Reproducing Random Forest Efficacy in Detecting Port Scanning

A Preprint

Jason M. Pittman
ORCID: 0000-0002-5198-8157

Abstract

Port scanning is the process of attempting to connect to various network ports on a computing endpoint to determine which ports are open and which services are running on them. It is a common method used by hackers to identify vulnerabilities in a network or system. By determining which ports are open, an attacker can identify which services and applications are running on a device and potentially exploit any known vulnerabilities in those services. Consequently, it is important to detect port scanning because it is often the first step in a cyber attack. By identifying port scanning attempts, cybersecurity professionals can take proactive measures to protect the systems and networks before an attacker has a chance to exploit any vulnerabilities. Against this background, researchers have worked for over a decade to develop robust methods to detect port scanning. One such method revealed by a recent systematic review is the random forest supervised machine learning algorithm. The review revealed six existing studies using random forest since 2021. Unfortunately, those studies each exhibit different results, do not all use the same training and testing dataset, and only two include source code. Accordingly, the goal of this work was to reproduce the six random forest studies while addressing the apparent shortcomings. The outcomes are significant for researchers looking to explore random forest to detect port scanning and for practitioners interested in reliable technology to detect the early stages of cyber attack.

Keywords random forest, machine learning, port scanning, cybersecurity, algorithms, training data

1 Introduction

Cybersecurity breaches are a persistent threat in the digital world.. Although phishing and malware account for a considerable number of cybersecurity incidents, 45% are caused by network-based cyber attacks [1]. Cyber attacks follow a set pattern or sequence of steps, although the number of steps varies among existing models and methodologies. Nonetheless, reconnaissance is universally recognized as the first step in these attacks. Reconnaissance typically involves performing some form of port scanning.

Port scanning is a method of identifying active services on a target system. The technique can be both a legitimate means to verify one's configurations [2] or as a precursor to an intrusion [3, 4]. Thus, a general problem is differentiating between what may be an authorized host enumeration versus a malicious scan. Moreover, if port scanning is a standard prelude to cyber attack, then we want to have a way to detect the technique with high certainty. To this end, there is a literature detailing a variety of port scan detection mechanisms. The literature covers intrusion detection systems relying on static signatures [5] to advanced, modern machine learning (ML) techniques [6].

Most recently, a systematic review [7] collected the body of work using ML to detect port scanning published since 2021. The review found 15 studies exploring five different algorithms. Those 15 studies demonstrated algorithm efficacy ranging from 37.92% to 100%. Based on the span of the algorithm efficacy, the review suggested follow-up replication of those studies be performed to determine why algorithm performance results varied so much across the literature. However, replication of all five algorithms is infeasible when accounting for the variations in hyperparameters, model training mechanisms, lack of available code, and so forth. Accounting for such is necessary because the majority (13 of 15) studies do not include or link to source code containing these implementation details. Therefore, we choose to focus on the random forest algorithm in this work because of its simple architecture but robust set of options.

The remainder of this work is organized in a way which (a) situates the reproduction of random forest research in existing knowledge and (b) maximizes understanding of the cutting edge. The first is achieved by discussing port scanning, the random forest algorithm, and port scan detection using random forest implementations. Thereafter, we discuss the research method and protocol used to generate and analyze data. Finally, we present the results of the analysis.

2 Related Work

The work most proximal to this study exists in two categories: scanning TCP/IP ports, random forest in ML, and the use of random forest to detect port scans. The following background discussion is not meant to be comprehensive, however. Rather, we present notable and relevant research. The intention is to construct a suitable conceptual framework for understanding how and why this work has significance to the field.

2.1 Port Scanning

The origin of the term *port scanning* in the academic literature can be traced back to the late 1980s and early 1990s as the Internet was growing and becoming more widely used. Port scanning involves using features of TCP/IP to gather information about computing systems on a network by identifying open or closed ports. The first reference to port scanning in the literature is found in the work of Fyodor [8] which described a method for determining the operating system of a remote host by sending probes to specific ports and analyzing the responses. This paper also introduced the first version of the open-source tool *nmap*. Later, De Vivo et al. [9] provided a classification of port scanning techniques and procedures into different types such as TCP connect scans, SYN scans, indirect scans, stealth scans, and others. Later, Barnett et al. [10] presented a classification system for network scanning techniques based on the level of interaction with the target system, the type of information gathered, and the purpose of the scan. Bou et al. [11] demonstrated a comprehensive overview of different types of cyber scanning techniques that are used to identify various features of networks, which can be divided into passive and active scanning techniques.

2.2 Random Forest

Random forest is a supervised ensemble ML algorithm that leverages a set (i.e., forest) of decision trees to make a prediction [12]. The name random forest refers to the fact that the each tree in the forest is trained on a random subset of the data and a random subset of the features. A final prediction is made by combining the predictions of all the trees in the forest, either by majority vote (for classification problems) or by averaging (for regression problems). While random forest can act as a regression tool and a classifier, the algorithm is particularly robust for classification problems [12, 13, 14, 15]. This is true particularly in cases where there are many features and interactions among features.

While the literature exhibits a variety of differing random forest algorithm implementations, all consist of four general steps:

- Select random samples from a given dataset.
- Construct a decision tree for each sample and get a prediction result from each decision tree.
- Perform a vote for each predicted result.
- Select the prediction result with the most votes as the final prediction.

2.2.1 Hyperparameters and Parameters

Hyperparameters and parameters are not synonymous. Most classification algorithms use hyperparameters. Hyperparameters are set before training a model and cannot be estimated from the data during training [12, 13]. The hyperparameters control the behavior and performance of the model as well as optimization. Conversely, parameters are estimated from the data during training of a machine learning model. These values are learned by the model during optimization and are input to generate predictions.

To that end, the set of common hyperparameters one can manipulate when using random forest includes:

- Estimators as the number of decision trees in the forest.
- Criterion as the measure used to split the tree at each node.
- The maximum depth of the decision trees.
- A minimum number of samples required to split an internal node.
- The minimum number of samples required to construct a leaf node.
- The number of features to consider when looking for the best split.
- Whether or not to sample with replacement.

A crucial hyperparameter to adjust is the number of random features to be taken into account at each split point. Another important aspect to adjust is the depth of the decision trees [13]. Although deeper trees tend to overfit the training data, they also exhibit lower correlation, which can potentially enhance the performance of the ensemble. Nonetheless, node size, tree quantity, and number of sampled features are generally accepted as mandatory to set [16, 17].

Still, even with perfect hyperparameters controlling training, random forest has advantages and disadvantages. Awareness of both can lead to more appropriate selection of the algorithm for a given problem. The same is true at the level of how the algorithm is configured through hyperparameters.

2.2.2 Advantages of random forest

The main advantage of Random Forest classifiers is that they are less likely to overfit than single decision trees, since the combination of multiple trees helps to average out the noise in the data [16, 17]. Random forest is also relatively easy to implement and use. This is because the algorithm produces trained models requiring little tuning compared to other machine learning algorithms [12, 13].

Another significant advantage of random forest is it is an ensemble method [13]. Other classifiers like logistic regression, support vector machine, and naive Bayes are single models. Random forest also uses the bagging technique to train each tree in the forest on a random subset of the data and a random subset of the features [12, 17, 13]. This helps to reduce the risk of overfitting and improve the stability of the model. As well, this allows random forest to compute feature importance, which is a measure of the contribution of each feature to the prediction [12, 13]. This can be useful for feature selection and interpretability.

2.2.3 Disadvantages of random forest

In terms of disadvantages, random forest classifiers may not perform as well on highly complex and non-linear problems [16, 17], where other algorithms like neural networks may be more suitable. On that note, random forest classifiers can be computationally expensive and slow to make predictions, especially for large datasets and many trees in the forest [13]. Further, random forest classifiers may not handle unbalanced data well, where one class is underrepresented [12, 17, 13]. This can lead to biased results towards the majority class. Finally, while random forest classifiers can provide feature importance such are not as interpretable as other models like linear regression where the coefficients have a clear meaning.

Regardless of advantages and disadvantages, random forest features prominently in port scan detection research. While collective results of the research since 2021 demonstrated high efficacy, the deviation of accuracy values appeared significant. The following section describes those values along with other pertinent random forest implementation details.

2.3 Detecting Port Scanning

Port scanning, as a reconnaissance technique, is detectable. ML is a compelling solution to detecting otherwise undetectable port scans because of its ability to correlate seemingly unrelated features across enormous datasets. Yet, not all ML algorithms work in the same way or have the capability to address the same problem.

Furthermore, there are a variety of ML algorithms types-classification, regression, deep learning, and so forth- with a diversity of implementation variations. This work focused on random forest and therefore we analyzed the following background works only for related details to that algorithm.

Table 1: Random Forest Algorithm Literature

Authors	Accuracy	Dataset
Algaolahi et al.	99.75	CICIDS2017
Baah et al.	99.98	CICIDS2017
Sirisha et al.	78.09	NSLKDD
Sirisha et al.	84.14	CICIDS2017
SaiKiran et al.	99.93	CICIDS2017
Mohseni et al.	99.94	CICIDS2017
Bertoli et al.	96.00	MAWILab
Bertoli et al.	100.00	Bonafide

The six studies [18, 19, 20, 21, 22, 23] in Table 1 experimented with port scan detection using different RF algorithm implementations. Overall, random forest performance ranged from 78.09% to 100% across those studies. One paper [19] included a link to source code in a public repository (e.g., GitHub). Four different datasets were used, three of which do not appear in other algorithm categories.

Two studies discussed the types of port scans present in training and evaluation data. SaiKiran et al. [23] mentioned *port sweep* but did not specify further. Bertoli et al. [19] conducted training and evaluating against the full spectrum of port scan types. Further, Bertoli et al. included port scan data from five different port scan tools.

Algaolahi et al. [18] used a 20% test, 80% training data split. Random forest hyperparameters were not provided. Baah et al. [22] also used a 20-80% data split. Additionally, the authors specified a 10-fold cross-validation value for training. Other hyperparameters were absent.

Meanwhile, Bertoli et al. [19] provided perhaps the most detailed description of their experimental configuration. Random forest grid-search was performed with estimators set to 10, 50, 100, and 200. The criteria were gini and entropy. A max depth of 5 and 10. Lastly, the authors set a class weight of *none, balanced, and balanced subsample*. The model evaluation used stratified k-fold (k=10) cross-validation.

Mohseni et al. [21] also described a set of hyperparameters for their implementation of the random forest algorithm. However, details were unclear. It is possible the class weight was balanced. Further, it seemed the estimators were set to 15, 25, 50, and 100. Training data split and cross-validation were not mentioned.

Less revealing, SaiKiran et al. [23] did not indicate any random forest hyperparameters, training data split, or cross-validation mechanism. Sirisha et al. [20] provided some hyperparameters such criterion as gini, estimator of 100, but nothing else. The authors used a data split of 25% test and 75% training.

3 Method

A recent systematic review [7] suggested follow-up replication be performed for research using ML to detect port scanning. Research replication and reproduction is often mentioned but seldom implemented. For completeness, replication duplicates a research study with the same methods and setup to verify and increase trust in results [24]. The feasibility of conducting replication depends on the degree to which instrumentation and experimental setups are discussed in the source studies [25, 26, 27]. Reproduction is related but different insofar as it recreates a study using different instruments or data sets to demonstrate generalizability in the source studies' results [24]. Reproduction does not depend on instrument or experimental setup details [25, 26, 27].

However, as noted both in the systematic review and the related work in this research, the body of random forest port scan detection literature lacks the necessary details to establish a replication. Thus, a reproduction is the most appropriate method [24].

Scientific reproduction is meaningful when such work follows the source literature as closely as possible, is fully transparent in its method and materials, and rigorously compares reproduction results to the source research. Thus, we pose a single research question to guide the reproduction process. The question is: to what extent does reproduction produce similar results to the original studies, and if not are the differences statistically significant? To generate an answer we measured the efficacy of the random forest algorithm under several implementations.

3.1 Protocol

With the above in mind, we constructed a reproduction protocol based on (a) extracted details from the source studies where available and (b) knowledge of how random forest can be implemented. Having a clear protocol ensured we handled important steps in the correct sequence. Moreover, a clear protocol is a key tenet in reproduction, especially when source work lacks clarity.

1. Initialize the environment

2. Select a dataset for random forest training and testing

3. Conduct necessary data pre-processing

4. Run a random forest trial for hyperparameter set A

5. Record the `score` output

6. Re-initialize the environment

7. Run a random forest trial for hyperparameter set B

8. Record the `score` output

9. Repeat steps 6 - 8 until all hyperparameter sets are trialed

3.2 Environment

We constructed the operational environment in two stages: hardware and software. The hardware consisted of an AMD Ryzen 9 5900X CPU, 128 gigabytes of RAM, and a Nvidia 3090 RTX GPU. The system ran Ubuntu 22.04 for an operating system. The ML stack included `Python 3.11.1`, `Numpy 1.24.1`, `Pandas 1.5.3`, and `Scikit-Learn 1.2.1`.

3.3 Datasets

The source studies used four different datasets. The CICIDS2017 dataset includes eight different files, one being explicitly labeled as having *port scan* material (Friday-WorkingHours-Afternoon-PortScan.pcap_ISCX). None of the studies using this dataset detail which of the eight files were used for training and testing. Further, recent work [28] discovered significant issues with the CICIDS2017 set including duplicate data and mislabeling. The NSLKDD dataset also includes multiple files. Sirisha et al. [20] used the KDDTrain+ and KDDTest+ files. However, Bertoli et al. [19] noted limitations of the data which rendered use of the data in ML-based solutions impractical.

In turn, Bertoli et al. used two different datasets. One dataset contained laboratory generated traffic (attack) while the other the contained traffic captured from an internet network segment (bonafide). We selected the bonafide dataset because it included both legitimate as well as port scan traffic and therefore more closely matched the contents of the CICISD2017 and NSLKDD datasets. Moreover, we manually inspected the bonafide dataset to ensure a variety of port scan tools and techniques were included. For tools, the dataset captured nmap, masscan, unicornscan, zmap, and hping.

We adhered to the data pre-processing process outlined by Bertoli et al. [19] with one deviation. We opted to save the final processed dataset to comma separated value file. Doing so saved time and reduced error surface given we were going to run multiple random forest trials. The dataset is available in the companion GitHub repository.

3.4 Hyperparameters

There are five parameters within the set of hyperparameters that can impact random forest accuracy based on Bertoli et al. [19] study implementation. We gave critical attention to three as we constructed the sets for reproduction trials. First, we let the two lists of `n_estimators` from source studies be hyperparameter sets A and B. This parameter defines the number of trees in the model. Then, we paired value definitions of `max_depth` also from source work to those sets. Here, the deeper the tree, the more splits can occur causing the model to capture more information about the data. Lastly, we isolated `min_samples_leaf` as the number of samples required

to be at a leaf node. No source work set this parameter so we let the default value hold for sets A and B.

The majority of source studies (4 of 6) did not specify hyperparameters. Thus, we took guidance from the literature to construct additional hyperparameter sets C and D. We used these sets as internal validity measures and for comparison of model efficacy under different random forest training conditions. Additionally, the sets C and D allow for comparison to the broader literature not reporting explicit implementation details.

Table 2: Hyperparameter sets and parameter values

Set	Estimators	Leaf	Depth
A	10, 50, 100, 200	1	5, 10
B	15, 25, 50, 100	1	5, 10
C	200, 500	14	4, 5, 6, 7, 8
D	200, 500	14	4, 5, 6, 7, 8

Set	Features	Criterion	Weight
A	sqrt	gini	balanced
B	sqrt	gini	balanced
C	sqrt	gini	balanced
D	log2	entropy	None

3.5 Instrumentation

We initially attempted to use the source code provided through Bertoli et al [19]. However, there were critical errors during runtime. We surmised these were due to version drift in core packages (e.g., `sklearn`). As well, because none of the other source studies offered code snippets or entire source, we opted to develop a standard random forest implementation based on available `Scikit-Learn` documentation. Further, in doing so we establish a base implementation for a known working algorithm and model. Further, doing so allowed us to modularize the code within a Jupyter notebook according to the protocol rather than modify an existing instrument.

We let the training and testing split be 70% and 30% across all reproduction trials. This seemed appropriate given the known configurations in the source studies. Further, each hyperparameter set included two tuning configuration methods: `RandomSearchCV` and `GridSearchCV`. Considering the majority of source studies did not include code or indicate which method was used, including both provides a comparative view for later analyses. Fit and scoring were achieved using the basic random forest methods.

We ran an internal pilot test to ensure the instrument functioned as expected. This included loading the data, preprocessing the data, instantiating the random forest algorithm with a hyperparameter set, and producing output. Lastly, we made the instrumentation and data used in this reproduction transparent. All related materials can be found in a public GitHub repository [29].

4 Results

We separate the results into four sections based on the set of random forest configurations. Each results table contains two rows. The first row demonstrates results associated with the `RandomSearchCV` hyperparameter set and the second shows results using `GridSearchCV`.

The first hyperparameter trial demonstrated macro averages of 99%. These results can be taken as a baseline because they closely align to the random forest implementation used in Bertoli et al. [19].

Table 3: Hyperparameter set A efficacy

Accuracy	Recall	Precision	F1
0.9970	0.99	0.99	0.99
0.9976	0.99	0.99	0.99

Changing the `n_estimators` to $[15, 25, 50, 100]$ produced the output in Table 4. While the macro averages remain at 99%, we observed a minor difference in the `GridSearchCV` results compared to Set A.

Table 4: Hyperparameter set B efficacy

Accuracy	Recall	Precision	F1
0.9970	0.99	0.99	0.99
0.9974	0.99	0.99	0.99

We changed the `n_estimators` for Set C as well, to 200 and 500. We also changed `max_depth` from $[5, 10]$ used in the prior two trials to $[4, 5, 6, 7, 8]$. As well, we moved from the default `min_samples_leaf` value of 1 to 14. Collectively, these tunings produced changes to accuracy but not recall, precision, or F1 (Table 5).

Table 5: Hyperparameter set C efficacy

Accuracy	Recall	Precision	F1
0.9939	0.99	0.99	0.99
0.9937	0.99	0.99	0.99

We modified the last hyperparameter set with `max_features` from sqrt to log2 and a `criterion` value from gini to entropy. We also let `class_weight` be None. We observed changes in efficacy across all four measures.

Table 6: Hyperparameter set D efficacy

Accuracy	Recall	Precision	F1
0.9947	0.98	0.99	0.98
0.9895	0.96	0.99	0.97

As a general reproduction of random forest efficacy in detecting port scans, we were interested in how the above results compare to the results from the source literature. To that end, we expanded Table 1 by adding results for Recall, Precision, and F1 when available. The NaN value indicates no data were available in that category from the source study.

Table 7: Efficacy comparison across source studies

Study	Accuracy	Recall	Precision	F1
Algaolahi	0.9975	0.9989	0.9975	0.9982
Baah	0.9998	0.9997	0.9999	0.9998
Sirisha	0.7650	0.6525	0.9721	0.7809
SaiKiran	0.9993	NaN	NaN	NaN
Mohseni	0.9964	NaN	NaN	NaN
Bertoli	NaN	NaN	NaN	1.0000

We statistically compared the `accuracy` results from our trials to the set of results from the source literature. We used a paired t-test and found no significant difference ($p = 0.4103$, $t = 0.9544$, $df = 3$, $std = 0.058$).

5   Conclusion

Port scanning is used by hackers to identify vulnerabilities in a network or system. Early detection of this step in the cyber kill chain can lend a sizable time advantage to network defense teams. Thus, much effort has been invested into developing techniques to do so. In that regard, researchers have put forth a variety of ML algorithms as solutions. Results from that work have shown promising results. However, whether such outcomes are generalizable has been unknown until now.

This study reproduced results from six studies published since 2021 using a random forest algorithm. We developed a new random forest instrument and ran it against one of the datasets detailed in one of the existing works. Further, to model potential variations in random forest implementation, we ran four trials with four different sets of hyperparameters.

Given the observable similarity of results, combined with the statistical comparison, the use of random forest to detect port scanning appears generalizable. However, a deeper inspection of the trial data revealed that two techniques- full connect and SYN scanning- from one of the tools (unicornscan) were detected with an `accuracy` well below 90%.

We assumed the source research implemented the random forest algorithm within a common, standard range of hyperparameters. We also assumed given all studies used comparable Python interpreters and core packages given the work occurred within the same timeframe. However, it is possible major version differences exist across the studies or abnormal hyperparameters were used. We also assumed the selected dataset from Bertoli et al. [19] did not contain duplicates, flaws, or mislabeled elements.

Based on the results, we recommend additional reproduction research be done with focus on the other ML algorithms. Doing so would expand the foundation of knowledge and provide research and practitioners with confidence in the applicability of using ML to detect port scans. Furthermore, development of a real-time engine using confirmed generalizable ML algorithms would allow

for experimentation with dynamic and distributed port scanning techniques. A component in such an engine might also give real-time measurements of sustainability or green compute impact. Lastly, deeper exploration of why unicornscan deviated from the high watermark predictions associated with other tools is warranted. Such investigation may explore novel port scanning techniques capable of evading ML-based solutions.

References

[1] Statista Research Department. Most common action varieties in data breaches worldwide in 2019, 2023.

[2] Weijie Wang, Baijian Yang, and Yingjie Victor Chen. Detecting subtle port scans through characteristics based on interactive visualization. In *Proceedings of the 3rd annual conference on Research in information technology*, pages 33–38, 2014.

[3] Habib Ullah Baig and Farrukh Kamran. Detection of port and network scan using time independent feature set. In *2007 IEEE Intelligence and Security Informatics*, pages 180–184. IEEE, 2007.

[4] Tarun Yadav and Arvind Mallari Rao. Technical aspects of cyber kill chain. In *International symposium on security in computing and communication*, pages 438–452. Springer, 2015.

[5] LT Heberlein, GV Dias, KN Levitt, B Mukherjee, J Wood, and D Wolber. A network security monitor. In *Proceedings. 1990 IEEE Computer Society Symposium on Research in Security and Privacy*, pages 296–304. IEEE, 1990.

[6] Azriel Henry, Sunil Gautam, Samrat Khanna, Khaled Rabie, Thokozani Shongwe, Pronaya Bhattacharya, Bhisham Sharma, and Subrata Chowdhury. Composition of hybrid deep learning model and feature optimization for intrusion detection system. *Sensors*, 23(2):890, 2023.

[7] Jason M Pittman. Machine learning and port scans: A systematic review. *arXiv preprint arXiv:2301.13581*, 2023.

[8] Gordon Lyon. Remote os detection via tcp/ip stack fingerprinting, Oct 1998.

[9] Marco De Vivo, Eddy Carrasco, Germinal Isern, and Gabriela O De Vivo. A review of port scanning techniques. *ACM SIGCOMM Computer Communication Review*, 29(2):41–48, 1999.

[10] Richard J Barnett and Barry Irwin. Towards a taxonomy of network scanning techniques. In *Proceedings of the 2008 annual research conference of the South African Institute of Computer Scientists and Information Technologists on IT research in developing countries: riding the wave of technology*, pages 1–7, 2008.

[11] Elias Bou-Harb, Mourad Debbabi, and Chadi Assi. Cyber scanning: a comprehensive survey. *Ieee com-*

munications surveys & tutorials, 16(3):1496–1519, 2013.

[12] Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.

[13] Matthias Schonlau and Rosie Yuyan Zou. The random forest algorithm for statistical learning. *The Stata Journal*, 20(1):3–29, 2020.

[14] Baoxun Xu, Xiufeng Guo, Yunming Ye, and Jiefeng Cheng. An improved random forest classifier for text categorization. *J. Comput.*, 7(12):2913–2920, 2012.

[15] Luka Čehovin and Zoran Bosnić. Empirical evaluation of feature selection methods in classification. *Intelligent data analysis*, 14(3):265–281, 2010.

[16] Khaled Fawagreh, Mohamed Medhat Gaber, and Eyad Elyan. Random forests: from early developments to recent advancements. *Systems Science & Control Engineering: An Open Access Journal*, 2(1):602–609, 2014.

[17] Naresh Kumar. Advantages and disadvantages of random forest algorithm in machine learning. *The Professionals Point*, 2019.

[18] Akram QM Algaolahi, Abdullah A Hasan, Amer Sallam, Abdullah M Sharaf, Aseel A Abdu, and Anas A Alqadi. Port-scanning attack detection using supervised machine learning classifiers. In *2021 1st International Conference on Emerging Smart Technologies and Applications (eSmarTA)*, pages 1–5. IEEE, 2021.

[19] Gustavo De Carvalho Bertoli, Lourenço Alves Pereira Júnior, Osamu Saotome, Aldri L Dos Santos, Filipe Alves Neto Verri, Cesar Augusto Cavalheiro Marcondes, Sidnei Barbieri, Moises S Rodrigues, and José M Parente De Oliveira. An end-to-end framework for machine learning-based network intrusion detection system. *IEEE Access*, 9:106790–106805, 2021.

[20] Aswadati Sirisha, Kosaraju Chaitanya, KVSSR Krishna, and Satya Sandeep Kanumalli. Intrusion detection models using supervised and unsupervised algorithms-a comparative estimation. *IJSSE*, 11(1):51–58, 2021.

[21] Mahsa Mohseni and Jafar Tanha. A density-based undersampling approach to intrusion detection. In *2021 5th International Conference on Pattern Recognition and Image Analysis (IPRIA)*, pages 1–7. IEEE, 2021.

[22] Emmanuel Kwesi Baah, Steven Yirenkyi, Dominic Asamoah, Stephen Opoku Oppong, Edward Opoku-Mensah, Benjamin Tei Partey, Anthony Kingsley Sackey, Oliver Kornyo, and Evans Obu. Enhancing port scans attack detection using principal component analysis and machine learning algorithms. In *International Conference on Frontiers in Cyber Security*, pages 119–133. Springer, 2022.

[23] N SaiKira, Pradeep Naidu PDS, K Harshini, M Venkateswarlu, et al. Detection of cyber attacks in network using machine learning techniques. *South Asian Journal of Engineering and Technology*, 12(3):138–145, 2022.

[24] Hans E Plesser. Reproducibility vs. replicability: a brief history of a confused terminology. *Frontiers in neuroinformatics*, 11:76, 2018.

[25] R Murray Lindsay and Andrew SC Ehrenberg. The design of replicated studies. *The American Statistician*, 47(3):217–228, 1993.

[26] Andrew Brooks, John Daly, James Miller, Marc Roper, and Murray Wood. Replication of experimental results in software engineering. *International Software Engineering Research Network (ISERN) Technical Report ISERN-96-10, University of Strathclyde*, 2, 1996.

[27] Omar S Gómez, Natalia Juristo, and Sira Vegas. Replications types in experimental disciplines. In *Proceedings of the 2010 ACM-IEEE international symposium on empirical software engineering and measurement*, pages 1–10, 2010.

[28] Maxime Lanvin, Pierre-François Gimenez, Yufei Han, Frédéric Majorczyk, Ludovic Mé, and Eric Totel. Errors in the cicids2017 dataset and the significant differences in detection performances it makes. 2023.

[29] `https://github.com/jasonmpittman/rf-port-scans`.