© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723220

Machine Learning-Based Enterprise Security Information and Event Management Systems: A Systematic Literature Review

OMOIGHE JOSHUA AHUOSE[1], TEMITOPE OLUFUNMI ATOYEBI[2], RIDWAN KOLAPO[3], PREMA KIRUBAKARAN[4]

[1, 2, 3, 4] *Information Technology Department, Nile University of Nigeria, Abuja, Nigeria.*

Abstract- ***Large firms utilize Security Information and Event Management (SIEM) systems that allow them to gather, normalize, correlate, and analyze security events that, in turn, come from endpoints, networks, identity platforms, cloud services, databases, and applications. Even if traditional rule-based SIEM is indeed useful for regulatory compliance and has an understanding of the attack patterns, it has its limitations due to the existence of high event volume, heterogeneous logs, false positive, alert fatigue, and multi-stage attacks. That's why this article is such a great read! They have done a deep analysis of the various machine learning-based enterprise SIEM systems available. PRISMA 2020 guided the reporting of study selection, while an adapted Waterfall process organized requirements definition, protocol design, search, screening, quality appraisal, extraction, synthesis, and reporting. A careful analysis of thirty studies was conducted with the help of both descriptive and thematic synthesis. The notable findings in the study indicate that the area of research is mostly concerned with log anomaly detection and threat detection. The best techniques for dealing with data that has correct labels & decisions by an analyst are supervised and ensemble methods; on the other hand, the semi-supervised, self-supervised, deep, and transformer-based methods are appropriate for the applications of the large unlabelled log streams; graph-based methods remain as the best option for the events that happen together; and the explainable artificial intelligence is the one that enables trust in analysts. The attention to incidents, system response and defense mechanisms as well as privacy, model drift, adversarial robustness, and the actual security operations center are less supported by involved statistics. It is stated in the article that no one specific machine learning technique is appropriate for each SIEM task. The success of an enterprise deployment is based on implementing the appropriate techniques depending on the security function, data quality, label availability, explanation requirements, and analyst workflow.***

Index Terms- ***Anomaly detection, Enterprise cybersecurity, Machine learning, Security Information and Event Management, Systematic literature review***

I. INTRODUCTION

Cybersecurity in enterprises is built on the interdependence of various networks, endpoints, identity services, cloud platforms, databases, business applications, and security tools. All security zones are potential loggers of events that show policy violations, attacks, successful attacks, or actions taken to respond to those activities. SIEM systems aggregate these documents by gathering, sorting, standardizing, storing, correlating, and presenting for surveillance, analysis, regulation, and incident management [1], [2]. From my perspective, Traditional SIEM detection does rely primarily on the method of detection and signature files such as thresholds, manually written rules, and correlative logic. These mechanisms are still thanks to the active known indices and the clear violation of the policy but are prone to issues when the attackers utilize legitimate accounts, low-and-slow activity, changing techniques, or actions distributed among multiple systems. One of the major challenges is that the threat landscape has changed, and malware, data threats, and availability attacks can often result in multiple datasets showing weak or incomplete signals [3].

One of the most crucial outcomes resulting from this situation is the alert fatigue. The use of general rules is a double-edge sword as it results in a high volume of duplicate alerts with either low priority or false reading whereas the use of very narrow rules may not signal the true risks. Due to the situation, the analysts are faced with the dilemma of determining which alerts are correlating, which assets are most valuable, and which incidents need a quick response. Proposals have been made for the use of an AI-assisted SIEM, machine-learning data analysis, risk thresholds, and ensemble alert management to streamline prioritization and lessen the workload of the analysts [4]-[7].

Practices for teaching reading and writing at the primary level. Machine learning actually is the chapter, which is a part of the extended scanner and is equipped with a computer to learn and adapt to the surrounding. It can not only classify but also detect abnormal log sequences, cluster related alerts, estimate risk, and prioritize investigations and hence the support response can be given. However, the machine learning is not fully independent of rule, threat intelligence, or human judgment as it also needs them. Its utility is proportional to the representative data, specific, well-defined, and measurable security tasks, unambiguous interpretability, continuous surveillance, and integration with Security Operations Centre daily business.

The acquisition and storage of this content remain scattered among SIEM, intrusion detection, system-log anomaly detection, event correlation, explainable artificial intelligence, adversarial robustness, and broader enterprise cybersecurity analytics. The present research unites this fragmentation by providing a systematic literature review of machine learning-based enterprise SIEM systems. It reveals the technologies and architectures employed, the supported SIEM capabilities, the dataset and assessment methods that have been applied, and the reported effectiveness, limitations, deployment challenges, and future research needs.

II. LITERATURE REVIEW

Supervised, unsupervised, semi-supervised, self-supervised, deep, ensemble, graph-based, and explainable methods are all machine learning approaches that serve different aspects of SIEM. For example, decision trees, random forests, support vector machines, logistic regression, gradient boosting, and neural networks are supervised methods that can be applied only if you have the appropriate tags or historical decisions by the analyst followed by the labels [8], [9]. Prevalently, the main issue is the fact that labelled enterprise attack data are usually rare, confidential, incomplete, or old.

The lack of labels in the data directly makes unsupervised, semi-supervised, and self-supervised learning very vital in the data acquisition process. The methods mentioned above such as clustering, one-class classification, isolation methods, autoencoders, contrastive learning, and transformer-based models are able to recognize and differentiate normal traffic from abnormal signs. LogBERT, SSDLog, MDFULog, CLDTLog, TeleDAL, OneLog, LogFiT, and their related studies trace the transformation from manually engineered templates to sequence, semantic, parser-light, and end-to-end log representations. [10]-[20]. These approaches are promising, but abnormal behaviour is not always malicious, so false positives and analyst interpretation remain important.

Both the correlation of events and the multi-stage studies need the models that illustrate the connections among users, hosts, processes, files, network connections, and alerts. Linking the weak signals into an attack chain and providing the incipient reconstruction of an incident that isolated events are not able to make known is hierarchical correlation and provenance-graph methods, respectively [21], [22]. Nevertheless, concrete proofs of the above from real enterprises are still scarce.

A comprehensive intrusion-detection and cybersecurity studies do not only address issues such as algorithm selection, data preparation, benchmark datasets, and generalization but also provide additional evidence on these topics [23]-[27]. Explainable AI studies point out that the models operating in a SOC (security operation center) must give reasons that analysts can both comprehend and deduce. Hence, LIME, SHAP, black-box explainer evaluation, and explainable deep learning have been prioritizing alert triage and response support [28]-[31]. Distribution shift, adversarial manipulation, and model drift on the other hand might quite often become the reasons for performance reduction (demonstrating none of the algorithms in deployment) [32]-[34].

Some reviews are showing the research on cybersecurity analytics, intrusion detection, and log anomaly detection, but they lack the regular connection between data collection, preprocessing, machine-learning analysis, correlation, alerting, investigation, and response as one enterprise SIEM capability [35]. Therefore, system-level synthesis is necessary to match techniques with SIEM functions and operational conditions.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723220

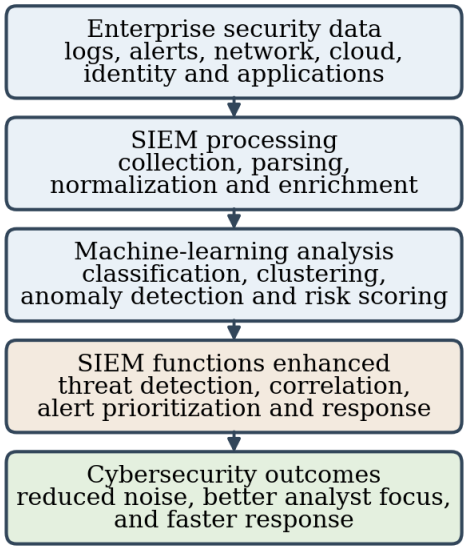

Figure 1: Conceptual framework for ML-enhanced SIEM

Fig. 1 is the overview of the systems used in the review. Enterprise logs, alerts, network records, cloud events, identity data, and application records, along with SIEM, pipeline feeds that are collected, parsed, normalized, enriched, and stored. Classification, clustering, anomaly detection, sequence analysis, or risk scoring are the different machine-learning models used. Their results provide the basis for threat detection, event correlation, alert prioritization, investigation, and response. The operational goals that need to be achieved are fewer irrelevant alerts, better analyst focus, quicker investigation, and enhanced response. The framework states that model accuracy is just one side of the value for enterprises; the appropriate utilization also depends on dependable data engineering, clear outputs, workflow integration, governance, and the constant monitoring of processes. This model goes on to validate an analyst to the detection process and also provides security feedback loop. The confirmed and dismissed alerts can be used to reevaluate the labels, decision on the thresholds, and retraining the models. The rules and signatures will be the same for the already known indicators, whereas the machine learning will help with the unknown or different ones. So, it becomes clear that addressability, the right to privacy, the watch of model-drift, and the human consent, which are the primary factors, control the analytical pipeline.

III. METHODOLOGY

The research implemented a systematic literature review design. PRISMA 2020 helped with unbiased presentation of the research processes, including identification, screening, eligibility assessment, and final inclusion [36]. The widely recognized research for systematic reviews in software engineering and information systems was the basis of the protocol, quality appraisal, extraction, and synthesis procedures [37]-[39]. A customized Waterfall sequence arranged the work into requirements definition, protocol and search design, literature search, screening and verification, data extraction, synthesis, and reporting [40].

Searched databases included IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, Scopus, Web of Science, Google Scholar, MDPI, Taylor & Francis Online, and Wiley Online Library.

We explored search strings that unite SIEM terms, machine-learning terms, and cybersecurity terms with the help of Boolean operators. Examples of how to do this are: ("Security Information and Event Management" OR SIEM) AND ("machine learning" OR "artificial intelligence" OR "deep learning") AND cybersecurity, as well as looking for log anomaly detection, event correlation, alert prioritization, false positives, and incident response.

The articles that were eligible for the research were those written in English, published within the 2021-2026 review time frame, and had a clear machine-learning or data-driven security element, included adequate methodology, and discussed enterprise SIEM or operational security monitoring. Blogs, marketing materials, opinion pieces, unrelated domains, and studies without adequate methods were excluded. The final synthesis contained 30 studies; the included academic studies were published from 2021 to 2025, with 2026 materials being used for contextual background.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723220

Table 1: Objective vs Methodology Matching

Objectives	Proposed Methodology	Expected Results
Objective (1)	Extract and classify ML techniques and architectures.	Taxonomy of methods and system designs.
Objective (2)	Map evidence to SIEM security functions.	Function-technique relationship.
Objective (3)	Code data sources, preprocessing, datasets, and metrics.	Data and evaluation profile.
Objective (4)	Synthesize effectiveness, limitations, and research gaps.	Evidence-based recommendations.

The application of the same extraction fields in all the studies included bibliographic information, SIEM function, machine-learning approach, model, data, preprocessing, evaluation, findings, limitations, and research gaps. Descriptive synthesis provided a summary of the information about the year of publication, focus area, technique, function, data type, and metric while thematic synthesis presented a comparison of operational issues.

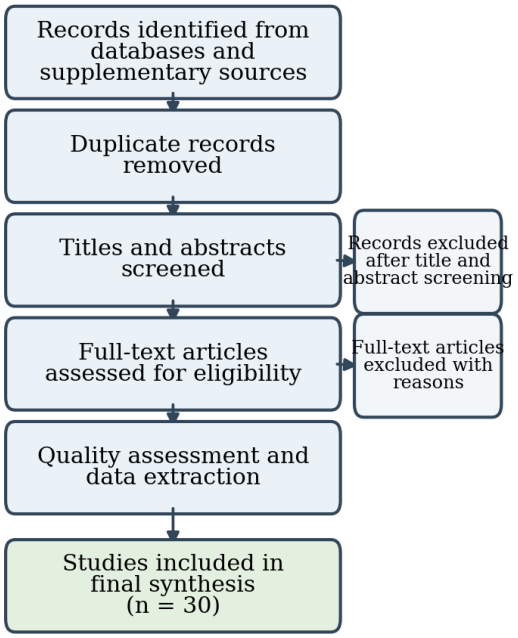

Figure 2: PRISMA-guided study-selection workflow

The passage is a brief outline of the entire work process that describes the whole cycle from detecting distinct and duplicate files to title and abstract screening, full-text eligibility assessment, quality appraisal, and finally the acceptance stage. The project results in the final summary of 30 studies. Quality evaluation also included the factors of relevance, transparency of the machine-learning approach, data description, preprocessing, and features, evaluation metrics, practical SIEM relevance, limitations, reproducibility, and contribution.

IV. RESULTS AND DISCUSSIONS

The 30 studies included for the analysis were mostly from the recent period. Out of these, four articles were published in 2021, while two were published in 2022, eight in 2023, and thirteen in 2024, with the remaining three in 2025. The highest number of articles published in 2023 and 2024 is the evidence of the researchers becoming increasingly interested in the issues like deep learning, transformer models, semantic log analysis, explainability, adversarial robustness, and operational alert management.

The findings were inconsistent for the various SIEM functions. The most favorable area was log anomaly detection and log analysis with 11 studies. SIEM alert management and prioritization were associated with three studies; intrusion detection and cybersecurity monitoring with four; explainable AI and analyst trust also with four; adversarial robustness and model security with three; event correlation and threat-informed detection with two; and enterprise cybersecurity analytics and review methods were also related with three. This distribution shows that current research is stronger at detecting suspicious activity than at proving end-to-end operational improvement in a live SOC.

The technique that is suitable should be chosen based on the specific task, the type of data, the quality of the label, the requirements for explanation, and the deployment environment. The most correct methods to use were the supervised and ensemble methods for known threat classification, and the labelled alert triage. The semi-supervised, self-supervised, and deep sequence models were the right choice for the cases where confirmed attack labels were infrequent. Character-level and parser-light language models were helpful for the logs that are unstable or semantically

rich, while the graph and provenance approaches were really good for correlation and attack reconstruction.

Table 2: Suitability of Machine Learning Approaches for Enterprise SIEM

ML Approach	Best-Suited SIEM Task	Main Limitation
Supervised and ensemble	Known threats and labelled alert triage	Label scarcity and drift
Unsupervised, semi- and self-supervised	Unknown behaviour and unlabelled log anomalies	False positives and baseline quality
Deep, transformer, and language models	Log sequence and semantic analysis	Compute, governance, and explainability
Graph and provenance learning	Multi-stage event correlation and investigation	Limited SIEM-scale validation
Explainable AI	Analyst trust, auditability, and response support	Explanation fidelity and usability

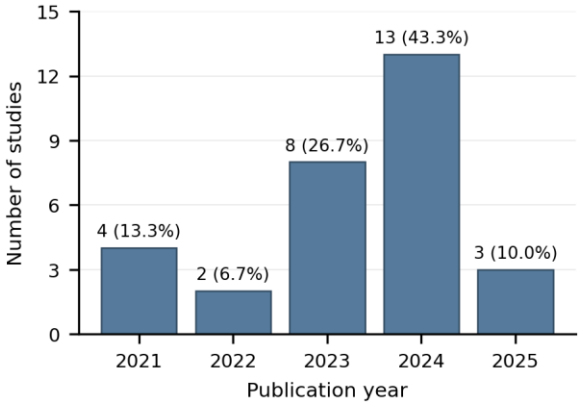

Figure 3: Distribution of selected studies by publication year

The data highlighted threat detection and anomaly detection with the highest degree of backing. Classification models used as a tool for detecting known-attack, where as unsupervised and sequence-based were used to identify the unusual log messages, user activities, authentication patterns and system behavior. A moderate level of evidence was provided concerning the prioritization of alerts and reduction of false positives: risk scoring, dynamic thresholds, and ensemble screening would be able to classify cases according to context and probable significance, but their worth is contingent on the trustworthy labels and workflow joining together.

Event correlation, incident reconstruction, and response support were still in their initial stages of development. Connection of events among users, hosts, processes, files, and network traffic is possible with the help of graph-based and hierarchical methods; however, limited number of studies evaluated analyst workload, investigation time, or response outcomes in the field of operational SIEM environments. Thus, explainability was not just a feature but a requirement that had to be fulfilled. A must for the analysts was comprehending the alert rise, evidences that influenced the score, and the following action required.

The research works exploited system and application logs, SIEM alerts, network traffic, intrusion-detection datasets, endpoint events, authentication records, threat intelligence, and cloud audit data. Common preprocessing techniques included parsing, tokenization, template extraction, normalization, sequence windows, embeddings, semantic representation, feature selection, and class balancing. Accuracy, precision, recall, F1-score, false positive rate, area under the curve, and detection time were the main technical metrics. Operational measures such as alert reduction, triage time, explanation usefulness, analyst acceptance, scalability, maintenance effort, and response speed were reported less consistently.
Admitted to obstacles were the main ones which primarily came from issues regarding data realism, privacy, explainability, model drift, adversarial manipulation, computational cost, and limited live SOC validation. A considerable number of tests used either public datasets or synthetic datasets as the primary source of data. These imitated conditions do not reflect the environment of enterprise logs that are attacked by variances of normality, orphan event types, data retention and privacy issues, as well as continually evolving conditions. The proof thus lays far beyond the idea of a hybrid model, where rules and signatures tackle the problems identified by the user and explicit policies, while machine learning is employed for anomaly discovery, risk ranking, semantic analysis, and correlation.

V. CONCLUSION

The 30 studies that were part of this literature review were the basis for the investigation of the software of machine learning-based Enterprise Security Information and Event Management Systems. It is revealed by the research results that machine learning can facilitate the tasks such as identifying the threats, detecting anomalies, log analysis, alert prioritization, event correlation, and incident-response support. In this regard, the most reliable findings involve the tasks of detecting and log anomaly analysis. Nevertheless, the ones on correlation, prioritization, response, and live deployment in the Security Operations Center are in need of additional operational validation.

Not all machine learning techniques can be applied for any SIEM task effectively. The supervised and ensemble techniques are generally best where validated labels and analyst decisions are available, while semi-supervised and self-supervised techniques are effective where there is limited labeling; Furthermore, deep learning and NLP models facilitate the processing of complex log streams; The combined graph and provenance models can provide the needed help in the steps investigation; the use of Explainable AI technology makes it possible for the analyst to develop their trust in it. The successful introduction of a new system is reliant on the provision of solid log-management skills, adequate representation of data, definite expression of operational goals, fair assessment, continuous monitoring and human intervention.

The future studies should emphasize on practical enterprise validation, achieving lower false-positive rates, quantifiable alert reduction, expediting triage, workflow-aware explanations, privacy-preserving learning, drift management, adversarial robustness, and coupling with analyst and incident-response processes. Instead of replacing them, machine learning should be the additional tool for SIEM rules, threat intelligence, cybersecurity governance, and professional judgment.

REFERENCES

[1] G. Gonzalez-Granadillo, S. Gonzalez-Zarzosa, and R. Diaz, "Security Information and Event Management (SIEM): Analysis, trends, and usage in critical infrastructures," *Sensors*, vol. 21, no. 14, p. 4759, 2021, doi: 10.3390/s21144759. MDPI

[2] K. A. Scarfone and M. P. Souppaya, *Cybersecurity Log Management Planning Guide*, NIST Special Publication 800-92r1 Initial Public Draft, National Institute of Standards and Technology, 2023, doi: 10.6028/NIST.SP.800-92r1.ipd. NIST

[3] European Union Agency for Cybersecurity, *ENISA Threat Landscape 2024*, 2024. ENISA

[4] T. Ban, T. Takahashi, S. Ndichu, and D. Inoue, "Breaking alert fatigue: AI-assisted SIEM framework for effective incident response," *Applied Sciences*, vol. 13, no. 11, p. 6610, 2023, doi: 10.3390/app13116610. MDPI

[5] N. Tendikov et al., "Security Information Event Management data acquisition and analysis methods with machine learning principles," *Results in Engineering*, vol. 22, p. 102254, 2024, doi: 10.1016/j.rineng.2024.102254. ScienceDirect

[6] A. Kapera and M. Niemiec, "Dynamic risk thresholds for SIEM alerting based on machine learning," *IEEE Access*, 2025, doi: 10.1109/ACCESS.2025.3588441. IEEE

[7] M. Khayat, E. Barka, M. A. Serhani, and F. Sallabi, "Advanced techniques for alert management in Security Information and Event Management systems with ensembled deep learning, hybrid optimization, and multi-feature extraction," *IEEE Open Journal of the Communications Society*, 2025, doi: 10.1109/OJCOMS.2025.3603000. IEEE

[8] H. Sarker, "Machine learning: Algorithms, real-world applications and research directions," *SN Computer Science*, vol. 2, p. 160, 2021, doi: 10.1007/s42979-021-00592-x. Springer

[9] X. Zhou, W. Liang, S. Shimizu, J. Ma, and Q. Jin, "Machine learning in cybersecurity: A survey," *ACM Computing Surveys*, 2022, doi: 10.1145/3495232.

[10] H. Guo, S. Yuan, and X. Wu, "LogBERT: Log anomaly detection via BERT," in *Proc. International Joint Conference on Neural Networks*, 2021, doi: 10.1109/IJCNN52387.2021.9534113. IEEE

[11] M. Landauer, S. Onder, F. Skopik, and M. Wurzenberger, "Deep learning for anomaly detection in log data: A survey," *Machine*

Learning with Applications, vol. 12, p. 100470, 2023, doi: 10.1016/j.mlwa.2023.100470. ScienceDirect

[12] S. Lu et al., "SSDLog: A semi-supervised dual branch model for log anomaly detection," *World Wide Web*, vol. 26, pp. 3137-3153, 2023, doi: 10.1007/s11280-023-01174-y. Springer

[13] M. Li, M. Sun, G. Li, D. Han, and M. Zhou, "MDFULog: Multi-feature deep fusion of unstable log anomaly detection model," *Applied Sciences*, vol. 13, no. 4, p. 2237, 2023, doi: 10.3390/app13042237. MDPI

[14] G. Tian, N. Luktarhan, H. Wu, and Z. Shi, "CLDTLog: System log anomaly detection method based on contrastive learning and dual objective tasks," *Sensors*, vol. 23, no. 11, p. 5042, 2023, doi: 10.3390/s23115042. MDPI

[15] G. Horvath, A. Meszaros, and P. Szilagyi, "TeleDAL: A regression-based template-less unsupervised method for finding anomalies in log sequences," *The Journal of Supercomputing*, vol. 79, pp. 18394-18416, 2023, doi: 10.1007/s11227-023-05379-w. Springer

[16] S. Hashemi and M. Mantyla, "OneLog: Towards end-to-end software log anomaly detection," *Automated Software Engineering*, vol. 31, p. 37, 2024, doi: 10.1007/s10515-024-00428-x. Springer

[17] S. Lupton, H. Washizaki, N. Yoshioka, and Y. Fukazawa, "Landscape and taxonomy of online parser-supported log anomaly detection methods," *IEEE Access*, 2024, doi: 10.1109/ACCESS.2024.3387287. IEEE

[18] C. Almodovar, F. Sabrina, S. Karimi, and S. A. Azad, "LogFiT: Log anomaly detection using fine-tuned language models," *IEEE Transactions on Network and Service Management*, 2024, doi: 10.1109/TNSM.2024.3358730. IEEE

[19] Z. Xu, Z. Wang, J. Xu, H. Shi, and H. Zhao, "Enhancing log anomaly detection with semantic embedding and integrated neural network innovations," *Computers, Materials & Continua*, vol. 80, no. 3, pp. 3991-4015, 2024, doi: 10.32604/cmc.2024.051620. Tech Science Press

[20] S. A. Mondal, P. Rv, S. Rao, and A. Menon, "LADDERS: Log based anomaly detection and diagnosis for enterprise systems," *Annals of Data Science*, vol. 11, pp. 1165-1183, 2024, doi: 10.1007/s40745-023-00471-7. Springer

[21] H. Maosa, K. Ouazzane, and M. C. Ghanem, "A hierarchical security event correlation model for real-time threat detection and response," *Network*, vol. 4, no. 1, pp. 68-90, 2024, doi: 10.3390/network4010004. MDPI

[22] Z. Cheng et al., "KAIROS: Practical intrusion detection and investigation using whole-system provenance," *arXiv:2308.05034*, 2023. arXiv

[23] S. Ali, C. Boufaied, D. Bianculli, P. Branco, and L. Briand, "A comprehensive study of machine learning techniques for log-based anomaly detection," *arXiv:2307.16714*, 2023. arXiv

[24] Z. A. Khan, D. Shin, D. Bianculli, and L. Briand, "Impact of log parsing on deep learning-based anomaly detection," *arXiv:2305.15897*, 2023. arXiv

[25] A. H. Ali et al., "Unveiling machine learning strategies and considerations in intrusion detection systems: A comprehensive survey," *Frontiers in Computer Science*, vol. 6, p. 1387354, 2024, doi: 10.3389/fcomp.2024.1387354. Frontiers

[26] Pinto, L. C. Herrera, Y. Donoso, and J. A. Gutierrez, "Survey on intrusion detection systems based on machine learning techniques for the protection of critical infrastructure," *Sensors*, vol. 23, no. 5, p. 2415, 2023, doi: 10.3390/s23052415. MDPI

[27] E. Gyamfi and A. Jurcut, "Intrusion detection in Internet of Things systems: A review on design approaches leveraging multi-access edge computing, machine learning, and datasets," *Sensors*, vol. 22, no. 10, p. 3744, 2022, doi: 10.3390/s22103744. MDPI

[28] M. Pawlicki, A. Pawlicka, R. Kozik, and M. Choras, "The survey on the dual nature of XAI challenges in intrusion detection and their potential for AI innovation," *Artificial Intelligence Review*, vol. 57, p. 330, 2024, doi: 10.1007/s10462-024-10972-3. Springer

[29] O. Arreche, T. R. Guntur, J. W. Roberts, and M. Abdallah, "E-XAI: Evaluating black-box explainable AI frameworks for network intrusion detection," *IEEE Access*, 2024, doi: 10.1109/ACCESS.2024.3365140. IEEE

[30] D. Gaspar, P. Silva, and C. Silva, "Explainable AI for intrusion detection systems: LIME and SHAP applicability on multi-layer perceptron,"

IEEE Access, 2024, doi: 10.1109/ACCESS.2024.3368377. IEEE

[31] N. H. A. Mutalib et al., "Explainable deep learning approach for advanced persistent threats detection in cybersecurity: A review," *Artificial Intelligence Review*, vol. 57, p. 297, 2024, doi: 10.1007/s10462-024-10890-4. Springer

[32] M. Wang, N. Yang, D. H. Gunasinghe, and N. Weng, "On the robustness of ML-based network intrusion detection systems: An adversarial and distribution shift perspective," *Computers*, vol. 12, no. 10, p. 209, 2023, doi: 10.3390/computers12100209. MDPI

[33] S. Sharma and Z. Chen, "A systematic study of adversarial attacks against network intrusion detection systems," *Electronics*, vol. 13, no. 24, p. 5030, 2024, doi: 10.3390/electronics13245030. MDPI

[34] A. Vassilev, A. Oprea, A. Fordyce, and H. Andersen, *Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations*, National Institute of Standards and Technology, 2024, doi: 10.6028/NIST.AI.100-2e2023. NIST

[35] T. D. Le, T. Le-Dinh, and S. Uwizeyemungu, "Cybersecurity analytics for the enterprise environment: A systematic literature review," *Electronics*, vol. 14, no. 11, p. 2252, 2025, doi: 10.3390/electronics14112252. MDPI

[36] M. J. Page et al., "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," *BMJ*, vol. 372, p. n71, 2021, doi: 10.1136/bmj.n71. BMJ

[37] Kitchenham and S. Charters, *Guidelines for Performing Systematic Literature Reviews in Software Engineering*, EBSE Technical Report EBSE-2007-01, Keele University and Durham University, 2007. EBSE

[38] Okoli, "A guide to conducting a standalone systematic literature review," *Communications of the Association for Information Systems*, vol. 37, pp. 879-910, 2015, doi: 10.17705/1CAIS.03743. CAIS

[39] H. Snyder, "Literature review as a research methodology: An overview and guidelines," *Journal of Business Research*, vol. 104, pp. 333-339, 2019, doi: 10.1016/j.jbusres.2019.07.039. ScienceDirect

[40] W. W. Royce, "Managing the development of large software systems," in *Proc. IEEE WESCON*, vol. 26, pp. 1-9, 1970.