

ARTICLE

A comparative study of non-deep learning, deep learning, and ensemble learning methods for sunspot number prediction

Yuchen Dang^a, Ziqi Chen^b, Heng Li^c and Hai Shu^{d,*}^aCenter for Data Science, New York University, New York, NY, USA^bKey Laboratory of Advanced Theory and Application in Statistics and Data Science-MOE, School of Statistics, East China Normal University, Shanghai, China^cDepartment of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, China^dDepartment of Biostatistics, New York University, New York, NY, USA

ABSTRACT

Solar activity has significant impacts on human activities and health. One most commonly used measure of solar activity is the sunspot number. This paper compares three important non-deep learning models, four popular deep learning models, and their five ensemble models in forecasting sunspot numbers. Our proposed ensemble model XGBoost-DL, which uses XGBoost as a two-level nonlinear ensemble method to combine the deep learning models, achieves the best forecasting performance among all considered models and the NASA's forecast. Our XGBoost-DL forecasts a peak sunspot number of 133.47 in May 2025 for Solar Cycle 25 and 164.62 in November 2035 for Solar Cycle 26, similar to but later than the NASA's at 137.7 in October 2024 and 161.2 in December 2034.

1. Introduction

Human activities and various events on Earth are strongly intertwined with solar activity (Pulkkinen 2007; Hathaway 2015). An increase in solar activity includes increases in extreme ultraviolet and X-ray emissions from the Sun toward Earth, resulting in the atmospheric heating that can be harmful to spacecrafts, satellites and radars. Increased solar flares and coronal mass ejections due to high solar activity can damage the communication and power systems on Earth. The approximately 11-year cyclic pattern of solar activity seems easily predictable, but the cycle varies in both amplitude and duration. Accurate prediction of solar activity is thus of great interest to estimate the expected impact of space weather on space missions and societal technologies.

Solar cycles are also considered to impact many aspects of human health. Juckett and Rosenberg (1993) observed longer human longevities during solar cycle minimums. Davis Jr and Lowell (2004) reported higher incidences of mental illness in chaotic solar cycles. Azcárate, Mendoza, and Levi (2016) found larger fluctuations in blood pressure during ascending phases of solar cycles. Qu (2016) concluded that most influenza pandemics occurred within one year of solar cycle peaks. Predictions of solar activity can hence assist people in taking the necessary precautions.

The number of the sunspots, which appear as dark areas on the solar disk, is the

*CONTACT Hai Shu. Email: hs120@nyu.edu

most commonly used measure of solar activity (Usoskin 2017). For one thing, the sunspot number is one directly visible characteristic of the Sun. For another, there is a publicly available record of sunspot numbers that can be traced back to as early as year 1749.

Prediction of sunspot numbers belongs to the scope of time-series forecasting that may be tackled by either non-deep learning or deep learning methods. Examples of popular non-deep learning methods in general time series forecasting are Auto-Regressive Moving Average (ARMA) models (Box and Jenkins 1976), Exponential Smoothing (Holt 1957; Winters 1960), and the Prophet model recently developed by Facebook (Taylor and Letham 2018). Deep-learning time-series forecasting methods, mainly prevailing in natural language processing, include Long Short-Term Memory (LSTM; Hochreiter and Schmidhuber 1997), Gated Recurrent Unit (GRU; Cho et al. 2014; Chung et al. 2014), Transformer (Vaswani et al. 2017), and the recent Informer (Zhou et al. 2021). Non-deep learning methods often have restrictive theoretical assumptions that limit their performance on real-world time series data (Lara-Benítez, Carranza-García, and Riquelme 2021). Deep learning methods are generally superior over non-deep learning methods due to the capability of extracting complex data representations at high levels of abstraction (Han et al. 2019).

Although there have been many studies on predicting the sunspot number by using non-deep learning (Xu et al. 2008; Hiremath 2008; Chattopadhyay, Jhajharia, and Chattopadhyay 2011; Tabassum, Rabbani, and Omar 2020) or deep learning forecasting methods (Pala and Atici 2019; Benson et al. 2020; Arfianti et al. 2021; Prasad et al. 2022), most are based on ARMA models or deep learning methods like LSTM or GRU. Little work has been done on the more recent time-series models Prophet, Transformer, and Informer for the sunspot number prediction. Moreover, the ensemble learning methods (Kuncheva 2014) are rarely considered in this forecasting problem.

Ensemble learning is widely used in machine learning to boost the performance by combining results from multiple models. Since the ensemble result is often better than that of a single model, this technique has been increasingly applied to time series forecasting (Wichard and Ogorzalek 2004; Qiu et al. 2014; Oliveira and Torgo 2015; Kaushik et al. 2020). Basic ensemble methods simply use the average or median of predictions from all base models. More sophisticated methods assign different weights to base models such as the error-based method and linear regression (Adhikari and Agrawal 2014).

Due to the regression nature of ensemble learning, any machine learning algorithm for regression can be used as an ensemble method. We thus propose to apply one state-of-the-art nonlinear method, Extreme Gradient Boosting (XGBoost; Chen and Guestrin 2016), to ensemble the predictions from multiple time-series forecasting models. XGBoost is a fast and accurate implementation of gradient boosted decision trees, which has recently been dominating applied machine learning for regression and classification tasks (Zhong et al. 2018; Chang, Chang, and Wu 2018; Pan 2018; Ogunleye and Wang 2019). The popularity of XGBoost mainly stems from its three modifications to traditional gradient boosted decision trees: the approximate greedy algorithm with weighted quantile sketch to fit large dataset, the sparsity-aware split finding algorithm to deal with missing values, and the cache-aware access technique to effectively utilize hardware resources. Besides these advantages, we essentially use XGBoost as a two-level ensemble method, because XGBoost itself is an ensemble of decision trees and each decision tree therein nonlinearly combines the forecasting models. The two-level nonlinear ensemble nature of our usage of XGBoost makes it potentially more powerful than single-level ensemble methods such as those aforementioned.

In this paper, we conduct a comparative study of non-deep learning and deep learning models as well as their ensemble models for the sunspot number prediction. We compare the three non-deep learning models, Seasonal Autoregressive Integrated Moving Average (SARIMA; Box and Jenkins 1976), Exponential Smoothing, and Prophet, and the four deep learning models, LSTM, GRU, Transformer, and Informer. We also consider their ensemble models from basic ensembles, the error-based method, linear regression, and XGBoost.

The contributions of the paper are summarized below.

- We compare three important non-deep learning models (SARIMA, Exponential Smoothing, and Prophet), four popular deep learning models (LSTM, GRU, Transformer, and Informer), and their five ensemble models (via mean, median, the error-based method, linear regression, and XGBoost) in predicting the sunspot number.
- We propose to use XGBoost as a two-level nonlinear ensemble method to combine the results from time-series forecasting models. Our XGBoost-DL model, which uses XGBoost to ensemble the four deep learning models, has the best performance in comparison with other considered base and ensemble models as well as the prediction from the National Aeronautics and Space Administration (NASA).
- We provide an open-source Python package of the XGBoost-DL model for the sunspot number prediction at https://github.com/yd1008/ts_ensemble_sunspot.
- We use the proposed XGBoost-DL model to forecast the Solar Cycles 25 and 26, and compare the result with the NASA’s prediction.

The rest of this paper is organized as follows. Section 2 introduces the seven aforementioned time-series forecasting methods. Section 3 describes the five ensemble learning methods including the proposed XGBoost-based ensemble. Section 4 compares all considered base and ensemble models as well as NASA’s report for the sunspot number prediction. Section 5 makes concluding remarks.

2. Forecasting Methods

In this section, we introduce seven time-series forecasting methods including the three non-deep learning methods SARIMA, Exponential Smoothing, and Prophet, and the four deep learning methods LSTM, GRU, Transformer, and Informer. We denote a univariate time series by $y = (y_1, y_2, \dots, y_t, \dots)$, where y_t is the observation at time t .

2.1. *Non-deep learning methods*

2.1.1. *SARIMA*

SARIMA (Box and Jenkins 1976), as an ARMA variant, is one most commonly used model in the past decades to forecast trend and seasonal time series. It uses a mix of autoregressive terms, moving average terms, and differencing procedures for both non-seasonal and seasonal components to represent the current value in a time series based

on prior observations. Specifically, an SARIMA(p, d, q)(P, D, Q) $_s$ model is defined by

$$\begin{aligned}\phi(L)\Phi(L^s)(1-L)^d(1-L^s)^D y_t &= c + \theta(L)\Theta(L^s)\epsilon_t \\ \phi(L) &= (1 - \phi_1 L^1 - \phi_2 L^2 - \dots - \phi_p L^p) \\ \Phi(L^s) &= (1 - \Phi_1 L^s - \Phi_2 L^{2s} - \dots - \Phi_P L^{Ps}) \\ \theta(L) &= (1 + \theta_1 L^1 + \theta_2 L^2 + \dots + \theta_q L^q) \\ \Theta(L^s) &= (1 + \Theta_1 L^s + \Theta_2 L^{2s} + \dots + \Theta_Q L^{Qs}),\end{aligned}$$

where $(\phi_i, \theta_i, \Phi_i, \Theta_i)$ and (p, q, P, Q) are the parameters and the orders of the non-seasonal autoregressive, the non-seasonal moving average, the seasonal autoregressive, and the seasonal moving average terms, respectively, ϵ_t is white noise, L is the lag operator, d is the order of non-seasonal differencing, D is the order of seasonal differencing, s is the span of the seasonality, and c is a constant.

2.1.2. Exponential Smoothing

Proposed in 1950s, Exponential Smoothing (Brown 1956; Holt 1957; Winters 1960) remains to be one of the widely used time series forecasting methods. Although there are several types of Exponential Smoothing, we focus on the Holt-Winters Exponential Smoothing method (Holt 1957; Winters 1960), which can model the trend and seasonality of a time series. The Holt-Winters method and SARIMA have shown comparable performances in a number of previous studies (Lidiema 2017; Liu et al. 2020; Rabbani et al. 2021).

Also known as Triple Exponential Smoothing, the Holt-Winters method comprises three smoothing equations for the level l_t , the trend b_t , and the seasonality s_t , respectively. There are two main models of this method, the additive seasonal model

$$\begin{aligned}y_{t+h} &= l_t + hb_t + s_{t-m+h_m^+} \\ l_t &= \alpha(y_t - s_{t-m}) + (1 - \alpha)(l_{t-1} + b_{t-1}) \\ b_t &= \beta(l_t - l_{t-1}) + (1 - \beta)b_{t-1} \\ s_t &= \gamma(y_t - l_{t-1} - b_{t-1}) + (1 - \gamma)s_{t-m},\end{aligned}$$

and the multiplicative seasonal model

$$\begin{aligned}y_{t+h} &= (l_t + hb_t)s_{t-m+h_m^+} \\ l_t &= \alpha \frac{y_t}{s_{t-m}} + (1 - \alpha)(l_{t-1} + b_{t-1}) \\ b_t &= \beta(l_t - l_{t-1}) + (1 - \beta)b_{t-1} \\ s_t &= \gamma \frac{y_t}{l_{t-1} + b_{t-1}} + (1 - \gamma)s_{t-m},\end{aligned}$$

where $\alpha, \beta, \gamma \in [0, 1]$ are smoothing parameters, m is the frequency of the seasonality, and $h_m^+ = [(h - 1) \bmod m] + 1$ for $h \geq 1$.

The two models differ in the nature of the seasonality. The additive model ought to be considered when the seasonal variations are stable over time, while the multiplicative model is used when the seasonal variations are changing proportional to the level of the time series. Due to the variability of the amplitude of sunspot cycles,

following Tabassum, Rabbani, and Omar (2020), we use the multiplicative model for the sunspot number prediction.

2.1.3. Prophet

Prophet (Taylor and Letham 2018) is a more recent time series forecasting algorithm compared to the previous two. Despite some commonalities with SARIMA and Exponential Smoothing, it provides a more intuitive approach to model the trend and seasonality of time series by incorporating more flexibilities in its configuration.

Prophet has three essential components: trend b_t , seasonality s_t , and holiday h_t . The holiday option allows Prophet to adjust the forecast that may be affected by holidays or major events. The full model with a logistic trend term is

$$\begin{aligned}
 y_t &= b_t + s_t + h_t + \epsilon_t \\
 b_t &= \frac{C}{1 + \exp(-k(t - m))} \\
 s_t &= \sum_{n=1}^N \left[\alpha_n \cos\left(\frac{2\pi nt}{P}\right) + \beta_n \sin\left(\frac{2\pi nt}{P}\right) \right] \\
 h_t &= [\mathbf{1}(t \in D_1), \dots, \mathbf{1}(t \in D_L)]\boldsymbol{\kappa},
 \end{aligned}$$

where ϵ_t is the error term, C is the carrying capacity that is the maximum value of the logistic curve, k is the growth rate that controls the steepness of the curve, m is an offset parameter corresponding to the curve’s midpoint, the seasonality s_t is expressed by a standard Fourier series with parameters $\{\alpha_n, \beta_n\}_{n=1}^N$ and a regular period P that the time series has, D_i is the set of dates for holidays, and $\boldsymbol{\kappa} \in \mathbb{R}^L$ is the change in the forecast caused by holidays.

2.2. Deep learning methods

Although non-deep learning methods can handle a wide range of time series forecasting tasks, their theoretical limitations often prevent them from being directly applicable to data or modeling complex non-linearity in data, as well as computational complexity makes them impractical to large datasets. Therefore, deep learning techniques such as Recurrent Neural Networks (RNNs) were introduced (Lara-Benítez, Carranza-García, and Riquelme 2021). RNNs have shown better performance than non-deep learning methods in time series forecasting due to its ability to deal with longer sequences and better capture temporal dependencies (Tsui et al. 1995; Zhang and Man 1998).

2.2.1. LSTM

The LSTM network (Hochreiter and Schmidhuber 1997) is one most popular model in the RNN family. The vanilla RNN suffers from vanishing gradients and is not capable to achieve desirable results for long sequence data (Hochreiter 1998). The LSTM network mitigates this issue rising from long-term dependencies by introducing a gated memory cell architecture, which controls the information flow with three gates: an input gate

i_t , an output gate o_t , and a forget gate f_t . The LSTM cell is formulated as follows:

$$\begin{aligned}
 i_t &= \sigma(W_i x_t + U_i h_{t-1} + b_i) \\
 f_t &= \sigma(W_f x_t + U_f h_{t-1} + b_f) \\
 o_t &= \sigma(W_o x_t + U_o h_{t-1} + b_o) \\
 \tilde{c}_t &= \tanh(W_c x_t + U_c h_{t-1} + b_c) \\
 c_t &= f_t \cdot c_{t-1} + i_t \cdot \tilde{c}_t \\
 h_t &= o_t \cdot \tanh(c_t),
 \end{aligned}
 \tag{1}$$

where x_t is the input which is the hidden state h_t from the previous layer and is y_t for the first hidden layer, $\sigma(\cdot)$ is the sigmoid function, $\tanh(\cdot)$ is the hyperbolic tangent function, h_t is the hidden state, c_t is the state of the memory cell, \tilde{c}_t is the candidate state of the memory cell, and W and U are the weights of the input and recurrent connections as well as b is bias with subscripts i , f , o and c for the input gate, forget gate, output gate, and memory cell, respectively.

In each hidden layer of the LSTM network, a sequence of LSTM cells are aligned side-by-side and input data are sequentially fed into each cell. LSTM's capability of communicating across multiple cells comes from the hidden states and the cell states. The hidden state carries over information from the previous cell to the next and a new hidden state is generated from each LSTM cell, while the cell state selectively stores the past information. The input, output, and forget gates generate a new hidden state and update the the cell state in the following procedure. The input gate determines whether new information will be added to the memory in two steps: i_t uses a sigmoid function to decide which information needs to be updated, and \tilde{c}_t utilizes a hyperbolic tangent function to select new candidate information to be added to the memory. The forget gate decides what information should be discarded from the memory by applying a sigmoid function to the previous hidden state h_{t-1} and current input value x_t . The output of the forget gate ranges between 0 and 1, with 0 indicating complete removal of the previously learnt value and 1 indicating retention of all information. The output gate determines what will be generated as the output. o_t acts like a filter by selecting the relevant information from the memory with a sigmoid function to generate an output value, which is then multiplied with the cell state passing through a hyperbolic tangent function to form the representations in the hidden state. The structure of the LSTM cell is illustrated in Figure 1.




Figure 1. The structure of the LSTM cell.

2.2.2. GRU

The GRU (Cho et al. 2014; Chung et al. 2014) network is another well-known RNN model using a gating mechanism similar to that in LSTM, but it has a simpler cell architecture and is computationally more efficient (Torres et al. 2021; Lara-Benítez, Carranza-García, and Riquelme 2021). The GRU cell only has two gates, an update gate z_t and a reset gate r_t . The update gate decides the amount of previous information to be passed to the next state, helping capture long-term dependencies in the sequence. The reset gate determines how much of the past information to neglect and is responsible to learn short-term dependencies. The GRU cell is formulated by

$$\begin{aligned} z_t &= \sigma(W_z x_t + U_z h_{t-1} + b_z) \\ r_t &= \sigma(W_r x_t + U_r h_{t-1} + b_r) \\ \tilde{h}_t &= \tanh(W_h x_t + U_h (r_t \cdot h_{t-1}) + b_h) \\ h_t &= (1 - z_t) \cdot h_{t-1} + z_t \cdot \tilde{h}_t, \end{aligned}$$

where the notation is similar to that in (1). The functionalities of the gates used in GRU are also similar to those in LSTM. Figure 2 shows the structure of the GRU cell.




Figure 2. The structure of the GRU cell.

2.2.3. Transformer

Since introduced in 2017, Transformer (Vaswani et al. 2017), a solely self-attention based encoder-decoder network, has become the state of the art model in natural language processing, along with a number of its variants (Devlin et al. 2018; Yang et al. 2019; Liu et al. 2019). Recently, Transformer and self-attention based models have gained increasing popularity in time series tasks (Wu et al. 2020; Zhou et al. 2021; Wen et al. 2020; Li et al. 2019). Transformer abandons the recurrent layers of RNNs that process data of the input sequence one after another. It instead uses a self-attention mechanism, which ditches the sequential operations and can access any part of the sequence, to capture global dependencies and enable parallel computation.

The core of Transformer is the scaled dot-product self-attention written as

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V, \quad (2)$$

where Q, K and V are matrices with the i -th rows being the i -th query, key and value vectors, and d_k is the dimension of the key vectors. In a self-attention layer, input data pass through three separate linear layers to form the query, key, and value matrices. Dot products of queries and keys are then calculated. In masked attention layers, a mask of the same size as the dot product matrix, with the upper triangle of the mask having values of $-\infty$ and 0's elsewhere, is added to the dot product to prevent values in a sequence from attending to succeeding ones. The dot product matrix scaled by $\sqrt{d_k}$ is then fed into the softmax function. The attention scores are calculated by the dot product between the output of the softmax function and the value matrix. In multi-head attention layers, multiple self-attention layers are stacked with each layer consisting of different sets of weights. The final attention scores are generated by combining attention scores calculated in parallel from each self-attention layer. To use the sequence-order information, Transformer adopts a positional encoding mechanism based on sine and cosine functions. The encoded values are added to the input data, indicating positional information of each value in the sequence, such that Transformer can distinguish values in one position from another without requiring specific order of the input data. Figure 3 illustrates Transformer’s encoder-decoder architecture.

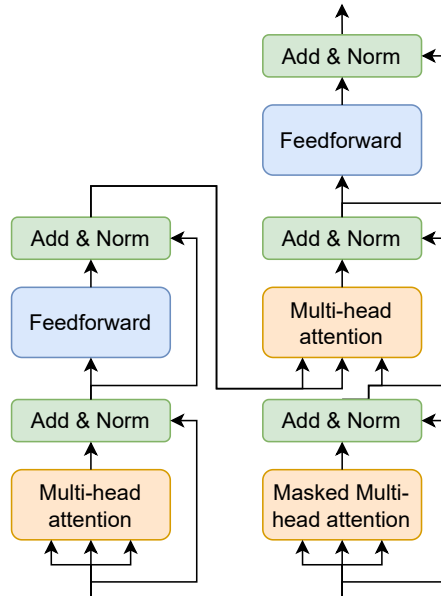


Figure 3. Transformer’s encoder-decoder architecture.

2.2.4. *Informer*

The quadratic computational complexity of dot-product self attention and the heavy memory usage in stacking layers are major concerns in dealing with long input sequences. There have been some attempts to address these issues (Child et al. 2019; Li et al. 2019; Beltagy, Peters, and Cohan 2020; Zhou et al. 2021). We focus on the award-winning method Informer (Zhou et al. 2021).

Informer is proposed as an enhancement of Transformer on long-sequence time series forecasting. Informer adopts a ProbSparse attention mechanism which selects only the

most dominant u queries for each key to attend to:

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

where \overline{Q} only contains the top- u queries chosen under the query sparsity measurement with Kullback-Leibler divergence which is further approximated efficiently with random sampling. The ProbSParse self-attention allows Informer to reduce both time complexity and memory usage from $O(L^2)$ down to $O(L \log L)$ for input length L . In addition, Informer allows longer input sequences by using the self-attention distilling:

$$X_{j+1}^t = \text{MaxPool}(\text{ELU}(\text{Conv1d}([\cdot]_{AB}^t))),$$

where $[\cdot]_{AB}$ is the attention block, X_{j+1}^t is the t -th sequence in the j -th attention layer, $\text{MaxPool}(\cdot)$ is a max-pooling layer with stride 2, $\text{ELU}(\cdot)$ is the ELU activation function, and $\text{Conv1d}(\cdot)$ is a 1-D convolutional filter with kernel size of 3. By distilling, a more condensed feature map is passed from one attention layer to the next, while information is largely preserved. The corresponding memory usage is $O((2 - \epsilon)L \log L)$ instead of $O(JL \log L)$ for J -stacking layers, where ϵ is a small number.

3. Ensemble Learning Methods

Ensemble learning combines the results from multiple models to produce predictions (Dong et al. 2020). The combined results usually exhibit better performance because ensemble methods can often decrease the likelihood of overfitting, reduce the chance of being trapped in a local minimum, and extend the size of the search space (Sagi and Rokach 2018). In time series forecasting tasks, commonly used ensemble methods include basic ensemble methods, error-based method, and machine learning methods. In the following sections, $\hat{Y} = (\hat{y}_{ij}) \in \mathbb{R}^{n \times m}$ denotes the matrix of predictions from m forecasting models at n time points, \hat{y}_{ij} is the predicted value from the j -th model at the i -th time point, and \hat{y}_i is the final prediction at the i -th time point.

3.1. Basic ensemble methods

Basic ensemble methods combine the predictions from all base models by some simple functions such as mean and median:

$$\begin{aligned} \hat{y}_i &= \text{mean}(\hat{y}_{i1}, \dots, \hat{y}_{im}), \\ \hat{y}_i &= \text{median}(\hat{y}_{i1}, \dots, \hat{y}_{im}). \end{aligned}$$

Although these methods are straightforward and simple to comprehend, they ignore the possible relationship among base models and thus are not adequate for combining nonstationary models (Allende and Valle 2017).

3.2. Error-based method

Unlike the mean ensemble method that weighs all base models equally, error-based method assigns weights inversely proportional to forecast errors (Adhikari and Agrawal

2014). Specifically, the data is divided into training and validation sets. A specific evaluation metric is selected to compute the forecast error e_j on the validation set for the j -th base model fitted on the training set, and then the model’s weight is

$$w_j = \frac{e_j^{-1}}{\sum_{j=1}^m e_j^{-1}}.$$

The final prediction of this method is

$$\hat{y}_i = \sum_{j=1}^m w_j \hat{y}_{ij}. \quad (3)$$

The error-based method improves the mean ensemble method by allowing different weights for base models, but its performance is still highly dependent on the forecasting performance of each base model.

3.3. *Linear regression*

The mean, median, and error-based ensemble methods are all linear ensemble methods in the form of (3). The linear ensemble problem can be written as the linear model:

$$y_i = \hat{y}_i + \varepsilon_i = \sum_{j=1}^m w_j \hat{y}_{ij} + \varepsilon_i, \quad (4)$$

where the forecasts of base models $\{\hat{y}_{ij}\}_{j=1}^m$ are the features, and the true observation y_i is the target. Linear regression, or the least squares method, provides the optimal weights that minimize the sum of squared errors $\sum_{i=1}^n \varepsilon_i^2$. Formulate (4) in the matrix form

$$y = \hat{Y}w + \varepsilon.$$

The optimal weight vector given by linear regression is $w = \hat{Y}^\dagger y$ with \hat{Y}^\dagger the Moore-Penrose pseudo-inverse of \hat{Y} (Penrose 1956).

3.4. *XGBoost*

As ensemble learning is essentially a regression problem with features $\{\hat{y}_{ij}\}_{j=1}^m$ and target y_i , any machine learning algorithm for regression is applicable as an ensemble method. In addition to linear regression, there are many nonlinear regression methods such as kernel regression (Li and Racine 2007), support vector regression (Cherkassky and Ma 2004), and tree-based regression algorithms (Breiman et al. 1984; Breiman 2001; Friedman 2001).

We consider the state-of-the-art nonlinear method, XGBoost (Chen and Guestrin 2016), a highly efficient and effective implementation of gradient boosted decision

trees. XGBoost aims to solve the objective function

$$\begin{aligned} \mathcal{L} &= \sum_{i=1}^n l(\hat{y}_i, y_i) + \sum_{k=1}^K \Omega(f_k) \\ \text{with } \hat{y}_i &= \sum_{k=1}^K f_k(\hat{y}_{i1}, \dots, \hat{y}_{im}) \\ \text{and } \Omega(f) &= \gamma T + \frac{1}{2} \lambda \|\omega\|_2^2, \end{aligned} \tag{5}$$

where l is a differentiable convex loss function that measures the difference between the prediction \hat{y}_i and the target y_i , $\{f_k\}_{k=1}^K$ are decision trees, and $\Omega(f)$ is the penalty term with tuning parameters γ and λ to control the model complexity in which T and ω are the number of leaves and the leaf weights in the tree. The first term is the objective of the traditional gradient tree boosting, while the second term added by XGBoost is to prevent overfitting.

We intrinsically use XGBoost as a two-level ensemble method here to combine the predictions of multiple forecasting methods. As shown in (5), XGBoost itself is an ensemble method that sums the results of K decision trees, each of which is a base model. Besides the ensemble nature of XGBoost, each decision tree therein nonlinearly ensembles the forecasts of base time-series models. The entire procedure turns out to be an ensemble model that has two levels of ensembles, which is expected to perform better than single-level ensemble methods such as those mentioned above.

4. Experiments

4.1. Data description

The sunspot number dataset is obtained from World Data Center SILSO, Royal Observatory of Belgium, Brussels¹. The dataset contains 3277 records of monthly averaged total sunspot number from January 1749 to January 2022. We also consider NASA’s past forecast (from April 1999 to January 2022) and most recent forecast (from February 2022 to October 2041)² for comparison with considered methods.

The observed sunspot number data is split in chronological order into training (2160 records from January 1749 to December 1928), validation (843 records from January 1929 to March 1999), and testing (274 records from April 1999 to January 2022) sets. The validation set monitors the training for hyperparameter tuning. The testing set evaluates the performance of each trained model on unseen data. The best model will be used to predict sunspot numbers from February 2022 to October 2041, covering the remaining portion of current Solar Cycle 25 and the coming Solar Cycle 26, in alignment with NASA’s current forecasting time range.

4.2. Implementation details

We compare the three non-deep learning models, SARIMA, Holt-Winters multiplicative Exponential Smoothing, and Prophet, and the four deep learning models, LSTM,

¹<https://www.bis.sidc.be/silso/datafiles>

²<https://www.nasa.gov/msfcsolar/archivedforecast>

GRU, Transformer, and Informer. We also consider their ensemble models from the mean and median ensemble methods, the error-based method, linear regression, and our proposed ensemble method via XGBoost.

We adopt the two evaluation metrics, the root mean squared error (RMSE) and the mean absolute error (MAE), which are widely-used in time series forecasting problems (Hyndman and Athanasopoulos 2018), to assess the prediction performance of considered methods:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2},$$

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i|,$$

where \hat{y}_i and y_i are the predicted value and the true value, respectively, and n is the number of predictions.

All the experiments are performed in Python 3.8 (Van Rossum and Drake 2009) environment. All deep learning methods are implemented with Pytorch 1.9.0 (Paszke et al. 2019). SARIMA, basic ensemble methods, and the error-based ensemble method are implemented with Merlion 1.0.0 (Bhatnagar et al. 2021). Exponential Smoothing and the linear regression ensemble method are performed with Darts 0.15.0 (Herzen et al. 2021). The XGBoost ensemble method is carried out using xgboost 1.5.1 (Chen and Guestrin 2016). Models are trained on two NVIDIA RTX8000 GPUs. We use Tune (Liaw et al. 2018) in Ray API 1.9.0 (Moritz et al. 2018) for tuning hyperparameters. All deep learning models are trained up to 200 epochs by the AdamW optimizer, and early-stopping is used to prevent overfitting. Source codes with pre-trained models, tuning ranges of model-specific hyperparameters, and best configurations for all models used in this study can be found at our Github repository³.

All ensemble methods are applied independently to ensemble the forecast outputs of the three non-deep learning models, the four deep learning models, and all the seven models, respectively. For basic ensemble methods, training is not required and performance is evaluated directly on the testing set. The ensemble weights from the error-based method are calculated by the inverse of the mean squared error (MSE) of each base model on the validation set. The linear-regression ensemble method computes the weights from the training and validation sets. MSE is chosen as the loss function for training XGBoost-based ensemble models.

4.3. Results

Table 1 presents the performance results of the seven base models on the testing set. All the four deep learning models outperform the three non-deep learning models. The two attention-based deep learning models (Transformer and Informer) exhibit better performance than the two RNN models (LSTM and GRU). In particular, Informer achieves the lowest RMSE and MAE with values 29.90 and 22.35, respectively. All the deep learning models except LSTM show more accurate results than NASA with Informer having 38.20% lower RMSE and 41.87% lower MAE. Figure 4 (a) displays the true sunspot numbers and their estimates from deep learning models and NASA

³https://github.com/yd1008/ts_ensemble_sunspot

for the testing set. It is observed that predictions of deep learning models generally follow along patterns in the ground truth data, while the LSTM and NASA have large deviations around years 2003 and 2020. Figure 4 (b) shows the forecasts from non-deep learning models. The three models fail to predict the true trend of the series, albeit have an approximately 11-year cycle. With the lowest RMSE and MAE among the three models, the estimates from SARIMA are rather accurate at the beginning of the testing time period but notably deviate from the actual values thereafter.

Table 1. Results of base models on the testing set of sunspot numbers.

	RMSE	MAE
SARIMA	54.11	45.51
Exponential Smoothing	61.41	49.76
Prophet	60.15	56.09
LSTM	46.14	39.44
GRU	37.14	26.77
Transformer	33.99	25.56
Informer	29.90	22.35
NASA	48.38	38.45

Table 2. Results of ensemble models on the testing set of sunspot numbers.

	Non-deep learning		Deep Learning		All Models	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
Mean	54.80	48.15	29.01	22.57	35.79	30.51
Median	53.45	46.02	29.96	21.75	39.52	33.88
Error-based	54.70	48.00	27.59	21.12	31.67	26.53
Regression	59.98	52.21	36.26	25.72	38.54	31.04
XGBoost	55.45	49.58	25.70	19.82	30.37	22.30

The improved performance of deep learning methods over non-deep learning ones can be attributed in part to the formers’ ability to effectively capture non-linearities, as observed in the complex variations in solar cycles, whereas the latters primarily use linear combinations. Furthermore, the two attention-based methods appear to better capture the trend using attention mechanisms that allow dependencies over a longer period of time to be included into forecasts.

Table 2 summarizes the testing results from the ensemble models. All ensembles of non-deep learning models and those of all base models fail to improve the performance since the forecasts are highly unstable and vary between non-deep learning models. All ensembles of deep learning models perform better than those of non-deep learning models due to the solid base on the well-performing deep learning models, and ensembles of all base models have performances between those of the former two.

In particular, the best ensemble model, our XGBoost-DL model, i.e., the XGBoost ensemble of deep learning models, significantly boosts the performance, with the smallest RMSE of 25.70 and MAE of 19.82 which are 14.05% and 11.32% lower than those of the best base model Informer. Figure 5 illustrates the forecasts from the three types of ensemble models. All deep-learning ensemble models show strong forecasting abilities in capturing the true trend and seasonality of the series.

We further consider the prediction of future sunspot numbers in current Solar Cycle 25 and the coming Solar Cycle 26. We fine-tune our XGBoost-DL model using the entire sunspot number data with a new training and validation split of ratio 7/3. Figure 6 shows the forecast from our XGBoost-DL model in comparison with NASA’s.

Our forecast indicates that the Solar Cycle 25 will reach a peak sunspot number of 133.47 in May 2025 and Solar Cycle 26 will have a peak number of 164.62 in November 2035. According to our prediction, the two Solar Cycles will be overall stronger than the past Solar Cycle 24. NASA’s forecast shows similar but earlier peak values of 137.7 for Solar Cycle 25 in October 2024 and 161.2 for Solar Cycle 26 in December 2034.

Although little work has been done on the sunspot number forecast for Solar Cycle 26, there are a number of recent forecasts for Solar Cycle 25 in the literature with the peak sunspot number ranging from 57.24 to 228.9 and occurring between 2022 and 2026. Our forecast is around the middle of the receptive ranges of magnitude and time. Labonville, Charbonneau, and Lemerle (2019) used a dynamo-based model and forecasted a weak Solar Cycle 25 with a maximum sunspot number of 89_{-14}^{+29} in $2025.3_{-1.05}^{+0.89}$. Covas, Peixinho, and Fernandes (2019) applied a feed-forward neural network and obtained a weaker Solar Cycle 25 with the peak sunspot number of 57.24 ± 16.76 in about 2022–2023. Han and Yin (2019) predicted a high peak value of 228.9 ± 40.5 around 2023.918 ± 1.64 years using the Vondrak smoothing method. Pala and Atici (2019) used the LSTM model and predicted a maximum sunspot number of 167.3 in July 2022. Benson et al. (2020) estimated the peak sunspot number in Solar Cycle 25 to be 106 ± 19.75 around $\text{March } 2025 \pm 1$ year by using a combination of the LSTM and WaveNet methods. Xiong et al. (2021) predicted with multiple regression a peak of 140.2 in March 2024. Prasad et al. (2022) used a stacked LSTM and predicted the cycle peak with value 171.9 ± 3.4 in around August 2023 ± 2 months.

5. Conclusion


We compare three non-deep learning models and four deep learning models to forecast sunspot numbers in this study. Deep learning models outperform non-deep learning models with lower RMSE and MAE. Additionally, five ensemble learning methods are applied to the forecast results of non-deep learning models, deep learning models, and all base models, separately. Ensemble models based on deep learning models have more accurate predictions than those based on non-deep learning models or all base models. Our proposed XGBoost-DL model that uses XGBoost to ensemble the four deep learning models achieves the best performance among all base and ensemble models as well as NASA’s forecast. Our prediction indicates that Solar Cycles 25 and 26 will be overall stronger than the most recent Solar Cycle 24, and will have the peak sunspot number of 133.47 in May 2025 and 164.62 in November 2035, respectively, which are similar to but later than the peaks forecast by NASA at 137.7 in October 2024 and 161.2 in December 2034.

Disclosure Statement


The authors report there are no competing interests to declare.

Funding

This work was supported by a startup fund from New York University.




(a) Forecasts from deep learning models and NASA.




(b) Forecasts from non-deep learning models.


Figure 4. Forecasts of sunspot numbers from base models for the testing set.



(a) Forecasts from ensembles of deep learning models and from NASA.



(b) Forecasts from ensembles of non-deep learning models.



(c) Forecasts from ensembles of all the seven base models.

Figure 5. Forecasts of sunspot numbers from ensemble models for the testing set.




Figure 6. Sunspot number forecast from February 2022 to October 2041.

References

- Adhikari, Ratnadip, and RK Agrawal. 2014. “Performance evaluation of weights selection schemes for linear combination of multiple forecasts.” *Artificial Intelligence Review* 42 (4): 529–548.
- Allende, Héctor, and Carlos Valle. 2017. “Ensemble methods for time series forecasting.” In *Claudio moraga: A passion for multi-valued logic and soft computing*, 217–232. Springer.
- Arfianti, Unix Izyah, Dian Candra Rini Novitasari, Nanang Widodo, Moh Hafiyusholeh, and Wika Dianita Utami. 2021. “Sunspot Number Prediction Using Gated Recurrent Unit (GRU) Algorithm.” *IJCCS (Indonesian Journal of Computing and Cybernetics Systems)* 15 (2): 141–152.
- Azcárate, T, B Mendoza, and JR Levi. 2016. “Influence of geomagnetic activity and atmospheric pressure on human arterial pressure during the solar cycle 24.” *Advances in Space Research* 58 (10): 2116–2125.
- Beltagy, Iz, Matthew E Peters, and Arman Cohan. 2020. “Longformer: The long-document transformer.” *arXiv preprint arXiv:2004.05150* .
- Benson, B, WD Pan, A Prasad, GA Gary, and Q Hu. 2020. “Forecasting solar cycle 25 using deep neural networks.” *Solar Physics* 295: 1–15.
- Bhatnagar, Aadyot, Paul Kassianik, Chenghao Liu, Tian Lan, Wenzhuo Yang, Rowan Cassius, Doyen Sahoo, et al. 2021. “Merlion: A machine learning library for time series.” *arXiv preprint arXiv:2109.09265* .
- Box, George E. P., and Gwilym M. Jenkins. 1976. *Time series analysis: forecasting and control*. Revised ed., Holden-Day Series in Time Series Analysis. Holden-Day, San Francisco, Calif.-Düsseldorf-Johannesburg.
- Breiman, Leo. 2001. “Random forests.” *Machine learning* 45 (1): 5–32.
- Breiman, Leo, Jerome H. Friedman, Richard A. Olshen, and Charles J. Stone. 1984. *Classification and regression trees*. Wadsworth Statistics/Probability Series. Wadsworth Advanced Books and Software, Belmont, CA.
- Brown, R.G. 1956. *Exponential Smoothing for Predicting Demand*. Arthur D. Little Inc. <https://www.industrydocuments.ucsf.edu/tobacco/docs/#id=jzlc0130>.
- Chang, Yung-Chia, Kuei-Hu Chang, and Guan-Jhih Wu. 2018. “Application of eXtreme gradient boosting trees in the construction of credit risk assessment models for financial institutions.” *Applied Soft Computing* 73: 914–920.
- Chattopadhyay, Surajit, Deepak Jhajharia, and Goutami Chattopadhyay. 2011. “Trend esti-

- mation and univariate forecast of the sunspot numbers: development and comparison of ARMA, ARIMA and autoregressive neural network models.” *Comptes Rendus Geoscience* 343 (7): 433–442.
- Chen, Tianqi, and Carlos Guestrin. 2016. “Xgboost: A scalable tree boosting system.” In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.
- Cherkassky, Vladimir, and Yunqian Ma. 2004. “Practical selection of SVM parameters and noise estimation for SVM regression.” *Neural networks* 17 (1): 113–126.
- Child, Rewon, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. “Generating long sequences with sparse transformers.” *arXiv preprint arXiv:1904.10509* .
- Cho, Kyunghyun, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. “On the Properties of Neural Machine Translation: Encoder–Decoder Approaches.” In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, Doha, Qatar, 103–111. Association for Computational Linguistics.
- Chung, Junyoung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2014. “Empirical evaluation of gated recurrent neural networks on sequence modeling.” In *NIPS 2014 Workshop on Deep Learning*, ArXiv preprint arXiv:1412.3555.
- Covas, Eurico, Nuno Peixinho, and João Fernandes. 2019. “Neural network forecast of the sunspot butterfly diagram.” *Solar Physics* 294 (3): 1–15.
- Davis Jr, George E, and Walter E Lowell. 2004. “Chaotic solar cycles modulate the incidence and severity of mental illness.” *Medical hypotheses* 62 (2): 207–214.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. “Bert: Pre-training of deep bidirectional transformers for language understanding.” *arXiv preprint arXiv:1810.04805* .
- Dong, Xibin, Zhiwen Yu, Wenming Cao, Yifan Shi, and Qianli Ma. 2020. “A survey on ensemble learning.” *Frontiers of Computer Science* 14 (2): 241–258.
- Friedman, Jerome H. 2001. “Greedy function approximation: a gradient boosting machine.” *Annals of statistics* 1189–1232.
- Han, YB, and ZQ Yin. 2019. “A decline phase modeling for the prediction of solar cycle 25.” *Solar Physics* 294 (8): 1–14.
- Han, Zhongyang, Jun Zhao, Henry Leung, King Fai Ma, and Wei Wang. 2019. “A review of deep learning models for time series prediction.” *IEEE Sensors Journal* 21 (6): 7833–7848.
- Hathaway, David H. 2015. “The solar cycle.” *Living reviews in solar physics* 12 (1): 4.
- Herzen, Julien, Francesco Lässig, Samuele Giuliano Piazzetta, Thomas Neuer, Léo Tafti, Guillaume Raille, Tomas Van Pottelbergh, et al. 2021. “Darts: User-friendly modern machine learning for time series.” *arXiv preprint arXiv:2110.03224* .
- Hiremath, KM. 2008. “Prediction of solar cycle 24 and beyond.” *Astrophysics and Space Science* 314 (1): 45–49.
- Hochreiter, Sepp. 1998. “The vanishing gradient problem during learning recurrent neural nets and problem solutions.” *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 6 (02): 107–116.
- Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. “Long short-term memory.” *Neural computation* 9 (8): 1735–1780.
- Holt, Charles C. 1957. *Forecasting seasonals and trends by exponentially weighted moving averages*. Office of Naval Research Memorandum 52.
- Hyndman, R.J., and G. Athanasopoulos. 2018. *Forecasting: principles and practice*. 2nd ed. Melbourne, Australia: OTexts.
- Juckett, David A, and Barnett Rosenberg. 1993. “Correlation of human longevity oscillations with sunspot cycles.” *Radiation research* 133 (3): 312–320.
- Kaushik, Shruti, Abhinav Choudhury, Pankaj Kumar Sheron, Nataraj Dasgupta, Sayee Natarajan, Larry A Pickett, and Varun Dutt. 2020. “AI in healthcare: time-series forecasting using statistical, neural, and ensemble architectures.” *Frontiers in big data* 3: 4.
- Kuncheva, Ludmila I. 2014. *Combining pattern classifiers: methods and algorithms*. John Wiley & Sons.

- Labonville, Francois, Paul Charbonneau, and Alexandre Lemerle. 2019. “A dynamo-based forecast of solar cycle 25.” *Solar Physics* 294 (6): 1–14.
- Lara-Benítez, Pedro, Manuel Carranza-García, and José C Riquelme. 2021. “An Experimental Review on Deep Learning Architectures for Time Series Forecasting.” *arXiv preprint arXiv:2103.12057* .
- Li, Qi, and Jeffrey Scott Racine. 2007. *Nonparametric econometrics: theory and practice*. Princeton University Press.
- Li, Shiyang, Xiaoyong Jin, Yao Xuan, Xiyong Zhou, Wenhui Chen, Yu-Xiang Wang, and Xifeng Yan. 2019. “Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting.” *Advances in Neural Information Processing Systems* 32: 5243–5253.
- Liaw, Richard, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E Gonzalez, and Ion Stoica. 2018. “Tune: A Research Platform for Distributed Model Selection and Training.” *arXiv preprint arXiv:1807.05118* .
- Lidiema, Caspah. 2017. “Modelling and Forecasting Inflation Rate in Kenya Using SARIMA and Holt-Winters Triple Exponential Smoothing.” *American Journal of Theoretical and Applied Statistics* 6 (3): 161–169.
- Liu, Huan, Chenxi Li, Yingqi Shao, Xin Zhang, Zhao Zhai, Xing Wang, Xinye Qi, et al. 2020. “Forecast of the trend in incidence of acute hemorrhagic conjunctivitis in China from 2011–2019 using the Seasonal Autoregressive Integrated Moving Average (SARIMA) and Exponential Smoothing (ETS) models.” *Journal of infection and public health* 13 (2): 287–294.
- Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. “Roberta: A robustly optimized bert pretraining approach.” *arXiv preprint arXiv:1907.11692* .
- Moritz, Philipp, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, et al. 2018. “Ray: A distributed framework for emerging {AI} applications.” In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, 561–577.
- Ogunleye, Adeola, and Qing-Guo Wang. 2019. “XGBoost model for chronic kidney disease diagnosis.” *IEEE/ACM transactions on computational biology and bioinformatics* 17 (6): 2131–2140.
- Oliveira, Mariana, and Luis Torgo. 2015. “Ensembles for time series forecasting.” In *Asian Conference on Machine Learning*, 360–370. PMLR.
- Pala, Zeydin, and Ramazan Atici. 2019. “Forecasting sunspot time series using deep learning methods.” *Solar Physics* 294 (5): 1–14.
- Pan, Bingyue. 2018. “Application of XGBoost algorithm in hourly PM_{2.5} concentration prediction.” In *IOP conference series: earth and environmental science*, Vol. 113, 012127. IOP publishing.
- Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, et al. 2019. “Pytorch: An imperative style, high-performance deep learning library.” *Advances in neural information processing systems* 32: 8026–8037.
- Penrose, Roger. 1956. “On best approximate solutions of linear matrix equations.” In *Mathematical Proceedings of the Cambridge Philosophical Society*, Vol. 52, 17–19. Cambridge University Press.
- Prasad, Amrita, Soumya Roy, Arindam Sarkar, Subhash Chandra Panja, and Sankar Narayan Patra. 2022. “Prediction of solar cycle 25 using deep learning based long short-term memory forecasting technique.” *Advances in Space Research* 69 (1): 798–813.
- Pulkkinen, Tuija. 2007. “Space weather: terrestrial perspective.” *Living Reviews in Solar Physics* 4 (1): 1–60.
- Qiu, Xueheng, Le Zhang, Ye Ren, Ponnuthurai N Suganthan, and Gehan Amaratunga. 2014. “Ensemble deep learning for regression and time series forecasting.” In *2014 IEEE symposium on computational intelligence in ensemble learning (CIEL)*, 1–6. IEEE.
- Qu, Jiangwen. 2016. “Is sunspot activity a factor in influenza pandemics?” *Reviews in medical virology* 26 (5): 309–313.

- Rabbani, Muhammad Babar Ali, Muhammad Ali Musarat, Wesam Salah Alaloul, Muhammad Shoaib Rabbani, Ahsen Maqsoom, Saba Ayub, Hamna Bukhari, and Muhammad Altaf. 2021. “A Comparison Between Seasonal Autoregressive Integrated Moving Average (SARIMA) and Exponential Smoothing (ES) Based on Time Series Model for Forecasting Road Accidents.” *Arabian Journal for Science and Engineering* 46 (11): 11113–11138.
- Sagi, Omer, and Lior Rokach. 2018. “Ensemble learning: A survey.” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 8 (4): e1249.
- Tabassum, Anika, Masud Rabbani, and Saad Bin Omar. 2020. “An Approach to Study on MA, ES, AR for Sunspot Number (SN) Prediction and to Forecast SN with Seasonal Variations Along with Trend Component of Time Series Analysis Using Moving Average (MA) and Exponential Smoothing (ES).” In *Advances in Electrical and Computer Technologies*, 373–386. Springer.
- Taylor, Sean J, and Benjamin Letham. 2018. “Forecasting at scale.” *The American Statistician* 72 (1): 37–45.
- Torres, José F, Dalil Hadjout, Abderrazak Sebaa, Francisco Martínez-Álvarez, and Alicia Troncoso. 2021. “Deep Learning for Time Series Forecasting: A Survey.” *Big Data* 9 (1): 3–21.
- Tsui, Fu-Chiang, Mingui Sun, Ching-Chung Li, and Robert J Scwabassi. 1995. “Recurrent neural networks and discrete wavelet transform for time series modeling and prediction.” In *1995 International Conference on Acoustics, Speech, and Signal Processing*, Vol. 5, 3359–3362. IEEE.
- Usoskin, Ilya G. 2017. “A history of solar activity over millennia.” *Living Reviews in Solar Physics* 14 (1): 1–97.
- Van Rossum, Guido, and Fred L. Drake. 2009. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. “Attention is all you need.” In *Advances in neural information processing systems*, 5998–6008.
- Wen, Qingsong, Liang Sun, Fan Yang, Xiaomin Song, Jingkun Gao, Xue Wang, and Huan Xu. 2020. “Time series data augmentation for deep learning: A survey.” *arXiv preprint arXiv:2002.12478* .
- Wichard, Jorg D, and Maciej Ogorzalek. 2004. “Time series prediction with ensemble models.” In *2004 IEEE International Joint Conference on Neural Networks (IEEE Cat. No. 04CH37541)*, Vol. 2, 1625–1630. IEEE.
- Winters, Peter R. 1960. “Forecasting sales by exponentially weighted moving averages.” *Management science* 6 (3): 324–342.
- Wu, Neo, Bradley Green, Xue Ben, and Shawn O’Banion. 2020. “Deep transformer models for time series forecasting: The influenza prevalence case.” *arXiv preprint arXiv:2001.08317* .
- Xiong, Yating, Jianyong Lu, Kai Zhao, Meng Sun, and Yang Gao. 2021. “Forecasting Solar Cycle 25 using comprehensive precursor combination and multiple regression technique.” *Monthly Notices of the Royal Astronomical Society* 505 (1): 1046–1052.
- Xu, Tong, Jian Wu, Zhen-Sen Wu, and Qiang Li. 2008. “Long-term sunspot number prediction based on EMD analysis and AR model.” *Chinese journal of astronomy and astrophysics* 8 (3): 337.
- Yang, Zhilin, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. “Xlnet: Generalized autoregressive pretraining for language understanding.” *Advances in neural information processing systems* 3232.
- Zhang, Jun, and Kim-Fung Man. 1998. “Time series prediction using RNN in multi-dimension embedding phase space.” In *SMC’98 Conference Proceedings. 1998 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No. 98CH36218)*, Vol. 2, 1868–1873. IEEE.
- Zhong, Jiancheng, Yusui Sun, Wei Peng, Minzhu Xie, Jiahong Yang, and Xiwei Tang. 2018. “XGBFEMF: an XGBoost-based framework for essential protein prediction.” *IEEE transactions on nanobioscience* 17 (3): 243–250.
- Zhou, Haoyi, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. “Informer: Beyond Efficient Transformer for Long Sequence Time-Series

Forecasting.” In *The Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Virtual Conference*, Vol. 35, 11106–11115. AAAI Press.