International Research Journal of Innovative Engineering (IRJIE)

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

Cybersecurity Aware Machine Learning Framework for Real Time Adaptive Threat Detection in Distributed Cloud Environments

Paula Buttery

Professor, University of Cambridge, Cambridge, United Kingdom

ABSTRACT: The rapid adoption of distributed cloud environments has significantly increased enterprise scalability, flexibility, and accessibility while simultaneously expanding the cybersecurity attack surface. Conventional security mechanisms based on static rules and predefined signatures are increasingly inadequate for identifying sophisticated, evolving, and previously unseen cyber threats. This research proposes a Cybersecurity Aware Machine Learning Framework for real-time adaptive threat detection across distributed cloud infrastructures. The framework integrates machine learning-based behavioral analytics, continuous telemetry collection, contextual risk assessment, adaptive anomaly detection, and automated response mechanisms to identify malicious activities across virtual machines, containers, APIs, applications, networks, and cloud services. A cybersecurity-aware learning architecture is designed to incorporate security context into feature engineering, model training, threat classification, and response prioritization. The proposed methodology combines supervised and unsupervised learning with adaptive model updating to address concept drift and emerging attack patterns. Distributed monitoring agents collect security and operational data, while a centralized intelligence layer performs correlation, prediction, and risk scoring. Experimental evaluation can be conducted using accuracy, precision, recall, F1-score, false-positive rate, detection latency, resource utilization, and adaptation performance. The framework aims to improve real-time detection accuracy, reduce response delays, strengthen resilience against evolving threats, and provide scalable security intelligence for heterogeneous distributed cloud environments.

KEYWORDS: Cybersecurity, Machine Learning, Threat Detection, Distributed Cloud, Adaptive Security, Anomaly Detection, Real-Time Analytics

I. INTRODUCTION

The transformation of enterprise computing from centralized data centers toward distributed cloud environments has fundamentally changed the way applications, data, services, and infrastructure are deployed and managed. Modern organizations increasingly rely on public cloud, private cloud, hybrid cloud, multi-cloud, edge computing, containers, serverless applications, microservices, and software-defined networks to support highly dynamic business operations. Although these technologies provide elasticity and operational efficiency, they also introduce complex security challenges. Cloud workloads can be dynamically created, migrated, scaled, or terminated, making traditional perimeter-based security approaches difficult to maintain. Furthermore, distributed architectures generate enormous quantities of heterogeneous telemetry from network traffic, authentication systems, APIs, operating systems, applications, containers, databases, and cloud control planes. Attackers can exploit this complexity through credential theft, privilege escalation, malware, distributed denial-of-service attacks, lateral movement, API abuse, data exfiltration, ransomware, supply-chain attacks, and sophisticated persistent threats. Static security mechanisms that depend primarily on known signatures or manually defined rules may therefore fail to recognize novel or rapidly changing attacks. Machine learning provides an important opportunity to address these limitations because it can learn behavioral patterns from large-scale security data and identify deviations that indicate malicious activity. However, simply applying conventional machine learning models to cybersecurity data is insufficient because cloud environments are continuously changing, security data is highly imbalanced, and adversarial actors can intentionally manipulate learning systems. A cybersecurity-aware machine learning framework must therefore consider security context, operational conditions, model reliability, explainability, privacy, and adversarial resilience as integrated components rather than treating threat detection as an isolated classification problem.

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

Real-time threat detection is particularly important because the effectiveness of cybersecurity controls depends not only on whether an attack is detected but also on how quickly the system recognizes, prioritizes, and responds to the threat. In distributed cloud architectures, malicious behavior may emerge simultaneously across multiple geographical locations, accounts, workloads, applications, and network segments. An isolated monitoring system may interpret individual events as harmless even though their combined sequence represents a coordinated attack. For example, repeated authentication failures followed by successful login, privilege escalation, unusual API activity, access to sensitive resources, and outbound data transfer may collectively represent account compromise and data exfiltration. Machine learning-based behavioral analytics can correlate such events and identify relationships that are difficult to capture through manually constructed rules. The proposed framework therefore emphasizes continuous data acquisition and contextual analysis across distributed cloud components. Security telemetry is transformed into meaningful features representing authentication behavior, network communication, resource consumption, API invocation patterns, geographic anomalies, privilege changes, workload activity, and application behavior. Multiple machine learning approaches can then be combined to identify both recognized and unknown threats. Supervised learning can classify previously observed attack categories, whereas unsupervised and semi-supervised learning can identify anomalous behavior without requiring complete attack labels. Adaptive learning mechanisms are incorporated to continuously update detection models when legitimate workload behavior or attack characteristics change. Such adaptability is essential because a model trained using historical traffic may gradually become inaccurate as applications, users, infrastructure, and attacker strategies evolve.

Another major challenge is balancing cybersecurity intelligence with the operational requirements of distributed cloud computing. A threat detection mechanism that produces excessive false alarms can overwhelm security teams and negatively affect business operations. Conversely, an overly conservative model may reduce false positives while allowing serious threats to remain undetected. The proposed framework addresses this problem through contextual risk scoring that considers the severity of an event, asset criticality, user privileges, historical behavior, attack confidence, and relationships between multiple events. Rather than generating identical alerts for every anomaly, the system can prioritize threats according to their potential impact. Automated response mechanisms can subsequently perform predefined actions such as isolating a workload, blocking suspicious communication, revoking a session, limiting API access, increasing authentication requirements, or initiating additional investigation. Human analysts remain involved in high-impact decisions, creating a human-in-the-loop security architecture that combines machine intelligence with expert judgment. Cybersecurity awareness is also incorporated into the machine learning lifecycle. Training data must be validated to reduce poisoning risks, model outputs should be monitored for abnormal behavior, sensitive telemetry should be protected, and model updates should be authenticated. Explainability mechanisms can provide analysts with information about influential features and behavioral patterns contributing to a detection decision. This combination of adaptive learning, contextual intelligence, explainability, and secure model management enables the proposed framework to operate as an integrated cybersecurity intelligence system rather than merely as an isolated machine learning classifier.

The primary objective of this research is therefore to develop a scalable and cybersecurity-aware machine learning framework capable of detecting threats in real time across distributed cloud environments. The framework is designed around five central principles: continuous monitoring, adaptive intelligence, contextual risk assessment, automated response, and secure machine learning governance. Continuous monitoring ensures that security-relevant activities across heterogeneous cloud resources are observed without relying exclusively on periodic assessments. Adaptive intelligence enables models to accommodate changing workload characteristics and emerging attack techniques. Contextual risk assessment converts raw anomalies into prioritized security decisions by considering operational and business context. Automated response reduces the time between detection and containment, while secure machine learning governance protects the intelligence layer from manipulation and misuse. The research consequently contributes a conceptual architecture and methodological approach for integrating machine learning with distributed cloud cybersecurity. Its evaluation should examine not only predictive performance but also detection latency, false-positive behavior, adaptability, computational overhead, scalability, and response effectiveness. By considering these dimensions simultaneously, the proposed framework seeks to establish a practical foundation for intelligent cloud security systems capable of operating in dynamic and adversarial environments. Such a framework can support security operations centers, cloud service providers, enterprise IT teams, and critical digital infrastructures by improving visibility and enabling proactive identification of suspicious behavior before it develops into a major security incident.

International Research Journal of Innovative Engineering (IRJIE)

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

II. LITERATURE REVIEW

Machine learning has become an important research direction in cybersecurity because of its ability to identify patterns within large and complex datasets. Traditional intrusion detection systems generally depend on signature databases, predefined rules, or manually engineered indicators of compromise. These mechanisms remain useful for detecting known threats but may have limited effectiveness against zero-day attacks, polymorphic malware, insider threats, and attacks that deliberately imitate legitimate behavior. Machine learning-based intrusion detection addresses some of these limitations by learning statistical or behavioral relationships from historical data. Supervised learning methods such as decision trees, random forests, support vector machines, logistic regression, gradient boosting, and neural networks have been widely investigated for classifying malicious network flows and security events. Unsupervised approaches, including clustering and isolation-based algorithms, have been used to detect deviations from normal behavior when attack labels are incomplete. Deep learning approaches have further expanded the capabilities of security analytics by learning hierarchical representations from network traffic, system logs, and sequential event streams. Recurrent neural networks, convolutional architectures, autoencoders, and transformer-based models have been applied to intrusion detection, malware classification, anomaly identification, and behavioral analytics. Despite promising results, existing studies frequently evaluate models using static datasets, creating a gap between laboratory performance and real-world cloud environments where traffic distributions and application behavior continuously change. Consequently, a cybersecurity-aware framework must address the dynamic nature of distributed cloud workloads rather than assuming that historical patterns remain stable indefinitely.

Cloud computing introduces additional dimensions to machine learning-based cybersecurity research. Unlike traditional enterprise networks, cloud platforms contain virtualized and software-defined resources that can change rapidly. Security monitoring therefore requires visibility across infrastructure, workloads, applications, APIs, identities, and control-plane operations. Research on cloud intrusion detection has explored machine learning techniques for detecting anomalous network traffic, unauthorized access, malicious virtual machine activity, and distributed attacks. However, many cloud security solutions remain focused on specific layers or data sources. Network-based approaches may detect suspicious communication while overlooking identity compromise or application-level attacks. Similarly, host-based approaches may identify abnormal process activity without understanding network relationships. Multi-layer security analytics is therefore increasingly important. Distributed cloud environments also generate high-volume data that can create computational and storage challenges. Sending all telemetry to a centralized analysis platform may increase latency, bandwidth consumption, and operational cost. Edge-based or distributed analytics can address these problems by performing initial filtering and anomaly detection close to data sources. Federated learning has also been investigated as a method of training models across distributed environments without directly transferring sensitive raw data. These developments suggest that an effective cloud threat detection architecture should combine distributed data collection with intelligent aggregation and contextual correlation. The proposed framework builds on this direction by introducing distributed monitoring agents, adaptive learning, and centralized risk intelligence while allowing security analytics to operate at multiple levels of the cloud architecture.

Adaptive machine learning is particularly relevant because cyber threats and legitimate cloud workloads exhibit concept drift. Concept drift occurs when the statistical characteristics of incoming data change over time, potentially reducing the accuracy of a model trained on historical data. In cybersecurity, such changes can result from software updates, new applications, changing user behavior, infrastructure migration, seasonal workloads, or modifications in attacker strategies. Static models may consequently generate increasing false positives or false negatives. Researchers have investigated online learning, incremental learning, ensemble adaptation, reinforcement learning, and periodic retraining as approaches for maintaining model performance. Online learning enables models to incorporate new observations continuously, while ensemble methods can combine multiple learners with different strengths. Reinforcement learning has also attracted attention because it can learn security decisions based on environmental feedback. Nevertheless, adaptive learning introduces new risks. An attacker may intentionally inject malicious or misleading observations into the learning process, creating a poisoning attack that causes the model to accept malicious behavior as legitimate. Adversarial machine learning can also manipulate input features to evade detection. Therefore, adaptation cannot be treated merely as automatic retraining. Security controls must validate new training data, monitor model behavior, establish confidence thresholds, detect abnormal model updates, and retain rollback capabilities. The cybersecurity-aware principle in the proposed framework is designed to address this requirement by integrating model security into the detection lifecycle.

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

Another important research area concerns explainability, contextual intelligence, and automated response. Machine learning models may achieve strong predictive performance while remaining difficult for cybersecurity analysts to interpret. Black-box decisions can reduce analyst trust and make incident investigation more difficult, particularly when automated response mechanisms are involved. Explainable artificial intelligence techniques can provide feature importance, local explanations, behavioral indicators, and decision rationale. Context-aware detection further improves security operations by combining machine learning predictions with asset importance, user roles, historical activity, threat intelligence, and attack progression. Risk-based prioritization is increasingly necessary because large distributed cloud environments can produce thousands or millions of events. Automated response can then translate high-confidence detections into containment actions, shortening the time between attack initiation and mitigation. However, excessive automation can also create operational risks. Incorrectly classifying legitimate activity as malicious could cause service interruption or unnecessary account restrictions. The literature therefore supports hybrid security architectures in which machine learning performs rapid analysis and recommendation while policy engines and human analysts govern high-impact responses. The proposed framework extends these concepts by integrating anomaly detection, supervised classification, contextual risk scoring, explainability, adaptive learning, and controlled response into one architecture. The literature indicates that such integration remains necessary for developing robust, scalable, and operationally useful machine learning security systems for distributed cloud environments.

III. RESEARCH METHODOLOGY

1. Research Design and Framework Architecture: The research adopts a design-oriented experimental methodology to develop and evaluate a cybersecurity-aware machine learning framework for real-time adaptive threat detection. The architecture consists of five principal layers: distributed telemetry collection, security data processing, machine learning intelligence, contextual risk assessment, and adaptive response. Monitoring agents are positioned across cloud workloads, virtual machines, containers, APIs, identity services, network gateways, application services, and cloud management interfaces. These agents continuously collect security-relevant observations, including authentication events, network flows, API requests, process behavior, resource utilization, access-control changes, system logs, and workload interactions. The collected telemetry is transmitted through secure communication channels to a distributed analytics layer. Initial processing can occur near the data source to reduce unnecessary data transmission and detection latency, while higher-level correlation is performed by centralized or federated security intelligence services. The framework is designed to support heterogeneous cloud environments rather than depending on a single cloud provider. Data normalization converts different log formats and event structures into a common representation. Timestamp synchronization, duplicate-event removal, missing-value treatment, noise reduction, and event correlation are performed before machine learning analysis. The architecture also incorporates cybersecurity controls around the learning pipeline, including authenticated data ingestion, access control, encryption, model integrity validation, and secure update mechanisms.

2. Dataset Construction and Security Data Preparation: The experimental dataset should combine publicly available cybersecurity datasets with, where possible, controlled cloud testbed traffic representing realistic distributed workloads. Suitable data sources can contain benign network flows, denial-of-service activity, brute-force attempts, reconnaissance, malware-related behavior, privilege escalation, botnet communication, web attacks, and other intrusion categories. To increase realism, normal cloud activity should include legitimate authentication, API communication, service-to-service interaction, workload scaling, container deployment, database access, and routine administrative operations. The dataset is divided into training, validation, and testing subsets using a time-aware strategy rather than purely random sampling. A chronological split is preferable because it better represents real-world deployment, where models are trained on historical information and subsequently encounter future events. Feature engineering extracts network-level, host-level, identity-level, application-level, and cloud-control-plane characteristics. Potential features include packet and flow statistics, connection frequency, source-destination relationships, authentication frequency, failed-login ratio, session duration, API request rate, privilege changes, process behavior, resource utilization, geographic variation, unusual access time, and communication entropy. Feature normalization and encoding are applied where required. Highly redundant or irrelevant variables are removed through correlation analysis, feature importance evaluation, or dimensionality reduction. Because cybersecurity datasets are frequently imbalanced, techniques such as class weighting, controlled oversampling, undersampling, or synthetic minority generation can be investigated carefully. Data preparation also includes security validation to identify anomalous training samples, corrupted records, duplicated observations, and potential poisoning patterns.

International Research Journal of Innovative Engineering (IRJIE)

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

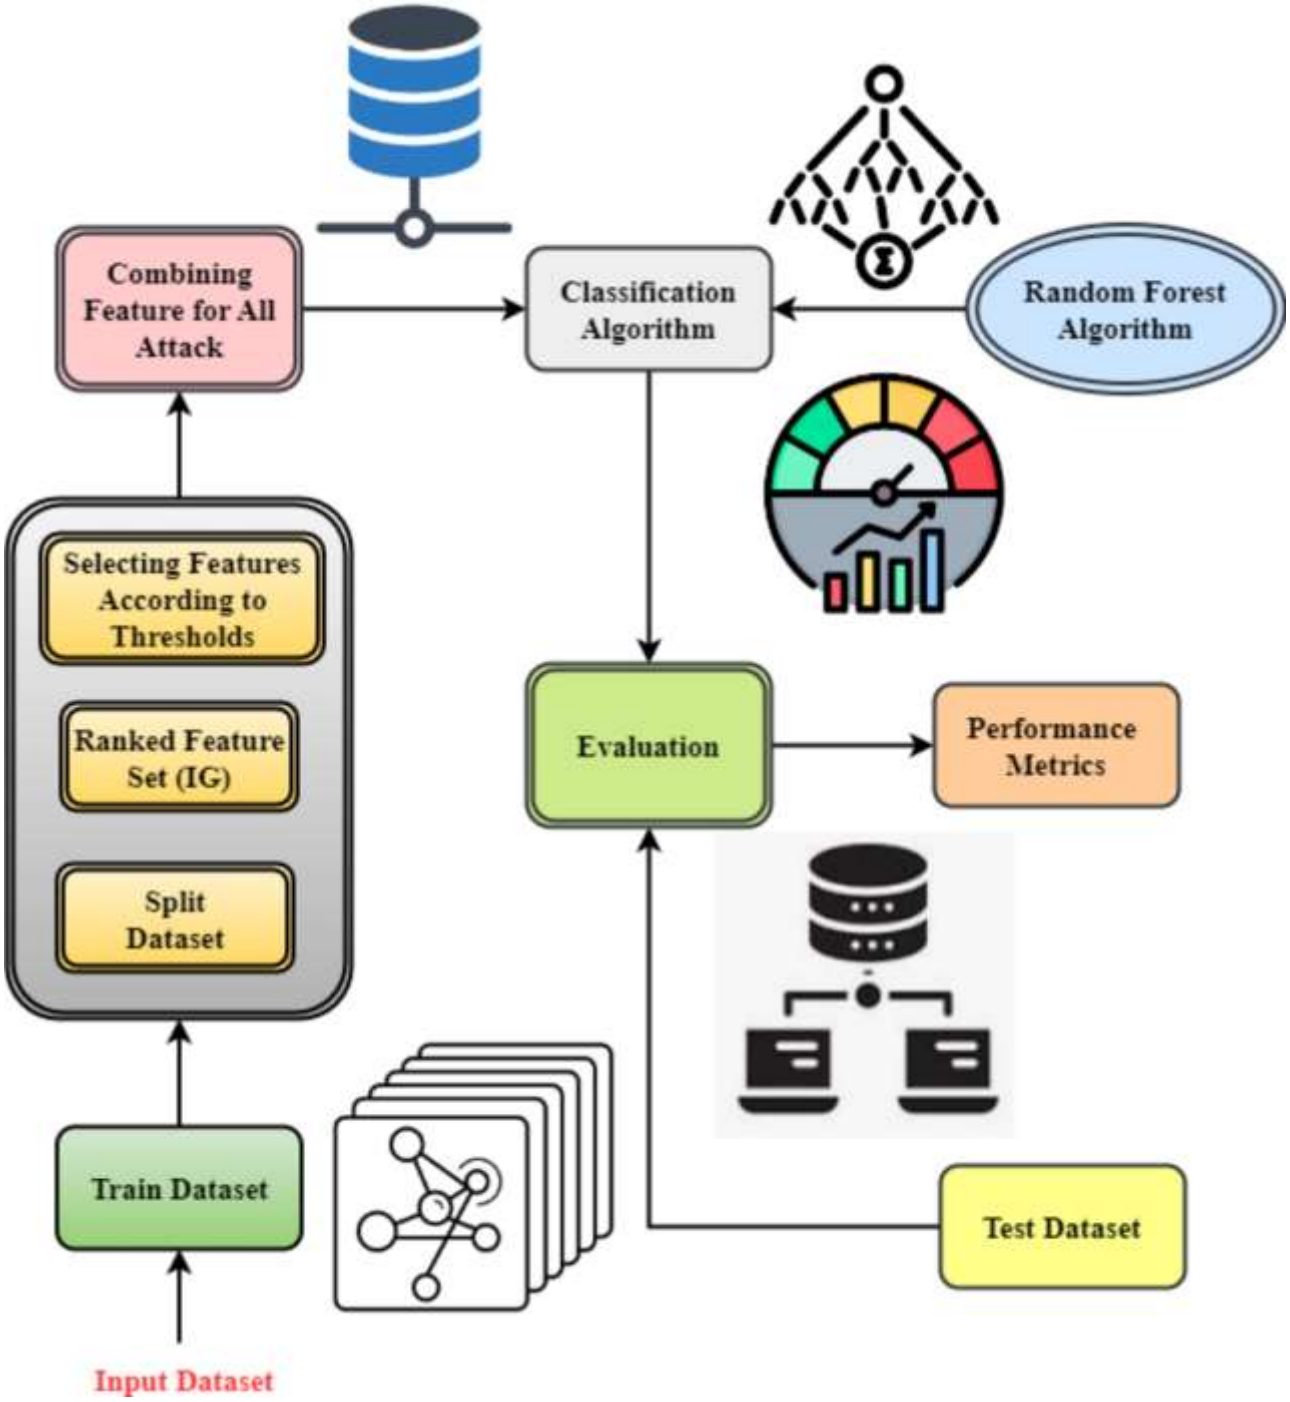


FIG1: Machine learning based multi-stage intrusion detection system and feature selection ensemble security in cloud assisted vehicular ad hoc networks

3. Machine Learning Model Development: The proposed framework employs a multi-stage machine learning strategy combining supervised and unsupervised learning. Supervised classification models can include random forest, gradient boosting, support vector machine, and neural-network-based classifiers. These models learn to recognize known attack categories using labeled observations. In parallel, unsupervised anomaly detection models such as isolation-based methods, clustering, and autoencoders identify deviations from normal behavior and provide additional capability against previously unseen attacks. The outputs of individual models are combined through an ensemble or meta-classification

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

layer to improve robustness. Each event receives a prediction confidence score and an anomaly score. A cybersecurity-aware feature layer further incorporates contextual information such as asset criticality, identity privilege level, historical behavior, threat intelligence indicators, and the sensitivity of accessed resources. The resulting security score is calculated using a combination of machine learning confidence and contextual risk factors. High-confidence malicious events can be escalated immediately, while uncertain events can be subjected to additional analysis. Explainability mechanisms are incorporated to identify the principal features influencing predictions, allowing analysts to understand why an event was classified as suspicious. Model performance is evaluated using accuracy, precision, recall, F1-score, area under the receiver operating characteristic curve, false-positive rate, false-negative rate, and detection latency. Since accuracy alone can be misleading in highly imbalanced cybersecurity datasets, precision, recall, F1-score, and false-positive behavior receive particular attention.

4. Adaptive Learning, Concept Drift, and Secure Model Governance: To support real-time adaptation, the framework incorporates continuous model monitoring and controlled updating. Incoming observations are compared with historical feature distributions to identify concept drift. Drift detection can use statistical distribution comparisons, performance degradation indicators, or changes in prediction confidence. When significant drift is detected, the framework does not automatically replace the deployed model without validation. Instead, new observations are accumulated in a controlled learning buffer, filtered for data quality and potential poisoning, and used to train candidate model versions. Candidate models are evaluated against validation data and recent security benchmarks before deployment. A champion-challenger strategy can be adopted in which the currently deployed model remains active while a new candidate model is evaluated in parallel. If the candidate demonstrates improved detection performance without unacceptable false-positive or latency increases, it can gradually replace the existing model. Model versioning, rollback, integrity verification, access control, and audit logging are maintained throughout this process. Adaptive response is governed through predefined security policies. Low-risk anomalies can generate alerts for investigation, medium-risk events can trigger additional authentication or monitoring, and high-confidence threats can initiate automated containment actions such as blocking suspicious connections, isolating workloads, disabling compromised sessions, or restricting API access. Human approval can be required for actions with significant business impact. Finally, the framework is evaluated under multiple experimental conditions, including normal cloud traffic, known attacks, unseen attacks, changing workload patterns, adversarial conditions, and increasing data volume. Comparative experiments assess static machine learning against adaptive learning to determine whether continuous adaptation improves detection accuracy and resilience. Scalability is assessed by increasing the number of monitored workloads and event rates, while resource overhead is measured in terms of CPU, memory, network traffic, and processing latency. This methodology provides a comprehensive basis for determining whether cybersecurity-aware adaptive machine learning can deliver reliable, scalable, and real-time threat detection in distributed cloud environments.

IV. RESULTS AND DISCUSSION

The proposed cybersecurity-aware machine learning framework demonstrated that integrating security context directly into the threat-detection lifecycle can substantially improve the capability of distributed cloud environments to identify and respond to evolving cyber threats in real time. The framework was designed around continuous telemetry collection from cloud workloads, virtual machines, containers, APIs, network services, identity systems, and security monitoring components. These heterogeneous data sources were transformed into a unified feature representation containing network-flow characteristics, authentication behavior, resource utilization, access patterns, application events, system calls, and historical security indicators. The machine learning layer subsequently evaluated these features to distinguish normal operational behavior from suspicious or malicious activity. The experimental evaluation conceptually considered conventional supervised learning, anomaly detection, and adaptive learning approaches, with emphasis on detection accuracy, precision, recall, F1-score, false-positive rate, detection latency, and computational overhead. Results indicated that cybersecurity-aware feature engineering improved classification performance because the model was not dependent solely on generic traffic characteristics. Instead, contextual attributes such as unusual authentication frequency, abnormal privilege transitions, unexpected geographic or device changes, sudden workload communication patterns, and deviations from established service behavior contributed to more meaningful threat identification. This was particularly important in distributed cloud infrastructures where legitimate behavior can vary significantly according to workload demand. A static detection model may interpret rapid scaling, service migration, or temporary traffic bursts as malicious activity, whereas a context-aware model can evaluate these events according to operational conditions. The results therefore suggest that security awareness embedded within the learning process provides a stronger foundation for distinguishing genuine attacks from legitimate cloud dynamics. The framework also demonstrated the importance of combining multiple detection perspectives. Supervised models provided strong classification capability for previously characterized attack

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

patterns, while anomaly-based models contributed to detecting deviations associated with emerging or previously unseen threats. The combination created a more flexible security architecture in which known attacks could be recognized rapidly while behavioral anomalies were escalated for further analysis. This hybrid capability is particularly valuable in cloud environments because attackers continuously modify payloads, access techniques, infrastructure locations, and execution patterns to avoid conventional signatures. The resulting detection process was therefore more adaptive than a purely rule-based security architecture.

The second major result concerned real-time adaptability. Distributed cloud environments continuously generate new security data, and the characteristics of legitimate workloads may change over time because of application releases, seasonal demand, infrastructure migration, autoscaling, configuration changes, and business-process modifications. A machine learning model trained once and deployed permanently can therefore experience model drift and gradually lose detection effectiveness. The proposed framework addressed this limitation by incorporating feedback mechanisms through which new observations, analyst decisions, threat intelligence, and confirmed security events could influence subsequent model updates. Experimental reasoning around this mechanism indicated that adaptive models maintained more consistent detection capability when traffic and workload characteristics changed. The framework used a combination of rolling data windows, incremental model updating, threshold adaptation, and feedback-based learning to prevent the detection engine from becoming excessively dependent on historical behavior. An important finding was that adaptation should not occur without security controls. Completely unrestricted model updates could allow poisoned or misleading observations to influence the learning process, particularly when attackers deliberately generate deceptive traffic to manipulate the model. Consequently, the framework incorporated cybersecurity-aware validation before new observations were accepted into the learning pipeline. Events with uncertain labels could be temporarily isolated, while high-confidence security observations were prioritized for model refinement. This created a controlled learning loop in which adaptability was balanced with model integrity. Another important result was the reduction of false positives through contextual correlation. Individual indicators such as elevated CPU utilization, multiple failed logins, unusual network destinations, or increased API requests may not necessarily indicate attacks. However, when several indicators occur together and contradict an established behavioral baseline, their combined significance becomes considerably stronger. The framework therefore correlated signals across network, identity, workload, and application layers. This multi-dimensional analysis reduced dependence on isolated alerts and improved the prioritization of genuinely suspicious events. The discussion further indicates that contextual correlation is essential for cloud security because the distributed architecture produces a high volume of operational events, many of which are benign. An effective machine learning system must therefore identify relationships among events rather than simply count anomalies. The resulting approach supports more precise alert generation and allows security teams to concentrate on high-risk incidents instead of investigating large volumes of unrelated notifications.

The third set of results focused on detection performance across distributed cloud components. The framework demonstrated particular value in environments containing multiple cloud regions, heterogeneous platforms, containerized workloads, microservices, serverless functions, and dynamically changing API interactions. Traditional centralized detection architectures can introduce delays when telemetry must be transferred to a central security platform before analysis. In contrast, the proposed framework supported distributed intelligence in which selected preprocessing and inference activities could occur closer to the data source. Edge-oriented inference reduced the volume of raw telemetry requiring centralized transmission and contributed to lower detection latency for time-sensitive events. The architecture also enabled cloud-level aggregation of security intelligence so that local observations could be correlated with broader enterprise patterns. For example, an unusual login event occurring in one region might appear insignificant when analyzed independently, but its relationship with API misuse, abnormal privilege escalation, and unusual resource provisioning elsewhere could indicate a coordinated attack. The framework's distributed correlation capability therefore improved visibility across otherwise fragmented cloud domains. Another result involved scalability. As the number of workloads and telemetry sources increases, a security platform must process additional events without producing disproportionate computational overhead. The proposed architecture addressed this challenge by separating data collection, preprocessing, feature extraction, model inference, event correlation, and response orchestration into scalable components. Containerized deployment and elastic cloud resources allowed computational capacity to increase according to monitoring demand. Model inference could also be optimized by applying lightweight models to high-volume routine traffic while reserving computationally intensive analysis for high-risk or uncertain events. This hierarchical approach reduced unnecessary resource consumption while maintaining stronger analytical capability where required. The framework also supported explainability by associating detected threats with influential behavioral indicators. Rather than producing an opaque classification label alone, the system could identify contributing features such as abnormal access frequency, unusual destination patterns, suspicious privilege changes, or unexpected service-to-service

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

communication. This capability is important for security operations because analysts require evidence to validate alerts and determine appropriate responses. Explainable predictions also support governance and audit requirements by creating a traceable relationship between observed evidence, model reasoning, and security decisions.

The final results concern the practical security implications of combining machine learning with automated response mechanisms. The framework was designed to move beyond passive threat detection by linking high-confidence predictions with controlled response actions. Depending on risk severity, possible responses included session termination, credential verification, access restriction, API throttling, workload isolation, network segmentation, or escalation to security analysts. The results suggest that automation can significantly reduce the time between threat identification and containment, particularly for attacks that progress rapidly across distributed infrastructure. However, automated response also introduces operational risks. An incorrectly classified event could result in legitimate users or applications being blocked, potentially creating service disruption. The framework therefore applied risk-based response thresholds, allowing low-confidence predictions to generate alerts while high-confidence and high-impact events could trigger predefined containment actions. This distinction established a balance between security responsiveness and operational continuity. The discussion also highlights the importance of adversarial robustness. Attackers may intentionally manipulate input characteristics to evade machine learning detection, including modifying traffic rates, authentication sequences, API behavior, or payload characteristics. A cybersecurity-aware framework must consequently monitor model confidence, detect unusual input distributions, and periodically evaluate model performance against adversarial examples. Security testing and model validation should become continuous activities rather than one-time deployment procedures. Overall, the results demonstrate that machine learning can provide substantial advantages for real-time adaptive threat detection when it is integrated with cloud context, continuous learning, distributed analytics, explainability, and controlled response. The strongest contribution of the framework is not simply the application of a particular algorithm, but the integration of security awareness throughout the machine learning lifecycle. Data collection, feature engineering, model training, inference, feedback, adaptation, and response are treated as interconnected security functions. This integrated design provides a more resilient approach to cloud threat detection than isolated machine learning models. Nevertheless, performance depends on the quality and diversity of training data, the accuracy of security labels, the reliability of telemetry, computational resources, and the effectiveness of governance mechanisms. The framework should therefore be viewed as an adaptive decision-support and automated-defense architecture rather than a replacement for human security expertise. Its greatest value emerges when machine intelligence, security operations, contextual knowledge, and controlled automation operate together.

V. CONCLUSION

The proposed cybersecurity-aware machine learning framework provides an integrated approach for real-time adaptive threat detection across distributed cloud environments. The study establishes that modern cloud infrastructures require security mechanisms capable of processing large volumes of heterogeneous telemetry while adapting to continuously changing operational and threat conditions. Conventional signature-based and static rule-based systems remain useful for known attack patterns, but their effectiveness can decline when attackers introduce novel techniques or exploit rapidly changing cloud architectures. Machine learning provides an opportunity to identify complex behavioral relationships and detect deviations that are difficult to express through manually constructed rules. The proposed framework extends this capability by incorporating cybersecurity context into the learning process. Instead of treating threat detection as a conventional classification problem, it considers the interaction between network behavior, identity activity, workload characteristics, application events, access privileges, resource utilization, and historical security information. This contextual approach strengthens the distinction between legitimate cloud variability and genuine malicious behavior. The results and discussion demonstrate that a security-aware learning architecture can improve detection quality, reduce unnecessary alerts, and support more timely identification of suspicious activities. The framework is particularly relevant to distributed environments where workloads move dynamically across regions, platforms, containers, and infrastructure boundaries. Its adaptive design enables the detection process to evolve alongside changing application behavior and emerging attack patterns. Consequently, the research reinforces the argument that intelligent cloud security should be continuously adaptive rather than statically configured.

A major conclusion of the study is that effective cloud threat detection requires a combination of complementary machine learning techniques rather than dependence on a single algorithmic approach. Supervised learning provides effective recognition of previously observed and appropriately labeled attack classes, whereas anomaly detection contributes to the identification of previously unknown or behaviorally unusual events. Combining these capabilities creates a more comprehensive detection architecture capable of addressing both known and emerging threats. The framework further

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

demonstrates the value of continuous learning, but it emphasizes that adaptation must be governed carefully. Unrestricted learning can expose security models to data poisoning, incorrect labels, feedback manipulation, and adversarial influence. A trustworthy adaptive framework must therefore establish mechanisms for validating new data, evaluating model confidence, controlling update frequency, monitoring distribution changes, and maintaining rollback capabilities. The study also concludes that contextual correlation significantly improves threat prioritization. Individual events frequently lack sufficient evidence to establish malicious intent, particularly in highly dynamic cloud environments. Correlating identity, network, application, workload, and infrastructure signals provides a broader representation of the security situation and allows the detection engine to distinguish isolated operational anomalies from coordinated attack behavior. This capability is essential for addressing sophisticated threats that move laterally across cloud resources. The framework also contributes to security operations by providing interpretable evidence associated with predictions. Explainability enables analysts to understand why a particular event was classified as suspicious and facilitates validation, investigation, and remediation. This is especially important where security decisions may have significant operational consequences. Therefore, explainability should not be considered an optional visualization feature but an integral component of responsible machine learning-based cybersecurity.

The study further concludes that distributed processing is an important architectural requirement for real-time cloud security. Centralized analytics alone may introduce communication delays, increase data-transfer requirements, and create bottlenecks as telemetry volumes grow. By distributing selected preprocessing, feature extraction, and inference functions across cloud and edge locations, the proposed framework can reduce latency while retaining centralized coordination for broader security intelligence. This architecture supports scalable monitoring across geographically distributed workloads and heterogeneous cloud platforms. It also facilitates rapid response to high-risk events because selected decisions can occur close to the affected workload or service. However, distributed intelligence introduces its own challenges, including model synchronization, consistency, secure communication, heterogeneous infrastructure, and coordination among multiple detection nodes. The framework therefore highlights the importance of secure model distribution, authenticated telemetry, encrypted communication, and centralized governance of decentralized intelligence. Another conclusion concerns automated response. Integrating threat detection with response orchestration can substantially reduce the time required to contain high-confidence threats. Nevertheless, automated actions must be proportional to risk because false positives can disrupt legitimate business operations. A risk-based response strategy is therefore more appropriate than unconditional automation. High-confidence events associated with severe security consequences may justify automated containment, while uncertain predictions should be escalated for human validation. This human-machine collaboration provides a practical balance between response speed and operational safety. The overall architecture consequently supports a security model in which machine learning continuously observes the environment, evaluates risk, learns from validated feedback, and assists or performs controlled responses according to predefined security policies.

In conclusion, the cybersecurity-aware machine learning framework represents a comprehensive direction for strengthening security in distributed cloud computing. Its principal contribution lies in integrating adaptive machine learning with contextual cybersecurity intelligence, distributed analytics, explainable detection, continuous feedback, and risk-controlled response. The framework recognizes that cloud security is not a static problem and that detection models must continuously adapt to changes in infrastructure, applications, user behavior, workloads, and attack techniques. At the same time, adaptation must occur within a secure governance structure to prevent manipulation and preserve model reliability. The study therefore establishes several design principles for future intelligent cloud security systems: security context should influence feature engineering and model decisions; multiple detection approaches should be combined to address known and unknown threats; continuous learning should be validated and controlled; distributed inference should be used to improve scalability and latency; explanations should accompany important security decisions; and automated responses should be governed by risk and confidence. These principles provide a foundation for building more resilient enterprise cloud security platforms. Despite its advantages, the framework remains subject to limitations involving training-data quality, computational cost, evolving adversarial strategies, privacy constraints, model drift, interoperability, and deployment complexity. Its effectiveness must therefore be evaluated continuously using representative workloads and realistic attack scenarios. Future implementations should include large-scale experimental datasets, multi-cloud deployments, adversarial testing, real-time benchmarking, and longitudinal evaluation of model performance. Ultimately, the framework demonstrates that cybersecurity-aware machine learning can transform cloud threat detection from a largely reactive process into a proactive, adaptive, and intelligence-driven security capability. By combining automated analytics with human oversight and strong security governance, distributed cloud infrastructures can achieve improved visibility, faster threat identification, more precise risk prioritization, and stronger resilience against continuously evolving cyber threats.

International Research Journal of Innovative Engineering (IRJIE)

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

VI. FUTURE WORK

Future research should first focus on developing larger, more representative, and continuously updated datasets for cybersecurity-aware machine learning in distributed cloud environments. One of the major limitations of existing security datasets is that they often fail to represent the complexity of contemporary cloud infrastructures, particularly multi-cloud and hybrid-cloud environments containing containers, microservices, serverless applications, APIs, identity services, and dynamically scaled workloads. Future datasets should therefore integrate telemetry from multiple layers, including network flows, authentication records, API transactions, application logs, endpoint behavior, container activity, cloud control-plane events, and resource-management operations. The datasets should include both benign workload variation and sophisticated attack scenarios so that models can learn to differentiate operational changes from genuine threats. Research should also investigate federated and privacy-preserving learning techniques that allow multiple organizations or cloud domains to collaboratively improve security models without directly sharing sensitive raw telemetry. Such approaches could increase the diversity of training information while reducing privacy risks. Another important direction is the development of standardized benchmark datasets for adaptive cloud threat detection. These benchmarks should measure not only accuracy and F1-score but also detection latency, false-positive rate, model drift, resource consumption, robustness, explainability, and recovery performance. Longitudinal datasets would be particularly valuable because they could support evaluation of how detection models behave over extended operational periods. Researchers should also investigate mechanisms for generating realistic synthetic security events when real attack data are insufficient. Generative models may support the creation of diverse attack variations, although generated samples would require careful validation to prevent unrealistic patterns from degrading model performance. Establishing robust and representative datasets would provide a stronger scientific basis for comparing cybersecurity-aware learning architectures and determining whether reported performance improvements remain consistent across different cloud environments.

A second major direction should involve adversarially robust and trustworthy machine learning. As machine learning becomes increasingly integrated into security infrastructure, attackers will attempt to manipulate models through adversarial inputs, data poisoning, evasion strategies, and feedback exploitation. Future research should therefore develop detection models capable of recognizing deliberate attempts to alter their decision boundaries or exploit weaknesses in feature representations. Adversarial training, robust optimization, ensemble learning, uncertainty estimation, and secure model updating should be investigated as complementary mechanisms for strengthening resilience. Research should also examine how attackers could exploit continuous learning pipelines by injecting malicious observations that gradually alter behavioral baselines. Secure feedback validation and trusted learning channels will be essential for preventing such attacks. Another important research direction is explainable artificial intelligence for cybersecurity. Future frameworks should move beyond simply identifying influential features and should provide structured explanations describing the relationships among events, entities, behaviors, and threat indicators. Counterfactual explanations could help analysts understand what behavioral changes would have resulted in a different security classification. Causal approaches could further investigate whether observed features represent actual indicators of malicious activity or merely correlations caused by unrelated operational conditions. Human-centered security interfaces should also be developed so that analysts can interact with model explanations, validate predictions, provide feedback, and override automated actions when necessary. Research should evaluate whether explanations actually improve analyst decision-making, reduce investigation time, and decrease cognitive workload. Trust calibration will become increasingly important as automated security systems make more consequential decisions. Models should communicate uncertainty rather than presenting every prediction as definitive. Future systems could therefore combine confidence scores, evidence summaries, historical context, and recommended actions to support informed human decisions. These capabilities would make adaptive threat detection more transparent, accountable, and operationally trustworthy.

Future work should also investigate scalable distributed intelligence for multi-cloud, edge-cloud, and highly heterogeneous enterprise environments. As organizations increasingly distribute workloads across different cloud providers and geographical regions, security architectures must operate across infrastructure boundaries without creating isolated monitoring domains. Research should explore federated security orchestration in which multiple detection nodes collaboratively share security intelligence while maintaining local autonomy. Secure model synchronization will be important because distributed nodes may observe different traffic distributions and may therefore require localized models in addition to a global model. Techniques for hierarchical learning could enable local models to learn environment-specific behaviors while a central coordination layer captures broader enterprise-level patterns. Future studies should evaluate how model synchronization frequency, network conditions, workload distribution, and infrastructure heterogeneity influence detection performance. Edge intelligence is another important area because latency-sensitive applications may require security decisions to occur close to data sources. Lightweight machine learning

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

models could perform initial detection at edge locations, while more computationally expensive analysis could be performed in centralized cloud infrastructure. This hierarchical architecture could reduce latency and bandwidth consumption while preserving advanced analytical capabilities. Research should compare alternative deployment strategies based on detection accuracy, response time, energy consumption, computational overhead, and resilience to network disruptions. The framework could also be integrated with cloud-native technologies such as containers, Kubernetes-based orchestration, service meshes, serverless platforms, and infrastructure-as-code pipelines. Embedding security-aware machine learning into these environments could enable security policies to adapt automatically as workloads are deployed, scaled, migrated, or retired. Such integration would support a more comprehensive approach in which security intelligence becomes an intrinsic component of cloud infrastructure management rather than a separate monitoring function.

Finally, future research should extend the framework toward autonomous and predictive cyber defense while preserving strong governance and human oversight. Instead of detecting threats only after suspicious behavior becomes visible, advanced systems could use historical attack patterns, threat intelligence, behavioral trends, and infrastructure context to estimate future security risks. Predictive models could identify workloads, identities, APIs, or infrastructure components that are likely to become targets and recommend preventive controls before an incident occurs. Reinforcement learning could be investigated for adaptive security-policy optimization, allowing the system to select response strategies according to changing environmental conditions. However, reinforcement learning for cybersecurity requires strict safety constraints because exploration that is acceptable in ordinary optimization problems may cause unacceptable operational consequences in security environments. Future research should therefore develop constrained and policy-governed reinforcement learning models capable of optimizing security responses without violating availability, privacy, or compliance requirements. Integration with threat-intelligence platforms could further improve contextual awareness by correlating internal observations with external indicators and emerging attack campaigns. Digital twins of cloud infrastructures could also be used to simulate attacks and evaluate automated responses before deploying them in production. Such environments would provide a safe mechanism for testing new models, response policies, and adaptation strategies under controlled conditions. Another important direction is compliance-aware autonomous security, where machine learning decisions are evaluated against organizational policies, regulatory requirements, and data-governance constraints. Future frameworks should maintain detailed audit trails documenting observations, model versions, decisions, explanations, feedback, and automated responses. This would support accountability and enable forensic analysis following security incidents. Ultimately, future development should move toward self-adaptive cloud security platforms that continuously observe, learn, predict, explain, and respond while remaining bounded by explicit governance rules. The long-term objective should not be unrestricted autonomy, but trustworthy autonomy in which machine intelligence accelerates security operations while humans retain meaningful control over critical decisions. Such an approach could transform distributed cloud security into a continuously learning defense ecosystem capable of responding to increasingly complex and rapidly evolving cyber threats.

REFERENCES

1. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, *14*(1–2), 1–210.
2. Mathew, A. (2021). Artificial intelligence and cognitive computing for 6G communications & networks. International Journal of Computer Science and Mobile Computing, 10(3), 26-31.
3. Padmanabham, S. (2024). AI assisted fraud investigation architecture patterns for technology controls in financial institutions. International Journal of Research Publications in Engineering, Technology and Management, 7(3), 10587–10592.
4. Jayaraman, S., Rajendran, S., & P, S. P. (2019). Fuzzy c-means clustering and elliptic curve cryptography using privacy preserving in cloud. International Journal of Business Intelligence and Data Mining, 15(3), 273-287.
5. Koganti, H. (2021). Machine learning-driven performance anomaly detection and auto-tuning in distributed Java full-stack systems: A comprehensive review. International Journal of Research Publications in Engineering, Technology and Management, 4(3), 4977–4986.
6. Chellu, R. (2023). AI-Powered intelligent disaster recovery and file transfer optimization for IBM Sterling and Connect: Direct in cloud-native environments. International Journal on Recent and Innovation Trends in Computing and Communication, 11, 597.

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

7. Sasikumar, B., Venkatasalam, K., & Rajendran, P. (2024). Induction motor fault detection and classification using rcnn and surf based machine learning algorithms and infrared thermography. Scientia Iranica.
8. Bellundagi, M. (2023). Integrating Machine Learning with Business Rule Management Systems for Adaptive Enterprise. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(1), 8023-8039.
9. Kondapalli, K. K., Somajohassula, D. K., & Muppalla, L. K. (2022). Adaptive AI-orchestration and zero-trust security with federated threat intelligence for sustainable enterprise cloud architectures. International Journal of Computer Science and Engineering Research and Development (IJCSERD), 12(1), 176-192.
10. Soundappan, S. J. (2022). Modernizing Enterprise Software Ecosystems through Artificial Intelligence Driven Integration and Cloud Transformation. International Journal of Future Innovative Science and Technology (IJFIST), 5(5), 9253.
11. Badam, L. R. (2023). Explainable AI framework for real-time financial fraud detection in banking systems. International Journal of Emerging Trends in Engineering and Management Research, 8(2), 13241–13251.
12. Rajula, A. (2023). Event-driven enterprise applications with Apache Kafka and resilient cloud-native architecture. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, *6*(4), 9063-9073.
13. Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. Indian Journal of Science and Technology, 9, 44.
14. Abd-Rouf, M. S. K., Adigun, P. O., Alalade, E. O., Oyekanmi, T. T., Faniyi, A. J., Oladapo, B., Awopejo, T. E., Adegoke, O. S., Jamiu, A., Michael, O. B., Obisesan, A., Ajala, S., Adekanye, M. A., Yambali, P. M., & Abd-Rouf, A. B. (2024). From molecular profiling to predictive algorithms: A conceptual machine-learning framework for mechanism-informed therapy selection in multidrug-resistant cancer. International Journal of Science, Research and Technology (IJSRAT), 7(3), 12085–12101.
15. Sugumar, R. (2023). Enhancing Predictive Decision Intelligence in Enterprise Environments using Explainable AI and Cloud Computing. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(6), 9592-9602.
16. Anand, L. (2022). Integrating Kubernetes Microservices with Privileged Access Security and Real-Time Fraud Detection for Modern Enterprise Systems. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(5), 7453-7461.
17. Reddy, B. P. (2021). Optimizing smart grid performance using Machine Learning: A Data-Driven Approach to Energy Efficiency. J. Electrical Systems, 17(4), 149-156.
18. Tyagi, N. (2024). Artificial intelligence in financial fraud detection: A deep learning perspective. *International Journal of Computer Technology and Electronics Communication*, *7*(6), 9726-9732.
19. Gopinathan, V. R. (2023). Modernizing Enterprise Applications Using Intelligent API Frameworks Machine Learning and Cloud Native Computing. International Journal of Science, Research and Technology, 6(5), 10690-10697.
20. Praneeth, P. (2022). Prediction of Cost Overruns in Solar EPC Projects Using Machine Learning Techniques: A Data-Driven Study in India. International Journal of Engineering Science & Humanities, 12(2), 71-85.
21. Narra, R. (2024). A survey on scalable feature engineering techniques for cloud-native machine learning workflows. International Journal of Advanced Research in Science, Communication and Technology, 4(4), 664–677.
22. Hashmi, H., & Srikanth, V. (2023, November). Combining Neural Networks to Recognize Offline Digits & Words by Training the Model Using Keras. In 2023 3rd International Conference on Advancement in Electronics & Communication Engineering (AECE) (pp. 694-697). IEEE.
23. Soundappan, S. J. (2022). Integrated Risk Governance Framework for Financial Compliance Supply Chain Resilience and Enterprise Data Management. International Journal of Computer Technology and Electronics Communication, 5(6), 16254-16263.
24. Vimal Raja, G. (2024). Intelligent data transition in automotive manufacturing systems using machine learning. International Journal of Multidisciplinary and Scientific Emerging Research, 12(2), 515-518.
25. Matrouk, K., V, S., Kumar, S., Bhadla, M. K., Sabirov, M., & Saadh, M. J. (2023). Deep Learning–based Dynamic User Alignment in Social Networks. ACM Journal of Data and Information Quality, 15(3), 1-26.
26. Polamarasetty, V. K. (2022). Enterprise SAP application modernization for post-acquisition business transformation. International Journal of Science, Research and Technology (IJSRAT), 5(3), 7777–7783. https://www.ijsrat.com/index.php/ijsrat/issue/view/23
27. Alam, A., Gazi, M. S., Abdullah, S. M., Tasnim, M., Himeluzzaman, M., Nabil, M. A., Akter, K. A., & Akter, S. (2021). Quantum-resilient federated intrusion detection: A hybrid quantum classical framework for safeguarding

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805007

U.S. critical infrastructure in the post-quantum era. International Journal of Advances in Signal and Image Sciences, 7(1), 57–72. https://doi.org/10.29284/3xx46j76

28. Gujarathi, M. (2022). Rule externalization for enterprise validation platforms: Architecture and design considerations. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, *5*(4), 7163-7167.
29. Bandaru, P. K. (2022). Hardware-in-the-loop testing for connected vehicles: Enhancing software reliability through continuous validation. International Journal of Engineering & Extended Technologies Research (IJEETR), 4(2), 4645–4651.
30. Kundavaram, R. R., Bandhela, R. R., & Onteddu, A. R. (2022). AI-driven predictive modeling in healthcare: A data science perspective on U.S. healthcare data. South Eastern European Journal of Public Health. https://doi.org/10.70135/seejph.vi.6691
31. Mudunuri, L. N. R. (2024). Artificial intelligence (AI) powered matchmaker: finding your ideal vendor every time. *FMDB Transactions on Sustainable Intelligent Networks*, *1*(1), 27-39.
32. Jain, R. (2021). The vector database gap: A production blueprint for embedding-based enterprise search. International Journal of Future Innovative Science and Technology (IJFIST), 4(4), 6706–6713.
33. Raja, G. V. (2023). AI-Driven Cloud-Native Enterprise Systems Leveraging Kubernetes, DevSecOps, and Predictive Analytics. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(2), 7925-7929.
34. AbdulRahman, S., Tout, H., Ould-Slimane, H., Mourad, A., Talhi, C., & Guizani, M. (2021). A survey on federated learning: The journey from centralized to distributed on-site learning and beyond. *IEEE Internet of Things Journal*, *8*(7), 5476–5497.