Physical Anomaly Detection in EV Charging Stations: Physics-based vs ResNet AE

Harindra S. Mavikumbure [1], Victor Cobilean [1], Chathurika S. Wickramasinghe [1], Tyler Phillips [2], *Member, IEEE*, Benny J. Varghese [2], Barney Carlson [2], Craig Rieger [2], *Senior Member, IEEE*, Timothy Pennington [2], and Milos Manic [1], *Fellow, IEEE*

[1]Department of Computer Science, Virginia Commonwealth University, Richmond, USA
(mavikumbureh, cobileanv, brahmanacsw)@vcu.edu, misko@ieee.org

[2] Energy and Environmental Science & Technology, Idaho National Laboratory, Idaho Falls, ID, USA
(tyler.phillips, benny.varghese, richard.carlson, craig.rieger, timothy.pennington)@inl.gov

***Abstract*—The number of electric car users has grown in recent years, increasing the demand for reliable electric vehicle charging stations (EVCS). The safety of EVCSs is very important as compromised charging stations can disrupt the grid, injure the end-users, and damage the vehicle. In this paper, we will focus on the physical security of EVCS, because physical attacks tend to be more harmful to the end user. Two anomaly detection approaches were presented for detecting physical anomalies: physics-based anomaly detection and deep learning-based anomaly detection (ResNet Autoencoder). The presented approaches were trained and tested using data collected from the EV Charging Station System testbed of the Idaho National Laboratory. Anomaly detection performance was evaluated on three different attack scenarios, targeting various parts of the system including power transfer subsystems and the cooling subsystem of the charger. The presented approaches were compared against two widely used unsupervised anomaly detection algorithms: OCSVM and LOF. Moreover, we evaluated the advantages and limitations of the physics-based vs ResNet Autoencoder approaches for each of the three attack scenarios. The ResNet Autoencoder approach showed the highest performance in terms of accuracy, F1, recall, and precision. Furthermore, this approach demonstrated a number of advantages including automated non-linear feature extraction and unsupervised learning.**

***Index Terms*—Deep Neural Networks, Autoencoders, Physics-based models, Unsupervised Learning, Anomaly Detection, Electric Vehicle Charging Systems**

I. Introduction

Electrified transportation is seen as a key driver for increasing energy efficiency and sustainable energy infrastructure. Electric vehicle (EV) offers from automakers are growing, and the infrastructure for car charging is quickly following. The United States experienced a 9.2% quarterly growth rate in public chargers in 2020 Q4 [1] and recently passed the 100,000 public charger mark in March 2021. It is expected that about 125 million EVs will be in operation on the road by 2030 [1]. The EV industry needs to facilitate adequate charging infrastructure for EV users to achieve this ambitious figure.

These charging ports, alternatively known as smart EV charging stations (EVCS), serve as the EVs' access points to the energy infrastructure (i.e., smart grid). Conventionally, the energy infrastructure and the smart EVs exchange information and energy through these EVCS [2]. Hence, appropriate management of EVCS (in terms of privacy protection) is inevitable as it may cause havoc on power grid infrastructure if any of these EVCS is compromised (by the adversary) or even remains unmanaged. Attackers who obtain control of an EVCS gain backdoor access to the grid. Because of the possibility of transferring a significant amount of energy, these attacks may cause network instability and cascade breakdowns [3]. The attackers can damage the charging vehicle by overcharging the battery and manipulating the charging profile. Moreover, a charging station is a cyber-physical system (CPS) whose operation is governed by the interactions between physical and cyber components. As a result, these systems are exposed to tampering with measurements and cyber-attacks. The charging station's cyber layer is vulnerable to many attacks such as DoS attacks, false data injections, spoofing, repudiation, and MITM attacks that can directly affect the physical layer of the system. [4].

In this study, we will present a physical anomaly detection system trained in an unsupervised manner for electric vehicle charging stations. We will provide a comparison of the two paradigms of the system modeling for anomaly detection: the physics-based system identification model and the data-driven deep learning model (ResNet Autoencoder). We implemented and compared the performance of the models based on evaluation metrics (F1, Precision, Recall, Accuracy) of test scenarios. We also provided a detailed analysis of each model's advantages, disadvantages, and limitations in relation to each test scenario and the overall physical security of EVCS.

The rest of the paper is organized as follows: Section II provides the Background and Related Work; Section III describes the experimental testbed; Section IV presents the Physics-based and ResNet AE based Methodology; Section V discusses the experimental setup, results, and discussion, and finally, Section VI concludes the paper.

II. Background and Related Work

As the scale of EVs (electric vehicles) continues to expand, the proper operation of EVs is essential. The normal operation of EVs is closely related to the daily operation of charging

The accepted version of the paper appearing in the 32nd International Symposium on Industrial Electronics (ISIE) 2023

stations and charging ports. Whether the EVs charging station, as the connection point between EVs users and the power grid, can work properly or not is the key to the normal operation of EVs, and the normal operation of the charging station is related to the normal operation and daily profit of the charging station. So ensuring the proper operation of the triad of EVs, charging stations, and charging ports is crucial to support EVs development, promote efficient EVs use, and ensure the operational efficiency and safety of the grid [5]. Currently, researchers are working towards megawatt charging systems and the increasing complexities in the infrastructure will increase vulnerabilities as well [6].

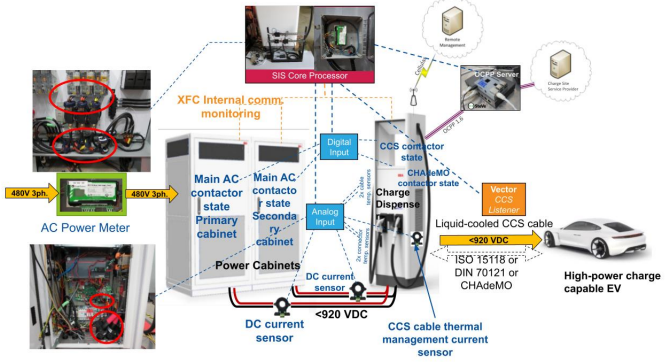

Fig. 1: EV Charging Station System testbed of the Idaho National Laboratory

Finding patterns in data that do not match expected behavior is known as anomaly detection. A decline in system performance, which might result in instability and failure, is strongly indicated by anomalies [7]. The use of data-driven machine learning algorithms is one of the promising approaches to detecting anomalies [8]. There are two main approaches for building data-driven anomaly detection systems, namely supervised and unsupervised. Supervised approaches require labeled data to train the algorithms. However, data labeling is time-consuming, and one cannot label all the anomalous behaviors of the system due to practical difficulties. Therefore, unsupervised techniques such as One-Class SVM, Local Outlier Factor, and Autoencoders have gained popularity over recent years.

Apart from machine learning models, physics-based system identification models have gained popularity over the past years. They analyze the physics-based dynamics of the system to monitor abnormal behaviors of the system. Two popular physics-based methods used for detection are safety limits and anomaly-based. Safety limits are a straightforward approach where constraints, or a bound, on values, are placed on measurements such as temperature, pressure, or flow rates. If the operation falls outside of the limits an alarm is raised. Anomaly-based detection is done with the use of a mathematical model which represents the dynamics of the system [9].

Recent literature shows exciting efforts that have been implemented to ensure the security of EVCSs. In [10], authors developed an isolation forest-based model with an anomaly score calculation method for detecting internal and external abnormalities in EVs and EVCSs. In [11], authors proposed a novel network traffic anomaly detection model based on Multi-Head Attentions (MHA) that takes into account the inherent correlations of traffic generated by ICSs. The MHA model is employed to substitute the traditional feature extraction and rule-making process with an acceptable computational cost for classifying traffic data. In [12], To identify DoS attacks in the EVCS, authors suggest new deep learning-based intrusion detection systems (IDS). To recognize and categorize DoS attacks in the EVCS, the deep neural network (DNN) and long-short-term memory (LSTM) algorithms are implemented. For both binary and multiclass classification, the proposed LSTM-based IDS outperformed a competing DNN-based IDS in terms of accuracy, precision, recall, and F1 score.

III. TESTBED DESCRIPTION

This section illustrates the INL's EV charging station system testbed system which is presented in Figure 1. INL's Electric Vehicle Infrastructure Lab (EVIL) facility consists of 350kW Extreme Fast Charging (XFC), 50kW Direct Current Fast Charging (DCFC), the Safety Instrumented System (SIS) core module integrated into each charger, several electric vehicle models (Nissan LEAF, BMW i3), and the CCS EV emulator. Numerous physical measurements are obtained from the XFC input, output, and internal components during operations. These measurements are acquired from an AC power meter, several DC hall-effect current sensors, and several thermistor temperature sensors. The AC power measures the real and apparent power for the entire XFC as well as the individual power cabinets. Hall-effect current sensors are used to measure the DC current output from the XFC power cabinets and the CCS liquid-cooled chiller electrical draw. Thermistor temperature sensors, installed with the CCS cable assembled by the manufacturer, are used to measure the operating temperature of the CCS cable and connector assembly.

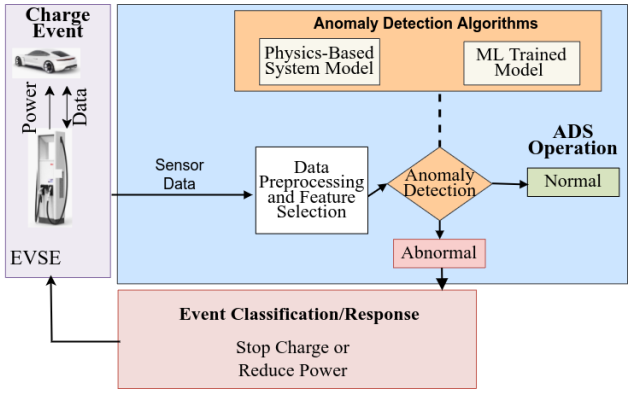

Fig. 2: Framework for Physical Anomaly Detection in EV Charging Station Systems

IV. PHYSICS-BASED AND RESNET AE-BASED ADS METHODOLOGY

This section describes the presented approach for data-driven and physics-based anomaly detection approach for physical health monitoring of the EVCS.

The presented approach analyzes physical data in order to identify internal anomalies through the use of anomaly detection algorithms. The overall architecture of the system is presented in Figure 2. The approach consists of a) Data pre-processing, b) Autoencoder based anomaly detection, and c) Physics-based anomaly detection. The following subsections describe the three main components mentioned above.

A. Data pre-processing

This section describes the data pre-processing techniques we used in this work. We used a window-based pre-processing approach [13] for eliminating the noise from the sensor data. The approach is averaging the sensor outputs within a defined-size window. The optimal window size (winSize) was chosen by experimenting with different sizes and the presented results in the paper are obtained using the best-found winSize=500 ms. The best performance was obtained by using overlapping windows because it increases the temporal resolution and the collected data points are analyzed multiple times with a temporal offset. Overlaps between two windows are kept at half of the window size.

Instead of using all the available sensor data (features) in the dataset, we performed feature selection in order to find the best feature combination which provides the most accurate AD results. Best feature selection is performed through a combined approach of Pearson correlation and grid search. Pearson correlation is used because it indicates the presence or absence of correlation between any two variables and determines the exact extent or degree to which they are correlated. Feature selection has advantages such as reduced training time as the dimensionality of the dataset decreases, removal of redundant features, and improved accuracy because of the removal of redundant data. Table II illustrates some of the important features extracted from the total collection of the sensors.

B. ResNet Autoencoder based anomaly detection

In this work, we used an Autoencoder (AE) Neural Network model to identify the abnormal behaviors of the system. We implemented a deep ResNet AE (RAE) architecture to detect any potential performance degradation and to improve feature learning capability [14]. Based on the traditional AE network, the ResNet-AE network replaces the linear structure in the AE network with the ResNet structure. There are multiple hidden levels in both the encoder and the decoder in AE. The two phases of the model's training are the encoding stage and the decoding stage. During the encoding step, the input data is transformed into an embedded representation, and during the decoding stage, the embedded representation is reconstructed to the original input record (reconstruction). The loss function (J_θ) of the AE model is computed using the difference between the input (x) and the reconstruction (x'). Thus reconstruction error of the AE is calculated as follows,

$$J_\theta = \frac{1}{T}\sum_{i=1}^{T} \|x_i - x_i'\|^2 \tag{1}$$

where x_i is the i th input sample, x_i' is the reconstruction for ith input sample, θ denotes the set of parameters of the AE (weights and biases).

The AE model is trained using only data from normal behaviors. Therefore, it only learns the possible normal behaviors of the system. When unseen records are presented to the trained AE, the amount of reconstruction error indicates how much the presented data differs from the learned normal behavior. A threshold value is defined to identify possible anomalies. The data records were detected as anomalies if the reconstruction error was higher than the defined threshold value. Thus, given data record x_i is detected as anomaly ($y = 1$) or normal ($y = 0$) as follows,

$$J_{\theta,i} = \|x_i - x_i'\|^2 \tag{2}$$

$$y = \begin{cases} 1 & : J_{\theta,i} \geq th \\ 0 & : J_{\theta,i} < th \end{cases} \tag{3}$$

where $J_{\theta,i}$ is the reconstruction error of ith data record, th denotes the threshold value, and y represents predicted label: anomaly or not. The threshold value is optimized based on the training baseline data, i.e., the threshold value should capture the baseline data boundary, capturing the normal behavior fluctuations.

The mean squared error was used as the loss function. A different number of hidden layer sizes were tested, and the paper presents the best results.

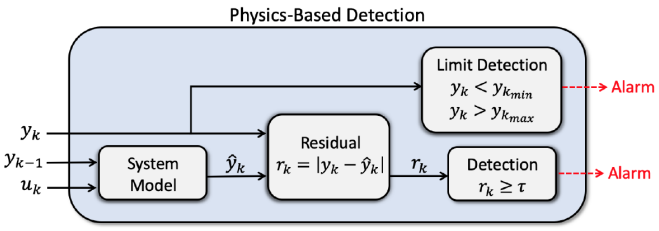

Fig. 3: Physics-based detection where both safety limits and anomaly-based detection are used for alarms. Image taken from [9].

C. Physics-based anomaly detection

The main idea behind a physics-based anomaly detection method is that physical systems must follow the imitable laws of nature. For example, the heat generation of a gas turbine generator follows thermodynamic properties, the flow rate of a hydropower generator follows fluid dynamic principles, and an electrical system follows electrodynamics.

Figure 3 illustrates the two approaches for physics-based anomaly detection. In this work, we utilized the anomaly-based detection method. Anomaly-based detection is done

TABLE I: ResNet-ADS anomaly detection comparison

Method	Accuracy	Precision	F1	Recall
LOF	87.50	100.00	80.00	66.67
OCSVM	93.86	86.36	92.68	100.00
Physics-based	87.61	88.00	93.61	100.00
ResNet AE	96.82	92.43	96.07	100.00

with the use of a mathematical model which represents the dynamics of the system. It is used to predict an expected measurement, $\hat{y}_k$, using the current control commands, u_k, and the previous sensor measurements, y_{k-1}.

The mathematical model of the system is derived through a data-driven process called system identification. In this work, the mathematical model is represented as a time-series ARX (Autoregressive-exogenous) model given as

$$y_k = \sum_{i=1}^{n_a} a_i y_{k-i} + \sum_{i=1}^{n_b} b_i u_{k-i} \tag{4}$$

Here, unknown constants a_i and b_i are solved using the Python package GEKKO [15] and the function *sysid()*. The number of previous outputs (n_a) and inputs (n_b) used are 1 and inputs of primaryDCA and chiller24A are used to predict the output of the cable temperature.

The anomaly detection test itself uses a time series of residual values, r_k. The residual is the difference between the real-time measured value and the predicted value from the mathematical model, given as

$$r_k = |y_k - \hat{y}_k| \tag{5}$$

The residuals can be used in either a stateless or stateful anomaly detection test. A stateless test raises an alarm every time a residual value reaches a set threshold, $r_k \geq \tau$. In this work, a stateful test was used. Here, the historical changes of the residual are kept as an additional statistic, denoted as S_k, and used to generate an alert if $S_k \geq \tau$. There are many ways to keep track of the residual for a stateful test, such as an exponential weighted moving average, using change detection statistics such as the non-parametric cumulative sum statistic, or tracking an average over a time window. The time window approach was selected for this work, using the previous 100 data points.

V. Experiments

This section presents the experimental results and analysis of the anomaly detection algorithms. We evaluate the performance of the ResNet Autoencoder and Physics-based model by comparing them with two benchmark algorithms: 1) One-Class Support Vector Machines (OCSVM), 2) Local Outlier Factor (LOF).

Data Collection: For experimental evaluation, data from normal and physical attack scenarios were collected. Physical sensor measurements collected from the testbed enable the direct detection of anomalous operating conditions as well as redundant calculations of various parameters to provide additional means to quickly determine anomalous operating conditions. The SIS system saves the attack timestamps in a log file, indicating when each attack starts and stops. These timestamps are used to label the dataset, indicating which data correspond to a normal state and which data correspond to physical anomalies. we run each scenario separately and keep the collected data in separate files.

TABLE II: Physical Feature List

Feature Name	Description
XFC Input Voltage	Voltage at the service panel feeding the XFC
XFC Input Current	Current at the service panel feeding the XFC
Primary Power Cabinet Voltage	Voltage at the input to the Primary Power Cabinet
Secondary Power Cabinet Voltage	Voltage at the input to the Secondary Power Cabinet
Primary Power Cabinet Power Factor	Calculated Power Factor of the Primary Power Cabinet
Secondary Power Cabinet Power Factor	Calculated Power Factor of the Secondary Power Cabinet
CCS Cable Temperature	Temperature Measurement within the CCS Liquid-Cooled Cable
CCS Connector Temperature	Temperature Measurement within the CCS Liquid-Cooled Connector
Primary Real Power Cabinet Current (Pri_Real_Power_i)	Current at the input to the Primary Power Cabinet
Secondary Real Power Cabinet Current (Sec_Real_Power_i)	Current at the input to the Secondary Power Cabinet
Primary Power Cabinet DC Output Current (Primary DCA)	DC Current Output from the Primary Power Cabinet
Secondary Power Cabinet DC Output Current (Secondary DCA)	DC Current Output from the Secondary Power Cabinet
CCS Cable Liquid Chiller Current (Chiller24A)	Current Draw of the CCS Chiller from the 24VDC Auxiliary System

In this work, we divided the initial dataset into two: 0.7 of the data for training and 0.3 for validation. The best models were evaluated using test data which included three attack scenarios. Before applying the algorithms the data were scaled to the 0-1 range.

We selected three scenarios for this paper to present the ADS output results. The selected scenarios are described below.

Scenario 1 - Power transfer system manipulation: This situation could be the consequence of an attack, as in this experiment, or it could be the result of power electronics failure. The power module controller is interrupted while charging, compromising the AC power quality at the grid connection. If the system controller is not sufficiently robust, this manipulation can harm equipment, start electrical fires, or disturb the stability of the power grid.

Scenario 2 - Liquid cooling system manipulation: This situation could be the consequence of an attack, as in this experiment, or it could be the result of a chiller pump failure. In the case of our scenario, the chiller was disabled and the temperature feedback from the thermal sensors was spoofed to appear normal. This scenario can be dangerous for the charging station equipment and present a burn hazard to the equipment user, causing injury to the user and damage to the charging equipment.

Scenario 3 - External manipulation of the charging profile: During vehicle charging, a malicious payload was injected to change the maximum power profile point. By manipulating the operating point, the attacker can harm the charger, and the vehicle, and endanger the user. It can cause the distribution to the grid by surging high amounts of power.

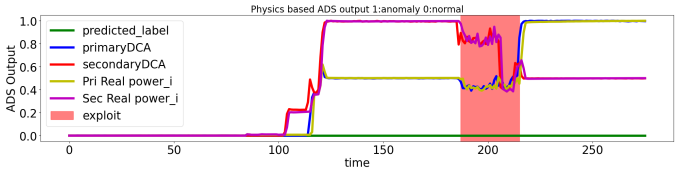

(a) Physics-based ADS output for scenario 1

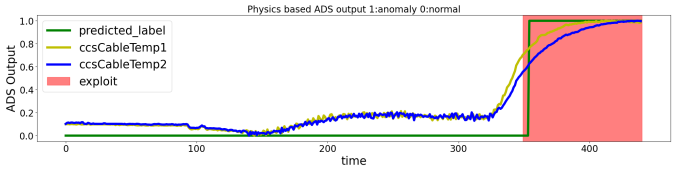

(b) Physics-based ADS output for scenario 2

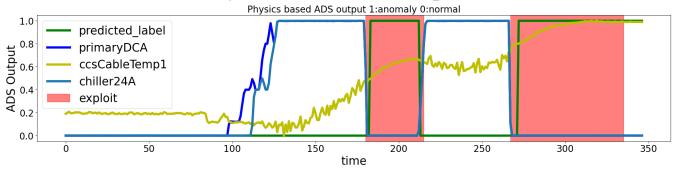

(c) Physics-based ADS output for scenario 3

Fig. 4: Physics-based anomaly detection results for the three scenarios.

We initialized our experiments with a physics-based system identification model. The following section describes the results of the physics-based model.

A. *Physics-based Anomaly Detection*

Results for the physics-based anomaly detection in each of the scenarios are shown in Figure 4. As shown in the figure the green line indicates the outcome of the ADS and the exploit duration is shown by the red area. Further, we have shown how different features change during the attack scenarios in order to explain the behavior of the system. All the features shown in the figure are scaled between 0-1. The accuracy, precision, F1, and recall are shown in Table I. Here it can be seen that physics-based anomaly detection performs similarly to LOF and OCSVM methods but not as well as the ResNet AE.

- **Scenario 1:** Figure 4a illustrates that the physics-based approach cannot detect the power module manipulation as the model input and output are still as expected. i.e. the heating of the cable is still as expected based on the measured primaryDCA (blue line) and chiller24A during the exploit period between time steps 187 - 215.
- **Scenario 2:** Figure 4b illustrates that the physics-based method detects the rapid heating of the cable during the chiller exploit from time steps 349 - 440. As shown in the Figure, CCS cable temperature feature values (blue line and yellow line) increased during the attack. Here, the attack manipulates the chiller operation and the model would not expect such rapid heating, and the alarm is triggered.
- **Scenario 3:** Figure 4c illustrates that the physics-based method detects false data attacks. Here, it can be seen that the increased heating of the cable remains similar during the attacks from time steps 187 - 215 and 275 - 335. In this figure, the heating of the cable is represented using the CCS cable temperature feature (yellow line). Based on the inputs to the physics-based model, this is unexpected; the residual window breaks the threshold and an alarm is raised.

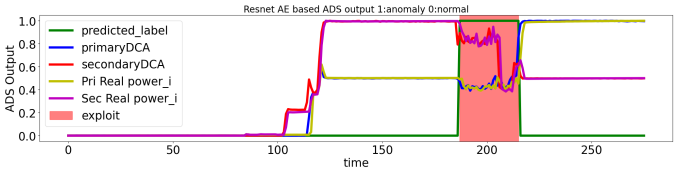

(a) ResNetAE-ADS output for scenario 1

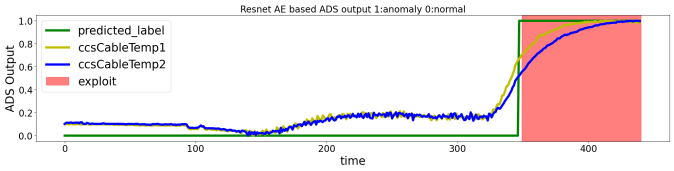

(b) ResNetAE-ADS output for scenario 2

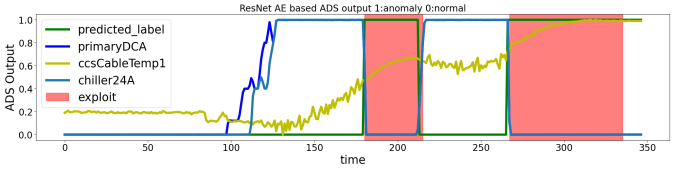

(c) ResNetAE-ADS output for scenario 3

Fig. 5: ResNet Autoencoder-based anomaly detection results for the three scenarios.

Results of the physics-based anomaly detection model confirm that it fails to detect scenarios where the physics match what is expected although there is an exploit. Due to the limitations of the physics-based system identification model, we experimented with our data with the ResNet Autoencoder model. The following section describes the results of the ResNet Autoencoder-based anomaly detection results.

B. *ResNet AE Anomaly Detection*

In this experiment, we evaluated the performance of the unsupervised ResNet AE to detect anomalies. Normal behavior corresponds to sections where no attack was being executed whereas an anomaly corresponds to any of the physical attacks executed during the experiment.

Results for the ResNet AE-based anomaly detection in each of the scenarios are shown in Figure 5. Table I shows the accuracy, precision, recall, and f1 scores of the ResNet AE method. The table shows that ResNet AE performs more accurately w.r.t. OCSVM, Physics-based, and LOF approaches. OCSVM and ResNet AE have the same recall while ResNet

AE has higher accuracy, precision, and F1 scores. The ResNet AE model shows higher accuracy, recall, F1, and precision than the physics-based model which confirms that the ResNet model performs better than the physics-based model in the current data setting. However, the results suggest room for improving the precision and F1 scores.

Figure 5a, 5b, and 5c show that the presented ResNet AE method is able to correctly identify abnormal scenarios.

- **Scenario 1:** Figure 5a illustrates the results of the ResNet AE for scenario 1. During the exploit period (time steps 187 - 215), the primaryDCA (blue line) and SecondaryDCA (red line) current decreased compared to the normal charging session. As shown in the Figure, our ADS detects the system's abnormal behavior when these drastic changes happen in the features mentioned above.
- **Scenario 2:** Figure 5b illustrates the results of the ResNet AE for scenario 2. During the exploit period (time steps 349 - 440), the ccs cable temperature values (yellow line and blue line) are observed to be increased since the attack manipulates the chiller operation. ResNet AE is able to detect the system's abnormal behavior when the temperatures start to show a drastic change and it correctly overlaps with the time that the attack launched.
- **Scenario 3:** Figure 5c illustrates the results of the ResNet AE for scenario 3. During the exploit period (time steps 187 - 215 and 275 - 335), the primaryDCA, Secondary-DCA, Primary Real power, and Secondary Real Power current features show a sudden drop due to the false data attack which set the "chargePointMaxProfile" to 100w while the correct value is 100kw. As shown in the Figure, our ResNet AE could detect the system's abnormal behavior when these drastic changes happen in the features mentioned above. As explained above, Figure 5c shows a sudden drop in the primaryDCA feature (blue line) during the exploit.

Pros/Cons of Physics-based model: As mentioned earlier, the physics-based system identification model fails to identify some anomalous scenarios accurately. Physics-based models are based on assumptions made by humans about the physical dynamics of the system. However, system complexity scales the number of physical variables exponentially. Therefore, most of the models will only cover a subset of scenarios of a system while it will fail in others. Moreover, as the system changes or with new scenarios the physics-based modeling requires to be calibrated and updated with new assumptions. However, deep learning model will fail in cases there is not enough data for learning normal behavior, but the physics-based model has a restricted optimization space that is conditioned by the imitable laws of nature.

Pros/Cons of ResNet AE model: In this study, we used a ResNet Autoencoder model which is a pure data-driven unsupervised algorithm it learns not only temperature-related features as done by the physics-based model but also overall feature deviations that happen during attack scenarios. Therefore, the ResNet Autoencoder model performs better than the physics-based model which is confirmed by the experimental results. Moreover, the Resnet AE algorithm is unsupervised meaning does not require labeled data. Data labeling is time-consuming, and one cannot label all the anomalous behaviors of the system due to practical difficulties Further effort is required in identifying the root causes when using a deep learning model to detect anomalies as opposed to a physics-based model.

VI. CONCLUSION AND FUTURE WORK

The objective of this paper is to develop a physical anomaly detection approach without using labeled data to improve the resiliency of electric vehicle charging stations. First, we explored a physics-based system identification model for the anomaly detection task. However, we demonstrate that there are limitations to this approach. Specifically, with attacks that alter the system but result in input and output data that follow the physics of the system. Therefore, we explored ResNet AE-based approach for unsupervised physical anomaly detection and our outcomes show that the ResNet AE approach has improved detection performance. More precisely, we show that by using a ResNet AE it is possible to detect the attacks with a High Accuracy (96.82), Recall(1.00), Precision (92.43), and F1 score (96.07). In the future, we intend to improve our anomaly detection system by incorporating cybersecurity-related scenarios of EV charging systems.

ACKNOWLEDGMENT

This work was supported in part by the Department of Energy through the U.S. DOE Idaho Operations Office under Contract DE-AC07-05ID14517, and in part by the Commonwealth Cyber Initiative, an Investment in the Advancement of Cyber Research and Development, Innovation and Workforce Development (cyberinitiative.org).

REFERENCES

[1] A. Brown, A. Schayowitz, and E. White, "Electric vehicle charging infrastructure trends from the alternative fueling station locator: Fourth quarter 2021," National Renewable Energy Lab.(NREL), Golden, CO (United States), Tech. Rep., 2022.

[2] J. Antoun, M. E. Kabir, B. Moussa, R. Atallah, and C. Assi, "A detailed security assessment of the ev charging ecosystem," *IEEE Network*, vol. 34, no. 3, pp. 200–207, 2020.

[3] S. Soltan, P. Mittal, and H. V. Poor, "{BlackIoT}:{IoT} botnet of high wattage devices can disrupt the power grid," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 15–32.

[4] J. Johnson, B. Anderson, B. Wright, J. Quiroz, T. Berg, R. Graves, J. Daley, K. Phan, M. Kunz, R. Pratt, T. Carroll, L. R. ONeil, B. Dindlebeck, P. Maloney, D. O'Brien, D. Gotthold, R. Varriale, T. Bohn, and K. Hardy, "Cybersecurity for electric vehicle charging infrastructure." 7 2022. [Online]. Available: https://www.osti.gov/biblio/1877784

[5] R. Metere, M. Neaimeh, C. Morisset, C. Maple, X. Bellekens, and R. M. Czekster, "Securing the electric vehicle charging infrastructure," *arXiv preprint arXiv:2105.02905*, 2021.

[6] A. Meintz, M. Starke, and T. Bohn, "Charging infrastructure technologies: Development of a multiport, >1 mw charging system for medium-and heavy-duty electric vehicles," National Renewable Energy Lab.(NREL), Golden, CO (United States), Tech. Rep., 2022.

[7] H. S. Mavikumbure, C. S. Wickramasinghe, D. L. Marino, V. Cobilean, and M. Manic, "Anomaly detection in critical-infrastructures using autoencoders: A survey," in *IECON 2022 – 48th Annual Conference of the IEEE Industrial Electronics Society*, 2022, pp. 1–7.

[8] J. Johnson, J. Hoaglund, R. Trevizan, and T. Nguyen, *Chapter 18: Physical Security and Cybersecurity of Energy Storage Systems*, 01 2021.

[9] T. Phillips, H. Mehrpouyan, J. Gardner, and S. Reese, "A covert system identification attack on constant setpoint control systems," in *2019 Seventh International Symposium on Computing and Networking Workshops (CANDARW)*, 2019, pp. 367–373.

[10] R. Jin, B. Wei, Y. Luo, T. Ren, and R. Wu, "Blockchain-based data collection with efficient anomaly detection for estimating battery state-of-health," *IEEE Sensors Journal*, vol. 21, no. 12, pp. 13 455–13 465, 2021.

[11] Y. Li, L. Zhang, Z. Lv, and W. Wang, "Detecting anomalies in intelligent vehicle charging and station power supply systems with multi-head attention models," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 1, pp. 555–564, 2020.

[12] M. Basnet and M. Hasan Ali, "Deep learning-based intrusion detection system for electric vehicle charging station," in *2020 2nd International Conference on Smart Power & Internet Energy Systems (SPIES)*, 2020, pp. 408–413.

[13] C. S. Wickramasinghe, D. L. Marino, H. S. Mavikumbure, V. Cobilean, T. D. Pennington, B. J. Varghese, C. Rieger, and M. Manic, "Rx-ads: Interpretable anomaly detection using adversarial ml for electric vehicle can data," 2022. [Online]. Available: https://arxiv.org/abs/2209.02052

[14] C. S. Wickramasinghe, D. L. Marino, and M. Manic, "Resnet autoencoders for unsupervised feature learning from high-dimensional data: Deep models resistant to performance degradation," *IEEE Access*, vol. 9, pp. 40 511–40 520, 2021.

[15] L. Beal, D. Hill, R. Martin, and J. Hedengren, "Gekko optimization suite," *Processes*, vol. 6, no. 8, p. 106, 2018.