



System-Wide Anomaly Detection of Industrial Control Systems via Deep Learning and Correlation Analysis

Gordon Haylett, Zahra Jadidi, Kien Nguyen Thanh

► To cite this version:

Gordon Haylett, Zahra Jadidi, Kien Nguyen Thanh. System-Wide Anomaly Detection of Industrial Control Systems via Deep Learning and Correlation Analysis. 17th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI), Jun 2021, Hersonissos, Crete, Greece. pp.362-373, 10.1007/978-3-030-79150-6_29 . hal-03287698

HAL Id: hal-03287698

<https://inria.hal.science/hal-03287698>

Submitted on 15 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

System-wide anomaly detection of industrial control systems via deep learning and correlation analysis

G. Haylett¹, Z. Jadidi¹, K. Nguyen¹

¹ Queensland University of Technology (QUT), Cyber Security Cooperative Research Centre

Abstract. In the last few decades, as industrial control systems (ICSs) became more interconnected via modern networking techniques, there has been a growing need for new security and monitoring techniques to protect these systems. Advanced cyber-attacks on industrial systems take multiple steps to reach ICS end devices. However, current anomaly detection systems can only detect attacks on individual local devices, and they do not consider the impact or consequences of an individual attack on the rest of the ICS devices. In this paper, we aim to explore how deep learning recurrent neural networks and correlation analysis techniques can be used collaboratively for anomaly detection in an ICS network on the scale of the entire systems. For each detected attack, our presented system-wide anomaly detection method will predict the next step of the attack. We use iTrust SWaT dataset and Power System Attack datasets from MSU national Labs to explore how the addition of correlation analysis to recurrent networks can expand anomaly detection methods to the system-wide scale.

Keywords: Anomaly Detection, Correlation Analysis, Deep Learning, Industrial Control System.

1 Introduction

Traditional industrial control systems (ICSs) were not designed for security as they were isolated networks running proprietary control protocols. However, due to the connection of current ICS networks to the Internet, the cybersecurity of ICSs becomes a growing concern. A cyber-attack in an ICS network can be caused by an attack spreading maliciously from information technology (IT) networks to operational technology (OT) networks [1].

Many anomaly detection methods have been proposed to detect anomalies against ICS devices. Existing work on anomaly detection in ICSs can be divided into two broad categories, network traffic-based and physical process-based [2]. Network traffic-based anomaly detection systems are based on the analysis of communication patterns between different devices in an ICS network [3], while physical process-based anomaly detection methods are built on the analysis of ICS device logs that record the state of physical devices such as sensors and actuators [4]. These anomaly detection methods only focus on monitoring a single source of data without considering the consequence of attacks on other devices. Analysis of correlation among various data sources is required to provide comprehensive and system-wide anomaly detection [5]. This enables

security experts to predict and prevent future steps of an attack. Correlation analysis was employed in some papers [6] to improve the security of enterprise networks. However, the novelty of our paper is presenting a combination of anomaly detection and correlation analysis to provide a system-wide anomaly detection in ICS networks. Recurrent networks will be coupled with correlation analysis to detect when and where an anomaly occurs and subsequently search correlated devices for signs of further attack or influence. The results will be evaluated using real-world ICS datasets to show the efficiency of the presented method.

The rest of this paper is organised as follows. Section 2 explains background and related works. Section 3 describes the system-wide anomaly detection method proposed in this paper. Evaluation datasets are explained in Section 4. Section 5 discusses the evaluation results of our method in ICS datasets. Section 6 analyses the results. Section 7 concludes the paper.

2. Background and related works

Historically ICSs were immune to cyber-attacks because they were “Air-Gapped” [7]. This is a term used to describe a physical disconnect between the ICS and the organization’s cyber network. Meaning there is no connection an attacker could use to get from the corporate network onto the ICS network. This changed near the end of the 1990s with the development of new IT systems and technologies that improved the ICS workflow but also made it less secure to outside threats [7]. This increasing trend of legacy industrial devices being connected to normal IT networks and the internet means these ICS networks are now vulnerable to new types of attacks. Zero-day attacks are almost inevitable in these ICS networks and are a serious concern. This is where anomaly detection and machine learning are used to detect anomalous events [8]. Different machine learning methods have been used for anomaly detection. However, research shows that deep learning methods outperform traditional methods as the scale of data increases, shown in Figure 1 [9, 19].




Fig. 1. Performance comparison of deep learning Vs traditional anomaly detection methods [9]

There are many papers that performed deep learning-based anomaly detection for ICSs [11, 19], and their results showed the high accuracy of deep learning-based

analysis. Motivated by the above discussions, a deep learning-based anomaly detection method is deployed in our paper to analyse ICS data.

Advanced ICS attacks have multiple steps to gain access to their target device. Therefore, detecting a single anomaly cannot identify the future steps of an advanced attack. Using correlation analysis, we can assess correlated features and identify if an anomaly is isolated or potentially part of a larger complex multi-part attack [5, 6, 20]. In this paper, we use correlation analysis to expand the anomaly detection system-wide.

When working with high dimensional data with low-value density, it is crucial to select correct features for analysis as a large amount of sensor data streams can lead to serious performance impacts [5, 6, 20]. Paper [10] used sensor data correlation changes to improve the performance of IoT equipment anomaly detection by correlating duplicate deployed sensors and clustering them. The IoT devices used in the paper are sensors measuring things much like those in ICS networks. The paper states “dynamic data correlations among industrial equipment sensors prevalently exist” is a key component of correlation analysis as if there is a change in a device then it is likely there will be impacts to correlated devices. Figure 2 that is an example from paper [10], shows a clear correlation when the coal flow dips power dips after a small delay.




Fig. 2. Correlated power and coal flow of Coal mil. [10]

The methods used in the paper [10] would be an effective way to perform correlation analysis, however; on low dimensionality systems with few weakly correlated nodes or features, such as a limited number of sensors that only monitor isolated features, this method has the potential to be ineffective. If the features are not correlated strongly enough to impact each other when there is a change, the analysis method would need to be highly tuned or sensitive to changes. This can lead to false positives and misleading correlations. The paper on correlation-change anomaly detection [10] provides important information and methods for effectively correlating data.

Inspired by the paper [10], we present our correlation-based solution to explore a method of system-wide anomaly detection in ICS networks. Using anomaly detection on a selection of key features, we will then use correlation analysis to explore other potentially anomalous features which may be related to a detected anomaly.

3. System-wide anomaly detection

In this section, we will be discussing the methods and processes used to perform correlation analysis, anomaly detection and the experimental process.

3.1 Process

The process involves identifying correlated features from a map of all features as nodes. Firstly, features that have the greatest impact on other features and features which are likely to change in clusters are identified using correlation analysis. This is done so that our single feature anomaly detection method can become system-wide. This means that it can comprehensively track anomalies throughout an entire system. This is advantageous as many detection systems focus on monitoring individual devices and features but here, we incorporate correlation analysis on top of that to create a broader system-wide detection system. Once several nodes are correlated then if an anomaly is detected in one the correlated nodes are checked for anomalies to see if the anomaly is isolated, spreading or a part of a larger chain of anomalies.

When dealing with multidimensional datasets, it is important to filter out non-correlated features. This can lead to difficulties with automated model-building methods struggling to select the important features from among thousands of candidates. As it is better to use fewer correlated features to train a model [12], the network map of correlated nodes can be used to identify nodes that are most correlated to other nodes, and it provides a good starting point for identifying high priority nodes for monitoring. Figure 3 provides an example of a correlation map.




Fig. 3. Example of two correlated features

3.2 Methodology

Different phases of this paper were as follows (Figure 4):

1. First feature reduction is performed to remove duplicate sensors or isolate key features.
2. Correlation analysis is performed on the system or dataset in question. This helps to identify correlated nodes and which features share the strongest relations to other features.
3. We identify the key features/nodes for anomaly detection. Then, we monitor the nodes for anomalies using Long Short-Term Memory (LSTM) models to detect changes in the expected behaviour of the nodes.

4. If an anomaly is detected in a node, then we perform correlation analysis and examine each of its correlated nodes in turn. If an anomaly is identified or detected in a correlated node, its correlated nodes are checked as well. This process continues until no further anomalies are found or until every unique node is checked.
5. Using the detected anomalies and their timestamps, it is then possible to show when the detected anomalies occurred in the system.

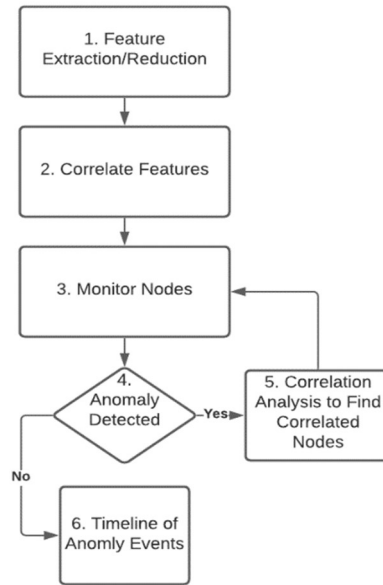


Fig. 4. Methodology flow

Our model is a combination of deep learning for anomaly detection and correlation analysis to expand it to be system-wide. This combination allows for the anomaly detection segment to focus on detection and the correlation analysis to provide the next target. This process allows us to follow the path an anomaly or attacker may have taken through a system and which devices may have been impacted and in what order. This can also be used to help identify attack insertion points and which nodes need hardening or further examination.

Anomaly Detection. For anomaly detection, an LSTM recurrent neural network was used as it is well suited for processing and making predictions based on time series data [9]. After correlation analysis is performed, nodes (aka features) are monitored by LSTM models of the features. If an anomaly is detected, then the correlated and adjacent nodes are also checked to see if the anomaly is isolated or if the anomaly is potentially part of greater concern. The LSTM used here is relatively simple and consists of an input and output layer and 4 LSTM layers using ‘relu’ activation method separated by an attention layer.

Correlation Analysis. When dealing with multidimensional datasets, it is important to filter out non-correlated features. This is done because two devices might monitor

the same sensor or device and would end up appearing highly correlated if there were a change in whatever device they were monitoring. This would create the impression these two sensors are linked or in some way impact each other when they do not. In this case, it may be better to use fewer highly correlated features to train a model [13]. For this paper, we use the Pearson correlation coefficient [18] as it indicates whether a statistically significant relationship exists between two continuous variables and whether a change in one variable may be associated with a proportional change in another variable.

When correlation analysis is performed, we can create a correlation matrix or correlation map to display the data in a more human-readable way as shown in Figure 5. The map on the right of Figure 5 shows the more highly correlated features from the matrix on the left in a more concise easy to read format. This map is beneficial when dealing with datasets which require some significant feature reduction or when it is unclear which parts of the system are the most interconnected. In highly connected systems, this may also be beneficial to reduce noise and prevent correlations being shown between every node and confusing the model with extraneous data. The size of the node in Figure 5 is indicative of how many edges (correlations) it has and the darker the edge the more highly correlated the nodes are. This map can also give a visual representation of clustering if there is any.

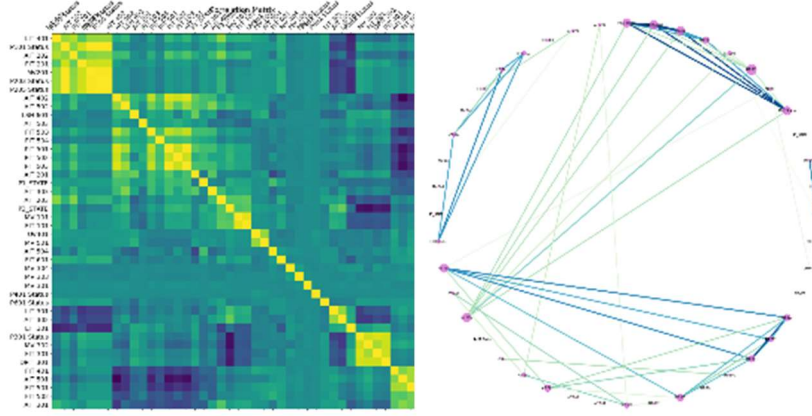


Fig. 5. SWaT Matrix and Map side by side comparison

4. Datasets

Our proposed model was evaluated using two ICS datasets, the Tommy Morris Power Systems dataset and the iTrust Secure Water Treatment (SWaT) dataset. These datasets were divided into training (70%), validation (10%) and testing sets (20%) for the evaluation of our model.

Power Systems. The ICS Cyber Attack Power System Datasets was developed in [15]. It contains 37 scenarios divided into 8 cases of Natural Events, 1 case of No Events and 28 Attack Events. The power system testbed used in generating the test scenarios

contains power generators, breakers, and Intelligent Electronic Devices (IEDs) that can switch the breakers on or off.

iTrust SWaT. The Secure Water Treatment (SWaT) dataset is a water treatment testbed for cybersecurity research [16]. The dataset consists of a modern six-stage process. There are 6 attacks performed on the dataset: Attack on FIT401 (Spoof value from 0.8 to 0.5), Attack on LIT301 (Spoof value from 835 to 1024), Attack on P601 (Switch from OFF to ON), Multi-point Attack (Switch from CLOSE to OPEN (MV201) and OFF to ON (P101)), Attack on MV501 (Switch from OPEN to CLOSE), and Attack on P301 (Switch from ON to OFF). FIT 401 sensor is a Flow Transmitter which controls the UV dechlorinator. LIT301 sensor is a Level Transmitter. P601 actuator pumps water from RO permeate tank to raw water tank. MV201 actuator is a motorized valve. P101 actuator pumps water from the raw water tank to the second stage. P301 actuator is a UF feed Pump.

5. Results

The experimental results are separated into two sections. The first section is for the Tommy Morris Power System Dataset, which is a simple case where the anomaly is fast and affects multiple devices in the same way. The second section is for the SWaT dataset.

5.1 Power Systems

The data in this section is from the line maintenance scenario where one or more relays are disabled on a specific line to do maintenance for that line [15]. Figure 6 shows the correlation matrix for the line maintenance scenario in this system. The correlation matrix uses Pearson correlation and shows which features have a correlation between each other, with highly correlated features being shown in yellow-green. This matrix has been sorted to show some clustering.

This matrix can be converted to a correlation map to show the correlations in a more readable format as shown in Figure 7. However, with such a large number of features so many correlations would make the map almost solid and impossible to read so the map is split into the strongest positive (>0.5 , left) and negative (<-0.5 , right) correlations. The first node monitored was R1-PA4:IH. Figure 7 shows the nodes readings in blue and the time periods where attacks were performed in red.

In Figure 8, the anomaly graph shows that several anomalies were detected in node R1-PA4:IH. These were determined to be anomalies by training a model of the data and taking the Mean Absolute Error (MAE) and lowering it by a small percentage to get a threshold for anomalies. As anomalies were detected in the initial node, as shown in Figure 8, all correlated nodes were also checked for anomalies. This process continues until no further anomalies were detected. This resulted in checking the following nodes:

R1-PA1:VH	R1-PA7:VH	R1-PA10:IH	R2-PA1:VH
R2-PA7:VH	R3-PA1:VH	R3-PA7:VH	R4-PA1:VH
R4-PA4:IH	R4-PA7:VH	R4-PA10:IH	R1-PA4:IH

All of which also had detected anomalies. This shows that if an anomaly is detected in any one node whether that be from a directed attack or other reasons the algorithm

will spread out from that node and search all adjacent and correlated nodes to find other anomalous behaviour. The accuracy of our LSTM in anomaly detection in the Power System dataset was 95%, and the precision was 95%.




Fig. 6. Power systems line maintenance correlation analysis matrix




Fig. 7. Power systems line maintenance positive and negative correlation network map




Fig. 8. R1-PA4:IH detected anomalies

5.2 iTrust SWaT

In the SWAT dataset, there are 6 different attacks performed. The first attack is on FIT401 trying to spoof value from 0.8 to 0.5 with the intent to stop de-chlorination by switching off UV401. Correlation analysis is performed to identify which nodes should be monitored most closely. The matrix in Figure 9 shows the correlations.

Using Figure 9 correlation Matrix, we construct the two maps in Figure 10 based on all the strongest positive (>0.5) and negative (<-0.5) correlations. Where a positive correlation represents a ‘perfect’ increasing relationship and negative correlations represent a ‘perfect’ decreasing relationship.

Using these maps, we can identify key nodes of interest with many strong positive correlations such as AIT402 and LIT401 as well as many strong negative correlations such as FIT401, AIT501 and LIT101. Any change in these nodes will likely have an impact on the network at large thus they are ideal for monitoring. The first node monitored for this dataset is FIT 401. This node is the first one attacked in the dataset. The anomaly graph in Figure 11 shows the anomalies detected in the node. The attacks are also highlighted in red. An anomaly that is detected within one of the red areas can be considered to be a true positive detection representing an anomaly that was detected as a part of an attack. An anomaly found outside of one of the red zones can be considered a false positive. The accuracy of our LSTM in detecting anomalies in the SWAT dataset was 97%, and the precision was 96%.




Fig. 9. SWaT correlation matrix

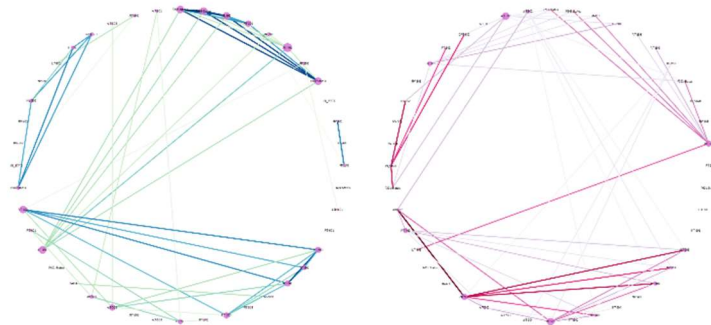


Fig. 10. SWaT Positive and Negative correlations




Fig. 11. FIT 401 detected anomalies

6. Analysis

The Power Systems dataset contains many different types of events which makes it challenging to perform anomaly detection. For R1-PA4:IH, there were 10 detected anomalies. An LSTM is used to predict the expected behaviour of each node (Figure 13), and a threshold value (Figure 12) helps to identify the anomaly area. The anomaly threshold is 10% in this paper to ensure as many anomalies are detected as possible. The area above this threshold shows anomalies, Figure 12. The LSTM models effectively predicted values closely aligned with the testing data (Figure 13).

The similar threshold value was used for the SWaT dataset. The testing and predicted data for FIT401 sensor in the SWaT dataset is shown in Figure 14. The accuracies of our LSTM in detecting anomalies in Power System and SWaT datasets were 95% and 97% respectively.

Existing anomaly detection methods can detect anomalies happening in each ICS device. However, advanced ICS attacks are mostly multi-step, and existing methods are not able to detect the correlation between different steps of an anomaly. In both Power System and SWaT datasets, our system-wide anomaly detection method provided high accuracy in anomaly detection, and it could detect the events correlated with each anomaly.




Fig. 12. Anomaly detection threshold (Power System Dataset)

Fig. 13. LSTM prediction of normal behaviour in R1-PA4:IH (Predicted values (red) and the actual values (green))




Fig. 14. FIT401 testing and predicted data (predicted (red) and the actual values (green))

7. Conclusion

Using deep learning-based anomaly detection in conjunction with correlation analysis, we could effectively expand single feature anomaly detection to be system-wide anomaly detection. This method of system-wide anomaly detection had the potential to accurately detect anomalies and then immediately check adjacent and correlated nodes for other anomalies. In this paper, our focus was on identifying other correlated nodes to the node with anomalous behaviour, and we wanted to identify the future steps of the ICS attacks. Our presented solution was evaluated using two real-world ICS datasets and it showed high accuracy in both datasets. In our future work, we will improve our method to include the time lag of anomalies during correlation analysis.

Acknowledgements

The authors acknowledge the support of the Commonwealth of Australia and Cybersecurity Research Centre Limited.

References

- [1] L. Maglaras, K. Kim, H. Janicke, M. Ferrag, S. Rallis, P. Fragkou, A. Maglaras, T. Cruz, "Cyber security of critical infrastructures", *ICT Express*, pp. 42-45, 2018.
- [2] Y. Hu, A. Yang, H. Li, Y. Sun, & L. Sun, A survey of intrusion detection on industrial control systems. *International Journal of Distributed Sensor Networks*. 1-14, 2018.
- [3] Dong, R., Wu, D., Zhang, Q., & Zhang, T., Traffic Characteristic Map-based Intrusion Detection Model for Industrial Internet. *International Journal of Network Security*. 359-370, 2018.
- [4] Hussain, M., Foo, E., & Suriadi, S., An Improved Industrial Control System Device Logs Processing Method for Process-Based Anomaly Detection. In *2019 International Conference on Frontiers of Information Technology (FIT)* (pp. 150-1505). IEEE.

- [5] Z. Tian, Y. Cui, L. An, S. Su, X. Yin, L. Yin, & X. Cui. A real-time correlation of host-level events in cyber range service for smart campus. *IEEE Access*, 6, 35355-35364 (2018).
- [6] F. Gottwalt, E. Chang & T. Dillon. CorrCorr: A feature selection method for multivariate correlation network anomaly detection techniques. *Computers & Security*, 83, 234-245, 2019.
- [7] T. Sommestad, G. N. Ericsson and J. Nordlander, "SCADA system cyber security - A comparison of standards," *IEEE PES General Meeting*, Providence, RI, pp. 1-8, 2010.
- [8] P. Parrend, J. Navarro, F. Guigou, A. Deruyver, P. Collet, "Foundations and applications of artificial Intelligence for zero-day and multi-step attack detection," *EURASIP Journal on Information Security*, 2018.
- [9] R. Chalapathy and S. Chawla, "Deep Learning for Anomaly Detection: A Survey," *arXiv:1901.03407*, 2019.
- [10] S. Su, Y. Sun, X. Gao, J. Qiu, Z. Tian, "A Correlation-Change Based Feature Selection Method for IoT Equipment Anomaly Detection," *Applied Sciences*, no. 437, p. 9, 2019.
- [11] D. Bhamare, M. Zolanvari, A. Erbad, R. Jain, K. Khan, N. Meskin, "Cybersecurity for industrial control systems: A survey," *Computers & Security*, p. 101677, 2020.
- [12] T. Tološi, T. Legauer "Classification with correlated features: unreliability of feature ranking and solutions," *Bioinformatics*, vol. 27, no. 14, pp. 1986-1994, 2011.
- [13] B. Ayinde, T. Inanc, J. Zurada, "On Correlation of Features Extracted by Deep Neural," *International Joint Conference on Neural Networks (IJCNN)*, 2019.
- [15] U. A. S. P. Tommy Morris, "Industrial Control System (ICS) Cyber Attack Datasets," Available: <https://sites.google.com/a/uah.edu/tommy-morris-uah/ics-data-sets>.
- [16] iTrust, "Secure Water treatment - iTrust," 2020. Available: <https://itrust.sutd.edu.sg/testbeds/secure-water-treatment-swat/>.
- [18] L. Wang, & R. Jones. Big data analytics in cyber security: Network traffic and attacks. *Journal of Computer Information Systems*, 1-8, 2020.
- [19] Z. Jadidi, A. Dorri, R. Jurdak and C. Fidge "Securing Manufacturing Using Blockchain," in 2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), Guangzhou, China, 2020 pp. 1920-1925.
- [20] T. R. B. Kushal & M. S. Illindala (2020). Correlation-based feature selection for resilience analysis of MVDC shipboard power system. *International Journal of Electrical Power & Energy Systems*, 117, 105742.