A Review of SUSAN: A Deep Learning based anomaly detection framework for sustainable industry

Ángel Luis Perales Gómez*[1], Lorenzo Fernández Maimó[1], Alberto Huertas Celdrán[2] and Félix J. García Clemente[1]
[1] Faculty of Computer Science, University of Murcia, 30100 Murcia, Spain
angelluis.perales@um.es; lfmaimo@um.es; fgarcia@um.es
[2]Communication Systems Group CSG, Department of Informatics IfI, University of Zurich UZH, CH-8050, Switzerland
huertas@ifi.uzh.ch

***Resumen*—Nowadays, sustainability is pivotal in green technologies, crucial for industries striving to cut carbon emissions and optimize energy use. Alongside this concern, cyberattacks impacting sustainability in industries are on the rise. These attacks target industrial systems managing processes, often needing specialized knowledge and evading traditional cybersecurity measures. To tackle this, SUSAN, a Deep Learning-based framework, is introduced, aimed at detecting cyberattacks on industrial sustainability. SUSAN's modular design allows combining multiple detectors for precise detection. Demonstrated in a water treatment plant using the SWaT testbed, SUSAN achieved a high recall rate (0.910) and acceptable precision (0.633), resulting in an F1-score of 0.747. It successfully detected all individual cyberattacks, surpassing related work with a worst recall rate of 0.64.**

***Index Terms*—anomaly detection, deep learning, industrial control systems, machine learning, sustainability**

Tipo de contribución: Investigación ya publicada

I. Introduction

Green technologies are increasingly vital in reshaping socio-economic growth worldwide, with sustainability being paramount for present and future generations' prosperity and environmental well-being. Industries, including energy, agriculture, and water treatment, are adapting production processes to align with sustainable practices, such as reducing carbon emissions and optimizing energy use. However, as technology evolves, Industry 4.0 emerges, integrating next-gen networks like 5G, big data, and Internet of Things (IoT), offering vast opportunities for sustainable manufacturing. Yet, this integration also exposes industrial systems to new cyberattacks, impacting process sustainability. Traditional cybersecurity methods are inadequate in detecting these threats, leading to the adoption of Anomaly Detection (AD) systems based on machine learning (ML) and deep learning (DL). However, existing AD systems in industrial control systems (ICS) lack focus on sustainability-related cyberattacks. Enhancing anomaly detectors to prioritize sustainability aspects becomes crucial as sustainability gains prominence in industrial settings, ensuring Industry 4.0 projects align with sustainable practices and improving factory sustainability overall.

This paper reviews [1] which introduces three main contributions: 1) a framework, called SUSAN, for building detectors of anomalies produced by cyberattacks that affect the sustainability of industrial processes; 2) A set of sustainability-based features that gather the resource consumption routines for modeling the behavior of industrial sensors and actuators related to sustainability; and 3) Validation of the proposal using the well-known Secure Water Treatment testbed, SWaT.

Tabla I
Recall comparison between different solutions

#	DNN	RNN	OCSVM	1D-CNN	DIF	Ours
6	0.95	0.72	0.72	0.90	**1**	0.92
11	0.99	0.98	**1**	**1**	**1**	0.96
19	**0.97**	0.12	0.13	0	0.45	0.64
20	0	0.85	0.85	**1**	0.45	**1**
22	0.98	0.99	**1**	**1**	**1**	0.97
24	**0.92**	0	0	0.17	0.34	0.78
28	0.03	0.94	0.94	**1**	**1**	0.93
38	0.77	0.92	0.93	0.86	**1**	0.69
40	0.78	0.93	0.93	**1**	**1**	0.74

II. SUSAN: Sustainability-aware Anomaly Detection Framework for ICS

This section details the design characteristics of the proposed DL-enabled framework, called SUSAN, for building detection systems of anomalies caused by cyberattacks that affect the sustainability of industrial processes. SUSAN is composed of four modules:

Data Preprocesing. This module includes the following components: data cleaning, feature encoding, dataset generation, and feature normalization. The data cleaning component explores the data and removes spurious or corrupted values. The Features Encoding component is in charge of transforming features to make them compatible with DL models. At this point, the component encodes categorical values of data by applying One-Hot Encoding (OHE). The Datasets Generation component is in charge of creating the three datasets that will be used to train and validate the model, i.e., training, validation, and test datasets. Finally, the Feature Normalization component studies the distributions of the features to perform the normalization process.

Sustainability-based Features Generation. This module includes the following components: feature filtering and feature extraction. Feature Filtering studies the features contained in the databases and remove those that do not provide enough information. In particular, this component conducts a correlation study to identify potential data leakage between features

and labels, a variance study to eliminate features with constant values, and ensures the preservation of distribution between datasets through the Kolmogorov-Smirnov test. Finally, the Feature Extraction component extracts higher-order features from the original ones, trying to add more information to discriminate between normal and cyberattack behavior.

Model Generation. The Model Generation comprises the following components: anomaly detection model selection, model fine-tune, and model trainer engine with validation. The anomaly detection model selection component aims to choose the most suitable model to detect sustainability anomalies. SUSAN proposes a range of DL regressor models applied to a feature window, as they are better suited to model time-series data with complex behavior compared to simpler ML techniques. The model fine-tune component selects a set of hyper-parameters and their respective ranges of values to be fine-tuned. The fine-tune process trains the model using the training dataset with a range of hyper-parameters values to determine the optimal values. Finally, the model trainer engine and validation component trains the selected model and validates it.

Sustainability Anomaly Detector. This module is composed of the following components: threshold monitor and anomaly classifier. The first component computes the error between the predicted values of the model and the new unseen values. If the error is higher than a given threshold, an anomaly is detected and passed on to the next component. The Anomaly Classifier component classifies the anomaly depending on the sensor/actuator where it impacts.

III. Sustainability-based Features for Industrial Environments

SUSAN utilizes basic features extracted from sensors and actuators to assess sustainability, generating higher-order features from these basic inputs. Specific features impacting sustainability are selected based on the industrial system's scenario, such as resource consumption or product dosing in water treatment plants. Time encoding is essential due to industrial systems' repetitive nature, achieved by computing two features representing time as sine and cosine values within a time window. SUSAN employs three techniques to extract statistical and time-series features within a defined window. The first technique consists in extracting features from the resource consumption habits. These new features allow us to discriminate between normal and abnormal behavior in resource consumption. The second technique is the autocorrelation which allows SUSAN to extract time patterns. Finally, the last technique is the Discrete Fourier Transform (DFT) which shares the same objective as autocorrelation.

IV. Experimental results

This section describes how SUSAN modules work in the SWaT scenario.

Data preprocessing. The first step was to select only those industrial processes containing sensors or actuators related to resource consumption. Furthermore, we removed the first 100 000 samples since they belong to the warm-up process. Then, we only selected abnormal samples from those cyberattacks that impacted the sustainability, i.e., cyberattacks 6, 11, 19, 20, 22, 24, 28, 38, and 40. Additionally, after completing an attack, the samples were labeled as normal. However, as the system status remained abnormal, we consequently removed 600 samples after each cyberattack. Finally, we encoded categorical features and normalized all required features.

Sustainability Features Generation. In this step, the feature filtering component did not remove any feature. Regarding the feature extraction component, the DFT, autocorrelation and statistical functions focused on consumption habits were applied over a windows of 120 samples. After the feature engineering process, the total number of features was 2 234.

Model Generation. The model selected was a Long Short-Term Memory (LSTM) that is suitable to be used with time-series data. Next, the Model Fine-tune component was in charge of selecting the optimal hyper-parameters values. The optimal values were 120 for the window length, 3 LSTM layers of 512, 256, and 128 neurons, 1 dense layer of 100 neurons and ReLU as activation function. Finally, the model was trained to find the proper threshold.

Validation. We used Precision, Recall, and F1-score, which are the common metrics for measuring detection performance in AD. Tabla I summarizes the recall rates achieved by the different approaches focused on the SWaT dataset. In general, our solution achieved the most balanced results for all cyberattacks that pose sustainability issues. Regarding global performance, SUSAN achieved a recall of 0.910, a precision of 0.633, and a F1-score of 0.747.

V. Conclusions

The paper introduces SUSAN, a flexible framework for building sustainability anomaly detectors in industrial processes, composed of three modules. It proposes techniques to extract features related to resource-consumption habits, aiding in distinguishing normal and abnormal behaviors. Experiments in a water treatment testbed demonstrate SUSAN's effectiveness. Notably, SUSAN achieved a recall rate of 1 for the best-performing cyberattack and 0.64 for the worst-performing one, with an overall recall of 0.910, precision of 0.633, and F1-score of 0.747.

Acknowledgements

This work has been funded under Grant TED2021-129300B-I00, by MCIN/AEI/10.13039/501100011033, NextGenerationEU/PRTR, UE, Grant PID2021-122466OB-I00, by MCIN/AEI/10.13039/501100011033/FEDER, UE, by the strategic project CDL-TALENTUM/DEFENDER from the Spanish National Institute of Cybersecurity (INCIBE), by the Recovery, Transformation and Resilience Plan, Next Generation EU, by the Swiss Federal Office for Defense Procurement (armasuisse) with the CyberForce (CYD-C-2020003), and by the University of Zurich (UZH).

Referencias

[1] Perales Gómez, Á. L., Fernández Maimó, L., Huertas Celdrán, A., & García Clemente, F. J. (2023). SUSAN: A Deep Learning based anomaly detection framework for sustainable industry. *Sustainable Computing: Informatics and Systems*, 37, 100842.