International Journal of Management, Knowledge and Learning | ISSN 2232-5697

Anomaly Detection in Time Series for Smart Agriculture

Vladislav Bína
Prague University of Economics and Business, Faculty of Management, Czechia
vladislav.bina@vse.cz

Jitka Bartošová
Prague University of Economics and Business, Faculty of Management, Czechia
jitka.bartosova@vse.cz

Vladimír Přibyl
Prague University of Economics and Business, Faculty of Management, Czechia
vladimir.pribyl@vse.cz

Purpose: Enormous amounts of data from sensors and automatic systems are available to contemporary farms and agricultural enterprise managers. Therefore, the automated monitoring of the fundamental processes is becoming more and more necessary. The manager of a modern agricultural company is responsible for his decisions but must have accurate and transparent support in the decision-making process. The reports based on wide-scale mathematical and statistical approaches integrated into the information system can provide a suitable hierarchy of alerts concerning important indicators. This paper provides a brief overview of methods used for anomaly detection in time series and basic (and open) implementation possibilities.

Study design/methodology/approach: A brief overview of the most essential and sufficiently flexible modern methods for detecting an anomaly in time series with a focus on agriculture.

Findings: The paper can serve as a starting point for creating alerting in agricultural information systems to support managerial decision-making. In connection to the high-speed development in data and information spheres in smart agriculture within the last two decades, the new approaches for prediction in time series and search for unexpected fluctuations in their course. Besides the traditional mathematical and statistical techniques, new methods emerge which are tailor-made to the internet environment and environment of social networks.

Originality/value: This basic overview can be used for a primary orientation of the problem and can serve as a basis for creating an automated alerting system. It is always necessary to remember the domain of use, economic relations, structure and hierarchy of time series and their periodicity, seasonality and cycles.

Introduction

The rapid development of modern information technologies and communication manifests in various fields of modern economies, e.g. in the form of Industry 4.0. But modern agriculture also increasingly relies on sensors, computers, data flow, clouds and communication.

Among many others, let us mention precision agriculture, satellite crop monitoring, automatic feeding, robotic milking or livestock tracking for reproduction and dairy management. The broad field of anomaly detection approaches can be divided into general methods employed on cross-sectional data and methods applied on time series. Nowadays, the abundance of sensors and data sources in agriculture is present in all production fields, at least in more developed countries, possibly only with the exception of small farms. To provide some illustration, the cross-sectional approach can comprise, e.g. systems of soil sensors for detecting the moisture and levels of nutrients using the IoT principles (see e.g. Jeba, Lingareddy, Kowsalya, Sree & Swetha, 2018). Finally, let us mention that this is a vast and vital field of precision farming, providing a possibility also to reduce the use of pesticides and employ the data from all technologies, including harvester sensors (for a survey, see Sinha & Dhanalakshmi, 2022).

Another important field is robotic milking, where during the process, a vast amount of data is collected and provides information about the condition and illnesses like mastitis (see e.g. Broucek & Tongel, 2017) together with walking and activity sensors detecting changes in gait and feeding behaviour indicating problems with legs (lameness, e.g. Haladjian, Hodaie, Nüske & Brügge, 2017) or reproduction stage (oestrus detection, e.g. Burnett, Madureira, Silper, Fernandes & Cerri, 2017).

The previous examples probably gave some impression of how broad this field can be. The sensors can provide a vast amount of data which is feasible to analyse only in an automated manner and to find outliers or anomalies which usually present some events of interest. Let us mention, in general, some approaches which are usable for anomaly detection in agriculture. They include simple interquartile range techniques, artificial neural networks, deep-learning methods, support vector machines, ARIMA models and long short-term memory, isolation forests and histogram-based outlier score (see, e.g. Moso, Cormier, de Runz, Fouchal & Wandeto, 2021) or possibly angle-based outlier detector, local outlier factor, k-nearest neighbours, cluster-based local outlier factor, etc. (see de Souza et al., 2020).

It appears that one of the most critical tasks is to provide some kind of warning or automated alerting when some characteristic changes in time after a rather long and successful run without any problems. Therefore, in the following text, we will provide a brief survey of the essential approaches usable for anomaly detection in time series.

Automated alerting using algorithms for outlier detection

The search for outliers and anomalies in data is a complex and non-trivial task not only from the perspective of human decision-makers but also in the case of computer support of managerial decision-making (Tabesh, Mousavidin & Hasani, 2019). Humans usually have difficulties handling massive amounts of data; on the other hand, an experienced person easily identifies "normal" values and distinguishes them from apparent "anomalies". In a way, the complementary algorithmic approach relatively easily tackles a significant amount of data, but the confidence of anomaly detection and its interpretation often lag. Automated detection of outliers is a field cultivated for decades. Among the first well, usable methods are the approaches of Holt-Winters exponential smoothing (Chatfield, 1978) and ARIMA methods (Box & Jenkins, 1970).

The growing demand for a solution from the field of anomaly detection caused significant development not only in theory but also in the domain of BI (business intelligence) components. Their approaches for anomaly detection were developed by major companies running social networks like Twitter (Kejariwal, 2015), LinkedIn (Luminol, 2018), Facebook (Taylor & Letham, 2018), and Google (Brodersen, Gallusser, Koehler, Remy & Scott, 2015), etc. These algorithms for detecting anomalies in time series were usually available at least in some preliminary variants and among promising and frequently used approach counts, namely Facebook's Prophet. Besides this, automated anomaly detection was integrated into the tools of several analytical BI companies. Let us mention Anodot (Toledano, Cohen, Ben-Simhon & Tadeski, 2018 and https://www.anodot.com/) or the part of the BI solution by Microsoft: Microsoft Cognitive Service based on Spectral Residuals (SR) and Convolutional Neural Networks (CNN) (see Ren et al., 2019 and https://azure.microsoft.com/en-us/services/cognitive-services/), or, e.g. products of Avora (https://avora.com/) and CrunchMetrics (https://www.crunchmetrics.ai/).

It appears that the emphasis or analysis of unusual and interesting events (which in a way deviate from previous (or expected) run of time series) becomes an important activity in all spheres where a vast amount of data is processed. This analysis focuses on metrics supporting the control of important processes in a company, customer relations or financial characteristics and determines the properties of the analyses time series. They usually aggregate data on a daily basis or sparsely and frequently show weekly and other (e.g., yearly) seasonal components.

Classification of anomaly detection types in time series

Detection methods can assess time series as single entities (as *a univariate time series*) or as a whole (multivariate time series). In this part, we will deal with univariate series; in the last section, we will sketch possibilities for the multivariate case. Besides this, we can consider anomalies of point type, some longer anomalous subsequence and the whole anomalous time series (in the case of considering multivariate time series). This paper will focus on point type anomalies, for other types refer, e.g. to a review paper Blázquez-García, Conde, Mori & Lozano (2021).

The anomaly detection methods can be further divided into approaches incorporating temporal information and approaches primarily not including the time information. Yet another attribute of considered approaches is the possibility of real-time anomaly detection during transaction processing, which in the vast majority of methods implies the necessity of re-estimating the time series model. This means that the detection is based on a prediction of future value in time series, contrary to the approaches based on a time series model that identifies anomalies by comparing the whole run of real-time series with its theoretical model.

The identification of anomalies in particular univariate time series thus can be performed using *model approaches* (modelling values within the series sequence or using prediction of new value) and *methods based on the density* of particular values. The first category contains methods based on regression, decomposition (e.g. STL and ETS) and ARIMA approaches. The members of the second category are, e.g. isolation forests or neural network approaches.

Statistical approaches to anomaly detection

The typical course of time series contains plenty of typical patterns, and basic approaches to time series modelling are based on a division of a sequel to essential components. These are trend, season and cycle components. The time series decomposition usually unites the trend and cycle into one "trend-cycle" component, which is generally (and somewhat inaccurately) denoted as a trend. The typical way is thus a decomposition into a trend (meaning trend-cycle), season and residual component (the part of data variability not explained by the model). In the case of series with higher frequencies (at least daily), the model can contain more seasonal components corresponding to different periodicities.

ETS decomposition models

ETS decomposition is based on the methodology developed at the end of the 1950s and uses exponential smoothing together with weighted averages of previous observations. The exponential smoothing itself (Brown, 1959) was extended by the inclusion of a linear model of trend (Holt, 1957), which was further supplemented by the modelling of seasonal components (Winters, 1960). During this time, a broad scale of ETS models was developed. Their summary can be found, e.g. in Hyndman & Khandakar (2008). Identification of anomalies employs a modelled value of time series and prediction intervals with a confidence level chosen by a user.

STL decomposition models

One of the universal and relatively robust approaches is an STL decomposition (Cleveland, Cleveland, McRae & Terpenning, 1990) which is based on the decomposition of seasonal and trend parts where the trend component is computed using the LOESS estimate (Cleveland, 1979). Unlike many other approaches (like X11 and SEATS decompositions), this decomposition can incorporate a more general type of seasonal components (not only quarterly and monthly) and is relatively robust with respect to the outliers. Therefore, the particular observations do not significantly change the components' estimates (Hyndman & Athanasopoulos, 2018). Simultaneously, the seasonal part can vary over time, and the speed of change and the smoothness of the trend curve can be set up in some implementations. Thanks to the LOESS estimate, the method is very successful in describing trend-cycle, i.e. when the long-run trend is apparent. These are the reasons for ranking this (and subsequent method) among the basic approaches for anomaly detection.

Twitter decomposition models

This approach is very similar to STL decomposition. It handles the seasonal component similarly and is implemented in the package AnomalyDetection by the Twitter company (Hochenbaum, Vallis & Kejariwal, 2017). However, unlike the STL approach, it estimates the trend part using the (local) median computation. Therefore, it works well and quickly in cases when the long-term trend is less important and short-term seasonal components dominate.

Both STL and Twitter decomposition can be used together with two types of approaches for anomaly point identification. These are IQR (Inner Quartile Range) a GESD (Generalized Extreme Studentized Deviate test). The IQR method uses an interquartile range and in basic settings are, anomaly points identified using the threshold given by three multiples of the interquartile range. This fast approach does not employ an iterative evaluation; thus, its weak adaptability ranks among the method's weaknesses. On the other hand, a more flexible GESD approach can handle significant anomalies that are iteratively eliminated during the computation. Therefore, this approach is computationally more demanding than IQR, which is substantial, particularly in the case of a higher amount of time series with many observations.

Automated ARIMA models

Similarly to exponential smoothing, the ARIMA ranks among the most extensively used methods in time series modelling. However, unlike the approaches based on decomposition, the ARIMA methods do not employ simple models of trend and season but the autocorrelation characteristics of modelled data.

Autoregressive Integrated Moving Average models incorporate the autoregressive character of the model. I.e., the dependence of actually observed value on the given number of older values, the integrated part provided by the differences in such a way that the character of the series is stationary, and also includes the part of moving averages employed for modelling of actual values using the linear combination of error terms from several previous values. Approaches for modelling seasonal time series similarly include the components mentioned above for the sake of modelling the seasonal component.

Since the intended purpose of employment demands the highest level of automation, the approach of "automated ARIMA" is used (see Hyndman & Khandakar, 2008). This approach is based on a unit root together with minimisation of AICc and MLE criteria for the choice of optimal ARIMA model. Again, it is a universal and broadly usable approach.

Prophet models

Facebook Prophet models (Taylor & Letham, 2018) also rank among the additive approaches based on decomposition. They employ piecewise linear curves and logistic transition for modelling non-periodic components with the model of the seasonal part and additional information inserted by the user and containing ranges of holidays leading to deviations in the run of a time series. The non-periodic component is estimated using the regression principles; the seasonal part is modelled using the Holt-Winters exponential smoothing (Chatfield, 1978). Since this approach serves to model and predicts time series, its use for anomaly detection is straightforward. Prophet ranks among very promising approaches because of its robustness in handling the missing values, significant seasonal changes, fluctuations during holidays, and significant shifts in trend.

TBATS models

The above-mentioned decomposition models considered only simple types of seasonal components. Among the approaches capable of simultaneously incorporating more different seasonal components (in our case, as a rule weekly and monthly periodic) is the TBATS model (de Livera, Hyndman & Snyder, 2011). A combination of exponential smoothing with Fourier constituents for modelling periodic components, Box-Cox transforms, and ARMA errors are used for modelling. Unlike the approaches of harmonic regression, the TBATS models allow for a slight change in seasonal components with the development of series in time. This complexity and flexibility of the model can be beneficial, but its usability is strongly limited by the fact that for more extended time series is, the estimation of the model computationally very demanding.

Approaches from the sphere of artificial intelligence

The anomaly detection tasks usually employ approaches of unsupervised learning, e.g. modern and frequently used isolation forests, specialised methods ranking among artificial neural networks approaches (recurrent neural networks and enhanced deep-learning LSTM autoencoders), or the google approach of CausalImpact and use of convolutional neural networks within the Microsoft solution.

A practical limitation of these methods can be a lack of open implementations usable for incorporation into own projects and problems in the compatibility of chosen software architecture.

Isolation forests

Isolation forests (see Liu, Ting & Zhou, 2012) comprise an approach of unsupervised learning generally applied to the sphere of anomaly detection. The methodological principle is based on building the set of isolation trees and averaging over this set to obtain a mean anomaly score of particular observations characterising the length of the path leading to this observation. The basic idea is that the "normal" values are similar to many other values and thus have many other values in their neighbourhood. To distinguish such similar values, we must pass through many branchings in the vast number of generated trees. Oppositely, the anomalous point can be usually separated during the passing of only several branchings. Mean depth of passage to the particular value after suitable normalisation provides a possibility to identify outliers or anomalies.

In case of search for anomalies in time series, usually, a suitable augmentation of the method must be used in order to incorporate the temporal character of data. One possibility is using a sliding window (see Ding & Fei, 2013) to consider the locality or temporal relation. However, the simplest variant is the use of one or more delayed (lagged) time series as an additional dimension for detection with the basic algorithm.

LSTM autoencoders

Artificial neural networks called autoencoders represent another approach using unsupervised learning and are employed as efficient methods for data compression. Besides this, the methodology is well usable for dimensionality reduction, generalisation and image processing. Anomaly detection using autoencoders is based on the simple idea that a well-trained autoencoder can reconstruct the data sufficiently. On the contrary, the autoencoder usually fails in reconstructing anomalous data.

In the case of time series and other sequential data, the LSTM (Long Short-Term Memory) autoencoders can be used (see Gers, Schmidhuber & Cummins, 1999, or Kulanuwat et al., 2021), which are not loaded by the drawbacks of recurrent neural networks lying in failures of learning methods on vanishing or growing gradients of long time series.

CausalImpact

CausalImpact (Brodersen et al., 2015) is a Google company approach based on econometric principles of Bayesian structural equations for time series. The approach uses time series from a control group and generates a basic synthetic time series. Anomalies are then found by comparing new (so-called test series) with this generated artificial series.

Modern implementation of selected approaches

In the previous section, we provided a brief survey of methods capable of detecting time series anomalies and selected a set of advanced and promising approaches. Here we bring at least basic possibilities of implementing these methods within the environment of statistical software R (R Core Team, 2021) with a focus on ETS, ARIMA, STL, Twitter and IF, partially also Prophet and TBATS.

Forecast library

A starting point for different types of models is a forecasting library (Hyndman et al., 2018). This library contains, among others also, modern and robust implementation of ARIMA approaches (including the automated creation of models by procedure auto.arima), ETS methods, STL and Twitter decomposition and implementation of TBATS models. Therefore, it is possible to use the estimated models of one-dimensional time series and generate prediction intervals of modelled values usable for identifying anomalies.

Anomalize library

The implementation of STL a Twitter approach is also available in the anomalize library (Dancho & Vaughan, 2020), which allows IQR and GESD variants and their parametrisation for detecting abnormal points.

Prophet library

The use of a robust Facebook Prophet approach capable of handling holidays and changes in trends is possible in the R environment thanks to the open library prophet (Taylor & Letham, 2020).

Solitude library as an implementation of isolation forests

An R package solitude (Liu, Ting & Zhou, 2012) can be successfully used to implement the isolation forest approach. It is a modern implementation well compatible with advanced R data structures and can employ parallelisation of isolation tree building. However, a significant drawback of this method is rather incorrect handling of the time dimension, which can be partially overcome using lagged time series as an additional dimension. Another weakness lies in a lack of standardisation of threshold score separation anomalous values. Therefore, it is necessary to propose suitable standardisation before the use of this method in the field of anomaly detection.

Modern handling of temporal data: lubridate and tsibble

Flexible and correct handling of different time and date formats together with conversion between text and date formats, extraction and setting of components, arithmetic operations and rounding, but also work with time zones, leap years, and daylight savings. These are the basic feature implemented in modern library lubridate (Grolemund & Wickham, 2011).

A modern method for manipulating dates within an R environment is a collection of tidyverse packages (see Wickham et al., 2019 or Wickham & Grolemund, 2016). It can import data and perform data manipulation, cleaning and transformation and visualisation with the aim of domain-specific packages from this collection. This set can be augmented by a package for effectively handling temporal data named tsibble (Wang, Cook & Hyndman, 2020). It extends the properties of general tibble data structures while using the temporal data as key data columns and thus uses the ordering given by the flow of time.

Basic possibilities for selection of significant anomalies and dimensionality reduction

The above-described automated approaches for anomaly detection and deployment on vast amounts of data created and archived in general practice in most of the bigger farms and companies, not only in agriculture, can be a source of many fluctuations and anomalies as a basis for user alert. This obviously brings an enormous workload in control of reported outputs or – on the contrary – the risk of overlooking some crucial alerts among the flood of less critical notifications. For clarification, higher user-friendliness and better usability of the system the following three possibilities can be used.

Economic interconnections

A starting point for grouping alerts based on economic indicators can be a range of economic relations and connections from basic to more complex and also from deterministic to stochastic. A significant disadvantage of this approach is that its strict and consistent application is frequently uneasy due to the specificity of the indicator set in different domains of use and thus the necessity to implement a system tailor-made. Nevertheless, typical and fundamental economic relations like sales fluctuations and unnecessarily reported fluctuations in VAT could be easily prevented.

Methods searching related anomalies based on statistical correlations

Statistical approaches can continuously search for closely related time series that cannot be discovered using well-known economic interconnections. The advantage of this approach is the possibility to automatise it to some extent. For example, the associated time series can be discovered using correlation and cointegration approaches or more sophisticated methods like Granger analysis (Granger, 1969) or Convergent cross-mapping (Cenys, Lasiene & Pyragas, 1991).

Hierarchical models of time series and rolling-up of interconnected anomalies

In most of the more prominent companies, some structure of centres and subsidiaries exists and, e.g. from the perspective of a range of products frequently, a (usually) multilevel hierarchical structure of products can be identified. Simple hierarchy or even multilevel grouping of time series can be modelled using the approaches of Wickramasuriya, Athanasopoulos & Hyndman (2019). A significant advantage of this methodology is that models of the respective time series are in correspondence, which means that, e.g. the sum of predicted values in lower levels of hierarchy fully corresponds to the predicted value on higher levels. These methods allow the user to go through the hierarchy from higher to deeper levels and gradually expand detected anomalies according to the demand. Thus, the higher-level manager can see a report with only substantial anomalies and an appropriate indication of the number of anomalies in nested levels.

Conclusion

This paper presents a basic summary of methods used for the automated detection of anomalies in smart agriculture. The paper can serve as a starting point for creating alerting in agricultural information systems to support managerial decision-making. In connection to the high-speed development in data and information spheres in smart agriculture within the last two decades, the new approaches for prediction in time series and search for unexpected fluctuations in their course. Besides the traditional mathematical and statistical techniques, new methods emerge which are tailor-made to the internet environment and environment of social networks. This basic overview can be used for a primary orientation in the problem and can serve as a basis for creating an automated alerting system. It is always necessary to remember the domain of use, economic relations, structure and hierarchy of time series and their periodicity, seasonality and cycles. The choice of methods must always be adapted to these circumstances.

Acknowledgement

This research was supported by the Technology Agency of the Czech Republic in the TREND programme under project FW01010606: Adaptive alerting business intelligence system.

References

Blázquez-García, A., Conde, A., Mori, U., & Lozano, J. A. (2021). A review on outlier/anomaly detection in time series data. *ACM Computing Surveys (CSUR), 54*(3), 1-33.

Box, G., & Jenkins, G. (1970). *Time Series Analysis: Forecasting and Control.* San Francisco: Holden-Day.

Brodersen, K. H., Gallusser, F., Koehler, J., Remy, N., & Scott, S. L. (2015). Inferring causal impact using Bayesian structural time-series models. *Annals of Applied Statistics, 9*(1), 247-274.

Broucek, J., & Tongel, P. (2017). Robotic Milking and Dairy Cows Behaviour. In Proceedings of 2017 *International Conference on Control, Artificial Intelligence, Robotics & Optimization, ICCAIRO, 2017* (pp. 33-38), Prague. doi: 10.1109/ICCAIRO.2017.16.

Brown, R. G. (1959). *Statistical forecasting for inventory control.* McGraw/Hill.

Burnett, T. A., Madureira, A. M., Silper, B. F., Fernandes, A. C. C., & Cerri, R. L. (2017). Integrating an automated activity monitor into an artificial insemination program and the associated risk factors affecting reproductive performance of dairy cows. *Journal of dairy science, 100*(6), 5005-5018.

Chatfield, C. (1978). The Holt□winters forecasting procedure. *Journal of the Royal Statistical Society: Series C (Applied Statistics), 27*(3), 264-279.

Cleveland, R. B., Cleveland, W. S., McRae, J. E., & Terpenning, I. (1990). STL: A seasonal-trend decomposition. *Journal of official statistics, 6*(1), 3-73.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American statistical association, 74*(368), 829-836.

Cenys, A., Lasiene, G., & Pyragas, K. (1991). Estimation of interrelation between chaotic observables. *Physica D: Nonlinear Phenomena, 52*(2-3), 332-337.

Dancho, M., & Vaughan, D. (2020). *anomalize: Tidy Anomaly Detection. R package version 0.2.2.* Retrieved from: https://CRAN.R-project.org/package=anomalize

Ding, Z., & Fei, M. (2013). An anomaly detection approach based on isolation forest algorithm for streaming data using sliding window. *IFAC Proceedings Volumes, 46*(20), 12-17.

Gers, F. A., Schmidhuber, J., & Cummins, F. (1999). Learning to forget: Continual prediction with LSTM. In *9th International Conference on Artificial Neural Networks: ICANN '99* (pp. 850-855).

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. Econometrica: *Journal of the Econometric Society, 37*(3), 424-438.

Grolemund, G., & Wickham, H. (2011). Dates and Times Made Easy with lubridate. *Journal of Statistical Software, 40*(3), 1-25. Retrieved from: https://www.jstatsoft.org/v40/i03/

Haladjian, J., Hodaie, Z., Nüske, S., & Brügge, B. (2017, November). Gait anomaly detection in dairy cattle. In *Proceedings of the Fourth International Conference on Animal-Computer Interaction* (pp. 1-8).

Hochenbaum, J., Vallis, O. S., & Kejariwal, A. (2017). *Automatic anomaly detection in the cloud via statistical learning.* Retrieved from: arXiv preprint arXiv:1704.07706

Holt, C. C. (1957). *Forecasting Trends and Seasonals by Exponentially Weighted Averages, 52*, Carnegie Institute of Technology, Pittsburgh, Pa, USA.

Hyndman, R. J., Athanasopoulos, G. (2018). *Forecasting: principles and practice.* OTexts. Retrieved from: https://otexts.com/fpp3/

Hyndman, R. J., Athanasopoulos, G., Bergmeir, C., Caceres, G., Chhay, L., O'Hara-Wild, M., Petropoulos, F., Razbash, S., Wang, E. Yasmeen, F. (2018). *forecast: Forecasting functions for time series and linear models.* Retrieved from: https://pkg.robjhyndman.com/forecast/

Hyndman, R. J., & Khandakar, Y. (2008). Automatic time series forecasting: the forecast package for R. *Journal of statistical software,* 27(3), 1-22.

Jeba, N., Lingareddy, S. C., Kowsalya, P., Sree, S. M., & Swetha, S. (2018). Anomaly detection to enhance crop productivity in smart farming. *International Journal of Pure and Applied Mathematics, 120*(6), 11-503.

Kejariwal, A. (2015). Introducing practical and robust anomaly detection in a time series. *Twitter Engineering Blog.* Retrieved from: https://blog.twitter.com/engineering/en_us/a/2015/introducing-practical-and-robust-anomaly-detection-in-a-time-series.html

Kulanuwat, L., Chantrapornchai, C., Maleewong, M., Wongchaisuwat, P., Wimala, S., Sarinnapakorn, K., Boonya-aroonnet, S. (2021). Anomaly detection using a sliding window technique and data imputation with machine learning for hydrological time series. *Water, 13*(13), 1862.

de Livera, A. M., Hyndman, R. J., & Snyder, R. D. (2011). Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American statistical association, 106*(496), 1513-1527.

Liu, F. T., Ting, K. M., & Zhou, Z. H. (2012). Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD), 6*(1), 1-39.

Luminol (2018). *Luminol: LinkedIn's Anomaly Detection and Correlation Library.* Retrieved from: https://github.com/linkedin/luminol/

Moso, J. C., Cormier, S., de Runz, C., Fouchal, H., & Wandeto, J. M. (2021). Anomaly Detection on Data Streams for Smart Agriculture. *Agriculture, 11*(11), 1083.

Ren, H., Xu, B., Wang, Y., Yi, C., Huang, C., Kou, X., Xing, T., Yang, M., Tong, J., & Zhang, Q. (2019). Time-series anomaly detection service at Microsoft. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 3009-3017).

Sinha, B. B., & Dhanalakshmi, R. (2022). Recent advancements and challenges of Internet of Things in smart agriculture: A survey. *Future Generation Computer Systems, 126*, 169-184.

de Souza, P. S. S., Rubin, F. P., Hohemberger, R., Ferreto, T. C., Lorenzon, A. F., Luizelli, M. C., & Rossi, F. D. (2020). Detecting abnormal sensors via machine learning: An IoT farming WSN-based architecture case study. *Measurement, 164*, 108042.

Tabesh, P., Mousavidin, E., & Hasani, S. (2019). Implementing big data strategies: A managerial perspective. *Business Horizons, 62*(3), 347-358.

Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician, 72*(1), 37-45.

Toledano, M., Cohen, I., Ben-Simhon, Y., & Tadeski, I. (2018). Real-time anomaly detection system for time series at scale. In *KDD 2017 Workshop on Anomaly Detection in Finance* (pp. 56-65).

Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L. D., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T. L., Miller, E., Bache, S. M., Müller, K., Ooms, J., Robinson, D., Seidel, D. P., Spinu, V., Takahashi, K., Vaughan, D., Wilke, C., Woo, K., & Yutani, H. (2019). Welcome to the tidyverse. *Journal of Open Source Software, 4*(43), 1686.

Wang, E., Cook, D., & Hyndman, R. J. (2020). A new tidy data structure to support exploration and modeling of temporal data., *Journal of Computational and Graphical Statistics, 29*(3), 466-478.

Wickham, H., & Grolemund, G. (2016). *R for data science: import, tidy, transform, visualise, and model data.* O'Reilly Media, Inc.

Wickramasuriya, S. L., Athanasopoulos, G., & Hyndman, R. J. (2019). Optimal forecast reconciliation for hierarchical and grouped time series through trace minimisation. *Journal of the American Statistical Association, 114.526,* 804-819.

Winters, P. R. (1960). Forecasting sales by exponentially weighted moving averages. *Management science, 6*(3), 324-342.