

Artificial Intelligence-Based Detection of Abnormal Network Activities Using Behavioral Data Analysis

Dr. Javier Martínez

Professor, Barcelona Institute of Network Security, Spain

Abstract

The increasing dependence on interconnected digital systems has made network security monitoring an important area of artificial intelligence research. Modern network environments generate substantial volumes of behavioral information reflecting communication patterns, traffic frequency, connection duration, resource utilization, and interactions among network entities. Conventional intrusion detection approaches often depend on previously identified attack patterns, which may restrict their ability to recognize unfamiliar or evolving network activities. Artificial intelligence-based behavioral analysis offers an alternative by examining statistical and contextual differences between established network behavior and potentially abnormal observations. This hypothetical methodological study investigates the potential application of artificial intelligence techniques for identifying abnormal network activities through behavioral data analysis. The proposed framework combines behavioral feature extraction, machine learning-based anomaly scoring, risk classification, and performance evaluation. A hypothetical dataset containing 1,000 network-flow observations is used to illustrate how different alert thresholds influence detection performance. Three assumed detection configurations are examined through confusion-matrix calculations, allowing comparison of precision, recall, F1 score, and false-positive rates.

The illustrative analysis demonstrates that increasing detection sensitivity may improve the identification of abnormal observations while simultaneously increasing the number of legitimate activities incorrectly classified as suspicious. Conversely, conservative alert configurations may reduce unnecessary warnings but fail to identify some abnormal events. These relationships emphasize the importance of balancing detection sensitivity and operational reliability when designing artificial intelligence-based network monitoring systems. The study contributes a structured methodological perspective on behavioral anomaly detection and identifies major challenges involving data quality, behavioral diversity, class imbalance, model adaptability, and evaluation reliability. It concludes that artificial intelligence-based behavioral analysis provides a promising foundation for network security monitoring, although practical effectiveness requires reproducible experimentation using representative datasets and appropriately validated predictive models.

Keywords: Artificial Intelligence, Network Anomaly Detection, Behavioral Data Analysis, Machine Learning, Intrusion Detection, Network Security, Predictive Analytics, Anomaly Classification, Traffic Monitoring, Cybersecurity.

1. Introduction

The rapid development of digital communication technologies has increased the complexity of modern computer networks and the volume of information exchanged between connected systems. Organizations increasingly depend on network infrastructures to support communication, information processing, cloud services, and distributed applications. These environments generate extensive records of network behavior that may provide useful evidence of abnormal activities. However, identifying meaningful irregularities within large volumes of routine traffic presents a significant analytical challenge. Artificial intelligence offers computational techniques capable of examining complex behavioral relationships and supporting automated network security monitoring.

Traditional network intrusion detection methods frequently rely on predefined signatures, fixed rules, or manually specified indicators of suspicious behavior. Although these approaches can identify activities matching recognized patterns, they may be less effective when unfamiliar conditions or previously unobserved behavioral deviations occur. García-Teodoro et al. (2009) examined the potential and limitations of anomaly-based intrusion detection, emphasizing the importance of flexible analytical methods. Behavioral data analysis provides an alternative perspective by studying how network entities normally communicate and identifying observations that depart from established operating characteristics.

Artificial intelligence-based detection techniques employ statistical learning, classification algorithms, clustering procedures, and anomaly-scoring methods to examine network behavior. These approaches may consider connection duration, traffic volume, packet frequency, communication regularity, and variations in interactions between network entities. Machine learning algorithms can use such characteristics to distinguish common behavioral patterns from observations requiring further examination. Chandola et al. (2009) established important theoretical foundations for anomaly detection, while Buczak and Guven (2016) examined the application of machine learning methods to intrusion detection research.

Behavioral anomaly detection nevertheless presents important methodological difficulties. Legitimate network activities may change because of software updates, operational schedules, new applications, or varying communication requirements. Such changes can resemble abnormal behavior even when no security incident has occurred. Furthermore, actual anomalous events may represent only a small proportion of total observations, creating difficulties for classification and performance evaluation. Sommer and Paxson (2010) highlighted practical limitations in applying machine learning to network intrusion detection, including the difference between experimental assumptions and operational network conditions.

The present study examines artificial intelligence-based detection through a hypothetical methodological framework centered on behavioral network information. It proposes a structured process for identifying behavioral characteristics, assigning anomaly scores, classifying potentially abnormal observations, and evaluating detection outcomes. A simplified analytical scenario is introduced to demonstrate the relationship between alert sensitivity and classification reliability. The study does not claim to have implemented a working intrusion detection system. Instead, it provides a foundation for understanding methodological requirements and designing future experiments involving behavioral network anomaly detection.

2. Objectives of the Study

The primary objective of this study is to examine the methodological potential of artificial intelligence-based behavioral analysis for identifying abnormal network activities. The investigation focuses on behavioral feature interpretation, anomaly classification, and detection performance. It also evaluates methodological limitations associated with false alerts, data diversity, and changing network conditions.

2.1 Examination of Behavioral Network Characteristics

The first objective is to examine behavioral characteristics that may support the identification of abnormal network activities. These characteristics include connection frequency, communication duration, packet distribution, traffic volume, and changes in interaction patterns. The study considers how such information can provide a statistical representation of routine network behavior. Identifying relevant behavioral variables is important because the quality of anomaly detection depends substantially on the information available to the learning algorithm and the extent to which selected features represent meaningful network activity.

The investigation also seeks to distinguish unusual behavior from confirmed security incidents. An observation may differ from established communication patterns without representing a malicious activity. Consequently, behavioral interpretation requires an understanding of network context, normal operational variation, and the limitations of automated classification. This objective emphasizes the importance of contextual feature selection and appropriate preprocessing. It establishes a conceptual basis for developing detection procedures that identify potentially suspicious activity while reducing unnecessary conclusions about legitimate network operations.

2.2 Assessment of Artificial Intelligence-Based Detection

The second objective is to examine how artificial intelligence techniques may support automated classification of potentially abnormal network observations. The study considers anomaly scoring, supervised classification, and threshold-based decision procedures as complementary components of a detection framework. A hypothetical analytical scenario is used to explore how different alert thresholds influence detection outcomes. The purpose is to clarify the relationship between sensitivity and classification precision without attributing illustrative calculations to a trained or experimentally tested artificial intelligence model.

Another aspect of this objective concerns the evaluation of detection reliability through quantitative performance measures. Precision, recall, F1 score, and false-positive rate provide different perspectives on classification behavior. The study examines how these measures respond to alternative hypothetical detection configurations when abnormal observations are relatively uncommon. Understanding their relationships is important because high detection sensitivity does not necessarily indicate satisfactory operational performance. The objective therefore emphasizes balanced evaluation rather than relying exclusively on overall classification accuracy.

2.3 Identification of Network Security Challenges

The third objective is to identify methodological challenges affecting artificial intelligence-based network anomaly detection. Particular attention is directed toward heterogeneous traffic, incomplete behavioral information, imbalanced datasets, and changes in legitimate network activity. These conditions may influence the reliability of anomaly scores and generate incorrect alerts. The investigation also considers the difficulty of distinguishing genuine behavioral irregularities from

temporary fluctuations. Identifying these challenges contributes to the development of more reliable and context-sensitive network monitoring procedures.

The study additionally examines the importance of reproducible evaluation and transparent reporting. Network anomaly detection systems should be assessed using clearly specified datasets, classification assumptions, and evaluation measures. Differences between hypothetical calculations and experimentally observed outcomes must remain explicit. This objective supports responsible interpretation of analytical findings and identifies the evidence required before an artificial intelligence-based detection framework can be considered operationally effective. It also establishes priorities for future research involving real network traffic and validated security monitoring applications.

3. Review of Literature

3.1 Development of Artificial Intelligence-Based Intrusion Detection

Lee and Stolfo (2000) developed a framework for constructing intrusion detection features and models from audit information. Their work demonstrated the importance of systematic feature extraction in learning-based security monitoring. Lazarevic et al. (2003) subsequently compared anomaly detection approaches for network intrusion detection, examining alternative methods for identifying unusual network observations. These contributions established that machine learning could support security analysis by identifying relationships within behavioral information rather than depending exclusively on manually defined detection rules.

Patcha and Park (2007) reviewed anomaly detection methods and discussed the development of network intrusion detection technologies. Tsai et al. (2009) examined machine learning approaches to intrusion detection, including individual and hybrid classification methods. Their work highlighted the importance of selecting appropriate learning techniques according to dataset characteristics and detection requirements. These studies collectively demonstrated that artificial intelligence provides multiple methodological approaches to network monitoring, although model performance remains sensitive to training information, algorithm selection, and evaluation conditions.

3.2 Behavioral Anomaly Identification and Classification

Breunig et al. (2000) introduced the local outlier factor method, providing a density-based approach for identifying observations that differ from neighboring data patterns. Schölkopf et al. (2001) developed an approach for estimating the support of high-dimensional distributions, contributing to one-class learning. Liu et al. (2008) introduced isolation forests as an alternative method for identifying unusual observations. These techniques provide relevant foundations for behavioral network analysis because they can represent abnormality without requiring every possible irregular activity to be specified in advance.

Chandola et al. (2009) reviewed anomaly detection methods across multiple application domains and established a useful classification of analytical approaches. García-Teodoro et al. (2009) examined anomaly-based network intrusion detection techniques and discussed their practical limitations. Ahmed et al. (2016) subsequently reviewed network anomaly detection methods and related evaluation challenges. Together, these publications demonstrate that behavioral anomaly detection depends on suitable feature representations, reliable baseline behavior, and careful interpretation of observations that deviate from expected patterns.

3.3 Network Datasets and Evaluation Reliability

Axelsson (2000) demonstrated the importance of false alarms when security incidents occur infrequently. Tavallaei et al. (2009) examined weaknesses in the KDD CUP 99 dataset and highlighted limitations affecting intrusion detection evaluation. Sommer and Paxson (2010) further questioned assumptions that may separate experimental machine learning performance from operational network environments. These contributions indicate that high classification accuracy on a particular benchmark does not necessarily establish reliable security detection under changing real-world conditions.

Liao et al. (2013) reviewed intrusion detection systems and their technological characteristics, while Moustafa and Slay (2015) introduced the UNSW-NB15 dataset for network security research. Buczak and Guven (2016) examined data mining and machine learning applications in intrusion detection. Saito and Rehmsmeier (2015) also emphasized the value of precision–recall measures for imbalanced classification. These studies support the use of representative datasets, transparent analytical procedures, and multiple performance measures when evaluating artificial intelligence-based detection systems.

4. Research Methodology

4.1 Research Design and Detection Framework

The present investigation adopts a hypothetical methodological research design supported by illustrative classification calculations. The proposed framework examines how behavioral network information may be transformed into features suitable for anomaly detection. It includes data preparation, behavioral feature extraction, model-based anomaly scoring, threshold selection, and performance assessment. The research does not claim that a machine learning detector has been trained or deployed. Instead, it develops an analytical example to illustrate methodological relationships relevant to artificial intelligence-based network monitoring and classification.

The detection framework assumes that network-flow observations can be represented using measurable behavioral characteristics. Potential variables include connection duration, transmitted and received data volumes, packet counts, communication frequency, and changes in interaction patterns. A learning algorithm could use these variables to estimate behavioral deviation from established network conditions. Depending on the available labels, supervised classifiers or unsupervised anomaly detection techniques may be appropriate. The present study does not select a verified optimal algorithm because experimental comparisons have not been performed.

4.2 Hypothetical Dataset and Classification Assumptions

The illustrative analytical dataset consists of 1,000 hypothetical network-flow observations. Of these observations, 950 are designated as normal and 50 as anomalous, representing an assumed abnormal-event prevalence of 5%. These quantities are selected exclusively for mathematical demonstration and do not represent collected network traffic. Three hypothetical alert configurations are considered: conservative, balanced, and sensitive. Each configuration is assigned an internally consistent set of classification outcomes to demonstrate the relationship between detection sensitivity and false-positive alerts.

The study assumes that the hypothetical ground-truth labels are correct and that all three configurations classify the same underlying set of 1,000 observations. A true positive represents an anomalous observation correctly identified, while a false positive represents normal activity incorrectly

flagged. A false negative indicates an undetected anomaly, and a true negative represents correctly classified normal behavior. The assumed confusion-matrix values are not produced by executing a detection model or adjusting actual learned thresholds.

4.3 Evaluation Procedure and Statistical Measures

The analytical procedure compares the three hypothetical configurations using their assumed true-positive, false-positive, false-negative, and true-negative counts. Performance measures are calculated directly from these values to demonstrate the consequences of different alerting decisions. Precision measures the proportion of alerts corresponding to actual anomalies, while recall represents the proportion of anomalous observations correctly identified. The F1 score combines precision and recall, whereas the false-positive rate describes the proportion of normal observations incorrectly flagged.

The principal evaluation formulas are:

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP})$$

$$\text{Recall} = \text{TP} / (\text{TP} + \text{FN})$$

$$\text{F1 Score} = 2\text{TP} / (2\text{TP} + \text{FP} + \text{FN})$$

$$\text{False-Positive Rate} = \text{FP} / (\text{FP} + \text{TN})$$

These measures provide complementary information about the hypothetical classification outcomes. Since the calculations are deterministic consequences of assumed counts, statistical significance tests and confidence intervals are not reported as experimental evidence. Future empirical validation would require actual model predictions, representative datasets, temporally appropriate evaluation, and repeated or independently replicated experiments.

5. Analysis

5.1 Hypothetical Classification Outcomes

The first analytical comparison considers classification outcomes for three hypothetical alert configurations. Each configuration evaluates the same assumed dataset containing 50 anomalous and 950 normal observations. The conservative configuration identifies fewer anomalies but produces relatively few false alerts. The sensitive configuration identifies more anomalies while also incorrectly flagging a larger number of normal observations. The balanced configuration represents an intermediate decision setting. These counts are illustrative assumptions rather than results generated by an artificial intelligence detection algorithm.

The hypothetical confusion matrices demonstrate that higher detection sensitivity can be accompanied by a larger false-alert burden. Under the conservative configuration, 30 anomalous observations are correctly identified and 10 normal observations are incorrectly flagged. The sensitive configuration identifies 47 anomalies but generates 95 false alerts. These assumed outcomes provide a mathematical basis for comparing detection effectiveness and alert quality. Their purpose is to illustrate evaluation principles, not to establish the superiority of a particular detection method.

Table 1: Hypothetical Network Anomaly Classification Outcomes

Classification Configuration	True Positives	False Positives	False Negatives	True Negatives
Conservative	30	10	20	940
Balanced	40	40	10	910
Sensitive	47	95	3	855

Note: Each configuration contains 1,000 observations, including 50 assumed anomalies and 950 normal records. The counts are illustrative and have not been obtained from a trained model.

5.2 Comparative Detection Performance

The second analytical comparison derives performance measures from the hypothetical confusion matrices. The conservative configuration produces precision of 75% and recall of 60%, indicating that its alerts are comparatively selective but some anomalies remain unidentified. The balanced configuration increases recall to 80% while reducing precision to 50%. The sensitive configuration reaches recall of 94%, although its precision decreases because the number of false-positive alerts becomes substantially larger relative to correctly identified anomalies.

These calculations illustrate why detection sensitivity should not be evaluated independently of alert reliability. A configuration identifying a larger proportion of anomalies may also generate additional warnings requiring investigation. Conversely, a configuration producing fewer false alerts may fail to identify a meaningful proportion of unusual observations. The preferred operating point therefore depends on the consequences of missed events, the cost of reviewing alerts, and the practical requirements of the monitored network environment.

Table 2: Hypothetical Evaluation Metrics for Network Anomaly Detection

Configuration	Precision (%)	Recall (%)	F1 Score (%)	False-Positive Rate (%)
Conservative	75.00	60.00	66.67	1.05
Balanced	50.00	80.00	61.54	4.21
Sensitive	33.10	94.00	48.96	10.00

Note: Percentages are rounded to two decimal places. These values are illustrative mathematical results, not empirical detection-system performance.

6. Discussion

6.1 Behavioral Analysis and Detection Effectiveness

The hypothetical analysis indicates that behavioral information can provide a methodological foundation for identifying observations that differ from expected network activity. Artificial intelligence-based techniques may represent such differences through statistical relationships, anomaly scores, or learned

classification boundaries. However, identifying unusual behavior is not equivalent to confirming a security incident. The interpretation of detection outcomes requires contextual information, reliable behavioral baselines, and appropriate procedures for evaluating whether an alert represents genuine security concern.

The proposed framework is consistent with anomaly detection principles discussed by Chandola et al. (2009) and García-Teodoro et al. (2009). Their research emphasizes the importance of identifying deviations while acknowledging limitations associated with false alarms and changing normal behavior. The present illustrative calculations extend this methodological discussion by demonstrating how alternative classification decisions influence precision and recall. Nevertheless, the hypothetical outcomes cannot establish the effectiveness of a specific artificial intelligence algorithm without actual training and testing.

6.2 Detection Sensitivity and False Alerts

The relationship between detection sensitivity and false alerts represents a central issue in behavioral network monitoring. The hypothetical results demonstrate that a sensitive configuration can identify a greater proportion of abnormal observations while generating more incorrect warnings. Such conditions may increase the workload associated with reviewing potential incidents. Axelsson (2000) highlighted the importance of low event prevalence in intrusion detection, demonstrating why false-positive behavior can substantially affect the practical usefulness of security monitoring systems.

Reducing unnecessary alerts is therefore important, but excessively conservative detection can also create difficulties. When a classification threshold is highly restrictive, some abnormal observations may remain unidentified. Network security monitoring requires an appropriate balance between identifying meaningful deviations and avoiding excessive investigation of ordinary activities. The preferred configuration cannot be determined solely through the illustrative values presented here. It must be evaluated according to the monitored environment, available investigative resources, and the consequences associated with different classification errors.

6.3 Artificial Intelligence and Network Monitoring Reliability

Artificial intelligence-based monitoring systems can examine large volumes of behavioral information and identify relationships that may be difficult to represent through fixed detection rules. Techniques such as isolation forests, one-class learning, and supervised classification provide different approaches to analyzing network observations. Their suitability depends on the availability of labeled data, feature quality, computational resources, and the distribution of abnormal events. Consequently, model selection should be informed by the characteristics of the specific monitoring environment.

The methodological limitations identified by Sommer and Paxson (2010) remain relevant to interpreting learning-based detection performance. Experimental datasets may contain simplified patterns or sampling conditions that do not represent ordinary network operation. A model performing well under a particular benchmark may encounter different traffic behavior after deployment. Reliable evaluation should therefore include independent observations, realistic class prevalence, repeated assessment, and transparent reporting of uncertainty. Such procedures are necessary before operational conclusions can be drawn from predictive performance measures.

7. Challenges

7.1 Behavioral Diversity and Data Quality

Network environments generate behavioral information from numerous applications, communication protocols, devices, and operational processes. Legitimate activity can vary substantially according to system purpose, workload, and user requirements. These variations complicate the construction of consistent behavioral baselines. Missing observations, inconsistent measurements, and incomplete network visibility may further reduce the reliability of anomaly detection. A learning system trained on incomplete or unrepresentative information may incorrectly classify ordinary activity as suspicious or fail to recognize meaningful deviations.

Data-quality challenges also influence the interpretation of machine learning features. Network-flow records may contain redundant measurements, highly correlated variables, or changes resulting from altered infrastructure configurations. These characteristics can affect anomaly scores without indicating genuine security problems. Effective behavioral analysis therefore requires careful preprocessing, feature validation, and awareness of the context in which observations were generated. Future experiments should examine whether selected behavioral variables remain informative across different network conditions and operational periods.

7.2 Class Imbalance and Detection Uncertainty

Abnormal network observations may represent a relatively small proportion of total monitored traffic. This imbalance creates difficulties when evaluating detection systems through aggregate accuracy alone. A classifier may correctly identify most ordinary observations while failing to detect many unusual events. He and Garcia (2009) examined broader challenges associated with imbalanced learning, while Saito and Rehmsmeier (2015) emphasized the importance of precision–recall evaluation. These findings support class-sensitive performance assessment in behavioral network detection research.

Another difficulty concerns uncertainty in the ground-truth classification of network activities. Unusual behavior does not always constitute malicious activity, and some abnormal observations may require contextual investigation before their status can be established. Incorrect or incomplete labels can distort performance estimates and influence supervised model development. Reliable evaluation therefore depends on appropriate labeling procedures and explicit handling of uncertain observations. Without these safeguards, apparent improvements in detection performance may reflect weaknesses in the evaluation data rather than genuine predictive capability.

7.3 Model Adaptability and Operational Constraints

Network behavior may change over time because of legitimate organizational activities, system updates, new applications, or alterations in communication requirements. A detection model developed using historical observations may become less representative of current conditions. Such changes can increase false alerts or reduce sensitivity to meaningful abnormalities. Adaptive learning procedures may help maintain model relevance, but unnecessary modification can also introduce instability. Reliable monitoring consequently requires mechanisms capable of distinguishing ordinary environmental variation from changes requiring model adjustment.

Operational constraints introduce additional challenges involving computational cost, inference latency, memory requirements, and the volume of alerts requiring review. Complex artificial intelligence models may provide sophisticated representations of behavioral patterns but demand substantial

processing resources. Furthermore, continuously updating models can complicate validation and maintenance. Effective detection architectures must therefore consider both analytical performance and deployment requirements. Resource-aware model selection, transparent monitoring procedures, and controlled updating policies are essential considerations for future practical implementation.

8. Conclusion

Artificial intelligence-based behavioral analysis provides a methodological foundation for examining abnormal activities in complex network environments. By representing communication patterns through measurable behavioral features, machine learning techniques may support the identification of observations that differ from established network conditions. The present hypothetical study developed a structured detection framework and examined the importance of anomaly scoring, classification decisions, and performance evaluation. Its findings remain illustrative and do not represent experimentally validated network security outcomes.

The hypothetical analysis demonstrated how alternative alert configurations influence precision, recall, F1 score, and false-positive rates. Increased detection sensitivity was associated with greater identification of anomalous observations but also a larger proportion of incorrect alerts. These relationships emphasize that reliable anomaly detection requires more than maximizing the number of identified irregularities. Appropriate evaluation must consider event prevalence, classification uncertainty, and the practical consequences of false-positive and false-negative decisions.

The study concludes that artificial intelligence-based network monitoring should combine behavioral feature analysis with transparent evaluation and context-sensitive detection decisions. Future research should implement and test suitable learning algorithms using representative network datasets and realistic monitoring conditions. Particular attention should be directed toward behavioral diversity, imbalanced observations, changing network patterns, and computational efficiency. Such investigations can establish whether the proposed methodological principles translate into measurable improvements in practical network anomaly detection.

References

1. Ahmad, Z., Khan, A. S., Shiang, C. W., Abdullah, J., & Ahmad, F. (2021). Network intrusion detection system: A systematic study of machine learning and deep learning approaches. *Transactions on Emerging Telecommunications Technologies*, 32(1), Article e4150. <https://doi.org/10.1002/ett.4150>
2. Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153–1176. <https://doi.org/10.1109/COMST.2015.2494502>
3. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15. <https://doi.org/10.1145/1541880.1541882>
4. Denning, D. E. (1987). An intrusion-detection model. *IEEE Transactions on Software Engineering*, SE-13(2), 222–232. <https://doi.org/10.1109/TSE.1987.232894>
5. Drewek-Ossowicka, A., Pietrolaj, M., & Rumiński, J. (2021). A survey of neural networks usage for intrusion detection systems. *Journal of Ambient Intelligence and Humanized Computing*, 12, 497–514. <https://doi.org/10.1007/s12652-020-02014-x>

6. Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, comparative study, and challenges. *Journal of Information Security and Applications*, 50, Article 102419. <https://doi.org/10.1016/j.jisa.2019.102419>
7. García-Teodoro, P., Díaz-Verdejo, J., Maciá-Fernández, G., & Vázquez, E. (2009). Anomaly-based network intrusion detection: Techniques, systems and challenges. *Computers & Security*, 28(1–2), 18–28. <https://doi.org/10.1016/j.cose.2008.08.003>
8. Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. *Cybersecurity*, 2, Article 20. <https://doi.org/10.1186/s42400-019-0038-7>
9. Lánský, J., Ali, S., Mohammadi, M., Majeed, M. K., Karim, S. H. T., Rashidi, S., Hosseinzadeh, M., & Rahmani, A. M. (2021). Deep learning-based intrusion detection systems: A systematic review. *IEEE Access*, 9, 101574–101599. <https://doi.org/10.1109/ACCESS.2021.3097247>
10. Meidan, Y., Bohadana, M., Mathov, Y., Mirsky, Y., Breitenbacher, D., Shabtai, A., & Elovici, Y. (2018). N-BaIoT: Network-based detection of IoT botnet attacks using deep autoencoders. *IEEE Pervasive Computing*, 17(3), 12–22. <https://doi.org/10.1109/MPRV.2018.03367731>
11. Mirsky, Y., Doitshman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: An ensemble of autoencoders for online network intrusion detection. In *Proceedings of the Network and Distributed System Security Symposium (NDSS 2018)*. Internet Society. <https://doi.org/10.14722/ndss.2018.23204>
12. Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). In *2015 Military Communications and Information Systems Conference (MilCIS)*. IEEE. <https://doi.org/10.1109/MILCIS.2015.7348942>
13. Patcha, A., & Park, J.-M. (2007). An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks*, 51(12), 3448–3470. <https://doi.org/10.1016/j.comnet.2007.02.001>
14. Ring, M., Wunderlich, S., Scheuring, D., Landes, D., & Hotho, A. (2019). A survey of network-based intrusion detection data sets. *Computers & Security*, 86, 147–167. <https://doi.org/10.1016/j.cose.2019.06.005>
15. Ruffo, V. G. da S., Lent, D. M. B., Komarchesqui, M., Schiavon, V. F., de Assis, M. V. O., Carvalho, L. F., & Proença, M. L., Jr. (2024). Anomaly and intrusion detection using deep learning for software-defined networks: A survey. *Expert Systems with Applications*, 256, Article 124982. <https://doi.org/10.1016/j.eswa.2024.124982>
16. Sharafaldin, I., Habibi Lashkari, A., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy* (pp. 108–116). SciTePress. <https://doi.org/10.5220/0006639801080116>
17. Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. In *2010 IEEE Symposium on Security and Privacy* (pp. 305–316). IEEE. <https://doi.org/10.1109/SP.2010.25>
18. Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. In *2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications* (pp. 1–6). IEEE. <https://doi.org/10.1109/CISDA.2009.5356528>

19. Vinayakumar, R., Alazab, M., Soman, K. P., Poornachandran, P., Al-Nemrat, A., & Venkatraman, S. (2019). Deep learning approach for intelligent intrusion detection system. IEEE Access, 7, 41525–41550. <https://doi.org/10.1109/ACCESS.2019.2895334>
20. Yin, C., Zhu, Y., Fei, J., & He, X. (2017). A deep learning approach for intrusion detection using recurrent neural networks. IEEE Access, 5, 21954–21961. <https://doi.org/10.1109/ACCESS.2017.2762418>