML-ready datasets. Curating such high-quality datasets for developing new models is challenging particularly because the modality of the input data varies significantly for different downstream tasks addressing different atmospheric scales (spatial and temporal). Here we introduce WxC-Bench (Weather and Climate Bench), a multi-modal dataset designed to support the development of generalizable AI models for various downstream use-cases in weather and climate research. WxC-Bench supports examining several atmospheric processes from meso- β (20 - 200 km) scale to synoptic scales (2500 km), such as aviation turbulence, hurricane intensity and track monitoring, weather analog search, gravity wave parameterization, and natural language report generation. The dataset and associated code have been made publicly available, and can be accessed using WxC-Bench Python package.

Datasets:

Repository Name	Dataset Title	Accession Number or DOI	URL to data record	Private reviewer access URL/code
WxC-Bench	Code Repository to create WxC-Bench Dataset		https://github.com/NASA-IMPACT/WxC-Bench	
WxC-Bench	Python Package		https://pypi.org/project/wxcbench/	
WxC-Bench	WxC-Bench	https://doi.org/10.57967/hf/7711	https://huggingface.co/datasets/nasa-impact/WxC-Bench	

WxC-Bench: A Novel Dataset for Weather and Climate Downstream Tasks

Rajat Shinde^{1,*}, Kumar Ankur^{1,†}, Christopher E. Phillips^{1,†}, Aman Gupta^{2,†}, Simon Pfreundschuh^{3,†}, Sujit Roy^{1,*}, Sheyenne Kirkland¹, Vishal Gaur¹, Venkatesh Kolluru¹, Amy Lin¹, Prajun Trital^{4,5}, Aditi Sheshadri², Udaysankar Nair¹, Manil Maskey⁶, and Rahul Ramachandran⁶

¹Earth System Science Center, The University of Alabama in Huntsville, Huntsville, AL, USA

²Stanford University, Stanford, CA, USA

³Department of Atmospheric Science, Colorado State University, Fort Collins, CO, USA

⁴Department of Computer Science, The University of Alabama in Huntsville, Huntsville, AL, USA

⁵Department of Space Science, The University of Alabama in Huntsville, Huntsville, AL, USA

⁶NASA Marshall Space Flight Center, Huntsville, AL, USA

*corresponding author(s): Rajat Shinde (rajat.shinde@uah.edu), Sujit Roy (sujit.roy@nasa.gov)

†these authors contributed equally to this work

ABSTRACT

High-quality openly-accessible machine learning (ML)-ready datasets play a foundational role in developing new artificial intelligence (AI) models or fine-tuning existing models for scientific applications such as weather and climate analysis. However, despite the growing development of new deep learning models for weather and climate, there is a scarcity of curated, pre-processed ML-ready datasets. Curating such high-quality datasets for developing new models is challenging particularly because the modality of the input data varies significantly for different downstream tasks addressing different atmospheric scales (spatial and temporal). Here we introduce WxC-Bench (Weather and Climate Bench), a multi-modal dataset designed to support the development of generalizable AI models for various downstream use-cases in weather and climate research. WxC-Bench supports examining several atmospheric processes from meso- β (20 - 200 km) scale to synoptic scales (2500 km), such as aviation turbulence, hurricane intensity and track monitoring, weather analog search, gravity wave parameterization, and natural language report generation. We provide a comprehensive description of the dataset and also present a technical validation for baseline analysis. The dataset and code to prepare the ML-ready data have been made publicly available on Hugging Face, and can be accessed using WxC-Bench Python package.

Background & Summary

figures/WxC-Bench.pdf




Figure 1. Conceptual overview of WxC-Bench [1] development and its role in end-to-end AI/ML workflows, illustrating the conversion of raw, non-ML-ready data into six task-specific, self-contained ML-ready datasets that can be used independently for training new models or fine-tuning pre-trained models across multiple spatial and temporal scales and resolutions.

Weather and climate variability influences all spheres of human activity including the economy, transport, and energy production. Extreme events such as hurricanes, wildfires, droughts, and floods often result in catastrophic damage to human lives and capital. Current forecast models can reliably predict atmospheric extremes only up to 14-15 days in advance with forecast skill decreasing with time. In the past, such limits of predictability were believed to be improved mostly by recording more comprehensive set of observations, and through improving traditional models. However, the past couple of years have witnessed the emergence of a radically new, Artificial Intelligence (AI) driven approach.

Table 1. Summary of recently proposed weather and climate data benchmarks and their key features.

Paper Title	Year of Publication	Key Features
WeatherBench [22], WeatherBench 2 [23]	2020, 2024	Dataset for data-driven medium-range weather forecasting (specifically 3–5 days). Dataset curated from ERA5 re-gridded to lower resolutions 5.625° (32×64 grid points), 2.8125° (64×128 grid points), and 1.40525° (128×256 grid points) to be used for deep learning (DL) approaches.
ClimateNet [24]	2020	Dataset trained on expert-labeled climate data via ClimateContours, enables precise weather pattern segmentation, which utilizes output data from the 25 km Community Atmospheric Model (CAM5.1). This is a curated expert-labeled dataset which is trained on the DL segmentation model.
ClimateBench [25]	2022	Dataset curated from a suite of Earth System Models participating in the Coupled Model Intercomparison Project 6 (CMIP6), covering a variety of global warming scenarios as yearly averages, at the native grid resolutions ($\sim 2^\circ$).
ClimSim [26]	2024	Dataset compiled using the Energy Exascale Earth System Model (E3SM)-multiscale modeling framework (MMF) or E3SM-MMF climate simulator and stored every 20 minutes for 10 years, at an approximate grid-spacing of 1.5° using a cubed-sphere unstructured grid. Curated for hybrid physics-based ML climate emulation.
Climatelearn [27]	2024	Dataset curated from the ERA5, extreme ERA5, CMIP6 and PRISM, and has been used for weather forecasting, downscaling and climate projection.

AI-driven weather forecasting models deliver accurate weather predictions quickly and cost-effectively than traditional numerical models. Notable examples of these models are FourCastNet [2], Graphcast [3], FengWu [4], Pangu [5], and MetNet [6]. Machine Learning (ML) models also hold potential for longer-term subseasonal-to-seasonal modeling through efficient ensemble generation or otherwise [7]. Additionally, ML-based methods are being explored to improve climate prediction [8], development of ML-driven parameterizations [9, 10, 11, 12], bias correction [13], and understanding climate change impacts [14, 15]. Recently, AI powered Numerical Weather Prediction (NWP) models, primarily using transformer-based architectures, blended with Fourier Neural Operators [2], Clifford Neural Operators [16] (NOs), and Graph Neural Networks have been introduced. These techniques aim to learn complex, nonlinear mappings required for approximating Partial Differential Equations (PDEs) that capture the inherent atmospheric physics.

Unlike traditional NWP models, which are widely used to study weather extremes, ozone recovery, air quality monitoring, parameterization development, boundary layer turbulence, etc., these pioneering AI models are typically trained on large volumes of data, but only to perform one specific task: medium-range weather forecasting. Thus, severely limiting their broader applicability, explaining the rising push to develop general-purpose AI models.

Recent models such as Prithvi WxC [17], ClimaX [18], Aurora [19], and Prithvi-EO-2.0 [20] have been developed as a generalist AI-system capable of performing multiple tasks. These generalist models are termed foundation models (FMs) for earth sciences. FMs [21], which were first proposed in the fields of natural language processing (NLP) and computer vision (CV), are general-purpose AI models that replace specialized task-specific models across a wide range of applications. FMs are trained in an unsupervised manner on a large volume of datasets. Within atmospheric sciences, Prithvi WxC FM [17], a 2.3 billion parameter model has been designed for weather and climate research with a focus on applying it to multiple downstream tasks.

The emergence of weather and climate FMs like Prithvi WxC, ClimaX, etc., also poses questions about their generalization over several scientific applications and inherent uncertainties. Therefore, to empower a systematic benchmarking and performance evaluation of such models, there is a need for a unified ML-ready dataset that spans over multiple spatial and temporal scales. WxC-Bench [1], the ML-ready dataset presented in this study, empowers such evaluations by facilitating the application of FMs to support downstream tasks and allowing a scoring of their performance.

In this context, Table 1 outlines a few prominent ML-ready datasets that already exist. These include WeatherBench [22] and WeatherBench2 [23] which are used exclusively for consistently scoring AI forecasts. Another dataset, ClimateNet [24], provides a collection of expert-labeled tropical cyclones and atmospheric rivers. On longer timescales, ClimateBench [25] and ClimSim [26] both provide curated climate model datasets for training climate emulators. Additionally, data processing tools for diverse ML tasks, such as Climatelearn [27] have also been developed. The main contribution of this paper is an extension of our previous work [28] with detailed analysis and comprehensive evaluation of the datasets.

Expanding the scope of the tasks and datasets used to evaluate the transfer learning capabilities of ML models is important for weather and climate applications, given their spatial and temporal complexities. To explore and assess the potential of generalizable AI models / FMs for meteorology research, it is imperative that the considered tasks (a) span multiple spatial and temporal scales, from forecasting, to seasonal prediction, and climate scale prediction (Figure 1 right), and (b) vary in input formats and output applications that require assimilating data from multiple sources.

The main contribution of this paper is WxC-Bench (See Figure 1, a multi-modal dataset that supports not only the evaluation of generalizable AI models for weather and climate research but also in the development of state-of-the-art ML benchmarks for scientific tasks. The dataset can be used for tasks spanning multiple spatio-temporal scales (Figure 1 right). While creating WxC-Bench [1], the selected tasks are designed to cover a diverse set of ML applications, including natural language forecast summarization, image classification and regression, to represent atmospheric processes ranging from meso- β (20 - 200 km) to synoptic scales (2500 km).

Downstream Task Descriptions and Datasets

This paper presents a compilation of 6 ML-ready datasets for developing task-specific ML models. The tasks are selected to represent a broad scale of the atmospheric processes and challenges (Table 2) while also having useful applications for research and day-to-day human activities (See Relevance column of Table 2). Some tasks, such as gravity wave (GW) parameterization, are considered relatively straightforward, while others like natural language forecasting are much more complicated attributing to learning from multi-modal features. This enables a comprehensive and complete assessment of a trained AI model's transfer learning capability for generalization across atmospheric science.

Table 2. Comparison of downstream machine learning (ML) tasks used in this study.

ML Tasks	Downstream Task	Temporal Scale	Spatial Scale	Relevance
Classification	Aviation Turbulence	Upto seconds	Microscale	Vital for flight safety and predicting turbulent regions
Parameterization	Atmospheric Gravity Waves	Upto Day	Mesoscale	Crucial for understanding momentum transfer and vertical couplings in the atmosphere
Search	Weather Analog Search	Hours	Meso-to-Synoptic scale	Helps in finding past weather patterns similar to current conditions for forecasting
Prediction	Long Term Precipitation Forecasting	Subseasonal	Meso-to-Synoptic scale	Essential for water resource management and agricultural planning
Prediction	Hurricane	Weeks	Mesoscale	Critical for predicting hurricane paths and impacts, saving lives and property
Multi-modal generation	Natural Language Forecasting	Daily	Synoptic	Communicates complex meteorological data into understandable forecasts for day-to-day activities

In the following sections, we elaborate on the dataset for the above-mentioned downstream tasks, each targeting a specific atmospheric process.

Aviation Turbulence Prediction

Aviation Turbulence in the lower and middle atmosphere presents risks for passenger and cargo airliners, especially when it is encountered unexpectedly [29, 30]. Turbulence, however, is a microscale feature that defies direct prediction by numerical weather prediction models with even relatively fine grid-spacing on the order of 1 Km. While machine learning techniques have been explored for turbulence prediction based on information from numerical weather prediction or onboard sensors [31, 32, 33], deep learning (DL) techniques remain under-investigated. Thus, there is a scope for sophisticated data-driven approaches to improve turbulence forecasts for aviation, which can be supported by this dataset.

Data Description The training dataset prepared for aviation turbulence detection is derived from the Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2) re-analysis data set [34]. MERRA-2 is available on a $0.625^\circ \times 0.5^\circ$ grid resolution with 72 hybrid model vertical levels and 42 output pressure levels. With this grid spacing, MERRA-2 resolves synoptic scale features such as frontal zones and midlatitude cyclones, which can impact turbulence production but not fine-scaled processes such as boundary layer growth. MERRA2 is used to provide temperature, relative humidity, zonal and meridional winds, vertical motion, orography, and geopotential height. These variables are retrieved for the

lowermost 40 pressure levels, spanning from the surface to 100 hPa. Each input feature is valid at 1800 UTC, corresponding to approximately local noon (Central Time) over the contiguous United States (CONUS).

The training labels for this task are Pilot Reports (PIREPs). PIREPs are point-location reports written by pilots that describe in-flight conditions and can include a variety of information, such as cloud layers, general weather conditions, icing, turbulence encountered, etc. The PIREPs are retrieved via the Iowa State University archives [35] and have a temporal availability starting January 1, 2003. In this work, we used a dataset spanning January 1, 2003, to November 6, 2023, which primarily includes PIREPs recorded over the CONUS. The spatial distribution of the turbulence reports are displayed in Figure 2.



Figure 2. Spatial distribution of PIREPs turbulence reports over CONUS from 2003-Present. Note that values are shown on a logarithmic scale.

The PIREP reports were first grouped by flight levels into bins of 0-15,000 ft (LOW), 15,000-30,000 ft (MIDDLE), and 30,000 ft or greater (HIGH). These levels were selected as they correspond approximately to the planetary boundary layer (PBL), free atmosphere, and tropopause respectively. Flight levels above 50,000 ft were removed as this is typically above the flight level of commercial aircraft.

Table 3. Distribution of turbulence instances across altitude levels.

Level	Turbulent	Non-turbulent	# Turbulences	# Data Points
LOW	10%	90%	38,520	371,473
MIDDLE	8.5%	91.5%	13,553	158,687
HIGH	9%	91%	22,563	243,388

To convert turbulence reports into binary classification labels, all reports are re-classified as moderate or greater (MODG) following prior studies [33, 36]. PIREPs are then binned into the nearest MERRA-2 grid cell by day, and any cell with MODG report more than 25% is considered "turbulent" for that day. Initially, the reports were binned on 3-hourly intervals to match

MERRA-2 output, but this was found to have very few reports within each cell for each time period. The distribution of labels for each level is shown in Table 3.

Gravity Wave (GW) Parameterization

Atmospheric mesoscale processes like clouds, gravity waves (GWs), and precipitation, have global impacts on the atmospheric flow but are often *parameterized* in weather and climate models [37, 38, 39]. Climate model parameterizations are typically designed to be strictly vertical i.e., single column [40, 41, 42], due to which they neglect the horizontal evolution of mesoscales. Particularly in the context of meso- α and meso- β GWs, this creates a major representation bias in weather and climate models, leading to global momentum imbalances such as colder-than-observed temperatures in the high latitudes (termed as cold-pole bias) or biases in global mass transport [43, 44, 45]. This means that the representation of horizontal propagation of atmospheric GWs, and in general, other parameterized meso- α and meso- β atmospheric processes, is crucial for both regional short-term weather forecasting and global long-term climate prediction [46, 47, 48].

GWs couple the different layers of the atmosphere and demonstrate significant horizontal motion as they propagate. This property makes them an excellent candidate for simulation using a specialized AI-model or a generalized FM, and for testing of AI's capability to learn regional variability and vertical coupling of the atmosphere. A successful ML-based prediction of subgrid-scale GW momentum fluxes opens avenues to first learn missing model physics from high-resolution observations and then represent the missing mesoscale physics in coarse-climate models using ML. The strategy can be applied to various other mesoscale processes, which are currently parameterized using (approximate) single-column schemes in weather and climate models. Availability of curated datasets, such as WxC-Bench, thus promotes the development of fast and potentially more physically accurate ML parameterizations of atmospheric processes.




Figure 3. Distribution of coarse-grained input features (scaled background atmospheric state) and output labels (scaled gravity-wave momentum fluxes) for January 2010, derived from ERA5.

Data Description For a given background state of the atmosphere, GW parameterizations in climate models predict the associated subgrid-scale momentum fluxes due to GWs [38]. Here, this problem is formulated as a regression task to predict the resolved GW fluxes associated with an input background state.

The subgrid-scale momentum fluxes, i.e., the output, are computed using Helmholtz decomposition. Helmholtz decomposition of the full horizontal velocity field $\vec{u} = (u, v)$ decomposes the full flow into its rotational (balanced) and divergent (unbalanced) parts, i.e., $(u, v) = \vec{u} = \vec{u}_{rot} + \vec{u}_{div}$. The divergent part is associated with GWs [49]. A high-pass filter is then applied to remove the first 21 harmonics of the divergent flow. The eddy vertical velocity, ω' , is computed by removing the zonal mean from full ω . The divergent part \vec{u}_{div} and ω' are then used to compute the fluxes, $(F_u, F_v) = (u_{div}\omega', v_{div}\omega')$.

The training data for global gravity wave parameterizations is prepared using the European Centre for Medium-Range Weather Forecasting (ECMWF) ReAnalysis v5 (ERA5) data product [50]. ERA5 resolves GWs with wavelengths exceeding 150-200 km and longer by accounting for grid-scale hyper-diffusion and other numerical effects [51, 52, 53]. The ERA5 data product is publicly available at a horizontal resolution of $0.25^\circ \times 0.25^\circ$ and 137 hybrid pressure levels at an hourly temporal resolution. The global wind and temperature data are retrieved on 122 vertical levels: from 1.45 hPa (model level 16) to the surface (model level 137), and are used to compute the training labels, i.e., the momentum fluxes due to GWs. The top 15 levels near the model top are removed to avoid the artificial sponge damping in the mesosphere.

The background winds and temperature at ERA5's native grid resolution are conservatively coarse-grained to a $2.8^\circ \times 2.8^\circ$ (64×128) Gaussian grid to obtain the background state of the atmosphere, i.e., the training input. Likewise, the output fluxes are computed on the native grid and then conservatively coarse-grained to the 2.8° grid, as denoted by the overbar. At this resolution, WxC-Bench provides a detailed, multi-year, global, evolution of subgrid-scale fluxes, far superseding other existing datasets, e.g., ClimSim [26]. Atmospheric GWs exist on spatial scales ranging from $O(1)$ km to $O(1000)$ km. We do not claim that ERA5, at 30 km horizontal resolution, completely resolves these scales. However, given the dearth of storm-resolving global climate simulations, ERA5 still presents as one of the most extensive long-term records to study atmospheric GWs with wavelengths ~ 150 km and above.

The input feature vector consists of three dynamical variables: zonal wind (u), meridional wind (v), and potential temperature, $\theta = T(p/p_0)^\gamma$, where T is the temperature, p is the pressure, $p_0 = 1000$ hPa is a reference surface pressure, and $\gamma = -0.286$ is a constant). Each quantity is a function of longitude, latitude, pressure (height), and time. Note that the input comprises the full winds, not just the divergent component. The output contains the vertical flux of zonal and meridional momentum carried by

GWs, i.e., $u'\omega', v'\omega'$. Similar to the input, the output is computed globally for each longitude, latitude, pressure (height), and time. For this downstream task, input features and labels are computed using hourly-ERA5 for four years, viz., 2010, 2012, 2014, and 2015. These years were chosen based on rich wintertime orographic GW activity in the Southern Hemisphere. Thus, the dataset comprises of a total of $64 \times 128 \times 24 \times 1461$ (~ 287 million) columns of training and validation data.

Additionally, the winds, temperature, etc., are normalized using their respective 4-year global mean and standard deviation as follows:

$$u_{scaled}(t, p, \phi, \lambda) = \frac{u'(t, p, \phi, \lambda)}{3\sigma_u} = \frac{u(t, p, \phi, \lambda) - \mu_u}{3\sigma_u} \quad (1)$$

where μ_u is the temporally and zonally averaged zonal wind and σ_u is the standard deviation. Identical normalization is performed for v . The scaling transforms the input values to a range within $[-2, 2]$. The potential temperature, θ , was scaled simply by dividing by 1000.

GW generation is highly intermittent; thus, the GW fluxes can have a symmetric exponential (Laplace) distribution with small values around zero. Therefore, the GW fluxes are scaled by their respective global means and standard deviations and then a cube root is applied, which transforms very small and very large values and maps them closer to 1. The normalized variables are then concatenated to form the input and output variables. Figure 3 illustrates the input and output variables distribution.

Weather Analog Search Dataset

Weather analogs, i.e. weather events and patterns which are similar to each other, serve several purposes in the Earth science community. For forecasters, past events provide insight into how current weather can evolve and serve as initial conditions for ensemble model systems [54, 55] and could also help train deep-learning models. For researchers, weather analogs make useful case studies for analyzing how particular weather systems evolve [56, 57]. Thus, a rapid analog search is useful for improving forecasts and aiding scientists in locating desirable historic records for their research.

Data Description For searching analogs over a desired dataset, e.g., MERRA-2, the data samples are required to be encoded to enable rapid search over the data space. There are several ways to accomplish this, including vector search based on approaches that minimize distances between embeddings generated by post-tokenization modules within machine learning architectures [58]. Also, lookup table-based searching over a database of images using transforms - Fast Fourier Transform [56] or Wavelet “fingerprints” as search keys [59, 60]. Searching analogs is also possible based on deep learning-based approach for generating latent features using all the temporal multivariate input predictors in this latent space rather than the original predictor space [61].

In this context of analog search, the searching approach must be spatially constrained as the global repetition of an atmospheric state is unlikely [55]. Thus, following scenarios emerge for potential applications from a user perspective:

- *A user searches for a specific phenomena, such as a mid-latitude cyclone, and is given multiple instances from across the globe.*

In this scenario, the model provides users with multiple instances of mid-latitude cyclones distributed globally. Each analog is spatially localized, increasing the likelihood of meaningful similarity to the user’s target case. Such an approach holds particular utility for individuals engaged in constructing climatologies or those endeavoring to identify case days for analyzing mid-latitude cyclone-driven weather events.

- *A user constrains the model to search a domain of interest.*

Alternatively, the model can be tailored to search within a designated domain of interest, restricting its focus to a specific region. This configuration is particularly advantageous for forecasters with a vested interest in the weather dynamics of a particular locale and for researchers and climatologists seeking targeted insights into regional phenomena.

figures/analog_fig04_updated.jpeg

ARTICLE IN PRESS

Figure 4. Example training inputs for 1 January 2019: (a) sea-level pressure and (b) 2-m air temperature.

The above-mentioned approaches may be handled simultaneously by breaking the larger grids into sub-grids and encoding each separately with location information. This allows the users to perform search operations across either all or some sub-grids. This sub-grid strategy has been demonstrated operationally in large-scale meteorological analogue-search workflows.

Using the same methodology, multiple MERRA-2 sub-grids have been generated with edge lengths of approximately 1,800 km (Figure 4). Note that the domain size in grid points is preserved, meaning that the edge length of the subgrids varies with latitude. Edge length of 1800 km has been shown to be adequate for forecasting features over a point of interest for at least 12 hours while maintaining a robust number of analogs available in the historical record [55]. Each subgrid overlaps by 10 points on the western and southern edges to mitigate issues where the point of interest is on the edge of a sub-domain. This leads to 623 subgrids for the entire globe at the MERRA-2 resolution of 0.5×0.625 degrees.

Long-range Precipitation Forecasting Dataset

The long-term precipitation or long-range precipitation forecast task consists of predicting global, daily-accumulated precipitation four weeks into the future. This forecasts range falls within the so-called *sub-seasonal to seasonal range* (S2S), which constitutes an important frontier in weather and climate research. Conventional numerical weather prediction (NWP) can produce skillful precipitation forecasts ten to twelve days into the future but beyond that their skill deteriorates to or below the level of a climatology [62]. While climate models can predict weather patterns on seasonal scales, the S2S range is currently considered a predictability desert [63]. Since many decisions in the management of agricultural and water resources fall into this two-weeks to two-months time scale, improving precipitation forecasts several weeks into the future has significant societal value [64]. The proposed dataset is targeted to address the S2S precipitation forecasting.

The input data for the long-range precipitation forecast consists of gridded, global satellite observations from a range of sensors providing observations in the visible, infrared, and microwave regions of the electromagnetic spectrum. The reference data is derived from satellite-based precipitation estimates that were corrected to match ground-based gauge measurements and arguably constitute the most reliable global estimates of precipitation currently available.

A related dataset that has been proposed is RainBench [65], which combines *simulated* satellite observations with satellite- and reanalysis-based precipitation estimates for machine learning-based precipitation forecasts. However, the simulated satellite observations in RainBench cover only a limited number of channels and are available only from 2016 onward, which makes the dataset less suitable for evaluating long-range forecasts of precipitation. Moreover, since the simulated observations are based on non-cloud-resolving numerical models they will inherit their limitations regarding the representation of cloud- and precipitation-related processes. Here, we propose a novel benchmark dataset for long-range precipitation forecasts based on almost forty years of satellite observations and corresponding precipitation estimates.

Data Description The input data for the long-range precipitation forecast consists of gridded satellite observations from three observation sources: (1) infrared and visible observations from geostationary satellites, (2) infrared and visible observations from polar-orbiting satellites, and (3) microwave observations from polar-orbiting satellites. The input observations were chosen to maximize the temporal coverage of the training dataset. The observations are extracted from three observational records. Geostationary observations are taken from the GridSat-B1 dataset [66]. Infrared and visible observations from polar-orbiting platforms are sourced from the Advanced Very High-Resolution Radiometer (AVHRR) and the High-Resolution Infrared Radiation Sounder (HIRS) within the PATMOS-x fundamental climate data record [67]. Microwave observations are taken from the Special Sensor Microwave/Imager (SSM/I) and the Advanced Microwave Scanning Radiometer 2 (AMSR2) products provided through the NOAA Climate Data Record (CDR) of microwave brightness temperatures [68]. Example observations from each of these records from selected channels for 1 January 2000 are shown in Figure 5.

The reference precipitation estimates are derived from the PERSIANN-CDR product [69] for the time range 1983 through June 2000 and the IMERG Final [70] product from June 2000 through 2024. In contrast to PERSIANN-CDR, which is based solely on geostationary visible and infrared observations, IMERG combines observations from passive microwave and geostationary sensors. Because of this, IMERG can be expected to provide better sensitivity to precipitation and is therefore used to evaluate the forecasts. However, the IMERG precipitation estimates are available only from the year 2000 onward. We therefore provide an extended training record based on the combined estimates from PERSIANN-CDR and IMERG for the case that the extended training period may prove beneficial for the training of the forecast models.

Both input and reference data are interpolated to a regular latitude-longitude grid with a zonal resolution of 0.625° and meridional resolution of 0.5° . The data are organized into separate files by day and input and target source.

figures/fig05.jpeg

ARTICLE IN PRESS

Figure 5. Example inputs for the long-range precipitation forecast task. Panels (a–c) show visible and infrared geostationary observations from GridSat-B1, panels (d–f) show corresponding polar-orbiting observations from PATMOS-x, and panels (g–i) show microwave observations from SSM/I.

Table 4. Number of daily training files for the long-range precipitation forecast task, by dataset split and input data source.

Split	Start	End	PATMOSx files	GridSat files	SSMI files	Precipitation files
Training	1983-01-01	2017-12-31	12 591	11073	12763	12784
Validation	2018-01-01	2018-12-31	358	358	358	358
Testing	2019-01-01	2022-12-31	1408	1442	1442	1444

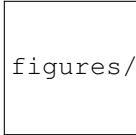

 figures/hurricane_tracks_2017.jpeg

Figure 6. Tracks of named tropical cyclones across all basins in 2017, categorized by intensity from tropical storm (TS) through Category 5 (Cat 5).

precipitation forecasts task. Years 1983 through 2018 of the extracted satellite observations and matched precipitation estimates are designated as training data, while year 2019 through 1 January 2023 are used for testing. Table 4 lists the number of files by input and target source for the three dataset splits. Scripts to evaluate the machine-learning-based forecasts against the conventional NWP baselines are provided together with datasets and section presents the technical validation of the proposed dataset.

Hurricane Prediction and Intensity Estimation Dataset

Hurricanes are one of the most destructive natural phenomena affecting human populations, infrastructure, and economies, causing approximately \$17 billion in damages annually in just the United States. Moreover, it has been observed that the frequency of hurricane occurrence is increasing over the years [71]. Further, studies show that the power dissipated by these hurricanes has doubled since the 1980s [72], likely due to rising sea surface temperatures [73]. Many of these hurricanes make landfall along the coastal regions. Thus, there is a need for accurate and timely predictions of hurricanes.

Data Description This dataset aggregates hurricane track data from the global basins spanning 1980 to 2024. Data is primarily sourced from the Hurricane Database (HURDAT) maintained by the United States National Hurricane Center for Atlantic and Pacific Oceans and International Best Track Archive for Climate Stewardship (IBTrACS) for rest of the basins. Hurricane details was accessed through the `tropical` package [74], and the track, intensity (10m sustained wind speed and minimum sea level pressure) were retrieved.


By compiling information on tropical cyclones (TCs) from all major ocean basins, this dataset enables models to learn from a spatially and dynamically diverse collection of storms, thereby improving their generalization across different regions and meteorological regimes. Figure 6 shows the global distribution of named TCs during 2017, categorized from tropical depressions (TD) and tropical storms (TS) to Category 5 (Cat 5) hurricanes. This year serves as a representative snapshot of global TC activity and the spatial variability of their intensities across the major basins. In 2017, a total of 95 named storms were recorded globally across seven basins. Among these, approximately 44% were tropical storms, while about 24% reached Category 1–2 intensity and around 22% intensified to major hurricane strength (Category 3 or higher). The Western Pacific was the most active basin, producing 33 named storms, followed by the Eastern Pacific (20 storms) and the North Atlantic (18 storms). The remaining basins, Australian Region (8), South Indian Ocean (6), South Pacific (6), and North Indian Ocean (4), had comparatively fewer but regionally significant systems. The global category distribution for 2017 was as follows: TD (10.5%), TS (44.2%), Cat 1 (11.6%), Cat 2 (12.6%), Cat 3 (7.4%), Cat 4 (10.5%), and Cat 5 (3.2%). This shows that while most storms remain at tropical storm intensity, a substantial fraction undergoes significant intensification to hurricane or typhoon strength. These statistics underscore the necessity of incorporating a wide range of storm intensities when training AI-based forecasting models, as both weaker and stronger storms exhibit distinct thermodynamic and dynamical signatures.

To provide a broader climatological perspective, Table 5 summarizes tropical cyclone counts by basin and intensity over the 1980–2024 period. Over these 45 years, a total of 4,500 named storms have been recorded globally, an average of roughly 100 storms per year. The Western Pacific and Eastern Pacific basins account for the largest shares (about 30% and 20%, respectively), followed by the North Atlantic. Globally, around 40% of all storms remain at tropical storm intensity, while approximately 35% strengthen to Category 1–2 hurricanes and about 25% reach major hurricane levels (Cat 3–5). These long-term statistics provide essential context for the 2017 dataset (Figure 6) and highlight the global diversity of tropical cyclone behavior. The inclusion of storms from multiple basins enables the dataset to capture a wide range of environmental conditions, from the warm, moisture-rich environments of the Western Pacific to the more shear-dominated Atlantic basin, thereby offering a robust foundation for machine learning models to understand and predict tropical cyclone formation, intensification, and evolution.

Figure 7 presents the interannual variability of tropical cyclone counts across global basins from 1980 to 2024. The global record (panel a) shows TC counts by years, with totals ranging between about 70 and 120 hurricanes annually. The Western Pacific remains the most active basin, followed by the Eastern Pacific and North Atlantic, each showing distinct variability linked to regional ocean–atmosphere conditions. The Indian Ocean basins generally produce fewer storms, while the Australian and South Pacific regions exhibit moderate and irregular activity. Overall, the figure highlights the large spatial and temporal

Table 5. Summary of global and basin-specific tropical cyclone counts by intensity, 1980–2024.

Basin	Total	Avg/Yr	TD	TS	Cat1	Cat2	Cat3	Cat4	Cat5
North Atlantic	706	15.7	110	293	126	52	48	56	21
Eastern Pacific	910	20.2	112	358	145	73	69	127	26
Western Pacific	1318	29.3	169	429	198	122	89	186	125
North Indian Ocean	242	5.4	9	146	35	10	14	19	9
South Indian Ocean	478	10.6	14	197	78	44	50	79	16
Australian Region	519	11.5	16	237	84	44	50	70	18
South Pacific	327	7.3	11	148	60	29	32	33	14
GLOBAL TOTAL	4500	100.0	441	1808	726	374	352	570	229



figures/hurricane_counts.jpeg

Figure 7. Annual tropical cyclone counts by basin for 1980–2024.

variability of tropical cyclone occurrence, with Pacific basins contributing the majority of global storms. This diverse record provides an essential foundation for training and evaluating machine learning models aimed at understanding global hurricane behavior.

Generation of Natural Language-based Weather Forecast Reports Dataset

While there have been many advances in data-driven weather and climate research and forecasting, there are still opportunities in leveraging AI to perform more human-oriented tasks. Further, natural language forecasts are often the primary way that a forecast is communicated to the public. Thus, natural language-based forecasts have a large impact on day-to-day life and play a crucial role in various applications including disasters and severe weather warnings, agricultural planning, and business and economic planning.

For enhanced forecast communication, broadcast meteorologists have been constantly exploring ways for generating interactive forecasts. Generating natural language weather discussions, however, combines computational linguistics with weather analysis, which is a difficult task. The Forecast Generator (FOG) based on the Forecast Production Assistant (FPA) uses rules and natural-language generator for converting weather maps into forecast text [75]. Similarly, GALiWeather [76], SUMTIME-MOUSAM [77], and other work [78], implement translating weather data into interactive reports and graphs. However, these frameworks lack standardization and generalizability in the implementation. With recent advances in large-language models (LLMs) and Vision-language models (VLMs), an AI-model can be used to generate forecast outputs after learning the lexical patterns and semantic present in the training data. The proposed novel dataset can be used to train such a model or fine-tune pre-trained FMs to generate textual weather forecast reports based on weather predictions.

Typically, the forecast parameters for generating forecast discussions include daily max and daily min temperatures (T_m), relative humidity (R_h), wind (u, v), cloud cover, probability of precipitation, and weather type. These parameters are used to develop grids of the expected weather conditions. The gridded forecasts are published to the local database and checked for quality. Subsequently, the human forecasters use specialized software to develop products such as the Zone Forecast Prediction, the Area Forecast Matrix, and the Point Forecast Matrix to develop the final textual forecast reports. These forecast reports are typically generated locally, state-wide, and country-wide.

The proposed dataset is focused on automating this process for generating end-to-end AI-based, natural language weather forecast reports.

Data Description Weather discussion reports involves textual descriptions of the weather, thus, leading to a multi-modal dataset. Table 6 shows certain quantitative stats for the proposed dataset.

To address the multi-modality present in the dataset due to text labels, we performed following pre-processing steps to the labels.

Table 6. Quantitative details of the proposed dataset, including subsetted atmospheric states and associated labels.

Number of samples	Temporal Span of the forecast labels	Geographical Extent
1249	August 2014 - December 2017	CONUS

- *Removing punctuation and special characters*: Punctuation characters are important for English grammar. However, they are avoided for text analysis. It is observed that special characters such as the newline character, $\backslash n$, are prominently present in the forecast reports. These characters do not enhance information in the context of weather conditions and are removed from the labels.
- *Expanding contractions* - Contractions are shortened representations of words such as *don't* for *do not*. The contractions are expanded to understand the full meaning of the text.
- *Unifying the word case* - Usually, the forecast reports contain the same words in different casing, such as *THUNDERSTORM* and *thunderstorm*. Since, from a natural language processing perspective, both words represent the same meaning, we unified the word case to lowercase in all the label forecast reports.
- *Removing numerals and stop words*: Numerals and stop words (words like *I, we, are, the*) in forecast reports decrease the importance of the domain-specific words. Additionally, stop words increase the dataset size and reduce the focus on informative words.

For this task, the High-Resolution Rapid Refresh (HRRR) is used as fine-tuning data along with weather reports as labels. The weather reports are scraped from the Storm Prediction Center's (SPC) archive provided by the National Oceanic and Atmospheric Administration (NOAA). These discussions are national forecasts of severe weather across CONUS.

Figure 8 illustrates the training data and corresponding forecast report. Each training data point, x_i , consists of an atmospheric analysis map whereas, weather report published on the same date as the corresponding label, y_i , where i represents an instance of the data sample. Naturally, generating weather forecast discussions is a challenging problem attributed to variability in natural language generation and finding accurate mapping with the corresponding atmospheric phenomenon.

Data Records

The WxC-Bench dataset is publicly available at [1] and the dataset repository hosted on <https://huggingface.co/datasets/nasa-impact/WxC-Bench> is organized by task directories as follows-

Table 7. Overview of the WxC-Bench data repository sub-directories and file types.

Sub-directory	ML Task	File type
aviation_turbulence	Aviation Turbulence Prediction	.nc
nonlocal_parameterization	Nonlocal Parameterization of Gravity Wave Momentum Flux	.nc
weather_analogs	Identifying Weather Analogs	.zip
hurricane	Hurricane Track and Intensity Prediction	.h5
weather_forecast_discussion	Generation of Natural Language Weather Forecasts	.grib2/ .csv
long_term_precipitation_forecast	Long-Term Precipitation Forecasting	.nc

Table 7 shows the overview of the WxC-Bench data directory. The task with .nc and .h5 files stores the data in variables. For identifying weather analogs task, a single .zip file stores the relevant data files. For generation of natural language weather forecast task, .grib2 files store the re-analysis data related to the weather forecasts whereas the corresponding captions are saved as .csv files.

Technical Validation

This section presents the technical validation of the datasets described above for baseline experiments, using both quantitative and qualitative metrics.

Aviation Turbulence Detection

The aviation turbulence detection dataset is validated by training an Artificial Neural Network (ANN) to perform a binary classification with two classes - turbulent and not-turbulent, for each of the LOW, MIDDLE, and HIGH levels. The trained ANN has 5 hidden layers of sizes 100, 60, 60, 40, and 20. For activation, Leaky ReLU is used within the hidden layers, and a sigmoid function is used for the classification layer. To mitigate the class imbalance, the loss is weighted by the class proportions. The ANN is evaluated based on the Overall Accuracy (OA), Probability of Detection (PoD) and False Alarm Rate (FAR) as performance metrics.

Table 8 presents the quantitative results based on the chosen performance metrics and Table 9 shows the comparative analysis with a baseline Random Forest (RF) model. The high overall accuracy with a lower probability of detection reflects the



Figure 8. (a) Word-frequency distribution in forecast reports. (b) Example training input showing surface temperature and (c) the corresponding forecast report for 31 December 2020, using HRRR 21Z analysis data.

Table 8. Performance of the ANN for aviation turbulence detection, evaluated using overall accuracy, probability of detection (PoD), false alarm rate (FAR), F1 score, and ROC AUC.

Level	Accuracy	PoD	FAR	F1	ROC AUC
LOW	80.0%	56.5%	17.8%	31.5%	0.73
MIDDLE	80.4%	38.9%	17.3%	17.5%	0.69
HIGH	80.6%	55.6%	17.6%	28.3%	0.71

Table 9. Performance comparison of the ANN and baseline Random Forest (RF) model for aviation turbulence detection, using overall accuracy, probability of detection (PoD), false alarm rate (FAR), F1 score, and ROC AUC.

Level	Model	Accuracy	PoD	FAR	F1	ROC AUC
LOW	Baseline	68.5%	32.1%	28.3%	18.2%	0.61
	ANN	80.0%	56.5%	17.8%	31.5%	0.73
MIDDLE	Baseline	66.2%	24.5%	29.7%	12.8%	0.58
	ANN	80.4%	38.9%	17.3%	17.5%	0.69
HIGH	Baseline	67.8%	31.8%	27.9%	17.1%	0.60
	ANN	80.6%	55.6%	17.6%	28.3%	0.71

large imbalances within the training dataset, as reports of turbulence are relatively rare compared to reports of non-turbulence. Due to the highly localized nature of turbulence compared to the relatively coarser grid spacing of MERRA-2, this downstream task is naturally expected to be challenging for deep learning.

Gravity Wave Momentum Fluxes

A proof-of-concept of the capability of DL models to learn and represent the nonlinear subgrid-scale GW evolution presents a new paradigm in their model parameterization development. This possibility is tested and demonstrated here using a baseline ML model. To test the AI-readiness of the data and to gauge the capability of ML models to learn from it, we used the WxC-Bench GW data to train an Attention-Unet [79] convolutional neural network.

figures/Gwflux_validation_200hpa.jpeg

Figure 9. Monthly mean normalized gravity-wave (GW) momentum flux at 200 hPa for May 2015. Panels (a) and (c) show the WxC-Bench zonal and meridional components, respectively, while panels (b) and (d) show the corresponding zonal and meridional fluxes predicted by a convolutional neural network trained on WxC-Bench. The model reproduces major stationary and transient hotspots as well as multi-scale features across latitudes, illustrating the potential of ML models to represent the nonlocal evolution of GWs (and potentially other mesoscale processes) in atmospheric models.

Figure 9 shows the CNN's performance in predicting the subgrid-scale fluxes when trained globally on the WxC-Bench GW momentum flux data. The ML model skillfully learns the global fluxes both over the land and the ocean. Over the land, the model learns fluxes over key hotspots including the Rocky mountains, Andes, and East Asian Mountains. Over the ocean, the mean predicted climatology captures the tropical convective fluxes associated with the ITCZ and the mid-latitude fluxes from storm tracks over the Southern Ocean. The predictability is notably higher in the mid-latitudes and over stationary structures like orography, with R^2 values of up to 0.6 in the mid-latitude troposphere vs. R^2 values of 0.3 in the tropical troposphere (see [80] for a more detailed analysis).

Current operational parameterizations struggle to accurately capture the global distribution of mesoscale momentum fluxes [81]. In principle, ML models such as this, trained on resolved GW fluxes provided by WxC-Bench, can be subsequently used to represent the learned subgrid-scale fluxes in coarser-climate models that cannot resolve these scales. The flux data can be further used to benchmark or compare various ML-driven GW parameterizations.

Weather Analog Search



Figure 10. Baseline model architecture for deep learning-based weather analog search.

Figure 10 shows the architecture for a convolutional encoder-decoder-based baseline model for implementing similarity search based on the MERRA-2 dataset. The model takes weather parameters (such as temperature, sea-level pressure, and wind vectors) as image inputs and searches for images closely representing similar conditions from the past based on cosine similarity. Currently, the technical validation experiment is designed to consider individual weather parameters as images for searching the archive for that particular parameter. In other words, the model can search over either temperature or surface pressure, but not both simultaneously. However, we envision a multi-parameter search for weather conditions by combining search results for individual parameters.

figures/weather_analog_updated.jpeg

ARTICLE IN PRESS

Figure 11. Top three temperature analog images retrieved for the query, showing their historical dates of occurrence and structural similarity (SSIM) scores relative to the query image.

The validation is done by training the encoder-decoder model for 100 epochs on input images of size 32x25. The learning rate is set to 0.001, and the batch size is set to 16 for training. Figure 11 presents the top five historical images with temperature patterns similar to the query image (Figure 11a), along with their structural similarity scores based on Structural Similarity Index (SSIM) [82]. The SSIM ranges from -1 to 1 , with 1 indicating identical structure, 0 no structural similarity, and -1 perfect anti-correlation. For this example query, the first three retrieved images match the query image ($SSIM > 0.5$).

Long-range Precipitation Forecasting

The baseline model for the precipitation forecast task is an auto-regressive forecast model based on a convolutional neural network. The architecture of the model is illustrated in Figure 12. The model input, which consists of the combined observations from the three satellite datasets from the eight days before the forecast initialization time, is first encoded spatially. The

spatial encoder encodes the input from each day independently using separate stems for each observations source. Each stem consists of a 3×3 convolution block with 64 output channels. The resulting stem-outputs from the three observation sources are then concatenated and fed into the shared spatial encoder, which encodes the combined observations from each day spatially, successively reducing the resolution by a factor of 4 in two stages. The spatially encoded observations are then concatenated and passed through two separate temporal encoders that map the encoded observations into the two initial latent forecast states. From these two states, successive latent forecast states are produced using a propagator module consisting of an encoder-decoder architecture that takes the states from two previous steps and maps them to a latent state representing the forecast for the following day. The predicted states are decoded using a shared decoder, which upsamples the forecasts back to the original resolution. Finally, a head consisting of a single point-wise convolution layer is used to produce precipitation forecast for all predicted forecast steps.

Table 10. Model configuration for the machine learning baseline of the long-range precipitation forecast. The table lists the stages for each component of the model. Each stage comprises one or multiple blocks, and the table lists their respective configuration and number of repetitions. k denotes the kernel size of the convolution operation, while b is the bottleneck width, and c the cardinality following [83].

Spatial encoder					
Stage	Input channels	Output channels	Up/Downsampling	Block type	# blocks
Head	24 / 48 / 14	64	-	Basic convolution, $k = 3$	1
Aggregator	192	128	-	Basic convolution, $k = 1, k = 1$	1
Stage 1	128	192	-	ResNeXt, $k = 3, b = 3, c = 32$	2
Stage 2	192	256	down, 2	ResNeXt, $k = 3, b = 4, c = 32$	2
Stage 3	256	256	down, 2	ResNeXt, $k = 3, b = 4, c = 32$	2
Temporal encoder					
Stage 1	2048	1024	-	ResNeXt, $k = 3, b = 32, c = 32$	1
Stage 2	1024	512	-	ResNeXt, $k = 3, b = 16, c = 32$	1
Propagator					
Stage 1	512	1024	-	ResNeXt, $k = 3, b = 32, c = 32$	3
Stage 2	1024	1024	down, 2	ResNeXt, $k = 3, b = 32, c = 32$	3
Stage 3	1024	1024	down, 2	ResNeXt, $k = 3, b = 32, c = 32$	3
Stage 4	1024	1024	up, 2	ResNeXt, $k = 3, b = 32, c = 32$	3
Stage 5	1024	512	up, 2	ResNeXt, $k = 3, b = 16, c = 32$	3
Decoder					
Stage 1	512	256	up, 2	ResNeXt, $k = 3, b = 8, c = 32$	3
Stage 2	256	128	up, 2	ResNeXt, $k = 3, b = 4, c = 32$	3
Head	128	32	-	Basic convolution, $k = 1$	1

All components of the model, except the stems and head, consist of ResNeXt [83] blocks. The exact configuration of components is provided in Table 10. Inputs to each convolution are padded to conserve the image size. Since the observations are global, periodic padding is applied along the longitudinal dimension, while reflect padding is applied along the latitudinal dimension.

The baseline is trained to predict 32 quantiles of the posterior distribution using a quantile loss function [84]. The training uses the AdamW optimizer [85] with an initial learning rate of 10^{-4} and a cosine annealing learning-rate schedule over 50 epochs. The model was trained to predict precipitation up to 16 days into the future. To make the model robust to changes in the input observations, individual input channels are dropped with a probability of 2% and all observations from a sensor are dropped with a probability of 5%.

The ML forecasts are compared to conventional NWP reference forecasts obtained from the S2S-database [86], which archives forecasts from various operational centers. For this evaluation, we consider only the forecasts from the ECMWF and the United Kingdom’s Met Office (UKMO) as these were found to be among the best-performing models in an evaluation of surface precipitation forecasts [87]. Examples of forecasts initialized on 10 January 2019 are shown in Figure 13. The ECMWF and UKMO forecasts yield good agreement with the reference estimates during the first week but yield a progressively larger error for the following weeks. Our machine learning-based baseline forecast also successfully captures the principal precipitation features during the first weeks but its predictions are generally less well-defined than those of the conventional models.

The results of a numerical evaluation of all forecasts from the evaluation period 1 January 2019 and 31 December 2022 are displayed in Figure 14. The forecasts are evaluated using the area-weighted forecast bias, linear correlation coefficient, and mean-squared error with respect to the reference precipitation estimates. In terms of bias, the ML-based model has a relatively



Figure 12. Architecture of the convolutional autoregressive forecasting model used as the machine learning baseline for the long-range precipitation forecast task.

small dry bias while the conventional models exhibit wet biases. The ML baseline forecasts is closest to the reference data, which is expected given that the precipitation estimates it was trained on were derived from the same precipitation product that is used for the evaluation. In terms of correlation coefficient, the conventional models yield higher correlations at shorter times but their advantage over the ML model decreases with longer lead times. The correlation coefficient of the conventional forecasts falls below that of the ML baseline after lead times exceeding 10 days but at this time the accuracy has decreased to or below that of a monthly precipitation climatology. Similar tendencies can be observed in the results of the mean-squared error.

Because machine-learning-based weather forecasts trained with pixel-wise losses tend to produce overly smooth fields, we further evaluate the spatial variability of the predictions using spherical-harmonic power spectra. The corresponding spectra for each of the four forecast weeks are shown in panel (d) of Figure 14. Although all models display a systematic underestimation of high-frequency variability, the ML forecasts show the strongest decay in spatial variability with increasing lead time. Whereas the power spectra of the two conventional baselines remain largely constant, the ML forecast exhibits a pronounced reduction in variability from week one to week two, followed by a continued but slower decline during weeks three and four.

Hurricane Forecasting based on FourCastNet

For validating the hurricane forecasting dataset, we gathered and analyzed a total of 1,459 named hurricanes that formed within the time frame of 1980 to 2020. This dataset includes 649 hurricanes from the Atlantic Ocean basin and 810 hurricanes from the Pacific Ocean basin.



figures/fig16_pcp_forecast.jpeg

Figure 13. Reference and predicted mean daily precipitation for forecasts initialized on 10 January 2019. The first row shows satellite-based reference precipitation estimates from the IMERG Final product. The second and third rows show corresponding precipitation forecasts from the control members of the ECMWF and UKMO S2S systems, respectively. The final row shows forecasts from the machine learning-based baseline model.




Figure 14. Skill of daily accumulated precipitation forecasts. Panel (a) shows forecast bias, defined as the area-weighted mean difference between predicted and reference precipitation. Panel (b) shows the area-weighted linear correlation coefficient. Panel (c) shows the area-weighted mean-squared error. Panel (d) shows spherical-harmonic power spectra for forecast weeks 1–4.

figures/Hurricane_correlation_rmse.jpeg

Figure 15. Time series of the correlation and root mean square error of wind speed (m/s), SLP (Pa), T2M (K) and Z500 (m) for the hurricane Michael (2018), evaluated over the hurricane Michael domain (105°W, 60°W, 20°N, 50°N).

To validate this task, FourCastNet has been used as the baseline model, which is a neural network architecture by combining elements from the Fourier Neural Operator (FNO) learning approach and the ViT architecture [2]. The model effectively addresses the computational challenges associated with high-resolution data by framing spatial token mixing as a continuous global convolution, which is efficiently implemented in the Fourier domain using Fast Fourier Transforms (FFTs). This innovative approach significantly reduces spatial mixing complexity, making FourCastNet well-suited for global forecasting.

Figure 15 shows an evaluation of the model's forecasting performance for key meteorological parameters over 5 days for Hurricane Michael initialized on 2018-10-08 00 UTC. Wind speed (WS) exhibits strong correlations, with values of 0.8 at a 2-day forecast, 0.7 at a 3-day forecast, and 0.4 at a 5-day forecast, indicating a significant decline in accuracy over the extended period. Sea level pressure (SLP) shows high correlations, particularly at shorter forecast intervals, with a remarkable 0.97 at a 2-day forecast, decreasing to 0.82 at 3 days, and further reducing to 0.3 at the 5-day mark. The temperature at 2 m above

ground level (T2M) displays diurnal variation in correlation, with daytime values peaking around 0.98 and nighttime values maintaining around 0.90. T2M outperforms other variables, consistently showcasing superior performance. Geopotential height at 500 hPa (Z500) is a reliable predictor, with correlations of 0.97 at 2 days, 0.92 at 3 days, and 0.75 at 4 days. The decreasing correlation values over the forecast period suggest a slow reduction in accuracy, highlighting the importance of considering shorter time frames for more reliable predictions. These findings collectively underscore the variable-dependent nature of the model's accuracy and emphasize the significance of temporal considerations. The time series correlation and root mean-squared error (RMSE) analysis show the nuanced forecasting capabilities of the model across different variables and forecast durations.



Figure 16. Spatial distribution of sea level pressure (Pa) for Hurricane Michael (2018) for 24H, 48H and 72H forecast from the ground truth, prediction, and the difference between the prediction and the ground truth.

The spatial distribution of model performance regarding the prediction of Hurricane Michael's intensity in terms of mean sea level pressure (MSLP) for 1-day, 2-day, and 5-day forecasts, as shown in Figure 16 from the MERRA trained model. The

RMSE for the 1-day forecast was 118.15 Pa, indicating a close alignment of the predicted hurricane center with the ground truth. However, for the 2-day and 3-day forecasts, the RMSE increased to 189.35 and 281.86 Pa, accompanied by correlation coefficients of 0.98 and 0.83, respectively.

Despite the spatial mean RMSE increasing to 281 for the 3-day forecast, it reveals that this error is most pronounced at the hurricane's center. This central discrepancy holds significant implications for hurricane prediction accuracy. The MSLP error for the 2-day and 3-day forecasts is quantified at 35 hPa and 52 hPa, respectively. While the tracking error remains minimal (below 40-50 km), the intensity error is notable, deviating by approximately 50-60 hPa.

Generation of Natural Language-based Weather Forecast Reports



Figure 17. High-level architecture for natural language weather forecast caption generation, based on an encoder–decoder framework with pre-trained vision and language components fine-tuned on the proposed dataset.

Figure 17 shows the high-level architecture for technical validation of the proposed dataset. The baseline experiment is performed by fine-tuning an encoder-decoder model for vision-language modeling. The encoder is based on the vision transformer (ViT) [88], whereas GPT-2 [89] is used as a decoder model for generating natural language-based forecast reports. The idea is to validate the proposed dataset as an ML-ready dataset capable of training on a multi-modal vision-language-based captioning problem. For the baseline model training, we have extracted the summary from the original forecast reports as labels for training instead of using the entire report. This summary is present in every forecast report under the section highlighted by "...SUMMARY...".

Table 11. Quantitative scores for the weather forecast report generation benchmark.

Date of the weather condition	ROUGE-L score
2017-03-30	0.30
2018-08-03	0.28
2019-08-30	0.42

We validate the model quantitatively based on the Recall-Oriented Understudy for Gisting Evaluation-Longest Common Subsequence (ROUGE-L) score. Table 11 highlights the quantitative evaluation of the proposed architecture based on the selected metrics. We observe a maximum ROUGE-L score of 0.42 across all the test samples. Typically, a ROUGE-L score of 0.4-0.6 is considered decent for image captioning task [90].

figures/nlp_discussion_updated.png




Figure 18. Performance of the natural language weather forecast captioning model trained on the proposed dataset. The top row shows a three-channel input stack (temperature and the u and v wind components), the second row shows the generated caption, and the third row shows the reference (original) caption.

Figure 18 shows the generated weather forecast reports along with the original reports. As evident from the results, the baseline model is capturing crucial information such as extreme events information (severe thunderstorms), geographical details (southern Plains and Ohio Valley), temporal details (this evening), catastrophic event details (damaging winds and hail), and relational details (across parts). However, comprehensive evaluation of the generated captions is being explored as a future work.

Discussion

Given the tremendous advancements unlocked by the AI-based FMs and LLMs in Natural Language Processing and Computer Vision, it seems reasonable to expect a similar impact within the earth and atmospheric sciences. In this context,

Machine Learning-ready datasets serve as a foundation for developing new models or fine-tuning existing models for downstream tasks. However, it is crucial that these datasets sample a wide range of realistic and relevant applications as well as data modalities to challenge AI model developers and demonstrate model capabilities to domain-specific scientists. The primary contribution of this paper - WxC-Bench, is a collection of high-quality datasets corresponding to a broad range of weather- and climate-related downstream tasks. WxC-Bench consists of datasets curated from satellite observations, reanalysis, high-resolution forecasting models, hurricane databases, and even pilot reports. Thus, the dataset can be used for benchmarking and testing the generalization of other specialized AI-models, and assess their ability to handle multi-modal data.

The selected six downstream tasks related to weather and atmospheric sciences are formulated as machine learning problems - classification, prediction, vector space search, and multi-modal language generation. The proposed WxC-Bench, thus, contains pre-processed ML-ready data for a range of atmospheric processes ranging from local scale (aviation turbulence) to synoptic scales (natural language forecasting). WxC-Bench is designed to be used for developing future weather and climate foundation models capable of working with multi-modal data, which can be used not only by researchers, but also by the broader scientific community. The openly-available data repository also promotes future additions of more downstream tasks to augment the capabilities of existing and new AI models trained on WxC-Bench.

Data availability

WxC-Bench is publicly available at <https://doi.org/10.57967/hf/7711> [1]. Additional details on the file formats and folder structure are provided in the Data Records section.

Code availability

The full codebase for dataset creation, pre-processing, and task-specific pipelines is openly available at: <https://github.com/NASA-IMPACT/WxC-Bench>. Additionally, a Python package providing programmatic access to WxC-Bench, along with helper utilities for loading and supporting documentation, is available on PyPI: <https://pypi.org/project/wxcbench/>.

References

1. Shinde, R. et al. Wxc-bench, <https://doi.org/10.57967/hf/7711> (2026).
2. Kurth, T. et al. Fourcastnet: Accelerating global high-resolution weather forecasting using adaptive fourier neural operators. In *Proceedings of the platform for advanced scientific computing conference*, 1–11, <https://doi.org/10.48550/arXiv.2208.05419> (2023).
3. Lam, R. et al. Graphcast: Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794* <https://doi.org/10.48550/arXiv.2212.12794> (2022).
4. Chen, K. et al. Fengwu: Pushing the skillful global medium-range weather forecast beyond 10 days lead. *arXiv preprint arXiv:2304.02948* <https://doi.org/10.48550/arXiv.2304.02948> (2023).
5. Bi, K. et al. Pangu-weather: A 3d high-resolution model for fast and accurate global weather forecast. *arXiv preprint arXiv:2211.02556* <https://doi.org/10.1038/s41586-023-06185-3> (2022).
6. Sønderby, C. K. et al. Metnet: A neural weather model for precipitation forecasting. *arXiv preprint arXiv:2003.12140* <https://doi.org/10.48550/arXiv.2003.12140> (2020).
7. Weyn, J. A., Durran, D. R., Caruana, R. & Cresswell-Clay, N. Sub-Seasonal Forecasting With a Large Ensemble of Deep-Learning Weather Prediction Models. *J. Adv. Model. Earth Syst.* **13**, e2021MS002502, [10.1029/2021MS002502](https://doi.org/10.1029/2021MS002502) (2021).
8. Mansfield, L. A. et al. Updates on Model Hierarchies for Understanding and Simulating the Climate System: A Focus on Data-Informed Methods and Climate Change Impacts. *J. Adv. Model. Earth Syst.* **15**, e2023MS003715, [10.1029/2023MS003715](https://doi.org/10.1029/2023MS003715) (2023).
9. Rasp, S., Pritchard, M. S. & Gentine, P. Deep learning to represent subgrid processes in climate models. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9684–9689, [10.1073/pnas.1810286115](https://doi.org/10.1073/pnas.1810286115) (2018).
10. Zanna, L. & Bolton, T. Data-Driven Equation Discovery of Ocean Mesoscale Closures. *Geophys. Res. Lett.* **47**, e2020GL088376, [10.1029/2020GL088376](https://doi.org/10.1029/2020GL088376) (2020).

11. Espinosa, Z. I., Sheshadri, A., Cain, G. R., Gerber, E. P. & DallaSanta, K. A Deep Learning Parameterization of Gravity Wave Drag Coupled to an Atmospheric Global Climate Model. *Geophys. Res. Lett.* (submitted, 2021).
12. Wang, P., Yuval, J. & O’Gorman, P. A. Non-Local Parameterization of Atmospheric Subgrid Processes With Neural Networks. *J. Adv. Model. Earth Syst.* **14**, e2022MS002984, [10.1029/2022MS002984](https://doi.org/10.1029/2022MS002984) (2022).
13. Bretherton, C. S. et al. Correcting Coarse-Grid Weather and Climate Models by Machine Learning From Global Storm-Resolving Simulations. *J. Adv. Model. Earth Syst.* **14**, e2021MS002794, [10.1029/2021MS002794](https://doi.org/10.1029/2021MS002794) (2022).
14. Davenport, F. V. & Diffenbaugh, N. S. Using Machine Learning to Analyze Physical Causes of Climate Change: A Case Study of U.S. Midwest Extreme Precipitation. *Geophys. Res. Lett.* **48**, e2021GL093787, [10.1029/2021GL093787](https://doi.org/10.1029/2021GL093787) (2021).
15. Diffenbaugh, N. S. & Barnes, E. A. Data-driven predictions of the time remaining until critical global warming thresholds are reached. *Proc. Natl. Acad. Sci.* **120**, e2207183120, [10.1073/pnas.2207183120](https://doi.org/10.1073/pnas.2207183120) (2023).
16. Roy, S. et al. Clifford neural operators on atmospheric data influenced partial differential equations. In *12th International Conference on Learning Representations* (2024).
17. Schmude, J. et al. Prithvi wxc: Foundation model for weather and climate. *arXiv preprint arXiv:2409.13598* (2024).
18. Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J. K. & Grover, A. ClimaX: A foundation model for weather and climate. *arXiv* (2023). [2301.10343](https://arxiv.org/abs/2301.10343).
19. Bodnar, C. et al. Aurora: A foundation model of the atmosphere. *arXiv preprint arXiv:2405.13063* (2024).
20. Szwarcman, D. et al. Prithvi-eo-2.0: A versatile multi-temporal foundation model for earth observation applications. *IEEE Transactions on Geosci. Remote. Sens.* <https://doi.org/10.1109/TGRS.2025.3642610> (2025).
21. Bommasani, R. et al. On the opportunities and risks of foundation models (2022). [2108.07258](https://arxiv.org/abs/2108.07258).
22. Rasp, S. et al. Weatherbench: a benchmark data set for data-driven weather forecasting. *J. Adv. Model. Earth Syst.* **12**, <https://doi.org/10.1029/2020MS002203> (2020).
23. Rasp, S. et al. Weatherbench 2: A benchmark for the next generation of data-driven global weather models (2024). [2308.15560](https://arxiv.org/abs/2308.15560).
24. Prabhat et al. Climateset: An expert-labelled open dataset and deep learning architecture for enabling high-precision analyses of extreme weather. *Geosci. Model. Dev. Discuss.* **2020**, 1–28, <https://doi.org/10.5194/gmd-14-107-2021> (2020).
25. Watson-Parris, D. et al. Climatebench v1.0: A benchmark for data-driven climate projections. *J. Adv. Model. Earth Syst.* **14**, e2021MS002954, <https://doi.org/10.1029/2021MS002954> (2022). E2021MS002954 2021MS002954, <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2021MS002954>.
26. Yu, S. et al. Climsim: A large multi-scale dataset for hybrid physics-ml climate emulation. In Oh, A. et al. (eds.) *Advances in Neural Information Processing Systems*, vol. 36, 22070–22084 (2023).
27. Nguyen, T., Jewik, J., Bansal, H., Sharma, P. & Grover, A. Climatelearn: Benchmarking machine learning for weather and climate modeling. In Oh, A. et al. (eds.) *Advances in Neural Information Processing Systems*, vol. 36, 75009–75025 (Curran Associates, Inc., 2023).
28. Shinde, R. et al. Windset: Weather insights and novel data for systematic evaluation and testing. In *The Twelfth International Conference on Learning Representations (ICLR)* (2024).
29. Ito, J., Niino, H. & Yoshino, K. Large Eddy Simulation on Horizontal Convective Rolls that Caused an Aircraft Accident during its Landing at Narita Airport. *Geophys. Res. Lett.* **47**, [10.1029/2020gl086999](https://doi.org/10.1029/2020gl086999) (2020).
30. Golding, W. Turbulence and Its Impact on Commercial Aviation. *J. Aviat. Educ. & Res.* [10.15394/jaaer.2002.1301](https://doi.org/10.15394/jaaer.2002.1301) (2000).
31. Emara, M. et al. Machine Learning Enabled Turbulence Prediction using Flight Data for Safety Analysis. In *32nd Congress of the International Council of the Aeronautical Sciences* (Shanghai, China, 2021).
32. Williams, J. K. Using random forests to diagnose aviation turbulence. *Mach. Learn.* **95**, 51–70, [10.1007/s10994-013-5346-7](https://doi.org/10.1007/s10994-013-5346-7) (2014).

33. Hon, K. K., Ng, C. W. & Chan, P. W. Machine learning based multi-index prediction of aviation turbulence over the Asia-Pacific. *Mach. Learn. with Appl.* **2**, 100008, [10.1016/j.mlwa.2020.100008](https://doi.org/10.1016/j.mlwa.2020.100008) (2020).
34. Gelaro, R. et al. The Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2). *J. Clim.* **30**, 5419–5454, [10.1175/jcli-d-16-0758.1](https://doi.org/10.1175/jcli-d-16-0758.1) (2017).
35. Iowa Environmental Mesonet (IEM) Download PIREPs, <https://mesonet.agron.iastate.edu/request/gis/pireps.php>.
36. Wolff, J. K. & Sharman, R. D. Climatology of Upper-Level Turbulence over the Contiguous United States. *J. Appl. Meteorol. Climatol.* **47**, 2198–2214, [10.1175/2008jamc1799.1](https://doi.org/10.1175/2008jamc1799.1) (2008).
37. Fritts, D. C. & Alexander, M. J. Gravity wave dynamics and effects in the middle atmosphere. *Rev. Geophys.* **41**, [10.1029/2001RG000106](https://doi.org/10.1029/2001RG000106) (2003).
38. Kim, Y.-J., Eckermann, S. D. & Chun, H.-Y. An overview of the past, present and future of gravity-wave drag parametrization for numerical climate and weather prediction models. *Atmosphere-Ocean* **41**, 65–98, [10.3137/ao.410105](https://doi.org/10.3137/ao.410105) (2003).
39. Achatz, U. et al. Atmospheric Gravity Waves: Processes and Parameterization. *J. Atmospheric Sci.* **-1**, [10.1175/JAS-D-23-0210.1](https://doi.org/10.1175/JAS-D-23-0210.1) (2023).
40. Sorbjan, Z. Improving Non-local Parameterization of the Convective Boundary Layer. *Boundary-Layer Meteorol* **130**, 57–69, [10.1007/s10546-008-9331-9](https://doi.org/10.1007/s10546-008-9331-9) (2009).
41. Chen, T.-C., Yau, M. K. & Kirshbaum, D. J. Assessment of Conditional Symmetric Instability from Global Reanalysis Data. *J. Atmospheric Sci.* **75**, 2425–2443, [10.1175/JAS-D-17-0221.1](https://doi.org/10.1175/JAS-D-17-0221.1) (2018).
42. Plougonven, R., de la Cámara, A., Hertzog, A. & Lott, F. How does knowledge of atmospheric gravity waves guide their parameterizations? *Q. J. Royal Meteorol. Soc.* **146**, 1529–1543, [10.1002/qj.3732](https://doi.org/10.1002/qj.3732) (2020).
43. McLandress, C., Scinocca, J. F., Shepherd, T. G., Reader, M. C. & Manney, G. L. Dynamical Control of the Mesosphere by Orographic and Nonorographic Gravity Wave Drag during the Extended Northern Winters of 2006 and 2009. *J. Atmos. Sci.* **70**, 2152–2169, [10.1175/JAS-D-12-0297.1](https://doi.org/10.1175/JAS-D-12-0297.1) (2012).
44. de la Cámara, A., Abalos, M. & Hitchcock, P. Changes in Stratospheric Transport and Mixing During Sudden Stratospheric Warmings. *J. Geophys. Res. Atmospheres* **123**, 3356–3373, [10.1002/2017JD028007](https://doi.org/10.1002/2017JD028007) (2018).
45. Gupta, A. et al. Estimates of Southern Hemispheric Gravity Wave Momentum Fluxes Across Observations, Reanalyses, and Kilometer-scale Numerical Weather Prediction Model. *J. Atmospheric Sci.* **-1**, [10.1175/JAS-D-23-0095.1](https://doi.org/10.1175/JAS-D-23-0095.1) (2024).
46. Kruse, C. G. et al. Observed and Modeled Mountain Waves from the Surface to the Mesosphere near the Drake Passage. *J. Atmospheric Sci.* **79**, 909–932, [10.1175/JAS-D-21-0252.1](https://doi.org/10.1175/JAS-D-21-0252.1) (2022).
47. Kim, Y.-H., Voelker, G. S., Bölöni, G., Zängl, G. & Achatz, U. Crucial role of obliquely propagating gravity waves in the quasi-biennial oscillation dynamics. *Atmospheric Chem. Phys.* **24**, 3297–3308, [10.5194/acp-24-3297-2024](https://doi.org/10.5194/acp-24-3297-2024) (2024).
48. Gupta, A., Sheshadri, A., Alexander, M. J. & Birner, T. Insights on Lateral Gravity Wave Propagation in the Extratropical Stratosphere from 44 Years of ERA5 Data. *Geophys. Res. Lett.* [10.1029/2024GL108541](https://doi.org/10.1029/2024GL108541) (2024).
49. Lindborg, E. A Helmholtz decomposition of structure functions and spectra calculated from aircraft data. *J. Fluid Mech.* **762**, R4, [10.1017/jfm.2014.685](https://doi.org/10.1017/jfm.2014.685) (2015).
50. Hersbach, H. et al. The ERA5 global reanalysis. *Q. J. Royal Meteorol. Soc.* **146**, 1999–2049, [10.1002/qj.3803](https://doi.org/10.1002/qj.3803) (2020).
51. Skamarock, W. C. Evaluating Mesoscale NWP Models Using Kinetic Energy Spectra. *Mon. Weather. Rev.* **132**, 3019–3032, [10.1175/MWR2830.1](https://doi.org/10.1175/MWR2830.1) (2004).
52. Gupta, A., Birner, T., Dörnbrack, A. & Polichtchouk, I. Importance of Gravity Wave Forcing for Springtime Southern Polar Vortex Breakdown as Revealed by ERA5. *Geophys. Res. Lett.* **48**, e2021GL092762, [10.1029/2021GL092762](https://doi.org/10.1029/2021GL092762) (2021).
53. Pahlavan, H. A., Fu, Q., Wallace, J. M. & Kiladis, G. N. Revisiting the Quasi-Biennial Oscillation as Seen in ERA5. Part I: Description and Momentum Budget. *J. Atmospheric Sci.* **78**, 673–691, [10.1175/JAS-D-20-0248.1](https://doi.org/10.1175/JAS-D-20-0248.1) (2021).

54. Chattopadhyay, A., Nabizadeh, E. & Hassanzadeh, P. Analog forecasting of extreme-causing weather patterns using deep learning. *J. Adv. Model. Earth Syst.* **12**, [10.1029/2019ms001958](https://doi.org/10.1029/2019ms001958) (2020).
55. Dool, H. M. v. d. A New Look at Weather Forecasting through Analogues. *Mon. Weather. Rev.* **117**, 2230–2247, [10.1175/1520-0493\(1989\)117<2230:anlawf>2.0.co;2](https://doi.org/10.1175/1520-0493(1989)117<2230:anlawf>2.0.co;2) (1989).
56. Yang, D. & Alessandrini, S. An ultra-fast way of searching weather analogs for renewable energy forecasting. *Sol. Energy* **185**, 255–261, <https://doi.org/10.1016/j.solener.2019.03.068> (2019).
57. Franch, G., Jurman, G., Coviello, L., Pendesini, M. & Furlanello, C. Mass-umap: Fast and accurate analog ensemble search in weather radar archives. *Remote. Sens.* **11**, [10.3390/rs11242922](https://doi.org/10.3390/rs11242922) (2019).
58. Ahn, H. et al. Searching similar weather maps using convolutional autoencoder and satellite images. *ICT Express* **9**, 69–75, <https://doi.org/10.1016/j.ict.2022.03.013> (2023).
59. Raoult, B., Di Fatta, G., Pappenberger, F. & Lawrence, B. Fast retrieval of weather analogues in a multi-petabytes archive using wavelet-based fingerprints. In Shi, Y. et al. (eds.) *Computational Science – ICCS 2018*, 697–710 (Springer International Publishing, Cham, 2018).
60. Demir, B. & Bruzzone, L. Hashing-based scalable remote sensing image search and retrieval in large archives. *IEEE Transactions on Geosci. Remote. Sens.* **54**, 892–904, [10.1109/TGRS.2015.2469138](https://doi.org/10.1109/TGRS.2015.2469138) (2016).
61. Hu, W., Cervone, G., Young, G. & Delle Monache, L. Machine learning weather analogs for near-surface variables. *Boundary-Layer Meteorol.* **186**, 711–735, <https://doi.org/10.1007/s10546-022-00779-6> (2023).
62. Vitart, F. et al. Outcomes of the wmo prize challenge to improve subseasonal to seasonal predictions using artificial intelligence. *Bull. Am. Meteorol. Soc.* **103**, E2878 – E2886, [10.1175/BAMS-D-22-0046.1](https://doi.org/10.1175/BAMS-D-22-0046.1) (2022).
63. Robertson, A. W., Vitart, F. & Camargo, S. J. Subseasonal to seasonal prediction of weather to climate with application to tropical cyclones. *J. Geophys. Res. Atmospheres* **125**, e2018JD029375, <https://doi.org/10.1029/2018JD029375> (2020). E2018JD029375 [10.1029/2018JD029375](https://doi.org/10.1029/2018JD029375), <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2018JD029375>.
64. Vitart, F., Robertson, A. W. & Anderson, D. Subseasonal to seasonal prediction project: bridging the gap between weather and climate (2012).
65. de Witt, C. S. et al. Rainbench: Towards global precipitation forecasting from satellite imagery (2020). [2012.09670](https://doi.org/10.21203/rs.3.rs-109670/v1).
66. Knapp, K. R. et al. Globally gridded satellite observations for climate studies. *Bull. Am. Meteorol. Soc.* **92**, 893 – 907, [10.1175/2011BAMS3039.1](https://doi.org/10.1175/2011BAMS3039.1) (2011).
67. Foster, M. J., Philips, C., Heidinger, A. K. & Program, N. C. Noaa climate data record (cdr) of advanced very high resolution radiometer (avhrr) and high-resolution infra-red sounder (hirs) reflectance, brightness temperature, and cloud products from pathfinder atmospheres - extended (patmos-x), version 6.0. <https://doi.org/10.7289/V5X9287> (2021).
68. Kummerow, C. D., Berg, W. K., Kuo, C.-P., & Program, N. C. Noaa climate data record (cdr) of advanced very high resolution radiometer (avhrr) and high-resolution infra-red sounder (hirs) reflectance, brightness temperature, and cloud products from pathfinder atmospheres - extended (patmos-x), version 6.0. <https://doi.org/10.25921/E3K5-BW77> (2022).
69. Ashouri, H. et al. Persiann-cdr: Daily precipitation climate data record from multisatellite observations for hydrological and climate studies. *Bull. Am. Meteorol. Soc.* **96**, 69–83 (2015).
70. Huffman, G., Stocker, E., Bolvin, D., Nelkin, E. & Jackson, T. Gpm imerg final precipitation l3 1 day 0.1 degree x 0.1 degree v07, greenbelt, md, goddard earth sciences data and information services center (ges disc).[accessed 2023 sep 14]. DOI: <https://doi.org/10.5067/GPM/IMERGDF/DAY/07> (2023).
71. Emanuel, K. Assessing the present and future probability of hurricane harvey’s rainfall. *Proc. Natl. Acad. Sci.* **114**, 12681–12684, <https://doi.org/10.1073/pnas.1716222114> (2017).
72. Bromirski, P. D. & Kossin, J. P. Increasing hurricane wave power along the us atlantic and gulf coasts. *J. Geophys. Res. Ocean.* **113**, <https://doi.org/10.1029/2007JC004706> (2008).
73. Emanuel, K. Increasing destructiveness of tropical cyclones over the past 30 years. *Nat.* **436**, 686–688, <https://doi.org/10.1038/nature03906> (2005).

74. Burg, T. & Lillo, S. P. Tropicool: A python package for analyzing tropical cyclones and more. In *34th Conference on Hurricanes and Tropical Meteorology* (AMS, 2021).
75. Goldberg, E., Driedger, N. & Kittredge, R. Using natural-language processing to produce weather forecasts. *IEEE Expert*, **9**, 45–53, [10.1109/64.294135](https://doi.org/10.1109/64.294135) (1994).
76. Ramos-Soto, A., Bugarín, A. J., Barro, S. & Taboada, J. Linguistic descriptions for automatic generation of textual short-term weather forecasts on real prediction data. *IEEE Transactions on Fuzzy Syst.* **23**, 44–57, [10.1109/TFUZZ.2014.2328011](https://doi.org/10.1109/TFUZZ.2014.2328011) (2015).
77. Reiter, E., Sripada, S., Hunter, J., Yu, J. & Davy, I. Choosing words in computer-generated weather forecasts. *Artif. Intell.* **167**, 137–169, <https://doi.org/10.1016/j.artint.2005.06.006> (2005).
78. Zhang, H.-P., Wu, H.-P., Gao, J., Zhao, Y.-P. & Lv, Z.-L. Meteorological bulletin automatic generation based on spatio-temporal reasoning. In *2011 International Conference on Machine Learning and Cybernetics*, vol. 4, 1927–1931, <https://doi.org/10.1109/ICMLC.2011.6016952> (IEEE, 2011).
79. Oktay, O. et al. Attention U-Net: Learning Where to Look for the Pancreas, [10.48550/arXiv.1804.03999](https://arxiv.org/abs/1804.03999) (2018). [1804.03999](https://arxiv.org/abs/1804.03999).
80. Gupta, A. et al. Machine Learning Global Simulation of Nonlocal Gravity Wave Propagation (2024). [2406.14775](https://arxiv.org/abs/2406.14775).
81. Geller, M. A. et al. A Comparison between Gravity Wave Momentum Fluxes in Observations and Climate Models. *J. Clim.* **26**, 6383–6405, [10.1175/JCLI-D-12-00545.1](https://doi.org/10.1175/JCLI-D-12-00545.1) (2013).
82. Bouzerdoum, A., Havstad, A. & Beghdadi, A. Image quality assessment using a neural network approach. In *Proceedings of the Fourth IEEE International Symposium on Signal Processing and Information Technology, 2004.*, 330–333, [10.1109/ISSPIT.2004.1433751](https://doi.org/10.1109/ISSPIT.2004.1433751) (2004).
83. Xie, S., Girshick, R., Dollar, P., Tu, Z. & He, K. Aggregated Residual Transformations for Deep Neural Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5987–5995, [10.1109/CVPR.2017.634](https://doi.org/10.1109/CVPR.2017.634) (IEEE, Honolulu, HI, 2017).
84. Pfreundschuh, S. et al. A neural network approach to estimating a posteriori distributions of bayesian retrieval problems. *Atmospheric Meas. Tech.* **11**, 4627–4643, [10.5194/amt-11-4627-2018](https://doi.org/10.5194/amt-11-4627-2018) (2018).
85. Loshchilov, I. & Hutter, F. Decoupled weight decay regularization (2019). [1711.05101](https://arxiv.org/abs/1711.05101).
86. Vitart, F. et al. The subseasonal to seasonal (s2s) prediction project database. *Bull. Am. Meteorol. Soc.* **98**, 163 – 173, [10.1175/BAMS-D-16-0017.1](https://doi.org/10.1175/BAMS-D-16-0017.1) (2017).
87. de Andrade, F. M., Coelho, C. A. & Cavalcanti, I. F. Global precipitation hindcast quality assessment of the subseasonal to seasonal (s2s) prediction project models. *Clim. Dyn.* **52**, 5451–5475, <https://doi.org/10.1007/s00382-018-4457-z> (2019).
88. Dosovitskiy, A. et al. An image is worth 16x16 words: Transformers for image recognition at scale (2021). [2010.11929](https://arxiv.org/abs/2010.11929).
89. Radford, A. et al. Language models are unsupervised multitask learners (2019).
90. Lin, C.-Y. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, 74–81 (Association for Computational Linguistics, Barcelona, Spain, 2004).

Acknowledgments

This work was supported by NASA’s Office of Chief Science Data Officer and Earth Science Division’s Earth Science Scientific Computing, Earth Science Data Systems Program, and the Earth Science Modeling and Analysis Program.

The long-term precipitation forecasting task uses S2S. S2S is a joint initiative of the World Weather Research Programme (WWRP) and the World Climate Research Programme (WCRP). The original S2S database is hosted at ECMWF as an extension of the TIGGE database. AG and AS were supported by Schmidt Sciences, LLC, a philanthropic initiative founded by Eric and Wendy Schmidt, as part of the Virtual Earth System Research Institute (VESRI). AS acknowledges support from the National Science Foundation through grant OAC-2004492. RS thanks Prajun Trital for working on the WxC-Bench Python Package during their internship. The authors thank the anonymous reviewers for their constructive feedback and Dr. Alireza Foroozani for editorial oversight and guidance throughout the review process.

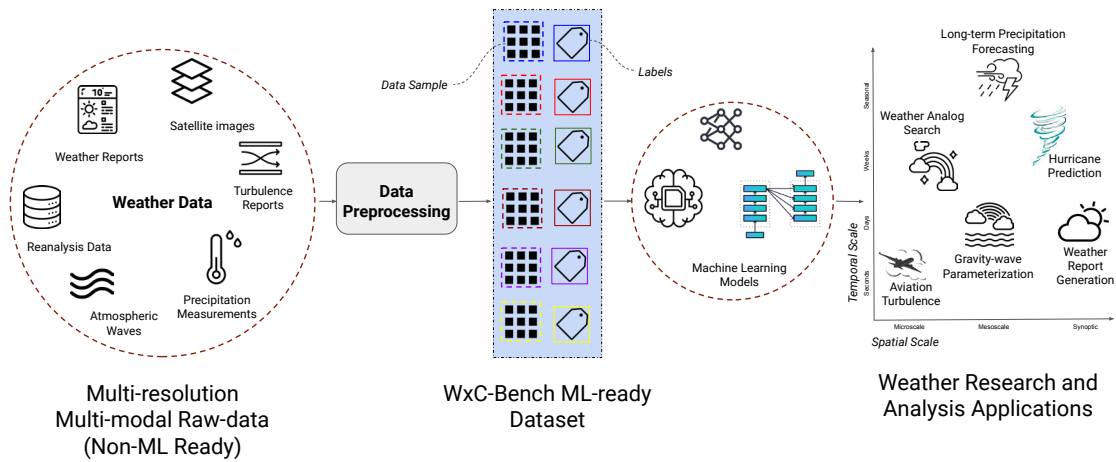
Author Contributions

R.S., C.E.P., K.A., S.R., A.G., S.P. compiled the manuscript. A.G., S.R., V.G., A.S. conceptualized the technical validation for Nonlocal Parameterization of Gravity Wave Momentum Flux task. S.P. performed the technical validation for Long-Term Precipitation Forecasting task. K.A. performed the technical validation for the Hurricane Track and Intensity Prediction task. R.S., S.K., and C.E.P. performed the technical validation for the Aviation Turbulence Prediction and Weather Analog Search task. R.S., K.A. curated the data whereas R.S. conceptualized the technical validation for the Generation of Natural Language Weather Forecasts task. V.K. and P.T. supported in development of supplementary materials and Python package for the final submission. A.L. validated the data files whereas U.N. supported in selection of downstream tasks. M.M., R.R. supervised the work and reviewed the manuscript. All the authors reviewed the manuscript.

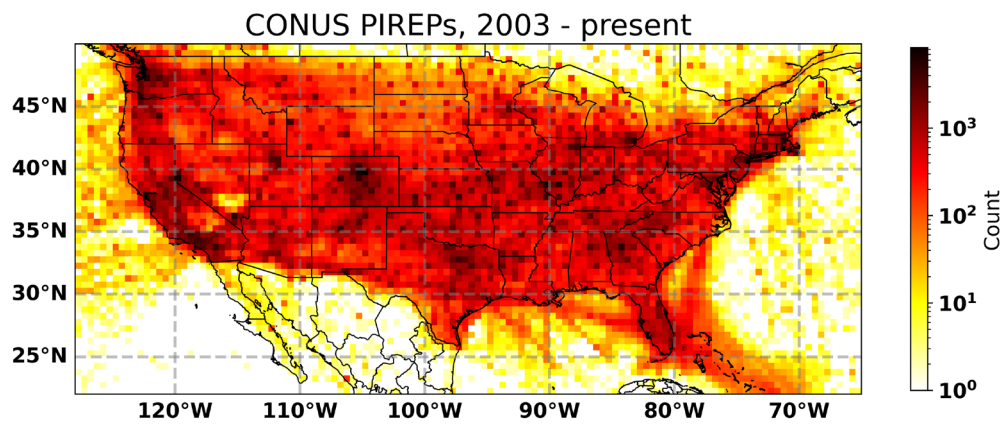
Competing interests


The author declares no competing interests.

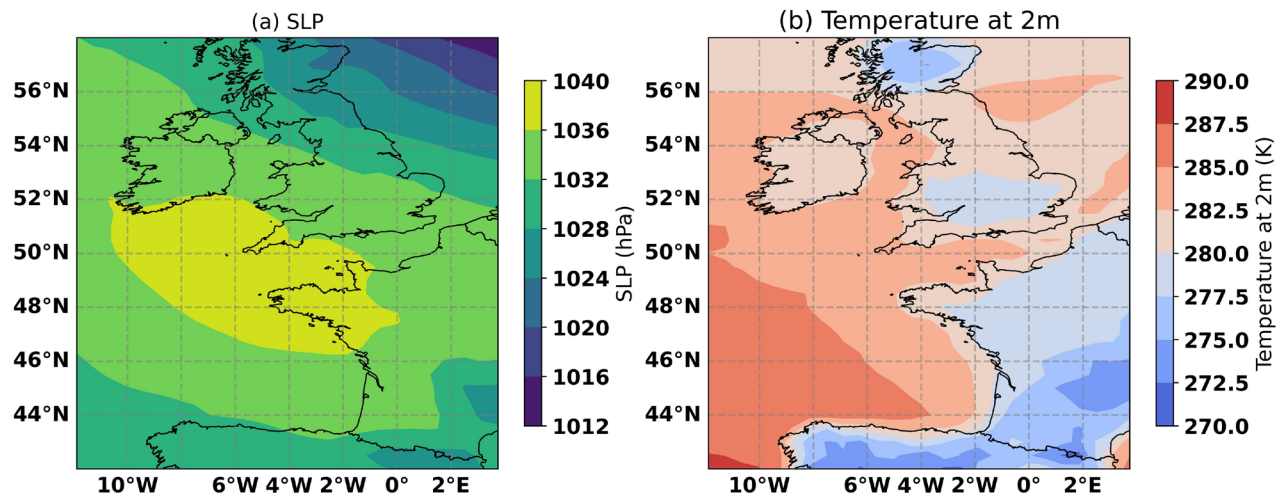
ARTICLE IN PRESS



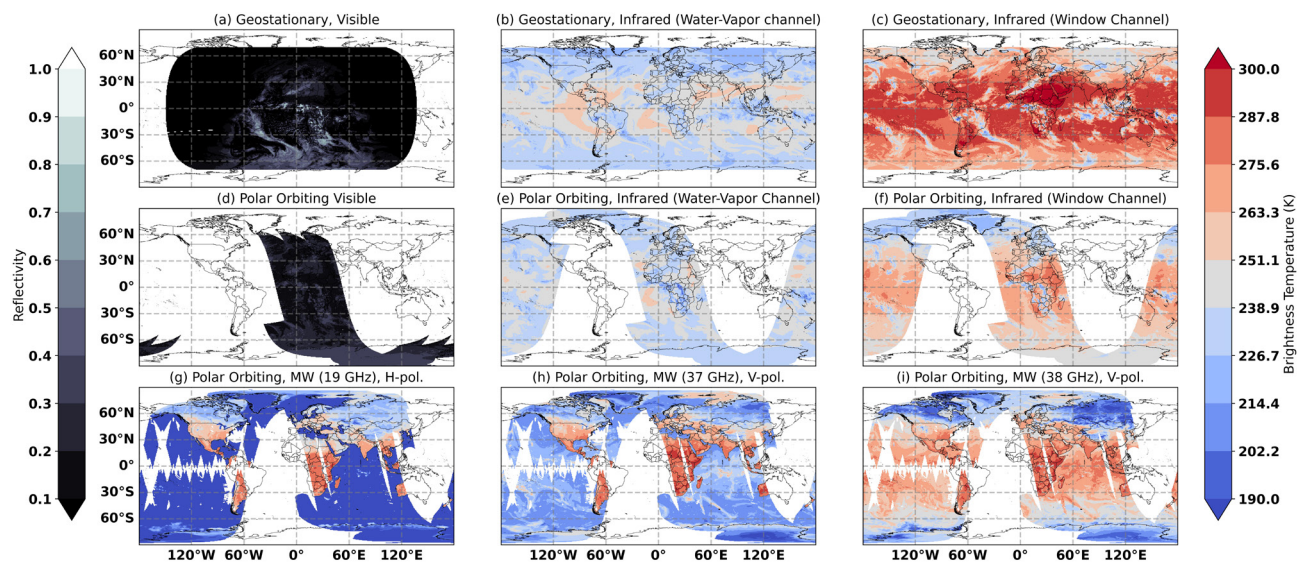
ARTICLE

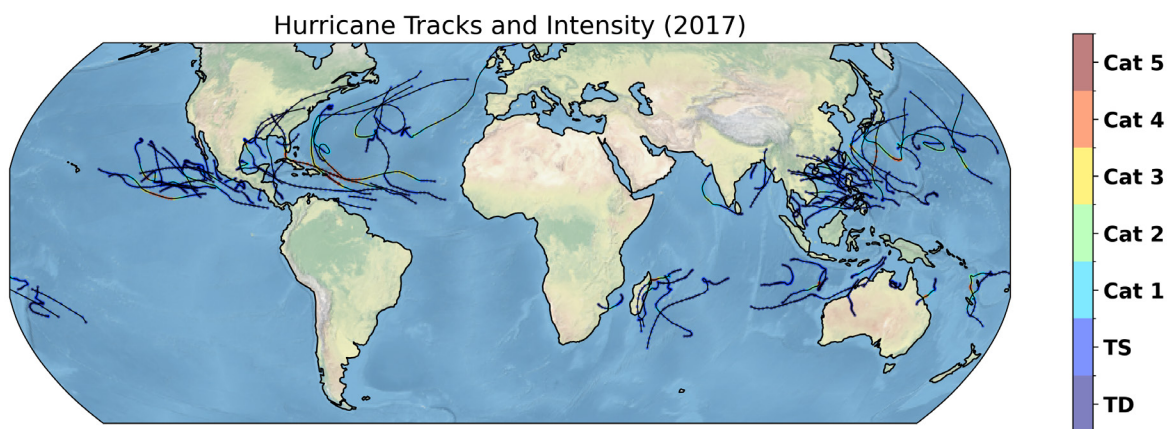








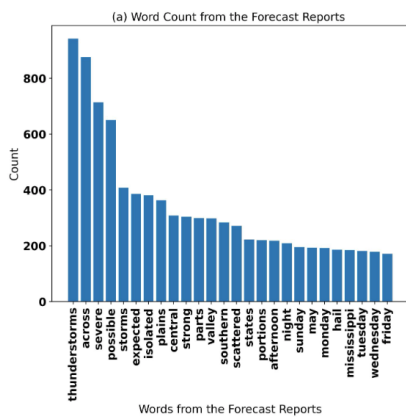
ARTICLE





ARTICLE





Day 1 Convective Outlook AMEND 1
 NWS Storm Prediction Center Norman OK
 0243 PM CST Thu Dec 31 2020

Valid 312000Z - 011200Z

...THERE IS AN ENHANCED RISK OF SEVERE THUNDERSTORMS FROM THE SABINE RIVER VALLEY AREAS EASTWARD/NORTHEASTWARD ACROSS THE CENTRAL GULF COAST STATES...

AMENDED TO ADD MRGL RISK AREA TO A PORTION OF THE EASTERN CAROLINAS

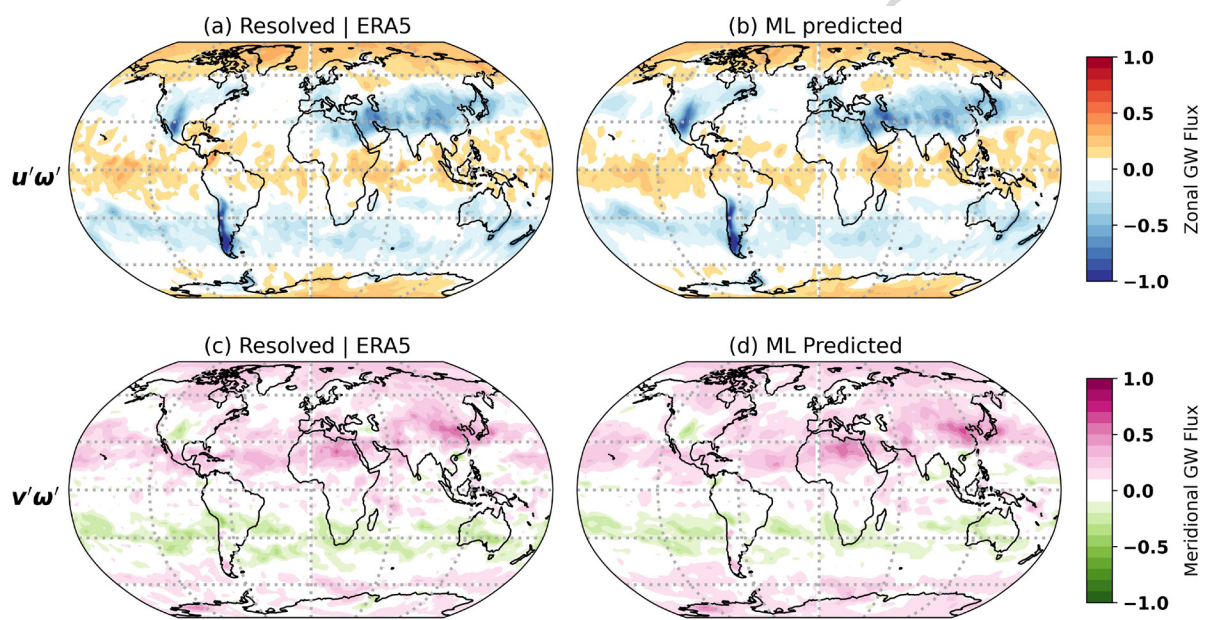
...SUMMARY...

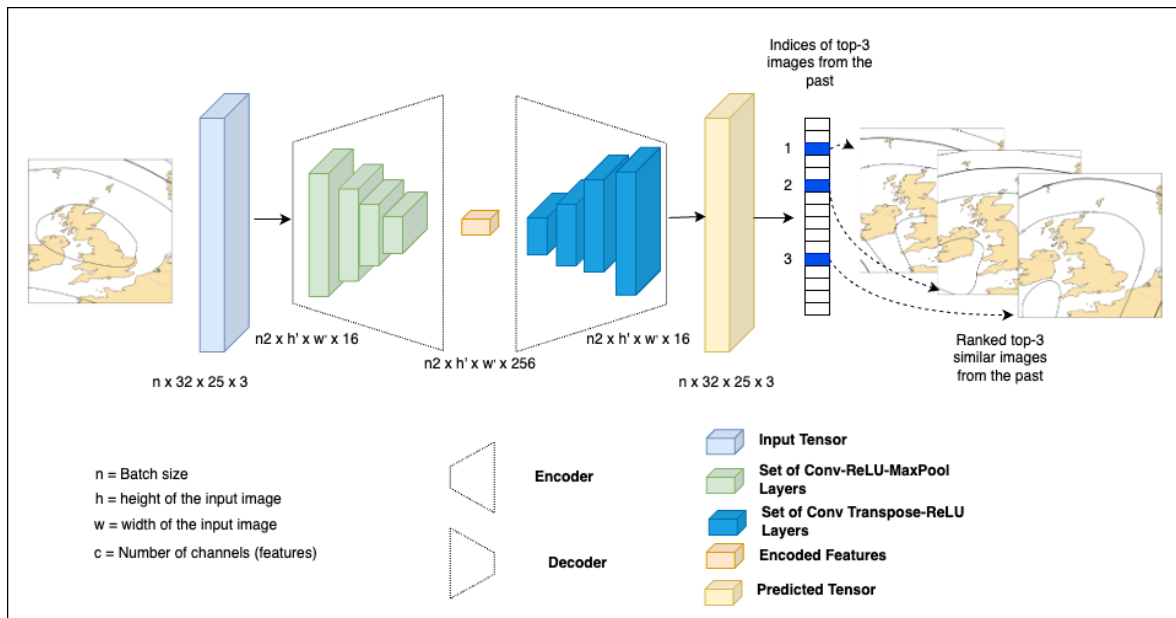
Thunderstorms -- a few of which will likely be severe, with tornadoes and damaging winds as the primary hazards -- will continue from the Sabine River Valley vicinity across southwest Louisiana through this afternoon, and then will expand across the rest of Louisiana into Mississippi and Alabama this evening and tonight.

...Discussion...

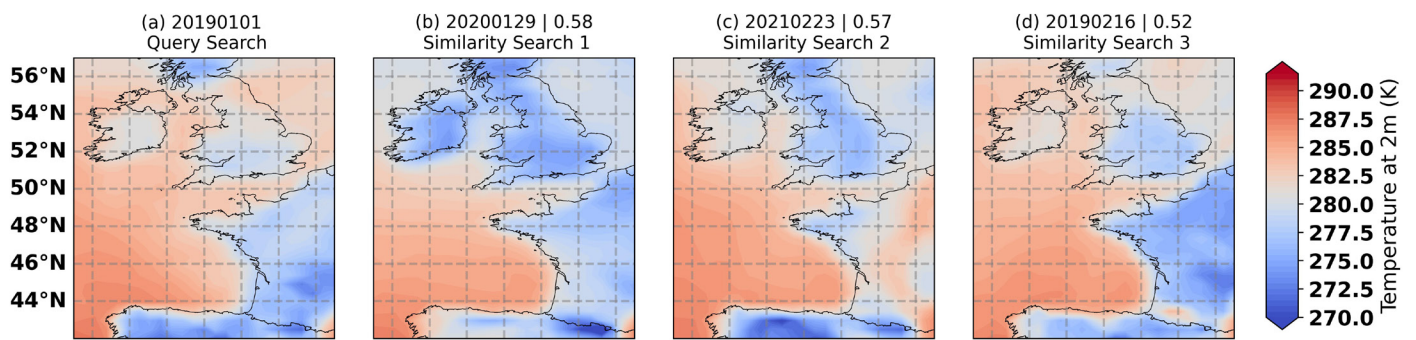
Thunderstorms continue across portions of east Texas at this time, and will continue spreading into western Louisiana over the next few hours. While severe risk has remained muted thus far, the environment downstream from the ongoing storms has become steadily more supportive of severe/rotating storms -- both thermodynamically and kinematically -- over the past few hours, as shown by changes in the LCH RAOBs from 12Z to 18Z.

Thus, the ongoing outlook reasoning and areas appear to remain representative of the evolving situation, with the only changes which appear necessary at this time being to trim western fringes of the areas to account for convective advance.

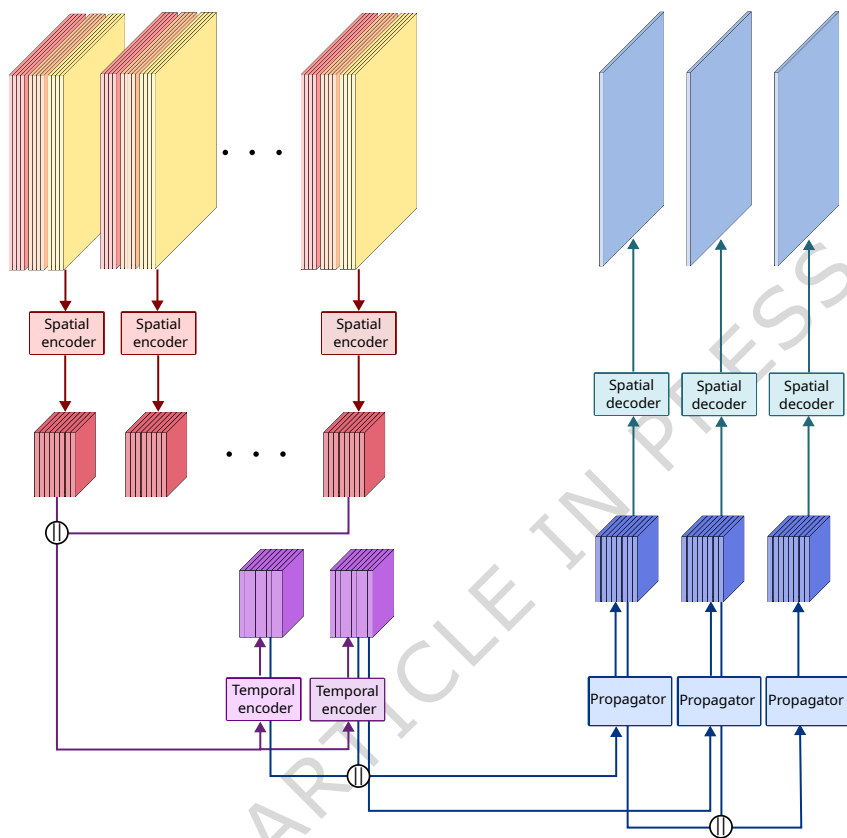


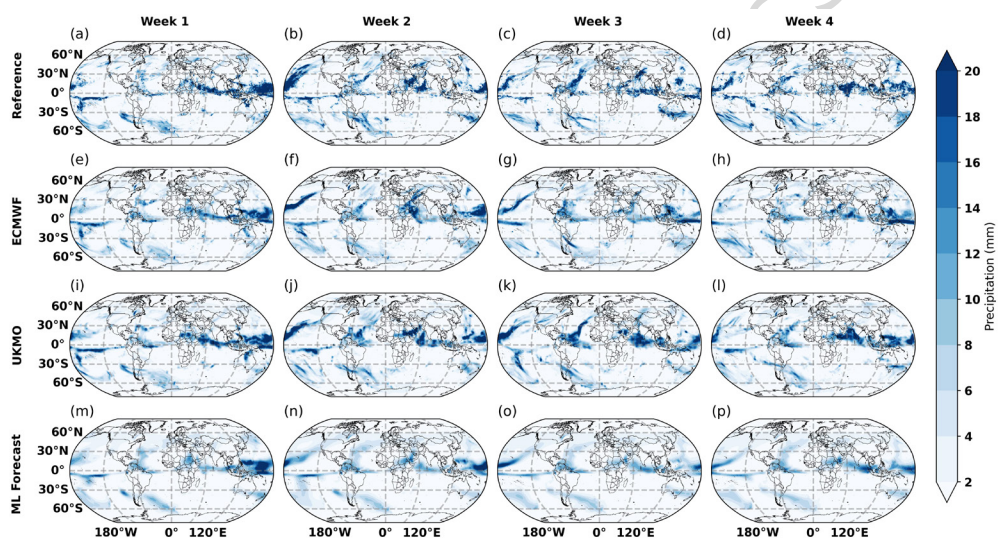



ARTICLE

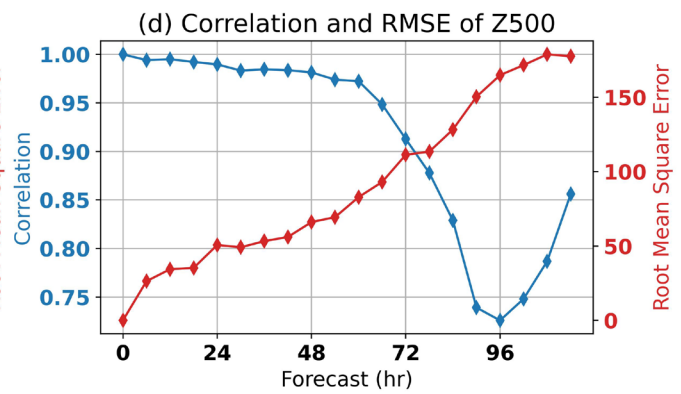
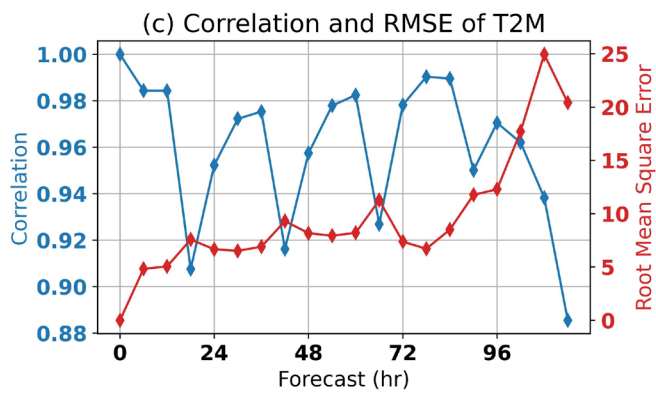
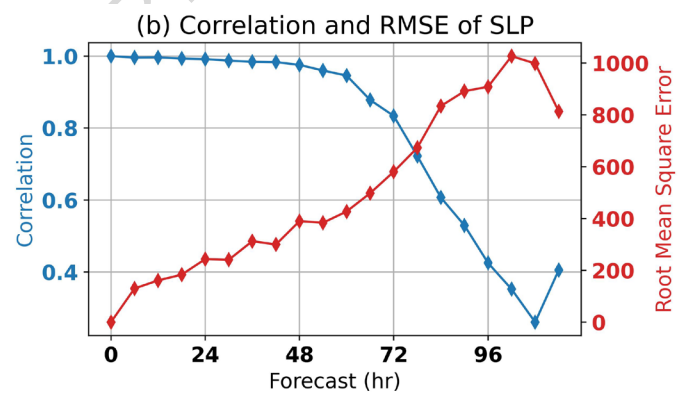
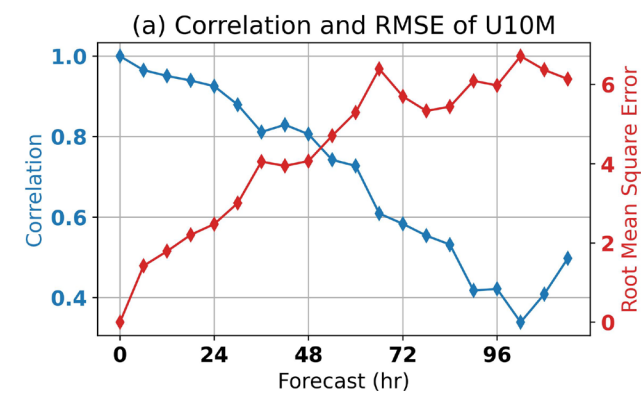


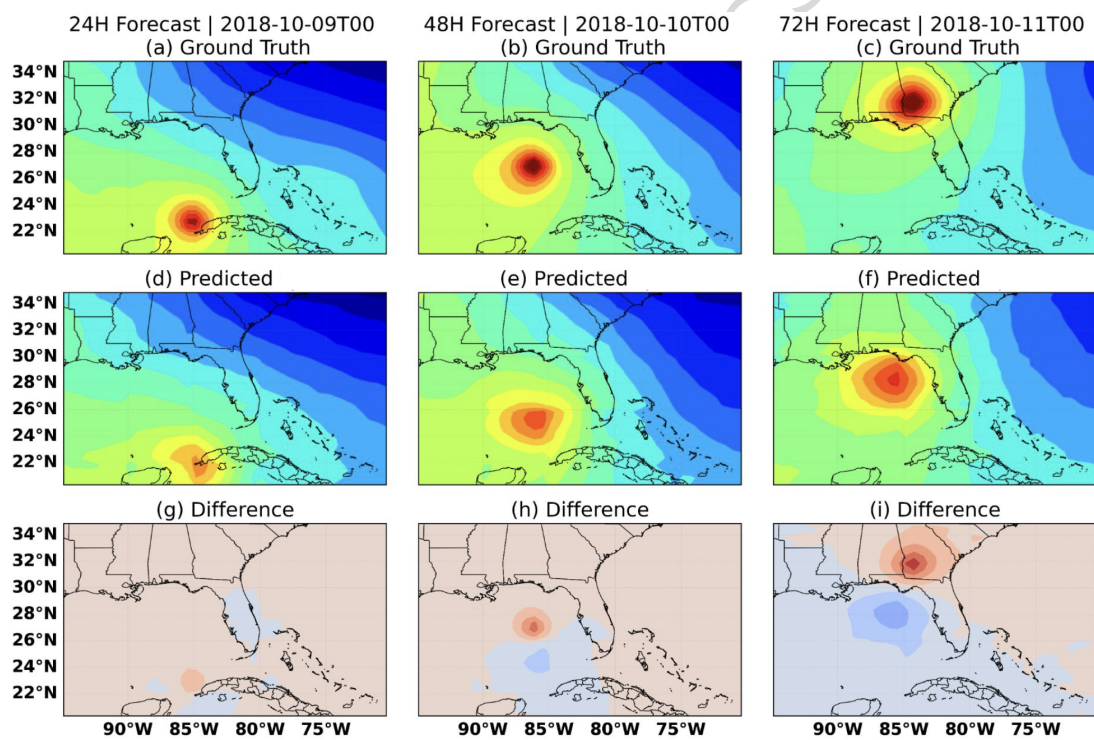
ARTICLE IN PRESS

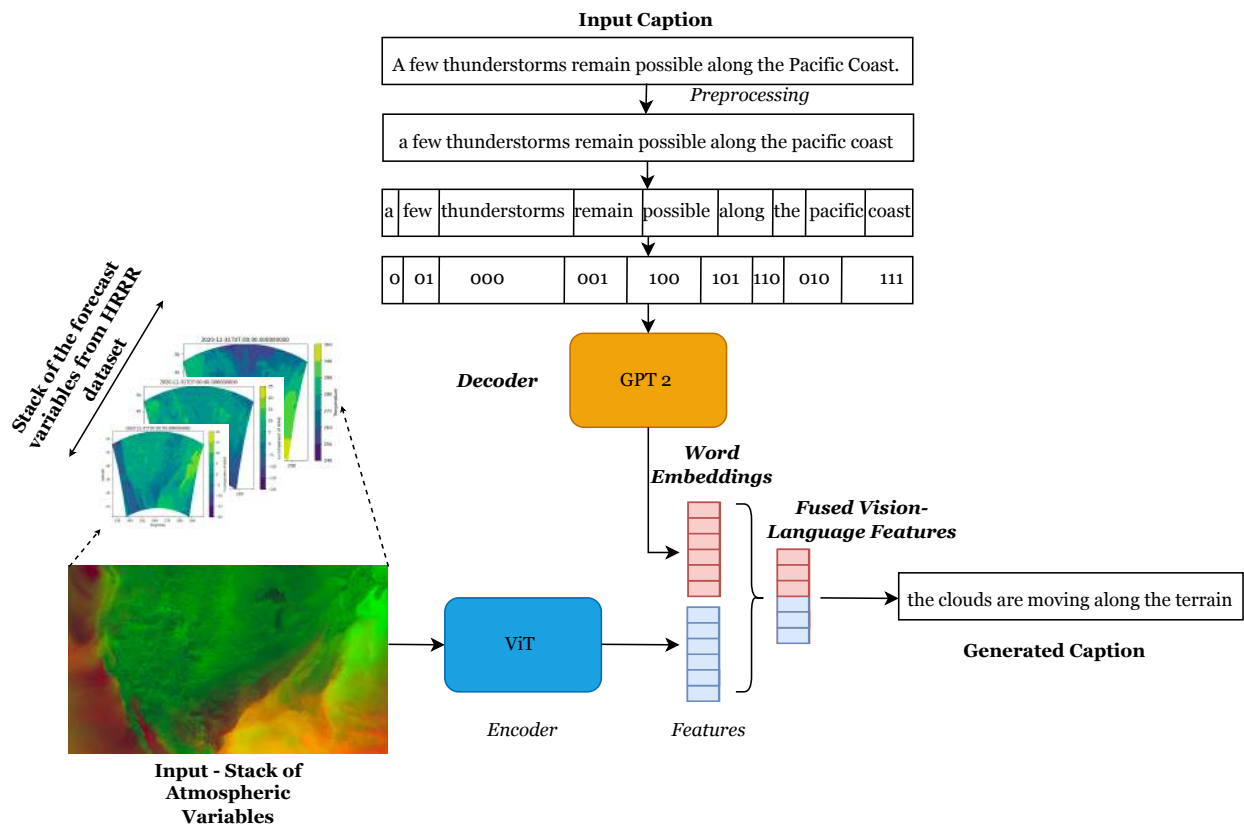















Generated: A few strong to severe thunderstorms are possible this evening across parts of the southern Plains and Ohio Valley.

Severe thunderstorms capable of large hail and damaging winds will continue across parts of the central and southern Plains this evening.

Severe thunderstorms capable of producing damaging winds and hail will continue this evening across parts of the central and southern Plains. A few strong storms will continue to impact parts of southern Nebraska and western Kansas into the overnight hours.

Original: Severe thunderstorms will continue tonight from parts of far East Texas, extreme eastern Oklahoma, and southeast Kansas eastward into the lower and middle Mississippi Valley. Large hail, damaging winds, and a few tornadoes are expected.

Severe storms remain possible across the central Plains and Upper Midwest through this evening, with a relative concentration expected over northern/central Kansas and vicinity.

Severe storms with hail and wind damage will be possible this evening from parts of northwest Kansas into west-central Missouri. A few marginally severe storms may also occur across parts of lower Michigan.

ARTICLE IN PRESS