



# Beyond Ensemble Averages: Leveraging Climate Model Ensembles for Subseasonal Forecasting

Elena Orlova,<sup>a</sup> Haokun Liu,<sup>a</sup> Raphael Rossellini,<sup>b</sup> Benjamin A. Cash,<sup>c</sup> and Rebecca Willett<sup>a,b</sup>

<sup>a</sup> *Department of Computer Science, University of Chicago, Chicago, USA*

<sup>b</sup> *Department of Statistics, University of Chicago, Chicago, USA*

<sup>c</sup> *Department of Atmospheric, Oceanic, and Earth Sciences, George Mason University, Fairfax, USA*

*Corresponding author:* Elena Orlova, eorlova@uchicago.edu

**Early Online Release:** This preliminary version has been accepted for publication in *Artificial Intelligence for the Earth Systems*, may be fully cited, and has been assigned DOI 10.1175/AIES-D-23-0103.1. The final typeset copyedited article will replace the EOR at the above DOI when it is published.

© 2024 American Meteorological Society. This is an Author Accepted Manuscript distributed under the terms of the default AMS reuse license. For information regarding reuse and general copyright information, consult the AMS Copyright Policy ([www.ametsoc.org/PUBSReuseLicenses](http://www.ametsoc.org/PUBSReuseLicenses)).

ABSTRACT: Producing high-quality forecasts of key climate variables, such as temperature and precipitation, on subseasonal time scales has long been a gap in operational forecasting. This study explores an application of machine learning (ML) models as post-processing tools for subseasonal forecasting. Lagged numerical ensemble forecasts (i.e., an ensemble where the members have different initialization dates) and observational data, including relative humidity, pressure at sea level, and geopotential height, are incorporated into various ML methods to predict monthly average precipitation and two-meter temperature two weeks in advance for the continental United States. For regression, quantile regression, and tercile classification tasks, we consider using linear models, random forests, convolutional neural networks, and stacked models (a multi-model approach based on the prediction of the individual ML models). Unlike previous ML approaches that often use ensemble mean alone, we leverage information embedded in the ensemble forecasts to enhance prediction accuracy. Additionally, we investigate extreme event predictions that are crucial for planning and mitigation efforts. Considering ensemble members as a collection of spatial forecasts, we explore different approaches to using spatial information. Trade-offs between different approaches may be mitigated with model stacking. Our proposed models outperform standard baselines such as climatological forecasts and ensemble means. In addition, we investigate feature importance, trade-offs between using the full ensemble or only the ensemble mean, and different modes of accounting for spatial variability.

Significance statement: Accurately forecasting temperature and precipitation on subseasonal time scales – two weeks to two months in advance – is extremely challenging. These forecasts would have immense value in agriculture, insurance, and economics. Our paper describes an application of machine learning techniques to improve forecasts of monthly average precipitation and 2-meter temperature using lagged physics-based predictions and observational data two weeks in advance for the entire continental United States. For lagged ensembles, the proposed models outperform standard benchmarks such as historical averages and averages of physics-based predictions. Our findings suggest that utilizing the full set of physics-based predictions instead of the average enhances the accuracy of the final forecast.

## 1. Introduction

High-quality forecasts of key climate variables such as temperature and precipitation on subseasonal time scales, defined here as the time range between two weeks and two months, have long been a gap in operational forecasting (Ban et al. 2016). Advances in weather forecasting on time scales of days to about a week (Lorenc 1986; National Academies of Sciences 2016; National Research Council 2010; Simmons and Hollingsworth 2002) or seasonal forecasts on time scales of two to nine months (Barnston et al. 2012) do not necessarily translate to the challenging subseasonal regime. Addressing the crucial need for forecasts on the seasonal-to-subseasonal (S2S) timescale, collaborative initiatives led by the World Weather Research Programme and the World Climate Research Programme aim to advance S2S forecasting by focusing on mesoscale–planetary-scale interactions, high-resolution simulations, data assimilation methods, and tailored socioeconomic support (Brunet et al. 2010). Skillful forecasts on subseasonal time scales would have immense value in agriculture, insurance, and economics (White et al. 2022; Mouatadid et al. 2023). The importance of improved subseasonal predictions is also detailed by Ban et al. (2016) and National Research Council (2010).

The National Centers for Environmental Prediction (NCEP), part of the National Oceanic and Atmospheric Administration (NOAA), currently issues a “week 3-4 outlook” for the contiguous United States (CONUS).<sup>1</sup> The NCEP outlooks are constructed using a combination of dynamical and statistical forecasts, with statistical forecasts based largely on how conditions in the past have varied (linearly) with indices of the El Niño-Southern Oscillation (ENSO), Madden-Julian

---

<sup>1</sup><https://www.cpc.ncep.noaa.gov/products/predictions/WK34/>

Oscillation (MJO), and global warming (i.e., the 30-year trend). There exists great potential to advance subseasonal forecasting (SSF) using machine learning (ML) techniques. Haupt et al. (2021) provides an overview of using ML methods for post-processing of numerical weather predictions. Vannitsem et al. (2021) highlight the crucial role of statistical post-processing techniques, including ML methods, in national meteorological services. They discuss theoretical developments and operational applications, current challenges, and potential future directions, particularly focusing on translating research findings into operational practices. A real-time forecasting competition called the Subseasonal Climate Forecast Rodeo (Hwang et al. 2018), sponsored by the Bureau of Reclamation in partnership with NOAA, USGS, and the U.S. Army Corps of Engineers, illustrated that teams using ML techniques can outperform forecasts from NOAA’s operational seasonal forecast system.

Here, we present work focused on developing ML-based forecasts that leverage lagged ensembles (i.e., an ensemble whose members are initialized from a succession of different start dates) of forecasts produced by NCEP in addition to observed data and other features. Previous studies, including successful methods in the Rodeo competition (e.g., Hwang et al. (2019)), incorporate the ensemble mean as a feature in their ML systems but do not use any other information about the ensemble. In other words, variations among the ensemble members are not reflected in the training data or incorporated into the learned model. In contrast, this paper *demonstrates that the full ensemble contains important information for subseasonal forecasting outside the ensemble mean*. Specifically, we consider the test case of predicting monthly 2-meter temperatures and precipitation two weeks in advance over 3000 locations over the continental United States using physics-based predictions, such as NCEP-CFSv2 hindcasts (Kirtman et al. 2014; Saha et al. 2014), using an ensemble of 24 distinct forecasts. We repeat this experiment for the Global Modeling and Assimilation Office from the National Aeronautics and Space Administration (NASA-GMAO) ensemble, which has 11 ensemble members (Nakada et al. 2018).

In this context, this paper makes the following contributions:

- We train a variety of ML models (including neural networks, random forests, linear regression, and model stacking) that input all ensemble member predictions as features in addition to observations of geopotential heights, relative humidity, precipitation, and temperature from past months to produce new forecasts with higher accuracy than the ensemble mean; forecast

accuracy is measured with a variety of metrics (Section 7). These models are considered in the context of regression, quantile regression, and tercile classification. Systematic experiments are used to characterize the influence of individual ensemble members on predictive skill (Section 8a(ii)).

- The collection of ML models employed allows us to consider different modes of accounting for spatial variability. ML models can account for spatial correlations among both features and targets; for example, when predicting Chicago precipitation, our models can leverage not only information about Chicago but also about neighboring regions. Specifically, we consider the following learning frameworks: (a) training a predictive model for each spatial location independently; (b) training a predictive model that inputs the spatial location as a feature and hence can be applied to any single spatial location; (c) training a predictive model for the full spatial map of temperature or precipitation – i.e., predicting an outcome for all spatial locations simultaneously. ML models present various ways to account for spatial variability, each with distinct advantages and disadvantages. Our application of model stacking (an ML technique where multiple models are combined, with their predictions used as input features for another model that produces the final prediction) allows our final learned model to exploit the advantages of each method.
- We conduct a series of experiments to help explain the learned model and which features the model uses most to make its predictions. We systematically explore the impact of using lagged observational data in addition to ensemble forecasts and positional encoding to account for spatial variations (Section 8c).
- The ensemble of forecasts from a physics-based model (e.g., NCEP-CFSv2 or NASA-GMAO) contain information salient to precipitation and temperature forecasting besides their mean, and ML models that leverage the full ensemble generally outperform methods that rely on the ensemble mean alone (Section 8a(i)).
- Finally, we emphasize that the final validation of our approach was conducted on data from 2011 to 2020 that was not used during any of the training, model development, parameter tuning, or model selection steps. We only conducted our final assessment of the predictive skill for 2011 to 2020 after we had completed all other aspects of this manuscript. Because of

this, our final empirical results accurately reflect the anticipated performance of our methods on new data.

This paper is organized as follows: Section 2 discusses related work, Section 3 introduces data used in the experiments, Section 4 describes forecasting problems, while baselines and learning-based methods are described in Section 5, experimental setup and evaluation metrics are given in Section 6. Finally, we present our results in Section 7 and discuss them in Section 8. Conclusions and directions for future work are given in Section 9.

## 2. Related work

While statistical models were common for weather prediction in the early days of weather forecasting (Nebeker 1995), forecasts using physics-based dynamic system models have been carried out since the 1980s and have been the dominant forecasting method in climate prediction centers since the 1990s (Barnston et al. 2012). Many physics-based forecast models are used both in academic research and operationally. Such systems often produce ensembles of forecasts – e.g., the result of running a physics-based simulation multiple times with different initial conditions or parameters, and are a mainstay of operational forecast centers around the globe.

Recently, skillful ML approaches have been developed for short-range weather prediction (Chen et al. 2023; Nagaraj and Kumar 2023; Frndá et al. 2022; Herman and Schumacher 2018; Ghaderi et al. 2017; Grover et al. 2015; Radhika and Shashi 2009; Cofino et al. 2002) and longer-term weather forecasting (Lam et al. 2023; Yang et al. 2023; Chen et al. 2023; Hewage et al. 2021; Cohen et al. 2019; Totz et al. 2017; Iglesias et al. 2015; Badr et al. 2014). However, forecasting on the subseasonal timescale, with 2-8 week outlooks, has been considered a far more difficult task than seasonal forecasting due to its complex dependence on both local weather and global climate variables (Vitart et al. 2012; Min et al. 2020). Seasonal prediction also benefits from targeting a much larger averaging period.

Some ML algorithms for subseasonal forecasting use purely observational data (i.e., not using any physics-based ensemble forecasts). He et al. (2020) focuses on analyzing different ML methods, including Gradient Boosting trees and Deep Learning (DL) for SSF. They propose a careful construction of feature representations of observational data and show that ML methods are able to outperform a climatology baseline, i.e., predictions corresponding to the 30-year mean at a

given location and time. This conclusion is based on comparing the relative  $R^2$  scores for the ML approaches and climatology. Srinivasan et al. (2021) proposes a Bayesian regression model that exploits spatial smoothness in the data.

Other works use the ensemble mean as a feature in their ML models. For example, in the subseasonal forecasting Rodeo (Hwang et al. 2018), a prediction challenge for temperature and precipitation at weeks 3-4 and 5-6 in the western U.S. sponsored by NOAA and the U.S. Bureau of Reclamation, simple yet thoughtful statistical models consistently outperform NOAA's dynamical systems forecasts. In particular, the winning approach uses a stacked model from two nonlinear regression models, a selection of climate variables such as temperature, precipitation, sea surface temperature, sea ice concentration, and a collection of physics-based forecast models including the ensemble mean from various modeling centers in the North American Multi-Model Ensemble (NMME). From the local linear regression with multitask feature selection model analysis, the ensemble mean is the first- or second-most important feature for forecasting, especially for precipitation. He et al. (2021) perform a comparison of modern ML models that use data from the Subseasonal Experiment (SubX) project for SSF in the western contiguous United States. The experiments show that incorporating the ensemble mean as an input feature to ML models leads to a significant improvement in forecasting performance, but that work does not explore the potential value of individual ensemble members aside from the ensemble mean. Grönquist et al. (2020) note that physics-based ensembles are computationally demanding to produce and propose an ML method that can input a subset of ensemble forecasts and generate an estimate of the full ensemble; they observe that the output ensemble estimate has more prediction skill than the original ensemble. Loken et al. (2022) analyze the forecast skill of random forests leveraging the ensemble members for next-day severe weather prediction compared to only using the ensemble mean. However, their results only cover forecasts with a lead time of up to 48 hours, so it is unclear if their methods would have succeeded in the tougher subseasonal forecasting setting.

This paper complements the prior work above by developing powerful learning-based approaches that incorporate both physics-based forecast models and observational data to improve SSF over CONUS.

### 3. Data

Table 1 describes variables used in the experiments. Climatological means of precipitation and temperature are calculated using 1971-2000 NOAA data (NOAA 2022). There are many ensembles of physics-based predictions produced by forecasting systems. NMME is a collection of physics-based forecast models from various modeling centers in North America, including NOAA/NCEP, NOAA/Geophysical Fluid Dynamics Laboratory (GFDL), International Research Institute for Climate and Society (IRI), National Center for Atmospheric Research (NCAR), NASA, and Canadian Meteorological Centre (Kirtman et al. 2014). NMME provides forecasts from multiple global forecast models from North American modeling centers (Kirtman et al. 2014). The NMME project has two predictive periods: hindcast and forecast. A hindcast period refers to when a dynamic model re-forecasts historical events, which can help climate scientists develop and test new models to improve forecasting and to evaluate model biases. In contrast, a forecast period has real-time predictions generated from dynamic models.

TABLE 1. **Description of climate variables and their data sources.** Our target climate variables for sub-seasonal forecasting are precipitation and 2-meter temperature. We use NOAA data to calculate the climatology from 1971 to 2000. We also perform linear spatial interpolation on the historical values to get values with the same resolution and support as target climate variables.

Type	Variable	Description	Unit	Spatial Coverage	Time Range	Data Source
Feature variable	tmp2m	Daily average temperature at 2 meters	°C	US mainland 0.5° × 0.5° grid	1985 to 2020	CPC Global Daily Temperature (Fan and Van den Dool 2008)
	precip	Daily average precipitation	mm	US mainland 0.5° × 0.5° grid	1985 to 2020	CPC Global Daily Precipitation (Xie et al. 2010)
	SSTs	Daily sea surface temperature	°C	Ocean only 0.25° × 0.25° grid	1985 to 2020	Optimum Interpolation SSTs High Resolution (OISST) (Reynolds et al. 2007)
	rhum	Daily relative humidity near the surface	Pa			
Climatology	slp	Daily pressure at sea level	%	US mainland and North Pacific & Atlantic Ocean 0.5° × 0.5° grid	1985 to 2020	Atmospheric Research Reanalysis Dataset (Kalnay et al. 1996)
	hgt500	Daily geopotential height at 500mb	m			
	tmp2m	Daily average temperature at 2 meters	K	Globally 1° × 1° grid	1971 to 2000	NOAA (NOAA 2022)
	precip	Daily average precipitation	mm	Globally 1° × 1° grid	1971 to 2000	NOAA (NOAA 2022)

In this manuscript, we use ensemble forecasts from the NMME’s NCEP-Climate Forecast System version 2 (CFSv2, Kirtman et al. (2014); Saha et al. (2014)), which has  $K = 24$  ensemble members at a  $1^\circ \times 1^\circ$  resolution over a 2-week lead time. NCEP-CFSv2 is the operational prediction model currently used by the U.S. Climate Prediction Center. The NCEP-CFSv2 model has two different products available in the NMME archive: we use its hindcasts from 1982 to 2010 for training and validation of our models, and we use its forecasts from April 2011 to December 2020 for the final evaluation of our models.

In order to ensure our results are not unique to a single forecasting model, we also analyze output from the NASA-Global Modeling and Assimilation (GMAO) from the Goddard Earth Observing System model version 5 (GEOS, Nakada et al. (2018)), which has  $K = 11$  ensemble members at a  $1^\circ \times 1^\circ$  resolution over a 2-week lead time. Similarly, we use its hindcasts from 1981 to 2010 for training and validation of our models, and we use its forecasts from January 2011 to January 2018 for final evaluation. The test periods of NCEP-CFSv2 and NASA-GMAO data differ due to data availability. Note that the identical version of each model is used to generate the test, train, and validation data.

Different ensemble members correspond to different initial conditions of the underlying physical model. The NCEP-CFSv2 forecasts are initialized in the following way: four initializations at times 0000, 0600, 1200, and 1800 UTC every fifth day, starting one month prior to the lead time of two weeks (Table B1 in Saha et al. (2014)). NASA-GMAO is a fully coupled atmosphere–ocean–land–sea ice model, with five forecasts initialized every five days. While additional members are generated through perturbation methods closest to the beginning of each month<sup>2</sup>, not all members are initialized on different dates, meaning that the ensemble is not strictly lagged. However, NASA-GMAO members are not interchangeable, as each is created using a distinct method.

All data are interpolated to lie on the same  $1^\circ \times 1^\circ$  grid, resulting in  $L = 3,274$  U.S. locations. Climate variables available daily (such as pressure at sea level or precipitation) are converted to monthly average values. When data are available as monthly averages only, we ensure that our forecast for time  $t + \delta_t$  does not use any information from the interval  $(t, t + \delta_t)$ .

---

<sup>2</sup><https://gmao.gsfc.nasa.gov/products/climateforecasts/GEOS5/DESC/init.php>

## 4. Forecast tasks

The learning task can be formulated as learning a model  $f_\theta : \mathbf{X} \rightarrow \mathbf{y}$  with parameters  $\theta$ . This model  $f_\theta$  can be a linear regression (where  $\theta$  is a set of regression weights), the mean of ensemble members (no  $\theta$  needs to be learned), a random forest (where  $\theta$  parameterizes the set of trees in the forest), a convolutional neural network (where  $\theta$  is the collection of neural network weights), or other learned models. We consider three forecasting tasks: regression, tercile classification, and quantile regression.

(i) *Regression* The goal of regression is to predict monthly average values of precipitation and 2-meter temperature two weeks in the future. These models are generally trained using the squared error loss function:

$$\ell_{\text{sq-err}}(\theta) = \mathbb{E}[(y - f_\theta(\mathbf{x}))^2]. \quad (1)$$

(ii) *Tercile classification* The goal of tercile classification is to predict whether the precipitation or 2-meter temperature will be “high” (above the 66th percentile, denoted  $q = 1$ ), “medium” (between the 33rd and the 66th percentiles, denoted  $q = 0$ ), or “low” (below the 33rd percentile, denoted  $q = -1$ ). We compute these percentile values using the 1971-2000 climatology (see Section 3 for details), and these percentiles are computed for each calendar month  $m$  and location  $l$  pair. These models are generally trained using the cross-entropy loss function:

$$\ell_{\text{CE}}(\theta) = \mathbb{E} \left[ \sum_{q=-1}^1 -\mathbb{I}_{\{y=q\}} \log(f_\theta(\mathbf{x}))_q \right], \quad (2)$$

where  $\mathbb{I}_{\{A\}} := \begin{cases} 1, & \text{if } A \text{ true} \\ 0, & \text{if } A \text{ false} \end{cases}$  is the indicator function and  $(f_\theta(x))_q$  is the predicted probability that the target  $y$  corresponding to feature vector  $x$  will be in tercile  $q$ .

(iii) *Quantile regression* For a given percentile  $\alpha$ , the goal of quantile regression is to predict the value  $z$  so that, conditioned on features  $x$ , the target  $y$  satisfies  $y \leq z$  with probability  $\alpha$ . When we set  $\alpha$  to a value close to one, such as  $\alpha = 0.9$ , this value  $z$  indicates what we can expect in “extreme outcomes”, not just on average. These models are generally trained using the pinball loss function:

$$\ell_{\text{quantile}}(\theta) = \mathbb{E}[\rho_\alpha(y - f_\theta(\mathbf{x}))] \quad (3)$$

where

$$\rho_\alpha(z) := z(\alpha - \mathbb{I}_{\{(z < 0)\}}) = \begin{cases} \alpha|z| & \text{if } z \geq 0 \\ (1-\alpha)|z| & \text{if } z < 0 \end{cases}. \quad (4)$$

## 5. Prediction methods

Our goal is to predict either the monthly average precipitation or the monthly average 2-meter temperature two weeks in advance (for example, we predict the average monthly precipitation for February on January 15). This section describes the notation used for features and targets, baselines and learning methods, and how spatial features are accounted for.

### a. Notation

We let  $T$  denote the number of time steps used in our analysis, and  $L$  denote the number of spatial locations. We define the following variables:

- $u_{t,l}^{(k)}$  is the  $k$ -th ensemble member at time  $t$  and location  $l$ , where  $k = 1, \dots, K$ ,  $t = 1, \dots, T$ ,  $l = 1, \dots, L$ . Every ensemble member represents the output of a given physics-based model forecast from different initial states.
- $v_{t,l}^{(p)}$  is the  $p$ -th observational variable, such as precipitation or temperature, geopotential height at 500mb, relative humidity near the surface, pressure at sea level and sea surface temperature, at time  $t$  and location  $l$ , with  $p = 1, \dots, P$ .
- $\mathbf{z}_l^{(1)}, \mathbf{z}_l^{(2)}$  represent information about longitude and latitude of location  $l$ , respectively; each is a vector of length  $d$ . More details about this representation, called positional encoding (PE), can be found in Section 6a.
- $\mathbf{x}_{t,l} := [u_{t,l}^{(1)}, \dots, u_{t,l}^{(K)}, v_{t,l}^{(1)}, \dots, v_{t,l}^{(P)}, \mathbf{z}_l^{(1)}, \mathbf{z}_l^{(2)}]$  is a set of features at time  $t$  and location  $l$ .
- $y_{t,l}$  is the target – the ground truth monthly average precipitation or 2-meter temperature at the target forecast time  $t + \delta_t$  at location  $l$ , where  $\delta_t = 14$  days is our forecast horizon. For simplicity, we use a subscript  $t$  for  $y_{t,l}$  instead of  $t + \delta_t$  to match with the input features notation. The same holds for our ensuing definitions.

- $\hat{y}_{t,l}$  is the output of a forecast model for a given task at target forecast time  $t + \delta_t$  and location  $l$ .
- $s_{m,l}$  – a 30-year mean (climatology) of an observed climate variable, such as precipitation or temperature, at a month  $m = 1, \dots, 12$  and location  $l$ .
- $\hat{s}_{m,l}$  – a 30-year climatology of a *predicted* climate variable, such as precipitation or temperature, at a month  $m = 1, \dots, 12$  and location  $l$ . For each location  $l$  and each month  $m$ , it is calculated as a mean of ensemble member predictions over the training period, as defined formally in Eq. (7).
- $y_{t,l}^{\text{anomaly}}$  and  $\hat{y}_{t,l}^{\text{anomaly}}$  are anomaly predictions and a true anomaly, at a month  $m = 1, \dots, 12$  and location  $l$ . They are used during evaluation. We define anomalies as

$$y_{t,l}^{\text{anomaly}} = y_{t,l} - s_{m(t),l}, \quad (5)$$

$$\hat{y}_{t,l}^{\text{anomaly}} = \hat{y}_{t,l} - \hat{s}_{m(t),l}. \quad (6)$$

For the special case of  $\hat{y}$  corresponding to the ensemble mean, the ensemble members may exhibit bias, in which case we also consider  $\hat{y}_{t,l}^{\text{anomaly}} = \hat{y}_{t,l} - \hat{s}_{m(t),l}$ , where  $\hat{s}_{m(t),l}$  is evaluated on the model’s (ensemble mean) predictions:

$$\hat{s}_{m,l} := \frac{1}{T} \sum_{t=1}^T \hat{y}_{t,l} \mathbb{I}_{\{m=m(t)\}}, \quad l = 1, \dots, L. \quad (7)$$

The model climatology  $\hat{s}_{m,l}$  is computed using the training data. Note that we do not subtract the climatology from the input features and target variables, i.e., precipitation and temperature, when training our ML models. We subtract climatology from the model outputs only when evaluating their performance, as including climatology in the inputs to our ML models during training improves performance. Section 6c provides more details on the model evaluation.

In our analyses, the number of locations is  $L = 3274$ , there are  $K = 24$  NCEP-CFSv2 ensemble members or  $K = 11$  NASA-GMAO ensemble members. The ensemble members are used as input features to the learning-based methods as they are, we do not perform any feature extraction from

them. The number of observational variables is usually  $P = 17$ . The details on these variables can be found in Section 3 and Section 6b.

The target variable  $y$  is observed from 1985 to 2020. Data from January 1985 to September 2005 are used for training (249 time steps), and data from October 2005 to December 2010 are used for validation and model selection (63 time steps). Data from 2011 to 2020 (or from 2011 to 2018 in the case of NASA-GMAO data) are used to test our methods after all model development, selection, and parameter tuning are completed.

### b. Baselines

(i) *Climatology* It is the fundamental benchmark for weather and climate predictability. In particular, for a given time  $t$ , let  $m(t) := (t \bmod 12)$  correspond to the calendar month corresponding to  $t$ ; then we compute the 30-year climatology of the target variable for a given location and time via

$$\hat{y}_{t,l}^{\text{hist}} = s_{m(t),l}, \quad t = 1, \dots, T, \quad l = 1, \dots, L. \quad (8)$$

(ii) *Ensemble mean* This is the mean of all ensemble members for each location  $l$  at each time step  $t$ :

$$\hat{y}_{t,l}^{\text{ens mean}} := \frac{1}{K} \sum_{k=1}^K u_{t,l}^{(k)}, \quad t = 1, \dots, T, \quad l = 1, \dots, L. \quad (9)$$

(iii) *Linear regression* Finally, we consider, as a baseline, a *linear regression* model applied to input features corresponding to ensemble member predictions:  $\mathbf{x}_{t,l} = [u_{t,l}^{(1)}, \dots, u_{t,l}^{(K)}]$ . Then, the model's output

$$\hat{y}_{t,l}^{\text{LR}} := \langle \boldsymbol{\theta}_l, \mathbf{x}_{t,l} \rangle + \theta_l^0, \quad (10)$$

where  $\boldsymbol{\theta}_l$  are the trained coefficients for input features for each location  $l$ , and  $\theta_l^0$  are the learned intercepts for each location  $l$ . Note that we train a different model for each spatial location, and the illustration for this model and its input's format is given in Figure 1(a).

### c. Learning-based methods

(i) *Linear regression (LR)* In contrast to the linear regression baseline, here other climate variables are added to the input features:  $\mathbf{x}_{t,l} = [u_{t,l}^{(1)}, \dots, u_{t,l}^{(K)}, v_{t,l}^{(1)}, \dots, v_{t,l}^{(P)}]$ . Then the model's output is defined with Eq. (10). Because the feature vector is higher dimensional here than for the baseline, the learned  $\theta_l$  is also higher dimensional. We train a different model for each spatial location. In our experiments with linear models, we do not include positional encoding  $(\mathbf{z}_l^{(1)}, \mathbf{z}_l^{(2)})$  as input features, since they would be constants for each location's linear model.

In the context of regression, we minimize the squared error loss. The linear quantile regressor (Linear QR) is a linear model trained to minimize the quantile loss

$$\ell^{\text{QR}} = \frac{1}{L} \sum_{l=1}^L \left[ \frac{1}{T} \sum_{t=1}^T \rho_\tau (y_{t,l} - \hat{y}_{t,l}) \right], \quad (11)$$

where  $\rho_\tau$  is defined in Eq. (4).

(ii) *Random forest* In the context of regression and tercile classification, we train a random forests that use ensemble predictions, the spatial location, and additional climate features to form the feature vector  $\mathbf{x}_{t,l} = [u_{t,l}^{(1)}, \dots, u_{t,l}^{(K)}, v_{t,l}^{(1)}, \dots, v_{t,l}^{(P)}, \mathbf{z}_l^{(1)}, \mathbf{z}_l^{(2)}]$  for all location  $l$  and time  $t$  pairs. One random forest is trained to make predictions for any spatial location. The illustration for RF and its input's format is given in Figure 1(b): we train one RF model for all locations, and the spatial information is encoded as input features via PE vectors  $\mathbf{z}_l^{(1)}, \mathbf{z}_l^{(2)}$ .

In the context of quantile regression, we train a random forest quantile regressor (RFQR, Meinshausen (2006)), which grows trees the same way as the original random forest while storing all training samples. To make a prediction for a test point, the RFQR computes a weight for each training sample that corresponds to the number of leaves (across all trees in the forest) that contain the test sample and the training sample. The RFQR prediction is then a quantile of the weighted training samples across all leaves that contain the test sample. We show a figure representation of the RFQR in Section D1b. With this formulation, training a single RFQR for all locations is computationally demanding, so we train individual RFQRs for every location.

Random forests are often referred as the best off-the-shelf classifiers (Hastie et al. 2009) even using the default hyperparameters (Biau and Scornet 2016). Our cross-validation (CV) and grid

search experiments show that the RFs hyperparameters have little impact on the accuracy, and thus, we use the default parameters for RFs from the Scikit-learn library (Pedregosa et al. 2011).

*(iii) Convolutional neural network* To produce a forecast map for the U.S., we adapted a U-Net architecture (Ronneberger et al. 2015), which has an encoder-decoder structure with convolutional layer blocks. The U-Net maps a stack of images to an output image; in our context, we treat each spatial map of a climate variable or forecast as an image. Thus, the input to our U-Net is can be represented as a tensor composed of matrices:  $\mathbf{X}_t = [\mathbf{U}_t^{(1)}, \dots, \mathbf{U}_t^{(K)}, \mathbf{V}_t^{(1)}, \dots, \mathbf{V}_t^{(P)}, \mathbf{Z}^{(1)}, \mathbf{Z}^{(2)}]$ .

Note that here, we use capital letters because the input to our U-Net consists of 2-D spatial maps, which are represented as matrices instead of vectors. The model output is a spatial map of the predicted target. This process is illustrated in Figure 1(c).

For the U-Net, we modify an available PyTorch implementation (Yakubovskiy 2020). The training set consists of 249 samples (images), which may be considered relatively limited for CNN training. To address this concern, we conduct bootstrapping experiments for the U-Net architecture, offering detailed insights into the impact of sample size on model performance. Further details are presented in Appendix C2. We use a 10-fold CV over our training data and grid search to select parameters such as learning rate, weight decay, batch size, and number of epochs. The Adam optimizer (Kingma and Ba 2014) is used in all experiments. After selecting hyperparameters, we train the U-Net model with those parameters on the full training dataset. The validation set is used to perform feature importance analysis. For regression, we train using squared error loss. In the context of quantile regression, we initialize the weights with those learned on squared error loss and then train on the quantile loss Eq. (11).

*(iv) Nonlinear model stacking* Model stacking can improve model performance by combining the outputs of several models (usually called base models) (Pavlyshenko 2018). In our case, linear regression, random forests, and the U-Net are substantially different in architecture and computation, and we observe that they produce qualitatively different forecasts. We stack the linear model, random forest, and U-Net forecasts using a nonparametric approach:

$$\hat{y}_{t,l} = h(\hat{y}_{t,l}^{\text{LR}}, \hat{y}_{t,l}^{\text{RF}}, \hat{y}_{t,l}^{\text{UNET}}), \quad (12)$$

where  $h$  is a simple feed-forward neural network with a non-linear activation and  $\hat{y}_{t,l}^{\text{LR}}, \hat{y}_{t,l}^{\text{RF}}, \hat{y}_{t,l}^{\text{UNET}}$  are the predictions of a linear model, random forest and the U-Net correspondingly and referred to as “base models”. One stacking model is trained to make predictions for any spatial location. Figure 1(b) with input features that are predictions from other ML models and no PE vectors demonstrates the stacking model’s framework. Model stacking can improve the forecast quality by combining predictions from three forecasting paradigms – spatial independence, conditional spatial independence, and spatial dependence (Section 5d), and is analogous to the multi-model ensemble approach commonly used in weather and climate forecasting. The architecture details can be found in Appendix D.

We apply the following procedure for model stacking: the base models are first trained on half of the training data, and predicted values on the second half are used to train the stacking model  $h$ . Then, we retrain the base models on all the training data and apply the trained stacked model to the outputs of the base models. The proposed procedure helps to avoid overfitting.

#### *d. Models of spatial variation*

We consider three different forecasting paradigms. In the first, which we call the **spatial independence** model, we ignore all spatial information and train a separate model for each spatial location. In the second, which we call the **conditional spatial independence** model, we consider samples corresponding to different locations  $l$  as independent conditioned on the spatial location as represented by features  $(\mathbf{z}_l^{(1)}, \mathbf{z}_l^{(2)})$ . In this setting, a training sample corresponds to  $(\mathbf{x}_i, y_i) = (\mathbf{x}_{t,l}, y_{t,l})$ , where, with a small abuse of notation, we let  $i$  index a  $t, l$  pair. In this case, the number of training samples is  $n = TL$ . In the third paradigm, which we call the **spatial dependence** model, we consider a single training sample as corresponding to full spatial information (across all  $l$ ) for a single  $t$ ; that is  $(\mathbf{X}_i, \mathbf{Y}_i) = ([x_{t,l}]_{l=1, \dots, L}, [y_{t,l}]_{l=1, \dots, L})$ , where now  $i$  indexes  $t$  alone. Models developed under the spatial dependence model account for the spatial variations in the features and targets. For instance, a convolutional neural network might input “heatmaps” representing the collection of physics-based model forecasts across the continental U.S. and output a forecast heatmap predicting spatial variations in temperature or precipitation instead of treating each spatial location as an independent sample.

Figure 1 shows general frameworks of these paradigms. All models combine information from all the different ensemble forecasts, and so in a broad sense, we can think of each prediction at a given time and location as a weighted sum of the ensemble forecasts across space, time, and ensemble members, where the weights are learned during the model training and may be data-dependent (i.e., nonlinear). From this perspective, we may think of different modeling paradigms as essentially placing different constraints on those weights:

- under spatial independence models, the weights may vary spatially but do not account for spatial correlations in the data;
- under conditional spatial independence models, the interpretation depends on the model being trained – linear models have the same weights on ensemble predictions regardless of spatial location, while nonlinear models (e.g., random forests) have weights that may depend on the spatial location;
- under spatial dependence models, the weights vary spatially, depend on the spatial location, and account for spatial correlations among the ensemble forecasts and other climate variables.

## 6. Experimental setup

This section provides details on the experimental setup, including positional encoding, removing climatology, and evaluation metrics. Data preprocessing details are presented in Appendix E.

### *a. Positional encoding and inputs*

Positional encoding (Vaswani et al. 2017) is a technique used in natural language processing (NLP) to inject positional information into data. In sequence-based tasks, such as language translation or text generation, the order of elements in the input sequence is important, but neural networks do not naturally capture this information. PE assigns unique encodings to each position in the sequence, which are then added to the original input before being processed by the model. This enables the model to consider the order and relative positions of elements, improving its ability to capture local and global context within a sequence and make accurate predictions (Devlin et al. 2018; Petroni et al. 2019; Narayanan et al. 2016). This technique is helpful to represent the positional information outside the original NLP tasks (Gamboa 2017; Gehring et al. 2017; Khan

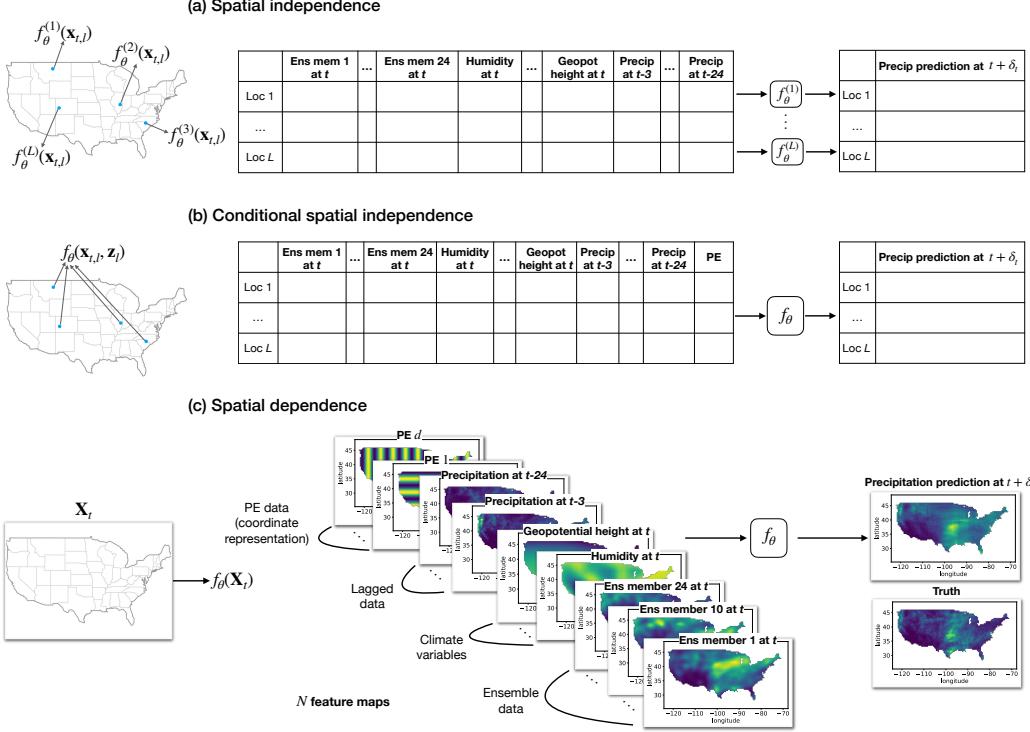


FIG. 1. **An illustration of different forecasting paradigms:** (a) spatial independence models with a model for each spatial location, no accounting for spatial information; (b) conditional spatial independence models with one model for all locations, might consider the spatial information; (c) spatial dependence models that account for the spatial information by design. We replace “precipitation” in the illustration with “temperature” for temperature prediction, but the overall structure remains the same.

et al. 2022). Several of our models use the spatial location as an input feature. Rather than directly using latitudes and longitudes, we use PE (Vaswani et al. 2017):

$$z_l^{(1)}(i) = \text{PE}(l, 2i) = \sin(l/10000^{2i/d}), \quad (13)$$

$$z_l^{(2)}(i) = \text{PE}(l, 2i+1) = \cos(l/10000^{2i/d}), \quad (14)$$

where  $l$  is a longitude or latitude value,  $d = 12$  is the dimensionality of the positional encoding, and  $i \in \{1, \dots, d\}$  is the index of the positional encoding vector. For the U-Net model, PE vectors are transformed into images in the following way: we take every value in the vector and fill the image

of the desired size with this value. So, there are  $d$  images with the corresponding PE values. For the RF models, PE vectors can be used as they are.

### *b. Models' inputs details*

Based on the available data, we use the following input features for our ML models:

- $K$  ensemble forecasts for the target month,
- four climate variables: relative humidity, pressure, geopotential height, and temperature (if the target is precipitation) or precipitation (if the target is temperature) two months before the target month,
- the lagged target variable (the target variable two, three, four, twelve and 24 months before the target date – five additional features),
- SSTs that are represented via principal components (PCs),
- and, finally, the positional embeddings.

SSTs are usually represented as eight PCs, and the embedding vector size is usually  $d = 12$  as we describe Section 6a. For example, using the NCEP-CFSv2 members, there are  $\underbrace{24}_{K} + \underbrace{4+5+8}_{P} + \underbrace{12}_{d} \times 2 = 65$  input features for every time step and location. Figure 1 provides an illustration of these input features.

### *c. Evaluation metrics*

(i) *Regression metrics* The forecast skill of our regression models is measured using the  $R^2$  value. For each location  $l$  and ground-truth values  $y_{t,l}$  and predictions  $\hat{y}_{t,l}$  at this location, we compute

$$R_l^2 = 1 - \frac{\sum_{t=1}^T (y_{t,l}^{\text{anomaly}} - \hat{y}_{t,l}^{\text{anomaly}})^2}{\sum_{t=1}^T (y_{t,l}^{\text{anomaly}} - \bar{y}_l^{\text{anomaly}})^2}, \quad (15)$$

where

$$\bar{y}_l^{\text{anomaly}} = \frac{1}{T} \sum_{t=1}^T \hat{y}_{t,l}^{\text{anomaly}}.$$

Then, the average  $R^2$  for all locations is calculated as

$$R^2 = \frac{1}{L} \sum_{l=1}^L R_l^2. \quad (16)$$

In addition to the average  $R^2$  on the test data, we also estimate the median  $R^2$  score across all U.S. locations.

We further report the *mean squared error* (MSE) of our predictions across all locations:

$$\text{MSE}_l := \frac{1}{T} \sum_{t=1}^T (y_{t,l} - \hat{y}_{t,l})^2,$$

for  $l = 1, \dots, L$ , and

$$\text{MSE} = \frac{1}{L} \sum_{l=1}^L \text{MSE}_l. \quad (17)$$

We also report the standard error (SE), median, and 90th percentile of  $\{\text{MSE}_l\}_l$ . We say the difference between the two models is significant if their  $\text{MSE} \pm \text{SE}$  intervals do not overlap. Note that the standard errors provided here should be used with caution since there are significant spatial correlations in the MSE values across locations, so we do not truly have  $L$  independent samples from an asymptotically normal distribution.

(ii) *Tercile classification metrics* We estimate the accuracy of our tercile classification predictions as the proportion of correctly classified samples out of all observations.

(iii) *Quantile regression metrics* For the quantile regression task, we report *mean quantile loss* from Eq. (11) across all locations.

## 7. Experimental results

In this section, we report the predictive skill of different models applied to SSF over the continental U.S. using NCEP-CFSv2 ensemble members for regression and quantile regression. Precipitation forecasting is known to be more challenging compared to temperature forecasting (Knapp et al. 2011). The results for the NASA-GMAO dataset are presented in Appendix A. The skill of different models on the tercile classification task is presented in Appendix B for both datasets. Recall that all methods are trained on data spanning January 1985–September 2005, with data spanning October

2005 - December 2010 used for validation (i.e., model selection and hyperparameter tuning). Test data spanning 2011 to 2020 was **not** viewed at any point of the model development and training process and only used to evaluate the predictive skill of our trained models on previously unseen data; we refer to this period as the “test period”. As a navigation tool for the reader, Table 2 gives references to the presented results for different tasks.

TABLE 2. A table with references to the main results.

Task	Data	Reference
Regression	precip NCEP-CFSv2	Table 3; Figure 2
	tmp NCEP-CFSv2	Table 4; Figure 3
	precip NASA-GMAO	Table A1; Figure A1
	tmp NASA-GMAO	Table A2; Figure A2
Quantile regression	precip NCEP-CFSv2	Table 5; Figure A3
	tmp NCEP-CFSv2	Table 6; Figure 4
	precip NASA-GMAO	Table A3; Figure A3
	tmp NASA-GMAO	Table A4; Figure A4
Feature importance	precip NCEP-CFSv2	Table 10
	tmp NCEP-CFSv2	Table 11
Tercile classification	precip NCEP-CFSv2	Table B1; Figure B1
	tmp NCEP-CFSv2	Table B2; Figure B3
	precip NASA-GMAO	Table B1; Figure B2
	tmp NASA-GMAO	Table B2; Figure B4

### a. Regression

(i) *Precipitation regression using NCEP-CFSv2* Precipitation regression results are presented in Table 3. While the individual ML approaches produce results generally similar to those of the baselines, the stacked ML model, in particular, outperforms the baseline models in almost all metrics. Note that the best  $R^2$  value, associated with the stacked model, is still near zero; while this is a significant improvement over, for example, the ensemble mean, which has an  $R^2$  value of -0.08, the low values for all methods indicate the difficulty of the forecasting problem. It is important to note that  $R^2$  measures the accuracy of a model relative to a baseline corresponding to the mean of the target *over the test period* – that is, relative to a model that could never be used in practice as a forecaster because it uses future observations. The best *practical* analog to

this would be the mean of the target *over the train period* – what we call the “historical mean” or climatology model. These two models are not the same, possibly because of the nonstationarity of the climate (Min et al. 2020). Thus, even when our  $R^2$  values are negative (i.e., we perform worse than the impractical mean of the target over the test period), we still perform much better than the practical climatology predictor. The model stacking approach is applied to the models trained on all available features (i.e., ensemble members, PE, climate variables; linear regression is trained on all features except PE). We decide what models to include in the stacking approach based on their performance on validation data. The low 90th percentile error implies that our methods not only have high skill on average but also that there are relatively few locations with large errors. While acknowledging the overall performance may not be exceptional, it is important to recognize the potential of machine learning methods in improving the quality of estimates relative to the standard baselines. To further evaluate the capabilities of the stacking approach, we also apply the approach to the baseline predictions, which include historical and ensemble means, as well as linear regression. The performance of the stacked baseline model exceeds that of any of the individual baseline models and is similar to the performance of the stacked ML approach in terms of the  $R^2$  metric. However, the stacked ML approach outperforms it in all MSE-based metrics, indicating that the ML techniques can still provide additional skill even for as notoriously challenging a quantity as precipitation.

TABLE 3. **Results for precipitation regression the using NCEP-CFSv2 ensemble, with errors reported over the test period.** LR refers to linear regression on all features, including ensemble members, lagged data, climate variables, and SSTs. ML model stacking is performed on models that are trained on all features. The **best** results are in bold. MSE is reported in squared mm.

Model	Features	Mean $R^2$ ( $\uparrow$ )	Median $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )	Median MSE ( $\downarrow$ )	90th prctl MSE ( $\downarrow$ )
Baseline	Climatology	-0.06	-0.01	$2.33 \pm 0.04$	1.59	4.96
	Ens mean	-0.08	0.01	$2.19 \pm 0.04$	1.55	4.57
	Linear Regr	-0.11	-0.07	$2.26 \pm 0.04$	1.54	4.72
	Baseline stacking	0.00	0.04	$2.15 \pm 0.04$	1.44	4.55
LR	All features	-0.33	-0.25	$2.71 \pm 0.05$	1.91	5.45
U-Net	All features	-0.10	-0.01	$2.18 \pm 0.03$	1.44	4.62
RF	All features	-0.11	-0.01	$2.17 \pm 0.05$	1.48	4.45
Stacked	LR, U-Net, RF outputs	<b>0.02</b>	<b>0.04</b>	<b><math>2.07 \pm 0.03</math></b>	<b>1.42</b>	<b>4.38</b>

Figure 2 illustrates performance of key methods with  $R^2$  heatmaps over the U.S. to highlight spatial variation in errors. The RF and U-Net  $R^2$  fields are qualitatively similar, but they are still quite different in certain states such as Georgia, North Carolina, Virginia, Utah, and Colorado. The LR map is noticeably poor across most of the regions. The stacked ML model’s heatmap reveals large regions where its predictive skill exceeds that of all other methods. Note that model stacking yields relatively accurate predictions even in regions where the three constituent models individually perform poorly (e.g., southwestern Arizona), highlighting the generalization abilities of our stacking approach. All methods tend to have higher accuracy on the Pacific Coast, in the Midwest, and in southern states such as Alabama and Missouri. The stacking model heatmaps both look similar. The stacking model applied to the baselines has better  $R^2$  scores in California compared to the stacked ML methods. However, the stacked ML model reveals larger positive  $R^2$  regions and fewer dark red spots, particularly evident in New Mexico, Minnesota, and Utah.

(ii) *Temperature regression* Table 4 shows results for 2-meter temperature regression. The learning-based models, especially the random forest and stacked model, significantly outperform the baseline models in terms of MSE and  $R^2$  score. The random forest also outperforms linear regression and the U-Net. Note that LR, U-Net, and RF are trained without using SST information since SST features yielded worse performance over the validation period. Figure 3 illustrates the performance of these methods with  $R^2$  heatmaps over the U.S. As expected, the model stacking approach shows the best results across spatial locations. We notice that there are still regions such as the West, some regions in Texas, Florida, and Georgia where all models tend to achieve a negative  $R^2$  score.

### *b. Quantile regression*

We explore the use of quantile regression to predict values  $z$  so that “there’s a 90% chance that the average temperature will be below  $z^\circ$  at your location next month” – or, equivalently, “there is a 10% chance that the average temperature will exceed  $z^\circ$  at your location next month.” In this sense, quantile regression focused on the 90-th percentile predicts temperature and precipitation extremes, a task highly relevant to many stakeholders. We train a linear regression model fitting the quantile loss (Linear QR), a random forest quantile regressor (RFQR), (Meinshausen 2006), a U-Net, and the stacked model. The Linear QR and the RFQR details are discussed in Section 5. The below

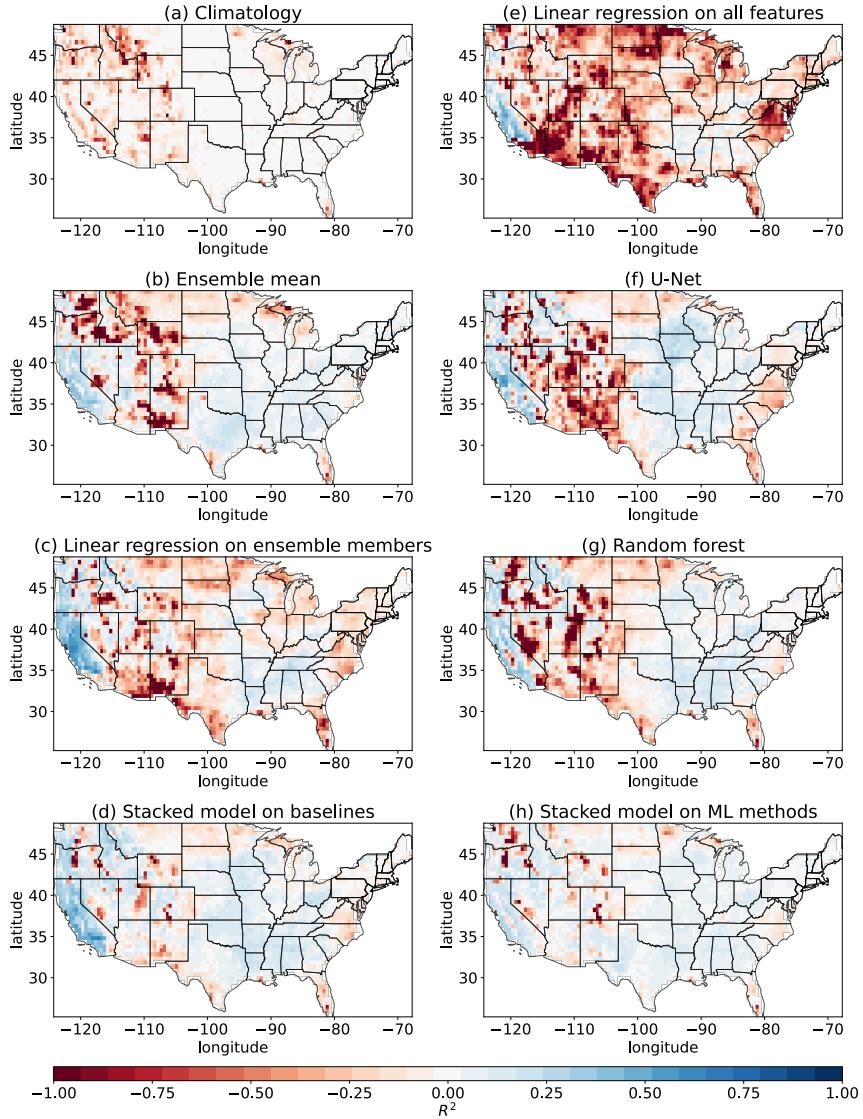


FIG. 2.  $R^2$  score heatmaps of baselines and learning-based methods for precipitation regression using NCEP-CFSv2 ensemble members; errors are computed over the test period. Positive values (blue) indicate better performance. See Section 7a(i) for details.

experimental results show that temperature extremes can be predicted with high accuracy by the learning-based models (particularly our stacked model), in stark contrast to historical quantiles or ensemble quantiles in the case of temperature quantile regression. The results for precipitation are less striking overall, though the learned models are significantly more predictive in some locations on this quantile regression task.

TABLE 4. **Results for temperature regression using the NCEP-CFSv2 ensemble, with errors reported over the test period.** LR refers to linear regression on all features including ensemble members, lagged data, land variables. Model stacking is performed on models that are trained on all features except SSTs. The **best** results are in bold. MSE is reported in squared  $^{\circ}\text{C}$ .

Model	Features	Mean $R^2$ ( $\uparrow$ )	Median $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )	Median MSE ( $\downarrow$ )	90th prctl MSE ( $\downarrow$ )
Baseline	Climatology	-0.66	-0.17	$6.57 \pm 0.11$	5.04	9.99
	Ens mean	-0.47	0.08	$5.51 \pm 0.10$	3.83	9.16
	Linear Regr	0.04	0.17	$3.60 \pm 0.03$	3.25	5.49
LR	All features w/o SSTs	0.05	0.16	$3.57 \pm 0.02$	3.33	5.41
U-Net	All features w/o SSTs	0.01	0.18	$3.65 \pm 0.02$	3.38	5.31
RF	All features w/o SSTs	0.16	0.25	$3.17 \pm 0.02$	2.99	4.63
Stacked	LR, U-Net, RF outputs	<b>0.18</b>	<b>0.27</b>	<b><math>3.11 \pm 0.02</math></b>	<b>2.93</b>	<b>4.56</b>

(i) *Quantile regression of precipitation* For each location, the 90th percentile value is calculated based on the historical data. For the ensemble 90th percentile, we simply take the 90th percentile of the  $K$  ensemble members. Table 5 summarizes results for precipitation quantile regression using the NCEP-CFSv2 ensemble. Our stacked model is able to significantly outperform all baselines. The performance illustration is given in Appendix A, Figure A3.

TABLE 5. **Test results for precipitation quantile regression using NCEP-CFSv2 dataset, with target quantile = 0.9.** Linear QR refers to a linear quantile regressor. RFQR corresponds to a Random Forest Quantile Regressor. Model stacking is performed on models that are trained on all features. The **best** results are in bold. Quantile loss is reported in mm.

Model	Features	Mean Qtr Loss ( $\downarrow$ )	Median Qtr Loss ( $\downarrow$ )	90th prctl Qtr Loss ( $\downarrow$ )
Baseline	Historical 90th percentile	$0.304 \pm 0.003$	0.278	0.504
	Ens 90th percentile	$0.311 \pm 0.003$	0.275	0.488
	Linear QR ens only	$0.310 \pm 0.003$	0.266	0.505
Linear QR	All features	$0.287 \pm 0.003$	0.248	0.463
U-Net	All features	$0.312 \pm 0.002$	0.281	0.504
RFQR	All features	<b><math>0.282 \pm 0.002</math></b>	0.257	<b>0.453</b>
Stacked	U-Net, RFQR, LQR outputs	<b><math>0.282 \pm 0.002</math></b>	<b>0.256</b>	0.457

(ii) *Quantile regression of temperature* Table 6 summarizes results for temperature quantile regression using the NCEP-CFSv2 ensemble. Note that we do not include SST features for

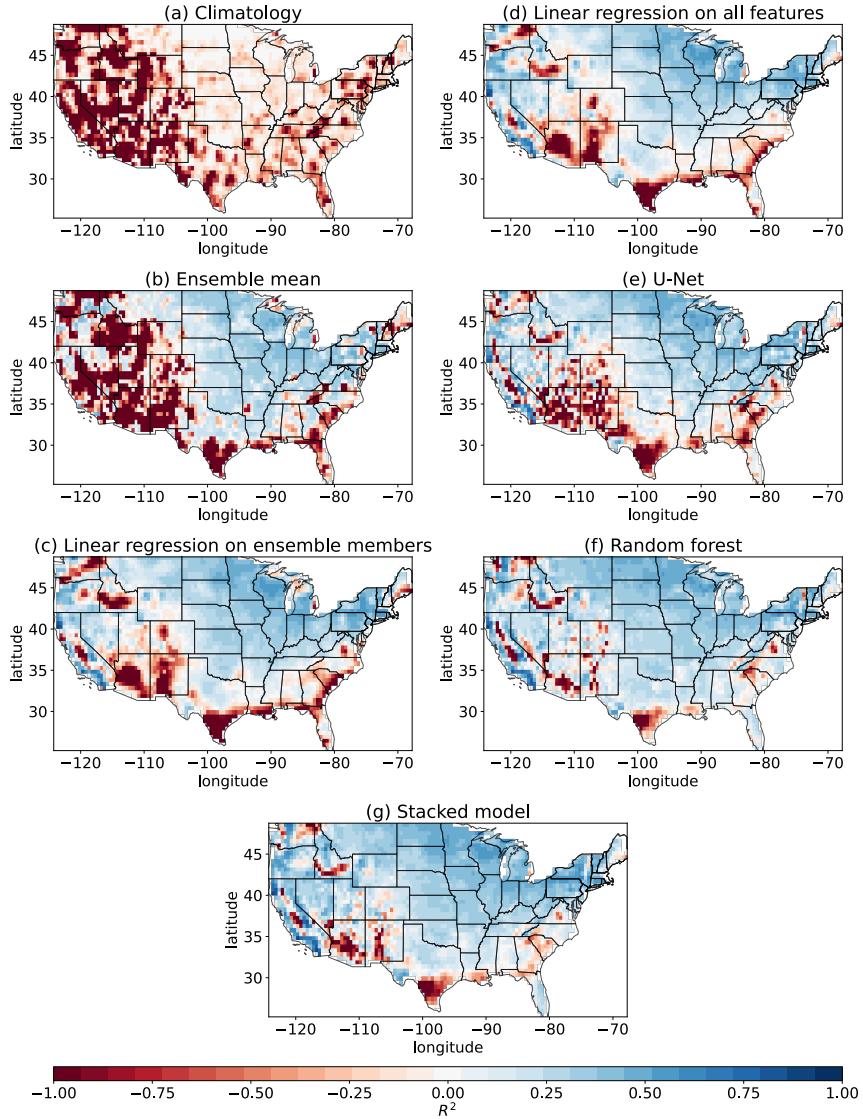


FIG. 3.  $R^2$  score heatmaps of baselines and learning-based methods for temperature regression using NCEP-CFSv2 ensemble members; errors are computed over the test period. Positive values (blue) indicate better performance. See Section 7a(ii) for details.

temperature quantile regression in our learned models. We observe that all of our learned models are able to significantly outperform all baselines. In Figure 4, we show the heatmaps of quantile loss of baselines and our learned models. We observe that the learned models produce predictions with varied quality, and the stacked model can pick up useful information from them. For example, in Arizona and Texas, the Linear QR, U-Net, and RFQR show some errors but in different locations, and the stacked model can exploit the advantages of each model.

TABLE 6. **Test results for temperature quantile regression using NCEP-CFSv2 dataset, with target quantile = 0.9.** Linear QR refers to a linear quantile regressor. RFQR corresponds to a Random Forest Quantile Regressor. Model stacking is performed on models that are trained on all features except for SSTs. Learned models can predict highly likely temperature ranges accurately, meaning there are fewer unpredicted temperature spikes. The **best** results are in bold.

Model	Features	Mean Qtr Loss (↓)	Median Qtr Loss (↓)	90th prctl Qtr Loss (↓)
Baseline	Historical 90th percentile	0.589 ± 0.008	0.435	0.980
	Ens 90th percentile	0.642 ± 0.009	0.468	1.196
	Linear QR ens only	0.336 ± 0.004	0.286	0.488
Linear QR	All features w/o SSTs	0.318 ± 0.002	0.301	0.407
U-Net	All features w/o SSTs	0.363 ± 0.003	0.329	0.488
RFQR	All features w/o SSTs	0.320 ± 0.002	0.307	0.384
Stacked	U-Net, RFQR, LQR outputs	<b>0.287 ± 0.001</b>	<b>0.285</b>	<b>0.344</b>

## 8. Discussion

### a. The efficacy of machine learning for SSF

Several hypotheses might explain why ML may be a promising approach for SSF, and we probe those hypotheses in this section.

(i) *Using full ensemble vs. ensemble mean* Past works use ensemble mean as an input feature to machine learning methods in addition to the climate variables (Hwang et al. 2018; He et al. 2021). Ensembles provide valuable information not only about expected climate behavior but also variance or uncertainty in multiple dimensions; methods that rely solely on ensemble mean lack information about this variance. Ensemble members may have systematic errors, either in the mean or the variability, arising from different initial conditions of the corresponding dynamic model that are not readily apparent to users. The more recently initialized an ensemble member is, the better it usually performs. While taking the average of these ensemble members may cancel out the deficiencies of each individual member, it is also possible that details of each member's systematic errors may be directly discovered and corrected independently by a machine learning model. Therefore, using a single ensemble statistic, such as the ensemble mean, as a feature may not fully capitalize on the information provided by using all members of the lagged ensemble as features.

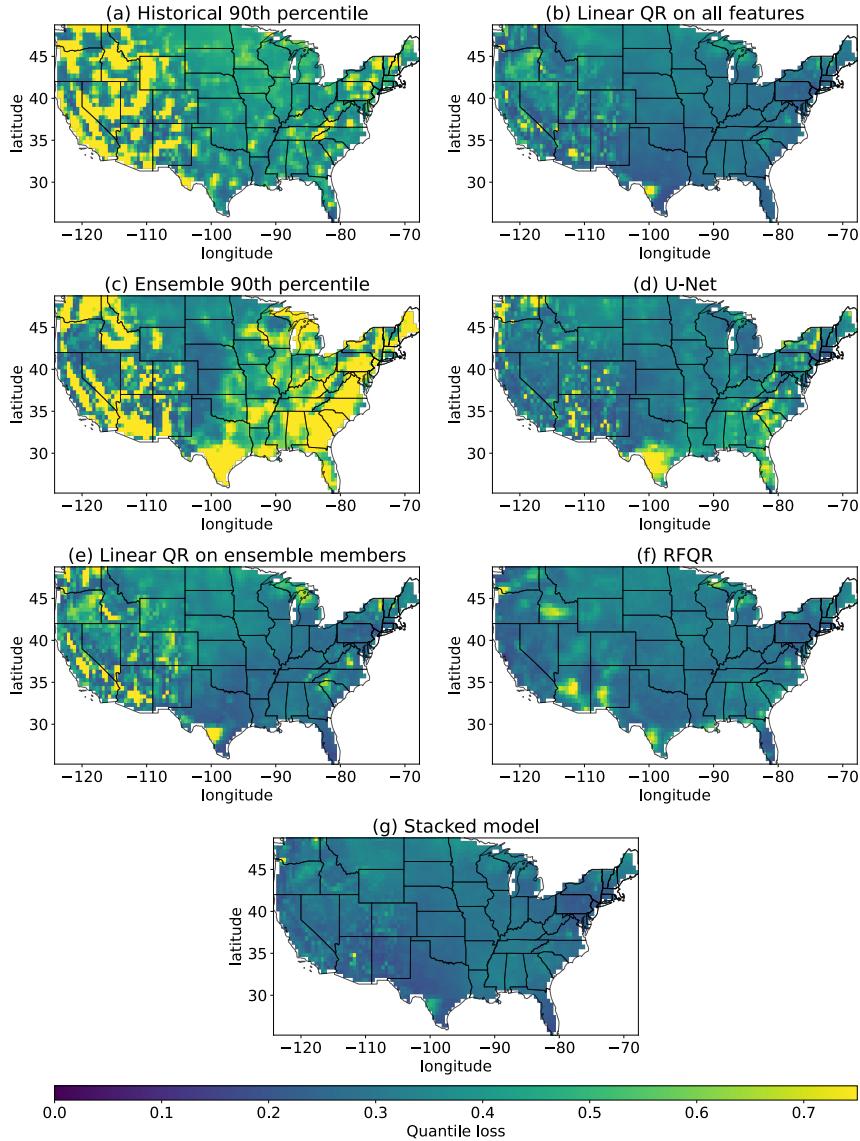


FIG. 4. **Test quantile loss heatmaps of baselines and learning-based methods for temperature quantile regression using NCEP-CFSv2 dataset.** Blue regions indicate smaller quantile loss. See Section 7b(ii) for details.

In our experiments, we find that using all available ensemble members enhances the prediction quality of our approaches. As an illustration, we show the results of the LR, RF, U-Net, and the stacked model trained on all ensemble members, compared to the ML models trained on the ensemble mean only. In addition to the full ensemble or the ensemble mean, we use other available features (as in our previous regression results). Table 7 and Table 8 demonstrate the precipitation

and temperature forecasting results. For the linear regression, utilizing the ensemble mean with all other features produces the best test performance compared to using the full ensemble. Such behavior is not surprising for the LR since the full ensemble incorporates large variance across ensemble members, which may result in a worse linear fit. For the U-Net, RF, and stacked model, we observe significant performance improvements, in terms of having at least one standard error smaller MSE, when using the full ensemble instead of using the ensemble mean. When we compare the performance of learned models using only the ensemble mean to that of the learned models that use both the ensemble mean and the ensemble standard deviation for each spatial location, we find that the addition of the standard deviation feature does not provide enough information to significantly improve the performance of ML models, and in fact the U-Net that exhibits a performance degradation – a potential a sign of overfitting. These observations are visually supported by Figure 5 and Figure 6, where the  $R^2$  heatmaps of our methods (except the U-Net) utilizing ensemble mean and standard deviation closely resemble the performance of methods solely relying on ensemble mean. We conclude that the full ensemble contains important information for SSF aside from the ensemble mean, and our models can capitalize on this information for precipitation and temperature forecasting.

We can perform a statistical test to verify that the performance discrepancies between using the ensemble mean and using the full ensemble are statistically significant for the stacked model. As before, let  $\hat{y}_{t,l}$  refer to the estimate under our usual stacked model (i.e., with all ensemble members). Let  $\hat{y}_{t,l}^{\text{SEA}}$  refer to a stacked model with just the ensemble mean as a feature, instead of all ensemble members. We can employ a sign test framework (DelSole and Tippett 2014; Cash et al. 2019) to compare model performance under minimal distributional assumptions. Namely, we only make the following i.i.d. assumption over the time dimension:

$$\mathbb{I}_{\{|\hat{y}_{t,l} - y_{t,l}| < |\hat{y}_{t,l}^{\text{SEA}} - y_{t,l}|\}} \stackrel{\text{iid}}{\sim} \text{Bernoulli}(p_l)$$

Intuitively, this corresponds to assuming it is a coin flip which model will perform better at each time point and location, and we would like to test whether each location’s “coin” is fair or not. We can then formulate our null and alternative hypotheses for each location  $l$  as follows:

$$H_{0,l} : p_l = 0.5, \quad H_{0,l} : p_l > 0.5$$

Thus, our overall test for significance is for the global null hypothesis  $H_0 = \cap_{l=1}^{3274} H_{0,l}$ . We calculate a p-value for each  $H_{0,l}$ , and then we check whether any of these p-values is below a Bonferroni-corrected threshold of  $0.05/3274 = 1.53 \times 10^{-5}$ , where 3274 refers to the number of locations. In fact, the minimum p-values for this test with precipitation and regression alike are far below this threshold ( $1.68 \times 10^{-10}$  and  $4.42 \times 10^{-24}$ , respectively). This allows us to reject the global null hypothesis for both temperature and precipitation, and we conclude that including the full ensemble in our stacked model significantly outperforms including just the ensemble mean.

TABLE 7. **Precipitation forecasting performance comparison of the LR, RF, U-Net, and stacked model trained using the ensemble mean, using the sorted ensemble members, or using the original ensemble, in addition to other features.** Scores on the test data are reported, and NCEP-CFSv2 data is used. The **best** results are in bold. MSE is reported in squared mm.

Model	Features	Mean $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )
LR	Ensemble mean + all features	<b>-0.28</b>	<b>2.59</b> $\pm$ 0.04
	Ensemble mean & std + all features	-0.29	2.61 $\pm$ 0.04
	Shuffled ensemble + all features	-0.41	2.84 $\pm$ 0.05
	Sorted ensemble + all features	-0.43	2.87 $\pm$ 0.05
	Full ensemble + all features	-0.33	2.71 $\pm$ 0.05
U-Net	Ensemble mean + all features	-0.45	2.76 $\pm$ 0.04
	Ensemble mean & std + all features	-0.25	2.65 $\pm$ 0.04
	Shuffled ensemble + all features	-0.27	2.77 $\pm$ 0.05
	Sorted ensemble + all features	-0.43	2.78 $\pm$ 0.04
	Full ensemble + all features	<b>-0.1</b>	<b>2.18</b> $\pm$ 0.03
RF	Ensemble mean + all features	-0.16	2.36 $\pm$ 0.04
	Ensemble mean & std + all features	-0.15	2.32 $\pm$ 0.04
	Shuffled ensemble + all features	-0.16	2.29 $\pm$ 0.04
	Sorted ensemble + all features	-0.18	2.30 $\pm$ 0.04
	Full ensemble + all features	<b>-0.11</b>	<b>2.17</b> $\pm$ 0.05
Stacked	Ensemble mean + all features	-0.08	2.28 $\pm$ 0.04
	Ensemble mean & std + all features	-0.05	2.26 $\pm$ 0.04
	Shuffled ensemble + all features	-0.04	2.25 $\pm$ 0.04
	Sorted ensemble + all features	-0.11	2.24 $\pm$ 0.04
	Full ensemble + all features	<b>0.02</b>	<b>2.07</b> $\pm$ 0.03

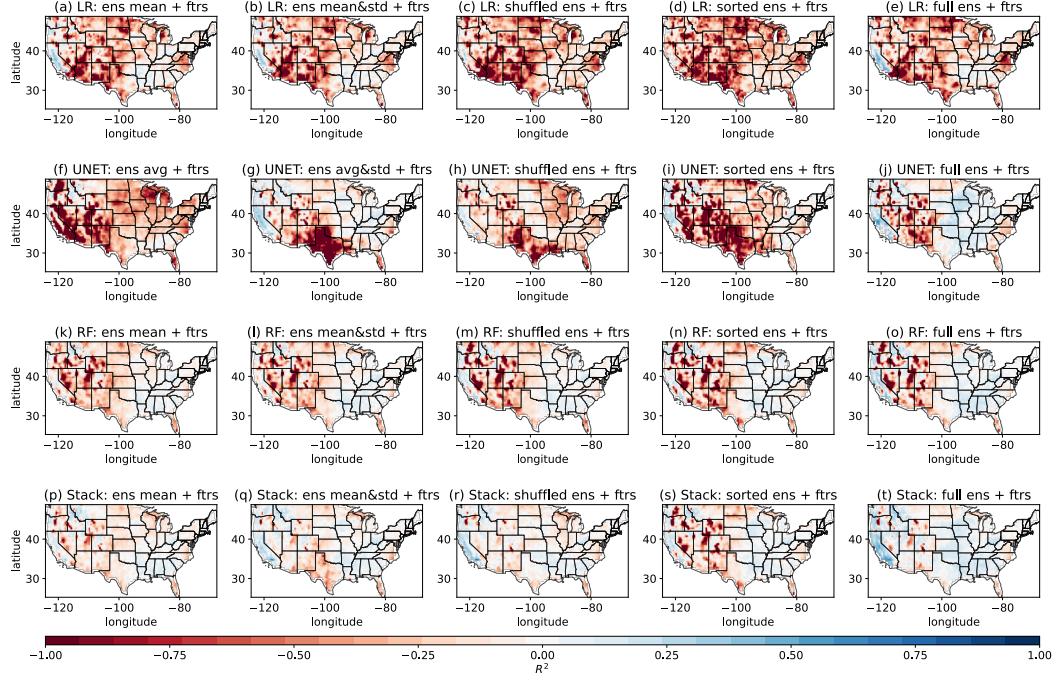


FIG. 5. **Precipitation regression test  $R^2$  heatmaps of LR, U-Net, RF, and stacked model trained using ensemble mean only, using sorted and shuffled ensemble, or using the full ensemble.** The NCEP-CFSv2 ensemble is used. See Section 8a for details.

(ii) *Sensitivity to ensemble formulation* We consider the hypothesis that there is a set of  $k$  ensemble members that are always best. To test this hypothesis, we use a training period to identify which  $k$  members perform best for each location, and then during the test period, compute the average of only these  $k$  ensemble members. The performance of this approach depends on  $k$ , the number of ensemble members we allow to be designated “good.” We have not found that the performance for any  $k$  exhibits a significant improvement over the ensemble mean.

If the ensemble members have different levels of accuracy over various seasons, locations, and conditions, then a machine learning model may be learning when to “trust” each member. We know that our ensemble members are lagged, meaning they are initialized at different times. We believe each ensemble member encapsulates valuable information derived from the underlying physical model during each initialization. To investigate the impact of ensemble member order, we perform the following experiment: we randomly permute ensemble members at every time step  $t$  for all locations (preserving the spatial information) and apply our ML models to these shuffled ensembles. From Table 7 and Table 8, this approach negatively affects the performance of the ML

TABLE 8. **Temperature forecasting performance comparison of the LR, RF, U-Net, and stacked model trained using the ensemble mean, using the sorted ensemble members, or using the original ensemble, in addition to other features.** Scores on the test data are reported and NCEP-CFSv2 data is used. The **best** results are in bold. MSE is reported in squared  $^{\circ}\text{C}$ .

	Model	Features	Mean $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )
LR	Ensemble mean + all features	<b>0.06</b>	<b>3.55</b> $\pm$ 0.03	
	Ensemble mean & std + all features	0.05	3.59 $\pm$ 0.03	
	Shuffled ensemble + all features	0.03	3.95 $\pm$ 0.03	
	Sorted ensemble + all features	-0.02	3.87 $\pm$ 0.03	
	Full ensemble + all features	0.05	3.57 $\pm$ 0.02	
U-Net	Ensemble mean + all features	0.00	3.77 $\pm$ 0.03	
	Ensemble mean & std + all features	0.19	4.61 $\pm$ 0.03	
	Shuffled ensemble + all features	0.29	4.75 $\pm$ 0.03	
	Sorted ensemble + all features	-0.94	6.51 $\pm$ 0.05	
	Full ensemble + all features	<b>0.01</b>	<b>3.65</b> $\pm$ 0.02	
RF	Ensemble mean + all features	0.10	3.57 $\pm$ 0.02	
	Ensemble mean & std + all features	0.10	3.56 $\pm$ 0.02	
	Shuffled ensemble + all features	0.05	3.65 $\pm$ 0.02	
	Sorted ensemble + all features	0.10	3.44 $\pm$ 0.02	
	Full ensemble + all features	<b>0.16</b>	<b>3.17</b> $\pm$ 0.02	
Stacked	Ensemble mean + all features	0.11	3.43 $\pm$ 0.02	
	Ensemble mean & std + all features	0.13	3.30 $\pm$ 0.02	
	Shuffled ensemble + all features	0.08	3.47 $\pm$ 0.02	
	Sorted ensemble + all features	0.03	3.70 $\pm$ 0.02	
	Full ensemble + all features	<b>0.18</b>	<b>3.11</b> $\pm$ 0.02	

models compared to using the full ensemble with the original order. One possible explanation is that the learned models lose the ability to learn which ensemble member to trust, as this information is tied to the initialization time of each ensemble member. Even though the spatial information remains intact after the shuffling, the models can no longer exploit dependencies associated with the original ensemble structure.

Additionally, we conduct an experiment designed to test whether it is important to keep track of which ensemble member made each prediction or whether it is the *distribution* of predictions that is important. The modeling approach for the former would be to feed in ensemble member 1's

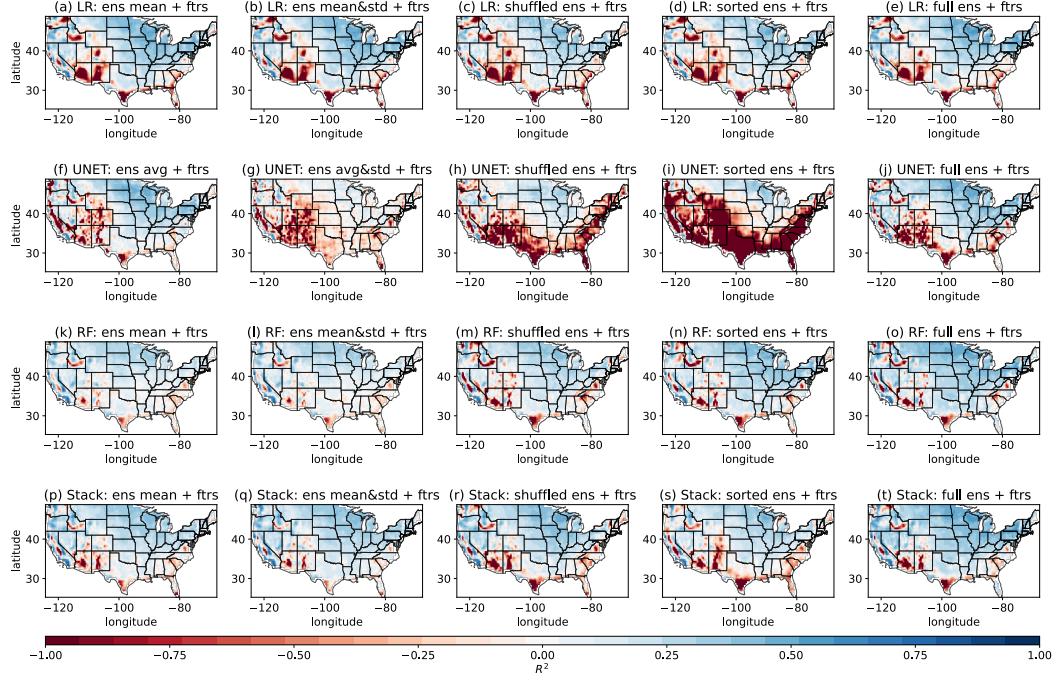


FIG. 6. Temperature regression test  $R^2$  heatmaps of LR, U-Net, RF and stacked model trained using ensemble mean only, using sorted and shuffled ensemble, or using the full ensemble. The NCEP-CFSv2 ensemble is used. See Section 8a for details.

forecast as the first feature, ensemble member 2's forecast as the second feature, etc. The modeling approach under the distributional hypothesis is to make the smallest prediction be the first feature, the second-smallest prediction be the second feature, and so on – i.e., we sort the ensemble forecasts for each location separately. Note that this entails treating the ensemble members symmetrically: the model would give the same prediction if ensemble member 1 predicted  $a$  and ensemble member 2 predicted  $b$  or if ensemble member 1 predicted  $b$  and ensemble member 2 predicted  $a$ . In statistical parlance, this is passing in the order statistics of the forecasts as the features rather than their original ordering. Note that for NCEP-CFSv2, ensemble forecasts are originally ordered according to the time their initial conditions are set (Saha et al. 2014). According to Table 7 and Table 8, using the sorted ensemble drastically degrades U-Net's performance, which is essentially because we sort the ensemble members for each location individually, and sorting the ensemble members individually for each location may hamper the ability for the U-Net to learn spatial structure. In the case of precipitation regression with the stacking model from Table 7, the MSE of the sorted approach is 2.24, which is worse than the 2.07 MSE for using the original ordering.

In the case of temperature forecasting from Table 8, the MSE of the sorted approach is 3.70, which is much worse than the 3.11 MSE for using the original ordering. The mean  $R^2$  of the sorted approach is also lower compared to the original ordering. In both cases, the performance is better when we feed in the features in such a way that the machine learning model has an opportunity to learn aspects of each ensemble member, not merely their order statistics. Therefore, imposing a symmetric treatment of ensemble members degrades performance. Figure 5 and Figure 6 shows the corresponding  $R^2$  heatmaps of our models for precipitation and temperature regression tasks.

### *b. Using spatial data*

There are several ways to incorporate information about location in our models. U-Net has access to spatial dependencies through its design. Specifically, our U-Net inputs the spatial location of each point in the map. Naively, we might represent each location using the latitude and longitude values. Alternatively, we may use positional encoding, which is known to be beneficial in many ML areas, not only in NLP (as we mention in Section 6a). PE captures the order (or position) and allows one to learn the contextual relationships (local context – relationships between nearby elements and global context dependencies across the entire sequence). We assume that the PE approach represents spatial information in a manner more accessible to our learned models.

As an illustration, Table 9 and Figure 7 demonstrate the performance of a stacked model using LR, RF, and U-Net trained using positional encoding, using latitude and longitude values and using no features representing the spatial information (no PE and no latitude or longitude values). Other inputs to the LR, RF, and U-Net models are ensemble member forecasts, lagged target variable, climate variables, and SSTs (except in the case of temperature forecasting). The results suggest that using PE enhances the predictive skill of our models, compared to using just the lat/lon values or no location information, especially for the temperature forecasting task. Using no information about locations hurts the performance of precipitation regression. Thus, our models can account for spatial dependencies using input features, and PE is more beneficial than the raw latitude and longitude information. These findings on PE effectiveness are consistent with prior findings in ML. For example, Wu et al. (2021) investigate the efficacy of PE in the context of the visual transformer model used for image classification and object detection. We show a more detailed analysis with results for the LR, RF, and U-Net in Section C1.

TABLE 9. **Test performance comparison of the stacked model of LR, RF, and U-Net trained using no spatial features, using latitude and longitude values, or using PE.** Utilizing spatial representations, including PE, latitude, and longitude values, helps advance the predictive skill. Furthermore, using positional encoding is more beneficial than using raw latitude and longitude values. The **best** results are in bold. MSE is reported in squared mm for precipitation and in squared  $^{\circ}\text{C}$  for temperature.

Target	Features	Mean $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )
Precip	All + no location info	-0.05	2.13 $\pm$ 0.03
	All + lat/lon values	-0.01	2.21 $\pm$ 0.04
	All + PE	<b>0.02</b>	<b>2.07</b> $\pm$ 0.03
Tmp	All + no location info	0.12	3.35 $\pm$ 0.02
	All + lat/lon values	0.12	3.33 $\pm$ 0.02
	All + PE	<b>0.18</b>	<b>3.11</b> $\pm$ 0.02

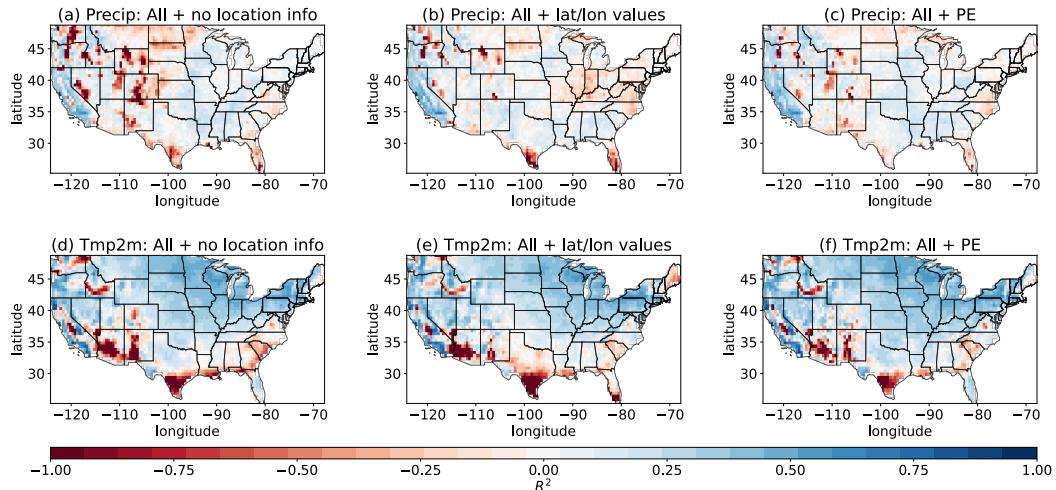


FIG. 7. **Test  $R^2$  heatmaps of the stacked model of LR, RF, and U-Net trained using no spatial features, using latitude and longitude values or using PE.** The NCEP-CFSv2 ensemble is used. See section 8b for details.

### c. Variable importance

One consideration when implementing ML for SSF is that ML models can incorporate side information (such as spatial information, lagged temperature and precipitation values, and climate variables). In this section, we explore the importance of the various components of side information. We see that including the observational climate variables improves the performance for the random

forest and the U-Net when doing precipitation regression. Furthermore, including positional encoding of the locations improves the performance of the U-Net, while the principle components of the sea surface temperature do not make a notable difference in the case of temperature prediction.

More specifically, Table 10 summarizes grouped feature importance of precipitation regression using the NCEP-CFSv2 ensemble. We observe that models, in particular random forest and U-Net, trained on all available data achieve the best performance. In the case of linear regression, the SSTs features are neither very helpful nor actively harmful. Therefore, to be consistent, we use predictions of these models trained on all features as input to the stacking model.

TABLE 10. **Grouped feature importance results on validation for precipitation regression task using NCEP-CFSv2 ensemble members.** The results suggest that using additional observational information helps to improve the performance of learning-based models for this task. “–” means a repetition of features that are used above. For example, in the U-Net part of the table, “– & lags” means that ensemble members, PE, and lags are used as features and “– & SSTs” means ensemble members, PE and lags, land features and SSTs are used as features. The **best** results are in bold. MSE is reported in squared mm.

Model	Features	Mean $R^2$ (↑)	Median $R^2$ (↑)	Mean Sq Err (↓)	Median MSE (↓)	90th prctl MSE (↓)
LR	Ens members	-0.13	-0.08	$2.11 \pm 0.03$	1.53	4.63
	– & lags	-0.11	-0.07	$2.10 \pm 0.03$	1.50	4.59
	– & climate variables (no SSTs)	<b>-0.09</b>	<b>-0.06</b>	<b><math>2.06 \pm 0.03</math></b>	<b>1.47</b>	<b>4.52</b>
	– & SSTs	-0.10	-0.07	$2.08 \pm 0.03$	1.47	4.61
U-Net	Ens members with PE	-0.13	-0.05	$2.01 \pm 0.03$	1.50	4.31
	– & lags	-0.08	-0.02	$1.92 \pm 0.03$	1.42	4.17
	– & climate variables (no SSTs)	-0.02	0.05	$1.86 \pm 0.03$	1.37	4.02
	– & SSTs	<b>0.00</b>	<b>0.05</b>	<b><math>1.83 \pm 0.03</math></b>	<b>1.34</b>	<b>3.94</b>
RF	Ens members with PE	-0.15	-0.04	$2.02 \pm 0.03$	1.49	4.34
	– & lags	-0.10	0.00	$1.96 \pm 0.03$	1.44	4.21
	– & climate variables (no SSTs)	-0.08	0.02	$1.93 \pm 0.03$	1.39	4.16
	– & SSTs	<b>-0.06</b>	<b>0.04</b>	<b><math>1.89 \pm 0.03</math></b>	<b>1.36</b>	<b>4.08</b>

Table 11 summarizes grouped feature importance of temperature regression using the NCEP-CFSv2 ensemble. In this case, adding some types of side information may yield only very small improvements to predictive skill, and in some cases, the additional information may decrease predictive skill. On the one hand, this effect can be explained by different training set sizes for different models: as we outline in Section 5d, the training set size for RF is  $n = TL$ , while

$n = T$  for U-Net. This effect also may be a sign of overfitting, as temperature forecasting presents a comparatively less complex challenge than precipitation forecasting. We also note that SSTs provide only marginal (if any) improvement in predictive skill, in part because Pacific SSTs are less helpful away from the western U.S. (Mamalakis et al. 2018; Seager et al. 2007). It could also be that information from the SSTs is already being well-captured by the output from the dynamical models, and thus, including observed SSTs does not provide much additional information. In order to be consistent, we use predictions of these models trained on all features except SSTs as input to the stacking model.

TABLE 11. **Grouped feature importance results on validation for temperature regression task using NCEP-CFSv2 ensemble members.** The results demonstrate that using some additional information may yield only very small improvements in predictive skill, and in some cases, the side information may decrease predictive skill. “–” means a repetition of features that are used above. For example, in the U-Net part of the table, “– & lags” means that ensemble members, PE, and lags are used as features and “– & SSTs” means ensemble members, PE and lags, land features and SSTs are used as features. The **best** results are in bold. MSE is reported in squared  $^{\circ}\text{C}$ .

Model	Features	Mean $R^2$ (↑)	Median $R^2$ (↑)	Mean Sq Err (↓)	Median MSE (↓)	90th prctl MSE (↓)
LR	Ens members	0.35	0.40	$2.19 \pm 0.02$	2.00	3.47
	– & lags	<b>0.37</b>	<b>0.40</b>	<b><math>2.12 \pm 0.02</math></b>	<b>1.94</b>	<b>3.30</b>
	– & climate variables (no SSTs)	0.36	0.39	$2.14 \pm 0.04$	1.94	3.40
	– & SSTs	0.34	0.38	$2.23 \pm 0.02$	1.99	3.73
U-Net	Ens members with PE	<b>0.33</b>	<b>0.41</b>	<b><math>2.22 \pm 0.04</math></b>	<b>2.02</b>	<b>3.47</b>
	– & lags	0.32	0.40	$2.24 \pm 0.02$	2.02	3.49
	– & climate variables (no SSTs)	0.31	0.41	$2.26 \pm 0.02$	2.08	3.48
	– & SSTs	0.28	0.38	$2.47 \pm 0.02$	2.20	3.95
RF	Ens members with PE	0.11	0.37	$2.85 \pm 0.04$	2.28	4.87
	– & lags	0.30	0.36	$2.35 \pm 0.02$	2.12	3.70
	– & climate variables (no SSTs)	<b>0.30</b>	<b>0.36</b>	<b><math>2.33 \pm 0.02</math></b>	<b>2.10</b>	<b>3.65</b>
	– & SSTs	0.28	0.34	$2.42 \pm 0.02$	2.17	3.83

## 9. Conclusions and future directions

This paper systematically explores the use of machine learning methods for subseasonal forecasting, highlighting several important factors: (1) the importance of using ensembles of physics-based

forecasts (as opposed to only using the mean, as in common practice); (2) the potential for forecasting temperature and precipitation extremes using quantile regression; (3) the efficacy of different mechanisms, such as positional encoding and convolutional neural networks, for modeling spatial dependencies; (4) the importance of various features, such as sea surface temperature and lagged temperature and precipitation values, for predictive accuracy; (5) model stacking provides substantial benefits by leveraging the different utilization of spatial data among contributing models. The stacking model probably capitalizes on this diversity, fostering performance enhancement. Together, these results provide new insights into using ML for subseasonal weather forecasting in terms of the selection of features, models, and methods.

Our results also suggest several important directions for future research. In terms of **features**, there are many climate forecasting ensembles computed by organizations such as NOAA and ECMWF. This paper focuses on ensembles in which ensemble members have a distinct ordering (in terms of lagged initial conditions used to generate them), but other ensembles correspond to initial conditions or parameters drawn independently from some distribution. Leveraging such ensemble forecasts and potentially jointly leveraging ensemble members from multiple distinct ensembles may further improve the predictive accuracy of our methods.

In terms of **models**, new neural architecture models such as transformers show remarkable performance on several image analysis tasks (Dosovitskiy et al. 2020; Carion et al. 2020; Chen et al. 2021; Khan et al. 2022) and have potential in the context of forecasting climate temperature and precipitation maps. A careful study is needed, as past image analysis work using transformers generally uses large quantities of training data, exceeding what is available in SSF contexts. Recent advancements in data-driven global weather forecasting models, such as Pangu-Weather (Bi et al. 2022), FourCastNet (Pathak et al. 2022), and GenCast (Price et al. 2023), demonstrate the potential of ML techniques to enhance forecasting capabilities across various timescales. These models outperform traditional numerical weather prediction approaches, suggesting that similar data-driven methods may hold promise for improving SSF quality.

In terms of **methods**, two outstanding challenges are particularly salient to the SSF community. The first is uncertainty quantification; that is, we wish not only to forecast temperature or precipitation but also to predict the likelihood of certain extreme events. Our work on quantile regression is an important step in this direction and statistical methods like conformalized quantile

regression (Romano et al. 2019) may provide additional insights. Second, we see in Figure C7 that, at least in some geographic regions, the distribution of ensemble hindcast and forecast data may be quite different. Employing methods that are more robust to *distribution drift* (Wiles et al. 2021; Subbaswamy et al. 2021; Zhu et al. 2021) is particularly important not only for handling forecast and hindcast data but also for accurate SSF in a changing climate.

*Acknowledgments.* The authors gratefully acknowledge the support of the NSF (OAC-1934637, DMS-1930049, and DMS-2023109) and C3.ai.

*Data availability statement.* All data used in the experiments are publicly available. NOAA data for historical averages is available at <https://www.ncei.noaa.gov/access/monitoring/climate-at-a-glance/national/time-series>. Other climate variables (tmp2m, precip, slp, hgt500, rhum, SSTs) are also public access at <https://sites.google.com/view/ssf-dataset>. The NCEP-CFSv2 ensemble members are available at <https://iridl.ldeo.columbia.edu/SOURCES/.Models/.NMME/.NCEP-CFSv2/>, The NASA-GMAO ensemble members are available at <https://iridl.ldeo.columbia.edu/SOURCES/.Models/.NMME/.NASA-GMAO/>. Our code is available at <https://github.com/eleno-orlova/SSF-project>.

## APPENDIX A

### Regression results for NASA-GMAO and NCEP-CFSv2

#### A1. Regression

(i) *Precipitation regression using NASA-GMAO* Precipitation regression results on the test data from NASA-GMAO are presented in Table A1. On this dataset, no learned method or method leveraging ensemble model forecasts significantly outperforms climatology. Note that the best  $R^2$  value associated with the climatology is still negative; the low values for all methods indicate the difficulty of the forecasting problem.

Figure A1 illustrates the test performance of key methods on NASA-GMAO data with  $R^2$  heatmaps over the U.S. Although the stacked model does not show the best performance in terms of mean  $R^2$  score, it has more geographic regions with positive  $R^2$  than any other method.

(ii) *Regression of temperature* Temperature regression results using NASA-GMAO ensemble members are presented in Table A2. The random forest and linear regression outperform all baselines in terms of both  $R^2$  score and MSE. However, the U-Net model's performance is lower compared to other learned methods, which might be a sign of overfitting. Despite this performance drop of U-Net, the model stacking approach still demonstrates the best predictive skill. Note that the model stacking approach is applied to the models that are trained on all available features except SSTs (similar to NCEP-CFSv2 data).

TABLE A1. **Test results for precipitation regression using NASA-GMAO dataset.** LR refers to linear regression on all features including ensemble members, lagged data, land variables, and SSTs. Model stacking is performed on models that are learned on all features. Bold values indicate the best performance for each statistic. MSE is reported in squared mm.

Model	Features	Mean $R^2$ ( $\uparrow$ )	Median $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )	Median MSE ( $\downarrow$ )	90th prctl MSE ( $\downarrow$ )
Baseline	Climatology	<b>-0.07</b>	<b>-0.02</b>	$2.14 \pm 0.04$	<b>1.51</b>	4.40
	Ens mean	-0.11	-0.06	$2.13 \pm 0.04$	1.52	4.31
	Linear Regr	-0.18	-0.14	$2.25 \pm 0.04$	1.62	4.68
LR	All features	-0.40	-0.29	$2.62 \pm 0.05$	1.93	5.42
U-Net	All features	-0.19	-0.09	$2.11 \pm 0.03$	1.56	<b>4.25</b>
RF	All features	-0.18	-0.11	$2.17 \pm 0.04$	1.55	4.44
Stacked	LR, U-Net, RF, outputs	-0.08	-0.06	<b><math>2.09 \pm 0.04</math></b>	1.52	4.27

TABLE A2. **Test results for temperature regression using NASA-GMAO dataset.** LR refers to linear regression on all features including ensemble members, lagged data, and land variables. Model stacking is performed on models that are learned on all features except SSTs. Bold values indicate the best performance for each statistic. MSE is reported in squared  $^{\circ}\text{C}$ .

Model	Features	Mean $R^2$ ( $\uparrow$ )	Median $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )	Median MSE ( $\downarrow$ )	90th prctl MSE ( $\downarrow$ )
Baseline	Climatology	-0.70	-0.20	$6.49 \pm 0.11$	5.06	9.72
	Ens mean	<b>-0.28</b>	0.12	$4.82 \pm 0.10$	3.43	7.82
	Linear Regr	0.12	0.14	$3.32 \pm 0.02$	3.11	4.70
LR	All features wo SSTs	0.17	0.17	$3.10 \pm 0.02$	3.05	4.26
U-Net	All features wo SSTs	0.06	0.12	$3.40 \pm 0.02$	3.27	4.52
RF	All features wo SSTs	0.20	0.22	$3.03 \pm 0.02$	2.94	4.25
Stacked	LR, U-Net, RF, outputs	<b>0.21</b>	<b>0.22</b>	<b><math>2.94 \pm 0.02</math></b>	<b>2.89</b>	<b>3.97</b>

Figure A2 illustrates the test performance of key methods on NASA-GMAO data with  $R^2$  heatmaps over the U.S. The stacked model shows the best performance across spatial locations. Similar to the NCEP-CFSv2 dataset, we notice that there are still regions where all models tend to exhibit a negative  $R^2$  score.

## A2. Quantile regression

(i) *Quantile regression of precipitation using NCEP-CFSv2 ensemble* In Figure A3, we show heatmaps of quantile loss using all locations in the U.S., where blue means smaller quantile loss

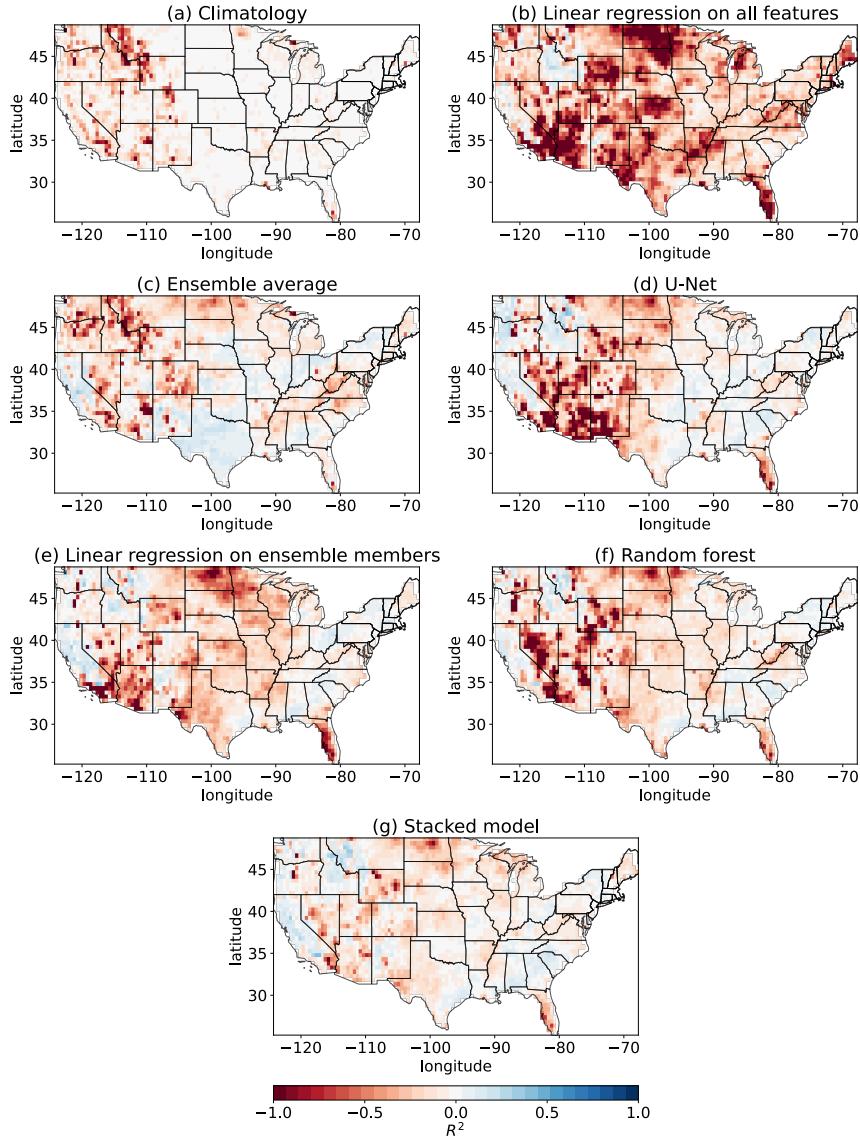


FIG. A1. Test  $R^2$  score heatmaps of baselines and learning-based methods for precipitation regression using the NASA-GMAO dataset. Positive values (blue) indicate better performance. See Section A1(i) for details.

and yellow means larger quantile loss. We observe that the learning-based models outperform the baselines, especially in Washington, California, Idaho, and near the Gulf of Mexico.

(ii) *Quantile regression of precipitation using NASA-GMAO ensemble* Table A3 summarizes results for precipitation quantile regression using the NASA-GMAO ensemble. The models are the same as the models applied to the NCEP-CFSv2 ensemble. Our best model shows similar

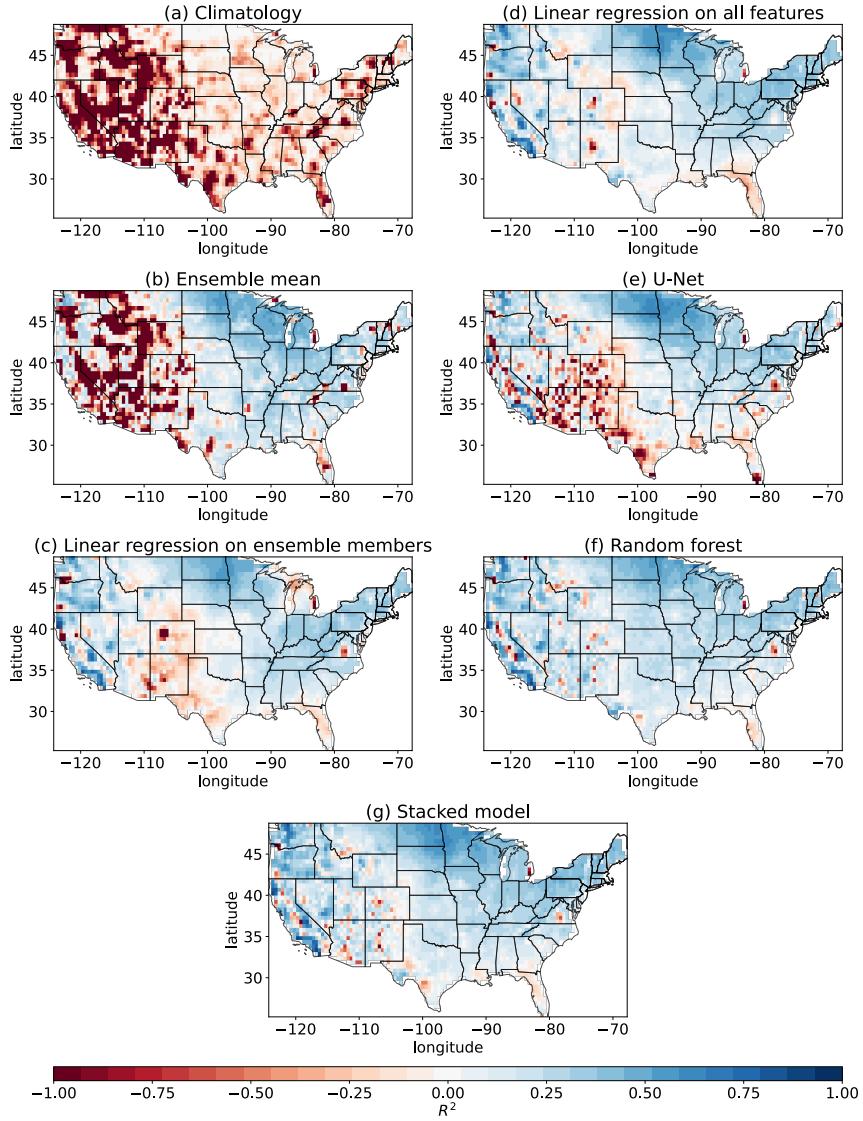


FIG. A2. **Test  $R^2$  score heatmaps of baselines and learning-based methods for temperature regression using NASA-GMAO dataset.** Positive values (blue) indicate better performance. See Section A1(ii) for details.

performance to that of the baselines. One possible reason is that the NASA-GMAO ensemble shows worse performance than the NCEP-CFSv2 ensemble empirically. Furthermore, according to the designs of ensemble members from both climate models, the NASA-GMAO ensemble has fewer ensemble members and, therefore, has less coverage on the distribution of precipitation than the NCEP-CFSv2 ensemble, so our learned model has access to less information about the true distribution of precipitation.

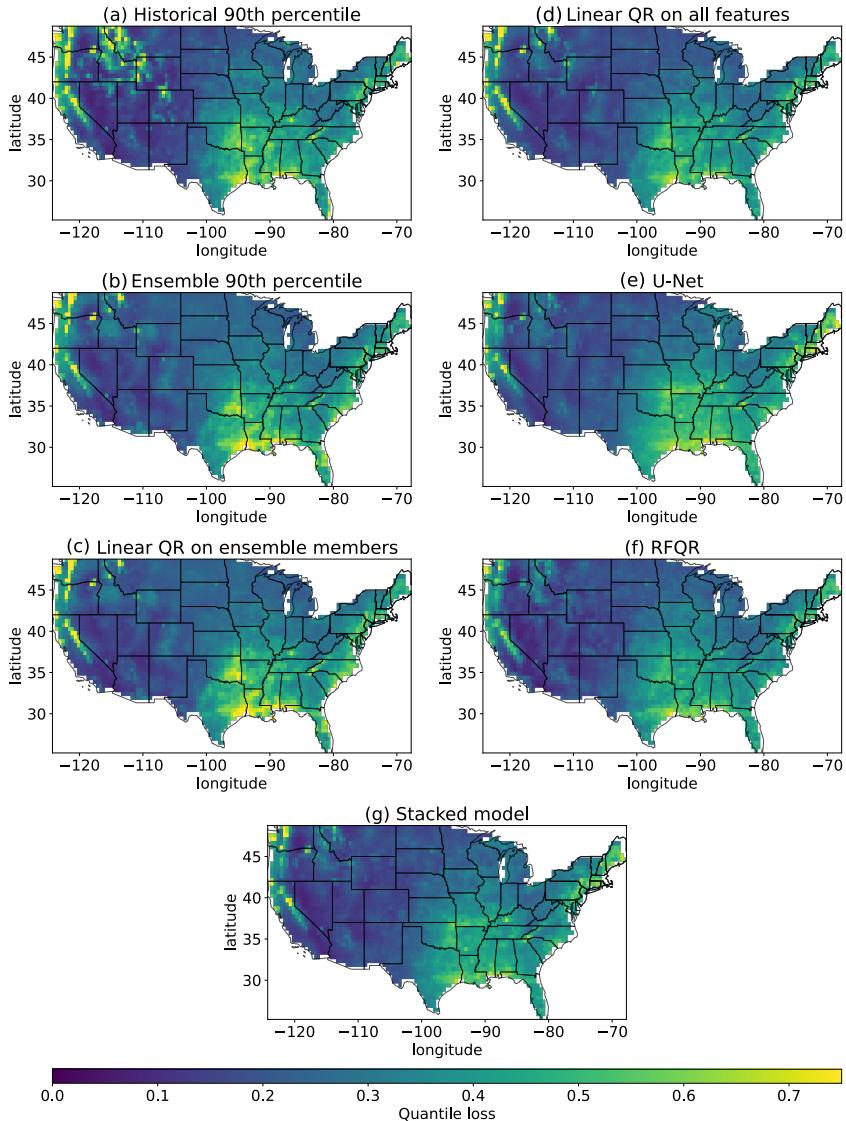


FIG. A3. **Test quantile loss heatmaps of baselines and learning-based methods for precipitation quantile regression using NCEP-CFSv2 dataset.** Blue regions indicate smaller quantile loss. See Section 7b(i) for details.

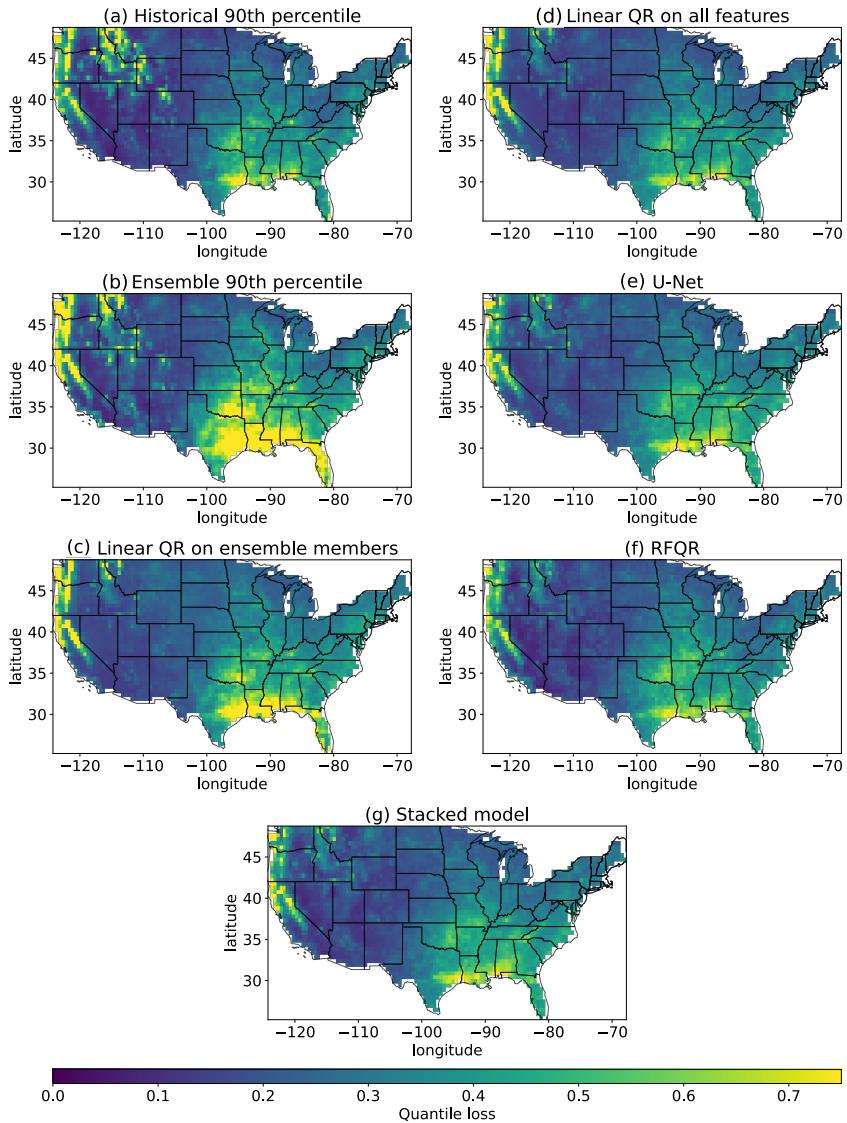
(iii) *Quantile regression of temperature* Table A4 summarizes results for temperature quantile regression using the NASA-GMAO ensemble. All of our learned models are able to outperform all baselines.

TABLE A3. **Test results for precipitation quantile regression using NASA-GMAO dataset, with target quantile = 0.9.** Model stacking is performed on models that are learned on all features.

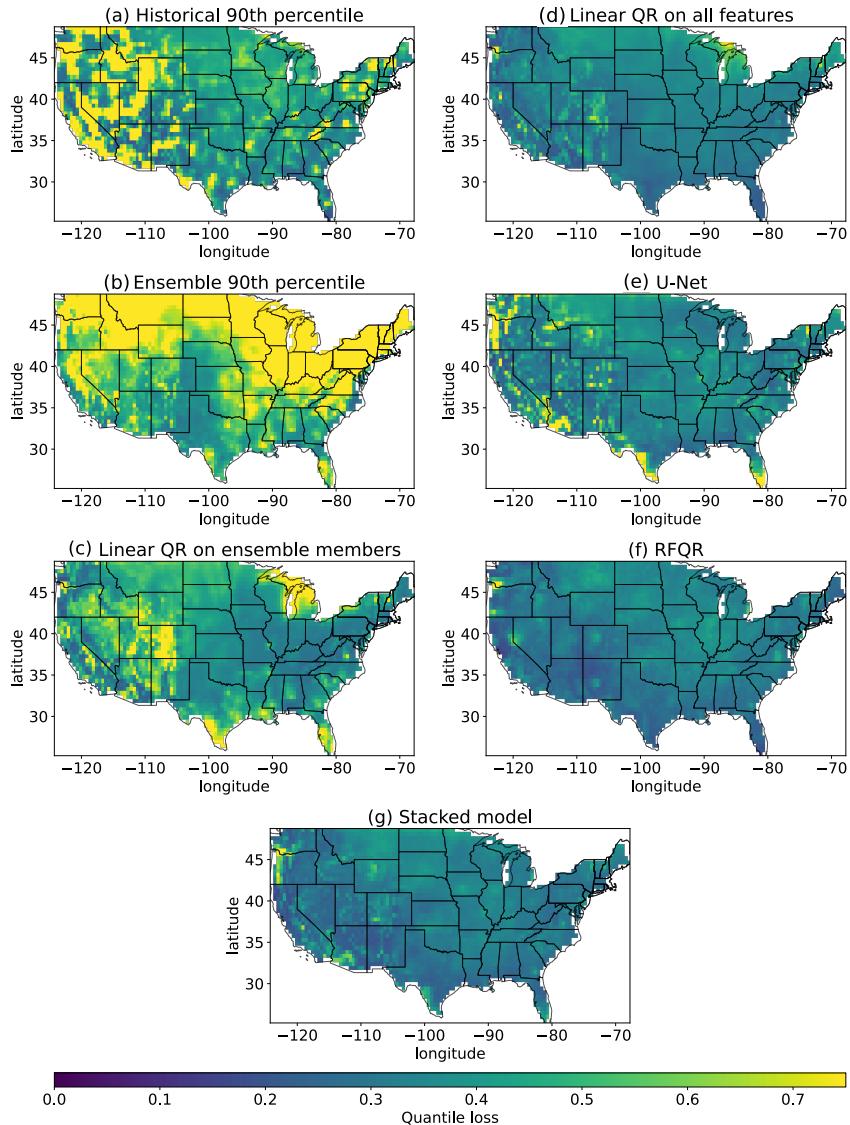
Model	Features	Mean Qtr Loss (↓)	Median Qtr Loss (↓)	90th prctl Qtr Loss (↓)
Baseline	Historical 90-th percentile	0.295 ± 0.003	0.263	0.484
	Ens 90-th percentile	0.378 ± 0.005	0.308	0.673
	Linear QR ens only	0.336 ± 0.004	0.286	0.531
Linear QR	All features	<b>0.290 ± 0.003</b>	<b>0.253</b>	<b>0.456</b>
U-Net	All features	0.310 ± 0.002	0.278	0.489
RFQR	All features	0.290 ± 0.002	0.265	0.471
Stacked	U-Net, RFQR, LQR outputs	0.296 ± 0.002	0.268	0.467

TABLE A4. **Test results for temperature quantile regression using NASA-GMAO dataset, with target quantile = 0.9.** Linear QR refers to a linear quantile regressor, i.e. linear regression fitting the quantile loss, using all features including ensemble members, positional encoding, lagged data, land variables. Model stacking is performed on models that are learned on all features except for SSTs .

Model	Features	Mean Qtr Loss (↓)	Median Qtr Loss (↓)	90th prctl Qtr Loss (↓)
Baseline	Historical 90-th percentile	0.596 ± 0.010	0.438	0.988
	Ens 90-th percentile	0.812 ± 0.009	0.646	1.493
	Linear QR ens only	0.445 ± 0.003	0.411	0.618
Linear QR	All features wo SSTs	0.341 ± 0.001	0.333	0.419
U-Net	All features wo SSTs	0.375 ± 0.003	0.347	0.477
RFQR	All features wo SSTs	0.318 ± 0.002	0.316	0.376
Stacked	U-Net, RFQR, LQR outputs	<b>0.315 ± 0.002</b>	<b>0.310</b>	<b>0.374</b>



**FIG. A4. Test quantile loss heatmaps of baselines and learning-based methods for precipitation quantile regression using NASA-GMAO dataset.** Blue regions indicate smaller quantile loss. See Section A2(ii) for details.



**FIG. A5. Test quantile loss heatmaps of baselines and learning-based methods for temperature quantile regression using NASA-GMAO dataset.** Blue regions indicate smaller quantile loss. See Section A2(iii) for details.

## APPENDIX B

### Tercile Classification

In this section, we present results for the tercile classification task for both climate variables and both datasets.

#### B1. Tercile classification of precipitation

In this case, the proposed learning-based methods are directly trained on the classification task. Predictions of baselines, such as the ensemble mean, are split into three classes according to the 33rd and 66th percentile values. Note that random forest and U-Net are trained for classification using all available features. We do not notice a significant difference in the performance of logistic regression on the validation if the inputs are ensemble members only or ensemble members with side information. So, we use logistic regression on ensemble members only. The model stacking is applied to the logistic regression, U-Net, and random forest outputs.

Table B1 summarizes results for NCEP-CFSv2 and NASA-GMAO datasets on the test data. For this task, the learning-based methods achieve the best performance in terms of accuracy for both datasets. In the case of NCEP-CFSv2 data, U-Net achieves the highest accuracy score, and the performance of the stacked model is comparable with it. For NASA-GMAO data, the stacked model shows the best performance.

The accuracy heatmaps over U. S. land are presented in the Figure B1 for NCEP-CFSv2 dataset. The plots corresponding to the learning-based methods show the best results, especially at the West Coast, Colorado and North America.

The accuracy heatmaps over the U. S. land are presented in Figure B2 for the NASA-GMAO dataset. The plots corresponding to the learning-based methods show the best results, the ensemble mean's figure has the most red regions.

#### B2. Tercile classification of temperature

The next task is tercile classification of 2-meter temperature. In this case, the threshold is applied to the regression predictions of all methods, meaning there is no direct training for a classification. Table B2 summarizes results for NCEP-CFSv2 and NASA-GMAO datasets on the test data. For this task, the learning-based methods achieve the best performance in terms of accuracy, stacked

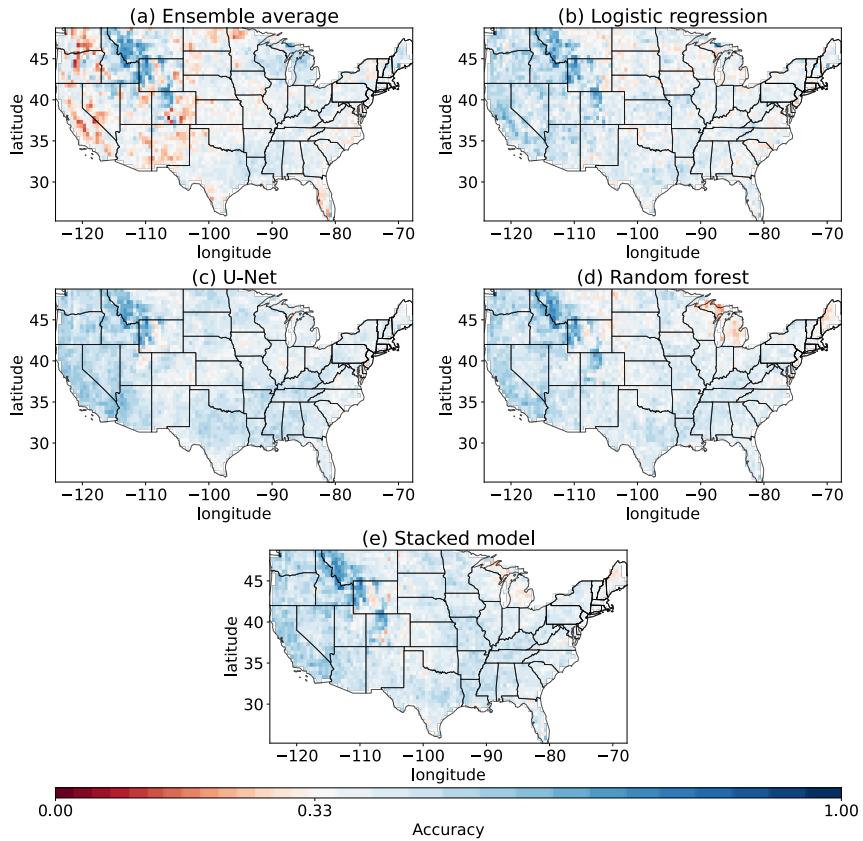
TABLE B1. **Test results for tercile classification of precipitation on different datasets.** Accuracy in % is reported. Note that for this task, our models are trained for classification directly while baselines perform regression, and a threshold for predicted values is applied. For stacking, logistic regression, U-Net and RF outputs are used.

Data	Model	Mean	Median
		accuracy ( $\uparrow$ )	accuracy ( $\uparrow$ )
NCEP-CFSv2	Ens mean	38.00 $\pm$ 0.16	37.61
	Logistic Regr	41.22 $\pm$ 0.14	40.17
	U-Net	<b>43.88</b> $\pm$ 0.12	<b>42.74</b>
	RF	42.38 $\pm$ 0.13	41.88
	Stacked	43.81 $\pm$ 0.13	<b>42.74</b>
NASA-GMAO	Ens mean	38.64 $\pm$ 0.14	37.65
	Logistic regr	41.51 $\pm$ 0.16	40.00
	U-Net	40.53 $\pm$ 0.11	40.00
	RF	40.79 $\pm$ 0.14	40.00
	Stacked	<b>42.08</b> $\pm$ 0.14	<b>41.18</b>

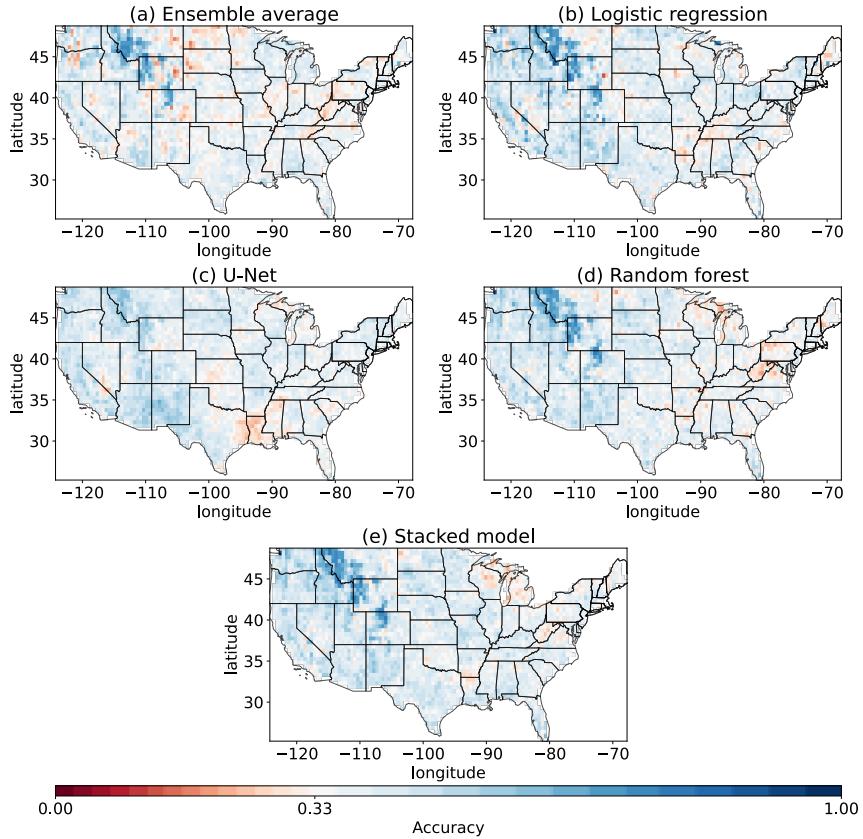
model using NCEP-CFSv2 data and linear regression using NASA-GMAO data and all additional features (except SSTs). In general, all learning-based models significantly outperform the ensemble mean.

Figure B3 shows accuracy heatmaps over the U.S. for different methods using NCEP-CFSv2 data. The stacked model shows the best performance across spatial locations. For example, the ensemble mean does not show great performance in the Southeast and Middle Atlantic regions, while learning-based methods demonstrate much stronger predictive skills in these areas. However, there are still some areas, such as Texas or South West region, with red pixels for all methods.

Figure B4 shows accuracy heatmaps over the U.S. for different methods using NASA-GMAO data. In this case, linear regression on all features achieves the best scores. Other learning-based methods outperform the ensemble mean too, especially in the West and in Minnesota.



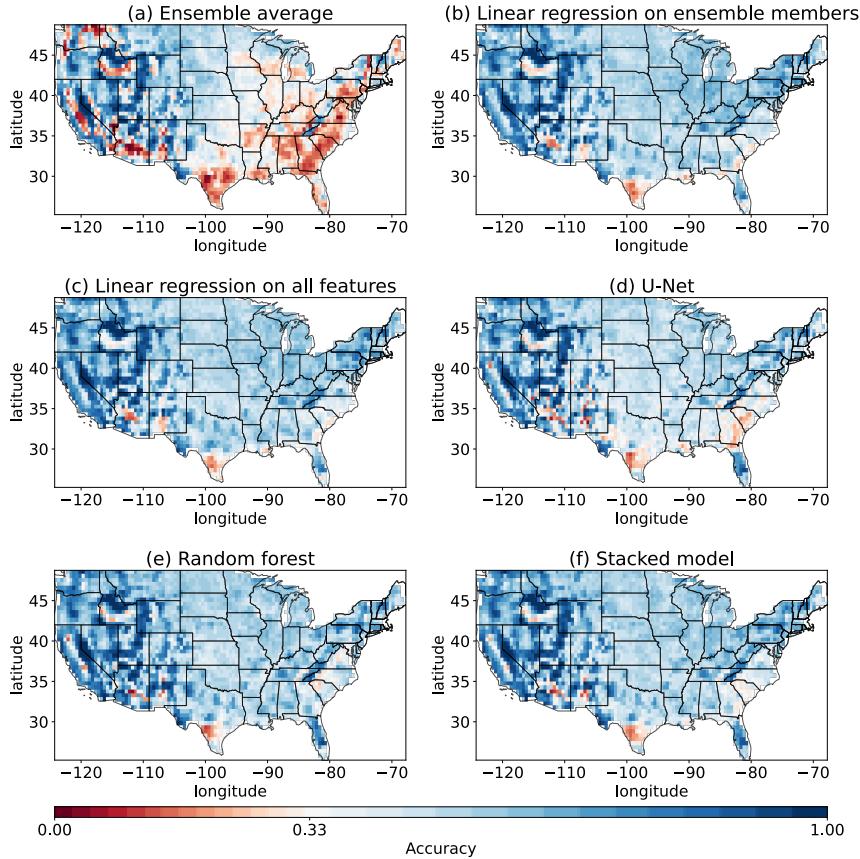
**FIG. B1. Test accuracy heatmaps of baselines and learning-based methods for tercile classification of precipitation using NCEP-CFSv2 dataset.** The accuracy colorbar is recentered to be white at  $\frac{1}{3}$ , what corresponds to a random guess score. Blue pixels indicate better performance, while red pixels correspond to performance that is worse than a random guess. See Section B1 for details.



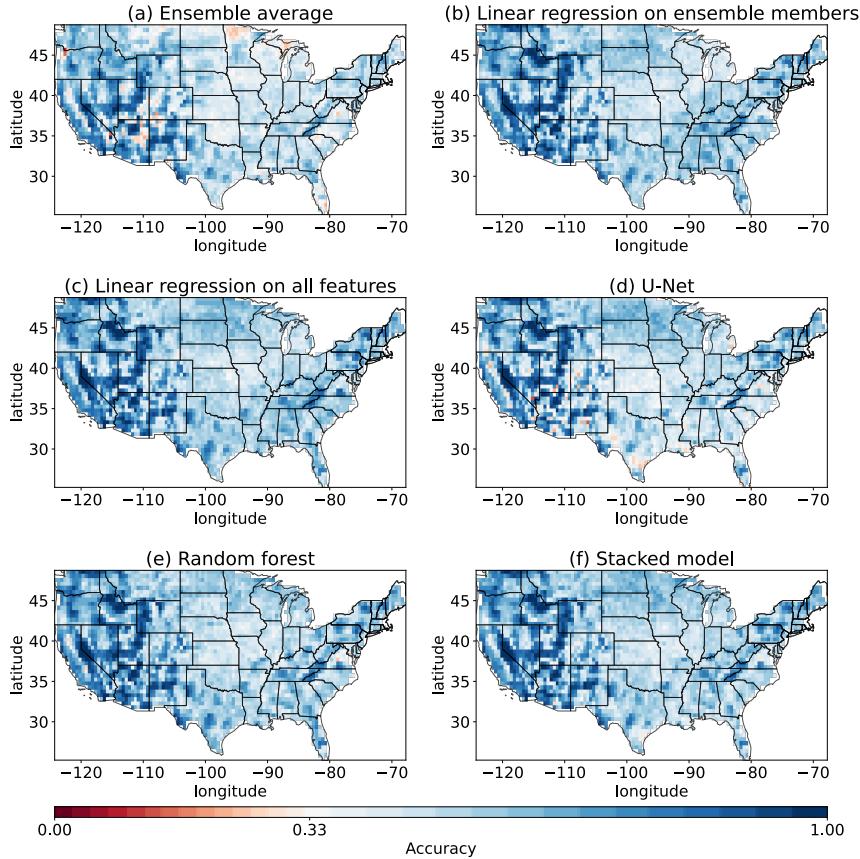
**FIG. B2. Test accuracy heatmaps of baselines and learning-based methods for tercile classification of precipitation using NASA-GMAO dataset.** The accuracy colorbar is recentered to be white at  $\frac{1}{3}$ , what corresponds to a random guess score. Blue pixels indicate better performance, while red pixels correspond to performance that is worse than a random guess. See Section B1 for details.

TABLE B2. **Test results for tercile classification of temperature on different datasets.** Accuracy in % is reported. Note that for this task, our models are trained for regression and the threshold for predicted values is applied.

Data	Model	Mean accuracy ( $\uparrow$ )	Median accuracy ( $\uparrow$ )
NCEP-CFSv2	Ens mean	$44.84 \pm 0.36$	42.74
	Linear Regr	$57.10 \pm 0.25$	54.69
	LR	$57.34 \pm 0.25$	<b>54.71</b>
	U-Net	$53.80 \pm 0.28$	50.43
	RF	$58.07 \pm 0.27$	54.70
NASA-GMAO	Stacked	<b>58.12 <math>\pm 0.20</math></b>	<b>54.71</b>
	Ens mean	$52.23 \pm 0.25$	49.41
	Linear Regr	$57.75 \pm 0.25$	54.11
	LR	<b>58.97 <math>\pm 0.25</math></b>	<b>55.29</b>
	U-Net	$55.64 \pm 0.27$	51.76
	RF	$58.78 \pm 0.26$	<b>55.29</b>
	Stacked	$58.72 \pm 0.25$	54.12



**FIG. B3. Test accuracy heatmaps of baselines and learning-based methods for tercile classification of temperature using NCEP-CFSv2 dataset.** The accuracy colorbar is recentered to be white at  $\frac{1}{3}$ , what corresponds to a random guess score. Blue pixels indicate better performance, while red pixels correspond to performance that is worse than a random guess. See Section B2 for details.



**FIG. B4. Test accuracy heatmaps of baselines and learning-based methods for tercile classification of temperature using NASA-GMAO dataset.** The accuracy colorbar is recentered to be white at  $\frac{1}{3}$ , what corresponds to a random guess score. Blue pixels indicate better performance, while red pixels correspond to performance that is worse than a random guess. See Section B2 for details.

## APPENDIX C

### Extended Discussions

In this section, we present more detailed results for Section 8a and Section 8d. We also provide additional experiments on the temperature forecasting analysis, and experiments on training set sizes and bootstrap.

#### C1. PE v.s. latitude/longitude values or no location information

In this section, we elaborate on our experiment with different uses of location information. Similar to Section 8d, we train and test our models with three different settings: using no location information, using latitude/longitude values, or using positional encodings. For the stacked model, we first train the LR, RF, and U-Net using these different settings before training the stacked model using the corresponding LR, RF, and U-Net outputs. Table C1 summarizes test performance for precipitation regression with these three settings of using location information. For linear regression, we observe that having latitude/longitude values or adding PE features does not improve its performance. One interesting result is that adding PE for the LR degraded its performance, which may be due to the fact that the PE features are non-linear transformations of latitude and longitude values, which is hard to fit with a linear model.

For the RF, U-Net, and the stacked model, we observe that adding PE improves their performance with significance, i.e. having at least one standard error smaller MSE. For the U-Net, using latitude/longitude values yields worse overall performance compared to using no location information. Figure C1 shows the test  $R^2$  heatmaps for LR, U-Net, RF, and stacked model under these three settings for precipitation regression. We observe that adding PE features not only improves performance for the U-Net and RF, the stacked model's performance also improves from having better predictions from the U-Net and RF.

Table C2 shows the test performances of LR, U-Net, RF and stacked model on temperature regression. Similar to precipitation regression, adding latitude/longitude values or adding PE does not help the LR, but we observe significant performance improvement when adding positional encoding features for the U-Net, RF, and the stacked model. Figure C2 then shows the test  $R^2$  heatmaps for temperature regression. We can see that adding PE to the U-Net, RF, and stacked

TABLE C1. **Precipitation regression test performance comparison of LR, U-Net, RF and stacked model trained using no spatial features, using latitude and longitude values or using PE.** The best results are in bold.

Model	Features	Mean $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )
LR	All + no location info	<b>-0.11</b>	<b>2.29</b> $\pm$ 0.04
	All + lat/lon values	<b>-0.11</b>	<b>2.29</b> $\pm$ 0.04
	All + PE	-0.33	2.71 $\pm$ 0.05
U-Net	All + no location info	-0.16	2.31 $\pm$ 0.04
	All + lat/lon values	-0.28	2.53 $\pm$ 0.05
	All + PE	<b>-0.10</b>	<b>2.18</b> $\pm$ 0.03
RF	All + no location info	-0.18	2.23 $\pm$ 0.04
	All + lat/lon values	-0.16	2.21 $\pm$ 0.04
	All + PE	<b>-0.11</b>	<b>2.17</b> $\pm$ 0.05
Stacked	All + no location info	-0.05	2.13 $\pm$ 0.03
	All + lat/lon values	-0.01	2.21 $\pm$ 0.04
	All + PE	<b>0.02</b>	<b>2.07</b> $\pm$ 0.03

model improves forecast performance, especially in regions like Arizona, New Mexico, and Texas.

TABLE C2. **Temperature regression test performance comparison of LR, U-Net, RF and stacked model trained using no spatial features, using latitude and longitude values or using PE.** The best results are in bold.

Model	Features	Mean $R^2$ ( $\uparrow$ )	Mean Sq Err ( $\downarrow$ )
LR	All + no location info	<b>0.05</b>	<b>3.57</b> $\pm$ 0.03
	All + lat/lon values	<b>0.05</b>	<b>3.57</b> $\pm$ 0.03
	All + PE	<b>0.05</b>	3.57 $\pm$ 0.03
U-Net	All + no location info	-0.35	4.81 $\pm$ 0.04
	All + lat/lon values	-0.21	4.47 $\pm$ 0.03
	All + PE	<b>0.01</b>	<b>3.65</b> $\pm$ 0.02
RF	All + no location info	0.11	3.37 $\pm$ 0.02
	All + lat/lon values	0.14	3.28 $\pm$ 0.02
	All + PE	<b>0.16</b>	<b>3.17</b> $\pm$ 0.02
Stacked	All + no location info	0.12	3.35 $\pm$ 0.02
	All + lat/lon values	0.12	3.33 $\pm$ 0.02
	All + PE	<b>0.18</b>	<b>3.11</b> $\pm$ 0.02

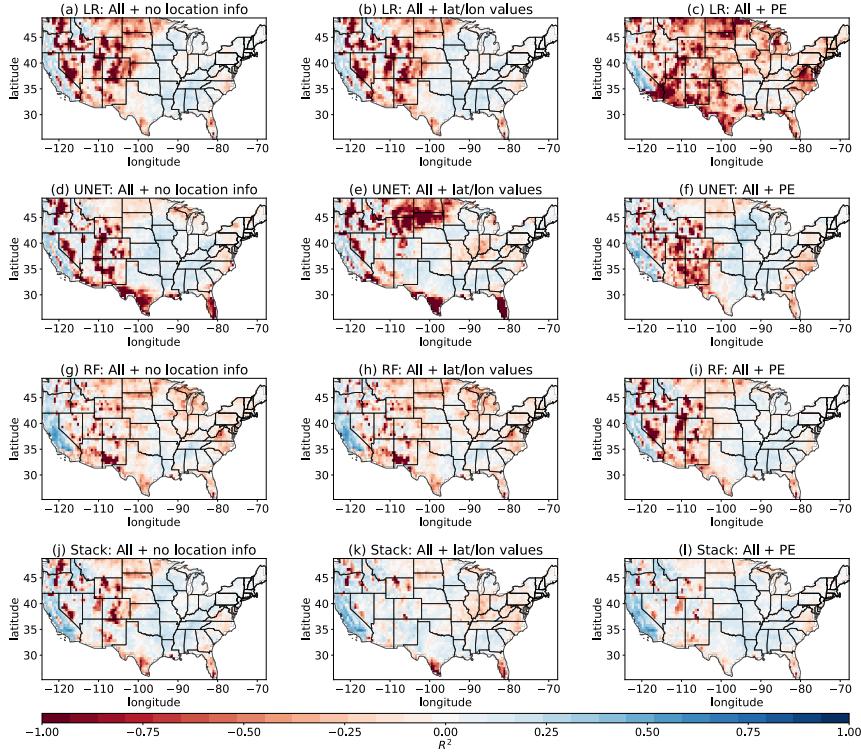


FIG. C1. **Precipitation regression test  $R^2$  heatmaps of LR, U-Net, RF and stacked model trained using no spatial features, using latitude and longitude values or using PE.** The NCEP-CFSv2 ensemble is used. See Appendix C1 for more details.

## C2. Bootstrap experiments

To evaluate the stability of our machine learning models with small sample sizes, we perform the following bootstrap experiments: We take bootstrap samples of size 200 from our training set and retrain our U-Net, RF, and LR. Then we evaluate these models on the test set. We repeat this process 50 times and show the results in Fig. C3. We observe from the plots that the U-Net performs consistently better than the LR in precipitation regression but not for temperature regression. This result is consistent with what we showed in Table 3 and Table 4.

We also observe that the U-Net is more sensitive to different bootstrap samples than the RF and LR, which is not surprising since for the U-Net, the bootstrap samples correspond to 200 different

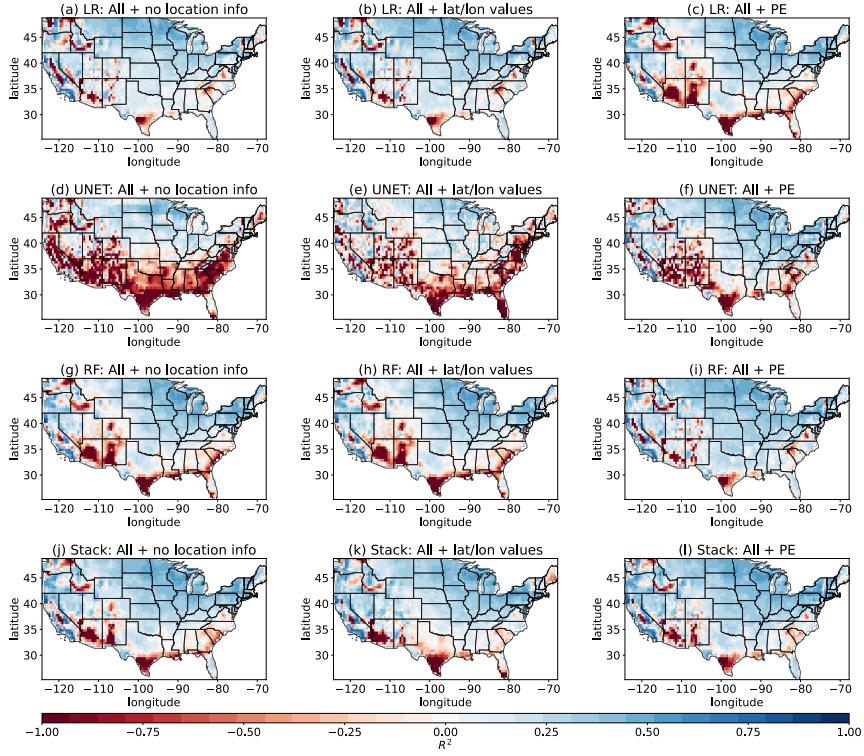


FIG. C2. **Temperature regression test  $R^2$  heatmaps of LR, U-Net, RF and stacked model trained using no spatial features, using latitude and longitude values or using PE.** The NCEP-CFSv2 ensemble is used. See Appendix C1 for more details.

spatial maps for training. In contrast, for the RF and LR, the bootstrap samples correspond to  $200 \cdot 3274$  training samples.

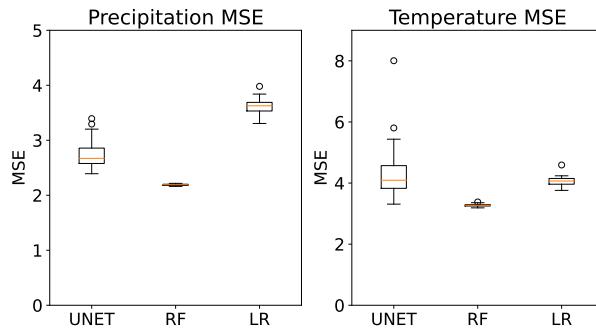


FIG. C3. **Box plots of MSEs for the U-Net, LR and RF trained on 50 set of different bootstrap samples, each with size 200.** The NCEP-CFSv2 ensemble is used. See Appendix C2 for more details.

### C3. Precipitation forecast example

While climate simulations and ensemble forecasts are designed to provide useful predictions of temperature and precipitation based on carefully developed physical models, we see that machine learning applied to those ensembles can yield a significantly higher predictive skill for a range of SSF tasks. Figure C4 illustrates key differences between different predictive models for predicting monthly precipitation with a lead time of 14 days. Individual ensemble members are predictions with high levels of spatial smoothness and more extreme values. Linear regression, the random forest, the U-Net, and the stacked model produce higher spatial frequencies. The linear regression result, which uses a different model trained for each spatial location separately, has the least spatial smoothness of all methods; this is especially visible in the southeast and potentially does not reflect realistic spatial structure. The learning-based models more accurately predict localized regions of high and low precipitation compared to the ensemble mean.

Figure C5 demonstrates differences between the ground truth and different model predictions. In this figure, the color white is associated with the smallest errors, while red pixels indicate overestimating precipitation and blue pixels indicate underestimating precipitation. The individual ensemble member in Figure C5(e) exhibits dark red regions across the West, while the ensemble mean in Figure C5(e) shows better performance in this area. The colors are more muted for the stacked model in Figure C5(h). The climatology in Figure C5(a) has the most neutral areas. However, its MSE is slightly higher than the stacked model's MSE. In general, all methods, including linear regression (b, d), U-Net (f), and random forest (g), tend to underpredict precipitation in the Southeast, Mid-Atlantic, and North Atlantic and predict higher precipitation levels in the West.

### C4. Temperature forecasting analysis

Figure 3 shows regions in Texas and Florida where the ensemble mean and linear regression performance is poor, while a random forest achieves far superior performance. We conduct an analysis of forecasts of the ensemble mean, linear regression, and random forests in these regions together with a region in Wisconsin where all methods show good performance. Figure C6 indicates these regions and Table C3 summarizes the performance of different methods in these regions: the ensemble mean prediction quality dramatically drops between the validation and test periods in Texas and Florida, which is not the case for the random forest.

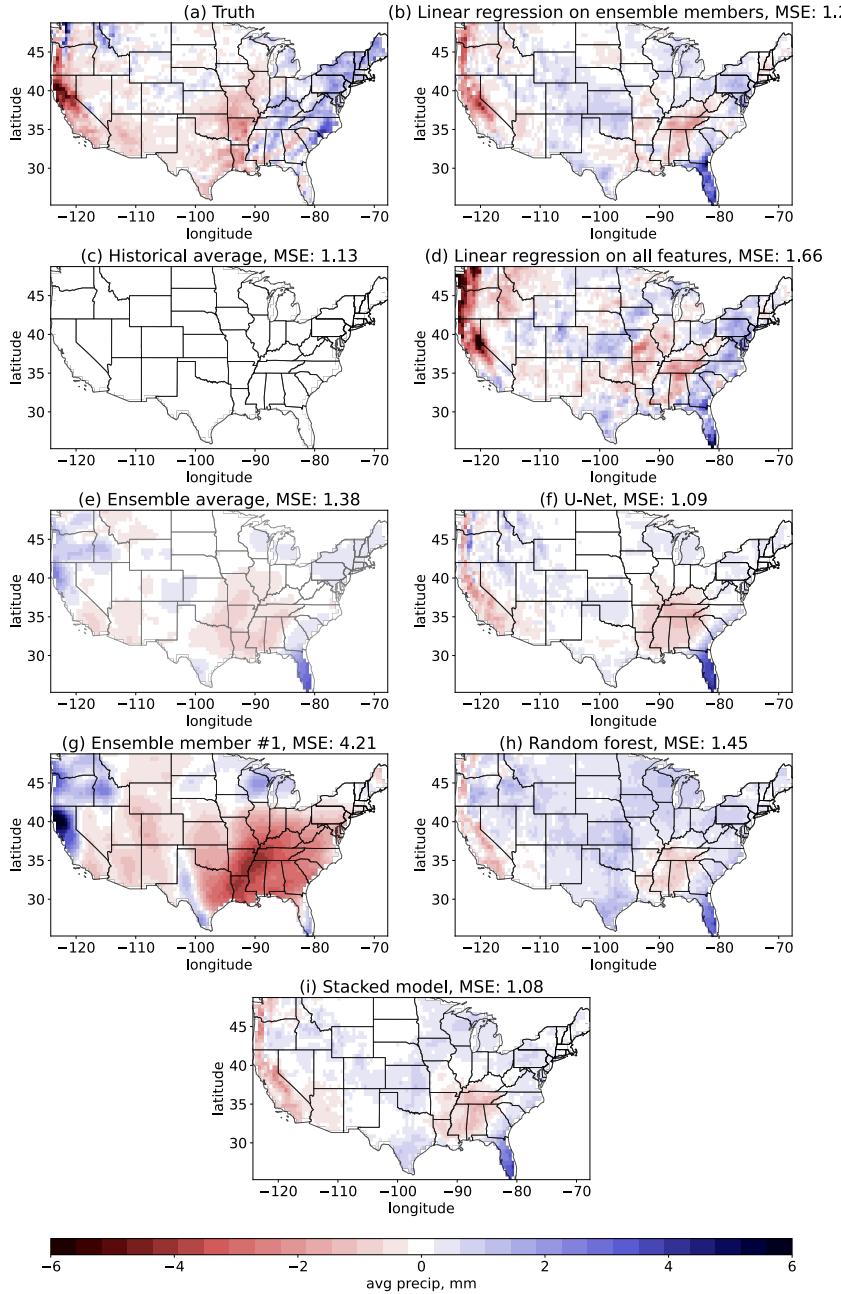


FIG. C4. An illustration of precipitation predictions  $\hat{y}_{t,l}^{\text{anomaly}}$  (in mm) of different methods for February 2016 (in test period). (a) True precipitation. (b) LR on ensemble members. (c) Climatology. (d) LR on all features. (e) Ensemble mean. (f) U-Net on all features. (g) Example single ensemble member. (h) Random forest on all features. (i) Stacked model. See Section C3 for details.

Why does RF perform so much better than simpler methods in some regions? One possibility is that the RF is a nonlinear model capable of more complex predictions. However, if that were the

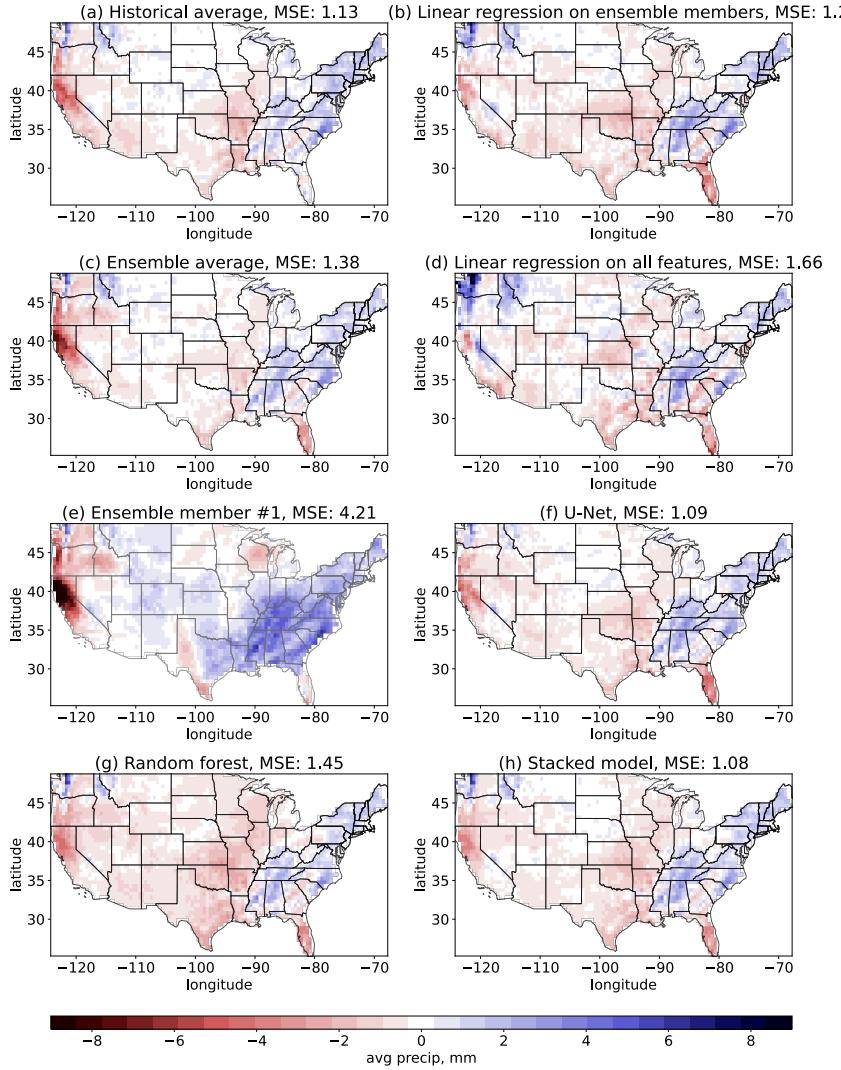


FIG. C5. An illustration of differences  $y_{t,l}^{\text{anomaly}} - \hat{y}_{t,l}^{\text{anomaly}}$  in precipitation predictions in mm of different methods for February 2016 (in test period). Red pixels indicate areas where a forecasting method predicts higher precipitation levels compared to the ground truth, blue pixels indicate an underestimation of the precipitation, and white pixels correspond to a precise prediction. See Section C3 for details.

only cause of the discrepancy in performance, then we would expect that the RF would be better not only during the test period, but during the validation period as well. Table C3 does not support this argument; it shows that the ensemble mean and linear regression have comparable, if not superior, performance to the random forest during the validation period. A second hypothesis is that the distribution of temperature is different during the test period than during the training and validation periods. This hypothesis is plausible for two reasons: (1) climate change, and (2) the training and

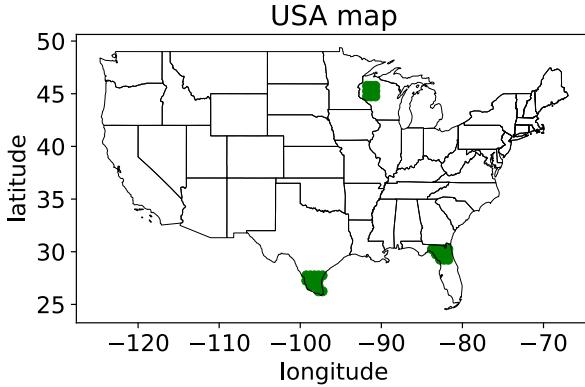


FIG. C6. **Regions where the temperature forecast is analyzed.** See Section C4 for details.

TABLE C3. **Train, validation, and test performance of different methods in Texas, Florida, and Wisconsin regions.** The task is temperature regression; NCEP-CFSv2 dataset is used. The performance of the ensemble mean and linear regression in the test period significantly decreases in Texas and Florida while the random forest is able to demonstrate reasonable results. All methods perform well in Wisconsin.

Region location	Model	Train mean $R^2$ ( $\uparrow$ )	Validation mean $R^2$ ( $\uparrow$ )	Test mean $R^2$ ( $\uparrow$ )
Texas	Ens mean	0.19	0.36	-1.55
	LR	0.53	0.49	-1.29
	RF	0.97	0.32	-0.33
Florida	Ens mean	0.11	0.34	-0.87
	LR	0.47	0.58	-0.56
	RF	0.97	0.36	0.11
Wisconsin	Ens mean	0.30	0.36	0.39
	LR	0.53	0.57	0.51
	RF	1.00	0.47	0.47

validation data use hindcast ensembles while the test data uses forecast ensembles. To investigate this hypothesis, in Figure C7 we plot the true temperature and ensemble mean in the training, validation, and test periods for the three geographic regions. The discrepancy between the true temperatures and ensemble means in the test period is generally greater than during the training and validation periods in Texas and Florida (though not in Wisconsin, a region where validation and test performance are comparable for all methods). This lends support to the hypothesis that hindcast and forecast ensembles exhibit distribution drift, and the superior performance of the RF during the test period may be due to a greater robustness to that distribution drift.

The hindcast and forecast ensembles may have different predictive accuracies because the hindcast ensembles have been debiased to fit past observations – a procedure not possible for forecast data. To explore the potential impact of debiasing, Figure C7 shows the “oracle debiased ensemble mean”, which is computed by using the test data to estimate the forecast ensemble bias and subtracting it from the ensemble mean. This procedure, *which would not be possible in practice and is used only to probe distribution drift ensemble bias*, yields smaller discrepancies between the true data and the (oracle debiased) ensemble mean than the discrepancies between the true data and the original (biased) ensemble mean. Specifically, the oracle ensemble member achieves -0.20 mean  $R^2$  score (TX) and -0.28 mean  $R^2$  score (FL) vs. -1.55  $R^2$  (TX) and -0.87  $R^2$  (FL) of the original forecast ensemble mean. The errors during the test period are generally larger than during the train and validation period, even after debiasing the ensemble members using future data. This effect may be attributed both to (a) the nonstationarity of the climate (note that there are more extreme values during the test period than during the training and validation periods, particularly in Texas and Florida) and (b) the fact that in the train and validation periods, we use hindcast ensemble members, whereas in the test period, we use forecast ensemble members.

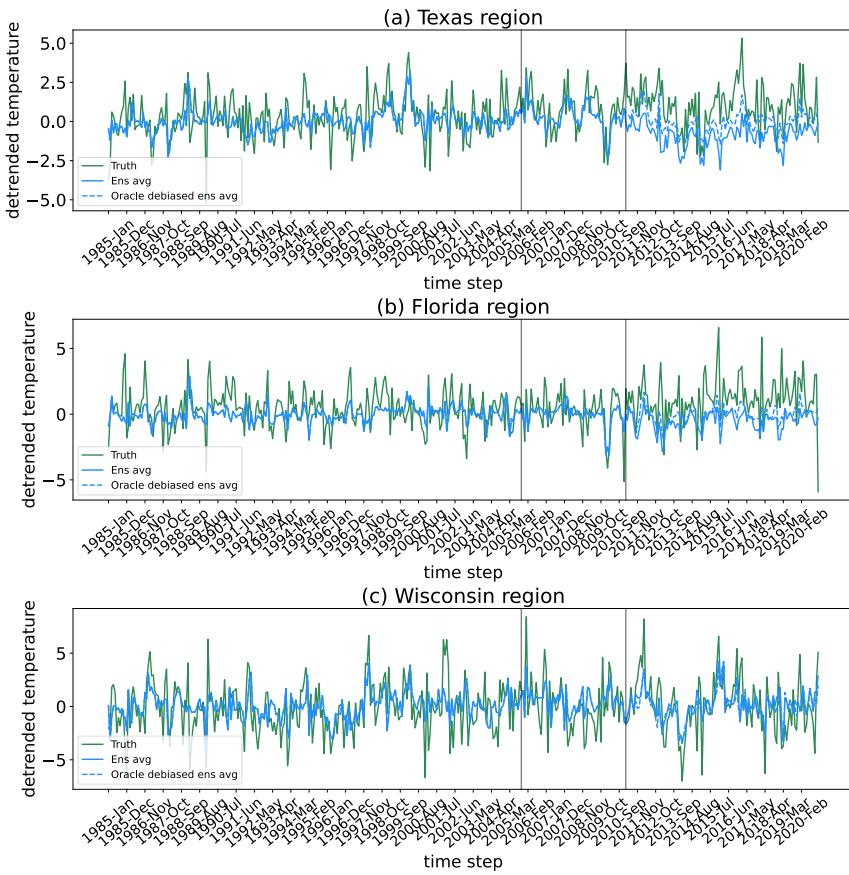


FIG. C7. Temperature predictions in  $^{\circ}\text{C}$  of different methods at Texas, Florida, and Wisconsin regions.

Black lines correspond to train/val and val/test splits; train and validation correspond to the hindcast regime of the ensemble, while test corresponds to the forecast regime. See Section C4 for details.

## APPENDIX D

### Architecture Details

#### D1. Machine learning architectures

##### a. U-Net details

The U-Net has residual connections from layers in the encoder part to the decoder part in a paired way so that it forms a U-shape. Figure D1 shows the architecture of the U-Net. The U-Net is a powerful deep convolutional network that is widely used in image processing tasks such as image segmentation (Ronneberger et al. 2015; Hao et al. 2020) or style transfer (Gatys et al. 2016; Jing et al. 2019).

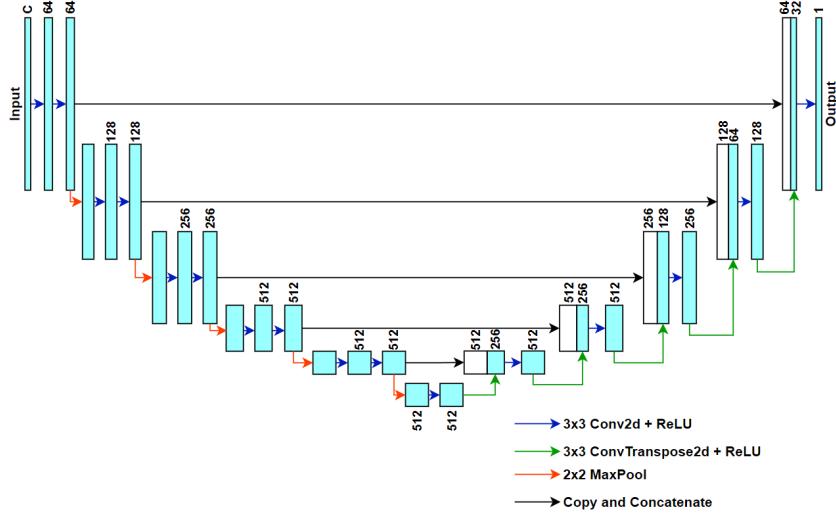


FIG. D1. U-Net architecture with input channels =  $C$ .  $C$  is the number of input channels, which, in our case, equals the number of ensemble members plus all climate data.

Our U-Net differs from the original U-Net by modifying the first 2D convolutional layer after input. Since our input channels can be different when we choose a different subset of features or different ensemble (NCEP-CFSv2 or NASA-GMAO), this 2D convolutional layer is used to transform our input with  $C$  channels into a latent representation with 64 channels. The number of channels  $C$  depends on which ensemble we are using and what task we are performing. For example, for precipitation tasks using the NCEP-CFSv2 ensemble, the input channels include 24

ensemble members, 5 lagged observations, 4 other observational variables, 8 principal components of SSTs, and 24 positional encodings, resulting in 65 channels in total. For temperature tasks using the NASA-GMAO ensemble, there are only 11 ensemble members, and we don't include SSTs information, hence there are only 44 channels in total. The other following layers use the same configurations with the standard U-Net Ronneberger et al. (2015).

We also perform careful hyperparameter tuning for the U-NET. In particular, we run a 10-fold cross-validation on our training set, and use grid search for tuning learning rate, batch size, number of epochs, and weight decay. Since we use different loss functions for different forecast tasks and different numbers of input channels for NCEP-CFSv2 and GMAO-GMAO ensemble, we run hyperparameter tuning with the same cross-validation scheme separately for these tasks. For instance, for precipitation regression, we choose from 100, 120, 150, 170, 200, and 250 epochs; batch size may be equal to 8, 16, 32; learning rate values are chosen from 0.0001, 0.001, 0.01; weight decay can be 0, 1e-3, 1e-4. In case of NCEP-CFSv2 precipitation regression, the optimal parameters are 170 epochs, batch size 16, learning rate 0.0001, and weight decay 1e-4. For temperature regression using the same data, the best parameters are 100 epochs, batch size 16, learning rate 0.001, and weight decay 1e-3. For tercile classification of precipitation, the best parameters are 80 epochs (we chose from 60, 70, 80, 90, and 100 epochs during classification), batch size 8, learning rate 0.001, and weight decay 1e-4.

#### *b. Random Forest Quantile Regressor details*

We show a figure representation of the RFQR in Fig. D2. The RFQR is essentially trained as a regular random forest, but it makes a quantile estimate by taking the sample quantile of the responses in all leaves associated with a new input.

#### *c. Stacking model details*

The stacking model is a simple one-layer neural network with 100 hidden neurons and a sigmoid activation function for regression and softmax for classification. We use an implementation from Scikit-learn library (Pedregosa et al. 2011). We choose 100 neurons based on the stacking model performance on the validation data (we also try 50, 75, 100 and 120 neurons). The stacking model demonstrates stable performance in general, but with 100 neurons it usually achieves the best

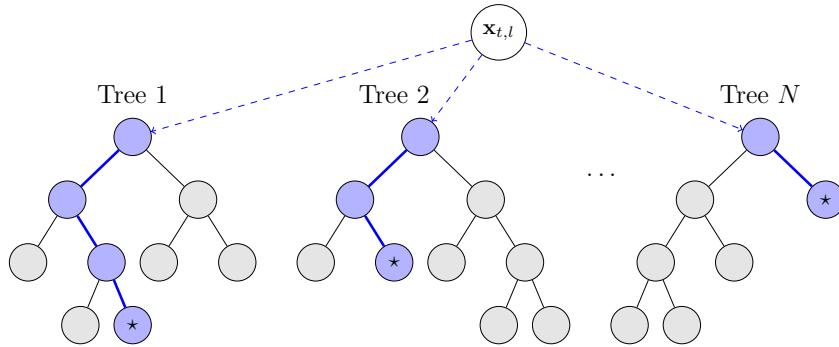


Fig. D2. Illustration of a random forest, which serves as a visual aid for our discussion. **Quantile regression forests** outputs the empirical  $\alpha$  quantiles of the collection of all responses in all the leaves associated with  $\mathbf{x}_{t,l}$  (marked with a star for each tree).

results. We use the “lbfgs” optimizer from quasi-Newton methods for the regression tasks, and the SGD optimizer for classification tasks.

## APPENDIX E

### Additional Preprocessing Details

Random forest and U-Net require different input formats. For U-Net, all input variables have natural image representation except SSTs and information about location. For example, ensemble predictions can be represented as a tensor of shape  $(K, W, H)$ , where  $K$  corresponds to the number of ensemble members (or number of channels of an image), and  $W$  and  $H$  are width and height of the corresponding image. In our case,  $W = 64$  and  $H = 128$ . For the U-Net model, we handle the missing land variables over the sea regions by the nearest neighbor interpolation of available values.

(i) *Sea surface temperatures* There are more than 100,000 SSTs locations available. We extract the top eight principal components. Principal component analysis fits on the train part and then is applied to the rest of the data. In the case of U-Net, we deal with PCs of SSTs by adding additional input channels that are constant across space, with each channel corresponding to one of PCs. Random forest can use PCs from SSTs directly with no special preprocessing.

(ii) *Normalization* We apply channel-wise min-max normalization to the input features at each location based on the training part of the dataset in the case of U-Net. As for normalization of the true values, min-max normalization is applied for precipitation, and standardization is applied for temperature. This choice affects the final layer of the U-Net model, too: for the precipitation regression task, a sigmoid activation is used, and no activation function is applied for temperature regression. For the stacking model, we apply min-max normalization to both input and target values.

## References

Badr, H. S., B. F. Zaitchik, and S. D. Guikema, 2014: Application of statistical models to the prediction of seasonal rainfall anomalies over the sahel. *Journal of Applied meteorology and climatology*, **53** (3), 614–636.

Ban, R. J., and Coauthors, 2016: *Next Generation Earth System Prediction: Strategies for Subseasonal to Seasonal Forecasts*. The National Academies Press, 276 pp.

Barnston, A. G., M. K. Tippett, M. L. L'Heureux, S. Li, and D. G. DeWitt, 2012: Skill of real-time seasonal enso model predictions during 2002–11: Is our capability increasing? *Bulletin of the American Meteorological Society*, **93** (5), 631–651.

Bi, K., L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian, 2022: Pangu-Weather: A 3d high-resolution model for fast and accurate global weather forecast. *arXiv preprint arXiv:2211.02556*.

Biau, G., and E. Scornet, 2016: A random forest guided tour. *Test*, **25**, 197–227.

Brunet, G., and Coauthors, 2010: Collaboration of the weather and climate communities to advance subseasonal-to-seasonal prediction. *Bulletin of the American Meteorological Society*, **91** (10), 1397–1406.

Carion, N., F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, 2020: End-to-end object detection with transformers. *European conference on computer vision*, Springer, 213–229.

Cash, B. A., J. V. Manganello, and J. L. Kinter, 2019: Evaluation of NMME temperature and precipitation bias and forecast skill for South Asia. *Climate dynamics*, **53**, 7363–7380.

Chen, H., and Coauthors, 2021: Pre-trained image processing transformer. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12 299–12 310.

Chen, L., B. Han, X. Wang, J. Zhao, W. Yang, and Z. Yang, 2023: Machine learning methods in weather and climate applications: A survey. *Applied Sciences*, **13** (21), 12 019.

Cofino, A. S., R. Cano, C. Sordo, and J. M. Gutierrez, 2002: Bayesian networks for probabilistic weather prediction. *15th European Conference on Artificial Intelligence (ECAI)*.

Cohen, J., D. Coumou, J. Hwang, L. Mackey, P. Orenstein, S. Totz, and E. Tziperman, 2019: S2s reboot: An argument for greater inclusion of machine learning in subseasonal to seasonal forecasts. *Wiley Interdisciplinary Reviews: Climate Change*, **10** (2), e00567.

DelSole, T., and M. K. Tippett, 2014: Comparing forecast skill. *Monthly Weather Review*, **142** (12), 4658–4678.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova, 2018: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv, 1810.04805.

Dosovitskiy, A., and Coauthors, 2020: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv, 2010.11929.

Fan, Y., and H. Van den Dool, 2008: A global monthly land surface air temperature analysis for 1948–present. *Journal of Geophysical Research: Atmospheres*, **113** (D1).

Frndá, J., M. Durica, J. Rozhon, M. Vojtekova, J. Nedoma, and R. Martinek, 2022: ECMWF short-term prediction accuracy improvement by deep learning. *Scientific Reports*, **12** (1), 7898.

Gamboa, J. C. B., 2017: Deep learning for time-series analysis. arXiv, 1701.01887.

Gatys, L. A., A. S. Ecker, and M. Bethge, 2016: Image style transfer using convolutional neural networks. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2414–2423.

Gehring, J., M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin, 2017: Convolutional sequence to sequence learning. *International conference on machine learning*, PMLR, 1243–1252.

Ghaderi, A., B. M. Sanandaji, and F. Ghaderi, 2017: Deep forecast: Deep learning-based spatio-temporal forecasting. arXiv, 1707.08110.

Grönquist, P., C. Yao, T. Ben-Nun, N. Dryden, P. Dueben, S. Li, and T. Hoefler, 2020: Deep learning for post-processing ensemble weather forecasts. *CoRR*, **abs/2005.08748**, URL <https://arxiv.org/abs/2005.08748>, 2005.08748.

Grover, A., A. Kapoor, and E. Horvitz, 2015: A deep hybrid model for weather forecasting. *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 379–386.

Hao, S., Y. Zhou, and Y. Guo, 2020: A brief survey on semantic segmentation with deep learning. *Neurocomputing*, **406**, 302–321.

Hastie, T., R. Tibshirani, J. H. Friedman, and J. H. Friedman, 2009: *The elements of statistical learning: data mining, inference, and prediction*, Vol. 2. Springer.

Haupt, S. E., W. Chapman, S. V. Adams, C. Kirkwood, J. S. Hosking, N. H. Robinson, S. Lerch, and A. C. Subramanian, 2021: Towards implementing artificial intelligence post-processing in weather and climate: proposed actions from the Oxford 2019 workshop. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **379 (2194)**, 20200091.

He, S., X. Li, T. DelSole, P. Ravikumar, and A. Banerjee, 2020: Sub-seasonal climate forecasting via machine learning: Challenges, analysis, and advances. arXiv, 2006.07972.

He, S., X. Li, L. Trenary, B. A. Cash, T. DelSole, and A. Banerjee, 2021: Learning and dynamical models for sub-seasonal climate forecasting: Comparison and collaboration. arXiv, 2110.05196.

Herman, G. R., and R. S. Schumacher, 2018: “Dendrology” in numerical weather prediction: What random forests and logistic regression tell us about forecasting extreme precipitation. *Monthly Weather Review*, **146 (6)**, 1785–1812.

Hewage, P., M. Trovati, E. Pereira, and A. Behera, 2021: Deep learning-based effective fine-grained weather forecasting model. *Pattern Analysis and Applications*, **24 (1)**, 343–366.

Hwang, J., P. Orenstein, J. Cohen, K. Pfeiffer, and L. Mackey, 2019: Improving subseasonal forecasting in the western us with machine learning. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2325–2335.

Hwang, J., P. Orenstein, K. Pfeiffer, J. Cohen, and L. Mackey, 2018: Improving subseasonal forecasting in the western us with machine learning. arXiv, 1809.07394.

Iglesias, G., D. C. Kale, and Y. Liu, 2015: An examination of deep learning for extreme climate pattern analysis. *The 5th International Workshop on Climate Informatics*.

Jing, Y., Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song, 2019: Neural style transfer: A review. *IEEE transactions on visualization and computer graphics*, **26 (11)**, 3365–3385.

Kalnay, E., and Coauthors, 1996: The NCEP/NCAR 40-year reanalysis project. *Bulletin of the American meteorological Society*, **77** (3), 437–472.

Khan, S., M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, 2022: Transformers in vision: A survey. *ACM computing surveys (CSUR)*, **54** (10s), 1–41.

Kingma, D. P., and J. Ba, 2014: Adam: A method for stochastic optimization. arXiv, 1412.6980.

Kirtman, B. P., and Coauthors, 2014: The North American multimodel ensemble: phase-1 seasonal-to-interannual prediction; phase-2 toward developing intraseasonal prediction. *Bulletin of the American Meteorological Society*, **95** (4), 585–601.

Knapp, K. R., and Coauthors, 2011: Globally gridded satellite observations for climate studies. *Bulletin of the American Meteorological Society*, **92** (7), 893–907.

Lam, R., and Coauthors, 2023: Learning skillful medium-range global weather forecasting. *Science*, eadi2336.

Loken, E. D., A. J. Clark, and A. McGovern, 2022: Comparing and interpreting differently designed random forests for next-day severe weather hazard prediction. *Weather and Forecasting*, **37** (6), 871 – 899, [https://doi.org/https://doi.org/10.1175/WAF-D-21-0138.1](https://doi.org/10.1175/WAF-D-21-0138.1), URL <https://journals.ametsoc.org/view/journals/wefo/37/6/WAF-D-21-0138.1.xml>.

Lorenc, A. C., 1986: Analysis methods for numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society*, **112** (474), 1177–1194.

Mamalakis, A., J.-Y. Yu, J. T. Randerson, A. AghaKouchak, and E. Foufoula-Georgiou, 2018: A new interhemispheric teleconnection increases predictability of winter precipitation in southwestern us. *Nature communications*, **9** (1), 2332.

Meinshausen, N., 2006: Quantile regression forests. *Journal of Machine Learning Research*, **7** (35), 983–999, URL <http://jmlr.org/papers/v7/meinshausen06a.html>.

Min, Y.-M., S. Ham, J.-H. Yoo, and S.-H. Han, 2020: Recent progress and future prospects of subseasonal and seasonal climate predictions. *Bulletin of the American Meteorological Society*, **101** (5), E640–E644.

Mouatadid, S., P. Orenstein, G. Flasphohler, J. Cohen, M. Oprescu, E. Fraenkel, and L. Mackey, 2023: Adaptive bias correction for improved subseasonal forecasting. *Nature Communications*, **14** (1), 3482.

Nagaraj, R., and L. S. Kumar, 2023: Univariate deep learning models for prediction of daily average temperature and relative humidity: The case study of chennai, india. *Journal of Earth System Science*, **132** (3), 100.

Nakada, K., R. M. Kovach, J. Marshak, and A. Molod, 2018: Global modeling and assimilation office - NASA. NASA, URL <https://gmao.gsfc.nasa.gov/pubs/docs/Nakada1033.pdf>.

Narayanan, A., M. Chandramohan, L. Chen, Y. Liu, and S. Saminathan, 2016: subgraph2vec: Learning distributed representations of rooted sub-graphs from large graphs. *arXiv*, 1606.08928.

National Academies of Sciences, 2016: *Next generation earth system prediction: strategies for subseasonal to seasonal forecasts*. National Academies Press.

National Research Council, 2010: *Assessment of intraseasonal to interannual climate prediction and predictability*. National Academies Press.

Nebeker, F., 1995: *Calculating the weather: Meteorology in the 20th century*. Elsevier.

NOAA, 2022: NOAA National Centers for Environmental information, Climate at a Glance: National Time Series. NOAA, <https://www.ncdc.noaa.gov/cag/>.

Pathak, J., and Coauthors, 2022: FourCastNet: A global data-driven high-resolution weather model using adaptive Fourier neural operators. *arXiv preprint arXiv:2202.11214*.

Pavlyshenko, B., 2018: Using stacking approaches for machine learning models. *2018 IEEE Second International Conference on Data Stream Mining & Processing (DSMP)*, IEEE, 255–258.

Pedregosa, F., and Coauthors, 2011: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, **12**, 2825–2830.

Petroni, F., T. Rocktäschel, P. Lewis, A. Bakhtin, Y. Wu, A. H. Miller, and S. Riedel, 2019: Language models as knowledge bases? *arXiv*, 1909.01066.

Price, I., and Coauthors, 2023: GenCast: Diffusion-based ensemble forecasting for medium-range weather. *arXiv preprint arXiv:2312.15796*.

Radhika, Y., and M. Shashi, 2009: Atmospheric temperature prediction using support vector machines. *International journal of computer theory and engineering*, **1** (1), 55.

Reynolds, R. W., T. M. Smith, C. Liu, D. B. Chelton, K. S. Casey, and M. G. Schlax, 2007: Daily high-resolution-blended analyses for sea surface temperature. *Journal of climate*, **20** (22), 5473–5496.

Romano, Y., E. Patterson, and E. J. Candès, 2019: Conformalized quantile regression. *NeurIPS*.

Ronneberger, O., P. Fischer, and T. Brox, 2015: U-Net: Convolutional Networks for Biomedical Image Segmentation. *MICCAI*.

Saha, S., and Coauthors, 2014: The NCEP climate forecast system version 2. *Journal of climate*, **27** (6), 2185–2208.

Seager, R., and Coauthors, 2007: Model projections of an imminent transition to a more arid climate in southwestern north america. *Science*, **316** (5828), 1181–1184.

Simmons, A. J., and A. Hollingsworth, 2002: Some aspects of the improvement in skill of numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society*, **128** (580), 647–677.

Srinivasan, V., J. Khim, A. Banerjee, and P. Ravikumar, 2021: Subseasonal climate prediction in the Western US using bayesian spatial models. *Uncertainty in artificial intelligence*, PMLR, 961–970.

Subbaswamy, A., R. Adams, and S. Saria, 2021: Evaluating model robustness and stability to dataset shift. *International Conference on Artificial Intelligence and Statistics*, PMLR, 2611–2619.

Totz, S., E. Tziperman, D. Coumou, K. Pfeiffer, and J. Cohen, 2017: Winter precipitation forecast in the European and Mediterranean regions using cluster analysis. *Geophysical Research Letters*, **44** (24), 12–418.

Vannitsem, S., and Coauthors, 2021: Statistical postprocessing for weather forecasts: Review, challenges, and avenues in a big data world. *Bulletin of the American Meteorological Society*, **102** (3), E681–E699.

Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, 2017: Attention is all you need. *Advances in neural information processing systems*, **30**.

Vitart, F., A. W. Robertson, and D. L. Anderson, 2012: Subseasonal to seasonal prediction project: Bridging the gap between weather and climate. *Bulletin of the World Meteorological Organization*, **61** (2), 23.

White, C. J., and Coauthors, 2022: Advances in the application and utility of subseasonal-to-seasonal predictions. *Bulletin of the American Meteorological Society*, **103** (6), E1448–E1472.

Wiles, O., S. Gowal, F. Stimberg, S. Alvise-Rebuffi, I. Ktena, T. Cemgil, and Coauthors, 2021: A fine-grained analysis on distribution shift. arXiv, 2110.11328.

Wu, K., H. Peng, M. Chen, J. Fu, and H. Chao, 2021: Rethinking and improving relative position encoding for vision transformer. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10 033–10 041.

Xie, P., M. Chen, and W. Shi, 2010: Cpc global unified gauge-based analysis of daily precipitation. *Preprints, 24th Conf. on Hydrology, Atlanta, GA, Amer. Metero. Soc*, Vol. 2.

Yakubovskiy, P., 2020: Segmentation models pytorch. GitHub, [https://github.com/qubvel/segmentation\\_models.pytorch](https://github.com/qubvel/segmentation_models.pytorch).

Yang, S., F. Ling, Y. Li, and J.-J. Luo, 2023: Improving seasonal prediction of summer precipitation in the middle–lower reaches of the yangtze river using a tu-net deep learning approach. *Artificial Intelligence for the Earth Systems*, **2** (2), 220 078.

Zhu, Q., N. Ponomareva, J. Han, and B. Perozzi, 2021: Shift-robust gnns: Overcoming the limitations of localized graph training data. *Advances in Neural Information Processing Systems*, **34**, 27 965–27 977.