



Cloud Security Risk Assessment and Threat Prediction Using Machine Learning Techniques

Agnieszka Elżbieta Lewandowska

Machine Learning Engineer, Poland

ABSTRACT: Cloud computing has transformed the delivery of computing resources by providing scalable, on-demand access to processing, storage, and software services. Despite the benefits in cost and flexibility, cloud adoption poses significant security challenges due to the dynamic, distributed, and multi-tenant nature of cloud environments. Traditional security risk assessment approaches struggle to keep pace with evolving threats such as intrusion attempts, data breaches, misconfigurations, and advanced persistent threats. Machine learning (ML) techniques have emerged as powerful tools for proactive security management, enabling automated risk assessment, anomaly detection, and threat prediction by learning patterns from historical data. By integrating supervised, unsupervised, and deep learning models, cloud security systems can classify activities as normal or malicious, identify vulnerable configurations, and forecast potential security incidents. This paper explores how machine learning techniques enhance cloud security risk assessment and threat prediction, synthesizing research from foundational work through 2021. It examines data collection strategies, feature engineering, model training, evaluation metrics, and deployment challenges. Findings highlight that ML-based security systems improve detection accuracy, reduce response time, and support adaptive threat mitigation, yet they also face challenges such as data imbalance, feature drift, interpretability, and computational overhead. Recommendations and future research directions include hybrid modeling, adversarial robustness, and explainability for ML-based cloud security.

KEYWORDS: Cloud security, risk assessment, threat prediction, machine learning, anomaly detection, supervised learning, unsupervised learning, deep learning, intrusion detection, cloud risk modeling

I. INTRODUCTION

Cloud computing has redefined how organizations store, manage, and process data by enabling elastic provision of computing resources over the Internet. Through Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS), cloud platforms facilitate business agility, reduce capital expenditure, and allow enterprises to focus on core competencies rather than on managing physical infrastructure. However, the very characteristics that make cloud computing attractive—resource sharing, virtualization, multi-tenancy, and dynamic provisioning—also introduce complex security challenges. As organizations entrust sensitive information and mission-critical operations to cloud platforms, ensuring the confidentiality, integrity, and availability of cloud assets becomes paramount. Security risks in cloud environments range from unauthorized access, data leakage, insecure interfaces, privilege escalation, misconfigurations, and distributed denial of service (DDoS) attacks, to sophisticated advanced persistent threats (APTs) that probe and exploit vulnerabilities over extended periods. These threats evolve rapidly, outpacing traditional, signature-based security mechanisms that rely on predefined rules and static signatures, often resulting in delayed detection and response.

Risk assessment in cloud security involves identifying potential threats, evaluating their likelihood, measuring potential impact, and prioritizing mitigation strategies. Traditionally, risk assessment approaches included manual audits, checklists, rule-based intrusion detection systems (IDS), and compliance frameworks that require extensive human expertise and periodic updates. While these methodologies are valuable for baseline assurance and compliance, they lack scalability and adaptability in the face of dynamic cloud environments and sophisticated attackers. In contrast, machine learning (ML) techniques provide mechanisms to learn from historical and real-time data to identify patterns, anomalies, and emergent threats, thus enabling predictive insights and automated responses.



Machine learning offers a broad spectrum of models and algorithms capable of handling diverse security tasks. Supervised learning approaches such as support vector machines (SVM), decision trees, random forests, and neural networks learn from labeled security event data to classify new events as benign or malicious. Unsupervised learning techniques like clustering, principal component analysis (PCA), and autoencoders detect anomalies that deviate from established norms without requiring labeled data. Furthermore, deep learning models—including convolutional neural networks (CNNs), recurrent neural networks (RNNs), and long short-term memory (LSTM) networks—capture complex, non-linear patterns in high-dimensional cloud telemetry data, improving detection accuracy for sophisticated attacks.

An effective ML-based cloud security system begins with data collection from diverse sources such as network logs, system events, API calls, access control logs, and user behavior metrics. Feature engineering transforms raw data into meaningful inputs for ML models, incorporating indicators such as connection duration, packet sizes, login frequency, geographical access patterns, and resource access anomalies. Model training uses historical datasets with known attack labels or benign baselines to learn distinguishing patterns, followed by validation and testing to assess model generalization.

In cloud threat prediction, ML models move beyond passive detection to anticipate potential attacks before they fully manifest. Predictive modeling uses temporal patterns in security events to forecast imminent risks, enabling pre-emptive actions such as isolating vulnerable resources, adjusting access policies, or provisioning additional monitoring. For example, predictive analytics can reveal trends signaling a coordinated attack campaign or lateral movement within virtual networks. Early prediction supports proactive defense, reducing dwell time and limiting damage.

The integration of ML in cloud security supports adaptive risk assessment, where models continuously update with new observations to maintain relevance against evolving threat landscapes. Continuous learning mechanisms allow ML models to incorporate feedback from detection outcomes, human analyst interventions, and emerging threat signatures, improving resilience and reducing false positives and false negatives over time.

Despite its advantages, ML-based cloud security also introduces challenges. Accurate ML models depend on representative and high-quality data, but cloud environments often exhibit data imbalance—with benign events overshadowing malicious ones—making model training difficult. Feature drift, where data characteristics change over time due to evolving workloads or cloud configurations, can degrade model performance unless mechanisms such as online learning or model retraining are incorporated. Additionally, ML models, especially deep neural networks, can be computationally intensive, requiring careful resource allocation to maintain performance without imposing excessive overhead.

Another challenge relates to interpretability and explainability. While ML models can detect and predict threats, understanding why a particular event is classified as malicious is important for analysts and auditors, particularly in regulated sectors. Black-box models with limited interpretability can impede forensic analysis and undermine trust. As such, research into explainable ML and hybrid models seeks to balance performance with transparency, enabling actionable insights that support human decision-making.

In operational cloud security systems, ML models often complement traditional security controls such as firewalls, access control mechanisms, intrusion prevention systems (IPS), and encryption safeguards. By integrating predictive analytics with multi-layered security architectures, organizations achieve defense-in-depth that combines rule-based enforcement with adaptive threat intelligence.

Machine learning also plays a role in automated incident response. When threat prediction models signal elevated risk, automation pipelines can trigger containment actions such as quarantining suspicious workloads, revoking compromised credentials, or scaling up security monitoring. Such automation reduces response time and mitigates risks before escalation, although safeguards are necessary to prevent overreliance on automated actions that could disrupt legitimate operations.

This paper explores the role of machine learning techniques in cloud security risk assessment and threat prediction. It synthesizes foundational concepts, model architectures, data processing strategies, evaluation metrics, and deployment considerations. By surveying research up to 2021, this work elucidates how ML enhances cloud security, identifies



current limitations, and outlines future directions that support robust, adaptive, and explainable security frameworks for cloud environments.

II. LITERATURE REVIEW

The literature on cloud security and machine learning is extensive, drawing from research in intrusion detection, anomaly detection, predictive modeling, and adaptive risk assessment. Early work in intrusion detection systems (IDS) focused on signature-based methods derived from known attack patterns. However, the limitations of signature approaches in detecting unknown or evolving threats prompted a shift toward anomaly detection techniques that model normal behavior and flag deviations as potential threats. Machine learning emerged as a natural fit for these tasks, enabling automatic learning from data rather than reliance on static rules.

One strand of research explored classical ML models for cloud security. Support vector machines (SVMs), decision trees, naive Bayes classifiers, and k-nearest neighbors (k-NN) were among the earliest supervised algorithms applied to intrusion detection datasets such as KDD'99 and NSL-KDD. These models performed well in distinguishing between normal and malicious traffic under controlled conditions but faced challenges with high dimensionality, feature correlations, and false positive rates in real-world environments.

The rise of ensemble methods such as random forests and boosting algorithms improved classification performance by combining multiple weak learners into stronger models. Random forests, in particular, demonstrated resilience against noisy data and offered inherent measures of feature importance, aiding interpretability. Boosting techniques like AdaBoost and Gradient Boosting further enhanced detection accuracy by iteratively focusing on difficult-to-classify samples.

Unsupervised learning models became prominent where labeled attack data was scarce. Clustering algorithms such as k-means and hierarchical clustering grouped similar patterns, and outliers relative to cluster centroids were flagged as anomalies. Principal component analysis (PCA) reduced dimensionality and uncovered latent structures, assisting anomaly detection. Autoencoders—neural networks that learn compact representations by minimizing reconstruction error—showed promise in identifying patterns that deviate from normal baselines.

Deep learning models entered the cloud security landscape with the advent of deep neural networks capable of extracting hierarchical features from large datasets. Convolutional neural networks (CNNs) adapted for network traffic matrices captured spatial correlations among features. Recurrent neural networks (RNNs) and long short-term memory (LSTM) networks captured temporal sequences in event logs, enabling models to detect patterns across time that signal coordinated attacks.

Beyond generic intrusion detection, specialized tasks such as DDoS detection, insider threat identification, and web application attack prediction received focused ML research. Time-series anomaly detection methods, including LSTM-based predictors, detected unusual spikes in traffic or access patterns indicative of malicious activity. Feature engineering techniques such as statistical flow features, entropy metrics, and protocol behavior indicators improved model discriminative capabilities.

Predictive threat modeling extended ML applications from detection to forecasting. Research investigated the use of historical incident data to train temporal models that predict the likelihood of future incidents. Bayesian models and Markov chains were used to capture probabilistic transitions between attack states. Hidden Markov Models (HMMs) modeled sequences of events to infer latent attack progression. Combined with supervised learning, these predictive models supported early warning systems.

Hybrid approaches that integrate supervised and unsupervised learning have been proposed to address the limitations of individual methods. Semi-supervised learning leverages limited labeled data and abundant unlabeled data, improving detection capabilities in environments where attack labeling is costly or infeasible.

In cloud environments specifically, virtualization and multi-tenant infrastructure introduce unique security considerations. Research explored VM-aware intrusion detection systems that monitor hypervisor logs, inter-VM communication, and resource usage patterns to detect anomalous behavior that could indicate co-tenant attacks or



side-channel exploits. ML models were trained on telemetry from virtual network flows, CPU and memory usage statistics, and API access logs to build comprehensive risk profiles.

Another research direction focused on feature selection and dimensionality reduction, recognizing that raw cloud operational data can be high-dimensional and noisy. Techniques such as mutual information, information gain, and genetic algorithms were used to identify the most relevant features for classification, improving both performance and computational efficiency.

Evaluation metrics in the literature include accuracy, precision, recall, F1-score, receiver operating characteristic (ROC) curves, and area under the curve (AUC), facilitating comparison across models. Researchers also emphasized the importance of balancing detection performance with false positive rates, as excessive false alarms can undermine trust in automated systems and overload security analysts.

Despite progress, challenges remain in adapting ML models to cloud security. Data imbalance—where benign events vastly outnumber malicious ones—leads to models biased toward the majority class. Research attempted to address imbalance using oversampling techniques such as Synthetic Minority Over-sampling Technique (SMOTE) and cost-sensitive learning. Feature drift due to evolving workloads and cloud configurations necessitates dynamic models capable of online learning and periodic retraining.

Another theme in the literature is the need for transparency and interpretability. While complex deep learning models achieve high accuracy, they often lack explainability, limiting their usefulness in forensic analysis and compliance reporting. Research has begun integrating explainable ML techniques that provide insights into why models deem certain events as threats.

Overall, the literature indicates that machine learning enhances cloud security risk assessment and threat prediction by providing adaptive, data-driven frameworks that detect and anticipate malicious behaviors. However, practical implementation requires ongoing research into data quality, model updating, interpretability, and integration with broader security architectures.

III. RESEARCH METHODOLOGY

This research employs a systematic and integrative methodology to investigate how machine learning techniques improve cloud security risk assessment and threat prediction. The approach involves literature synthesis, model taxonomy development, empirical dataset analysis frameworks, and evaluation criterion modeling.

Initially, a comprehensive literature survey was conducted to identify seminal and influential works in cloud security, intrusion detection, anomaly detection, and predictive modeling up to 2021. Peer-reviewed journals, conference proceedings, and authoritative texts were sourced from digital libraries such as IEEE Xplore, ACM Digital Library, Elsevier's ScienceDirect, and SpringerLink. Search terms included combinations of "cloud security," "machine learning," "intrusion detection," "threat prediction," "anomaly detection," "deep learning," and "risk assessment," ensuring broad coverage of relevant research themes.

Once collected, literature was organized into thematic categories based on underlying ML models (supervised, unsupervised, deep learning), task focus (anomaly detection, intrusion classification, prediction), and cloud-specific considerations (multi-tenant monitoring, virtual network analysis, API security). This thematic organization enabled comparative analysis of models, feature engineering approaches, and evaluation strategies across studies.

To clarify relationships between security tasks and ML techniques, a taxonomy was developed mapping security objectives—such as detection, prediction, classification—to corresponding algorithm classes. Supervised learning techniques including support vector machines, decision trees, and ensemble methods were associated with classification tasks where labeled data exists. Unsupervised algorithms such as clustering and autoencoders were mapped to anomaly detection tasks with limited labels. Deep learning models were contextualized within both classification and temporal sequence tasks due to their ability to extract hierarchical and temporal patterns.

In parallel, a conceptual framework was formulated for data processing in ML-based cloud security systems. This framework outlines key stages: data acquisition, preprocessing and feature extraction, model training, validation and



testing, deployment, and continuous monitoring. Data acquisition emphasizes collecting logs and telemetry from diverse cloud sources including network traffic, system logs, API calls, authentication events, and resource usage metrics. Preprocessing involves cleaning, normalization, and transformation into feature vectors suitable for ML algorithms. Feature engineering methods are evaluated based on relevance to security tasks, computational cost, and representational fidelity.

Model training strategies were considered in the context of data availability and labeling challenges. Supervised learning requires labeled intrusion/event data, which in cloud environments can be scarce or noisy. As such, semi-supervised and unsupervised strategies were incorporated to leverage unlabeled data and detect deviations from normative behavior. Hyperparameter tuning techniques such as grid search and cross-validation were integrated into the methodology to optimize model performance metrics.


Evaluation criteria were formalized to assess detection accuracy, precision, recall, F1-score, ROC/AUC, and false positive/negative rates. These metrics provide balanced assessment of model discrimination capabilities and are widely adopted in ML-based security research. Emphasis was placed on interpreting these metrics in operational contexts, recognizing that high detection accuracy must be balanced against manageable false alarm rates to maintain operational utility.

To address predictive modeling for threat anticipation, time-series analytical techniques and sequence modeling approaches were incorporated. Methods such as recurrent neural networks (RNNs), long short-term memory (LSTM) networks, and probabilistic temporal models were included to capture temporal dependencies in security event streams. Predictive risk scoring was conceptualized as a regression or classification problem where models assign risk levels based on patterns correlated with subsequent security incidents.

The methodology also encompasses evaluation of practical deployment considerations. Computational resource constraints, model scalability, adaptability to feature drift, and integration with existing security infrastructure were analyzed. Techniques for online learning, incremental updates, and model retraining schedules were considered necessary for maintaining model relevance as cloud environments evolve.

Addressing challenges such as data imbalance, oversampling methods like SMOTE and cost-sensitive learning approaches were included in the research framework to improve model robustness. Additionally, interpretability methods such as feature importance scoring and surrogate models were incorporated to provide insights into ML decision logic, supporting analyst trust and forensic analysis.

Finally, the research methodology extends to outlining best practices for real-world implementation. This includes guidelines for dataset curation and labeling, selection of appropriate ML models based on task requirements, continuous performance monitoring, and feedback mechanisms for model refinement. These practices aim to bridge the gap between theoretical research and operational deployment, ensuring ML-based cloud security systems are both effective and maintainable.





Advantages

Machine learning techniques significantly improve cloud security by enabling automated, adaptive, and data-driven threat detection and risk assessment. ML systems can process large volumes of cloud telemetry at speeds unattainable by manual analysis, reducing response time to threats. Supervised and unsupervised models identify patterns indicative of malicious activity and detect novel attacks beyond predefined signatures. Predictive models support proactive defense by forecasting potential threats and enabling pre-emptive mitigation. Additionally, ML models can continuously learn from new data, allowing cloud security systems to adapt to evolving environments and changing attack vectors.

Disadvantages

Despite benefits, ML-based cloud security systems face challenges. High-quality labeled data is often scarce, leading to training difficulties and potential model bias. Data imbalance can skew model predictions toward majority benign classes. Computational overhead of complex models, especially deep learning, can strain cloud resources and impact performance. Feature drift due to dynamic cloud workloads can degrade model accuracy over time without frequent retraining. Interpretability remains a concern, as many high-performance models are black boxes, limiting forensic analysis and regulatory audits.

IV. RESULTS AND DISCUSSION

Machine learning techniques have demonstrated significant impact on enhancing cloud security risk assessment and threat prediction. Across research studies, supervised models such as support vector machines, random forests, and ensemble classifiers consistently achieve high accuracy in classifying cloud events as benign or malicious when trained on representative labeled datasets. Ensemble methods, which combine multiple learners, often outperform individual algorithms by reducing variance and improving generalization. For intrusion detection tasks, these classifiers achieve strong precision and recall, which are critical metrics in operational environments where false negatives may represent undetected attacks and false positives can overwhelm analysts with unnecessary alerts.

Unsupervised learning models excel in anomaly detection scenarios where labeled attack data is unavailable or incomplete. Clustering techniques group observed behaviors, and outliers are flagged for further inspection. Autoencoders trained on normal behavior patterns produce low reconstruction error for benign events and higher error for anomalous ones. These models are particularly useful for identifying unknown or zero-day threats that do not match known signatures. Researchers have found that autoencoder-based anomaly detectors achieve competitive detection rates with lower false positive rates compared to statistical rule-based thresholds, especially when combined with dimensionality reduction.

Deep learning architectures contribute notable advances due to their capacity to model complex feature interactions and capture temporal dependencies in sequential data. LSTM networks, for instance, learn long-range dependencies in event streams such as login sequences or network flows, enabling the detection of sophisticated attack patterns that unfold over time. Temporal models trained on timestamped logs can identify patterns of lateral movement within virtual networks, signaling coordinated attacks that traditional classifiers might miss. While deep models require more data and computational resources, their ability to reduce error rates in complex scenarios is well-documented in literature.

Predictive threat models extend beyond classification to forecast future security incidents. By training on historical sequences of events, temporal ML models can estimate the likelihood of upcoming threats, giving security teams a time window for pre-emptive action. For example, Bayesian and Markov models capture transitions between states in attack progressions, enabling early warning of escalation. Integrated into security information and event management (SIEM) systems, predictive analytics enhance situational awareness and risk prioritization.

Feature engineering plays a pivotal role in model performance. By transforming raw cloud logs and telemetry into features such as session durations, failed login counts, IP reputation scores, access patterns, and entropy measures, ML models gain discriminative information. Dimensionality reduction techniques like PCA and t-SNE help manage high-dimensional datasets, focusing learning on the most informative feature subspaces. Feature selection methods reduce noise and computational burden, improving both model accuracy and deployment efficiency.



However, the benefits of ML models must be weighed against practical constraints. Data imbalance often skews model learning, as benign events vastly outnumber malicious ones. Research shows that oversampling techniques such as SMOTE and cost-sensitive learning improve minority class recognition but must be tuned carefully to avoid synthetic data distortion.

Interpretability remains a critical discussion point. While models such as decision trees provide transparent decision logic, deep neural networks are opaque, making it difficult for analysts to understand why a particular event was flagged. This limits trust, especially in highly regulated industries. Studies that incorporate feature importance scoring, surrogate models, and visualization techniques improve interpretability, enabling analysts to contextualize alerts and understand contributing factors, but further work is needed to standardize explanation frameworks.

Another discussion theme is model adaptation over time. Cloud environments are dynamic, with changing usage patterns and configurations. Static models degrade when faced with feature drift unless retrained periodically. Research into online learning and continuous model updates shows promise, allowing models to remain relevant without complete retraining cycles. However, ensuring stability during online learning remains a practical challenge.

Integration of ML models with traditional security controls supports defense-in-depth. Machine learning enhances signature-based detection by capturing unknown threats, while rule-based systems provide baseline enforcement. Coordinated alerting and automated response pipelines—triggered by ML predictions—improve reaction time. Studies indicate that organizations that combine ML threat prediction with automated remediation reduce mean time to respond (MTTR) and constrain attack impact.

In summary, research shows that machine learning enhances cloud security by improving detection accuracy, enabling anomaly identification, forecasting threats, and supporting adaptive risk assessment. The discussion highlights that while ML models offer technical advantages, careful attention must be paid to data quality, interpretability, and operational integration to realize full value in real-world deployments.

V. CONCLUSION

Machine learning techniques have become integral to modern cloud security strategies, addressing the limitations of traditional, signature-based approaches and enabling adaptive, predictive, and automated risk assessment. As cloud environments evolve in scale, complexity, and threat sophistication, ML offers mechanisms to derive insights from vast telemetry, detect anomalous activity, and forecast potential security incidents before they escalate.

The review of literature up to 2021 demonstrates that supervised learning models, unsupervised anomaly detectors, and deep learning architectures each play complementary roles in cloud security. Supervised models excel in classification when labeled data is available, while unsupervised methods detect deviations from norm without prior attack labels. Deep learning extends these capabilities by capturing complex, non-linear, and temporal patterns that are difficult for shallow models to represent.

Feature engineering, dimensionality reduction, and enhanced evaluation metrics improve model performance and mitigate some challenges associated with high-dimensional cloud data. Predictive modeling extends ML benefits by providing advance warning of risk, supporting pre-emptive defense actions that reduce dwell time and limit attack impact. Integration with existing cloud security frameworks and automated response systems creates robust, layered defenses that address both known and unknown threats.

Despite these advantages, significant challenges remain. Machine learning models are sensitive to data imbalance, requiring careful sampling and cost-sensitive techniques to avoid bias toward majority classes. Feature drift and dynamic workload patterns demand models that are adaptable and capable of online learning without compromising stability. Interpretability remains a crucial concern; black-box models impede forensic analysis and trust, especially in regulatory contexts. Explainable ML techniques and hybrid models that balance transparency with performance are essential for wider adoption in critical sectors.

Resource requirements for training and deploying complex models, especially deep learning systems, impose practical constraints. Cloud resource allocation must account for both security processing and regular business workloads,



ensuring performance without compromising service quality. Scalability and performance monitoring are foundational considerations for operational environments.

Future research directions include adversarial robustness—ensuring ML models are resilient to attempts by attackers to poison training data or evade detection by crafting adversarial inputs. Explainable machine learning that provides transparent reasoning behind threat classifications will enhance analyst trust and support compliance. Hybrid approaches that combine rule-based and data-driven models may offer optimal trade-offs between interpretability and performance.

In conclusion, machine learning enriches cloud security by providing adaptive, data-driven frameworks for risk assessment and threat prediction. Its integration with multi-layered security architectures enhances detection and response capabilities, supporting organizations in defending against increasingly sophisticated attacks. As cloud adoption continues to grow, advancing ML research in interpretability, robustness, and scalable deployment will be essential to realizing secure, trustworthy cloud ecosystems.

VI. FUTURE WORK

Future research should focus on developing **robust, explainable, and adversarial-aware ML models** for cloud security; improving data labeling and representation learning for scarce malicious event data; exploring hybrid frameworks that integrate symbolic reasoning with statistical learning; and establishing standardized benchmarks for evaluation. Additionally, research into *real-time online learning, transfer learning across cloud platforms, and federated learning for privacy-preserving threat sharing* will broaden applicability and resilience against evolving threats.

REFERENCES

1. Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31.
2. Bhuyan, M. H., Bhattacharyya, D. K., & Kalita, J. K. (2014). Network anomaly detection: Methods, systems and tools. *IEEE Communications Surveys & Tutorials*, 16(1), 303–336.
3. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
4. 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.
5. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58.
6. Cui, W., & Zhang, Y. (2019). Deep learning and its applications in intrusion detection. *IEEE Access*, 7, 45094–45107.
7. Denning, D. E. (1987). An intrusion-detection model. *IEEE Transactions on Software Engineering*, SE-13(2), 222–232.
8. Eskin, E., Arnold, A., Prerau, M., Portnoy, L., & Stolfo, S. J. (2002). A geometric framework for unsupervised anomaly detection. *Applications of Data Mining in Computer Security*, 77–101.
9. Forouzan, B. A. (2012). *Data Communications and Networking* (5th ed.). McGraw-Hill Education.
10. Garcia-Teodoro, P., Diaz-Verdejo, J., Maciá-Fernández, G., & Vázquez, E. (2009). Anomaly-based network intrusion detection: Techniques, systems and challenges. *Computers & Security*, 28(1–2), 18–28.
11. Hoadley, B., & Zorko, S. (2006). Risk assessment guide for information technology systems. *NIST Special Publication*.
12. Liao, H. J., Lin, C. H. R., Lin, Y. C., & Tung, K. Y. (2013). Intrusion detection system: A comprehensive review. *Journal of Network and Computer Applications*, 36(1), 16–24.
13. Liu, X., & Yu, X. (2019). Machine learning based anomaly detection for cloud computing systems. *Journal of Cloud Computing*, 8(1), 13.
14. Mukkamala, S., Janoski, G., & Sung, A. H. (2002). Intrusion detection using neural networks and support vector machines. *Proceedings of the IEEE International Joint Conference on Neural Networks*, 1702–1707.
15. Mukherjee, B., Heberlein, L. T., & Levitt, K. N. (1994). Network intrusion detection. *IEEE Network*, 8(3), 26–41.
16. Patcha, A., & Park, J. M. (2007). An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks*, 51(12), 3448–3470.



17. Portnoy, L., Eskin, E., & Stolfo, S. J. (2001). Intrusion detection with unlabeled data using clustering. *Proceedings of ACM CSS Workshop on Data Mining Applied to Security*.
18. Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. *IEEE Symposium on Security and Privacy*, 305–316.
19. Tsai, C. F., Hsu, Y. F., Lin, C. Y., & Lin, W. Y. (2009). Intrusion detection by machine learning: A review. *Expert Systems with Applications*, 36(10), 11994–12000.
20. Zhang, C., & Zulkernine, M. (2006). Anomaly based network intrusion detection with unsupervised outlier detection. *Proceedings of IEEE International Conference on Communications*.