Anomaly detection to predict failures in server systems

Mikhail Dubrovin[0000-0002-8580-1303], Igor Gluhih[0000-0002-0683-6138], and Yuri Karyakin[0000-0003-2346-402X]

Tyumen State University, 6, Volodarskogo ave., Tyumen, 625003, Russian Federation
mikle1203@yandex.ru

Abstract. The article discusses the problem of predicting failures to ensure the uninterrupted state of the server and the application of anomaly detection methodology to solve this problem. Some solutions in this area are briefly analyzed and the advantage of the method based on the Holt-Winters forecasting model is described. A mathematical formulation of the problem of detecting anomalies and a formalized description of the method for solving the research objectives are carried out. Several shortcomings of the standard method for detecting anomalies in the operation of server systems have been identified. Several additions are proposed that allow you to adapt the method for solving work goals and reduce the number of false positives. To improve the forecasting accuracy, point anomalies entering the model input are additionally smoothed by a weighted moving average. To eliminate redundant detections associated with the simultaneous appearance of anomalies, a comprehensive assessment of the server state is introduced. To exclude false alarms associated with noisy data, only those anomalous events are recognized, in which an abnormal state of the server remains for a period of time. Computational experiments were carried out to evaluate the resulting improved method. It is concluded that the proposed additions make it possible to improve the forecasting accuracy of the model and reduce the number of false positives of the method, and the method can be used for early detection of gradual failures in the operation of server systems.

Keywords: Anomaly detection, Proactive monitoring, Failure prediction, Holt-Winters model, Brutlag method, Server health.

1 Introduction

Server systems are often one of the central elements in the corporate information system of enterprises. Server failures can lead to the loss of valuable information and significant financial costs, which necessitates ensuring the smooth operation of such systems. One of the relevant approaches for these purposes is the introduction of proactive monitoring systems [1], which are aimed at predicting failures in the operation of server systems, which allows you to eliminate possible problems at the stage of their inception.

One of the tools for implementing proactive monitoring systems are anomaly detection methods. An anomaly is understood as a piece of data in which the behavior of an object significantly differs from the standard or expected behavior. In this context, the task is reduced to the detection of contextual and collective new anomalies in real time [2]. This article discusses only hardware and software anomalies that are not caused by deliberate actions of intruders. Prediction of gradual failures in this case is based on the assumption that if an anomaly occurs on the server, after a while, its functioning may be disrupted.

Analysis of the works shows that there is no universal method for detecting anomalies for solving any problem [2]. To improve the accuracy of anomaly recognition, it is necessary to take into account the peculiarities of the problem and the data under investigation. Anomaly detection methods based on tagged data [3] are less suitable in this area, since the change in the functional load on the server necessitates constant retraining of the classifier. One Class SVM is used to detect anomalies [4], but for its application it is necessary to have an idea of the percentage of anomalies in datasets, which is not always possible. Choosing this option unjustifiably can lead to gaps or false detections. Methods based on artificial neural networks are quite widespread, from simple perceptrons to complex recurrent networks, for example, LSTM [5]. The disadvantages of ANNs in this area are high computational complexity and black box operation. Because of the latter, it becomes difficult to interpret the results, which is why the localization of the problem in the server operation requires a separate analysis. A number of methods based on forecasting models are also widespread [6]. The disadvantage of models from the ARIMA family [7] is that the models work with stationary data, and converting the server parameter values to a stationary form according to the Box-Jenkins methodology can be quite laborious.

The paper considers an anomaly detection method based on the Holt-Winters forecasting model [8]. To use the method, no tagged data is required, which expands the possible options for its use. The method does not require preliminary processing of the input data, which is an advantage for analyzing many server parameters. The method has a relatively low computational complexity, which makes it possible to process information about a large number of server parameters in short periods and thereby obtain up-to-date information about the state of the object. The results of the method are clear and amenable to interpretation. More details on the application of the Holt-Winters model for predicting the state of servers are described in another article by the authors [9], this work is based on the results of the previous study.

The purpose of this work is to study and supplement the anomaly detection method based on the Holt-Winters model for predicting failures in server systems. Within the framework of this goal, the article first provides a mathematical formulation of the problem of detecting anomalies and a formalized description of the application of the problem to the server's operation. Then, a number of additions to the standard anomaly detection method are proposed, which allow it to be adapted to achieve the set goal and to reduce the number of redundant detections. Taking into account the additions, a modified method for detecting anomalies is described. After that, computational experiments are carried out to evaluate the proposed additions. At the last stage, the results were discussed and tasks for further research were proposed.

2 Materials and methods

Each corporate server S can be described with a list of parameters P characterizing its state, $P = \{x_1, x_2, .., x_n\}$. Then the state of the server Z_t at a moment in time t is described by a vector of values of its parameters.

$$Z_t = \begin{pmatrix} x_{1,t} \\ \vdots \\ x_{n,t} \end{pmatrix} \quad (1)$$

The standard formulation of the problem of detecting anomalies consists in constructing a functional dependence $f : x_{i,t} \to \{-1,1\}$ such that:

$$x_{i,t} = \begin{cases} normal, f(x_{i,t}) = 1 \\ anomaly, f(x_{i,t}) = -1 \end{cases}, \quad (2)$$

Where $x_{i,t}$ is the value of the *i* parameter at the moment *t*.

To analyze the state of the server, it is necessary to collect statistical information about its work. The values of each parameter $x_i \in P$ known at discrete points in time $t = \overline{1,T}$ represent a time series $X_i = \{x_{i,1}, x_{i,2}, .., x_{i,T}\}$.

To analyze the state of the server, it is necessary to collect statistical information about its work. The values of each parameter known at discrete points in time represent a time series.

The basic technique for detecting anomalies in time series, based on the forecasting model, consists in generating a forecast for each parameter value and calculating the forecast error. A significant deviation of the forecast from the actual value indicates the presence of an anomaly [10].

To simplify further notation, information about an arbitrary parameter $x \mathrm{O} P$ is described without specifying an index. The Holt-Winters multiplicative model was used as a forecasting model. The model is described by the following system of equations [8]:

$$\begin{cases} R_t = a * \dfrac{x_t}{S_{t-L}} + (1-a) * (R_{t-1} + T_{t-1}) \\ T_t = \beta * (R_t - R_{t-1}) + (1-\beta) * T_{t-1} \\ S_t = \gamma * \dfrac{x_t}{S_t} + (1-\gamma) * S_{t-L} \\ \overline{x}_{t+m} = (S_t + m * T_t) * S_{t-L+m} \end{cases}, \quad (3)$$

Where R_t is an exponentially smoothed series; T_t - trend value; S_t - the seasonal component of the series; $\overline{x}_{t+m}$ - forecast for m steps; a, b, g - coefficients of smoothing of the series, trend and seasonality; L - seasonality period. Training of the model (3) consists in the selection of the optimal values a, b, g. The selection is carried

out using an enumeration of possible combinations of parameters in order to minimize the loss function for cross-validation [11].

To detect anomalous values, the Brutlag method is used [12]. At the first stage, the measure of deviation is determined $d_t : d_t = \gamma^* | x_t - \overline{x_t} | + (1-\gamma) * d_{t-L}$. Values $x_t, \overline{x_t}$ and parameters g, L are contained in model (3). Next, the upper and lower boundaries of the predicted value are calculated:

$$\begin{cases} \overline{x}_{\max t} = \overline{x_t} + m * d_{t-L} \\ \overline{x}_{\min t} = x_t - m * d_{t-L} \end{cases}, \tag{4}$$

Where m is the coefficient of the interval width.

As a result, the value x_t is considered abnormal if it is outside the boundaries of the predicted value (4), and normal in the opposite case:

$$f(x_t) = \begin{cases} 1, x_t \in [\overline{x}_{\min t}, \overline{x}_{\max t}] \\ -1, x_t \notin [\overline{x}_{\min t}, \overline{x}_{\max t}] \end{cases}. \tag{5}$$

Using the Brutlag method to detect anomalies has a number of advantages: the administrator does not need to manually create thresholds, the threshold values are not static and are adjusted to the server state, and this method allows you to detect many anomalies. But when applying the method to solve research problems, a number of shortcomings were identified. Below is a number of author's additions that allow us to eliminate the disadvantages indicated below and to adapt the use of the Holt-Winters model and the Brutlag method for detecting anomalies in the operation of server systems and some other types of complex information systems.

Appendix 1. It has been investigated that the Holt-Winters model (3) shows less accurate forecast results on the series with an increased noise level [9]. Reducing the accuracy of predictions can lead to gaps or false positives in the anomaly detection method. To eliminate this disadvantage, the following solution is proposed.

Parameter values that are recognized by the Brutlag method as anomalous, and do not retain their anomalous value in the future, are considered single random outliers in the data. When forming the next forecast, such emissions are additionally smoothed by replacing the actual value with a weighted moving average [13] based on K the previous normal values:

$$x_t^* = \frac{\sum_{i=0}^{K-1} (K-i) * x_{t-i}}{\sum_{i=1}^{K} i},$$

Where x_t^* is the value of the point outlier supplied to the input of the Holt-Winters model. This addition presumably leads to a decrease in the influence of noise on the formation of a forecast and an increase in its accuracy.

Appendix 2. The server is a multi-parameter system; the occurrence of anomalies in its operation can be caused by the exit of several parameters at once outside the permissible intervals. The standard approach identifies each deviation as a new anomaly.

To exclude redundant detections, not the value of an individual parameter is analyzed, but the state of the object as a whole. The state of the server Z_t is considered normal Z_N if all values of the server parameters are in the normal range, and abnormal Z_A if at least one parameter is in the abnormal range.

$$Z_t = \begin{cases} Z_N, \forall x \in P, x_t \in [\underline{x}_{\min t}, \overline{x}_{\max t}] \\ Z_A, \exists x \in P, x_t \notin [\underline{x}_{\min t}, \overline{x}_{\max t}] \end{cases}. \quad (6)$$

Appendix 3. Time series describing server states are not stationary and can be quite noisy [9], which leads to false positives of the anomaly detection method. Such redundant discoveries are not of interest for analyzing server health.

To exclude such alarms, only those anomalous events *A* are recognized, in which the abnormal state of the server remains at the time interval *p*:

$$A = \left[Z_1, .., Z_p \right], \forall Z_t, Z_t = Z_A, t = \overline{1, p}. \quad (7)$$

Taking into account the additions, the modified anomaly detection method consists of the following steps:

- Collect statistical data on the values of parameters that characterize the state of the server;
- Determine the value of seasonality *L* and smoothing parameters *a*, *b*, *g* for model (3), the coefficient of the width of the interval for detecting anomalies;
- Collect new actual values of the server parameters that require analysis;
- For each parameter, generate a forecast based on the history of previous values, taking into account the condition: if there are point anomalies in the history, replace each outlier value with a weighted moving average based on *K* of previous normal values;
- Determine the measure of deviation of the predicted values from the actual ones;
- For each forecast, determine the boundaries of the area of normal values (4);
- For each parameter value, determine the belonging to the area of normal or abnormal data (5);
- Record the anomalous state of the server if there is at least one parameter that is in the area of anomalous values (6);
- Record the occurrence of an anomalous event, while maintaining an anomalous state of the server during *p* time steps (7);
- Repeat from step 3.

To assess the correctness of the proposed improved method for detecting anomalies, computational experiments were carried out, the results of which are presented in the next section of the article.

3 Results

To prepare an experimental study, the values of 4 parameters were recorded from 2 computer servers within 3 weeks with a discreteness of 1 minute. The sufficiency of these parameters for assessing the state of servers is not considered in the article. The statistics for the first two weeks served as a training sample for further experiments. The selection of the optimal smoothing coefficients *a*, *b*, *c* for the Holt-Winters model was made on the training sample. Seasonality parameter *L* corresponds to a weekly period and is equal to 1440 measurements. The parameter of the interval width *m* for the Brutlag method is empirically selected as 6 units. A smaller value leads to a larger number of detections that slightly deviate from the normal range.

On the control sample, for each parameter value, a short-term forecast one step ahead was formed based on the previous history. The accuracy of the obtained model results was assessed using the MAPE metric [14]:

$$MAPE = \frac{1}{N}\sum_{t=1}^{N}\frac{|x_t - \overline{x_t}|}{x_t} * 100, \quad (8)$$

Where x_t is the actual value of the series, $\overline{x_t}$ is the predicted value of the series, N is the number of measurements.

Two options for the formation of forecasts are considered. In the first version, the average accuracy of which is indicated in the table as $MAPE_1$, the classical application of the Holt-Winters model was carried out. The implementation of the second version of the forecasts ($MAPE_2$) was carried out taking into account the first addition described in the previous section. The previously recognized point anomalous values, when fed into the model (3), are replaced by a weighted moving average of the three previous normal values. Information about server parameters, smoothing coefficients and experiment results is presented in Table 1.

According to the results of the experiment, it can be seen that, in general, the Holt-Winters model shows a fairly high accuracy in predicting the values of server parameters. For parameters with a high random component, the forecast accuracy is slightly lower. For such parameters, an increase in accuracy is observed when applying the principle described in the first supplement of the previous section.

For the first server:

$$\begin{gathered} 2.65 \le MAPE_1 \le 9.5, \overline{MAPE_1} = 4.88. \\ 2.65 \le MAPE_2 \le 9.08, \overline{MAPE_2} = 4.77. \end{gathered} \quad (9)$$

For the second server:

$$\begin{gathered} 2.4 \le MAPE_1 \le 12.7, \overline{MAPE_1} = 5.5. \\ 2.4 \le MAPE_2 \le 11.29, \overline{MAPE_2} = 5.13. \end{gathered} \quad (10)$$

Table 1. Results of using the Holt-Winters model.

No	Server	Parameter	L	a	β	γ	$MAPE_1$	$MAPE_2$
1	1	CPU usage	1440	0.7	0.2	0.1	9.5	9.08
2	1	Memory usage	1440	0.55	0.05	0.05	2.65	2.65
3	1	Net. traffic in	1440	0.75	0.05	0.25	3.32	3.32
4	1	Net. traffic out	1440	0.7	0	0.3	4.06	4.04
5	2	CPU usage	1440	0.5	0.1	0.1	12.7	11.29
6	2	Memory usage	1440	0.7	0.1	0.1	4	3.94
7	2	Net. traffic in	1440	0.85	0.05	0.05	2.4	2.4
8	2	Net. traffic out	1440	0.9	0.1	0.1	2.9	2.9

The second part of the pilot study is to analyze the method for detecting anomalies, taking into account the three additions described in the previous section. In the table below, the improved method is indicated as $Method_1$. To test the method empirically, it is laborious and inefficient to expect anomalous events in the server's operation. In this regard, three long-term anomalies with a duration of 4 minutes were programmatically simulated on the control sample for each server. Simulation of the occurrence of anomalous events was carried out by combining the actual values of the parameters and the values of the linear generator function $\widehat{x_t} = k * x_t + b$.

To evaluate the method, the following indicators were used: total number of detected anomalies (All); the number of correctly detected anomalies (TP); the number of missed anomalies (FN); number of false positives on normal data (FP).

Based on the indicators, the following criteria were calculated [15]. Detection accuracy, $Precision = \frac{TP}{TP + FP}$. Completeness, $\mathrm{Re}call = \frac{TP}{TP + FN}$. Integral indicator F-score, $F\text{-}score = 2 * \frac{\mathrm{Pr}ecision * \mathrm{Re}call}{\mathrm{Pr}ecision + \mathrm{Re}call}$.

To compare the results, on the same data, we calculated the criteria for the classical application of the Brutlag method, designated as. The results are presented in Table 2.

Table 2. Results of using the anomaly detection method.

Method	No	Server	All	TP	FN	FP	Precision	Recall	F-score
$Method_1$	1	1	12	12	0	0	1.00	1.00	1.00
	2	2	12	12	0	0	1.00	1.00	1.00
$Method_2$	3	1	21	12	0	9	0.57	1.00	0.73
	4	2	19	12	0	7	0.63	1.00	0.77

Using the described additions, the method recognized only those anomalous values that were simulated by software. The Brutlag method additionally recognized point

random outliers that are caused by the partially stochastic behavior of the server and are not of interest for diagnosing the state of the object.

The occurrence of an anomalous event on the Memory Usage graph of the first server is shown in Figure 1. The blue line on the graph denotes actual values, orange - predicted values. The green and brown lines indicate the upper and lower boundaries of the normal range. The occurrence of an anomalous event is marked in the figure with red dots.

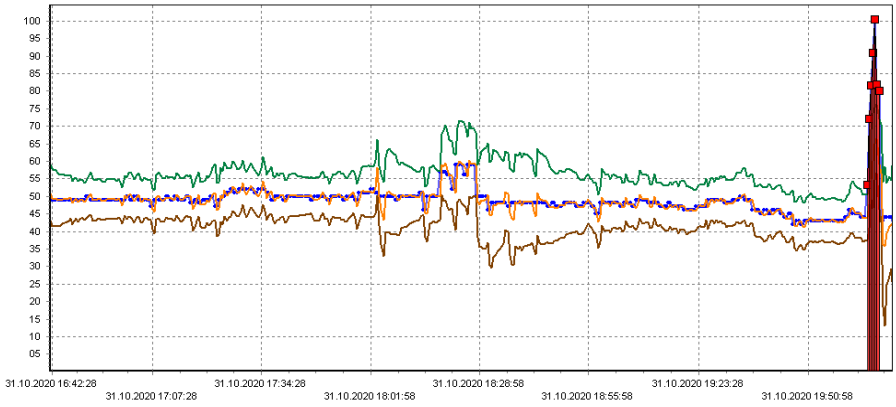

Fig. 1. Occurrence of an anomalous event on the Memory usage graph.

4 Discussion

The development of anomaly detection methods based on the Holt-Winters forecasting model is of practical and research interest, which is confirmed by research in this direction [16-17]. The authors of the work made the following contribution in this area.

Based on the results of the experiment, it can be argued that the Holt-Winters model shows good results when predicting the values of server parameters. A study on the analysis of the SARIMA model for predicting server load demonstrates less accurate forecast results [18]. But in the presence of noise and emissions in the structure of time series, the accuracy of the forecasts formed by the model (3) decreases [9]. To reduce the impact of the random component, the following was undertaken. Recognition of point anomalous values in datasets is performed. At the next generation of the forecast, the detected outliers are additionally smoothed by replacing the actual value with a weighted moving average. The results of the experiment demonstrated that such an addition allows one to obtain more accurate forecasts for time series with an increased level of noise. The improvement in the accuracy of the model for a given loss function for individual parameters is up to 12.5%, the average forecast accuracy increases to 7.2%.

When using the Brutlag method, it is recommended [12] to choose the coefficient of the width of the interval m (4) in the range from 2 to 3 units. Application of the method in practice has shown that setting the coefficient in such a range leads to a large number

of method triggers, slightly deviating from the range of normal values. In order to reduce the excess operation during the experiments, the coefficient was set equal to 6 units.

It was revealed that the Brutlag method generates redundant triggers in the case when several server parameters are simultaneously abnormal. To exclude redundant detections, not the value of an individual parameter is analyzed, but the state of the object as a whole. In addition, the method detects point random outliers that are caused by the partially stochastic behavior of the server and are not of interest for diagnosing its state. To exclude such alarms, only those anomalous events are recognized, in which the anomalous state of the server persists for a time interval. The proposed additions to the method make it possible to reduce the number of false positives, which is demonstrated using a computational experiment.

The application of the modified method makes it possible to recognize contextual collective anomalies, which may indicate the occurrence of gradual failures in the functioning of the server, which makes it possible to eliminate possible problems even at the stage of their inception.

5 Conclusion

The work is based on the results of a previous study by the authors [9], which considered the application of the Holt-Winters model to predict the state of server systems. The aim of the current work is to investigate and supplement an anomaly detection method based on the Holt-Winters model for predicting failures in server systems. At the first stage of the work, a mathematical formulation of the problem of detecting anomalies and a formalized description of the application of the problem to the work of the server are carried out. The existing method of detecting anomalies has a number of shortcomings, on the basis of which a number of additions are proposed that allow it to be adapted to achieve the set goal and reduce the number of redundant detections. To improve forecasting accuracy, single outliers entering the model input are smoothed with a weighted moving average. To eliminate redundant detections associated with the simultaneous appearance of anomalies, a comprehensive assessment of the server state is introduced. To exclude false alarms associated with noisy data, only those anomalous events are recognized, in which an abnormal state of the server remains for a period of time. Further, computational experiments are carried out to evaluate the resulting improved method. Based on the results of the experiments, it was concluded that the proposed additions make it possible to improve the forecasting accuracy of the model and reduce the number of false positives of the method, and the method can be used for early detection of gradual failures in the operation of server systems. Further prospects for the development of the research area include the classification of recognized anomalies and support for the administrator's decision-making when such events are detected.

References

1. Kothamasu, R., Huang, S., VerDuin W.: System health monitoring and prognostics – a review of current paradigms and practices. The International Journal of Advanced Manufacturing Technology, 28(9-10), 1012-1024 (2006).
2. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM computing surveys (CSUR), 41(3), 1-58 (2009).
3. Omar, S., Ngadi, A., Jebur, H.: Machine learning techniques for anomaly detection: an overview. International Journal of Computer Applications, 79(2) (2013).
4. Li, K.L. et al.: Improving one-class SVM for anomaly detection. In: Proceedings of the 2003 International Conference on Machine Learning and Cybernetics, 5, 3077-3081, IEEE, Xi'an (2003).
5. Ergen, T., Kozat S.: Unsupervised Anomaly Detection With LSTM Neural Networks. In: IEEE Transactions on Neural Networks and Learning Systems, 31, 8, 3127-3141, IEEE (2019).
6. Gupta, M. et al.: Outlier detection for temporal data: A survey. In: IEEE Transactions on Knowledge and data Engineering, 26, 9, 2250-2267, IEEE (2013).
7. Moayedi, H.Z., Masnadi-Shirazi, M.A.: Arima model for network traffic prediction and anomaly detection. International Symposium on Information Technology, 4, 1-6, IEEE, Kuala Lumpur (2008).
8. Kalekar, P.S.: Time series forecasting using holt-winters exponential smoothing. Kanwal Rekhi School of Information Technology, 4329008(13) (2004).
9. Dubrovin, M.G., Gluhih, I.N., Karyakin, I.Y.: Forecasting the server status using the triple exponential smoothing model. J. Phys.: Conf. Ser., 1661, 012031 (2020).
10. Jinka, P.: Anomaly detection for monitoring: A statistical approach to time series anomaly detection. O'Reilly Media, USA, 15-31 (2015).
11. Hyndman, R. J.: Measuring forecast accuracy. Business forecasting: Practical problems and solutions, 177-183 (2014).
12. Brutlag, J.D. Aberrant Behavior Detection in Time Series for Network Monitoring. LISA, 14(2000), 139-146 (2000).
13. Zhuang, Y. et al.: A weighted moving average-based approach for cleaning sensor data. In: 27th International Conference on Distributed Computing Systems (ICDCS'07), 38-38, IEEE, Toronto (2007).
14. Shcherbakov, M. et al: A survey of forecast error measures. World Applied Sciences Journal, 24, 171-176 (2013).
15. Elmrabit, N. et al.: Evaluation of machine learning algorithms for anomaly detection. In: 2020 International Conference on Cyber Security and Protection of Digital Services (Cyber Security), 1-8, IEEE, Dublin, Ireland (2020).
16. Szmit, M. et al.: Implementation of Brutlag's algorithm in Anomaly Detection 3.0. In: 2012 Federated Conference on Computer Science and Information Systems (FedCSIS), 685-691, IEEE, Wroclaw (2012).
17. Ekberg, J., Ylinen, J., Loula, P.: Network behaviour anomaly detection using Holt-Winters algorithm. International Conference for Internet Technology and Secured Transactions, 627-631, IEEE, Abu Dhabi (2011).
18. Debusschere, V.: Hourly server workload forecasting up to 168 hours ahead using seasonal ARIMA model. In: 2012 IEEE international conference on industrial technology, 1127-1131, IEEE, Athens (2012).