
DUNE: A MACHINE LEARNING DEEP UNET++ BASED ENSEMBLE APPROACH TO MONTHLY, SEASONAL AND ANNUAL CLIMATE FORECASTING

Pratik Shukla, Milton Halem

Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
Baltimore, MD, USA
{pratiks2, halem}@umbc.edu

ABSTRACT

Capitalizing on the recent availability of ERA5 monthly averaged long-term data records of mean atmospheric and climate fields based on high-resolution reanalysis, deep-learning architectures offer an alternative to physics-based daily numerical weather predictions for subseasonal to seasonal (S2S) and annual means. A novel Deep UNet++-based Ensemble (DUNE) neural architecture is introduced, employing multi-encoder-decoder structures with residual blocks. When initialized from a prior month or year, this architecture produced the first AI-based global monthly, seasonal, or annual mean forecast of 2-meter temperatures (T2m) and sea surface temperatures (SST). ERA5 monthly mean data is used as input for T2m over land, SST over oceans, and solar radiation at the top of the atmosphere for each month of 40 years to train the model. Validation forecasts are performed for an additional two years, followed by five years of forecast evaluations to account for natural annual variability. AI-trained inference forecast weights generate forecasts in seconds, enabling ensemble seasonal forecasts. Root Mean Squared Error (RMSE), Anomaly Correlation Coefficient (ACC), and Heidke Skill Score (HSS) statistics are presented globally and over specific regions. These forecasts outperform persistence, climatology, and multiple linear regression for all domains. DUNE forecasts demonstrate comparable statistical accuracy to NOAA’s operational monthly and seasonal probabilistic outlook forecasts over the US but at significantly higher resolutions. RMSE and ACC error statistics for other recent AI-based daily forecasts also show superior performance for DUNE-based forecasts. The DUNE model’s application to an ensemble data assimilation cycle shows comparable forecast accuracy with a single high-resolution model, potentially eliminating the need for retraining on extrapolated datasets.

Keywords AI Ensemble Forecasting · Seasonal-to-Annual Forecasts · 2m-Temperature · Sea Surface Temperature

1 Introduction

Traditional numerical weather prediction (NWP) models employed by operational weather forecasting centers for subseasonal to seasonal (S2S) forecasts are based on ensembles of physics-based daily weather forecasting models. Operational models aim to overcome extended-range predictability limitations by using the mean of ensembles of forecast models with either different physical parameterizations or perturbations of the initial forecast conditions. All such ensemble members contain variations in their numerical approximations of fluid flows and physical parameterizations of sub-scale meteorological processes. Nevertheless, NWP has been shown to have limited forecast skill beyond two weeks, regardless of model parameterization and resolution of initial observations (Palmer, 2024; Charney, 1966; Lorenz, 1963c; Halem and Jastrow, 1970; Lorenz, 1982; Zhang et al., 2019). Machine learning has now demonstrated similar 2-4 week forecast skills (Shen et al., 2023; Udel, 2019; Weyn et al., 2021; Bi et al., 2023; Nguyen et al., 2023) based on training with hourly to daily reanalysis data, yet it suffers from the same limitations in extended predictability beyond two weeks (Palmer, 2024). Since boundary forcings govern long-term seasonal and annual climates based on

mean properties of the atmosphere, the problem of S2S prediction is addressed using a deep learning self-generating AI ensemble forecasting model, Deep UNet++ Ensemble (DUNE), trained on very long records of reanalysis observations of monthly mean ERA5 data. This study focuses on Subseasonal-to-Seasonal and Annual (S2SA) machine learning model forecasting.

Reliable predictive models of Subseasonal to Seasonal (S2S) temperatures at 2 meters above the surface (T2m) are helpful and often urgently needed for many environmental, economic, and social applications. Temperatures, precipitation, and potential evaporation are indices of drought and significant components influencing wildfire behavior. To enhance mega wildfire likelihood forecasts, one must predict surface drought indices such as annual surface temperature heat anomalies, potential evapotranspiration influenced by T2m, and vegetation dryness due to extended dry precipitation anomalies. Such accurate predictions are vital for anticipating wildfire likelihood and severity. Predicting seasonal and annual surface temperature variations is equally critical. Higher summer temperatures in Boreal regions signal broader climate change trends. Understanding the complex relationship between drought conditions and wildfires is crucial for refining predictive models. Droughts create favorable ignition conditions and contribute to vegetation drying, enhancing wildfire intensity and extent. Accurate drought predictions will enable actionable strategies for proactive wildfire management. Thus, as a vital first step in implementing sustainable strategies to address the challenges of global warming's impact on wildfire risk, we focus in this paper mainly on developing a comprehensive approach to predicting summer and annual T2m. Such an AI model is applicable to the operational requirement of subseasonal to seasonal prediction.

In the **Related Work** section 2, we review relevant ML papers, including numerical weather prediction (NWP) forecasts and various short-term and long-term AI-based approaches. The **Data Description** section 3 details the datasets utilized for training our model. An outline of the methods employed to assess our model's performance are presented in the **Evaluation** section 4. The **Methodology** section 5 provides an in-depth description of the DUNE AI model architecture and its training procedures. In the **Results** section 6, forecasts for monthly, seasonal, and yearly means, as well as moving window forecasts are presented. The **Discussions** section 7 offers insights and reflections on the experiments conducted. Finally, the **Conclusions** section 8 provides closing remarks and summaries. The Appendix provides supplementary graphs, tables, and plots.

2 Related Work

Weather forecasts play a crucial role in our modern world, impacting various sectors of the economy and society. Accurate temperature predictions are critical for agriculture, water management, energy production, extreme heat, drought and wildfire risk.

Neural networks are emerging as complementary models and even as an alternative to operational numerical weather prediction (NWP) models. The recent availability of reanalysis observational data along with computational capabilities has facilitated the development of ensemble probabilistic forecasting methodologies for NWP models (Li et al., 2022; Rasp and Lerch, 2018; Grönquist et al., 2021). These methodologies leverage deep neural networks to generate and improve the representation of forecast uncertainty, which is critical for predictability under uncertainty.

Larraondo et al. (2019) explored various convolutional neural network (CNN) architectures for predicting precipitation fields from geopotential height fields using reanalysis data. Their findings showed the superiority of the U-Net architecture in generating more accurate precipitation forecasts, thus suggesting this CNN algorithm for our weather prediction endeavors.

Additionally, (Weyn et al., 2019, 2020, 2021) utilized the Resnet CNN to forecast atmospheric variables over different time scales, surpassing climatological benchmarks for weather prediction. Similar experiments involving CNNs by Taylor and Feng (2022) highlighted the role of land surface temperatures (T2m) and sea surface temperatures (SST) in advancing temperature forecasting. Their study primarily focused on analyzing sea surface temperature anomalies using monthly mean data to predict ocean temperatures. It is well-established that weather forecasts, particularly for T2m, tend to exhibit lower accuracy beyond two weeks. Inspired by Taylor and Feng (2022), we departed from the conventional approach of using daily data and adopted monthly mean data for our experiments. Moreover, we incorporated a mixed field of T2m and SST, aligning with the approach presented by (Taylor and Feng, 2022).

In recent years, several innovative approaches have emerged for weather forecasting. However, most of the papers focused on short-term forecasts ranging from a few hours to a few weeks. In their paper, "FourCastNet," Pathak et al. (2022) implemented the Adaptive Fourier Neural Operator (FNO) framework by Guibas et al. (2021) for weather forecasting, efficiently capturing complex spatiotemporal patterns. Lam et al. (2022) introduced "GraphCast," a high-resolution graph neural network model designed for 10-day global weather forecasting, outperforming conventional models. In "ClimaX," (Nguyen et al., 2023) utilized the transformer architecture with innovative encoding and

aggregation blocks for weather forecasting, enabling accurate predictions across different geographic regions and time scales.

Numerous researchers have demonstrated proficiency in utilizing the UNet architecture for weather forecasting. Unal et al. (2023) utilized the UNet++ architecture proposed by (Zhou et al., 2018) and emphasized the significance of using anomalies over raw data to enhance forecasts specifically for long-term weather forecasting for multi-annual forecasts. The UNet++ architecture developed by (Zhou et al., 2018), was shown to significantly advance in medical applications over the UNet model. UNet++ incorporates skip connections at multiple scales, which helps in better integrating feature maps and allows information flow between different layers. This hierarchical connectivity is essential as it enables the model to capture both local and global context, improving its ability to understand complex patterns in the data. Additionally, UNet++ introduces a nested architecture where each encoder block is paired with a corresponding decoder block, further enhancing feature extraction and semantic segmentation performance.

Thus, the UNet++ architecture was tailored to suit the DUNE inference needs. Finally, a new ensemble mean forecast approach was employed by averaging four intermediate forecast outcomes generated by the DUNE AI model. This ensemble technique proved effective, especially during fall and winter, resulting in better forecast accuracy where there was considerable variability in the observations. See the Methodology section 5.

Key Contributions

- Assembled a novel deep learning architecture called the Deep UNet++-based Ensemble (DUNE), specifically designed to incorporate intermediate ensemble inferences (i.e., forecasts) to account for the natural variability in S2S predictions for a range of global climate parameters, including T2m, SST, U10, V10, and more.
- Presented AI-based forecasts for monthly, seasonal, and annual means, validated against monthly ground truth (Reanalysis) data for the past five years (2019-2023). These forecasts encompass global regions, land areas, Oceans, the US, Australia, and the global Boreal regions. Additionally, we have exhibited the feasibility of using DUNE in an operational setting by producing forecasts for 2024 and 2025.
- Produced the first AI-based 12-month forecast for 2024, comparable statistically to NOAA’s operational 12-month probabilistic forecasts, 1000 times faster.
- Trained the DUNE AI model on T2m over land and SST over oceans, enabling simultaneous forecasting of both SST and T2m in a single forecast.
- Employed the DUNE AI model, explicitly trained on anomalies, to predict anomalies rather than actual temperatures.
- Delivered the first Boreal summer anomaly forecast for the mega fire years 2021 and 2023.
- Demonstrated the use of intrinsic ensembles within the DUNE AI model, which improves forecast accuracy during the Fall and Winter months when temperature variability is high.
- Illustrated the global temperature increase over the past 5 years by presenting month-by-month forecasts relative to 20-year temperature intervals.
- Experimented with integrating wavelet transforms into the downsampling component of our DUNE AI architecture. Given their promising results, we plan to implement trainable wavelet transforms in the DUNE AI model’s downsampling and upsampling phases. Due to increased training time, wavelet transforms were not used in the current DUNE AI architecture.

3 Data Description

The historical atmospheric data for training the AI prediction model is sourced from the European Center for Medium-Range Weather Forecasts (ECMWF) monthly mean ReAnalysis version 5, which we obtained from the Copernicus Climate Change Service (C3S) (Hersbach et al., 2020). This ERA5 reanalysis dataset offers monthly estimates of various atmospheric, land, and oceanic variables worldwide, a spatial resolution of 0.25°, roughly equivalent to 30 km. It also includes a 10-member ensemble at roughly 60km. ECMWF developed ERA5 with a global 4-D variational data assimilation cycle using a short-term physical integration as a background and the assimilation of all the available ground and space-borne observations and their error statistics, creating a globally comprehensive and consistent dataset based on current physical and chemistry parameterizations. The dataset covers the Earth and has dimensions of 721 x 1440 for latitude and longitude, respectively. The ensemble data set is 360 x 720. The monthly mean data is available from January 1940 onwards with a temporal resolution of 1 month.

A combination of sea surface temperature (ERA5 label “sst”) and 2m temperature (ERA5 label “t2m”) variables from the ERA5 data set is used as input for training. The land-sea mask (ERA5 label “lsm”) parameter is provided by the ERA5 data, which was used to combine the SST and T2m data sets.

The following five constants plus solar radiation fields are used as input to train the AI model as suggested by Weyn et al. (2020). The constants are lsm, orography (ERA5 label “z”/9.80665), Soil Type (ERA5 label “slt”), High Vegetation Cover (ERA5 label “cvh”) and Low Vegetation Cover (ERA5 label “cvl”). Lastly, we used the Top of the Atmosphere Incident Solar Radiation (ERA5 label “tisir”), which shows the incoming solar radiation from the sun at the top of the atmosphere. This is the only value that varies monthly, although it remains constant over the years.

The complete ERA5 monthly data set was downloaded from January 1940 to June 2024. To calculate the climatology for each month (January to December), the temperature data used spans the 30-year period, from January 1950 to December 1979. An anomaly temperature data set is formed by subtracting the climatology of the respective month from the respective monthly value for each month. Instead of predicting the actual temperature values, the model is trained to predict the temperature anomalies. The actual temperature data are recovered by adding the climatological averages for the respective month.

In the following experiments, data from January 1980 to December 2016 (444 months) are used as input for training, January 2017 to December 2018 (24 months) is used for validation, and January 2019 to December 2023 (60 months) is for conducting the experiments. All the inputs are normalized. A standard normalization technique was applied as shown in equation (1). The following scaling formula was used for the six constant fields to normalize each variable separately. Since the model predicts the scaled temperature, an inverse scaling was applied to the model’s output. The unscaled output was used for model validation.

$$z_i = \frac{x_i - x_{\min}}{x_{\max} - x_{\min}} \quad (1)$$

In equation (1), z_i is the normalized parameter value, x_i is the raw or the input parameter value, $x_{\min}$ is the minimum parameter value, and $x_{\max}$ is the maximum parameter value.

4 Evaluation

The loss function is adopted from Rasp et al. (2020). The root mean squared error (RMSE) is employed as our primary evaluation metric because it penalizes large differences in the predictions.

$$\text{RMSE} = \frac{1}{N_{\text{forecasts}}} \sum_{i=1}^{N_{\text{forecasts}}} \sqrt{\frac{1}{N_{\text{lat}} N_{\text{lon}}} \sum_{j=1}^{N_{\text{lat}}} \sum_{k=1}^{N_{\text{lon}}} L(j) (\text{forecast}_{i,j,k} - \text{truth}_{i,j,k})^2} \quad (2)$$

In equation (2), *forecast* is the model forecast, *truth* is the ERA5 ground truth, N_{lat} represents the number of latitude values, N_{lon} represents the number of longitude values, $N_{\text{forecasts}}$ represents the total number of forecasts made, and $L(j)$ represents the mean latitude weighting factor.

A mean-latitude-weighting factor is used because the area covered by a $0.25^\circ \times 0.25^\circ$ grid box decreases as we move from the tropics to the polar regions. It’s important to note that the grid box area remains constant for different longitude values at each latitude. The mean-latitude-weighting factor is defined in (3).

$$L(j) = \frac{\cos(\text{lat}(j))}{\sum_{j=1}^{N_{\text{lat}}} \cos(\text{lat}(j))} \quad (3)$$

Model forecasts were also evaluated using the mean-latitude-weighted anomaly correlation coefficient (ACC). This metric ranges from -1 to 1, where a higher ACC value indicates better agreement between the model’s predictions and the actual observations, while a lower ACC value indicates poorer agreement. The ACC assesses the consistency of spatial variations, regardless of their magnitudes.

$$\text{ACC} = \frac{\sum_{i,j,k} L(j) \text{forecast}'_{i,j,k} \text{truth}'_{i,j,k}}{\sqrt{\sum_{i,j,k} L(j) \text{forecast}'_{i,j,k}{}^2 \sum_{i,j,k} L(j) \text{truth}'_{i,j,k}{}^2}} \quad (4)$$

In equation (4), *forecast'* and *truth'* denote the difference to the climatology. We used the climatology as defined in section 3.

The Heidke Skill Score (HSS), as defined by NOAA, is used to compare the DUNE AI inferences with the NOAA operational forecasts. To evaluate forecast accuracy using the Heidke Skill Score (HSS), a contingency table is

employed to categorize forecast outcomes into hits (H), misses (M), false alarms (F), and correct negatives (C). Hits (H) represent correct forecasts where predicted events match observed events. The total number of valid forecast-observation pairs (T) is the sum of all categories in the contingency table (H + M + F + C). More details on this can be found at https://www.cawcr.gov.au/projects/verification/#Contingency_table. The expected number of correct forecasts (E), assuming random chance, is typically one-third of T. The HSS is then calculated using the formula:

$$\text{HSS\%} = 100 \times \frac{H - E}{T - E} \quad (5)$$

In equation (5), H denotes the number of correct forecasts, E denotes the expected number of correct forecasts (1/3 of total), and T is the total number of valid forecast-observation pairs.

The HSS measures how often the forecast category correctly matches the observed category above the number of correct hits expected by chance alone. HSS utilizes the number of correct and incorrect category hits, with values ranging from -50 to 100. A score of 100 indicates a perfect forecast, while a score of -50 indicates a perfectly incorrect forecast. Scores greater than 0 indicate improvement compared to a random forecast, thus indicating skill (NOAA Heidke Skill Score Explanation, 2024).

We categorized the data using percentiles based on the climatology (1991-2020). For any given month, if the inference/observation value is less than the lowest 33rd percentile, it is classified as "below normal"; if it is higher than the 66th percentile, it is classified as "above normal"; and if it falls between the 33rd and 66th percentiles, it is categorized as "near normal." The results with the Heidke Skill Score (HSS) are in Section 6.7.

5 Methodology

The Deep UNet++ Ensemble (DUNE) AI forecast model is a specific encoder-decoder architecture based on the UNet++ architecture (Zhou et al., 2018) proposed initially for medical image segmentation tasks. The UNet++ architecture is a variation of the famous U-Net architecture (Ronneberger et al., 2015), which is also a variation of the traditional encoder-decoder networks (Baldi, 2012) that have succeeded in image segmentation tasks. Larraondo et al. (2019) investigated the application of several CNNs, such as SegNet, VGG16, and U-Net, to predict precipitation fields from geopotential height fields using the reanalysis data and found that U-Net delivered the best forecasts.

The DUNE AI forecast model, illustrated in Figure 1, uses a fully convolutional neural network (FCN) (Long et al., 2015) to predict the atmospheric state from one monthly mean to the next. The convolution operations within the CNN learn a specified number of filters, each with a shape of $M \times N$. These filters are translated across the input atmospheric state, producing an output state for each $M \times N$ input image multiplied by the filters. The expanded architecture is illustrated in Figure A1.

The initial weights for these filters are generated using the Kaiming initialization method (He et al., 2015), and the model learns their weights through gradient back-propagation during the training phase (Shukla, 2022). These filters extract complex patterns and features from the input states, and their weights are shared across the input domain. This design allows the DUNE AI model to perform pixel-wise semantic calculation, enabling the generation of quantitative estimates of meteorological parameters such as T2m and SST at each latitude-longitude grid point.

Notably, the DUNE AI model has a relatively low number of trainable parameters, approximately 45,843,720. In contrast, a neural network that fully connects all input features to all output features would have a trillion parameters, which is impractical.

The DUNE architecture was based on the UNet++ architecture for a prediction task. The architecture employs convolution operations to learn filter weights and to increase the number of channels. Throughout the network, the same filter size (3 x 3) is used with padding to maintain dimensionality at each level. Two residual CNN blocks are utilized at each level, with the model updating filter weights iteratively through back-propagation. As the number of channels progressively increases, this augmentation enables the network to capture and learn more complex patterns and features from the input data and to better generalize to unseen data.

These convolutional operations are followed by a spatial pooling operation, which reduces the dimensionality of the atmospheric state by a factor of 2. This reduction is achieved by calculating the average value within each 2 x 2 sub-grid. After the pooling layer, convolutional operations are again applied with the same filter size (3 x 3) on states with progressively coarser spatial resolution. Through this process, the CNNs learn filters that represent larger-scale atmospheric patterns. This sequence constitutes the contracting path, or the encoder part of the architecture, effectively

summarizing the information by reducing the dimensionality of the atmospheric state while increasing the number of channels.

In a CNN, low-level feature maps at higher resolutions capture detailed texture information, such as fine-scale temperature variations and localized anomalies. This allows the model to detect subtle patterns and gradients in temperature distribution. In contrast, high-level feature maps at lower resolutions encode more abstract semantic information, capturing broader climatic trends and large-scale patterns such as seasonal cycles and long-term temperature changes. We often deepen the network to enhance our model’s feature extraction ability. Inspired by the DenseNet architecture (Huang et al., 2017), we implemented an extreme form of skip connections to supply low-level features to deeper network layers. Specifically, we passed each layer’s features to all subsequent layers and then fused them through concatenation. This was achieved by applying pooling operations of varying sizes to the higher-dimensional feature maps before integrating them with the lower-dimensional ones. Including additional feature maps from higher dimensions to help the model capture more accurate and subtle patterns.




Figure 1: The DUNE AI Architecture

As the network gets deeper, one often faces the problem of vanishing gradients and network degradation (Hochreiter, 1998). To avoid these problems and ensure a faster model convergence, we used a residual CNN block (He et al., 2016) instead of a CNN block at each level. To create a residual block, we applied a linear transformation using a CNN with filter size 1×1 to the original input so that it has the same number of channels required for the summation operation at the end. The residual CNN block is shown in Figure 2(a), and a regular CNN block is shown in Figure 2(b).

In the expanding path or the decoder part of the architecture, transposed convolutions Zeiler et al. (2010) are used to increase the dimensionality of the state of the atmosphere. At each up-sampling layer, the spatial dimensionality is increased by a factor of 2. This operation is continued until one reaches the dimensions of the original input.

However, down-sampling operations lead to the loss of spatial information, potentially resulting in overly smooth predictions. U-Net addresses this by concatenating the tensor state of the CNN at each encoding step with the tensor state immediately following each up-sampling operation in the decoding phase. These skip connections reduce information loss after the bottleneck layer and aid in recovering fine-grained details.




Figure 2: (a) Residual Block (b) Regular CNN Block

In UNet++, these skip connections are redesigned to minimize the semantic gap between the features incoming from the sub-networks of the encoder and the decoder, making learning more effective. By doing so, the network can learn more abstract and meaningful representations of the input data. Additionally, integrating additional convolutional layers further enhances the network's ability to extract and refine features, resulting in a significantly improved feature representation that benefits prediction tasks.

Four predictions were generated by trimming our architecture to various depths, using the power of ensemble forecasts within a single architecture. Next, convolutional layers are applied to each output with a filter size of 1×1 to consolidate multiple channel outputs into a single channel. Finally, a final prediction is created by averaging these four ensemble predictions from the four channels.

Except for the final four convolutions, which utilize a filter size of (1×1) , each convolutional operation in the architecture is followed by a ReLU nonlinear activation function Agarap (2018).

The model is trained using the Adam optimizer Kingma and Ba (2014), a version of the stochastic gradient descent optimization algorithm Loshchilov and Hutter (2016), with a learning rate of $1e-3$, a batch size of 4, and a weight decay of $1e-4$. Additionally, the CosineAnnealing learning rate scheduler is used with a maximum of 225 iterations to adjust the learning rate. It uses the latest learning rate after the 225th epochs. A mean latitude-weighted root-mean-squared-error is used as the loss function (see section 4). The model is trained for 500 epochs to ensure effective loss minimization, with early stopping based on the validation set loss. If no new validation loss minimum is observed within 100 epochs during training, the training stops, and the model weights yielding the smallest validation loss are restored for further testing.

The model was constructed using the open-source PyTorch library (Paszke et al., 2019) from Meta AI in Python. For most of the data pre-processing pipeline, the open-source xarray project was utilized (Hoyer and Hamman, 2017). The model training was conducted on 8 Tesla A100-40 GB GPUs on the NASA Center for Climate Simulation (NCCS) servers at the Goddard Space Flight Center. The Matplotlib library (Hunter, 2007) was used for all visualizations, employing Basemap for map projections.

Long-Range Forecast using the Moving Window Approach: The DUNE AI model can forecast the monthly mean, seasonal mean, or annual mean temperatures one monthly/seasonal/annual mean in advance. For example, given the prior month's monthly mean temperature, it will forecast the next month's monthly mean temperature. Inspired by (Tripura and Chakraborty, 2023), a moving window approach was used to extend the model's capabilities to forecast more than one month in advance.

In the moving window approach, the prior x months are used to forecast the next x months. For example, if one wants to forecast the temperature for the next 12 months, the prior 12 months' temperatures and constants can be used. Note

that there are 12 inputs for TISR and 5 channels for the other constants. The results of DUNE are presented in section 6.6. More information on this can be found in the Appendix.

6 DUNE Forecast Results

This section presents the monthly, seasonal, and annual mean forecast experiments using the DUNE AI method. These forecasts are compared to three other baseline methods: persistence, climatology, and multiple linear regression. Regional performances, derived from the global forecast, are separately analyzed with respect to performance statistics.

6.1 Global Forecast Comparison: DUNE AI versus Other AI Methods

Table 1 presents a comparative analysis of the DUNE AI method against other AI-based forecasting approaches using monthly data. The results indicate that when trained on monthly mean data, the DUNE AI model yields superior global forecast accuracy compared to other AI models that utilize hourly or daily data for monthly forecasting (shown in table 1). Specifically, the DUNE AI method demonstrates a twofold reduction in RMSE relative to other AI methods when employing monthly mean data as input.

Evaluation Method	Forecast Time	Climax (5.625°)	Climax (1.40625°)	ResNet (5.625°)	UNet (5.625°)	DUNE AI (0.25°)
RMSE ^a (↓)	720hrs	2.67	2.74	2.86	2.81	1.07^c
ACC ^b (↑)	720hrs	0.85	0.84	0.83	0.83	0.95/0.74^{b, c}

Table 1: Comparison of DUNE AI monthly forecasts (bold) with other AI hourly/daily averaged methods

^a The DUNE AI method generates monthly mean forecasts, whereas the other methods in the table produce hourly forecasts up to 720 hours in advance. This highlights the advantage of using monthly mean data.

^b To calculate the ACC, the DUNE AI method used climatology data from 1950 to 1979. In contrast, the other AI methods in the table used climatology data from their respective test periods. If the DUNE AI method used the climatology from the test period, it gives the ACC of 0.95, as shown in the table. It is worth noting that the DUNE AI forecast is at a significantly finer resolution. The table values are adopted from (Nguyen et al., 2023). ^c Results using monthly mean data.

6.2 Global Monthly Mean Forecasts

This experiment uses the inferred weights computed from training the DUNE AI model to use the prior month to forecast the successive monthly mean temperatures using the ERA5 input data from 1980 to 2016. Given the temperature anomalies and the six constants, DUNE infers the temperature anomalies for the next month. In particular, using data from 1980 to 2016 (444 months) for training successive months, and then data from the years 2017 and 2018 (24 months) to validate the forecast inference, forecasts (i.e., inferences) are then made employing the derived weights to infer monthly data from all months between 2019 and 2023 (60 months). The wall clock time required for processing and inferring forecasts over 60 months was approximately 20 seconds.

Table 2 compares the DUNE AI model’s forecast skills (inferences) with the baselines, persistence, climatology, and multiple linear regression. The average of RMSE, ACC, and HSS is presented in table 2 over the testing period from 2019 to 2023 (5*12 = 60 months). The DUNE AI method outperforms all the baselines for monthly mean forecasts. We utilized the Scikit-learn library (Pedregosa et al., 2011) to calculate statistics for the multiple linear regression baseline. Since multiple linear regression consistently yields poorer performance than other baselines, they were not included in further regional results.

Method	Global			Global Land			Global Oceans			US			Australia			Boreal Forests		
	RMSE (↓)	ACC (↑)	HSS (↑)	RMSE (↓)	ACC (↑)	HSS (↑)	RMSE (↓)	ACC (↑)	HSS (↑)	RMSE (↓)	ACC (↑)	HSS (↑)	RMSE (↓)	ACC (↑)	HSS (↑)	RMSE (↓)	ACC (↑)	HSS (↑)
Persistence (Prior Month)	2.76	0.33	19.39	4.41	0.27	11.46	1.16	0.59	23.46	4.68	0.19	13.01	2.96	0.13	5.87	6.82	0.16	5.38
Persistence (Prior Year Same Month)	1.52	0.53	9.71	2.28	0.51	7.16	0.90	0.61	11.02	2.39	0.42	5.54	1.82	0.22	1.15	3.25	0.46	4.36
Climatology (1950-1979)	1.60	0	-11.24	2.38	0	-12.30	1.03	-	-10.70	2.26	0	-10.40	1.51	0	-5.20	3.18	0	-9.10
Multiple Linear Regression (2° resolution)	1.96	0.42	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
DUNE AI Method	1.07	0.74	26.84	1.65	0.69	10.95	0.52	0.85	34.95	1.80	0.59	9.65	1.32	0.45	8.13	2.41	0.65	18.71

Table 2: Comparison of five-year average (2019-2023) monthly mean forecasts using RMSE, ACC, and HSS metrics. The results demonstrate that DUNE AI inferences significantly improve performance compared to baseline models, including multiple linear regression. The DUNE AI model outperforms the baselines in all the regions.

Figure 3 and Figure 4 show the global RMSE and global ACC for the test period from 2019 to 2023, respectively. For RMSE, the lower, the better; for ACC, the higher, the better. The DUNE AI model generated better forecasts for all the months compared to the baselines.




Figure 3: Global RMSE of Monthly Mean Forecasts: (a) Persistence (PYSM), (b) Climatology, and (c) DUNE AI Method. The comparison shows that the DUNE AI model consistently achieves lower RMSE values, indicating superior forecasting skills. Notably, during the record warm year of 2023, the DUNE AI model exhibited relatively small variability compared to the baseline models. Over the five years, the DUNE AI model displays slightly higher variance in January, February, and December but maintains a uniformly small variance for the remaining months.




Figure 4: Global ACC of Monthly Mean Forecasts: (a) Persistence (PYSM), (b) Climatology, and (c) DUNE AI Method. The comparison shows that the DUNE AI model consistently achieves higher ACC values, indicating superior forecasting skills. Notably, during the record warm year of 2023, the DUNE AI model exhibited relatively small variability compared to the baseline models. Over the five years, the DUNE AI model displays higher ACC value of 0.8 compared to 0.5 (baselines) in addition to maintaining a smaller variance.

Since December 2023 exhibits the highest temperature variability, Figure 5 compares forecast errors (prediction-truth) from the DUNE AI model with baseline methods for this month. In the figure, forecast errors closer to zero (depicted in green) indicate improved forecast accuracy. The visual assessment demonstrates that the DUNE AI model's forecast errors are substantially lower than those from the baseline methods.




Figure 5: Visualization of forecast errors for December 2023: (a) Persistence (PM), (b) Persistence (PYSM), (c) Climatology, (d) DUNE AI Method. RMSE and ACC values are displayed below each plot. (a) and (b) highlight significant prediction errors using the Persistence method over North America and Eurasia. In contrast, the DUNE AI model demonstrates substantial performance improvements, particularly in the El Niño region, South America, Africa, North America, and Boreal regions across North America and Eurasia.

Figure 6 compares ERA5 ground truth anomalies with the DUNE AI forecast anomalies for December 2023, which shows the DUNE AI forecast is very close to the ERA5 ground truth anomalies.




Figure 6: A visual comparison of ERA5 ground truth anomalies and the DUNE AI forecast anomalies for December 2023: (a) ERA5 ground truth anomalies for December 2023, (b) DUNE AI forecast anomalies for December 2023. This comparison illustrates the alignment between the observed data and the AI model’s predictions, highlighting the DUNE AI’s capability to replicate the observed anomalies for December 2023 accurately. The DUNE AI model notably captured anomalies in the El Niño region near South America. It also successfully predicted anomalies in North America, though these predictions were less pronounced than the observed anomalies. Additionally, the model accurately predicted the Western Australian Current and the Equatorial Countercurrent. However, the DUNE AI model underpredicted anomalies over Australia and failed to capture the lower anomalies in the Amazon Rainforest.

6.3 Global Seasonal Mean Forecasts

The following experiments were designed to predict seasonal mean temperatures one season in advance. A season is defined as the average of the temperatures from the three months comprising one of the four seasons. Using data from January 1980 to December 2016, the first season available for prediction is Summer 1980, as Spring would require data from December 1979. The last season available for prediction is Fall 2016, as Winter would require data from January 2017. Consequently, for the 36 years of training data, there were 142 seasons available for training seasonal predictions, 6 seasons for validation, and 18 seasons for testing.

Model prediction statistics were computed for the following baselines: Persistence using Prior Season (PS), Persistence using Prior Year Same Season (PYSS), and Climatology from 1950-1979. The table shows the average RMSE and ACC over 18 seasons from 2019 to 2023. Table 3 presents global forecast comparisons. Table 3 shows the average RMSE and ACC over the testing period from 2019 to 2023. Some supplementary graphs and visualizations for the seasonal mean experiment are in the Appendix.

Improved results for seasonal forecasts can be attributed to the reduced variability in seasonal mean temperatures compared to monthly means.

Method	Global		Global Land		Global Oceans		USA		Australia		Boreal Forests	
	RMSE (↓)	ACC (↑)	RMSE (↓)	ACC (↑)	RMSE (↓)	ACC (↑)	RMSE (↓)	ACC (↑)	RMSE (↓)	ACC (↑)	RMSE (↓)	ACC (↑)
Persistence (Prior Season)	6.24	0.12	10.12	0.09	2.26	0.35	11.18	0.005	6.74	0.09	16.34	-0.006
Persistence (Prior Year Same Season)	1.09	0.67	1.52	0.69	0.76	0.65	1.35	0.22	1.35	0.33	2.18	0.67
Climatology (1950-1979)	1.37	0	1.94	0	0.95	0	1.65	0	1.23	0	2.57	0
Multiple Linear Regression (2°resolution)	1.65	0.48	-	-	-	-	-	-	-	-	-	-
DUNE AI Method	0.87	0.77	1.20	0.79	0.63	0.74	1.12	0.79	1.10	0.44	1.66	0.77

Table 3: Comparison of five-year average (2019-2023) seasonal mean forecasts using RMSE and ACC metrics. The results demonstrate that DUNE AI inferences significantly improve performance compared to baseline models, including multiple linear regression. The DUNE AI model outperforms the baselines in all the regions.

6.4 Global Yearly Mean Forecasts

Similarly, predicting yearly mean temperatures a year in advance leaves 37 years for training, 2 years for validation, and 5 years for testing. One could increase the number of seasons and years for predictions by creating sliding 3-month

seasons and years different from the calendar, but this was not done for this study. Table 4 compares global yearly mean RMSE and ACC results. Visualization results are in the Appendix.

Enhanced accuracy in yearly mean forecasts is attributed to the reduced variability in annual mean temperatures compared to seasonal or monthly means.

Method	Global		Global Land		Global Oceans		USA		Australia		Boreal Forests	
	RMSE ($\downarrow$)	ACC ($\uparrow$)	RMSE ($\downarrow$)	ACC ($\uparrow$)	RMSE ($\downarrow$)	ACC ($\uparrow$)	RMSE ($\downarrow$)	ACC ($\uparrow$)	RMSE ($\downarrow$)	ACC ($\uparrow$)	RMSE ($\downarrow$)	ACC ($\uparrow$)
Persistence	0.67	0.83	0.86	0.86	0.52	0.78	0.91	0.83	0.76	0.64	1.33	0.84
Climatology (1950-1979)	1.20	0	1.66	0	0.83	0	1.49	0	0.90	0	2.25	0
Multiple Linear Regression (2°resolution)	1.40	0.55	-	-	-	-	-	-	-	-	-	-
DUNE AI Method	0.63	0.85	0.77	0.89	0.53	0.77	0.79	0.85	0.74	0.64	1.08	0.89

Table 4: Comparison of five-year average (2019-2023) annual mean forecasts using RMSE and ACC metrics. The results demonstrate that DUNE AI inferences significantly improve performance compared to baseline models, including multiple linear regression. The DUNE AI model outperforms the baselines in all the regions.

6.5 Validation and Testing Period Statistics Comparison

DUNE AI Forecast	Monthly Mean		Seasonal Mean		Annual Mean	
	RMSE ($\downarrow$)	ACC ($\uparrow$)	RMSE ($\downarrow$)	ACC ($\uparrow$)	RMSE ($\downarrow$)	ACC ($\uparrow$)
Validation Period (2017-2018)	1.03	0.71	0.80	0.76	0.50	0.87
Testing Period (2019-2023)	1.07	0.74	0.87	0.77	0.63	0.85

Table 5: Comparison of DUNE AI monthly/seasonal/annual mean forecasts for validation (2017-2018) and testing (2019-2023) period.

Figure 7 shows the training, validation, and testing loss.




Figure 7: Monthly Mean Forecast Training, Validation, and Testing loss

6.6 Monthly Mean Extended-Range Forecast Results

In this section, DUNE monthly mean forecast results are presented using the moving window approach for extended periods. Using this approach, monthly mean forecasts were conducted for 1, 2, 3, 4, 6, and 12 months in advance.

Table 6 compares global monthly mean forecast results. Although the forecast skills of the DUNE model decrease with an extended forecast period, as anticipated, they remain superior to the baseline methods. Additionally, the RMSE and ACC plots are provided in the appendix for further reference.

Moving Window Frame	Method	Global		Global Land	
		RMSE ($\downarrow$)	ACC ($\uparrow$)	RMSE ($\downarrow$)	ACC ($\uparrow$)
-	Persistence (Prior Month)	2.76	0.33	4.30	0.27
-	Persistence (Prior Year Same Month)	1.52	0.53	2.23	0.51
-	Climatology (1950-1979)	1.60	0	2.29	0
-	Multiple Linear Regression (2°resolution)	1.96	0.42	-	-
One Month	DUNE AI Method	1.07	0.74	1.65	0.69
Two Months	DUNE AI Method	1.13	0.70	1.66	0.68
Three Months	DUNE AI Method	1.16	0.69	1.67	0.68
Four Months	DUNE AI Method	1.17	0.68	1.68	0.67
Six Months	DUNE AI Method	1.20	0.66	1.69	0.67
Twelve Months	DUNE AI Method	1.27	0.61	1.75	0.64

Table 6: A five-year (2019-2023) comparison of average monthly mean forecasts using the Moving Window Approach reveals an expected increase in error statistics and a decrease in forecasting skills as the moving window extends from 1 month to 12 months. Notably, the decline in skill over land is minimal and remains superior to all baseline results.

6.7 Comparison with NOAA’s Monthly Forecasts

This section provides a comparative analysis between the forecasts generated by the DUNE AI model and those issued by the National Oceanic and Atmospheric Administration (NOAA). The NOAA forecasts are designed explicitly for the United States, whereas the DUNE AI model is trained and inferred from training to predict T2m globally. For a direct comparison, statistics for the United States were extracted from the global output of the DUNE AI model. Notably, NOAA’s forecasts are produced at a spatial resolution of $2^\circ \times 2^\circ$, whereas the DUNE AI model provides inferences at a significantly finer resolution of $0.25^\circ \times 0.25^\circ$. The NOAA Climate Prediction Center Gridded Monthly Forecasts and an explanation of their products can be accessed at: https://www.cpc.ncep.noaa.gov/products/predictions/long_range/tools/briefing/mon_veri.grid.php.

For this experiment, forecast and observation categories (i.e., above, near, and below normal as defined by NOAA) were based on the percentile values of the climatology (1991-2020). For each month, if the forecast or observation value falls below the 33rd percentile of the climatology for that month, it is classified as below normal (indicated in blue). If the value exceeds the 66th percentile, it is classified as above normal (shown in red). Values between the 33rd and 66th percentiles are classified as near normal. These classifications for the DUNE forecasts have been validated by comparing ERA5 observations. NOAA categories are validated by NOAA relative to the NOAA observations and taken from the above NOAA Climate Center link. The climatology and categories used for the DUNE results are intended to be consistent with those described by the NOAA CPC (NOAA Operation Forecasts Blogs, 2024).

The following figures selectively show monthly forecast comparisons of DUNE and the CPC methodology to illustrate the difference between observations at 0.25 and 2.0 degrees in performing HSS accuracy. Figure 8 provides a comparative analysis of the DUNE AI inferences and NOAA official forecasts relative to observations for August 2023. The analysis reveals that the DUNE AI inference, with a Heidke Skill Score (HSS) of 34.18, substantially outperformed the NOAA official forecast for this case, which has an HSS of 23.49, as the NOAA measure of predictive accuracy.




Figure 8: Comparison of forecasts and observations for August 2023: (a) DUNE AI Inference with an HSS of 34.18, (b) ERA5 Observations, (c) NOAA Operational Forecast with an HSS of 23.49, (d) NOAA Observations. The DUNE AI model successfully inferred above-normal temperatures on the west coast, midwest, and southeast coast. Additionally, in the northeast, the DUNE AI model predicted near-normal temperatures, while NOAA forecasted higher temperatures. However, the actual temperatures in the northeast were below normal. Overall, the DUNE AI model performed very well, demonstrating superior performance compared to the NOAA forecast.

Figure 9 presents a comparative analysis of forecast and observation data for October 2022. The DUNE AI inference achieved a Heidke Skill Score (HSS) of 29.97, whereas the NOAA operational forecast had an HSS of 7.33. This analysis highlights the performance of the DUNE AI model compared to the NOAA forecasts for this month.




Figure 9: Comparison of forecasts and observations for October 2022: (a) DUNE AI Inference with an HSS of 29.97, (b) ERA5 Observations, (c) NOAA Operational Forecast with an HSS of 7.33, (d) NOAA Observations. The DUNE AI model successfully captured the higher temperature anomalies in the western region of the United States. However, in the southeastern region, it failed to capture the below-normal anomalies. Comparing the DUNE AI and NOAA forecasts in the southeastern region, the DUNE AI model performed comparatively better. It predicted near-normal temperatures, whereas the NOAA forecast predicted above-normal temperatures, even though the observed values were below normal.

Figure 10 presents a comparative analysis of forecast and observation data for May 2021. The DUNE AI inference achieved a Heidke Skill Score (HSS) of 15.71, whereas the NOAA operational forecast had an HSS of -16.59. This analysis highlights the performance of the DUNE AI model compared to the NOAA forecasts for this month.




Figure 10: Comparison of forecasts and observations for May 2021: (a) DUNE AI Inference with an HSS of 15.71, (b) ERA5 Observations, (c) NOAA Operational Forecast with an HSS of -16.59, (d) NOAA Observations. In May 2021, both DUNE AI and NOAA performed poorly in temperature forecasting. However, the DUNE AI model managed to capture above-normal temperatures in Florida, the northeastern coast, and parts of the western coast. Notably, in the central region, where below-normal temperatures were observed, the DUNE AI model inferred near-normal temperatures, while NOAA predicted above-normal temperatures.

Figure 11 presents a comparative analysis of the Heidke Skill Score (HSS) between NOAA operational forecasts and DUNE AI inferences over 5 years (2019-2023). This comparison highlights performance differences between the two forecasting approaches. The blue curve represents inferences from global DUNE AI results tailored to the United States. In contrast, the green curve represents the DUNE AI model trained by setting all global values over oceans and land, excluding the US, to zero. This is intended to focus training exclusively on the US. Over the five years, the DUNE AI model demonstrates superior performance in February, May, and October but shows reduced skill in March, June, November, and December. The results for January, April, July, August, and September are comparable between the two methods. It is worth noting that NOAA operational forecasts are provided at a $2^\circ \times 2^\circ$ resolution, while DUNE AI inferences are available at a finer $0.25^\circ \times 0.25^\circ$ resolution.




Figure 11: Heidke Skill Score (HSS) Comparison for Monthly Averaged Forecasts (2019-2023) (a) Comparison of monthly averaged HSS for NOAA operational forecasts, DUNE AI global inferences calculated just for the US, and DUNE AI inferences trained just for the US (b) Heatmap of the HSS for NOAA operational forecasts and DUNE AI inferences.

Since the DUNE global model was trained on ERA5 from 1980 to 2016, an additional experiment was performed by training DUNE from 1980 to 2020. Global inferences were carried out for 2021 and 2023 without any validation. Figure 12 compares the monthly Heidke Skill Score (HSS) for each month of the 3 years 2021, 2022, and 2023 between NOAA CPC operational Gridded Monthly Forecasts and DUNE AI inferences. The blue curve represents inferences derived from global DUNE AI results tailored to the United States, with the model trained on data up to 2016. The green curve indicates inferences derived from global DUNE AI results tailored to the United States, with the model trained on data up to 2020. The results suggest that training the DUNE AI model with more recent data improves the HSS. This may be a result of continued increased global monthly warming from the period 2017-2020. Based on these results, the DUNE AI model demonstrates comparable performance to NOAA operational forecasts in February, April, June, August, and December. The DUNE model exhibits improved performance in May, October, and November. However, it shows a lack of skill in March, June, and September.




Figure 12: Heidke Skill Score (HSS) Comparison for Monthly Averaged Forecasts (2021-2023) (a) Comparison of monthly averaged HSS for NOAA operational forecasts, DUNE AI inferences trained on data up to 2016, and DUNE AI inferences trained on data up to 2020. (b) Heatmap of the HSS for NOAA operational forecasts and DUNE AI inferences.

6.8 Ensemble Members Comparison

The ECMWF ERA5 monthly dataset comprises an ensemble of ten averaged monthly members, each with a spatial resolution of $0.5^\circ \times 0.5^\circ$ for all available years. To achieve a finer resolution of $0.25^\circ \times 0.25^\circ$, bilinear interpolation was employed to upscale the data. The DUNE AI model was trained on the original monthly-averaged t2m reanalysis dataset. After training, the model's inferred weights were applied to the respective extrapolated dataset for each ensemble member.

The testing period spanned from 2019 to 2023, covering 60 months. Given that there are 10 ensemble members per month and data for 5 years, the total number of sample months for training amounts to 600. The RMSE and ACC were computed for each ensemble member relative to the ERA5 monthly observations and plotted, showing the mean and standard deviation of the ensemble alongside the monthly mean of the original DUNE-inferred RMSE and ACC. Figure 13 illustrates that the variance in RMSE and ACC of the ensemble forecast is only marginally different from the original stand-alone forecasts, with some noticeable deviations exceeding one standard deviation from July to September for ACC and from August to October for RMSE. The conclusion drawn from this analysis is that interpolating ensemble member data from a lower to a higher spatial resolution does not significantly differ from the mean sample of five

years. This suggests that extrapolating T2m from the ERA5 resolution to 0.01 degrees may not substantially differ from observations at that higher resolution.




Figure 13: (a) Comparison of RMSE: 10-Member Ensemble vs. Standalone Inference for Averaged Inferences (2019-2023). (b) Comparison of ACC: 10-Member Ensemble vs. Standalone Inference for Averaged Inferences (2019-2023)

7 Discussions

Examining ERA5 data from 1940, observations indicate there is less global warming in temperature from 1940-1979 than appears in the data from 1990 to 2023. Based on these trends, it was not clear which portion of the dataset should be chosen as climatology for use in forming the anomalies in the training of DUNE. As a result, we conducted experiments, not shown herein, that trained the DUNE model on data from the following periods: 1940-2016, 1950-2016, 1960-2016, 1970-2016, and 1980-2016. The results showed that training from 1980-2016 yielded the most accurate global statistical results using the climatology (1950-1979).

Combining SST data over oceans with T2M data over land produced improved DUNE forecast performance, particularly in coastal regions. This improvement can be attributed to SST data providing more accurate and relevant information over oceans, capturing oceanic influences and variations more effectively. Conversely, T2M data is more precise over land areas, reflecting terrestrial temperature changes more accurately. By leveraging the strengths of both datasets in their respective domains, the model better accounts for the complex interactions between oceanic and terrestrial climates, leading to enhanced prediction accuracy. As a result, SST over oceans was adopted in all subsequent experiments reported.

Incorporating constant fields into the model training significantly enhanced the model's overall forecasting capability. These constant fields provide essential static information that aids the model in understanding the underlying physical geography and land surface characteristics, which are crucial for accurate climate predictions. To understand their individual effects, each constant was run as a separate experiment with a separate channel. This approach allowed us to isolate and evaluate the contribution of each field. It was concluded that using the six constant fields helped DUNE better generalize and improve its forecasting results by providing these static environmental factors.

Experiments using the sliding window data from the prior 12, 6, 4, 3, 2, and 1 months to predict the next month's mean temperature did not significantly deteriorate model performance. This method allows the model to capture temporal patterns and trends over longer periods, which is beneficial for multi-month forecasting. As a result, it was included in the presentation here of the DUNE model. This served as an addition of an aspect of attention that is planned for the next version of DUNE. Based on this, it was decided to use only the prior month to forecast the next month, as it provided the most pertinent and actionable data for our predictions. However, for the moving window approach, we used data from the prior x months to forecast the temperatures for the next x months.

Incorporating the average of the four internal ensembles into the DUNE AI model significantly improved forecast accuracy, particularly during winter months when temperature variability is the highest over land. The availability of ten ensemble forecasts allowed for testing the impact of external ensemble forecasting in minutes, which we demonstrated helps address the impact of model uncertainty errors and captures a broader range of possible outcomes. Adding this capability in the future will enhance the robustness and reliability of our forecasts.

In analyzing the monthly mean results published by Taylor and Feng (2022), it was found that the DUNE model performed equally well, if not better, in predicting SST over the oceans. This comparison highlights the robustness and accuracy of the model in handling oceanic temperature predictions. It's important to note that while Taylor and Feng

(2022)’s research focused solely on forecasting SST, the results presented here include both land and oceans, offering a more comprehensive approach. This dual focus allows our model to capture interactions between terrestrial and marine climates, further enhancing its predictive capabilities.

In comparing the forecast results of Dune with those of the NOAA Climate Prediction Center Gridded Monthly Temperature Outlook Product, the extensive methodology involved in its operational production should be noted. For example, NOAA monthly forecasts are based on a number of tools, including ensembles of dynamical models, statistical regression forecasts, pattern correlation anomalies, extensive supplementary data sets, and hybrid models involving many organizational outputs. Based on the favorable comparisons of the simple DUNE regional monthly forecast results over the US shown in section 6.6 at considerably reduced time and resources with those of the operational products, it would indicate a strong potential need to include the DUNE AI product or an equivalent AI method among the complement of other components going into the NOAA Gridded Monthly Temperature Outlook.

An interesting serendipitous ensemble forecast result was the conduct of an experiment with an ensemble of 10 DUNE forecasts from 10 ERA5 monthly reanalyses at ERA5 reanalysis 0.5 degrees extrapolated to 0.25 degrees that exhibited comparable forecast accuracy with forecasts initialized from the ERA5 0.25 degrees. This might indicate the ability of DUNE AI to be independent of the need for resolution retraining when applied to extrapolated higher-resolution data sets.

In a note of caution, it should be pointed that a limitation inherent to AI/ML climate forecasting relative to physics-based forecasting is its dependence on the training data set selected. DUNE employs atmospheric states of monthly anomalies, which depend on the choice of a climatology period. However, monthly atmospheric states are not in equilibrium, as Figure 13 shows that all months are experiencing global warming. This is particularly true when applied to regional predictions for different domains whose inputs were trained as anomalies based on a specific set of years for defining a climate. The experiment shown in Figure 12 reflects the improved HSS accuracy that can be obtained by selecting more recent climatology for training. In addition, the impact on regional forecast performance can be improved substantially in terms of HSS over the US when trained with data just over the US rather than globally. Future regional forecasts will need to be separately trained.




Figure 14: Mean-Cosine-Weighted Monthly averaged global temperatures derived from ERA5. It shows the increase in global temperatures over the years.

8 Conclusion

A deep learning UNet++-based self-forming ensemble AI model, DUNE, was developed to reduce forecast uncertainty during periods of atmospheric variability in monthly or seasonal near-surface air temperature (T2m) forecasts over global land-based regional domains. The efficacy of DUNE was assessed by comparing its monthly and seasonal RMSE, ACC, and HSS statistics with those of persistence, climatology, and other published physics-based and AI forecast models over multiple years. In all cases, DUNE’s error statistics were comparable to or outperformed those of other forecast methods. Additionally, this study presents favorable comparative RMSE, ACC, and HSS error statistics for monthly gridded 2-meter temperature forecasts over the US and other regional domains for the most recent five years, at a resolution of 0.25 degrees compared with the NOAA operational gridded monthly Outlook at 2-degree resolution. The results show that the DUNE AI model’s performance is comparable to NOAA’s forecasts, as evaluated against

ERA5 monthly observations. In addition, for the first time, AI-based Seasonal and Annual forecasts are presented for 2024 pending verification.

Data Availability Statement

The ERA5 monthly mean data used in our study is available for download from Copernicus Climate Change Service <https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-pressure-levels-monthly-means?tab=overview>. This dataset includes all the variables required for our DUNE AI model. The NOAA monthly forecasts are available on their website at https://www.cpc.ncep.noaa.gov/products/predictions/long_range/tools/briefing/mon_veri.grid.php. The code for training the monthly mean, seasonal mean, annual mean, and moving window approaches is available on GitHub: https://github.com/Pratik-Shukla-1/DUNE_AI_Model.

Acknowledgments

We want to express our gratitude to the NASA/ESTO Firesense Program Manager, Haris Riris, and his staff for their support on our grant number 80NSSC22K1405, which has allowed graduate student support to pursue this critical research project. Their belief in the significance of this AI/ML study in advancing knowledge in seasonal and annual prediction will be instrumental in addressing the impact and risk of wildfires on climate change.

We would also like to acknowledge the support from the NASA HPC office in making the GSFC/NCCS computing facility available to this grant, without which these breakthrough findings would not have been possible. Moreover, we would like to thank the NCCS staff for their professional management in providing reliable system support and access to their advanced machine learning system. That support has been vital in facilitating our data aggregation, training, and predictions of the extensive number of experiments needed to test and evaluate this unique DUNE AI/ML Earth system forecasting model.

Finally, we wish to recognize Prof. Karuna Joshi and the support provided by the UMBC NSF-funded Center for Accelerated Real Time Analytics (CARTA) for providing their Computing and Laboratory resources to support this machine learning collaborative research study.

References

- Abien Fred Agarap. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018.
- HM Arnold, IM Moroz, and TN Palmer. Stochastic parametrizations and model uncertainty in the lorenz’96 system. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1991): 20110479, 2013.
- Pierre Baldi. Autoencoders, unsupervised learning, and deep architectures. In *Proceedings of ICML workshop on unsupervised and transfer learning*, pages 37–49. JMLR Workshop and Conference Proceedings, 2012.
- Peter Bauer, Alan Thorpe, and Gilbert Brunet. The quiet revolution of numerical weather prediction. *Nature*, 525 (7567):47–55, 2015.
- Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.
- Canadian Meteorological Centre (CMC). *Global Environmental Multiscale (GEM) Model*. Canadian Meteorological Centre, 2020. URL https://weather.gc.ca/grib/gem_e.html.
- WE Chapman, AC Subramanian, L Delle Monache, SP Xie, and FM Ralph. Improving atmospheric river forecasts with machine learning. *Geophysical Research Letters*, 46(17-18):10627–10635, 2019.
- Jule Gregory Charney. the feasibility of a global observation and analysis experiment,. *Bulletin of the American Meteorological Society*, 47(3):200 – 221, 1966. doi: 10.1175/1520-0477-47.3.200. URL https://journals.ametsoc.org/view/journals/bams/47/3/1520-0477-47_3_200.xml.
- Ashesh Chattopadhyay, Ebrahim Nabizadeh, and Pedram Hassanzadeh. Analog forecasting of extreme-causing weather patterns using deep learning. *Journal of Advances in Modeling Earth Systems*, 12(2):e2019MS001958, 2020.
- Lei Chen, Fei Du, Yuan Hu, Zhibin Wang, and Fan Wang. Swinrdm: integrate swinrrnn with diffusion model towards high-resolution and high-quality weather forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 322–330, 2023a.

- Lei Chen, Xiaohui Zhong, Feng Zhang, Yuan Cheng, Yinghui Xu, Yuan Qi, and Hao Li. Fuxi: a cascade machine learning forecasting system for 15-day global weather forecast. *npj Climate and Atmospheric Science*, 6(1):190, 2023b.
- PHILIPPE Courtier, J-N Thépaut, and Anthony Hollingsworth. A strategy for operational implementation of 4d-var, using an incremental approach. *Quarterly Journal of the Royal Meteorological Society*, 120(519):1367–1387, 1994.
- Matteo De Felice, Andrea Alessandri, and Franco Catalano. Seasonal climate forecasts for medium-term electricity demand forecasting. *Applied Energy*, 137:435–444, 2015.
- F Doblas-Reyes, Anice Garcia, James Hansen, Luigi Mariani, Ajeet Nain, Kulasekaran Ramesh, and R Venkataraman. Weather and climate forecasts for agriculture. *Guide to agricultural, meteorological practices*, 57, 2003.
- Peter D Dueben and Peter Bauer. Challenges and design choices for global weather and climate models based on machine learning. *Geoscientific Model Development*, 11(10):3999–4009, 2018.
- Elizabeth Ebert, Laurie Wilson, Andreas Weigel, Marion Mittermaier, Perti Nurmi, Philip Gill, Martin Göber, Susan Joslyn, Barbara Brown, Tressa Fowler, et al. Progress and challenges in forecast verification. *Meteorological Applications*, 20(2):130–139, 2013.
- European Centre for Medium-Range Weather Forecasts (ECMWF). ECMWF Models. *European Centre for Medium-Range Weather Forecasts*, 2020. URL <https://www.ecmwf.int/en/forecasts/datasets>.
- David L Feldman and Helen M Ingram. Making science useful to decision makers: climate forecasts, water management, and knowledge networks. *Weather, Climate, and Society*, 1(1):9–21, 2009.
- Brigitta Goger, Mathias W Rotach, Alexander Gohm, Ivana Stiperski, and Oliver Fuhrer. Current challenges for numerical weather prediction in complex terrain: Topography representation and parameterizations. In *2016 International Conference on High Performance Computing & Simulation (HPCS)*, pages 890–894. IEEE, 2016.
- Peter Grönquist, Chengyuan Yao, Tal Ben-Nun, Nikoli Dryden, Peter Dueben, Shigang Li, and Torsten Hoefler. Deep learning for post-processing ensemble weather forecasts. *Philosophical Transactions of the Royal Society A*, 379(2194):20200092, 2021.
- John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, and Bryan Catanzaro. Adaptive fourier neural operators: Efficient token mixers for transformers. *arXiv preprint arXiv:2111.13587*, 2021.
- Guy Hadash, Einat Kermany, Boaz Carmeli, Ofer Lavi, George Kour, and Alon Jacovi. Estimate and replace: A novel approach to integrating deep neural networks with existing applications. *arXiv preprint arXiv:1804.09028*, 2018.
- Milton Halem and Robert Jastrow. Analysis of garp data requirements. *Journal of Atmospheric Sciences*, 27(1):177–177, 1970.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, et al. The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020.
- Sepp Hochreiter. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116, 1998.
- S. Hoyer and J. Hamman. xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software*, 5(1), 2017. doi: 10.5334/jors.148. URL <https://doi.org/10.5334/jors.148>.
- Yuan Hu, Lei Chen, Zhibin Wang, and Hao Li. Swinvrnn: A data-driven ensemble forecasting model via learned distribution perturbation. *Journal of Advances in Modeling Earth Systems*, 15(2):e2022MS003211, 2023.
- Di Huang, Shuaian Wang, and Zhiyuan Liu. A systematic review of prediction methods for emergency management. *International Journal of Disaster Risk Reduction*, 62:102412, 2021.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.

- James W Jones, James W Hansen, Frederick S Royce, and Carlos D Messina. Potential benefits of climate forecasting to agriculture. *Agriculture, ecosystems & environment*, 82(1-3):169–184, 2000.
- Eugenia Kalnay, Stephen J Lord, and Ronald D McPherson. Maturity of operational numerical weather prediction: Medium range. *Bulletin of the American Meteorological Society*, 79(12):2753–2770, 1998.
- Soo Kyung Kim, Sasha Ames, Jiwoo Lee, Chengzhu Zhang, Aaron C Wilson, and Dean Williams. Massive scale deep learning for detecting extreme climate events. *Climate Informatics*, 5, 2017.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- George Kour and Raid Saabne. Fast classification of handwritten on-line arabic characters. In *Soft Computing and Pattern Recognition (SoCPar)*, 2014 6th International Conference of, pages 312–318. IEEE, 2014a.
- George Kour and Raid Saabne. Real-time segmentation of on-line handwritten arabic script. In *Frontiers in Handwriting Recognition (ICFHR)*, 2014 14th International Conference on, pages 417–422. IEEE, 2014b.
- T. N. Krishnamurti, C. M. Kishtawal, Z. Zhang, T. E. LaRow, D. R. Bachiochi, E. C. Williford, and S. Gadgil. Multimodel ensemble forecasts for weather and seasonal climate. *Journal of Climate*, 16(12):1821–1836, 2003.
- Robert J Kuligowski and Ana P Barros. Localized precipitation forecasts from a numerical weather prediction model using artificial neural networks. *Weather and forecasting*, 13(4):1194–1204, 1998.
- Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, et al. Graphcast: Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794*, 2022.
- Pablo Rozas Larraondo, Luigi J Renzullo, Inaki Inza, and Jose A Lozano. A data-driven approach to precipitation parameterizations using convolutional encoder-decoder neural networks. *arXiv preprint arXiv:1903.10274*, 2019.
- Philippe Lauret, Hadja Maïmouna Diagne, and Mathieu David. A neural network post-processing approach to improving nwp solar radiation forecasts. *Energy Procedia*, 57:1044–1052, 2014.
- Wentao Li, Baoxiang Pan, Jiangjiang Xia, and Qingyun Duan. Convolutional neural network-based statistical post-processing of ensemble precipitation forecasts. *Journal of hydrology*, 605:127301, 2022.
- Yunjie Liu, Evan Racah, Joaquin Correa, Amir Khosrowshahi, David Lavers, Kenneth Kunkel, Michael Wehner, William Collins, et al. Application of deep convolutional neural networks for detecting extreme weather in climate datasets. *arXiv preprint arXiv:1605.01156*, 2016.
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- E. N. Lorenz. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2):130–141, 1963a.
- Edward N Lorenz. Deterministic nonperiodic flow. *Journal of atmospheric sciences*, 20(2):130–141, 1963b.
- Edward N Lorenz. The predictability of hydrodynamic flow. *Trans. NY Acad. Sci*, 25(4):409–432, 1963c.
- Edward N Lorenz. Atmospheric predictability experiments with a large numerical model. *Tellus*, 34(6):505–513, 1982.
- Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- M. P. Mittermaier. Skill and predictability in weather forecasting: a probabilistic approach. *Meteorological Applications*, 15(1):21–32, 2008.
- Mrutyunjay Mohapatra. Climate and weather forewarning systems for disaster preparedness and response. In *Disaster Risk and Management Under Climate Change*, pages 39–54. Springer, 2024.
- National Centers for Environmental Prediction (NCEP). Global Forecast System (GFS). Technical report, National Oceanic and Atmospheric Administration (NOAA), 2024. URL <https://www.emc.ncep.noaa.gov/gfs/>.
- Munir Ahmad Nayak and Subimal Ghosh. Prediction of extreme rainfall event using weather pattern recognition and support vector machine classifier. *Theoretical and applied climatology*, 114:583–603, 2013.
- Tung Nguyen, Johannes Brandstetter, Ashish Kapoor, Jayesh K Gupta, and Aditya Grover. Climax: A foundation model for weather and climate. *arXiv preprint arXiv:2301.10343*, 2023.
- NOAA Heidke Skill Score Explanation. NOAA Heidke Skill Score Explanation. Technical report, National Oceanic and Atmospheric Administration (NOAA), 2024. URL <https://www.cpc.ncep.noaa.gov/products/verification/summary/index.php?page=tutorial>.
- NOAA Operation Forecasts Blogs. NOAA Operational Forecasts Blog. Technical report, National Oceanic and Atmospheric Administration (NOAA), 2024. URL <https://www.climate.gov/news-features/blogs>.

- LA Ogallo, MS Boulahya, and T Keane. Applications of seasonal to interannual climate prediction in agricultural planning and operations. *Agricultural and Forest Meteorology*, 103(1-2):159–166, 2000.
- Anton Orlov, Jana Sillmann, and Ilaria Vigo. Better seasonal forecasts for the renewable energy industry. *Nature Energy*, 5(2):108–110, 2020.
- T. N. Palmer, R. Buizza, F. J. Doblas-Reyes, T. Jung, M. Leutbecher, and G. J. Shutts. Stochastic parametrization and model uncertainty. *ECMWF Seminar on Predictability*, 3:3–7, 2005.
- Tim Palmer. The real butterfly effect and maggoty apples. *Physics Today*, 77(5):30–35, 05 2024. ISSN 0031-9228. doi: 10.1063/pt.eike.hsbz. URL <https://doi.org/10.1063/pt.eike.hsbz>.
- Timothy Noel Palmer. Predicting uncertainty in numerical weather forecasts. In *International Geophysics*, volume 83, pages 3–13. Elsevier, 2002.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, et al. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Di Qi and Andrew J Majda. Using machine learning to predict extreme events in complex systems. *Proceedings of the National Academy of Sciences*, 117(1):52–59, 2020.
- Stephan Rasp and Sebastian Lerch. Neural networks for postprocessing ensemble weather forecasts. *Monthly Weather Review*, 146(11):3885–3900, 2018.
- Stephan Rasp and Nils Thuerey. Data-driven medium-range weather prediction with a resnet pretrained on climate simulations: A new model for weatherbench. *Journal of Advances in Modeling Earth Systems*, 13(2):e2020MS002405, 2021.
- Stephan Rasp, Peter D Dueben, Sebastian Scher, Jonathan A Weyn, Soukayna Mouatadid, and Nils Thuerey. Weatherbench: a benchmark data set for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, 12(11):e2020MS002203, 2020.
- Steve Rayner, Denise Lach, and Helen Ingram. Weather forecasts are for wimps: why water resource managers do not use climate forecasts. *Climatic change*, 69(2):197–227, 2005.
- Eduardo Rocha Rodrigues, Igor Oliveira, Renato Cunha, and Marco Netto. Deepdownscale: A deep learning strategy for high-resolution weather forecast. In *2018 IEEE 14th International Conference on e-Science (e-Science)*, pages 415–422. IEEE, 2018.
- Adrian Rojas-Campos, Martin Wittenbrink, Pascal Nieters, Erik J Schaffernicht, Jan D Keller, and Gordon Pipa. Postprocessing of nwp precipitation forecasts using deep learning. *Weather and Forecasting*, 38(3):487–497, 2023.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III* 18, pages 234–241. Springer, 2015.
- Sebastian Scher. Toward data-driven weather and climate forecasting: Approximating a simple general circulation model with deep learning. *Geophysical Research Letters*, 45(22):12–616, 2018.
- Sebastian Scher and Gabriele Messori. Predicting weather forecast uncertainty with machine learning. *Quarterly Journal of the Royal Meteorological Society*, 144(717):2830–2841, 2018.
- Bo-Wen Shen, Roger A Pielke Sr, Xubin Zeng, and Xiping Zeng. Lorenz’s view on the predictability limit of the atmosphere. *Encyclopedia*, 3(3):887–899, 2023.
- Pratik Shukla. The gradient descent algorithm, 2022. URL <https://pub.towardsai.net/the-gradient-descent-algorithm-defddd1d312e>.
- Adrian J Simmons and Anthony Hollingsworth. Some aspects of the improvement in skill of numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, 128(580):647–677, 2002.
- L. A. Smith, H. Stern, and F. Giorgi. Improvements in weather prediction skill resulting from co2-induced climate changes. *Science*, 299(5615):1556–1559, 2017.

- John Taylor and Ming Feng. A deep learning model for forecasting global monthly mean sea surface temperature anomalies. *Frontiers in Climate*, 4:932932, 2022.
- Zoltan Toth and Roberto Buizza. Weather forecasting: What sets the forecast skill horizon? In *Sub-Seasonal to Seasonal Prediction*, pages 17–45. Elsevier, 2019.
- Tapas Tripura and Souvik Chakraborty. Wavelet neural operator for solving parametric partial differential equations in computational mechanics problems. *Computer Methods in Applied Mechanics and Engineering*, 404:115783, 2023.
- Diana Udel. Um-led experiment closes critical gap in weather forecasting., 2019. URL <https://news.miami.edu/rosenstiel/stories/2019/12/um-led-experiment-closes-critical-gap-in-weather-forecasting.html>.
- Alper Unal, Busra Asan, Ismail Sezen, Bugra Yesilkaynak, Yusuf Aydin, Mehmet Ilicak, and Gozde Unal. Climate model-driven seasonal forecasting approach with deep learning. *Environmental Data Science*, 2:e29, 2023.
- United Kingdom Met Office. *Global and Regional Ensemble Prediction System (MOGREPS)*. United Kingdom Met Office, 2020. URL <https://www.metoffice.gov.uk/research/weather/ensemble-forecasting>.
- H. M. Van den Dool. Searching for analogues, how long must we wait? *Tellus A*, 41(4):324–335, 1989.
- Christian Viel, Anne-Lise Beaulant, Jean-Michel Soubeyroux, and Jean-Pierre Céron. How seasonal forecast could help a decision maker: an example of climate service for water resource management. *Advances in Science and Research*, 13:51–55, 2016.
- Bin Wang, Jie Lu, Zheng Yan, Huaishao Luo, Tianrui Li, Yu Zheng, and Guangquan Zhang. Deep uncertainty quantification: A machine learning approach for weather forecasting. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2087–2095, 2019.
- Jonathan A Weyn, Dale R Durran, and Rich Caruana. Can machines learn to predict weather? using deep learning to predict gridded 500-hpa geopotential height from historical weather data. *Journal of Advances in Modeling Earth Systems*, 11(8):2680–2693, 2019.
- Jonathan A Weyn, Dale R Durran, and Rich Caruana. Improving data-driven global weather prediction using deep convolutional neural networks on a cubed sphere. *Journal of Advances in Modeling Earth Systems*, 12(9):e2020MS002109, 2020.
- Jonathan A Weyn, Dale R Durran, Rich Caruana, and Nathaniel Cresswell-Clay. Sub-seasonal forecasting with a large ensemble of deep-learning weather prediction models. *Journal of Advances in Modeling Earth Systems*, 13(7):e2021MS002502, 2021.
- Nedjeljka Žagar. A global perspective of the limits of prediction skill of nwp models. *Tellus A: Dynamic Meteorology and Oceanography*, 69(1):1317573, 2017.
- Matthew D Zeiler, Dilip Krishnan, Graham W Taylor, and Rob Fergus. Deconvolutional networks. In *2010 IEEE Computer Society Conference on computer vision and pattern recognition*, pages 2528–2535. IEEE, 2010.
- Fuqing Zhang, Y Qiang Sun, Linus Magnusson, Roberto Buizza, Shian-Jiann Lin, Jan-Huey Chen, and Kerry Emanuel. What is the predictability limit of midlatitude weather? *Journal of the Atmospheric Sciences*, 76(4):1077–1091, 2019.
- Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 3–11. Springer, 2018.

A DUNE AI Architecture




Figure A1: DUNE AI Extended Architecture: The input number of channels is adjusted to forecast the monthly mean temperature one month in advance, depending on the number of time steps (months) to be forecasted. Details are provided in Table B1. We utilize seven input channels for monthly mean predictions starting one month in advance: temperature, LSM, SLT, Orography, TISR, CVH, and CVL. At the initial level, the input dimensions are 720 x 1440. As the network progresses deeper, the dimensions are halved while the number of channels increases. In the decoding phase of the architecture, the input dimensions are doubled at each level, and the number of channels at each level is indicated within the rectangular boxes.

B Extended-Range Forecast Results

Figure B1 illustrates the application of the moving window approach to forecast temperatures three months in advance. This methodology can be similarly extended to forecast temperatures any number of months in advance. The comparative results are presented in section 6.6.




Figure B1: The moving window approach shows the forecast three months in advance. It demonstrates how we can use the ground truth of months 1, 2, and 3 to forecast the temperature for month 4*. Then, it uses months 2, 3, and 4* to forecast the temperature for month 5*. Finally, it uses months 3, 4*, and 5* to forecast the temperature for month 6*. The moving window approach demonstrates how it uses the ground truth of months 1, 2, and 3 to forecast the temperature for months 4*, 5*, and 6*. In this figure, * represents the forecasted month.

Table B1 shows the input number of channels with the respective number of months to be forecasted.

Forecast Window (Months)	Input Channels	Description
1	7	1: temperature, 1: tisir, 5: constants
2	9	2: temperature, 2: tisir, 5: constants
3	11	3: temperature, 3: tisir, 5: constants
4	13	4: temperature, 4: tisir, 5: constants
6	17	6: temperature, 6: tisir, 5: constants
12	29	12: temperature, 12: tisir, 5: constants

Table B1: Forecast Window (Months) and No. of Input Channels

Figures B2 and B3 display the land RMSE and land ACC for the 12-month moving window over the test period from 2019 to 2023. The DUNE AI model continues to outperform the baseline methods.

Figure B4 shows the forecast of T2m over the US for JJA and SON 2024. These forecast results are pending verification.




Figure B2: Monthly Mean Forecast (12-Month Moving Window Approach) Land RMSE: (a) Persistence (PYSM), (b) Climatology, and (c) DUNE AI Method. The comparison reveals that the DUNE AI model consistently achieves lower RMSE values, indicating superior forecasting accuracy. Notably, for the record warm year of 2023, the DUNE AI model exhibited relatively small variability in its year-ahead forecasts compared to the baseline models.




Figure B3: Monthly Mean Forecast (12-Month Moving Window Approach) Land ACC: (a) Persistence (PYSM), (b) Climatology, and (c) DUNE AI Method. The comparison demonstrates that the DUNE AI model consistently achieves higher ACC values, indicating superior forecasting skills. Notably, for the record warm year of 2023, the DUNE AI model exhibited relatively small variability in its forecasts made a year in advance compared to the baseline models.




Figure B4: DUNE AI Inference using the 12-month moving window approach for 2024. (a) DUNE AI inference for June-July-August (JJA) 2024 (6 months in advance). (b) DUNE AI inference for September-October-November (SON) 2024 (9 months in advance).

C Seasonal Mean Forecasts

Figure C1 and Figure C2 show the global RMSE and global ACC for our test period from 2019 to 2023.




Figure C1: Global RMSE of Seasonal Mean Forecasts: (a) Persistence (PYSS), (b) Climatology, and (c) DUNE AI Method. The comparison shows that the DUNE AI model consistently achieves lower RMSE values, indicating superior forecasting skills. Notably, during the record warm year of 2023, the DUNE AI model exhibited relatively small variability compared to the baseline models.




Figure C2: Global ACC of Seasonal Mean Forecasts: (a) Persistence (PYSS), (b) Climatology, and (c) DUNE AI Method. The comparison shows that the DUNE AI model consistently achieves higher ACC values, indicating superior forecasting skills. Notably, during the record warm year of 2023, the DUNE AI model exhibited relatively small variability compared to the baseline models.

Figure C3 shows the visual comparison of the DUNE AI model forecast errors with the baseline methods for summer 2023.




Figure C3: Visualization of forecast errors for summer (JJA) 2023 (a) Persistence (PS) (b) Persistence (PYSS) (c) Climatology (d) DUNE AI Method. The DUNE AI model outperforms other baselines in North America, Africa, Australia, Eurasia, and the El Niño region near South America.

D Yearly Mean Forecasts

Figure D1 and Figure D2 show the global RMSE and global ACC for our test period from 2019 to 2023.




Figure D1: Yearly Mean Forecast Global RMSE (a) Persistence (b) Climatology (c) DUNE AI Method

Figure D3 shows the visual comparison of the DUNE AI model forecast errors with the baseline methods for 2021 globally.

Figure D4 shows the annual mean forecast errors for the boreal region in 2023.




Figure D2: Yearly Mean Forecast Global ACC (a) Persistence (b) Climatology (c) DUNE AI Method




Figure D3: Visualization of forecast errors for 2021: (a) Persistence, (b) Climatology, and (c) DUNE AI Method. The DUNE AI method shows reduced errors in Eurasia, Africa, and North America. Using annual mean data results in better performance compared to monthly mean forecasts averaged over 12 months due to the reduced variability in annual mean predictions.




Figure D4: Visualization of annual mean forecast errors for 2023 in the Boreal region (a) Persistence, (b) Climatology, and (c) DUNE AI Method. The DUNE AI method shows the reduced errors in the Boreal region across Eurasia and North America. It is worth noting that 2023 was the warmest year on record.