

Benchmark Evaluation of Anomaly-Based Intrusion Detection Systems in the Context of Smart Grids

Ömer Sen^{*†}, Simon Glomb^{*}, Martin Henze^{‡§}, Andreas Ulbig^{*†}

^{*}IAEW, RWTH Aachen University, Aachen, Germany — [†]Digital Energy, Fraunhofer FIT, Aachen, Germany

Email: {o.sen, a.ulbig}@iaew.rwth-aachen.de, simon.glomb@rwth-aachen.de, {oemer.sen, andreas.ulbig}@fit.fraunhofer.de

[‡]SPICE, RWTH Aachen University, Aachen, Germany — [§]CA&D, Fraunhofer FKIE, Wachtberg, Germany

Email: henze@cs.rwth-aachen.de, martin.henze@fkie.fraunhofer.de

Abstract—The increasing digitization of smart grids has made addressing cybersecurity issues crucial in order to secure the power supply. Anomaly detection has emerged as a key technology for cybersecurity in smart grids, enabling the detection of unknown threats. Many research efforts have proposed various machine-learning-based approaches for anomaly detection in grid operations. However, there is a need for a reproducible and comprehensive evaluation environment to investigate and compare different approaches to anomaly detection. The assessment process is highly dependent on the specific application and requires an evaluation that considers representative datasets from the use case as well as the specific characteristics of the use case. In this work, we present an evaluation environment for anomaly detection methods in smart grids that facilitates reproducible and comprehensive evaluation of different anomaly detection methods.

Index Terms—Intrusion Detection, Anomaly Detection, Smart Grids, Benchmarking, Explainability

I. INTRODUCTION

The evolution of power grids into smart grids (SGs) has brought changes to the information and communication technology (ICT) infrastructure. Traditional SGs used to have isolated and proprietary components, creating barriers to external third-parties. However, with the increasing penetration of ICT, the enhanced interconnection of SGs has resulted in new cybersecurity threats in the form of cyberattacks [1]. To better detect these threats, detection mechanisms are needed to observe and monitor the cybersecurity posture of SGs. Intrusion detection systems (IDSs) have become a standard security component for identifying unusual activities and cyber-threats in the system [2]. An IDS is a software designed to detect malicious activity on a network or system [3].

Despite significant development over the years, existing IDSs still face challenges in improving detection accuracy, reducing false positive rates, and detecting unknown attacks. Machine learning has emerged as a crucial research topic to address these issues, as it has shown the potential to achieve high accuracy and generalizability in detecting new attacks [4]. However, these approaches often suffer from high false positive rates [5] and lack explainability in issued alerts [6].

To minimize the effort and expertise required for post-processing a large volume of issued alerts and reduce the time needed to correctly identify and understand the situation, an anomaly detection system with high precision and explainable

alerts is necessary. However, this requires an anomaly detection approach with reliable and explainable detection quality that needs to be verified and validated.

To evaluate the properties of such a system, experimental studies are necessary, where approaches are examined based on different criteria, scenarios (datasets), and use cases. However, such studies often rely on the evaluation of a single dataset and lack the use of in-depth metrics to comprehensively evaluate a system [7]. Thus, a reproducible and comprehensive benchmark environment is needed to support research in this direction.

In this work, we address the challenges of reproducibility of performance evaluation using different datasets, evaluation of the explanatory power of different machine learning approaches, and comparison of different anomaly detection systems in the same environment, with the goal of studying and evaluating machine-learning based anomaly detection for SGs. Our contributions in this paper include presenting the current state of the art in machine learning testing and highlighting the benchmarking issues (Section III), describing the overall benchmark approach covering dataset selection to performance evaluation of predictions (Section IV), and demonstrating and discussing the capabilities of our proposed approach through a case study (Section V).

II. BACKGROUND & RELATED WORK

In this section, we describe IDS, its function in the context of anomaly detection, and its main features in terms of explainability. We also present related work in the context of machine learning-based IDS and benchmarking.

A. Intrusion Detection Systems

An IDS is a security solution, either software or hardware-based, that monitors and analyzes network or system activity to detect potential cyber threats like unauthorized access, misuse, or disruption. It can be implemented in different ways, including individual devices or network-wide, and can monitor various data types like device logs, network traffic, and process data to identify suspicious or malicious activity.

Anomalies in Industrial Control Systems (ICSs) like SGs can be caused by various factors, including attacks on different segments, altered communication functions, and faults like protocol errors or missing messages. Machine learning-based

IDSs use complex methods to detect anomalies, but explainability is an increasingly important feature to verify results and maintain confidence. Explainability refers to the ability to provide qualitative understanding of the relationship between model input and prediction in a user-friendly manner. Local explanations explain a model's behavior on one input instance, while global explanations understand the inner workings of a model's behavior. Explainability methods can be model-agnostic or model-specific.

B. Purdue Model

It is crucial to utilize realistic datasets that accurately reflect the complexity and diversity of real-world ICS environments when evaluating machine learning-based IDSs. The Purdue Reference Model, which is a component of Purdue Enterprise Reference Architecture (PERA) [8], provides a helpful framework for comprehending the various layers and components of ICS networks. Hence, it can be utilized as a guide for selecting datasets for benchmarking purposes.

The Purdue Model is frequently utilized as a reference model for designing, analyzing, and securing ICS networks, as it helps understand the architecture and components of an ICS. It splits the ICS architecture into two main zones: the Operational Technology (OT) zone, which comprises the physical equipment and operational processes monitoring and control systems, and the Information Technology (IT) zone, which encompasses the systems used for data management and communication.

The Purdue Model also includes a Demilitarized Zone (DMZ) between the IT and OT zones to control access between the two areas. Therefore, the Purdue Model divides the ICS architecture into six levels, with each level representing a different layer of the ICS system.

C. Related Work

The machine learning community has been exploring ways to create benchmark environments for evaluating the performance of different machine learning algorithms in detecting anomalies and intrusions in ICS. One approach is to create benchmark libraries such as the Penn Machine Learning Benchmark [9] and the Scientific Machine Learning Benchmark suite [10]. These libraries are intended to provide valuable resources for testing and comparing the performance of different algorithms. Another approach is to create datasets from actual ICS environments to evaluate the performance of algorithms. Researchers have proposed using datasets such as the Cyber-kit datasets [11] and the Numenta anomaly benchmark [12] to evaluate unsupervised anomaly detection techniques or IDSs in ICS. In addition, studies have focused on the use of datasets specific to SGs for evaluating the performance of machine learning algorithms in detecting anomalies in SG systems [13].

Several works have evaluated the accuracy and false positive rate of different anomaly detection approaches using datasets from real ICS environments [14]. Similarly, studies have proposed methods to assess the performance of machine

learning algorithms in detecting anomalies in SGs and evaluate their performance based on various metrics such as accuracy, precision, and recall [15]. Overall, benchmarking and creating datasets are important approaches to evaluate the effectiveness of machine learning algorithms in anomaly detection for ICS and SGs.

D. Problem Analysis

The field of machine learning-based anomaly detection for ICS is rapidly growing and diverse, with many datasets and benchmarks proposed to evaluate machine learning algorithms for this purpose. However, there is still a need for more specialized and comprehensive approaches that consider the unique characteristics and constraints of this domain. Specifically, existing benchmarks may not cover all relevant evaluation methods or consider all important factors, such as the explainability of algorithms and the statistical confidence of findings. Furthermore, SGs may have unique characteristics in terms of their inputs, metrics, or other domain-specific factors, which are not always considered in benchmarks. To address these issues, this paper aims to provide a comprehensive benchmark environment that evaluates machine learning-based anomaly detection for SGs using a diverse range of metrics and criteria, with a particular focus on the explainability of models.

III. ANALYSIS

In this section, we analyze requirements of a benchmarking machine learning-based anomaly detection in SGs.

A. Requirement

The reproducibility of a research study refers to the ability to obtain the exact same results. To ensure high reproducibility, various aspects need to be considered, such as datasets, data preprocessing, model training, model evaluation, and model software and hardware environment. In addition to these requirements for individual studies, there are also requirements for benchmarks in general, including reporting, coverage, and extensibility.

The requirements include using publicly available datasets related to ICS, integrating common preprocessing methods, allowing customization of preprocessing methods and training data, comparing multiple algorithms, providing analysis for binary classification, using statistical methods, controlling randomization, providing a software environment as a container, and investigating new methods. The environment should also provide information about the CPU and RAM, store necessary information for the reproduction of the study, and ensure that the datasets and software environment are accessible.

B. Datasets

By referring to the Purdue model (cf. Section II-B), datasets reflecting real-world attack scenarios can be used. It also allows the selection of datasets that represent specific levels or components of the ICS, thus making it more relevant to evaluate and compare machine learning-based IDSs for ICS environments. These datasets with realistic and actual attacks

enable accurate evaluation of IDSs' performance and reflect real-world scenarios, making the benchmark more realistic and applicable for the ICS environment.

Various datasets are available, each useful for benchmarking different types of IDSs. The Bot IoT dataset includes anomalies related to botnets categorized as Level 0, which represent physical components of a system [16]. The Power System dataset [17] focuses on anomalies at Level 1 and Level 2, which monitor and control physical equipment and operational processes, while the CICIDS-2017 [18] and CIDDS-002 datasets [19] provide a range of intrusion and attack scenarios assigned to Levels 3 and 4. The UNSW-NB15 [20] and NSL-KDD datasets [21] are focused on anomalies at Levels 4 and 5, managing logistics of productive operations and providing communications and data. These datasets provide a diverse set of anomalies, enabling evaluation of IDSs to detect and classify various types of network-based threats in ICS.

To evaluate machine learning-based IDSs in ICS, including SG research, it is important to have a benchmark environment that includes a diverse set of anomalies representative of the type of attacks that these systems may encounter in real-world scenarios. Selecting datasets that include anomalies representative of the different layers of the Purdue Model ensures that the machine learning-based IDSs are tested against a diverse set of attack scenarios that are relevant to ICSs. Publicly available, labeled, and network-based datasets are considered for the evaluation of multiple datasets. Additionally, it is critical to have explainable methods to understand the reasoning behind an IDS decision, building trust in these systems, and mitigating the impacts of false positives.

C. Machine Learning Model

The paper by Davis et al. [22] lists several important preprocessing steps for network-based IDSs. These steps include data transformations, data cleaning, data reduction, discretization, and re-labeling. However, feature construction and dataset creation are excluded as they are highly dependent on the specific protocols and features of the datasets, and this work aims to abstract from such specifics. It is recommended to consider all other preprocessing steps, and multiple preprocessing methods should be usable without impairing other methods, i.e. allowed to be omitted in evaluation.

Ensemble learning, which combines multiple learning algorithms, is often used in recent works to improve classification performance. In this paper, the term "stacking" is used synonymously with ensemble methods, and the models being combined are referred to as first-level models, while the model that combines their predictions is the second-level model. For the proposed framework, any stacking algorithm should be analyzable, and it should be possible to choose the number and type of base classifiers arbitrarily. Additionally, we want to be able to compare two arbitrary algorithms, and the data on which the base classifiers are trained should be customizable.

Evaluation metrics for IDSs are highly dependent on the specific application. High detection accuracy, runtime efficiency, and explainability are important criteria, and enhancing

the robustness of IDSs is a novel way to improve detection accuracy. Furthermore, diversity in ensemble methods is a fundamental issue and should be evaluated. The proposed framework should support research on these main directions for IDSs. Binary classification is chosen as the evaluation task, and multiple datasets should be used for evaluation. Statistical tests and confidence intervals should be used to quantify the generalizability of results in a reproducible manner (i.e. control of randomization).

D. Environment

To ensure reproducibility, it is recommended to use containers, as stated in [23]. We require that the container used in the proposed environment should be versioned to make the environment accessible and enable the creation of reproducible results. It is also important to report hardware information, specifically CPU and RAM, as relevant hardware information. This is necessary to assess studies conducted on different devices and potentially use them in the future.

All information pertaining to each study must be stored in a database to ensure reproducibility. This includes information about the dataset, preprocessing strategy, model training with hyperparameters, training method, model evaluation, as well as software and hardware environment. Our goal is to ensure that the reporting of each examination is reproducible and provides the context for reproducing the results.

Our work is limited to anomaly-based network IDSs with datasets specific to the SGs domain. We focus exclusively on binary classification and machine learning approaches, both supervised and unsupervised. However, we require the investigation of new methods, and the implementation and evaluation of new approaches should be possible. We will present a methodology as a reference implementation, which incorporates the requirements we have analyzed.

IV. METHODOLOGY

In this section, we present a concrete methodology on how the requirements can be fulfilled. With the methodology, we also present how our tool is constructed and what it provides.

A. Overview

We present a component diagram in Figure 1 which provides an overview of our design. Each methodology component can be associated with a specific component, thereby offering a tool that meets our established requirements for evaluating IDSs. Because machine learning algorithms can have hyperparameters that are arbitrarily nested, a relational database is unsuitable for this purpose. Therefore, we have chosen to use a document-oriented database to store the results from the environment, specifically the MongoDB database management system. BSON (Binary JSON) is the format for data storage, which is a superset of JSON that includes additional data types. Additionally, we provide a BSON-Schema to establish the structure for the results. The schema includes JSON-LD to provide linked data for the results, which we have used as an example and as a proof of concept that the proposed architecture is fully compatible with JSON-LD.

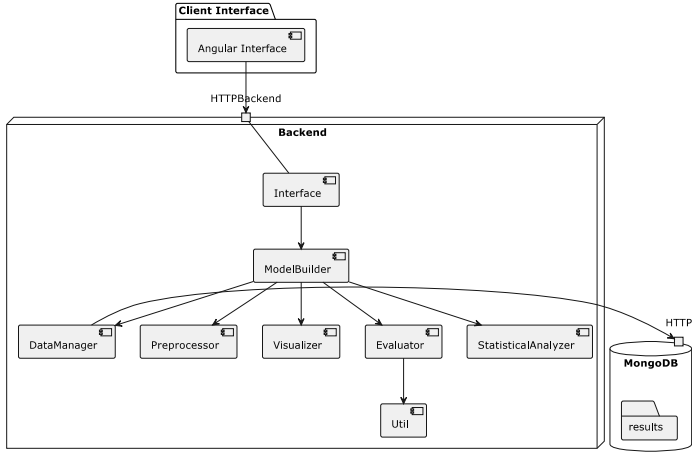


Fig. 1. UML Component-Diagram for the Core Components.

B. Preprocessing and Evaluation Strategy

In the following section, we will explain how we preprocess our datasets prior to training, and how we select the appropriate approach for assessing the explainability of our method.

The first step in preprocessing is data transformation. Normalization is considered to be a crucial technique to enhance model performance and computational efficiency for network-based IDSs. Moreover, we need to perform discretization, as symbolic features must be transformed into numerical features. As we cannot guarantee that the features are ordinal, we will use the commonly used One-Hot-Encoding method. For data reduction, we will focus on dimensionality reduction, and we will choose the widely used principal component analysis (PCA) method.

We have chosen the SHAP framework to generate explanations [24]. SHAP can explain the output of any machine learning model, and is based on game-theoretic concepts. The reason for selecting SHAP is its ability to produce global explanations, which can help to build trust in the IDS. Additionally, we deemed it necessary for the method to also generate local explanations for specific intrusion detection cases, as this can help to improve trust in the IDS when used in practice. This is because the same method can be used to evaluate a model and a particular case when the system is put into practice. Another mandatory requirement was that the method had to be model-agnostic, as this allows us to apply the method to all models and are not restricted on the IDSs that can be analyzed. Furthermore, SHAP is already being studied for use in IDSs.

The metrics typically used to measure the performance of IDSs include training time, prediction time, and explanation time. The literature suggests using the following metrics for classification in the intrusion detection context: False Positive Rate, AUC Score, Balanced Accuracy, and MMC. In the following, we will cover additional metrics for evaluating anomaly detection.

C. Measuring Explainability

We will evaluate the quality of explanations based on known properties as described in [25]. Our evaluation will cover three aspects: accuracy, stability, and degree of importance. Importantly, these quality measures are not specific to any particular explanation method, including SHAP. For SHAP, we will compute the SHAP values using test data as suggested by [26], which is relevant for computing all the explanation metrics.

To measure the accuracy of explanations, we aim to quantify the discrepancy between the output of a classifier predicted by explanations and the actual output of the model.

For measuring the stability of explanations, we will compute the maximum sensitivity of the explanation to input perturbations as suggested in [27]. The sensitivity should be as low as possible, and we will use the definition and approach provided by [28], which can be found in Equation 1.

$$\text{SENS}_{\text{MAX}}(\Phi, f_c, x, r) = \max_{\|x' - x\|_{\infty} \leq r} \|\Phi(f_c, x') - \Phi(f_c, x)\|_F \quad (1)$$

Here, $\|\cdot\|_F$ denotes the Frobenius norm, f_c is the classifier, Φ is the explanation method, x is the unlabeled data, and r is the perturbation radius. We found a perturbation radius of 0.01 suitable for normalized data through experiments.

$$\text{AUC}_{\text{MoRF}}(\Phi, f_c, x, r) = \sum_{k=2}^D \frac{f_c(x_{\text{MoRF}}^{(k-1)}) + f_c(x_{\text{MoRF}}^{(k)})}{2} \quad (2)$$

We will also use the Area Under the Most Relevant First Perturbation Curve (AUC-MORF) metric to evaluate the degree of importance of the features declared by the explanation $\phi(f_c, x)$. The metric measures how fast the output of the classifier decreases as we progressively remove information from the input x , based on the importance declared by the explanation. We will use the definition of Area Under Perturbation Curve (AUPC) from [28]. The AUC-MORF metric is computed by perturbing the most important features of the data $x_{\text{MoRF}}^{(k)}$, where k denotes the number of most important features perturbed. We will use the formula presented in Equation 2 to compute the area under the perturbation curve. The smaller the area under the curve, the more reliable the information is, as it confirms that the declared most important features are indeed the most important ones.

D. Measuring Robustness

The stability of explanations has been measured using the concept of perturbations and sensitivity. Both classifiers' decisions and interpretability methods are sensitive to input perturbations and susceptible to adversarial attacks. Adversarial robustness, which represents the worst-case scenario for any unforeseen randomness, is focused on instead of robustness to random noise. The robustness we investigate is the robustness of an already learned classifier against evasion attacks. Poisoning attempts, where malicious training data is injected into the learning system, are not included in our measurement. The adversarial robustness of a data instance

x on the classifier f_c is defined as the minimum perturbation r that causes f_c to misclassify x .

$$\Delta_{\text{adv}}(x; f_c) = \min_{\substack{r \in \mathbb{R}^d \\ f_c(x) - f_c(x+r) \neq 0}} \|r\|_2 \quad (3)$$

We measure the sensitivity of the classifier using the Lipschitz constant, which is the maximum ratio between output and input variations of f_c .

$$L(f_c) = \max_{\substack{x_i, x_j \\ \in \mathcal{T}}} \frac{\|f_c(x_i) - f_c(x_j)\|_2}{\|x_i - x_j\|_2} \quad (4)$$

We use a lower bound for the true global Lipschitz constant, and it is analyzed with respect to the L_2 -norm. Other measures and approximations of the Lipschitz constant are available, but they are model-specific and, therefore, disregarded.

E. Statistical Methods

This section focuses on the utilization of statistical methods for quantifying the generalizability of results. The work addresses the scenario of statistical comparison across multiple datasets, where sample size refers to the number of datasets used. Nonparametric tests are preferred in this case due to the lack of assumptions about homogeneity of variance. When comparing learning algorithms, the Wilcoxon signed ranks test is recommended for two-algorithm comparisons, while the Friedman test is suggested for comparisons involving more than two classifiers. Since our evaluation requires comparing two arbitrary algorithms, we will use the Wilcoxon signed ranks test. However, it cannot be computed if all data points for both groups are equal.

Estimation statistics, which uses effect sizes and confidence intervals, can be used in addition to traditional hypothesis testing and is sometimes recommended to replace null hypothesis significance testing. Effect size is used to quantify the magnitude of an effect in our case of comparing algorithm metrics, and various measurements such as percentage, median, correlation, or Cohen’s d (recommended for comparing learning algorithms) can be used as effect size. A value of 0.2 or lower is considered a small effect, 0.5 is considered medium, and 0.8 or higher is considered large for Cohen’s d . We will use this terminology in our scenario evaluation. Cohen’s d can have values in the range of $(-\infty, \infty)$, but it cannot be computed if both classifiers have no deviation.

Confidence intervals are used to estimate the range of an unknown parameter. In our case, we are interested in confidence intervals for the difference of sample means. A sample size of 30 is often suggested as the minimum requirement for assuming a normal distribution of mean measurements. However, since this assumption does not hold for our analysis, we will use bootstrap confidence intervals [29] based on the implementation described in [30].

The Gardner-Altman estimation plot [31] is considered suitable and is a standard plot for utilization with estimation statistics. It presents all data points as a swarm plot to display all observed values, and shows the point estimate (e.g. mean difference) with its 95% bootstrap confidence interval

on a separate difference-axis. Figure 2 shows a Gardner-Altman estimation plot for the mean difference of explanation sensitivity in a comparison of two explanation methods applied on ensembles.

V. RESULT

In this section, we present the general procedure of using our environment, and then present a case study on the investigation of a method for the use in the intrusion detection.

A. Objective

We have chosen the aspect of explainability from the topics we have covered, and demonstrated that our evaluation goes beyond classical metrics. Moreover, we have shown that our evaluation can be more comprehensive than evaluating on a single dataset. When investigating a specific IDS, it is useful to compare it with another system. We abstract from the entire IDS and only consider the underlying detection method. The user inputs two learning algorithms that must inherit the interface from the sklearn class `sklearn.base.BaseEstimator`. Without using our environment, there would be more work to do to evaluate the algorithms. We encapsulate the knowledge presented in the methodology to facilitate comparable research results. The environment adds value by providing a state-of-the-art evaluation process that reduces the workload for researchers. New insights can be incorporated into this environment by extending the presented components.

B. Investigation Subject

We investigate a new method called basic join, which focuses on the explainability of ensembles of classifiers. We aim to explore the possibility of utilizing this method for intrusion detection systems in SGs using the proposed environment. The issues that our proposed environment addresses are as follows: dataset integration, which is usually time-consuming but made easier with the proposed environment; selection of domain-specific metrics and evaluation measures, which are already done and do not need to be re-implemented; storage of results in a reproducible way for future re-evaluation without loss of configurations; and provision of confidence estimation of results for the domain of SG. The basic join, which we want to evaluate, provides an explanation for the entire model. Unlike the traditional approach of treating the whole model as a black box, the basic join determines the importance of features of all first-level models and then computes the dot product with the vector containing the importance of features of the second-level model. The features of the second-level model are the predictions of the first-level models, so the importance can be seen as the contribution of the first-level models to the second-level output. As models, we chose two ensembles, one using the traditional approach and the other using the basic join. Both ensembles consist of three learning algorithms: Logistic Regression, Decision Tree, and Multi-Layer Perceptron. For preprocessing, we used normalization and did not include methods like principal component analysis as it impacts explainability. Our models will be evaluated

on all included datasets using the metrics proposed in the methodology, with a focus on the explainability aspect.

C. Case Study




Fig. 2. Gardner-Altman Estimation Plot for the Mean Difference of the Explanation Sensitivity (Expl.Sens.) for two Ensembles (both consisting of the same Logistic Regression, Decision Tree and Multi-Layer Perceptron).

Regarding the explainability metrics, we observed the following: There was a small difference in explanation error (Cohen’s $d = -0.14$). No difference was observed in the area under the precision-recall curve (AUPC), which means that at least the first five features have the same order of relevance for both approaches (as we chose to evaluate five features for AUPC). The most significant effects were seen in explanation time (Cohen’s $d = -0.35$) and explanation sensitivity (Cohen’s $d = 0.31$). We present a Gardner-Altman Estimation plot for the latter metric in Figure 2. From the results, we can see that on average, the traditional approach has slightly lower explanation error and requires less time for explanation creation compared to the basic join. However, the new method appears to have benefits in terms of the robustness of explanations on some datasets and may be advantageous for dealing with data noise. The datasets where the basic join performs better are the Power System dataset and the CIDD5-002 dataset. By utilizing our proposed environment, we can infer that the method under scrutiny has a lower susceptibility to noise, but it has certain downsides concerning computational resources. Overall, the investigated method has this trade-off between noise sensitivity and computational requirements.

VI. CONCLUSION

SG cybersecurity is crucial for securing SG tasks, and anomaly detection is a critical technology in addressing this challenge. This study introduces an evaluation environment for anomaly detection methods in SG that allows for reproducible and comprehensive evaluation of different techniques, providing a general framework for evaluations and comparisons. Our contributions to the energy community include conducting a detailed analysis of benchmark environment requirements for anomaly-based network IDS, providing a reference implementation for building such an environment for the energy sector by aggregating results from various research branches, and

demonstrating a proof of concept for using statistical methods for cross-dataset evaluation in the context of IDS. Future work includes incorporating measures of classifier diversity and combining datasets to enhance evaluation metrics, towards standardized benchmark environments for evaluating IDSs.

REFERENCES

- [1] S. Pan *et al.*, “Developing a hybrid intrusion detection system using data mining for power systems,” *IEEE Transactions on Smart Grid*, 2015.
- [2] J. Shun *et al.*, “Network intrusion detection system using neural networks,” in *IEEE ICNC*, 2008.
- [3] J. M. Wagle, “Utilizing the shap framework to bypass intrusion detection systems,” Master’s thesis, The University of Bergen, 2021.
- [4] H. Liu *et al.*, “Machine learning and deep learning methods for intrusion detection systems: A survey,” *applied sciences*, 2019.
- [5] R. Samrin *et al.*, “Review on anomaly based network intrusion detection system,” in *IEEE ICEECCOT*, 2017.
- [6] D. Kus *et al.*, “A False Sense of Security? Revisiting the State of Machine Learning-Based Industrial Intrusion Detection,” in *ACM CPSS*, 2022.
- [7] K. Wolsing *et al.*, “IPAL: Breaking up Silos of Protocol-dependent and Domain-specific Industrial Intrusion Detection Systems,” in *RAID*, 2022.
- [8] T. Williams, “An overview of pera and the purdue methodology,” *Architectures for Enterprise Integration*, 1996.
- [9] R. S. Olson *et al.*, “Pmlb: a large benchmark suite for machine learning evaluation and comparison,” *BioData mining*, 2017.
- [10] J. Thiyyagalingam *et al.*, “Scientific machine learning benchmarks,” *Nature Reviews Physics*, 2022.
- [11] S. Mubarak *et al.*, “Ics cyber attack detection with ensemble machine learning and dpi using cyber-kit datasets,” in *IEEE ICCCE*, 2021.
- [12] A. Lavin *et al.*, “Evaluating real-time anomaly detection algorithms—the numenta anomaly benchmark,” in *IEEE ICMLA*.
- [13] G. Bernieri *et al.*, “Evaluation of machine learning algorithms for anomaly detection in industrial networks,” in *IEEE M&N*, 2019.
- [14] S. Liyakkathali *et al.*, “Validating anomaly detection mechanisms in industrial control systems,” in *Proceedings of TMCE*, 2020.
- [15] M. Mohammadpourfard *et al.*, “Benchmark of machine learning algorithms on capturing future distribution network anomalies,” *IET Generation, Transmission & Distribution*, 2019.
- [16] N. Koroniotis, “Designing an effective network forensic framework for the investigation of botnets in the internet of things,” Ph.D. dissertation, UNSW Sydney, 2020.
- [17] Mississippi State University and Oak Ridge National Laboratory, “Power system attack datasets,” 2014.
- [18] I. Sharafaldin *et al.*, “Toward generating a new intrusion detection dataset and intrusion traffic characterization,” *ICISSp*, 2018.
- [19] M. Ring *et al.*, “Creation of flow-based data sets for intrusion detection,” *JIW*, 2017.
- [20] N. Moustafa *et al.*, “The evaluation of network anomaly detection systems: Statistical analysis of the unsw-nb15 data set and the comparison with the kdd99 data set,” *ISJ*, 2016.
- [21] M. Tavallaee *et al.*, “Nsl-kdd dataset,” 2023, canadian Institute for Cybersecurity, University of New Brunswick.
- [22] J. J. Davis *et al.*, “Data preprocessing for anomaly based network intrusion detection: A review,” *computers & security*, 2011.
- [23] R. Tatman *et al.*, “A practical taxonomy of reproducibility for machine learning research,” *openreview.net*, 2018.
- [24] G. Van den Broeck *et al.*, “On the tractability of shap explanations,” *Journal of Artificial Intelligence Research*, 2022.
- [25] S. R. Islam *et al.*, “Towards quantification of explainability in explainable artificial intelligence methods,” in *FLAIRS-32*, 2020.
- [26] C. Molnar, *Interpretable machine learning*, 2020.
- [27] F. VANNI, “Design and validation of a comparison methodology for explainable ai techniques,” *webthesis.biblio.polito.it*, 2022.
- [28] I. Kakogeorgiou *et al.*, “Evaluating explainable artificial intelligence methods for multi-label deep learning classification tasks in remote sensing,” *IJAEOG*, 2021.
- [29] B. Efron *et al.*, *An introduction to the bootstrap*. CRC press, 1994.
- [30] J. Ho *et al.*, “Moving beyond p values: data analysis with estimation graphics,” *Nature methods*, 2019.
- [31] M. J. Gardner *et al.*, “Confidence intervals rather than p values: estimation rather than hypothesis testing,” *Br Med J (Clin Res Ed)*, 1986.