

# Arbitrated Ensemble for Solar Radiation Forecasting

Vítor Cerqueira<sup>1,2(✉)</sup>, Luís Torgo<sup>1,2</sup>, and Carlos Soares<sup>1,2</sup>

<sup>1</sup> INESC TEC, Porto, Portugal

{vmac,ltorgo}@inesctec.pt, csoares@fe.up.pt

<sup>2</sup> University of Porto, Porto, Portugal

**Abstract.** Utility companies rely on solar radiation forecasting models to control the supply and demand of energy as well as the operability of the grid. They use these predictive models to schedule power plant operations, negotiate prices in the electricity market and improve the performance of solar technologies in general. This paper proposes a novel method for global horizontal irradiance forecasting. The method is based on an ensemble approach, in which individual competing models are arbitrated by a metalearning layer. The goal of arbitrating individual forecasters is to dynamically combine them according to their aptitude in the input data. We validate our proposed model for solar radiation forecasting using data collected by a real-world provider. The results from empirical experiments show that the proposed method is competitive with other methods, including current state-of-the-art methods used for time series forecasting tasks.

**Keywords:** Solar radiation forecasting · Renewable energy · Ensemble methods · Metalearning · Time series

## 1 Introduction

The sun is a fundamental source of energy to our planet. Over the years many technologies have been developed to take advantage of the solar energy. We use this energy to produce electricity, heating water and buildings, lighting and destroying toxic waste.

In this paper we present a novel method for solar radiation forecasting. Specifically, we focus on daily predictions of Global Horizontal Irradiance (GHI). Utility companies use solar radiation forecasting systems to support decision making process in several ways. They use them to predict if the energy produced by a given solar technology can meet the daily electricity demand, to balance electricity market prices and schedule power plant operations. Engineers leverage solar radiation forecasting systems to improve the performance and economics of solar radiation technologies, for example photovoltaic devices. Moreover, solar radiation forecasts enable a dynamic configuration of air-conditioning systems within buildings to optimise their efficiency. In sum, a data-driven solar radiation forecasting system can provide a way to maximise the performance of solar

technologies while reducing operating costs. This paper presents a novel time series model for GHI forecasting. The proposed method leverages the predictive power of ensemble methods, combining individual learning models with different inductive bias using a metalearning strategy. We explore ways of combining the predictions of forecasters in a dynamic – online – fashion. In time evolving environments the process causing the underlying data is prone to change over time and the combined model should adapt accordingly.

We use metalearning (e.g. [4]) to analyse the expertise of each individual forecaster across the time series of solar radiation. We can then use this meta knowledge to dynamically weight the predictions of base learners according to their competence in a given observation. If we expect a given forecaster to perform poorly in some subset of the data we render it a low weight in the combination rule. On the other hand, if we are optimistic about some learner in our ensemble we augment its weight with respect to the other learners.

The intuition behind our approach is that different learning models may have different areas of expertise across the input space. That is, in a given test observation, some individual learner may be better than the combined model and different individual learners will be better on different cases. In effect, we can learn about the learning process of each base-learner. In fact, we hypothesise that the underlying process causing the series of solar radiation follows a recurring pattern due to seasonal factors [7]. Consequently, we hypothesise that the metalearning layer enables the combined model to better detect changes between different regimes and quickly adapt itself to the environment.

Our metalearning strategy follows an Arbitrating scheme [13,22], in which we build a meta-learner for each base-learner comprising the ensemble. Each meta-learner is specifically designed to assess how apt its base counterpart is to make a prediction in a given observation. This is accomplished by analysing how the error incurred by a given learning model relates to the characteristics of the data. At test time, the base-learners are weighted according to their degree of competence in the input observation, estimated by the predictions of the meta-learners.

Our goal is to predict the next value of the series of solar radiation. We use regression models as our base-learners by transforming the solar radiation time series into an Euclidean space using time delay embedding [29]. Furthermore, in order to augment the information about the data we also use external predictors such as weather reports. In summary, the contributions of this paper are the following:

- An arbitrated ensemble for GHI forecasting. The ensemble includes a meta-learning layer based on an arbitrating scheme, used to dynamically combine individual models;
- We use the Arbitrating strategy to dynamically weight individual models, whereas typical applications select the most reliable model at each test query.

We start by outlining the related work in Sect. 2. The methodology is addressed in Sect. 3, where we formalise and explain our contributions. The Case Study is briefly described in Sect. 4, along with the pre-processing steps

and descriptive statistics. The experiments and respective results are presented and discussed in Sect. 5. Finally, the conclusions are produced in Sect. 6, along with some remarks about future work and reproducibility of the proposed methods.

## 2 Related Work

In this paper we focus on ensembles with self-adapting mechanisms to predict solar radiation, which is a time series with a numerical outcome. Ensemble methods for numerical predictions problems have a vast literature. We refer to the survey written in [19] for a complete overview on ensemble approaches for these tasks.

Building adaptable models is important in dynamic real-world environments in which data is constantly changing over time due to several factors, for example, seasonality. Our proposed method is motivated by the core concepts behind Arbitrating classifiers [13, 22]. Arbitrating is an ensemble method used to combine classifiers according to their expertise on the input data. The expertise of a base-learner is learnt with a corresponding meta-model, which learns the loss of its base counterpart according to a set of meta-features. At test time, the classifier with greatest confidence in the input data-point is selected to make a prediction. The authors reason that each meta-model holds information about which part of the data its base counterpart works best and considers when it can make a reliable prediction.

Other forms for dynamically combining models for time series forecasting with numerical outcome were proposed in the literature. In [26], the authors use Zoomed Ranking [28] approach to rank and select time series forecasting models. MetaStream is proposed in [25]. The authors summarise the dynamics of the recent and upcoming observations in a data stream to either select or combine regression models. In [15, 32] presented other two approaches that use the characteristics of time series in a meta-level to improve the combination of individual forecasters. They use these characteristics to induce several rules to weight or select between different models.

Our approach is different from the existing literature in the sense that we apply an arbitrating scheme to meta-learn and weight the individual base-learners. To the best of our knowledge, this is the first application of an Arbitrating scheme for time series prediction with numerical outcome, particularly solar radiation forecasting.

### 2.1 Solar Radiation Forecasting

Several solar radiation forecasting models have been proposed in the literature. The most typical approaches rely on regression and time series analysis models (e.g. [3, 8, 23]). The connectionist approach of Artificial Neural Networks is also of common use, for example in [11, 18] or [27]. In our paper we focus on daily

forecasts, but the temporal granularity typically ranges from hourly to weekly forecasts.

Many approaches also incorporate external features in their methodology, such as [16] or [1]. These typically include weather information.

### 3 Global Horizontal Irradiance Forecasting

A given solar technological device collects solar radiation in two ways: direct radiation and diffuse radiation. In this paper we aim at predicting the Global Horizontal Irradiance (GHI), which can be derived by summing direct radiation with diffuse radiation and accounting for the sun’s position.

GHI forecasting is a particular instance of time series forecasting tasks. We start addressing the methodology by presenting the main notation employed throughout this section:

**Time Series:** A time series is a temporal sequence of values  $Y = \{y_1, y_2, \dots, y_n\}$ , where  $y_i$  is the value of  $Y$  at time  $i$  and  $n$  is the length of  $Y$ ;

**Embedded Time Series:**  $Y^K$  denotes the embedded time series with embedding dimension  $K$ . We use time delay embedding to represent  $Y$  in an Euclidean space with embedding dimension  $K$ , according to [29]. In effect, we generate the following matrix:

$$Y^K = \begin{bmatrix} y_1 & y_2 & \dots & y_{K-1} & y_K \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ y_{i-K+1} & y_{i-K+2} & \dots & y_{i-1} & y_i \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ y_{n-K+1} & y_{n-K+2} & \dots & y_{n-1} & y_n \end{bmatrix} \tag{1}$$

Each row denotes an embedding vector  $v_r, \forall r \in \{1, \dots, n - K + 1\}$ . Our goal is to predict the next point in the series, represented by the last column in Matrix 1;

**External Predictors:**  $Y^{ext}$  denotes the set of external predictors computed for each embedding vector  $v \in V$ . These include external information (e.g. weather data) which helps to model the target concept;

**Base-Learners:** We denote as  $M$  the set of  $m$  base-learners comprising the ensemble  $S$ ;

**Meta-Learners:**  $\overline{M}^j$  is a meta-learner for  $M^j$ , with  $j \in \{1, \dots, m\}$ ;

**Base-Learners Loss:**  $e_i^j$  represents the absolute loss of  $M^j$  in the observation  $y_i$ ;

**Base-Learners weights:**  $w_i^j$  denotes the weights assigned to  $M^j$  for predicting the value of  $y_i$ .

Our methodology for GHI forecasting settles on the three main steps: An offline (i) training step of  $M$  and the online iterative steps: (ii) Meta-learning of  $\overline{M}$  and (iii) prediction of  $y_{t+1}$  using  $M$  which is dynamically weighted according to  $\overline{M}$ .

### 3.1 Learning $M$

In the first step we train the learning models  $M$  which are then combined to make a prediction. Concretely, each  $M^j, \forall j \in \{1, \dots, m\}$  is individually trained using the available  $Y_{tr}^K$ , the embedded time series combined with the external predictors  $Y^{ext}$ .  $M$  is composed of individual regression models with different inductive bias. Different models (e.g. Gaussian Processes and Neural Networks) hold different assumptions regarding the underlying data. This divergence across base-learners comprising  $S$  encourages diversity in the ensemble – a fundamental ingredient in the ensemble recipe [5].

### 3.2 Metalearning $\overline{M}$

The metalearning step of our methodology is an online process run at test time. Our objective in applying this metalearning strategy is to extract the information about the expertise of each individual model in  $M$  across the series of water consumption.

We use a metalearning layer for arbitrating among competing individual learners. However, instead of selecting the most reliable model (as in [13, 22]), we use the meta-knowledge to weight the base learners according to their expertise in the input signal.

Formally, each meta-learner  $\overline{M}^j, \forall j \in \{1, \dots, m\}$ , is trained to build a model for  $e^j = f(\overline{X})$ , where  $f$  denotes the regression function.  $\overline{X}$  represent the meta-features, i.e., the set of **features** used in the **meta**-level by the meta-learners in  $\overline{M}$ .  $\overline{X}$  is composed by the primitive features used by  $M$  along with some summary statistics. These statistics are computed for each embedding vector and characterise the recent dynamics of the series as well as its structure.

We conduct this meta regression analysis to understand how the loss of a given base-learner relates to the different dynamics of the series. In effect, we can explore forms of capitalising from these relationships. Specifically, we use the information from  $e$  to dynamically weight the base-learners  $M$ .

### 3.3 Predicting $y_{t+1}$

When a new observation  $y_{t+1}$  arrives for prediction we combine the predictions of  $M$  with the meta information from  $\overline{M}$ . The arbitrating layer composed by  $\overline{M}$  is able to predict how well each base learner in  $M$  will perform with respect to each other. If  $\overline{M}^j$  predicts that its counterpart  $M^j$  will make a large error ( $\hat{e}_{t+1}^j$ ) relative to the other base learners ( $\hat{e}_{t+1}^l, \forall l \in \{1, \dots, m\} \setminus \{j\}$ ) then  $M^j$  will be assigned a small relative weight in the final prediction. Conversely, if  $\hat{e}_{t+1}^j$  is predicted to be small (also with respect to the loss of other base learners),  $M^j$  will be important for the upcoming prediction. Even though the learning models comprising  $M$  are trained in a batch way, the models in  $\overline{M}$  are updated after every test observation. Moreover, the predictions by  $\overline{M}$  are produced for

each test observation, rendering an online nature to our method. Formally, we measure the weights of each base-learner using the following equation:

$$w_{t+1}^j = \frac{erfc(\hat{e}_{t+1}^j)}{\sum_{i \in \overline{M}} erfc(\hat{e}_{t+1}^i)} \tag{2}$$

where  $\hat{e}_{t+1}^j$  is the prediction made by  $\overline{M}^j$  for the absolute loss that  $M^j$  will incur in  $y_{t+1}$ . The function *erfc* denotes the complementary Gaussian error function which is formalised as follows:

$$erfc(x) = \frac{4}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt \tag{3}$$

The final prediction is a weighted average of the predictions made by the base-learners  $\hat{y}_{t+1}^j$  with respect to  $w_{t+1}^j$  computed according to Eq. 4:

$$\hat{y}_{t+1} = \sum_{j=1}^m \hat{y}_{t+1}^j \times w_{t+1}^j \tag{4}$$

The proposed methodology is summarised in Algorithm 1.

## 4 Case Study

Our study was conducted using data collected by the Oak Ridge National Laboratory [17], in Tennessee, USA. The solar radiation data includes global horizontal radiation, direct radiation and diffuse horizontal radiation. These were harvested using a rotating shadow-band radiometer, a low-cost equipment for measuring solar radiation.

The data is collected in an hourly basis. Our sample ranged from 19-01-2009 to 19-01-2017, totalling 70151 observations. Additionally, other external variables were collected: the average air temperature, relative humidity, average wind speed and precipitation levels. These follow the same granularity and temporal scope as the solar radiation data.

### 4.1 Pre-processing

We focused our work on daily forecasts so we aggregated the data by day, reaching a total of 2922 observations across the above-mentioned time-span. The units of the solar radiation levels are in watts per square meter ( $W/m^2$ ).

Direct radiation and diffuse radiation levels are used as predictor variables in our model. Concretely, we use the information of these attributes from the previous day as well as the mean of the last  $K$  days. To augment the information of solar radiation levels we also include the mean and standard deviation of the embedding vectors described in Matrix 1 as predictors. Moreover, from the hourly average air temperature we design two features: max and mean air

---

**Algorithm 1:** Arbitrated Ensemble for Solar Radiation Forecasting

---

```

Input:
– Global Horizontal Irradiance time series  $Y$ ;
– External predictors  $Y^{ext}$ ;
– embedding dimension  $K$ .
/* Learning Step */
Embed  $Y$  and attach  $Y^{ext} \rightarrow Y_{tr}^K$ 
foreach  $M^j \in M$  do
| train  $M^j$  using  $Y_{tr}^K$ 
end

/* Metalearning Step */
for all available test observations  $Y_{ts}$  do
| Embed  $Y_{ts}$  and attach  $Y^{ext} \rightarrow Y_{ts}^K$ 
|  $\bar{X} \leftarrow$  Extract Metafeatures from  $Y_{ts}^K$ 
| foreach  $\bar{M}^j \in \bar{M}$  do
| | train  $\bar{M}^j$  to model:  $e^j = f(\bar{X})$ 
| end
end

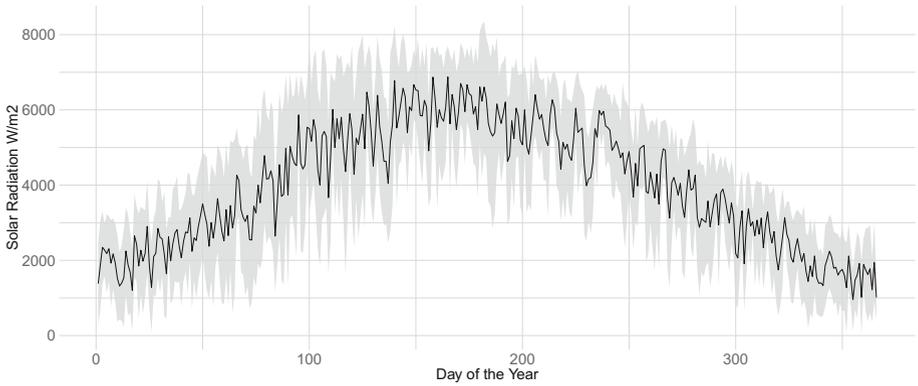
/* Predicting new observations */
foreach upcoming observation  $y_{t+1}$  do
| Embed  $y_{t+1}$  along with the  $K - 1$  most recent observations and attach  $y_{t+1}^{ext}$ 
| Compute predictions  $\hat{e}_{t+1}^j$  from models  $\bar{M}^j \in \bar{M}$ 
| Compute weights  $w_{t+1}^j$ 
| Compute predictions  $\hat{y}_{t+1}^j$  from models  $M^j \in M$ 
| Compute final prediction  $\hat{y}_{t+1} = \sum_{j=1}^m \hat{y}_{t+1}^j \cdot w_{t+1}^j$ 
| Add  $y_{t+1}$  to  $Y_{ts}$  and return to metalearning step
end

```

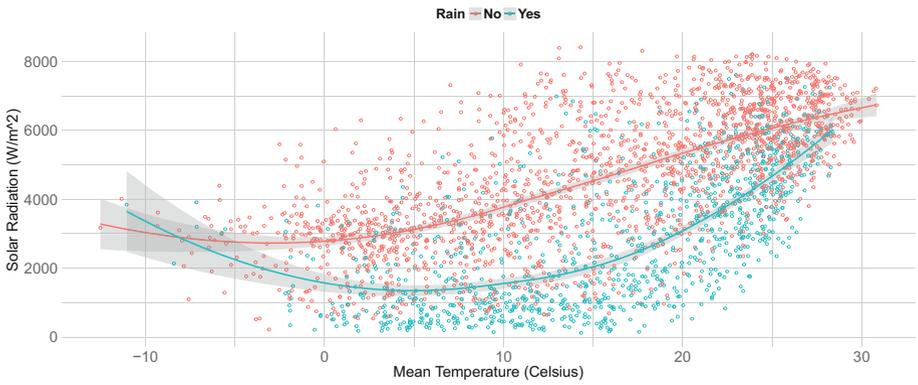
---

temperature of a given day. From the precipitation levels we create a logical variable that describes if it did or did not rain in a given day.

In Figs. 1 and 2 we present a view to the solar radiation dynamics. Figure 1 shows the mean and respective deviation of solar radiation levels per day of the year. As expected, solar radiation is higher in the warmer seasons of the year. Nonetheless, it also presents a complex structure with several peaks across the days. Figure 2 illustrates the solar radiation per mean temperature, grouped by days that had and had not any rainfall. It also contains LOESS curves that indicate a positive correlation between temperature and solar radiation. Moreover, days without rainfall have considerably higher solar radiation than rainy ones.



**Fig. 1.** Mean and standard deviation of solar radiation per day of the year (in  $W/m^2$ )



**Fig. 2.** Solar radiation by mean temperature, grouped by days that had rainfall and days that had not any rainfall

## 5 Empirical Experiments

In this section we present the empirical experiments carried out to validate the proposed method for solar radiation forecasting. These address the following research questions:

- Q1:** Is it beneficial to weight individual forecasters according to an Arbitrating scheme for solar radiation forecasting tasks?;
- Q2:** How does the performance of the proposed method relates to the performance of the state-of-the-art methods for solar radiation forecasting tasks?

The experiments were carried out using performance Estimation [30] *R* package. The methods used in the experiments were evaluated using the Root Mean Squared Error (RMSE) and the Mean Absolute Error (MAE) using a Monte Carlo procedure with 10 repetitions. For each repetition, a random point is

picked from the time series. The previous window comprising **40%** of  $t$  is used for training and the following window of **20%** of  $t$  is used for testing.

The metafeatures used by  $\bar{M}$  are the primitive ones, described previously in the Sect. 4.1, together with the following characteristics computed at each embedding vector: **(i) kurtosis**, which is a measure of flatness of the data distribution with respect to a Gaussian distribution; **(ii) skewness**, which measures the symmetry of the distribution; **(iii) series trend**, calculated according to the ratio between the standard deviation of the series and the standard deviation of the differenced series; **(iv) serial correlation**, estimated using a Box-Pierce test statistic; and **(v) long-range dependence**, using a Hurst exponent estimation with a wavelet transform. These statistics summarise the overall structure of the time series of water consumption. For a comprehensive description of each statistic see [32].

We estimate the optimal embedding dimension ( $K$ ) using the method of False Nearest Neighbours [12]. This method analyses the behaviour of the nearest neighbours as we increase  $K$ . According to the authors from [12], with a low sub-optimal  $K$  many of the nearest neighbours will be false. Then, as we increase  $K$  and approach an optimal embedding dimension those false neighbours disappear.  $K$  is set to **6** in our experiments.

The base-learners  $M$  comprising the ensemble are the following: MARS [20], Generalized Linear Models [6], Random Forest [34], SVM [10], Rule-based regression [14], Generalized Boosted Regression [24], Gaussian Processes [10] and Feed Forward Neural Networks [31]. Each of the individual learners is composed of 6 different parameter settings adding up to **48** learning models. We use a Random Forest as a meta-learner model.

We compare the proposed method to the following four baselines:

**ARIMA:** The state-of-the-art ARIMA model, using the function *auto.arima* from the forecast R package [9]. This function automatically tunes ARIMA to an **optimal parameter setting**;

**ARIMAX:** Similar to the one above, but augmented with the external features outlined in the case study section;

**NN:** A feed forward neural network with a single hidden layer. The neural network was optimized using a grid search procedure using a total of 56 parameter combinations. The final parameter setting was 7 hidden units and a weight decay of 0.2;

**BT:** Bagged Trees from [21]. This bagging approach is specifically designed for time series forecasting tasks;

**S:** This is a variant of the proposed method, but stripped of the metalearning layer. That is,  $M$  is trained in advance and their predictions are simply averaged at run-time using the arithmetic mean;

**Blending:** We use a metalearning technique called Blending to combine the individual learning models. Blending was introduced in [2] in their winning solution for the well known Netflix prize. In practice, it is a variant of Stacking [33] in which out-of-bag predictions are produced with a holdout strategy.

The results from the Monte Carlo experiments are reported in Table 1. Besides the baselines described above, **AE** denotes the proposed method for GHI forecasting tasks.

Relative to **S**, our approach is able to overcome its performance, with similar deviance across Monte Carlo repetitions, both in terms of RMSE and MAE. Overall we conclude that indeed our metalearning approach is beneficial for solar radiation forecasting tasks (**Q1**).

The performance of the proposed method are slightly better than the ARIMAX approach, which shows the competitiveness of our method. Other state-of-the-art baselines such as NN and ARIMA (without external features) perform clearly worse than our method. In effect, we also conclude that our hypothesis **Q2** is valid. Our method was also able to overcome the Blending approach to model combination, which is a widely used technique.

In summary, our experiments validate our hypothesis that our proposed method is able to model the different dynamics of solar radiation with a competitive performance relative to state-of-the-art methods.

**Table 1.** Average results from the methods using RMSE and MAE

Method	RMSE	MAE
ARIMA	1407.561 ± 88.06	1129.49 ± 76.32
ARIMAX	728.43 ± 35.46	579.18 ± 29.03
BT	1068.96 ± 52.90	856.57 ± 44.85
NN	1232.26 ± 610.20	1000.47 ± 533.57
Blending	731.49 ± 40.67	566.80 ± 32.93
S	808.82 ± 46.99	649.36 ± 39.12
<b>AE</b>	<b>721.31 ± 41.47</b>	<b>566.22 ± 35.49</b>

## 6 Conclusions

In this paper we presented a new method for GHI forecasting tasks. We argued that the planning of operations related to solar radiation is an important topic with economical and social impact. Our proposed method settles on a metalearning scheme called Arbitrating, introduced before by [13, 22]. We extend their ideas to GHI forecasting tasks.

We leverage the Arbitrating strategy to dynamically weight individual models in an ensemble. We reasoned that the series of water consumption follows a recurring pattern with different regimes. In effect, our approach allows a fast detection and adaptation to the different regimes causing the data.

Results from numerical experiments suggest that our metalearning is worthwhile. Moreover, we empirically demonstrate that the proposed method is competitive with other state-of-the-art techniques for GHI forecasting tasks, such as Neural Networks and the classical time series model ARIMA.

Future work includes: (i) generalise the proposed methodology for other time series forecasting tasks; (ii) Compare the proposed method against a time-dependent combining heuristic (e.g. recent performance of individual learners).

In the interest of reproducible research our methods are publicly available as an **R** package called **tsensembler**<sup>1</sup>.

**Acknowledgements.** This work is financed by the ERDF - European Regional Development Fund through the Operational Programme for Competitiveness and Internationalisation - COMPETE 2020 Programme within project POCI-01-0145-FEDER-006961, and by National Funds through the FCT - Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) as part of project UID/EEA/50014/2013; Project “NORTE-01-0145-FEDER-000036” is financed by the North Portugal Regional Operational Programme (NORTE 2020), under the PORTUGAL 2020 Partnership Agreement, and through the European Regional Development Fund (ERDF).

## References

1. Bacher, P., Madsen, H., Nielsen, H.A.: Online short-term solar power forecasting. *Sol. Energy* **83**(10), 1772–1783 (2009)
2. Bell, R.M., Koren, Y., Volinsky, C.: The bellKor 2008 solution to the netflix prize. Statistics Research Department at AT&T Research (2008)
3. Boland, J.: Time series modelling of solar radiation. In: Badescu, V. (ed.) *Modeling Solar Radiation at the Earths Surface*, pp. 283–312. Springer, Heidelberg (2008)
4. Brazdil, P., Carrier, C.G., Soares, C., Vilalta, R.: *Metalearning: Applications to Data Mining*. Springer Science & Business Media, Heidelberg (2008)
5. Brown, G.: Ensemble learning. In: Sammut, C., Webb, G.I. (eds.) *Encyclopedia of Machine Learning*, pp. 312–320. Springer, Boston (2010)
6. Friedman, J., Hastie, T., Tibshirani, R.: Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* **33**(1), 1–22 (2010)
7. Gama, J., Kosina, P.: Tracking recurring concepts with meta-learners. In: Lopes, L.S., Lau, N., Mariano, P., Rocha, L.M. (eds.) *EPIA 2009*. LNCS, vol. 5816, pp. 423–434. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-04686-5\\_35](https://doi.org/10.1007/978-3-642-04686-5_35)
8. Huang, J., Korolkiewicz, M., Agrawal, M., Boland, J.: Forecasting solar radiation on an hourly time scale using a coupled autoregressive and dynamical system (cards) model. *Sol. Energy* **87**, 136–149 (2013)
9. Hyndman, R.J., Athanasopoulos, G., Razbash, S., Schmidt, D., Zhou, Z., Khan, Y., Bergmeir, C., Wang, E.: *Forecast: Forecasting functions for time series and linear models*, R package version 5.6 (2014)
10. Karatzoglou, A., Smola, A., Hornik, K., Zeileis, A.: kernlab - an S4 package for kernel methods in R. *J. Stat. Softw.* **11**(9), 1–20 (2004)
11. Kemmoku, Y., Orita, S., Nakagawa, S., Sakakibara, T.: Daily insolation forecasting using a multi-stage neural network. *Sol. Energy* **66**(3), 193–199 (1999)

<sup>1</sup> <https://github.com/vcerqueira/tsensembler>.

12. Kennel, M.B., Brown, R., Abarbanel, H.D.: Determining embedding dimension for phase-space reconstruction using a geometrical construction. *Phys. Rev. A* **45**(6), 3403 (1992)
13. Koppel, M., Engelson, S.P.: Integrating multiple classifiers by finding their areas of expertise. In: *AAAI-1996 Workshop On Integrating Multiple Learning Models*, pp. 53–58. Citeseer (1996)
14. Kuhn, M., Weston, S., Keefer, C., Quinlan, R.: *Cubist: Rule- and Instance-Based Regression Modeling*, R package version 0.0.18 (2014)
15. Lemke, C., Gabrys, B.: Meta-learning for time series forecasting and forecast combination. *Neurocomputing* **73**(10), 2006–2016 (2010)
16. Lorenz, E., Hurka, J., Heinemann, D., Beyer, H.G.: Irradiance forecasting for the power prediction of grid-connected photovoltaic systems. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2**(1), 2–10 (2009)
17. Maxey, C., Andreas, A.: Oak ridge national laboratory (ORNL); rotating shadow-band radiometer (RSR); oak ridge, tennessee (DATA). Technical report, National Renewable Energy Laboratory (NREL), Golden, CO, USA (2007)
18. Mellit, A., Pavan, A.M.: A 24-h forecast of solar irradiance using artificial neural network: application for performance prediction of a grid-connected PV plant at Trieste, Italy. *Sol. Energy* **84**(5), 807–821 (2010)
19. Mendes-Moreira, J., Soares, C., Jorge, A.M., Sousa, J.F.D.: Ensemble approaches for regression: a survey. *ACM Comput. Surv. (CSUR)* **45**(1), 10 (2012)
20. Milborrow, S.: *Earth: Multivariate Adaptive Regression Spline Models*. Derived from *mda:mars* by Trevor Hastie and Rob Tibshirani (2012)
21. Oliveira, M., Torgo, L.: Ensembles for time series forecasting. In: *ACML Proceedings of Asian Conference on Machine Learning. JMLR: Workshop and Conference Proceedings* (2014)
22. Ortega, J., Koppel, M., Argamon, S.: Arbitrating among competing classifiers using learned referees. *Knowl. Inf. Syst.* **3**(4), 470–490 (2001)
23. Reikard, G.: Predicting solar radiation at high resolutions: a comparison of time series forecasts. *Sol. Energy* **83**(3), 342–349 (2009)
24. Ridgeway, G.: *GBM: Generalized Boosted Regression Models*, R package version 2.1.1 (2015)
25. Rossi, A.L.D., de Leon Ferreira, A.C.P., Soares, C., De Souza, B.F., et al.: *Metastream: a meta-learning based method for periodic algorithm selection in time-changing data*. *Neurocomputing* **127**, 52–64 (2014)
26. dos Santos, P.M., Ludermir, T.B., Prudencio, R.B.C.: Selection of time series forecasting models based on performance information. In: *Fourth International Conference on Hybrid Intelligent Systems, HIS 2004*, pp. 366–371. IEEE (2004)
27. Sfetsos, A., Coonick, A.: Univariate and multivariate forecasting of hourly solar radiation with artificial intelligence techniques. *Sol. Energy* **68**(2), 169–178 (2000)
28. Soares, C., Brazdil, P.B.: Zoomed ranking: selection of classification algorithms based on relevant performance information. In: Zighed, D.A., Komorowski, J., Żytkow, J. (eds.) *PKDD 2000. LNCS*, vol. 1910, pp. 126–135. Springer, Heidelberg (2000). doi:[10.1007/3-540-45372-5\\_13](https://doi.org/10.1007/3-540-45372-5_13)
29. Takens, F.: Detecting strange attractors in turbulence. In: Rand, D., Young, L.-S. (eds.) *Dynamical Systems and Turbulence, Warwick 1980. LNM*, vol. 898, pp. 366–381. Springer, Heidelberg (1981). doi:[10.1007/BFb0091924](https://doi.org/10.1007/BFb0091924)
30. Torgo, L.: *An Infra-Structure for Performance Estimation and Experimental Comparison of Predictive Models*, R package version 0.1.1 (2013)
31. Venables, W.N., Ripley, B.D.: *Modern Applied Statistics with S*, 4th edn. Springer, New York (2002). ISBN 0-387-95457-0

32. Wang, X., Smith-Miles, K., Hyndman, R.: Rule induction for forecasting method selection: meta-learning the characteristics of univariate time series. *Neurocomputing* **72**(10), 2581–2594 (2009)
33. Wolpert, D.H.: Stacked generalization. *Neural Netw.* **5**(2), 241–259 (1992)
34. Wright, M.N.: Ranger: A Fast Implementation of Random Forests, R package version 0.3.0 (2015)