



Original software publication

tsrobbrep – an R package for robust preprocessing of time series data

Michał Narajewski^{*}, Jens Kley-Holsteg, Florian Ziel

University of Duisburg–Essen, Germany



ARTICLE INFO

Article history:

Received 1 October 2020

Received in revised form 13 July 2021

Accepted 8 September 2021

Keywords:

Data cleaning

Missing values

Outliers

Anomaly detection

Time series

Robust preprocessing

Energy system data

ABSTRACT

Data cleaning is a crucial part of every data analysis exercise. Yet, the currently available R packages do not provide fast and robust methods for cleaning and preparation of time series data. The open source package `tsrobbrep` introduces efficient methods for handling missing values and outliers using model based approaches. For data imputation a probabilistic replacement model is proposed, which may consist of autoregressive components and external inputs. For outlier detection a clustering algorithm based on finite mixture modelling is introduced, which considers time series properties in terms of the gradient and the underlying seasonality as features. The procedure allows to return a probability for each observation being outlying data as well as a specific cause for an outlier assignment in terms of the provided feature space. The methods work robust and are fully tunable. Moreover, by providing the `auto_data_cleaning` function the data preprocessing can be carried out in one cast, without comprehensive tuning and providing suitable results. The primary motivation of the package is the preprocessing of energy system data. We present application for electricity load, wind and solar power data.

© 2021 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Code metadata

Current code version	0.3.1
Permanent link to code/repository used for this code version	https://github.com/ElsevierSoftwareX/SOFTX_2020_123
Legal Code License	MIT license (MIT)
Code versioning system used	git
Software code languages, tools, and services used	R
Compilation requirements, operating environments & dependencies	R ($\geq 3.2.0$)
If available Link to developer documentation/manual	https://cran.r-project.org/web/packages/tsrobbrep/tsrobbrep.pdf
Support email for questions	michal.narajewski@uni-due.de

Software metadata

Current software version	0.3.1
Permanent link to executables of this version	https://cran.r-project.org/web/packages/tsrobbrep
Legal Software License	MIT license (MIT)
Computing platforms/Operating Systems	Linux, OS X, Microsoft Windows
Installation requirements & dependencies	R ($\geq 3.2.0$)
If available, link to user manual - if formally published include a reference to the publication in the reference list	https://cran.r-project.org/web/packages/tsrobbrep/tsrobbrep.pdf
Support email for questions	michal.narajewski@uni-due.de

1. Introduction

The increasing availability of high frequency data measured, recorded and processed in real-time leverage the potential benefit

^{*} Corresponding author.

E-mail address: michal.narajewski@uni-due.de (Michał Narajewski).

of data analytics in day to day business. However, as the data is usually measured by sensors and real measurement stations, measurement and transmission failures are inevitable so that a data preprocessing is required. Focusing on energy systems, two aspects are crucial, firstly, the accrued data is often immediately processed so that fast and automatic algorithms are needed. Secondly, the data is usually unlabelled so that the detection requires purely data-driven approaches. Hence, the data cleaning and preprocessing requires an automatic, robust and mainly data-driven procedure, which is easily tunable and applicable in real-time to provide verified data for a safe operation of the energy system.

In the literature several R packages are available, which address the task of data cleaning and preparation. On the one hand, there are several packages specifically designed for dealing with missing values [1]. Here, the packages are mainly focused either on data imputation in multivariate data by utilizing inter-attribute correlations as done in Amelia [2], mice [3] and VIM [4], or on univariate data by employing time dependencies as done in imputeTS [5], zoo [6] and forecast [7]. On the other hand, there are several packages specifically designed for dealing with outlying data. Here, among others HDoutliers [8], mvoutlier [9], and tsoutliers [10] can be named. However, packages combining both outlier detection and data imputation methods for time series data in a robust and efficient way, as implemented e.g. in forecast [7] and imputeFin [11] for univariate data, are rare.

The presented package tsrobprep is specialized on dealing with outliers and missing values in time series data in a robust and simple applicable way. It offers as core functions an algorithm for modelling missing values and an algorithm for detecting outliers. To enhance usability the package comes with three additional side functions, a robust decomposition algorithm, an algorithm to impute the modelled values, and a wrapper function to apply the former named functions at once. The first core function for modelling missing values is based on a two-step approach, which allows to deal with autoregressive effects and given external inputs. The second core function for outlier detection is based on a feature based clustering approach with finite (Gaussian) mixture models. Both core functions provide a probabilistic framework such that for data imputation a whole distribution can be modelled and for outlier detection the probability being an outlying data point can be provided. Focused on robustness, efficiency and usability the package tsrobprep is perfectly suited for cleaning various time series data sets in a univariate and multivariate setting, as done e.g. in [12]. The package is available online in the CRAN repository.

The paper is structured as follows. Section 2 describes the utilized methods in detail and Section 3 presents the usability and data preprocessing potential. Section 4 compares the methods to competitors from literature. Finally, Section 5 discusses the impact of the package and concludes the paper.

2. Software description

For dealing with missing values, the `model_missing_data` function is proposed for handling univariate and multivariate time series data. It takes the time series data to be cleaned as an argument and returns the modelled missing values to be imputed. The function uses a two-step model-based approach to replace missing values. First, the data is decomposed to trend, seasonal and external input components, as well as to their interactions. This is done using the `robust_decompose` function that is available in the package. Then, on the remainder we utilize the following regression

$$\text{Target}_\tau(Y_t) = \beta_{0,\tau} + \sum_{p \in \mathcal{L}} \beta_{p,\tau} Y_{t-p} + \sum_{i \in \mathcal{I}} \beta_{i,\tau} X_{i,t} + \varepsilon_t = \beta_\tau \mathbb{X}_t + \varepsilon_t, \quad (1)$$

where Target_τ is either the expected value or a τ -quantile, Y_t is the time series remainder with the missing data to be replaced, and $X_{i,t}$ are the decomposition components (including external regressors) to allow in a multivariate framework for utilizing beside time-dependencies also inter-attribute correlations. Additionally, \mathcal{L} is a set of lags for the autoregression, and \mathcal{I} is a set of the decomposition components. By allowing for positive and negative lags in \mathcal{L} , past and forthcoming (if available) values can be utilized for modelling. By default, the set of lags \mathcal{L} takes reasonable values, depending on the provided vector of seasonalities $\mathbf{S} = (S_1, \dots, S_K)$ and the input time series data. Precisely, \mathcal{L} contains all lags except those which involve multiple seasonal parameters and their negative values of a multi-seasonal $\text{AR}(p+1, p, \dots, p)$ with seasonalities \mathbf{S} and p depending on data size. For example for $\mathbf{S} = (24, 168)$ and $p = 1$, this is $\mathcal{L} = \{1, 2, 24, 25, 168, 169, -1, -2, -24, -25, -168, -169\}$.

If the user provides a matrix of external inputs, the algorithm allows to select the most appropriate inputs by user-tunable thresholds based on the linear correlation between the externals and the explained data. To utilize not only the inter-attribute information and autoregressive effects also lags for the externals can be applied for modelling. In case of the Target being the expected value the weighted lasso implemented in function `glmnet` in the `glmnet` package [13] can be applied. Otherwise, a quantile regression implemented with the fast Frisch–Newton algorithm in function `rq.fit.fnb` in `quantreg` package [14] can be applied. Moreover, the algorithm uses recursive replacement by default, which improves the modelling accuracy in the case of larger gaps, but it may also cause an error accumulation, so we give the user the option to switch it off.

For dealing with outlying data, the second core function `detect_outliers` is proposed for handling univariate data. It takes the time series data to be cleansed and returns the position of the detected outliers, the probability of being an outlying instance, the applied feature matrix and the specific cause for an outlier assignment in terms of the applied features. By recursive application within the wrapper function `auto_data_cleaning` also multivariate data sets can be handled. The algorithm can be tuned to control robustness, computational time and detection sensitivity. By utilizing specifically designed features for time series data, various patterns can be identified. As features, the gradient, absolute gradient, relative gradient, deterministic seasonal trend, seasonal gradient and absolute seasonal gradient are applied. For example, the absolute gradient tracks two-sided jumps with regard to the closest neighbours, the absolute seasonal gradient works the same, however, applied to the closest neighbours in terms of the provided seasonality \mathbf{S} . The relative gradient in turn tracks two sided jumps in relation to the local seasonal variance to account for moderate spikes in low volatility periods. By applying an unsupervised learning approach the procedure is well suited to handle unlabelled and unstructured data. The approach is based on model-based clustering with finite mixture models as outlined in [15]. It is assumed that the multivariate distribution of the feature space is a mixture of G components, the likelihood for a mixture model can be defined as

$$\ell_{\text{MIX}}(\theta_1, \dots, \theta_G | \mathbf{y}) = \prod_{i=1}^n \sum_{k=1}^G \pi_k f_k(\mathbf{y} | \theta_k), \quad (2)$$

where \mathbf{y} is the observation vector y_1, \dots, y_n , f_k is the density (assumed to be multivariate normal (Gaussian)), θ_k are the parameters of the k th component in the mixture and π_k is the probability that an observation belongs to the k th component. Note that $\sum_{k=1}^G \pi_k = 1$. Within the mixture one component is intended to represent outlying data. This component is introduced by artificially augmenting the number of outliers in the

data set by blowing up the covariance matrix of the mixture. For modelling the data as a Gaussian finite mixture with different covariance structures and numbers of mixture components the `Mclust` function in `mclust` package [16] is applied. For estimation the expectation-maximization algorithm initialized by hierarchical model-based agglomerative clustering is used as outlined in [15]. To speed up the computational time and robustness of the method, the estimation is carried out repeatedly on varying data subsets. To provide the specific cause for an outlier assignment, each observation, which has been identified as an outlying data point is shifted to the corresponding feature and multivariate feature combination mean, respectively. Then, the probability of being outlying data is recomputed allowing to name a specific feature or a combination of features, which caused the outlier assignment.

The auxiliary function `impute_modelled_data` eases the usage of the package. It takes as argument a `tsrobprep` object, i.e. the object returned by the `model_missing_data` function, imputes the replaced values for given quantiles τ and returns a list of matrices (each element of the list corresponds to each quantile). The `auto_data_cleaning` function combines everything described above. First, it replaces missing values (if there are any) with a median forecast. This step is required as the outlier detection algorithm does not allow for missing values. Second, unreliable outliers are detected so that in the third step, the detected outliers and the initially modelled missing values are modelled and imputed in one cast. The repetitive modelling of the initially modelled missing values is reasonable as potential outliers might have negatively influenced the modelling process in the first run. The workflow of the function is presented in Fig. 1.

3. Illustrative examples

We present the potential of the package by using its core functions on three time series data sets: the electricity load in Great Britain, the solar generation in Germany (TransnetBW control zone), and the wind onshore generation in Italy (Sardinian control zone). All these data sets span the date range from January 01, 2015 to June 30, 2020. The first time series is half-hourly, i.e. contains 48 values per day, the second quarter-hourly, i.e. contains 96 values per day, and the third is hourly, i.e. contains 24 values per day. The utilized data is freely available at ENTSOE Transparency platform [17]. The considered time series are plotted over time in Fig. 2. We can clearly see that the time series of electricity load in the United Kingdom exhibits many outliers that are very likely unreliable. The other data sets do not seem to contain such values, at least based on the first sight. All data sets contain missing

values, what is hard to notice in Fig. 2 due to the high frequency of the data.

For a better illustration, a subset of each series containing either missing values, outliers or both is depicted in Fig. 3. The time series of electricity load in Great Britain exhibits both missing values and unreliable outliers. The latter ones are quite easy to spot by eyeballing – they are the two negative spikes visible in the plot. The time series of solar generation in Germany contains unreliable outliers, however, much harder to detect. Namely, all the observations of February 07, 2015 are set to zero, what is impossible even for a very cloudy day. Such outliers are especially challenging as half of the observations are equal to zero. The third example, the wind onshore generation in Italy, contains no unreliable values but a clear gap in the data on November 09, 2017.

The cleansing of the named series was done with the `auto_data_cleaning` function. The first time series, electricity load in Great Britain, and the third time series, wind onshore generation in Italy, were handled in a univariate setting without external inputs with $S = (48, 7 \cdot 48)$ and $S = 24$, respectively. To account for the unreliable zeros during the day shown in Fig. 3 the `consider.as.missings` parameter has been used to recognize zeros as missing values only if the whole day consists of zeros. The second series, the solar generation in Germany, has been handled with $S = 96$ in a multivariate setting such that external regressors as plotted in Fig. 4 were utilized for data modelling and imputation. Additionally, we set the `whole.period.missing.only` parameter to `TRUE` to consider zeros as missings only if the whole day is zero, and we set `replace.recursively` to `FALSE` to avoid error accumulation during night hours. All the other arguments were set to their default values. The models were estimated for quantiles $\tau \in \{0.025, 0.05, 0.25, 0.5, 0.75, 0.95, 0.975\}$.

The results provided in Fig. 3 are very promising. This accounts in particular for the first data set, where only autoregressive effects were used. It can be seen that the outliers were smoothly replaced, i.e. not only the outliers themselves were replaced, but also their direct neighbours. The cleaning of the other two data sets is also satisfying and the values seem to be reasonable.

Although the algorithm works well and robust, it has to be noted that specific causes of an anomaly (e.g. technical failures, true behavioural issues or nature phenomenons) are not always clear distinguishable and detectable. An example of such situation is presented in Fig. 5. Here, the algorithm is able to detect the anomalous negative spike. However, due to the fact that this spike corresponds to a partial solar eclipse, the detection is considered as false positive. By providing external inputs, i.e. generation in the other German control zones, the algorithm is able to correct this misclassification by imputing very similar values. However,

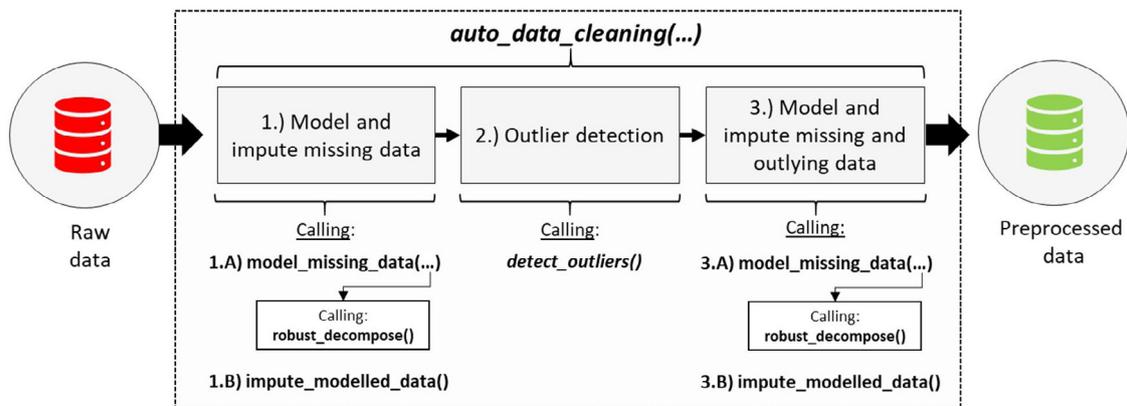


Fig. 1. Workflow of the auto data cleaning function.

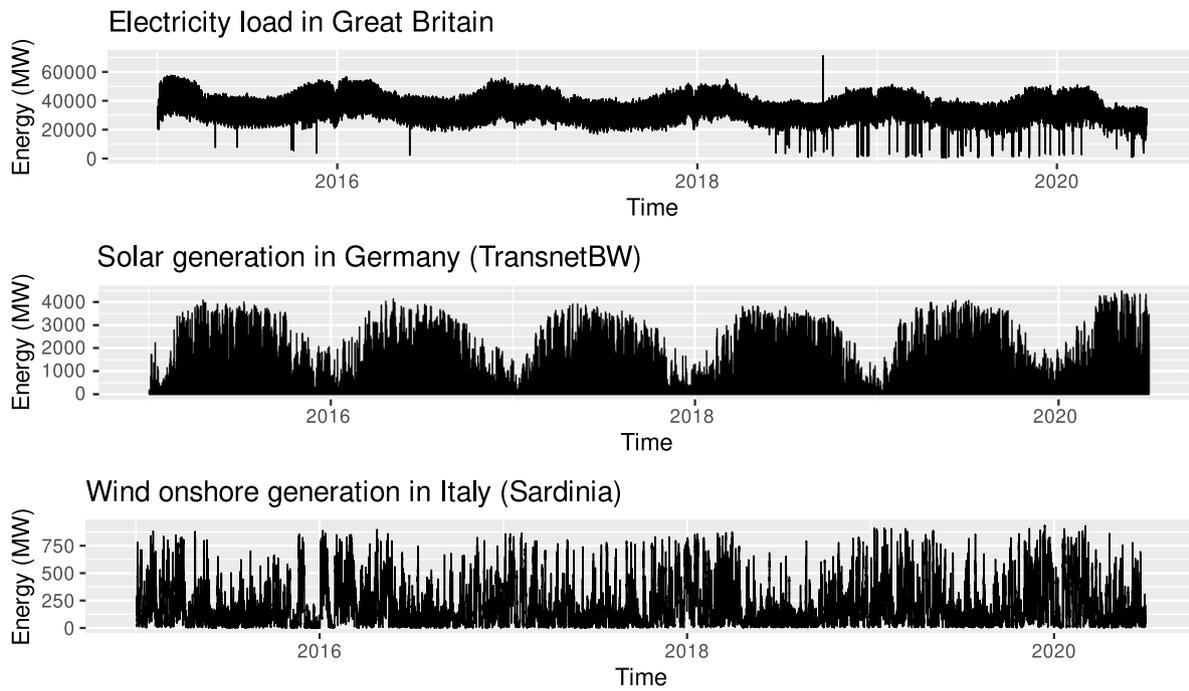


Fig. 2. Three exemplary data sets containing both missing values and unreliable outliers.

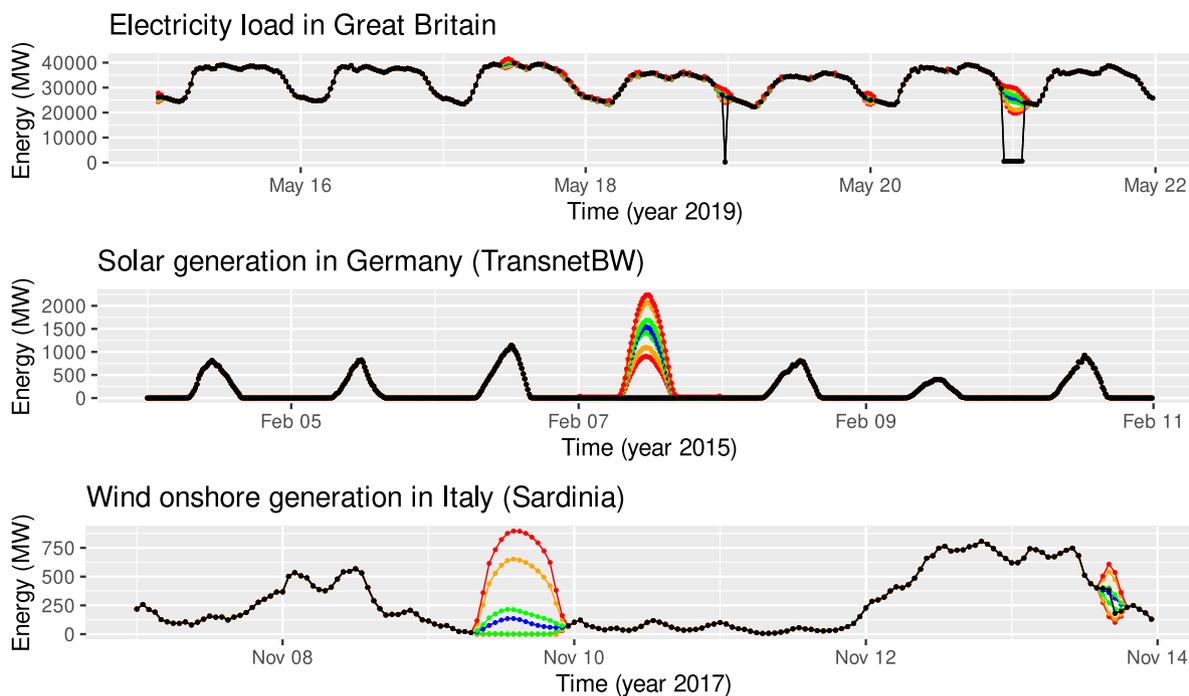


Fig. 3. Chosen parts of the data that illustrate the problem with replaced missing values and unreliable outliers. Red lines indicate 95% confidence intervals, orange 90%, green 50%, and blue indicates the median. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

this example shows us that even though the provided package works robust, a supervision by humans, in particular when dealing with outliers and extreme values, is recommended.

4. Performance evaluation against competitors

To verify the potential of the proposed package, its functions are compared and evaluated against benchmarks from the

literature. The performance of `model_missing_data` is evaluated in an imputation study. Here, the electricity load in Great Britain ranging from January 01, 2019 to December 31, 2019 is used. To allow a performance evaluation the named data set is manipulated by replacing existing values by missing entries. Observations are removed in continuous blocks. Assuming that values are missing at random (MAR) such that the likelihood of an observation missing depends on whether observations closer

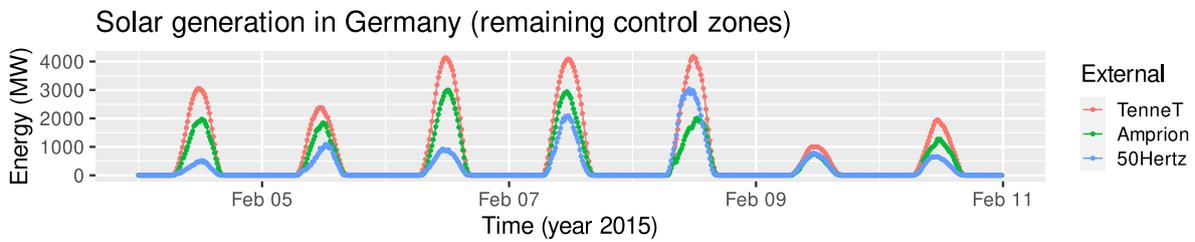


Fig. 4. The solar generation in Germany in remaining control zones used as external regressors in our exercise.

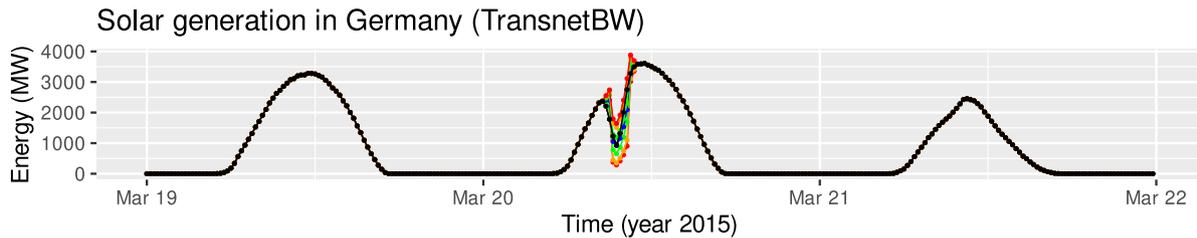


Fig. 5. A drop in the solar generation in Germany that may seem to be unreliable while it was caused by a natural phenomenon, i.e. the partial solar eclipse. Red lines indicate 95% confidence intervals, orange 90%, green 50%, and blue indicates the median. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

in time have also been missing. This sampling approach is implemented by a log-normal distribution parametrized with $\mu = 12$ and $\sigma = 6$. To account for a varying number of missing values, a share of 0.01, 0.05, 0.1, 0.2 and 0.5 missing values is considered. Each simulation is repeated 20 times to allow a robust evaluation. As most methods in the literature are restricted to univariate time series data, the imputation study is implemented in a univariate setting. As competitors the `impute_AR1_t` function of the `imputeFin` package [11], the `na.StructTS`, `na.locf` and `na.aggregate` functions of the `zoo` package [6], the `na_ma`, `na.kalman` and `na_interpolation` functions of the `imputeTS` package [5], the `prophet` function of the `prophet` package [18] as well as the `ts_impute_vec` of the `timetk` [19] package are applied. All methods are provided with default settings and named in accordance to the package and function name. As evaluation measures the mean absolute error (MAE) and the computational time (single-core on Intel Core i5-8250U CPU 1.80 GHz) are used. The results are shown in Table 1. It can be seen that the proposed function `model_missing_data` dominates all considered benchmarks for each simulation in terms of the MAE. In terms of the computational time, the `model_missing_data` function performs moderate in comparison to the benchmarks.

Focusing on the second core function `detect_outliers` and the wrapper function `auto_data_cleaning`, a confident evaluation is not straight-forward as the definition of outliers cannot be easily narrowed down in a generally valid way. Hence, the performance is presented illustratively by comparing the cleansed data set returned by the `auto_data_cleaning` function with the

results of two benchmark procedures from the literature, namely the `tsclean` function of the `forecast` package [7] and the `impute_AR1_t` function of the `imputeFin` package [11]. As data set the electricity load in Great Britain without any manipulation ranging from January 01, 2018 to December 31, 2019 is used. Two examples are shown in Fig. 6. It is observable that the proposed procedure `auto_data_cleaning` detects all three anomalies in the data. The outlying instances on the left side of the Figure are clearly technical failures and the outlying data on the right side of the Figure is at least questionable. By considering the lower plots of Fig. 6, the feature space of the `detect_outliers` function is displayed. The vertical lines indicate the probability of each observation being outlying data. The results provided in Table 1 and Fig. 6 are promising and underline the suitability, especially, for energy system data.

5. Impact and conclusions

So far the available R packages for time series data preprocessing are based on either simple and not robust or on complex and very slow algorithms. The introduction of the `tsrobprep` package fills this gap by providing a robust and efficient method to replace missing values and handle unreliable outliers. Moreover, by providing the function `auto_data_cleaning`, the package is simple applicable and provides already with the default parameter setting satisfying results. More advanced users can take advantage of the big set of tuning parameters to further adapt the functions to their needs.

Table 1
Results of the imputation study for varying shares of missing values in terms of the MAE in MW (and the computational time in seconds). Colour indicates the performance column-wise (the greener, the better). With bold, we depicted the best values in each column.

	0.01	0.05	0.1	0.2	0.5
<code>tsrobprep::model_missing_data</code>	623 (4.634)	603 (8.907)	631 (14.837)	657 (23.302)	828 (43.821)
<code>imputeTS::na_seadec.interpolation</code>	926 (0.062)	849 (0.063)	880 (0.063)	932 (0.06)	1317 (0.06)
<code>prophet::prophet</code>	1822 (58.115)	1779 (61.191)	1849 (56.287)	1828 (53.033)	1809 (61.058)
<code>imputeTS::na_kalman.auto.arima</code>	1887 (33.148)	2112 (27.029)	2451 (18.33)	2488 (11.76)	3235 (9.526)
<code>imputeTS::na_interpolation.stine</code>	2242 (0.005)	2542 (0.006)	2720 (0.006)	2828 (0.006)	3574 (0.006)
<code>timetk::ts_impute_vec</code>	2749 (1.525)	3016 (1.567)	3209 (1.492)	3323 (1.397)	4003 (1.589)
<code>imputeTS::na_ma.exponential</code>	3263 (0.002)	3460 (0.005)	3653 (0.01)	3721 (0.02)	4357 (0.058)
<code>imputeFin::impute_AR1_t</code>	3417 (13.768)	3730 (16.703)	3838 (19.709)	4000 (19.611)	4658 (27.522)
<code>zoo::na.locf</code>	4881 (0.001)	5285 (0.001)	5268 (0.003)	5537 (0.001)	6000 (0.002)
<code>zoo::na.aggregate.median</code>	5758 (0.002)	5968 (0.002)	6163 (0.006)	6120 (0.002)	6119 (0.004)

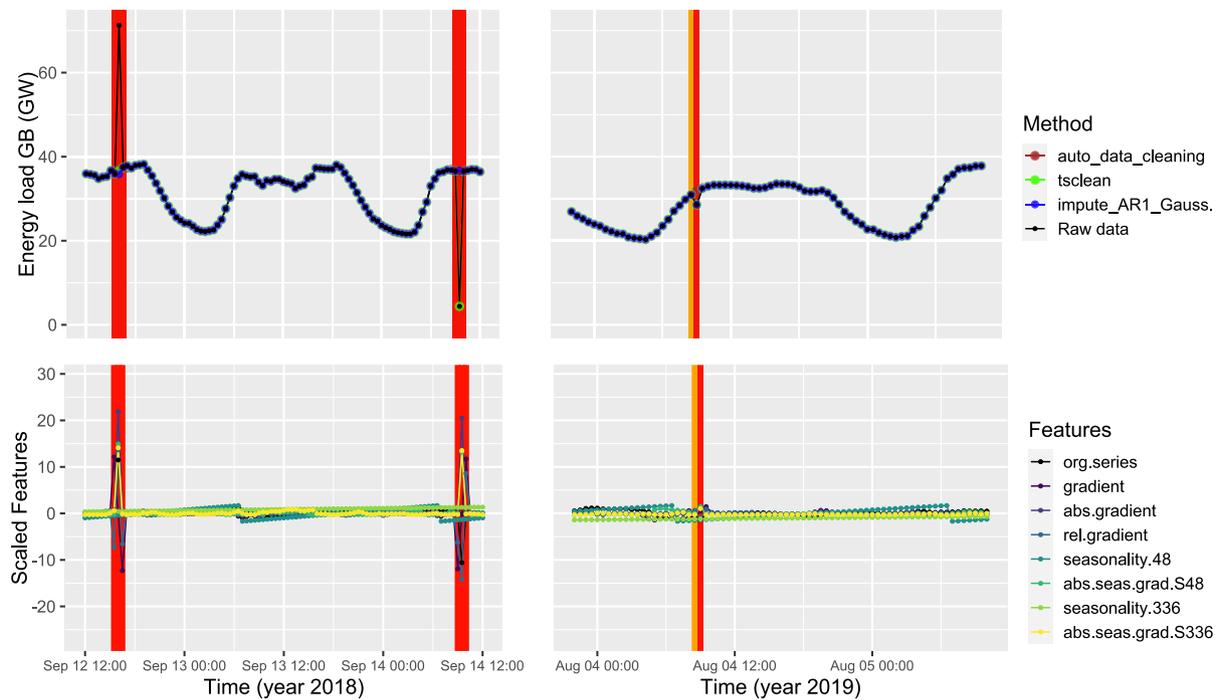


Fig. 6. Exemplary comparison of the proposed function with competitors from literature on the electricity load in Great Britain. The upper plots display the cleaned data of the considered methods and the lower plots the feature space used for outlier detection of the `detect_outliers` function. The coloured vertical lines indicate the probability of a specific observation being outlying data, ranging from dark red (high probability) to yellow (low probability). Probabilities below 0.001 are not coloured. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The software has been already used in the purpose of academic research [12] and industrial projects in the area of power plant electricity generation. Moreover, the software is written in such a manner that it can be utilized on a daily basis with almost no effort. The package is open-source, available online in the CRAN repository, well-documented and constantly developed. The presented examples have shown that the package might be a powerful tool in various time series data preprocessing exercises.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This article was partially supported by the German Research Foundation (DFG, Germany) and the National Science Center (NCN, Poland) through BEETHOVEN grant no. 2016/23/G/HS4/01005. Moreover, we acknowledge support by the Open Access Publication Fund of the University of Duisburg-Essen.

References

[1] Moritz S, Sardá A, Bartz-Beielstein T, Zaefferer M, Stork J. Comparison of different methods for univariate time series imputation in R. 2015, [arXiv: 1510.03924](https://arxiv.org/abs/1510.03924).
 [2] Honaker J, King G, Blackwell M. Amelia II: A program for missing data. *J Stat Softw* 2011;45(7):1–47, <http://www.jstatsoft.org/v45/i07/>.
 [3] van Buuren S, Groothuis-Oudshoorn K. mice: Multivariate imputation by chained equations in R. *J Stat Softw* 2011;45(3):1–67, <https://www.jstatsoft.org/v45/i03/>.
 [4] Kowarik A, Templ M. Imputation with the R package VIM. *J Stat Softw* 2016;74(7):1–16. <http://dx.doi.org/10.18637/jss.v074.i07>.

[5] Moritz S, Bartz-Beielstein T. imputeTS: Time Series Missing Value Imputation in R. *R J* 2017;9(1):207–18. <http://dx.doi.org/10.32614/RJ-2017-009>.
 [6] Zeileis A, Grothendieck G. zoo: S3 infrastructure for regular and irregular time series. *J Stat Softw* 2005;14(6):1–27. <http://dx.doi.org/10.18637/jss.v014.i06>.
 [7] Hyndman RJ, Khandakar Y. Automatic time series forecasting: The forecast package for R. *J Stat Softw* 2008;26(3):1–22, <http://www.jstatsoft.org/article/view/v027i03>.
 [8] Fraley C. HDoutliers: Leland Wilkinson’s Algorithm for Detecting Multidimensional Outliers, R package version 1.0.3. 2020, <https://CRAN.R-project.org/package=HDoutliers>.
 [9] Filzmoser P, Gschwandtner M. mvoutlier: Multivariate Outlier Detection Based on Robust Methods, R package version 2.0.9. 2018, <https://CRAN.R-project.org/package=mvoutlier>.
 [10] de Lacalle JL. tsoutliers: Detection of Outliers in Time Series, R package version 0.6-8. 2019, <https://CRAN.R-project.org/package=tsoutliers>.
 [11] Liu J, Palomar DP. imputeFin: Imputation of Financial Time Series with Missing Values, R package version 0.1.2. 2021, <https://CRAN.R-project.org/package=imputeFin>.
 [12] Narajewski M, Ziel F. Ensemble forecasting for intraday electricity prices: Simulating trajectories. *Appl Energy* 2020;279:115801. <http://dx.doi.org/10.1016/j.apenergy.2020.115801>.
 [13] Friedman J, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *J Stat Softw* 2010;33(1):1–22, <http://www.jstatsoft.org/v33/i01/>.
 [14] Koenker R. quantreg: Quantile Regression, R package version 5.61. 2020, <https://CRAN.R-project.org/package=quantreg>.
 [15] Fraley C, Raftery AE. Model-based clustering, discriminant analysis density estimation. *J Amer Statist Assoc* 2002;97(458):611–31. <http://dx.doi.org/10.1198/016214502760047131>.
 [16] Scrucca L, Fop M, Murphy TB, Raftery AE. mclust 5: Clustering, classification and density estimation using Gaussian finite mixture models. *R J* 2016;8(1):289–317. <http://dx.doi.org/10.32614/RJ-2016-021>.
 [17] ENTSOE. ENTSOE Transparency; 2020. <https://transparency.entsoe.eu/>. [Accessed: 2020-09-30].
 [18] Taylor S, Letham B. prophet: Automatic Forecasting Procedure, R package version 0.6.1. 2020, <https://CRAN.R-project.org/package=prophet>.
 [19] Dancho M, Vaughan D. timetk: A Tool Kit for Working with Time Series in R, R package version 2.6.1. 2021, <https://CRAN.R-project.org/package=timetk>.

DuEPublico

Duisburg-Essen Publications online

UNIVERSITÄT
DUISBURG
ESSEN

Offen im Denken

ub | universitäts
bibliothek

This text is made available via DuEPublico, the institutional repository of the University of Duisburg-Essen. This version may eventually differ from another version distributed by a commercial publisher.

DOI: 10.1016/j.softx.2021.100809

URN: urn:nbn:de:hbz:465-20220720-135239-0



This work may be used under a Creative Commons Attribution 4.0 License (CC BY 4.0).