

Lawrence Berkeley National Laboratory

LBL Publications

Title

Long-term missing value imputation for time series data using deep neural networks

Permalink

<https://escholarship.org/uc/item/63f2185k>

Journal

Neural Computing and Applications, 35(12)

ISSN

0941-0643

Authors

Park, Jangho

Müller, Juliane

Arora, Bhavna

et al.

Publication Date

2023-04-01

DOI

10.1007/s00521-022-08165-6

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



Long-term missing value imputation for time series data using deep neural networks

Jangho Park¹ · Juliane Müller² · Bhavna Arora³ · Boris Faybishenko³ · Gilberto Pastorello² · Charuleka Varadharajan³ · Reetik Sahu⁴ · Deborah Agarwal²

Received: 5 March 2022 / Accepted: 6 December 2022 / Published online: 23 December 2022

This is a U.S. Government work and not under copyright protection in the US; foreign copyright protection may apply 2022

Abstract

We present an approach that uses a deep learning model, in particular, a MultiLayer Perceptron, for estimating the missing values of a variable in multivariate time series data. We focus on filling a long continuous gap (e.g., multiple months of missing daily observations) rather than on individual randomly missing observations. Our proposed gap filling algorithm uses an automated method for determining the optimal MLP model architecture, thus allowing for optimal prediction performance for the given time series. We tested our approach by filling gaps of various lengths (three months to three years) in three environmental datasets with different time series characteristics, namely daily groundwater levels, daily soil moisture, and hourly Net Ecosystem Exchange. We compared the accuracy of the gap-filled values obtained with our approach to the widely used R-based time series gap filling methods *ImputeTS* and *mtsdi*. The results indicate that using an MLP for filling a large gap leads to better results, especially when the data behave nonlinearly. Thus, our approach enables the use of datasets that have a large gap in one variable, which is common in many long-term environmental monitoring observations.

Keywords Missing value imputation · Environmental data · Machine learning · Hyperparameter optimization · Derivative-free optimization · Surrogate models

1 Introduction and literature review

Time series data are recorded in diverse application areas ranging from earth sciences to healthcare, finance, traffic, etc. Time series often have missing values (gaps) due to, for example, sensor failures, collection errors, or lack of

resources [1–4]. However, often complete time series datasets are required for analysis or when these data are used as inputs to numerical models. For environmental science applications, such time series gaps could coincide with crucial times where hydrological or biogeochemical fluxes or rates are high. As an example, high stream stage conditions or fluctuations in temperature can promote carbon and nutrient cycling in hyporheic zone environments [5, 6]. Thus, it becomes crucial to impute missing time series values with highly accurate estimates. There are two categories of time series missing value imputation methods, namely multivariate and univariate methods [7, 8]. The first class estimates missing values by exploiting relationships between variables. The second class solely relies on available observations of the time series whose gaps must be filled. In this paper, we study a special case where a large number of consecutive values are missing in a single variable of a multivariate time series.

Many different methods have been used to deal with different characteristics of missing values. For example, for small gaps in univariate time series (individual missing

This research was conducted when Jangho Park and Juliane Müller were at Lawrence Berkeley National Laboratory.

✉ Juliane Müller
Juliane.Mueller@nrel.gov

¹ Artificial Intelligence Team, Teikametrics, 280 Summer St, Boston, MA 02210, USA

² Computational Science Center, National Renewable Energy Laboratory, 15013 Denver W Pkwy, Golden, CO 80401, USA

³ Earth and Environmental Sciences Area, Lawrence Berkeley National Laboratory, Calvin Rd, Berkeley, CA 94705, USA

⁴ Water Security Research Group, International Institute for Applied Systems Analysis, Schloßpl. 1, 2361 Laxenburg, Austria

values), the last observation before the gap can be carried forward or the next observation after the gap backward to estimate missing values. Other approaches, which are implemented in *ImputeTS* [9] provide linear, spline interpolation, or Kalman smoothing methods. Machine learning-based algorithms have also been studied for gap filling purposes. For example, [10] estimated missing values in univariate CO₂ concentrations and air temperature data. This approach, however, is not able to take advantage of multidimensional data.

For the imputation of missing values in the multivariate case, the *K*-nearest neighbor method has been successfully applied for imputing data in the field of medical science [11]. However, this and other such applications have been limited to imputation problems where values are missing only at random. Recently, recurrent neural networks [12, 13] and generative adversarial networks [14, 15] have been used for imputing values in health-care and public air quality datasets. These studies focused exclusively on the performance and accuracy of the proposed approaches for filling individual randomly missing values.

MLPs have also been successfully used for imputing missing time series values across areas of differing complexity and applications. For example, [16, 17] showed that time delayed deep neural network models can impute missing values in univariate hourly traffic volumes. The referenced works focus on filling relatively short gaps (up to 168 data points) and a genetic algorithm was employed for optimizing the hyperparameters of the used gap filling methods. [18] used an MLP to fill missing values in 16 meteorological drivers and three key ecosystem fluxes in the regional Australian and New Zealand flux tower network. The authors analyzed the performance of gap filling methods by manually changing the lags (the number of historical time steps used for making the next time step prediction). To estimate the missing flux values, they used multiple meteorological drivers such as air temperature, wind speed, soil water content, ground heat flux, and net radiation. The architecture of their MLP was determined by using a grid search method provided in scikit-learn [19]. In another study, [20] focused on filling the gaps in methane flux measurements using multiple input variables including net ecosystem exchange (NEE) of CO₂, latent heat flux, sensible heat flux, global radiation, outgoing long wave radiation, air temperature, soil temperature, relative humidity, vapor pressure deficit, friction velocity, air pressure, and water table height. They used an MLP with two hidden layers and determined the number of nodes by using a three-fold cross-validation method.

Unlike the referenced studies, the main contribution of our work is a method for filling a long continuous gap (e.g., multiple continuous months of missing daily observations) in a single variable of a multivariate time series dataset.

Another key difference between the previous studies and our research is the size of the datasets used. Our focus is on smaller datasets that have fewer variables than the cited works. Usually, these types of datasets are hard to use and tend to be discarded in analysis and modeling efforts. Our approach would enable greater use of time series datasets that have multi-year gaps in a variable, which are common in many long-term environmental monitoring observations. In contrast to prior approaches using MLPs, our proposed gap filling method optimizes the MLP model architecture and the length of the lag with an efficient method that is based on surrogate models. Our method requires fewer resources because it is automated and uses adaptive sampling rather than evolutionary or trial-and-error approaches.

In this paper, we present a gap filling approach that is based on the concept of Time-Delayed Deep Neural Networks (TDNN) [21] with the goal of filling a long continuous gap in one variable of a multivariate time series. The concept of TDNN is suitable for handling temporal time series that are a characteristic feature of environmental datasets. This research extends our previous work [22] in which we used deep neural networks to predict groundwater levels for multiple years into the future using supporting variables such as temperature, precipitation, and river flux data. In the cited work, we optimized the network architectures with an automated hyperparameter tuning method, which we adopt and modify in the present work to enable gap filling. We assume that the target time series variable has one large gap and that the supporting variables that explain the target variable are fully observed. Our previous work showed that MLPs are successful in making accurate predictions, and therefore, we use an MLP for gap filling in the present study. Our research will have an impact on application areas where gap-free long-term observations are particularly useful.

The remainder of this paper is organized as follows. Section 2 provides a brief description of MLP models and hyperparameter optimization (HPO) used for gap filling. Section 3 presents the results of our numerical study in which we compare the performance of our proposed method with *ImputeTS* [9] and *mtsdi* [23] for filling gaps of various lengths (three months to three years) of three different environmental datasets. Finally, Sect. 4 concludes the paper and provides future research directions.

2 Multilayer perceptrons and hyperparameter optimization for gap filling

In this section, we provide brief descriptions of MLPs and the hyperparameter optimization method used to determine its architecture.

2.1 Multilayer perceptron (MLP)

Deep neural networks are known to have the ability to approximate nonlinear functions well. In this paper, we use MLPs with nonlinear activation functions for regression purposes to map input to output features. An MLP is a feed-forward artificial neural network. A generic MLP consists of an input layer, at least one hidden layer, and an output layer with each layer containing multiple nodes. Each node in the input and output layers corresponds to an input feature and target variable, respectively. In order to construct an input layer that encodes temporal information, we structure each input such that it contains multiple time steps. We define the amount of historical information in an input as lag l . Thus, with T time periods, an input is defined as $\{x_{t-l}, \dots, x_{t-1}, x_t\}$, where $t = l + 1, \dots, T$. Our goal is to predict the value at the next time step, and thus the output layer has only a single node, which corresponds to the target variable Y that needs to be gap-filled. With this architecture, the MLP is trained by minimizing over a set of weights and biases using a predefined loss function, such as a root-mean-squared error or other metric that reflects how well the MLP approximates the data. The root-mean-squared error is the square root of the average difference between the actual (ground truth) and predicted values. In the first step of this training, linear combinations of input nodes with initial weights are constructed. Each linear combination is passed to all nodes in the first hidden layer. An activation function in the nodes receives this value and returns a transformed value. The transformed values from all nodes in the first hidden layer are given to the nodes in the second hidden layer for further transformation. This process repeats until the output layer returns the final transformed value. The loss function to be minimized reflects the error between the transformed output values and the true values. By iteratively adjusting the weights and biases associated with each node, the loss function is minimized. We use the Rectified Linear Unit (ReLU) activation function, $f(\cdot) = \max(0, \cdot)$ [24], which allows us to approximate nonlinear functions. Our training objective function (loss) is the mean squared error (MSE). The iterative optimization of the weights and biases is performed with the Adam optimizer [25], which is a type of backpropagation algorithm based on gradient descent and the chain rule [26].

Other essential hyperparameters that influence the performance of the MLP, such as the batch size (size of batched samples in the gradient update), the number of epochs (number of iterations used in training), the dropout rate [27], the number of hidden layers and nodes in each layer, and the lag value l , are determined by the automated optimization method described in the next section. We

assume that the learning and decay rates of the stochastic gradient descent method used for training the MLP are known and fixed. For further details about MLPs, we refer the reader to [28].

2.2 Optimization of MLP architectures for gap filling

In order to improve the predictive performance of MLPs (and deep learning models in general), the hyperparameters (model architecture) should be carefully chosen [29]. Different hyperparameter optimization (HPO) methods have been developed and used in the literature, including random search [30], a combination of grid search and trial and error [31], Bayesian optimization [32], and genetic algorithms [33].

In this work, we adapt the HPO method proposed in [22], which interprets the training of a given architecture as an expensive black-box function evaluation of the architecture's performance. Thus, the goal is to find the best hyperparameters within as few trials as possible by employing computationally efficient methods. To this end, surrogate models (in particular radial basis functions and Gaussian process models) are used to map the hyperparameters to their respective performance. An initial experimental design with n_0 different hyperparameter sets is created for which the performance is evaluated by training the respective models. This initial design is used to create the computationally cheap surrogate models. In each iteration of the optimization, an auxiliary optimization problem is solved on the computationally fast to evaluate surrogate model to determine the next hyperparameters to be tried. The surrogate model is updated each time a new hyperparameter set has been evaluated (see Fig. 1 and Reference [22] for more algorithmic details).

In order to adapt the HPO algorithm described in [22] for our gap filling purposes, we only have to change the way we split the training and validation datasets by taking into account the fact that some data are missing. Since we include the lag (the number of historical time steps for predicting the value at the next time step) in our list of MLP hyperparameters to be optimized, we have to create a lagged table of the data for each set of hyperparameters we want to evaluate.

We denote the observation matrix with N variables and time length S as $X = (x_1, x_2, \dots, x_S)^T \in \mathbb{R}^{S \times N}$, where $x_s = (x_s^1, x_s^2, \dots, x_s^N) \in \mathbb{R}^N$ is an observation vector at time s , $1 \leq s \leq S$. Each x_s^d is an observation of the d^{th} variable, where $d = 1, \dots, N$. We assume that the first column in the matrix X corresponds to the target vector Y that contains missing values, and the remaining columns are the supporting variables. In the following, we denote $x_s^d = \text{NaN}$ if

the value is missing. For example, assume we have nine observations ($S = 9$), one target variable, and two supporting variables ($N = 3$), then the observation matrix and target value vector are, respectively,

$$X = \begin{matrix} & 1 & 2 & 3 \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \end{matrix} & \begin{pmatrix} x_1^1 & x_1^2 & x_1^3 \\ x_2^1 & x_2^2 & x_2^3 \\ x_3^1 & x_3^2 & x_3^3 \\ x_4^1 & x_4^2 & x_4^3 \\ \text{NaN} & x_5^2 & x_5^3 \\ \text{NaN} & x_6^2 & x_6^3 \\ x_7^1 & x_7^2 & x_7^3 \\ x_8^1 & x_8^2 & x_8^3 \\ x_9^1 & x_9^2 & x_9^3 \end{pmatrix} & , & Y = \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \end{matrix} \begin{pmatrix} x_1^1 \\ x_2^1 \\ x_3^1 \\ x_4^1 \\ \text{NaN} \\ \text{NaN} \\ x_7^1 \\ x_8^1 \\ x_9^1 \end{pmatrix} .$$

Here, the values for x_5^1 and x_6^1 are missing and the goal is to estimate their values using an optimized MLP. Then, the lagged table for, e.g., lag $l = 1$ becomes

$$\text{Input} = \begin{matrix} & 1 & 2 & 3 & 4 & 5 & 6 \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \end{matrix} & \begin{pmatrix} x_1^1 & x_1^2 & x_1^3 & x_1^4 & x_2^2 & x_2^3 \\ x_2^1 & x_2^2 & x_2^3 & x_3^1 & x_3^2 & x_3^3 \\ x_3^1 & x_3^2 & x_3^3 & x_4^1 & x_4^2 & x_4^3 \\ x_4^1 & x_4^2 & x_4^3 & \text{NaN} & x_5^2 & x_5^3 \\ \text{NaN} & x_5^2 & x_5^3 & \text{NaN} & x_6^2 & x_6^3 \\ \text{NaN} & x_6^2 & x_6^3 & x_7^1 & x_7^2 & x_7^3 \\ x_7^1 & x_7^2 & x_7^3 & x_8^1 & x_8^2 & x_8^3 \end{pmatrix} & , & \text{Output} = \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \end{matrix} \begin{pmatrix} x_3^1 \\ x_4^1 \\ \text{NaN} \\ \text{NaN} \\ x_7^1 \\ x_8^1 \\ x_9^1 \end{pmatrix} .$$

After the lagged tables have been constructed, we determine if they contain any missing values (NaN). Any row that contains at least one NaN entry in the input or output can be removed because the rows in the lagged table are independent. Thus, for the above example, we obtain the following reduced inputs and outputs

$$\text{Reduced Input} = \begin{matrix} & 1 & 2 & 3 & 4 & 5 & 6 \\ \begin{matrix} 1 \\ 2 \\ 7 \end{matrix} & \begin{pmatrix} x_1^1 & x_1^2 & x_1^3 & x_2^1 & x_2^2 & x_2^3 \\ x_2^1 & x_2^2 & x_2^3 & x_3^1 & x_3^2 & x_3^3 \\ x_7^1 & x_7^2 & x_7^3 & x_8^1 & x_8^2 & x_8^3 \end{pmatrix} & , & \text{Reduced Output} = \begin{matrix} 1 \\ 2 \\ 7 \end{matrix} \begin{pmatrix} x_3^1 \\ x_4^1 \\ x_9^1 \end{pmatrix} .$$

Once the reduced lagged input and output are obtained, they are divided into training and validation datasets. For any architecture, the reduced lagged training dataset is used to train the corresponding neural network model. Since this training is done by stochastic gradient descent (and thus the final trained model can be interpreted as the realization of a

stochastic process), we train the model M times, each with a different random number seed. Each of the M trained models is validated against the reduced lagged validation dataset, and the architecture’s performance is calculated as the sample average over all M values.

Finally, we impute the missing values sequentially starting from the first missing value to the last missing value, using the previous estimated missing value to impute the next missing value. In the small example above, we impute the first missing value x_5^1 and denote it as \hat{x}_5^1 by using $(x_3^1, x_3^2, x_3^3, x_4^1, x_4^2, x_4^3)$. Then, the next missing value (x_6^1) is estimated based on \hat{x}_5^1 , i.e., $(x_4^1, x_4^2, x_4^3, \hat{x}_5^1, x_5^2, x_5^3)$. In this way, all missing values are estimated.

3 Numerical experiments

In this section, we demonstrate that our proposed method for gap filling of multivariate time series data is general

enough to be used in different applications that have varying observation frequencies and data characteristics. We describe three time series datasets for which we assume that the target variable has one large gap that must be filled. These time series include observed groundwater levels, simulated soil moisture, and derived net ecosystem exchange from flux tower measurements. Below, we pro-

vide details of the setup of the numerical experiments and alternative gap filling algorithms that we use for comparison.

Fig. 1 Steps of the surrogate model algorithm used to tune the MLP hyperparameters. For further details, see [22]

Input: n_0 (the size of the initial experimental design); n (maximum number of hyperparameter sets to be tried); Ω (hyperparameter space).
Output: Best hyperparameters θ^* found.
Step 1: Create initial experimental design by randomly sampling $\theta_1, \dots, \theta_{n_0}$ from Ω . Prepare the data and train MLP for all n_0 architectures to obtain losses $L_1 = L(\theta_1), \dots, L_{n_0} = L(\theta_{n_0})$.
Step 2: Set $k \leftarrow n_0$; $\mathcal{L} = \{L_1, \dots, L_k\}$; $\mathcal{S} = \{\theta_1, \dots, \theta_k\}$.
Step 3: Build surrogate model using all pairs $\{\theta_i, L_i\}_{i=1}^k$.
Step 4: Use surrogate model to optimize an auxiliary problem and select new hyperparameter set θ_{new} and evaluate $L_{new} = L(\theta_{new})$ by training the MLP for the new architecture θ_{new} .
Step 5: Update $k \leftarrow k + 1$; $\mathcal{L} \leftarrow \mathcal{L} \cup \{L_{new}\}$; $\mathcal{S} \leftarrow \mathcal{S} \cup \{\theta_{new}\}$.
Step 6: If $k < n$: go to Step 3.
Step 6: Else: Stop.

3.1 Description of the datasets

3.1.1 Measured groundwater levels

Our first dataset contained groundwater level time series measured in a groundwater monitoring well in Butte County in California (CA), United States. The monitoring well¹ is 200 m deep from the surface, and provides the groundwater elevation above the mean sea level at daily time resolution. We obtained the data from the California Natural Resources Agency. Our goal was to gap fill the time series of groundwater levels using measurements of the supporting variables - temperature, precipitation and river discharge. The nearest discharge monitoring station (Butte Creek Durham) measures the daily discharge rate at the Butte Creek about 8 km from the well². Temperature and precipitation data were obtained from the Chico weather station located 7 km from the well site³. The dataset used in the study consisted of approximately eight years (2010-2018) of daily observations. For a more detailed description of the dataset and the data sources, we refer the reader to [22, 34]. Figure 2 shows the time series plots for all variables.

3.1.2 Simulated soil moisture

Our second dataset was created with the HYDRUS-1D model [35, 36], a soil hydrologic simulation model. The HYDRUS-1D model is a modular, freely-available, state-of-the-art, and widely used soil hydrologic model with many advanced and coupled water and reactive chemical transport features including equivalent continuum and dual permeability modeling approaches for preferential flow and transport. In particular, the model simulates water, heat and

solute transport in variably saturated and saturated porous media. The model has been extensively applied for modeling across scales—from laboratory cores to watersheds [37, 38].

Here, we used the HYDRUS-1D model to simulate a representative soil column in the Butte County in California to determine the changes in soil moisture profiles over approximately seven years (2011-2018). As a result, our target variable was the daily total water content (TWC) in the upper 80 cm of the simulated soil column. The net amount of water entering the soil column from the top (vTop) and the groundwater recharge were used as the supporting variables for calculations. Figure 3 shows the time series plots for all variables. Because simulation models usually do not lead to missing values, we use it here as an ideal case study to investigate the general applicability of our MLP-based gap filling approach.

3.1.3 Net ecosystem exchange data

Our third example application was the most recently produced FLUXNET dataset, that is FLUXNET2015 dataset. It includes data on CO₂, water, energy exchange, and other meteorological and biological measurements [39, 40]. The eddy-covariance method is used to allow for the non-destructive estimation of fluxes between atmosphere and biosphere for a single site to global scale. FLUXNET datasets have been used in a wide range of research areas ranging from soil microbiology to validation of large-scale earth system models.

The target and supporting variables for the Morgan Monroe State Forest site, Indiana, United States for 16 years (1999-2014) [41] are described in Table 1 and the time series are illustrated in Fig. 4. Note that none of the variables in the studied date range have missing values.

In all three case studies, the supporting variables had been identified with the help of domain science experts, but filter methods, recursive feature elimination, or sensitivity analyses [34] may prove useful to further downselect the most important features. For all test cases, temporal information such as the time of day, the day, and the month

¹ 22N01E28J001M, California Natural Resources Agency, Periodic Groundwater Level Measurements, <https://data.cnra.ca.gov/dataset/periodic-groundwater-level-measurements>.

² California Data Exchange Center, Butte Creek Durham station, <https://cdec.water.ca.gov/webgis/?appid=cdecstation &sta=BCD>.

³ California Data Exchange Center, Chico Station, <https://cdec.water.ca.gov/webgis/?appid=cdecstation &sta=CHI>.

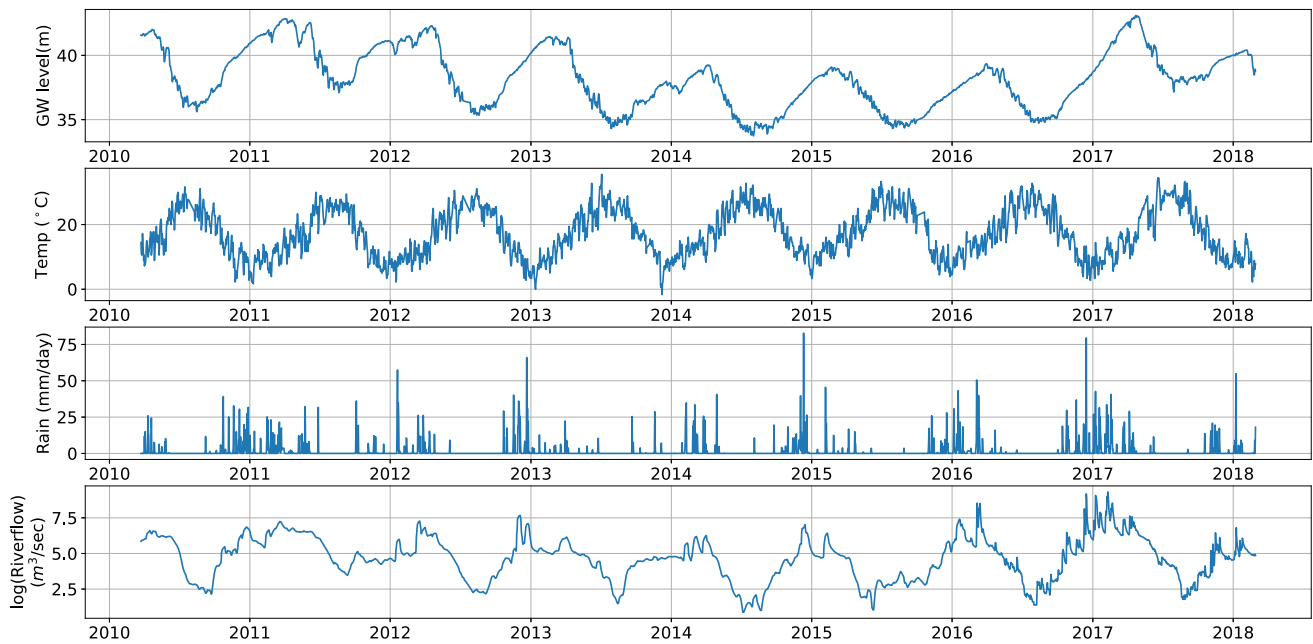


Fig. 2 Time series plots for the groundwater level imputation test case. Shown are (top to bottom) the daily groundwater levels (GW), temperature (Temp), precipitation (Rain), and log-transformed riverflux

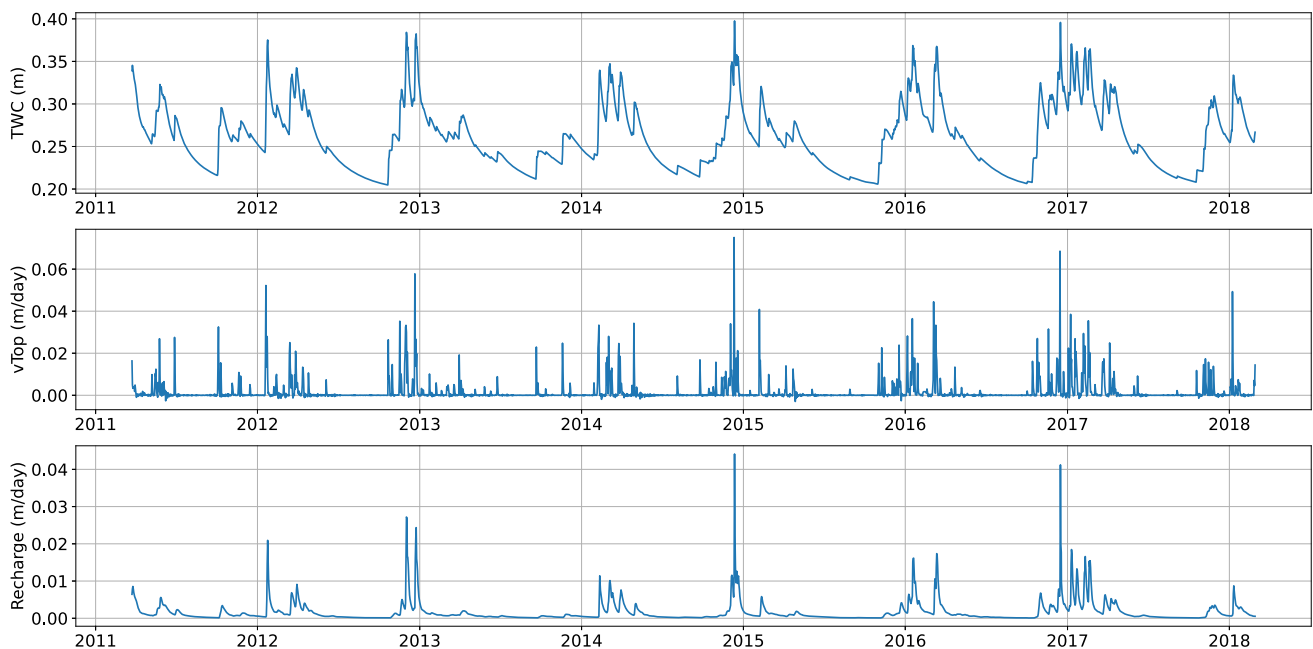


Fig. 3 Time series plots for the soil moisture imputation test case. Shown are (top to bottom) the daily total water content (TWC), the net amount of water entering the soil column from the top (vTop), and the groundwater recharge

of the year the measurements were taken were also included as additional supporting variables. Also, all values of each feature were rescaled with the min-max normalization, i.e., for each variable x^i , to transform all variables to have the same scale. We found the minimum value x_{\min}^i and the maximum value x_{\max}^i , and we rescaled the values according to $\hat{x}^i = \frac{x^i - x_{\min}^i}{x_{\max}^i - x_{\min}^i}$.

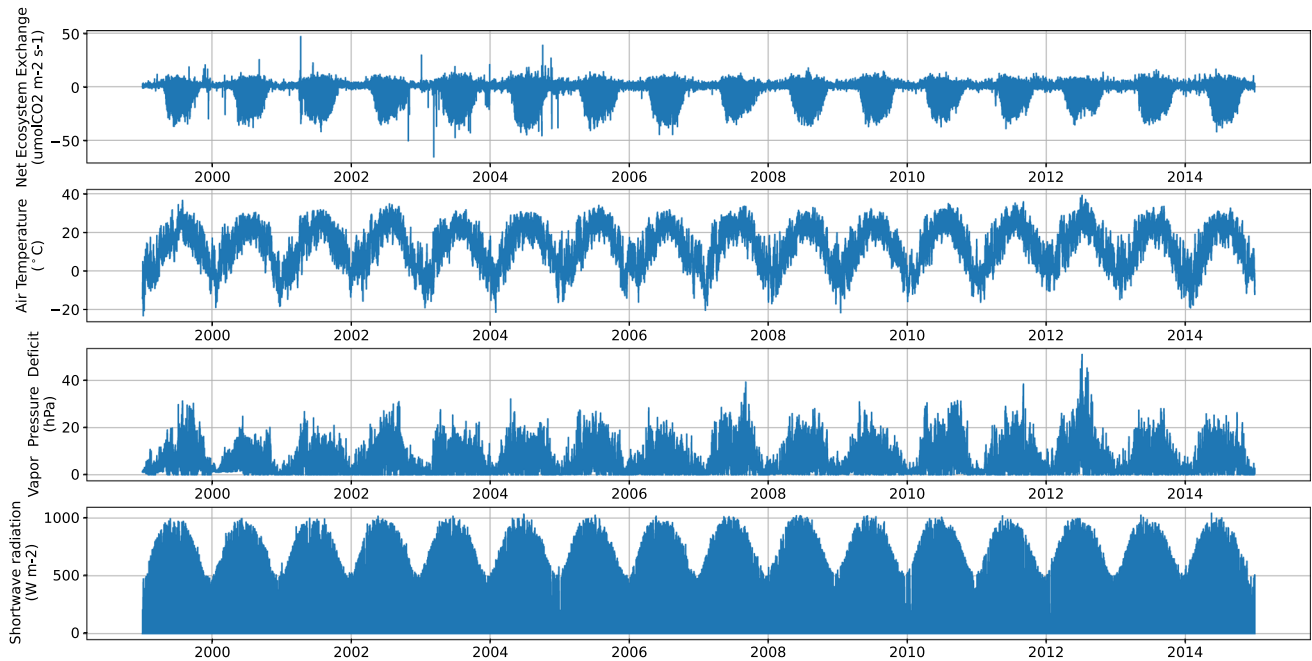
3.2 Setup of numerical experiments

3.2.1 HPO details

We implemented the HPO algorithm for gap filling described in Sect. 2.2 in python (*version 3.7*) using PyTorch (*version 1.4.0*) [42]. All experiments were run on

Table 1 Target and supporting variables with description of FLUXNET2015 dataset

	Variable	Units	Description
Target variable	NEE_VUT_USTAR50	$\mu\text{molCO}_2/\text{m}^2\text{s}$	Net Ecosystem Exchange, using Variable Ustar Threshold (VUT) for each year, from 50 percentile of USTAR threshold
Supporting variables	TA_F	$^{\circ}\text{C}$	Air temperature
	VPD_F	hecto Pascals	Vapor Pressure Deficit
	SW_IN_F	Watt/m^2	Shortwave radiation, incoming

**Fig. 4** Time series plots for the net ecosystem exchange test case. Shown are (top to bottom) the hourly Net Ecosystem Exchange (NEE_VUT_USTAR50), Air Temperature (TA_F), Vapor Pressure Deficit (VPD_F), and incoming shortwave radiation (SW_IN_F)

Ubuntu 16.04 with Intel[®] Xeon(R) CPU E3-1245 v6 @ 3.70GHz \times 8, and 31.2 GB memory. We used the same parameter settings as in [22] and we optimized over the following hyperparameters and their ranges:

- Batch size $\in \{10, 15, \dots, 195, 200\}$,
- Epochs $\in \{50, 100, \dots, 450, 500\}$,
- Layers $\in \{1, 2, \dots, 5, 6\}$,
- Number of hidden nodes $\in \{5, 10, \dots, 45, 50\}$,
- Dropout rate $\in \{0.0, 0.1, \dots, 0.4, 0.5\}$,
- Lags $\in \{30, 35, \dots, 360, 365\}$.

The optimization hyperparameters and their possible values were the same for all three test cases. Note that all parameters were mapped to consecutive integers for the optimization, and there is a finite number of possible combinations (architectures). With the provided options, there are 7,588,800 possible MLP architectures, and it is

computationally intractable to train that many networks, thus motivating the use of automated HPO.

We assumed that each hidden network layer has the same number of hidden nodes and the same dropout rate. We use the MSE as performance measure in tuning the hyperparameters and we split the data into 85% training data and 15% validation data.

In the HPO algorithm, we set $n_0 = 10$ as the initial experimental design size (we used 10 different hyperparameter sets and we trained the models for each to obtain the model performance, which we use to initialize the surrogate model). We stopped the algorithm after $n = 50$ hyperparameter sets had been tried. In order to take into account the stochasticity that arises from using the stochastic gradient descent method for training the models, we computed the hyperparameter performance as the sample average over multiple training trials. For example,

for the daily groundwater level application and the HYDRUS-1D application, we trained the model for each hyperparameter set five times using the same training dataset, and the performance was defined as the average MSE over those trials. Due to the computational expense associated with training the MLP models for the FLUXNET data (because of the large amount of training data), we computed the performance only over one training trial for each hyperparameter set.

3.2.2 Description of gaps in time series data

In order to test our developed gap filling method, we created an artificial continuous gap in the target time series. We studied different instances in which the gaps have varying lengths and are located at different times of the year. By using artificial gaps, we have the ability to compare the gap-filled values to the true held-out observations as well as analyze and compare the performance of different gap filling methods. In all three test cases, the supporting variables did not have gaps.

For the groundwater application, we created gaps of three, six and twelve months of daily observations. For each case, we used 12 different problem instances in which we removed the data of different months of the year, which allowed us to analyze a possible dependence of the performance of the gap filling methods on the trend characteristics of the missing data. For example, for the three month gap, we assumed that groundwater levels were missing (i) from January through March 2012, (ii) from February through April 2012, (iii) from March through May 2012, etc. For the HYDRUS-1D application, we assumed that the total water content values were missing for 12 and 15 months, respectively. For the 12-month gap, we assumed that the total water content values were missing (i) from January 2013 through December 2013, (ii) from February 2013 through January 2014, and (iii) from March 2013 through February 2014. For the 15 months gap, we considered 9 cases where the missing values were starting from the months of April, May, till December. Finally, for the hourly FLUXNET data, we considered between 2,159 hours (3 months) and 35,063 hours (36 months) of missing values in the Net Ecosystem Exchange data (NEE_VUT_USTAR50).

3.2.3 Gap filling algorithms used in the numerical comparison

We compared our proposed gap filling method with state-of-the-art imputation packages, including `ImputeTS` [9] and `mtsdi` [23]. `ImputeTS` is a univariate time series missing value imputation package implemented in R. We used the `na.seadec` function with an interpolation

algorithm. The `na.seadec` function imputes missing values with deseasonalized time series data and adds the seasonality again. It has the option to identify the number of observations before the seasonal pattern repeats. However, in our applications, this feature did not yield satisfactory results, and thus we manually set the daily time series frequency, i.e., the frequency of 365. `mtsdi` is an expectation-maximization algorithm-based multivariate time series missing value imputation package in R. We used the default settings, i.e., the `mnimput` function with the spline method, and we set the spline smooth control to 7. With `mtsdi`, we included indicator variables such as the month, week, and day.

3.3 Results and discussion

In this section we provide a description of the outcomes of our numerical experiments for all three use cases. We present results in the form of tables and illustrations of the gap filled time series.

3.3.1 Groundwater level predictions

Tables 2, 3, and 4 present a comparison of the root-mean-squared error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE) computed between the predicted values and the true (held-out) values for the groundwater use cases with three, six, and twelve months of missing daily data values, respectively. Shown are the three metrics for our proposed MLP-based gap filling method, `ImputeTS`, and `mtsdi`. The RMSE and MAE are given in meters (m), and the MAPE is given in percentage. The lower tables show the optimal hyperparameters that were identified with HPO and led to the MLP outcomes.

For three months of missing daily observations (top section of Table 2), our results show that the best results were achieved by different methods depending on the date range for which the data are missing. For example, for the three months that span November 2012–January 2013 (see also Fig. 5), the groundwater levels show a linear growth.

All gap filling methods captured this increasing trend and `ImputeTS` yielded the best performance in terms of the three comparison metrics. This is due to the fact that `ImputeTS` uses `na.seadec` with an interpolation algorithm. In other words, `ImputeTS` performed seasonally decomposed missing value imputation by linear interpolation which yielded excellent agreement with the held-out data for the missing data segment that had linear behavior. Also the MLP-based gap filling methods were able to capture the linear trend, but the predictions were not as accurate. `mtsdi` introduced strong oscillations around the

Table 2 (a) RMSE, MAE, and MAPE values for the different gap filling methods for *three* months of missing values in daily groundwater levels. (b) Optimal hyperparameters ([batch size, # epochs, # layers, # nodes per layer, dropout rate, # lags]) for both RBF- and GP-based HPO

Missing date range (yyyy-mm-dd)		#obs	RMSE, MAE, MAPE			
From	To		MLP RBF	MLP GP	ImputeTS	mtsdi
(a)						
2012-01-01	2012-03-31	91	0.34, 0.28, 0.69	0.38, 0.29, 0.70	0.58, 0.46, 1.12	1.25, 1.02, 2.50
2012-02-01	2012-04-30	90	0.45, 0.32, 0.75	0.56, 0.43, 1.04	0.48, 0.36, 0.88	1.14, 0.85, 2.03
2012-03-01	2012-05-31	92	0.54, 0.41, 0.99	0.61, 0.50, 1.20	1.00, 0.78, 1.87	1.35, 1.05, 2.54
2012-04-01	2012-06-30	91	0.56, 0.43, 1.06	0.62, 0.49, 1.22	0.66, 0.51, 1.25	1.38, 1.11, 2.73
2012-05-01	2012-07-31	92	0.61, 0.51, 1.33	0.62, 0.52, 1.36	0.63, 0.50, 1.32	1.16, 0.99, 2.62
2012-06-01	2012-08-31	92	0.42, 0.32, 0.87	0.41, 0.33, 0.90	0.38, 0.32, 0.86	1.60, 1.50, 4.12
2012-07-01	2012-09-30	92	0.26, 0.22, 0.61	0.33, 0.26, 0.73	0.29, 0.24, 0.67	1.74, 1.71, 4.74
2012-08-01	2012-10-31	92	0.35, 0.27, 0.73	0.35, 0.27, 0.74	0.47, 0.40, 1.12	1.56, 1.47, 4.08
2012-09-01	2012-11-30	91	0.49, 0.39, 1.04	0.80, 0.62, 1.62	0.28, 0.22, 0.60	1.03, 0.87, 2.37
2012-10-01	2012-12-31	92	0.80, 0.71, 1.85	0.93, 0.86, 2.23	0.25, 0.20, 0.52	0.74, 0.61, 1.59
2012-11-01	2013-01-31	92	0.74, 0.70, 1.77	0.75, 0.73, 1.83	0.23, 0.19, 0.49	0.79, 0.68, 1.72
2012-12-01	2013-02-28	90	0.69, 0.63, 1.54	0.67, 0.56, 1.36	0.55, 0.47, 1.16	0.97, 0.86, 2.10

Missing date range (yyyy-mm-dd)		#obs	Hyperparameters	
From	To		MLP RBF	MLP GP
(b)				
2012-01-01	2012-03-31	91	145, 150, 1, 45, 0, 330	190, 450, 4, 50, 0, 270
2012-02-01	2012-04-30	90	120, 400, 1, 25, 0, 250	95, 450, 3, 40, 0, 210
2012-03-01	2012-05-31	92	195, 400, 1, 45, 0, 250	170, 350, 2, 20, 0, 290
2012-04-01	2012-06-30	91	175, 450, 1, 40, 0, 240	70, 500, 1, 25, 0, 345
2012-05-01	2012-07-31	92	130, 450, 4, 45, 0, 330	170, 350, 1, 30, 0, 240
2012-06-01	2012-08-31	92	170, 250, 1, 45, 0, 265	70, 500, 1, 25, 0, 345
2012-07-01	2012-09-30	92	170, 250, 1, 45, 0, 265	70, 500, 1, 25, 0, 345
2012-08-01	2012-10-31	92	170, 250, 1, 45, 0, 265	70, 500, 1, 25, 0, 345
2012-09-01	2012-11-30	91	125, 400, 2, 45, 0, 230	50, 150, 4, 50, 0, 290
2012-10-01	2012-12-31	92	170, 250, 1, 45, 0, 265	70, 500, 1, 25, 0, 345
2012-11-01	2013-01-31	92	170, 250, 1, 45, 0, 265	70, 500, 1, 25, 0, 345
2012-12-01	2013-02-28	90	135, 450, 3, 45, 0, 305	110, 200, 1, 10, 0, 255

Bold numbers indicate the **best value** for a given problem instance. The column “#obs” indicates the number of missing observations. “MLP RBF” and “MLP GP” indicate the methods using the MLP with the RBF and GP surrogate during HPO, respectively

overall growing trend. Moreover, *mtsdi* performed overall the worst for all date ranges for the three month test case.

In contrast, Figs. 6 and 7 show scenarios where the data are behaving nonlinearly. In Fig. 6, the interpolation used by *imputeTS* failed to capture the trend and produced an almost flat prediction. Both MLP-based methods captured the nonlinear trend in the missing data better, yet the deviation over the last ~ 3 weeks was larger. *mtsdi* made again predictions with large oscillations around the true data, and for the last 6 weeks of missing data, it deviated significantly from the truth. In Fig. 7, *mtsdi* shows again an oscillating pattern, whereas *imputeTS* underestimated the

groundwater level. The MLP solutions are closest to reality. Regarding the optimal MLP architectures that correspond to the best solutions, Table 2 shows that there are only few differences. Most optimal architectures have only 1 layer and 40-45 nodes in that layer, often with a relatively large number of training epochs and no dropout. Similarly, the number of lags used to make predictions are 240 days or more, indicating the importance of at least 8 months worth of data points to make accurate predictions for the gap filling. We can also see that there are a few architectures that are more complex with more layers but that achieve a somewhat similar solution quality as the smaller

Table 3 (a) RMSE, MAE, and MAPE values for the different gap filling methods for six months of missing values in daily groundwater levels. (b) Optimal hyperparameters ([batch size, # epochs, # layers, # nodes per layer, dropout rate, # lags]) for both RBF- and GP-based HPO

Missing date range (yyyy-mm-dd)		#obs	RMSE, MAE, MAPE			
From	To		MLP RBF	MLP GP	ImputeTS	mtsdi
(a)						
2012-01-01	2012-06-30	181	0.54, 0.42, 1.04	0.48, 0.37, 0.90	0.62, 0.46, 1.13	1.21, 0.96, 2.37
2012-02-01	2012-07-31	181	0.44, 0.36, 0.91	0.45, 0.37, 0.94	0.60, 0.42, 1.04	1.12, 0.89, 2.26
2012-03-01	2012-08-31	183	0.46, 0.37, 0.94	0.47, 0.38, 0.98	0.82, 0.58, 1.45	1.31, 1.12, 2.93
2012-04-01	2012-09-30	182	0.48, 0.40, 1.05	0.91, 0.75, 2.00	0.55, 0.44, 1.15	1.50, 1.36, 3.62
2012-05-01	2012-10-31	183	0.57, 0.42, 1.12	0.60, 0.46, 1.23	1.05, 0.95, 2.58	1.63, 1.49, 4.07
2012-06-01	2012-11-30	182	0.54, 0.40, 1.07	0.70, 0.60, 1.63	0.76, 0.65, 1.80	1.64, 1.48, 4.06
2012-07-01	2012-12-31	183	0.80, 0.69, 1.81	0.77, 0.65, 1.71	0.70, 0.61, 1.66	1.43, 1.23, 3.37
2012-08-01	2013-01-31	183	0.64, 0.53, 1.36	0.55, 0.45, 1.16	0.53, 0.46, 1.22	1.17, 0.98, 2.65
2012-09-01	2013-02-28	180	0.44, 0.37, 0.93	1.12, 1.02, 2.57	0.45, 0.38, 0.95	0.94, 0.82, 2.11
2012-10-01	2013-03-31	181	0.59, 0.51, 1.27	1.03, 0.98, 2.44	1.17, 1.08, 2.68	1.19, 1.00, 2.50
2012-11-01	2013-04-30	180	0.85, 0.80, 2.00	0.78, 0.72, 1.77	1.67, 1.46, 3.57	1.61, 1.40, 3.44
2012-12-01	2013-05-31	181	0.78, 0.61, 1.55	0.81, 0.62, 1.58	1.09, 0.93, 2.28	1.63, 1.42, 3.52

Missing date range (yyyy-mm-dd)		#obs	Hyperparameters	
From	To		MLP RBF	MLP GP
(b)				
2012-01-01	2012-06-30	181	145, 400, 5, 45, 0, 265	185, 450, 1, 40, 0, 250
2012-02-01	2012-07-31	181	140, 400, 2, 45, 0, 310	50, 200, 3, 40, 0, 310
2012-03-01	2012-08-31	183	190, 500, 1, 50, 0, 325	190, 450, 4, 50, 0, 270
2012-04-01	2012-09-30	182	155, 400, 2, 45, 0, 245	200, 400, 4, 45, 0, 285
2012-05-01	2012-10-31	183	195, 450, 1, 40, 0, 335	50, 250, 1, 45, 0, 320
2012-06-01	2012-11-30	182	90, 400, 1, 45, 0, 310	145, 500, 4, 20, 0, 245
2012-07-01	2012-12-31	183	85, 250, 1, 45, 0, 305	70, 500, 1, 25, 0, 345
2012-08-01	2013-01-31	183	195, 450, 1, 35, 0, 260	160, 400, 1, 35, 0, 245
2012-09-01	2013-02-28	180	65, 450, 4, 45, 0, 215	60, 400, 4, 25, 0, 285
2012-10-01	2013-03-31	181	195, 450, 1, 30, 0, 230	155, 400, 2, 35, 0, 305
2012-11-01	2013-04-30	180	120, 300, 1, 25, 0, 275	60, 250, 4, 50, 0, 255
2012-12-01	2013-05-31	181	145, 400, 5, 45, 0, 265	80, 150, 2, 35, 0, 255

Bold numbers indicate the **best value** for a given problem instance. The column “#obs” indicates the number of missing observations. “MLP RBF” and “MLP GP” indicate the methods using the MLP with the RBF and GP surrogate during HPO, respectively

architectures (e.g., missing date range 2012-01-01 to 2012-03-31). This indicates that different architectures may lead to similar results. In this case, choosing the smaller architecture may be beneficial because they tend to train faster.

The results for the case where six months of daily observations were missing are summarized in Table 3. Except for two cases (missing values from July 2012 to December 2012 and from August 2012 to January 2013), our proposed approach with either RBF or GP made more accurate predictions than imputeTS and mtsdi in terms of the comparison metrics. Similarly to the three-month

missing data case, architectures of different complexities can lead to similar performance, e.g., missing date range 2012-03-01 to 2012-08-31. The optimal values for the number of lags shows again that more than 200 days are needed to make good gap-filling predictions. This indicates that information about the ups and downs of groundwater levels over the course of several months are needed for good gap filling performance.

Figure 8 shows a comparison between all methods with missing values ranging from November 2012 to April 2013.

Table 4 (a) RMSE, MAE, and MAPE values for the different gap filling methods for *twelve* months of missing values in daily groundwater levels. (b) Optimal hyperparameters ([batch size, # epochs, # layers, # nodes per layer, dropout rate, # lags]) for both RBF- and GP-based HPO

Missing date range (yyyy-mm-dd)		#obs	RMSE, MAE, MAPE			
From	To		MLP RBF	MLP GP	ImputeTS	mtsdi
(a)						
2012-01-01	2012-12-31	366	0.91, 0.75, 1.95	0.68, 0.55, 1.42	1.03, 0.90, 2.37	1.52, 1.30, 3.42
2012-02-01	2013-01-31	366	0.71, 0.56, 1.44	0.74, 0.62, 1.59	1.13, 0.96, 2.54	1.32, 1.11, 2.93
2012-03-01	2013-02-28	365	0.74, 0.61, 1.54	0.71, 0.57, 1.50	0.77, 0.66, 1.71	1.26, 1.06, 2.79
2012-04-01	2013-03-31	365	0.82, 0.70, 1.78	0.73, 0.58, 1.53	0.91, 0.78, 2.05	1.28, 1.11, 2.88
2012-05-01	2013-04-30	365	0.71, 0.53, 1.40	1.10, 0.92, 2.45	1.35, 1.13, 2.87	1.35, 1.17, 3.02
2012-06-01	2013-05-31	365	0.64, 0.50, 1.29	0.62, 0.47, 1.24	1.05, 0.81, 2.05	1.36, 1.17, 3.06
2012-07-01	2013-06-30	365	0.76, 0.62, 1.62	0.78, 0.65, 1.70	1.09, 0.83, 2.09	1.31, 1.12, 2.92
2012-08-01	2013-07-31	365	0.86, 0.70, 1.84	0.92, 0.82, 2.12	1.15, 0.89, 2.23	1.25, 1.05, 2.73
2012-09-01	2013-08-31	365	0.85, 0.75, 1.97	0.95, 0.86, 2.24	1.01, 0.83, 2.11	1.20, 1.01, 2.60
2012-10-01	2013-09-30	365	0.88, 0.79, 2.07	0.88, 0.81, 2.11	1.17, 1.08, 2.81	1.19, 0.98, 2.54
2012-11-01	2013-10-31	365	0.82, 0.73, 1.93	0.85, 0.72, 1.92	0.87, 0.70, 1.81	1.19, 0.97, 2.49
2012-12-01	2013-11-30	365	0.81, 0.67, 1.80	1.02, 0.82, 2.22	0.84, 0.71, 1.83	1.17, 0.93, 2.40

Missing date range (yyyy-mm-dd)		#obs	Hyperparameters	
From	To		MLP RBF	MLP GP
(b)				
2012-01-01	2012-12-31	366	80, 500, 2, 35, 0, 355	70, 500, 1, 25, 0, 345
2012-02-01	2013-01-31	366	155, 450,4, 45,0, 350	70, 500,1, 25,0, 345
2012-03-01	2013-02-28	365	160, 350,1, 40,0, 305	70, 300,2, 15,0, 260
2012-04-01	2013-03-31	365	85, 450,1, 40,0, 320	120, 400,4, 25,0, 270
2012-05-01	2013-04-30	365	185, 450,1, 45,0, 255	185, 100,2, 45,0, 270
2012-06-01	2013-05-31	365	185, 450,1, 45,0, 255	140, 500,1, 15,0, 235
2012-07-01	2013-06-30	365	185, 450,1, 45,0, 255	85, 400,5, 25,0, 240
2012-08-01	2013-07-31	365	110, 400,1, 45,0, 260	170, 300,2, 40,0, 315
2012-09-01	2013-08-31	365	145, 450,5, 40,0, 265	170, 300,2, 40,0, 315
2012-10-01	2013-09-30	365	145, 450,5, 40,0, 265	140, 300,1, 45,0, 260
2012-11-01	2013-10-31	365	145, 450,5, 40,0, 265	195, 500,5, 35,0, 255
2012-12-01	2013-11-30	365	185, 450,1, 45,0, 255	70, 500,1, 25,0, 345

Bold numbers indicate the **best value** for a given problem instance. The column “#obs” indicates the number of missing observations. “MLP RBF” and “MLP GP” indicate the methods using the MLP with the RBF and GP surrogate during HPO, respectively

We can see that our proposed method captured the high peak better than the other methods and *imputeTS* made again a fairly flat prediction. Similarly to the cases of three months of continuous gaps, *mtsdi* provided the lowest accuracy. Note that the increased size of the continuous gap allowed for a larger probability that seasonality effects will be included in the missing value ranges.

The results for the case where 12 months of daily observations were missing are summarized in Table 4. For all test cases, the MLP was able to fill the gaps better than *imputeTS* and *mtsdi*. For most cases, the performance of *imputeTS* and *mtsdi* does not come close to MLP. From

the optimal architectures (lower table) we can see again that different model complexities can lead to similar results, e.g. missing values from 2012-10-01 to 2013-09-30. Whether the possibility of small improvements is worth a significantly larger amount of training time depends in the end on the user and their insights on how these small differences will effect follow-on analyses. Similarly to the previous cases, the lag should comprise at least several months worth of data for achieving good MLP performance.

Figure 9 shows that the MLP-based gap filling method captured the nonlinearities in the missing data significantly

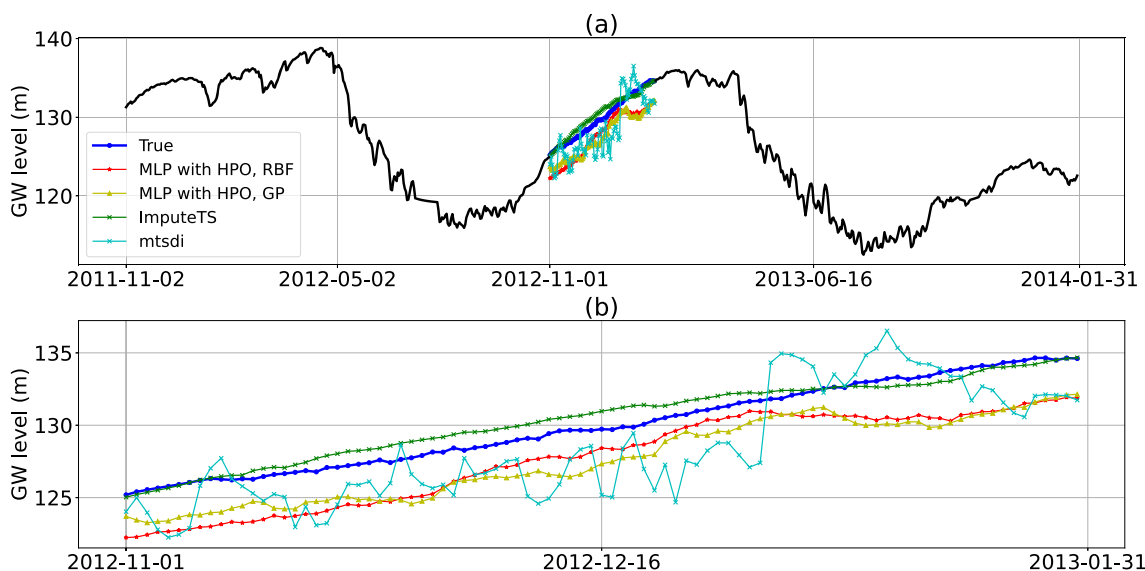


Fig. 5 Daily groundwater level imputation for *three* months of missing daily measurement data (November 2012 - January 2013). The top figure shows the missing time series in the context of

available data, and the bottom figure shows a zoom onto the missing period. The missing data have a linear trend which is captured by *imputeTS* and the MLP-based methods

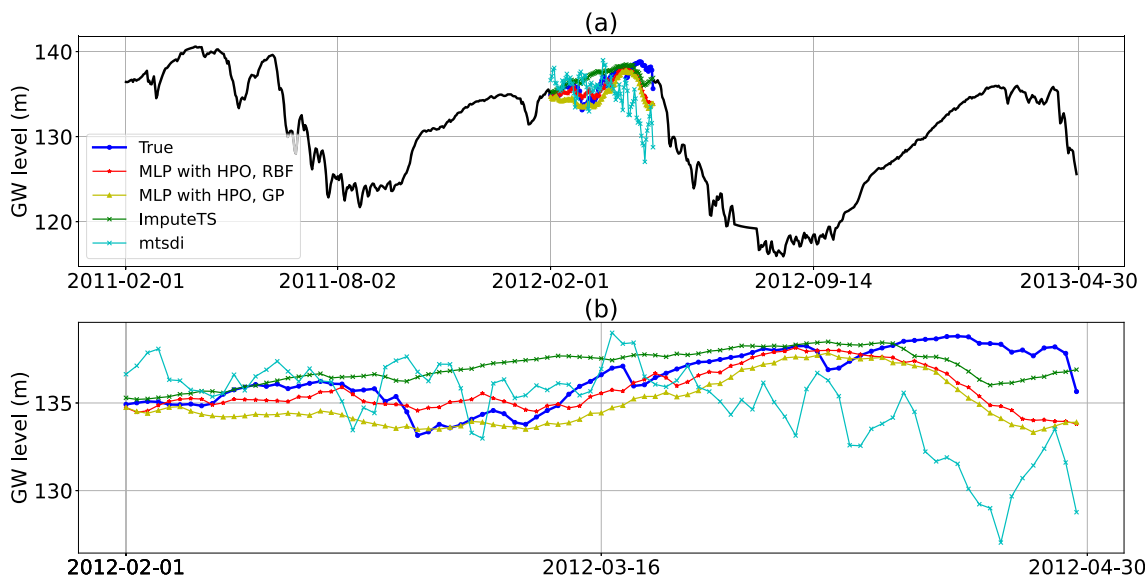


Fig. 6 Daily Groundwater level imputation for *three* months of missing daily measurement data (February 2012–April 2012). The top figure shows the missing time series in context of available data, and

the bottom figure shows a zoom onto the missing period. The missing data have a nonlinear trend and non-interpolating methods like our MLP approach perform better for these instances

better than *mtsdi* (which introduced strong oscillations and did not capture the extrema of the missing values) and *imputeTS* (which failed to capture the winter data (November 2012–April 2013) and appeared like a piecewise linear approximation of the missing data). The MLP-based imputation method showed stronger deviations from the true data for the last ~ 1.5 months of gap filling, which is possibly due to the fact that our method estimates the missing values sequentially based on the previously

estimated values, and thus it is possible that errors started to accumulate.

Comparing the optimal architectures across Tables 2, 3, and 4 shows that the optimal hyperparameters for the MLP depended on the date range for which the data were missing, and thus the same hyperparameters may not be optimal for all date ranges making tuning necessary. Some “preferences” toward architectures with only one layer and lags larger than 200 exist, indicating that the search range over these hyperparameter could potentially be decreased.

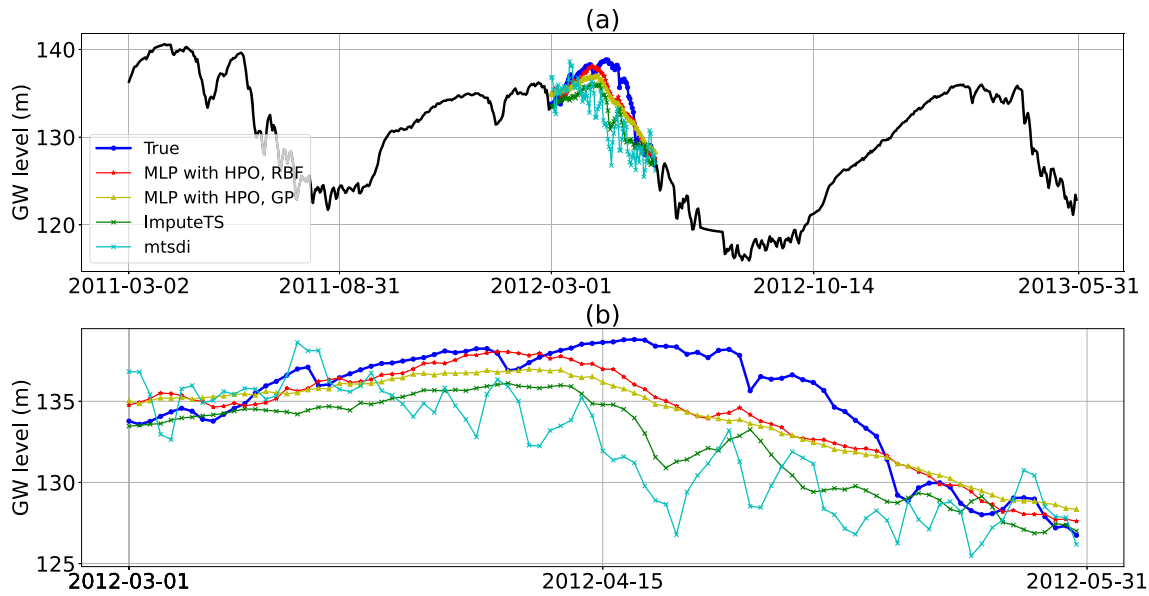


Fig. 7 Daily Groundwater level imputation for *three* months of missing daily measurement data (March 2012–May 2012). The top figure shows the missing time series in context of available data, and the bottom figure shows a zoom onto the missing period. The missing data have a nonlinear trend and non-interpolating methods like our MLP approach perform better for these instances

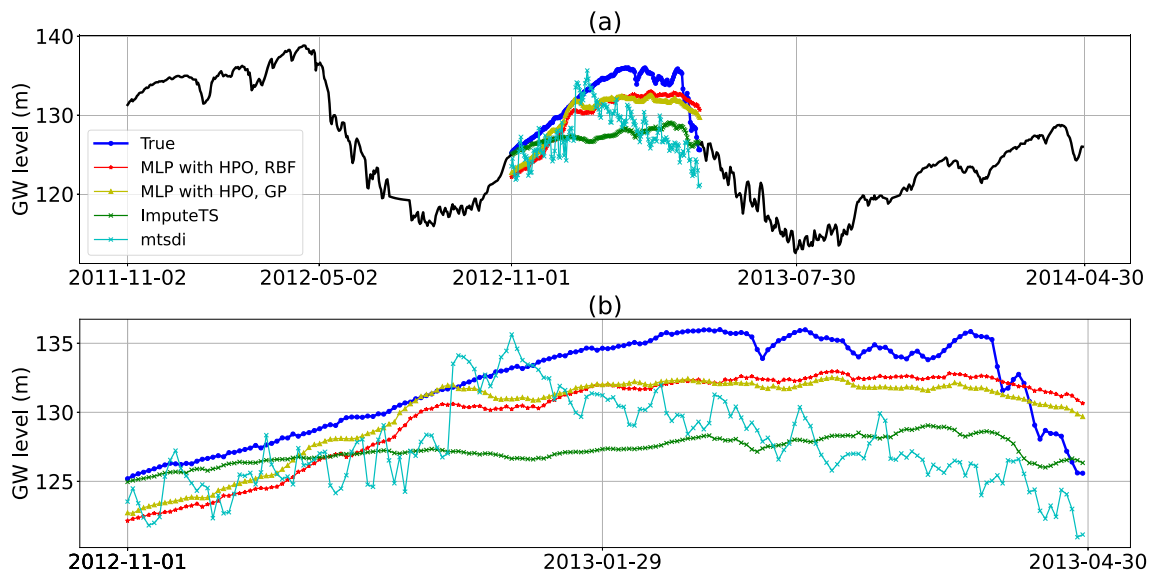


Fig. 8 Daily Groundwater level imputation for *six* months of missing daily measurement data (November 2012–April 2013). The top figure shows the missing time series in context of available data, and the bottom figure shows a zoom onto the missing period

However, a detailed sensitivity analysis of the MLP performance with respect to the hyperparameters and their search ranges is beyond the scope of this paper. The MLP performance in terms of RMSE, MAE, and MAPE for a specific architecture was impacted by the data that were missing. For example, in Table 2, we see that the same optimal hyperparameters were obtained for the date ranges June 2012–August 2012, July 2012–September 2012, August 2012–October 2012, October 2012–December 2012, and November 2012–January 2013 when using either the

RBF or the GP during hyperparameter tuning, respectively. However, the corresponding RMSE metrics ranged between 0.26 and 0.8 for RBF, and between 0.33 and 0.93 for the GP. We observe a similar behavior for the cases of six and twelve months of missing data.

3.3.2 Total water content predictions

Tables 5 and 6 show the results of our data imputation for the TWC simulated with HYDRUS-1D. Shown are the

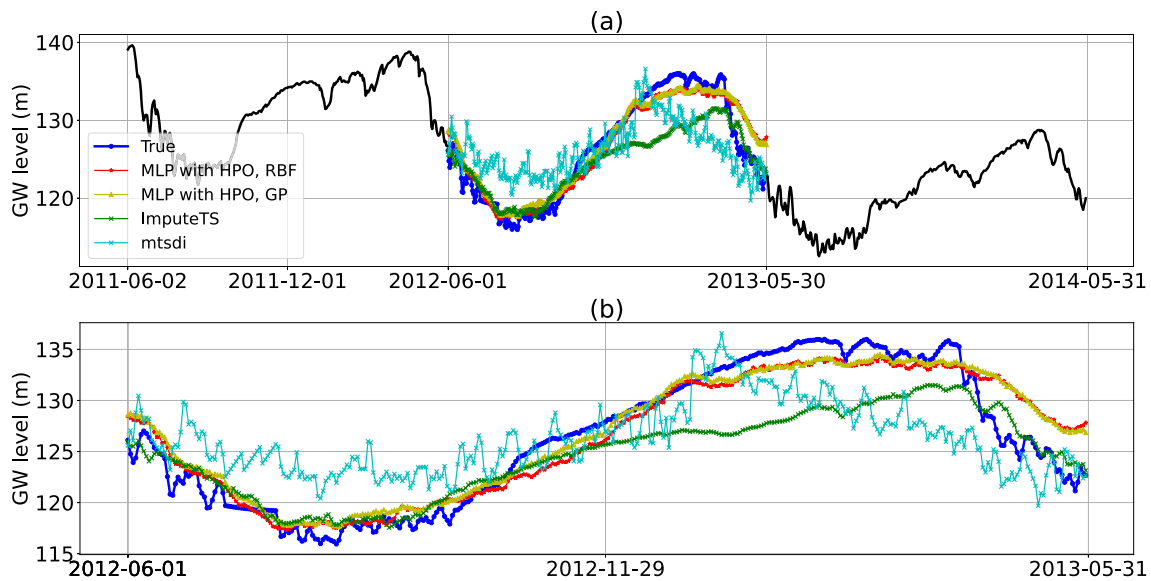


Fig. 9 Daily Groundwater level imputation for twelve months of missing daily measurement data (June 2012–May 2013). The top figure shows the missing time series in context of available data, and

the bottom figure shows a zoom onto the missing period. Although all imputation methods try to capture the nonlinearities, the MLP-based imputation agrees best with the true values

RMSE, MAE, and MAPE values between the true and estimated values obtained with the different gap filling methods for the 12- and 15-month cases. The optimal hyperparameters for the MLP-based approaches are also shown. Our MLP-based gap filling method achieved better performance than *ImputeTS* for all three of the tested 12-month date ranges, but it was outperformed by *mtsdi* for two of the test cases. Notably, for the two cases where *mtsdi* outperformed the MLP, the largest indicator for MLP performance is the number of lags, which were significantly larger for the two outperformed cases. This may indicate that a too large lag may be detrimental to predictive performance. Compared to the groundwater level time series, the TWC time series has several pronounced small scale features (peaks and troughs) in addition to larger scale trends. Basing predictions on a larger amount of data may not be the best course of action for TWC. In fact, we observe the same pattern (small lags lead to good performance, large lags lead to bad performance) also in the 15-month missing data case in Table 6. Again, the MLP-based gap filling method outperformed *imputeTS* for all instances, but it was outperformed by *mtsdi* for most cases. The two instances for which the MLP performed better than *mtsdi* had a small lag (30 and 120, respectively).

Figure 10 shows the true simulated TWC values and the values as estimated by the gap filling methods for a 12-month case.

We can see that all methods had difficulties capturing the true data, and our proposed method and *imputeTS* did a better job than *mtsdi* at capturing the first seven months

of missing data, but then both methods generated larger errors during the remaining five months. In contrast, *mtsdi* did not capture the trend in the data of the first seven months well, but the resulting RMSE is smaller overall regardless.

Figure 11 shows the results for the gap September 2013–November 2014 (a case in which the MLP-based method performed better than the other methods).

We can see that the MLP that used the RBF in the HPO led to predictions that almost perfectly matched the true data. *mtsdi* appeared to be able to capture some of the trends in the data, in particular the large peaks with higher frequency oscillations, but it failed to approximate the “less rugged” and lower values well. We can also observe from Fig. 11 how important the choice of hyperparameters is for the MLP. The MLP using the GP during HPO did not perform nearly as well as the MLP that used the RBF during HPO. The main differences between both architectures were in the batch size (50 vs. 70), the number of nodes per layer (50 vs. 25), and the number of lags (30 vs. 345). On the other hand, in Fig. 10, the GP-based and RBF-based MLP predictions were approximately equal, and for this case the optimal architectures differed only with respect to the number of nodes per layer (45 vs. 25), with the larger architecture (RBF solution) being marginally better than the GP-based solution. Thus, we can recognize the importance of HPO when using DL models and the sensitivity of the MLP’s performance with regard to the architecture.

Table 5 (a) RMSE, MAE, and MAPE for the different gap filling methods for 12 months of missing values in daily TWC. (b) Optimal hyperparameters ([batch size, # epochs, # layers, # nodes per layer, dropout rate, # lags]) for both RBF- and GP-based HPO

Missing date range (yyyy-mm-dd)		#obs	RMSE, MAE, MAPE			
From	To		MLP RBF	MLP GP	ImputeTS	mtsdi
(a)						
2013-01-01	2013-12-31	365	0.008, 0.007, 2.651	0.012, 0.008, 3.302	0.031, 0.025, 10.006	0.016, 0.013, 5.299
2013-02-01	2014-01-31	365	0.025, 0.019, 7.607	0.023, 0.018, 7.303	0.045, 0.041, 16.860	0.018, 0.015, 6.543
2013-03-01	2014-02-28	365	0.023, 0.017, 6.784	0.023, 0.017, 6.857	0.031, 0.022, 8.686	0.019, 0.016, 6.702
Missing date range (yyyy-mm-dd)		#obs	Hyperparameters			
From	To		MLP RBF	MLP GP		
(b)						
2013-01-01	2013-12-31	365	50, 500, 1, 50, 0, 45	55, 300, 6, 50, 0, 55		
2013-02-01	2014-01-31	365	60, 500, 1, 35, 0, 355	70, 500, 1, 25, 0, 345		
2013-03-01	2014-02-28	365	70, 500, 1, 45, 0, 345	70, 500, 1, 25, 0, 345		

Bold numbers indicate the **best value** for a given problem instance. The column “#obs” indicates the number of missing observations. “MLP RBF” and “MLP GP” indicate the methods using the MLP with the RBF and GP model during HPO, respectively

3.3.3 FLUXNET predictions

Table 7 shows the results for the FLUXNET2015 dataset. Note that we show the symmetric mean absolute percentage error (SMAPE) rather than MAPE because the true target variable, NEE VUT USTAR50, includes values close to zero. Because the number of available observations was large, we stress-tested the performance of the gap filling methods for very large continuous data gaps, using between 2,159 (3 months) and 35,063 (36 months) missing observations. The results show that the MLP-based gap filling method attained generally better performance than *imputeTS* and *mtsdi*. Unlike the previous two case studies, the optimal hyperparameters do not show a pattern in terms of the best MLP performance. In fact, the three metrics (RMSE, MAE, and SMAPE) are fairly similar for MLP RBF and MLP GP, but the corresponding architectures differ a lot. For example, for the case where 2,159 observations were missing, the difference in RMSEs was only 0.02, but the MLP GP used an additional hidden layer and a significantly larger number of lags than the MLP RBF. Thus, for the FLUXNET use case, there may be nonlinear interactions between hyperparameters present, and thus multiple different architectures are able to achieve similar performance. A detailed study of these hyperparameter interactions and performance sensitivities is, however, out of the scope of this paper.

Figure 12 shows the results for filling the 8,759 missing value gap of NEE_VUT_USTAR50 with the different methods. The top panel shows the true and the estimated values of NEE_VUT_USTAR50 in the context of the

longer time series, the middle panel zooms in on the filled gap, and the lower panels show scatter plots of the true values versus the predicted values. If the methods were able to exactly repredict the missing values, all points would lie on the diagonal. The bottom panel of Fig. 12 clearly shows that *imputeTS* completely failed to capture the variability in the data and its predictions are almost a constant value. The MLP-based methods on the other hand made predictions that were close to the true values (points lie close to the diagonal). The point cloud obtained with *mtsdi* shows a larger variability around the diagonal, and for this use case *mtsdi* made less reliable predictions.

Figure 13 plots the results for the case when 17,543 observations are missing. Again, *imputeTS* failed to capture the trend in the data completely. *mtsdi* captured the seasonal trend but it failed to accurately predict the large and small values. Our proposed MLP-based imputation methods estimated the high and low values significantly better. This is also reflected in the scatter plots which show that the point clouds of our proposed method lie closer to the diagonal than those of *imputeTS* and *mtsdi*.

3.3.4 Discussion

The three gap filling case studies showed that there was no one single method that performed best for all problems and all instances. We found that for the two examples with real observation data (as opposed to simulation-created data), our proposed gap filling method tended to outperform both *imputeTS* and *mtsdi* in terms of RMSE, MAE and MAPE

Table 6 (a) RMSE, MAE, and MAPE for the different gap filling methods for 15 months of missing values in daily TWC. (b) Optimal hyperparameters ([batch size, # epochs, # layers, # nodes per layer, dropout rate, # lags]) for both RBF- and GP-based HPO

Missing date range (yyyy-mm-dd)		#obs	RMSE, MAE, MAPE			
From	To		MLP RBF	MLP GP	ImputeTS	mtsdi
(a)						
2013-04-01	2014-06-30	456	0.020, 0.015, 6.074	0.021, 0.016, 6.239	0.029, 0.020, 7.334	0.017, 0.015, 5.834
2013-05-01	2014-07-31	457	0.024, 0.017, 6.678	0.020, 0.015, 5.791	0.028, 0.019, 6.914	0.018, 0.016, 6.553
2013-06-01	2014-08-31	457	0.023, 0.017, 6.575	0.021, 0.015, 5.780	0.028, 0.018, 6.926	0.019, 0.017, 6.912
2013-07-01	2014-09-30	457	0.024, 0.017, 6.824	0.021, 0.015, 6.097	0.028, 0.021, 8.332	0.020, 0.017, 7.300
2013-08-01	2014-10-31	457	0.024, 0.018, 7.096	0.021, 0.016, 6.348	0.027, 0.019, 7.169	0.020, 0.017, 7.167
2013-09-01	2014-11-30	456	0.005, 0.003, 1.31	0.024, 0.019, 7.430	0.030, 0.021, 8.100	0.018, 0.015, 6.245
2013-10-01	2014-12-31	457	0.028, 0.022, 8.366	0.028, 0.020, 8.055	0.039, 0.030, 11.822	0.023, 0.017, 6.32
2013-11-01	2015-01-31	457	0.029, 0.022, 8.518	0.027, 0.019, 7.560	0.050, 0.041, 15.037	0.023, 0.017, 6.230
2013-12-01	2015-02-28	455	0.014, 0.010, 3.613	0.020, 0.016, 5.882	0.031, 0.022, 7.822	0.023, 0.017, 6.196

Missing date range (yyyy-mm-dd)		#obs	Hyperparameters	
From	To		MLP RBF	MLP GP
(b)				
2013-04-01	2014-06-30	456	50, 500, 1, 45, 0, 345	70, 500, 1, 25, 0, 345
2013-05-01	2014-07-31	457	195, 450, 1, 40, 0, 335	70, 500, 1, 25, 0, 345
2013-06-01	2014-08-31	457	195, 450, 1, 40, 0, 335	70, 500, 1, 25, 0, 345
2013-07-01	2014-09-30	457	195, 450, 1, 40, 0, 335	70, 500, 1, 25, 0, 345
2013-08-01	2014-10-31	457	195, 450, 1, 40, 0, 335	70, 500, 1, 25, 0, 345
2013-09-01	2014-11-30	456	50, 500, 1, 50, 0, 30	70, 500, 1, 25, 0, 345
2013-10-01	2014-12-31	457	195, 450, 1, 40, 0, 335	70, 500, 1, 25, 0, 345
2013-11-01	2015-01-31	457	195, 450, 1, 40, 0, 335	70, 500, 1, 25, 0, 345
2013-12-01	2015-02-28	455	50, 450, 1, 45, 0, 120	70, 500, 1, 25, 0, 345

Bold numbers indicate the **best value** for a given problem instance. The column “#obs” indicates the number of missing observations. “MLP RBF” and “MLP GP” indicate the methods using the MLP with the RBF and GP model during HPO, respectively

(or SMAPE) in particular when the gaps were large. As described in the previous subsections, *ImputeTS* performed well when the missing data can be approximated by a somewhat (piecewise) linear function. However, for the FLUXNET2015 data, *ImputeTS* failed, which may be due to the amount of missing data, or perhaps due to the high frequency oscillations in the data. *mtsdi* performed better than our method for the simulated total water content data in terms of our metrics, although it did not seem to capture much of the details of the lowest values or the dynamics in the data. Figure 11 shows that *mtsdi* seemed to capture the large faster oscillating peaks quite well, but for the slower changes of TWC, it appeared to use a mean value instead of capturing the slower variabilities in the data.

Although our proposed imputation approach performed generally well and often outperformed the other methods, its computing time requirement was significantly larger than that of *imputeTS* and *mtsdi*. This is due to the

hyperparameter optimization trying many different sets of hyperparameters and each training is computationally nontrivial. On the other hand, if the values to be gap filled behave nonlinearly, and if high prediction accuracy is desired, we recommend using the DL-based approach due to the limitations of the other methods in these cases. In practice, one would not conduct a detailed study as presented here with artificially created gaps, but rather optimize and train the MLP once for the time series to be gap filled. One way to speed up the DL approach is to either parallelize the repetitions of the training if compute resources are available or to reduce the number of repetitions. However, the latter option may lead to models that are less reliable in terms of predictive performance. In the FLUXNET test cases, we used only one training repetition (instead of five as done with the two other use cases). The MLP-based gap filling method always performed significantly better than the other two methods, indicating that it

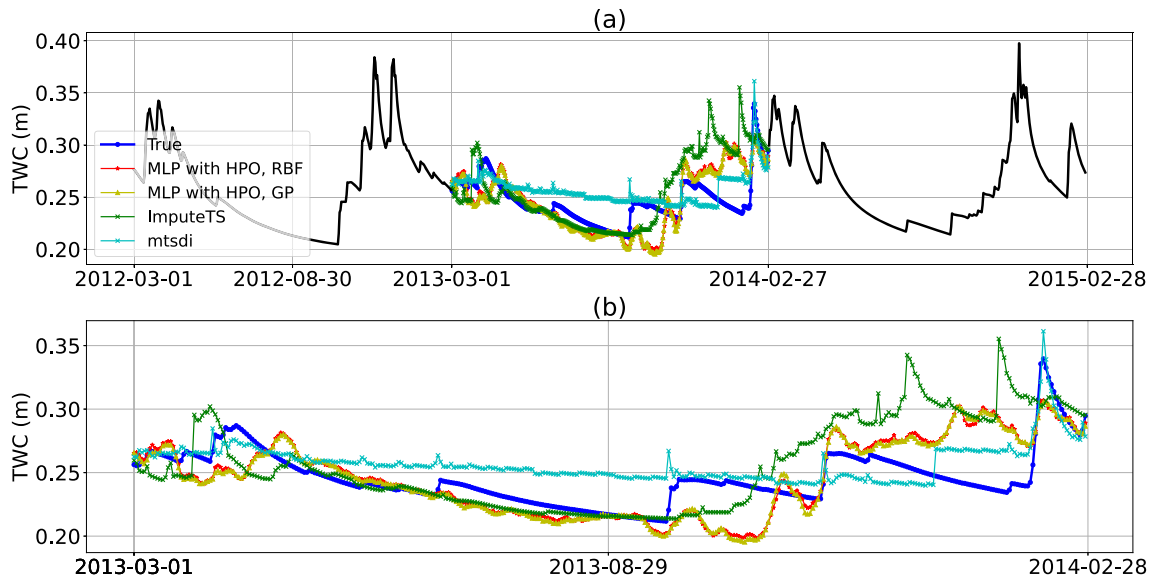


Fig. 10 Comparison of gap filling methods for daily TWC with missing values between March 2013 and February 2014. The top figure shows the missing time series in context of available data, and the bottom figure shows a zoom onto the missing period

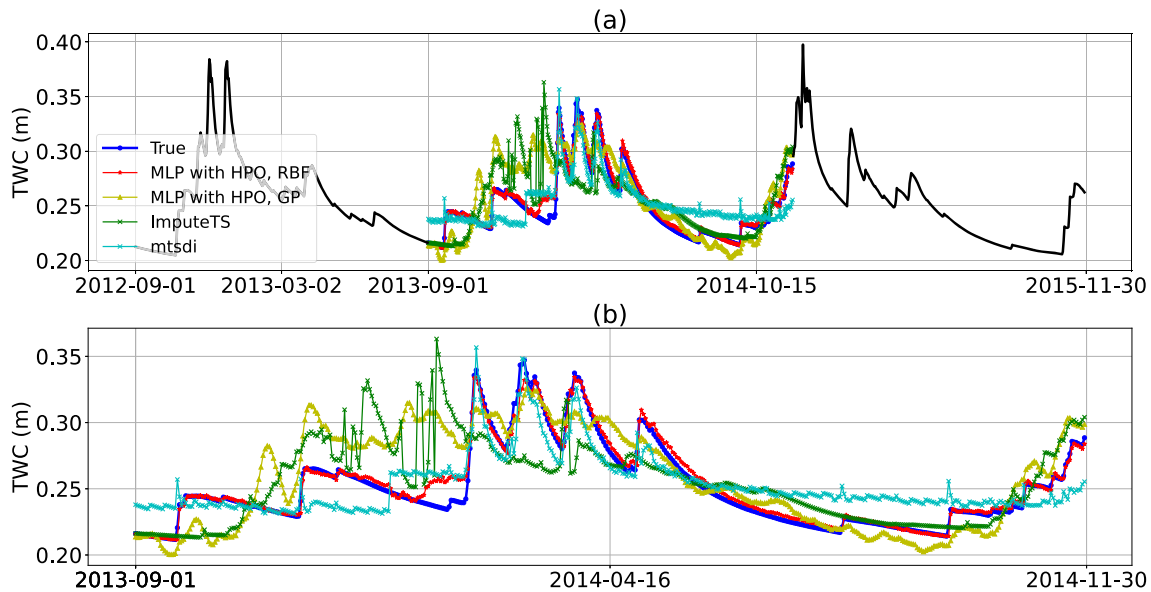


Fig. 11 Comparison of gap filling methods for daily TWC with missing values between September 2013 and November 2014. The top figure shows the missing time series in context of available data, and the bottom figure shows a zoom onto the missing period

was not simply a “lucky” draw. At the same time, we saw that different architectures were able to perform similarly well for this case, indicating that there could be regions in the hyperparameter search space where the objective function is flat and thus finding good hyperparameters and well-performing MLPs may be easier.

Many environmental datasets contain functional relationships between data features. However, many time series imputation methods (including the methods `ImputeTS` and `mtsdi` discussed in this paper) do not take advantage of this additional information and they use only statistical

information contained in the time series data. In other words, they are purely data-driven approaches. Unlike these approaches, our proposed DL method can be extended to include any known physical relationships and constraints in addition to the statistical information. For example, in the `HYDRUS-1D` case, we know that there is a mass conservation constraint present between the features and it is possible to include this constraint, e.g., by adding a penalty term to the objective function.

In our study, we compared the different methods based on their `RMSE`, `MAE`, and `MAPE` (or `SMAPE`)

Table 7 (a) RMSE, MAE, and SMAPE for different gap filling methods for hourly FLUXNET dataset. (b) Optimal hyperparameters ([batch size, # epochs, # layers, # nodes per layer, dropout rate, # lags]) for both RBF- and GP-based HPO

Missing date range (yyyy-mm-dd)		#obs	RMSE, MAE, SMAPE			
From	To		MLP RBF	MLP GP	ImputeTS	mtsdi
(a)						
2007-01-01	2007-03-31	2,159	0.98, 0.7, 44.58	0.96, 0.68, 41.98	1.62, 1.40, 47.42	4.83, 3.51, 65.61
2007-01-01	2007-12-31	8,759	2.72, 1.68, 29.85	2.68, 1.73, 30.75	7.92, 4.82, 69.10	5.17, 3.80, 51.10
2007-01-01	2008-12-31	17,543	2.86, 1.77, 34.59	2.96, 1.89, 56.29	8.13, 4.79, 57.69	5.14, 3.76, 49.85
2007-01-01	2010-12-31	35,063	2.75, 1.68, 28.58	2.94, 1.84, 33.35	8.66, 5.17, 55.44	5.60, 4.03, 52.05

Missing date range (yyyy-mm-dd)		#obs	Hyperparameters	
From	To		MLP RBF	MLP GP
(b)				
2007-01-01	2007-03-31	2,159	190, 450, 2, 45, 0, 50	180, 100, 3, 40, 0, 240
2007-01-01	2007-12-31	8,759	55, 400, 4, 40, 0, 30	150, 200, 1, 35, 0, 355
2007-01-01	2008-12-31	17,543	60, 350, 5, 45, 0, 90	55, 200, 2, 35, 0, 185
2007-01-01	2010-12-31	35,063	60, 350, 5, 45, 0, 90	70, 500, 1, 25, 0, 345

Bold numbers indicate the **best value** for a given problem instance. The column “#obs” indicates the number of missing observations. “MLP RBF” and “MLP GP” indicate the methods using the MLP with the RBF and GP model during HPO, respectively

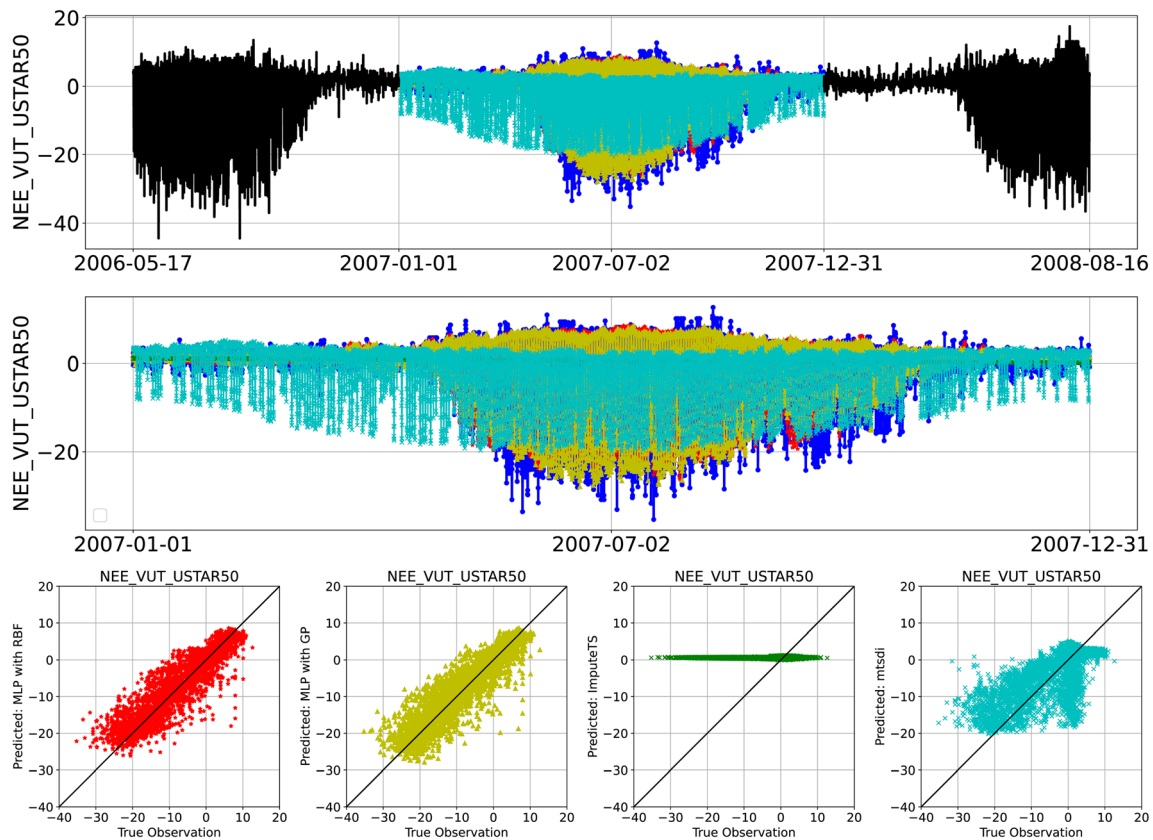


Fig. 12 Hourly NEE_VUT_USTAR50 (FLUXNET) imputation comparison: missing value range is January 1, 2007 to December 31, 2007. The top panel shows the true and the predicted values in the context of the longer time series. The middle panel shows a zoom

onto the filled gap. The bottom panels show scatter plots for each method (ideally, prediction and true data are identical and points would lie on the diagonal)

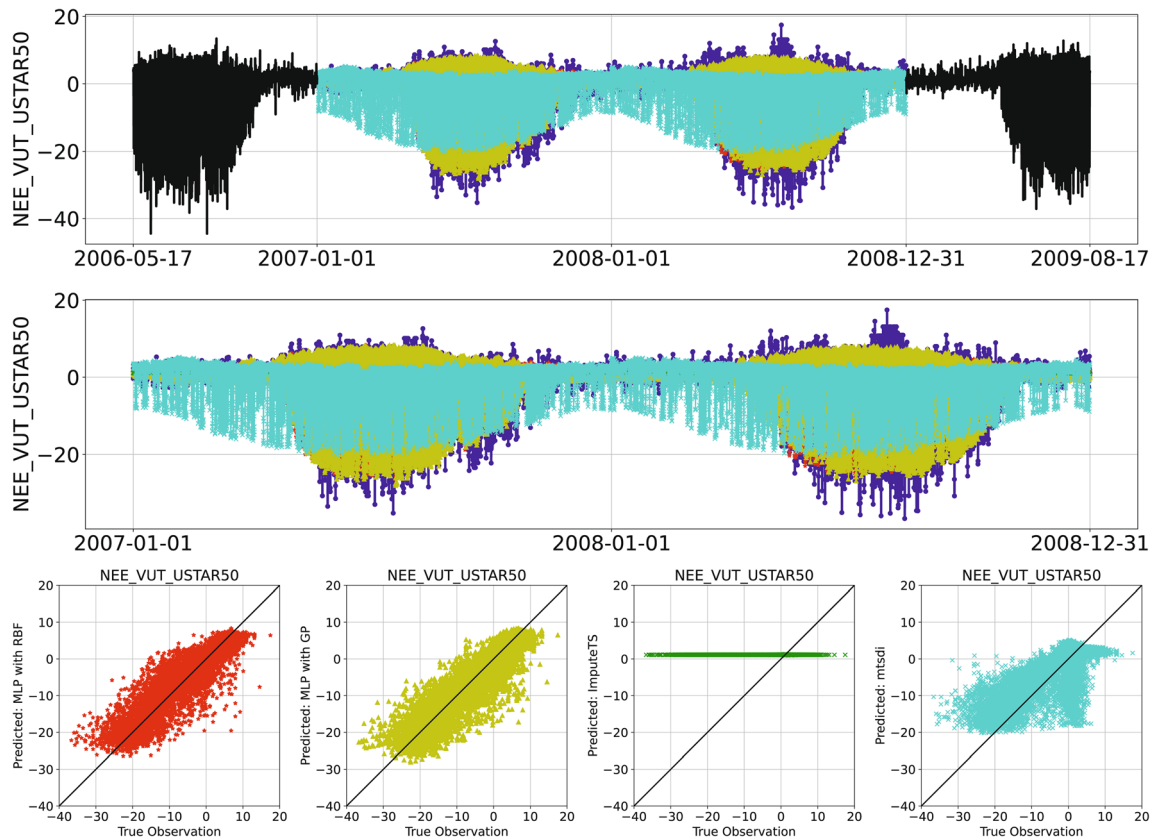


Fig. 13 Hourly NEE_VUT_USTAR50 imputation comparison: missing value range is January 1, 2007 to December 31, 2008. The top panel shows the true and the predicted values in the context of the longer time series. The middle panel shows a zoom onto the filled

performance. Although these measures are widely used, they do not reflect the individual errors well. For example, in the daily TWC imputation comparison shown in Fig. 10, one could argue that the proposed MLP-based methods produce fairly accurate and better predictions than *metsdi* for 50% of the missing values, but they accumulate larger errors for the remaining values. On the other hand, *metsdi* does not match most of the data well, but its almost flat prediction accrues overall a lower error. Thus, it is possible that optimizing for another performance measure may lead to better outcomes for the MLP methods.

4 Summary and conclusions

In this paper, we proposed a time series missing value imputation method that uses DL models. Our focus was on time series imputation problems where a long continuous gap was present in one variable, and where all other supporting variables are fully observed. Our proposed method uses an MLP whose architecture is tuned with a derivative-free surrogate model based optimization approach. To this

gap. The bottom panels show scatter plots for each method (ideally, prediction and true data are identical and points would lie on the diagonal)

end, we modified the hyperparameter optimization approach proposed in [22] in order to facilitate the imputation of missing values in time series.

After training the MLP, we imputed the missing values sequentially from the first missing value to the last missing value, basing predictions on previous predictions as we go along. Our sequential estimation of missing values allowed us to fill long-term gaps. We chose MLPs as DL models because they have previously shown good performance for time series data and because the time needed to optimize the model is significantly shorter than for other types of DL models (see [22]). Generally, however, any feedforward artificial neural network can be used within our proposed method such as convolutional neural networks [43–45].

We performed numerical experiments with three different test cases in which we created artificial gaps which allowed us to assess the quality of the gap filled values. We used two test cases with observed values and one test case with simulated data. The numerical results showed that for many problem instances, our proposed approach was able to fill the gaps with highly accurate values and it was able to capture trends and dynamics in the data better than other

data-driven methods such as *ImputeTS* and *mtsdi*. Using the RBF as surrogate model during HPO appeared to yield more stable performance than using the GP. Therefore, we recommend using our proposed approach with the RBF surrogate during HPO.

The numerical results also showed that the optimal MLP hyperparameters depended on which data were missing (gap size and location in the time series). Therefore, hyperparameter optimization is essential before using any ML model to impute the missing values.

The results of this paper can lay the foundation for studying new approaches for filling long gaps in time series. Although the proposed method was able to fill large gaps with seasonal trends well, including capturing large and small values, it sometimes failed to correctly estimate the missing values toward the end of the data gap, indicating a potential for accumulation of errors as we inform predictions with previously predicted values. A “backward imputation approach” might be helpful to alleviate this drawback. In this approach, we could modify the MLP inputs such that we start imputing data from the end of the gap backwards in time. Combining both forward and backward imputation may allow us to reap the benefits of both.

Our study focused on time series that have one large gap. However, we expect the method to perform reasonably well when we have multiple large gaps in the time series. Since the inputs to the MLP are designed such that they are independent and the HPO is general, we expect that filling multiple gaps with this method is straight-forward. For time series with a mix of large and small gaps, a combination of our DL model-based method and other methods that are aimed at filling small gaps can be used. Incorporation of physics constraints in the DL model predictions has the potential to improve prediction accuracy. Moreover, if values are missing in several variables of the multivariate time series, it may be possible to modify the proposed approach for this scenario. This is a future research topic.

This paper studied the time-series imputation with the MLP which is the most basic feedforward neural network. It would also be interesting to study the proposed time-series imputation approach with other advanced feedforward neural networks such as convolutional neural network [45] and AR-NET [46] as well as recently developed forecasting software such as *NeuralProphet* [47].

Acknowledgements This work was supported by Laboratory Directed Research and Development (LDRD) funding from Berkeley Lab, provided by the Director, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

Data availability The datasets used in this study have been deposited in the GitHub public repository: https://github.com/JanghoPark-LBL/Dataset_Missing_Value_Imputation_with_Deep_Neural_Networks.

Declarations

Conflict of interest There is no conflict of interest to disclose. This research does not involve Human Participants and/or Animals. All authors are satisfied with the submitted version of the manuscript.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- García-Laencina PJ, Sancho-Gómez J-L, Figueiras-Vidal AR (2010) Pattern classification with missing data: a review. *Neural Comput Appl* 19(2):263–282
- Yozgatligil C, Aslan S, Iyigun C, Batmaz I (2013) Comparison of missing value imputation methods in time series: the case of Turkish meteorological data. *Theoret Appl Climatol* 112(1–2):143–167
- Kalteh AM, Hjorth P (2009) Imputation of missing values in a precipitation-runoff process database. *Hydrol Res* 40(4):420–432
- Aissia M-AB, Chebana F, Ouada TB (2017) Multivariate missing data in hydrology-review and applications. *Adv Water Resour* 110:299–309
- Gu C, Anderson W, Maggi F (2012) Riparian biogeochemical hot moments induced by stream fluctuations. *Water Resour Res* 48(9)
- Arora B, Wainwright HM, Dwivedi D, Vaughn LJ, Curtis JB, Torn MS, Dafflon B, Hubbard SS (2019) Evaluating temporal controls on greenhouse gas (ghg) fluxes in an arctic tundra environment: An entropy-based approach. *Sci Total Environ* 649:284–299
- Thi-Thu-Hong Phan B, Caillaud EP, Bigand A (2019) edtwbi: Effective imputation method for univariate time series. In: *Advanced Computational Methods for Knowledge Engineering: Proceedings of the 6th International Conference on Computer Science, Applied Mathematics and Applications, ICCSAMA 2019*, vol. 1121, p 121. Springer Nature
- Moritz S, Sardá A, Bartz-Beielstein T, Zaefferer M, Stork J (2015) Comparison of different methods for univariate time series imputation in R. *arXiv preprint arXiv:1510.03924*
- Moritz S, Bartz-Beielstein T (2017) *imputeTS*: time series missing value imputation in R. *The R J* 9(1):207–218
- Phan T-T-H (2020) Machine learning for univariate time series imputation. In: *2020 International Conference on Multimedia Analysis and Pattern Recognition (MAPR)*, pp 1–6
- Batista GEAPA, Monard MC (2002) A study of K-nearest neighbour as an imputation method. *Hybrid Intell Syst* 87(251–260):48

12. Che Z, Purushotham S, Cho K, Sontag D, Liu Y (2018) Recurrent neural networks for multivariate time series with missing values. *Sci Rep* 8(1):1–12
13. Cao W, Wang D, Li J, Zhou H, Li L, Li Y (2018) Brits: bidirectional recurrent imputation for time series. In: *NeurIPS*
14. Luo Y, Cai X, Zhang Y, Xu J (2018) Multivariate time series imputation with generative adversarial networks. In: *Advances in Neural Information Processing Systems*, pp 1596–1607
15. Zhang Y, Zhou B, Cai X, Guo W, Ding X, Yuan X (2021) Missing value imputation in multivariate time series with end-to-end generative adversarial networks. *Inf Sci* 551:67–82
16. Lingras P, Zhong M, Sharma S (2008) Evolutionary regression and neural imputations of missing values. In: *Soft Computing Applications in Industry*, pp 151–163
17. Zhong M, Sharma S, Lingras P (2007) Rationalizing reliable imputation durations of genetically designed time delay neural network and locally weighted regression models. *Transp Plan Technol* 30(6):609–626
18. Mahabbati A, Beringer J, Leopold M, McHugh I, Cleverly J, Isaac P, Izady A (2020) A comparison of gap-filling algorithms for eddy covariance fluxes and their drivers. *Geosci Instrum Methods Data Syst Discuss* 10:1–31
19. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E (2011) Scikit-learn: machine learning in Python. *J Mach Learn Res* 12:2825–2830
20. Kim Y, Johnson MS, Knox SH, Black TA, Dalmagro HJ, Kang M, Kim J, Baldocchi D (2020) Gap-filling approaches for eddy covariance methane fluxes: a comparison of three machine learning algorithms and a traditional method with principal component analysis. *Glob Change Biol* 26(3):1499–1518
21. Waibel A, Hanazawa T, Hinton G, Shikano K, Lang KJ (1989) Phoneme recognition using time-delay neural networks. *IEEE Trans Acoust Speech Signal Process* 37(3):328–339
22. Müller J, Park J, Sahu R, Varadharajan C, Arora B, Faybisenko B, Agarwal D (2020) Surrogate optimization of deep neural networks for groundwater predictions. *J Glob Optim* 10–100710898020009120
23. Junger W Ponce de Leon, A.: Package 'mstsi'
24. Nair V, Hinton GE (2010) Rectified linear units improve restricted boltzmann machines. *ICML'10*, pp 807–814. *Omni-press, Madison, WI, USA*
25. Kingma DP, Ba J (2014) Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*
26. Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. *Nature* 323(6088):533–536
27. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R (2014) Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res* 15(1):1929–1958
28. Goodfellow I, Bengio Y, Courville A, Bengio Y (2016) *Deep learning*, vol 1. MIT press Cambridge, Cambridge
29. Feurer M, Hutter F (2019) *Hyperparameter optimization. Automated machine learning*. Springer, Berlin, pp 3–33
30. Bergstra J, Bengio Y (2012) Random search for hyper-parameter optimization. *J Mach Learn Res* 13(2)
31. Larochelle H, Erhan D, Courville A, Bergstra J, Bengio Y (2007) An empirical evaluation of deep architectures on problems with many factors of variation. In: *Proceedings of the 24th International Conference on Machine Learning*, pp 473–480
32. Snoek J, Larochelle H, Adams RP (2012) Practical Bayesian optimization of machine learning algorithms. *arXiv preprint arXiv:1206.2944*
33. Xie L, Yuille A (2017) Genetic CNN. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp 1379–1388
34. Sahu RK, Müller J, Park J, Varadharajan C, Arora B, Faybisenko B, Agarwal D (2020) Impact of input feature selection on groundwater level prediction from a multi-layer perceptron neural network. *Front Water* 2:46
35. Šimůnek J, Šejna M, Saito H, Sakai M, van Genuchten MT (2013) *The HYDRUS-1D Software Package for Simulating the One-Dimensional Movement of Water, Heat, and Multiple Solutes in Variably-Saturated Media*, Version 4.17, edn. Riverside, California, Riverside, California
36. Šimůnek J, van Genuchten MT (2008) Modeling nonequilibrium flow and transport processes using hydrus. *Vadose Zone J* 7(2):782–797
37. Arora B, Mohanty BP, McGuire JT (2015) An integrated Markov chain Monte Carlo algorithm for upscaling hydrological and geochemical parameters from column to field scale. *Sci Total Environ* 512:428–443
38. Baek S, Ligaray M, Pachepsky Y, Chun JA, Yoon K-S, Park Y, Cho KH (2020) Assessment of a green roof practice using the coupled SWMM and HYDRUS models. *J Environ Manage* 261:109920
39. Pastorello G, Trotta C, Canfora E, Chu H, Christianson D, Cheah Y-W, Poindexter C, Chen J, Elbashandy A, Humphrey M (2020) The FLUXNET2015 dataset and the ONEFlux processing pipeline for eddy covariance data. *Scientific Data* 7(1):1–27
40. United States Department of Energy, T.O.o.S.: FLUXNET2015 Dataset. <https://fluxnet.org/data/fluxnet2015-dataset/>. Last accessed: 2021-1-1
41. Novick K, Phillips R (2016) FLUXNET2015 US-MMS morgan monroe state forest. FLUXNET; Indiana Univ. <https://doi.org/10.18140/flux/1440083>
42. Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, Lin Z, Desmaison A, Antiga L, Lerer A (2017) Automatic differentiation in PyTorch. In: *31st Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, USA
43. Fukushima K (1980) Neocognitron: a self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol Cybern* 36(4):193–202
44. LeCun Y, Boser B, Denker JS, Henderson D, Howard RE, Hubbard W, Jackel LD (1989) Backpropagation applied to handwritten zip code recognition. *Neural Comput* 1(4):541–551
45. LeCun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. *Proc IEEE* 86(11):2278–2324
46. Triebe O, Laptev N, Rajagopal R (2019) Ar-net: A simple autoregressive neural network for time-series. *arXiv preprint arXiv:1911.12436*
47. Triebe O, Hewamalage H, Pilyugina P, Laptev N, Bergmeir C, Rajagopal R (2021) Neuralprophet: Explainable forecasting at scale. *arXiv preprint arXiv:2111.15397*

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.