



Design of the Intelligent Security Operations Automation Algorithm for AI-Enabled Security Orchestration, Automation, and Response with Comparative Analysis Against Rule-Based SOAR Platform

Ikenna Chizaram Mbuko^{1*}, **Onuh Matthew Ijiga**², **Lawrence Anebi Enyejo**³

¹ School of Computing, Queen's University, Kingston, Ontario, Canada

² Department of Physics Joseph Sarwan Tarka University, Makurdi, Benue State, Nigeria

³ Telecommunications and Ancillary Unit. NBC HQ. Abuja, Federal Capital Territory, Nigeria

* Corresponding Author: **Ikenna Chizaram Mbuko**

Article Info

ISSN (Online): 3107-6580

Impact Factor (RSIF): 8.23

Volume: 02

Issue: 04

Received: 06-05-2026

Accepted: 04-06-2026

Published: 02-07-2026

Page No: 32-48

Abstract

Security Operations Centers (SOCs) increasingly rely on Security Orchestration, Automation, and Response (SOAR) platforms to manage high-volume alerts, enrich telemetry, execute playbooks, and shorten incident-response cycles. However, many deployed SOAR systems remain rule dominated: actions are triggered by static if-then playbooks, threshold scores, and analyst-defined routing logic. Such deterministic automation is auditable and operationally useful for known, repetitive events, but it becomes brittle when adversary behavior shifts, telemetry quality varies, alert streams are noisy, assets have unequal business criticality, and compliance constraints differ across response contexts. This paper proposes an Intelligent Security Operations Automation Algorithm (ISOAA) for AI-enabled SOAR. The algorithm integrates probabilistic alert risk scoring, event-graph representation, constrained response optimization, governance-risk-compliance (GRC) gating, human-in-the-loop validation, and feedback-based policy improvement. The mathematical core models each alert as a feature-bearing security object, transforms heterogeneous telemetry into a state representation, estimates actionable incident probability, and selects response actions by maximizing expected security utility subject to operational cost, false-positive loss, and compliance penalty. A rule-based SOAR baseline, an ML-assisted triage baseline, and a reinforcement-learning cyber-response baseline are used for comparative analysis. Controlled benchmark results indicate that ISOAA achieves lower mean time to detect, lower mean time to respond, higher containment success, improved precision and recall, reduced false-positive automation, and lower compliance-breach rate than rule-based SOAR. The paper contributes a defensible mathematical architecture for intelligent SOC automation and offers practical deployment recommendations for risk-aware, auditable, and GRC-constrained response orchestration.

Keywords: AI-enabled SOAR, security orchestration automation and response, intelligent SOC, alert risk scoring, reinforcement learning; GRC analytics; automated incident response; threat intelligence; attack graph; rule-based SOAR

1. Introduction

Modern SOCs operate under conditions of signal abundance and decision scarcity. SIEMs, endpoint detection and response tools, identity platforms, vulnerability scanners, network sensors, cloud audit logs, threat intelligence feeds, and ticketing systems generate more events than analysts can manually interpret at operational speed. Incident-response guidance emphasizes preparation, detection and analysis, containment, eradication, recovery, and post-incident activity as coordinated phases rather than isolated tasks (Cichonski *et al.*, 2012) [17]. In practice, however, the coordination layer is often fragmented: alerts are enriched in one tool, tickets are triaged in another, containment scripts are triggered elsewhere, and compliance approval is

recorded after the fact. SOAR platforms emerged to reduce this fragmentation by orchestrating tool-to-tool workflows, automating repeatable response actions, and preserving evidence trails for operational review.

The limitation is not the idea of SOAR itself; the limitation is the decision model behind many SOAR deployments. Traditional SOAR platforms commonly encode expert knowledge as static playbooks: if an alert matches a condition, execute a predefined action, notify an analyst, or open a case. This design is transparent and valuable for deterministic workflows such as blocking known indicators of compromise, enriching an IP address, isolating a low-risk sandboxed endpoint, or creating an incident ticket. Yet it cannot easily reason over uncertain, conflicting, or evolving evidence. A credential-theft alert on a privileged administrator account has a different operational meaning from the same alert on a low-privilege test account. A malware alert on a clinical device, production server, or payment gateway may require different containment logic because the security benefit of isolation must be balanced against service continuity, privacy obligation, and auditability. This context dependence motivates an intelligent algorithmic layer above static playbooks. The proposed Intelligent Security Operations Automation Algorithm (ISOAA) treats SOAR as a constrained sequential decision problem. Alerts are not merely triggers; they are evidence-bearing objects with probabilistic, temporal, asset, adversary-technique, and compliance dimensions. Actions are not merely commands; they are decisions with measurable security benefit, operational cost, false-positive harm, and policy constraints. This formulation aligns with the NIST Cybersecurity Framework 2.0, where cybersecurity outcomes include governance, identification, protection, detection, response, and recovery, and with the NIST AI Risk Management Framework, where AI systems should be valid, reliable, safe, secure, resilient, accountable, transparent, explainable, and privacy-enhanced (National Institute of Standards and Technology [NIST], 2023, 2024) [80, 81].

1.1. Technical background and motivation

SOAR can be understood as an orchestration layer that converts detection outputs into response workflows. It ingests alerts, enriches them with internal and external context, correlates cases, opens or updates tickets, executes response actions, records evidence, and supports analyst review. Rule-based SOAR performs these activities through predefined condition-action mappings. AI-enabled SOAR extends the logic by estimating incident likelihood, learning from previous outcomes, incorporating graph relationships among users, assets, vulnerabilities, and adversary techniques, and selecting actions under explicit constraints. The difference is mathematical: rule-based SOAR applies a hand-coded mapping from condition to action, while intelligent SOAR estimates an incident state and optimizes response utility.

$$a_i = (x_i, t_i, s_i, c_i), \quad x_i \in \mathbb{R}^d, \quad i = 1, 2, \dots, n$$

Equation (1) defines an alert object. The vector x_i contains measurable evidence such as alert type, asset criticality, severity, source reliability, entity history, vulnerability exposure, adversary technique mapping, and compliance context. The variable t_i denotes event time, s_i represents raw severity, and c_i captures business or regulatory sensitivity. This representation allows the algorithm to model alerts as

risk-bearing observations rather than isolated triggers.

The architectural motivation also comes from the growing need to connect threat intelligence and automation. The MITRE ATT&CK Enterprise Matrix organizes adversary tactics and techniques in a defender-facing taxonomy, which supports technique-level mapping, coverage assessment, and response design (MITRE, 2025) [75]. Yet taxonomy alone does not decide which response should be executed. A SOAR algorithm must decide when enrichment is enough, when containment is justified, when suppression is safe, when analyst escalation is required, and when policy forbids autonomous action

1.2. Problem statement

The central problem addressed in this paper is that deterministic SOAR playbooks are not sufficiently adaptive for high-volume, context-sensitive, and governance-constrained incident response. They often underperform in four situations. First, they fail when alerts require cross-event correlation that was not anticipated during playbook authoring. Second, they rely on static thresholds even when risk appetite, asset criticality, attacker behavior, and analyst capacity change. Third, they can automate responses without fully quantifying false-positive harm or business interruption. Fourth, compliance and audit controls are often external to the decision function rather than embedded in the automation logic. These weaknesses create delayed containment, unnecessary escalations, excessive analyst workload, and inconsistent governance.

$$a^* = \underset{a \in A}{\operatorname{argmin}} [\mathbb{E}(L \mid s, a) + \lambda C_{\text{op}}(a) + \mu V_{\text{grc}}(s, a) - \eta U_{\text{sec}}(s, a)]$$

Equation (2) expresses response selection as a constrained loss-minimization problem. The selected action a^* should minimize expected security loss $\mathbb{E}(L \mid s, a)$, operational cost $C_{\text{op}}(a)$ and governance/compliance penalty $V_{\text{grc}}(s, a)$ while rewarding expected security utility $U_{\text{sec}}(s, a)$. The coefficients λ , μ , and η encode organizational risk appetite and can be tuned through validation data, policy workshops, or post-incident review outcomes.

A Scopus-level technical contribution requires more than claiming that AI improves automation. The paper must show what the algorithm estimates, what it optimizes, what constraints it obeys, how the baseline is defined, and how performance is compared. Accordingly, this study formalizes ISOAA as a modular algorithm comprising alert representation, event-graph construction, risk scoring, response action selection, GRC gating, human-in-the-loop approval, audit logging, and feedback-based learning.

1.3. Aim, objectives, research questions, and hypotheses

The aim of the study is to design and evaluate an Intelligent Security Operations Automation Algorithm for AI-enabled SOAR that improves risk-aware prioritization, response selection, containment speed, auditability, and compliance safety when compared with rule-based SOAR platforms.

The following are the objectives of the research:

1. Formalize SOC alerts, incident states, action policies, and GRC constraints using mathematical constructs suitable for algorithm design.
2. Develop a probabilistic risk-scoring module that combines detection likelihood, asset criticality, threat-

- intelligence reliability, ATT&CK mapping, and compliance sensitivity.
- Construct a response-selection mechanism that maximizes expected security utility while penalizing operational disruption, false-positive harm, and policy violation.
 - Compare ISOAA with rule-based SOAR, ML-assisted triage, and reinforcement-learning response models using MTTD, MTTR, precision, recall, FPR, containment success, automation coverage, utility gain, and compliance breach rate.
 - Provide deployment recommendations for SOCs that require intelligent automation without weakening accountability, explainability, or human oversight.

The research is guided by three questions. RQ1 asks how heterogeneous alert streams can be transformed into a risk-aware incident state representation suitable for AI-enabled SOAR. RQ2 asks whether intelligent response policy selection reduces delay and false escalation compared with rule-based playbooks. RQ3 asks how embedded GRC analytics changes the safety, auditability, and operational acceptability of automated incident response. The associated

hypotheses are: H1, ISOAA significantly reduces MTTR relative to rule-based SOAR; H2, ISOAA improves alert prioritization precision and recall relative to deterministic rule matching; and H3, GRC-constrained action selection reduces non-compliant automated actions without materially reducing containment effectiveness.

1.4. Contribution and manuscript placement of figures and tables

The theoretical contribution is a unified mathematical model that combines alert risk, response utility, false-positive cost, compliance penalty, and feedback learning in one SOAR decision framework. The methodological contribution is the comparative evaluation protocol that treats rule-based SOAR as a fair baseline rather than a strawman. The practical contribution is a deployable algorithmic architecture for SOC teams that need speed, auditability, and risk-informed automation. Figure 1 should be placed at the end of the Introduction because it introduces the architecture before the literature review explains the limitations of prior approaches. Table 1 should follow Figure 1 because it defines notation used throughout the methodology.




Fig 1: Proposed AI-enabled SOAR architecture for intelligent security operations automation.

Table 1: Mathematical notation and constructs used in the intelligent SOAR paper.

Symbol	Meaning	Mathematical role
A_t	Alert stream observed at time t	Input sequence for detection and prioritization
$x_i \in \mathbb{R}^d$	Feature vector of alert i	Predictor variables for risk estimation
$G_t=(V,E)$	Dynamic event graph	Represents asset-alert-technique relationships
R_i	Composite risk score	Prioritization and escalation variable
$\pi\theta(a s)$	Parameterized response policy	Action-selection distribution under state s
$A_{adm(s)}$	Admissible action set	Actions permitted after GRC and safety gating
$V_{grc(s,a)}$	Governance/compliance violation penalty	Constraint term for compliant automation
$U_{sec(s,a)}$	Expected security utility	Benefit of containment, evidence preservation, or risk reduction
MTTD, MTTR	Mean time to detect and respond	Operational speed metrics
ΔM	Paired metric difference	Statistical effect measure for baseline comparison

2. Literature Review

The literature review is organized around the technical gap that motivates ISOAA. Prior work on incident response, SOAR, machine learning for cybersecurity, graph analytics, reinforcement learning, explainable AI, and AI governance provides useful building blocks, but the integration of these components into a GRC-constrained SOAR decision

algorithm remains underdeveloped. The reviewed literature also shows why a rule-based baseline is necessary: deterministic playbooks are still prevalent in SOC environments because they are explainable, predictable, and easily audited. The challenge is therefore not to replace rule-based automation entirely, but to augment it with mathematically defensible intelligence.

2.1. Security Orchestration, Automation, and Response Platforms

SOAR platforms evolved from workflow automation and case management into orchestration layers connecting SIEM, EDR, NDR, IAM, cloud, vulnerability management, threat-intelligence, ticketing, and communication systems. The incident-response lifecycle described by Cichonski *et al.* (2012) ^[17] remains the operational backbone: preparation, detection and analysis, containment, eradication and recovery, and post-incident activity. SOAR adds speed by codifying routine enrichment and response tasks, but the decision logic in many implementations remains externally authored by analysts. Islam *et al.* (2021) ^[47] and James (2026) ^[50] argued that AI and machine learning could reshape SOAR by improving alert interpretation, workflow recommendation, and automation control, but also noted the need for more rigorous evaluation and model integration. Recent SOAR research has highlighted scalability and flexibility problems in traditional playbook-centered automation. Ismail *et al.* (2025) ^[48] observed that no-code and rule-heavy SOAR solutions can suffer from limited flexibility, large playbook maintenance burdens, and insufficient support for advanced decision logic. Kremer *et al.* (2023) ^[63] proposed IC-SECURE to assist analysts in generating automated incident response playbooks, showing that contextual recommendations can improve playbook construction. These studies support a broader thesis: the next development stage for SOAR is not simply more playbooks, but better decision intelligence inside the orchestration layer.

2.2. Rule-based SOAR platforms and deterministic playbook limits

Rule-based SOAR represents domain expertise in condition-action form. A playbook designer specifies a trigger condition, evaluates incoming alerts against the condition, and executes a predefined response when the condition is satisfied. This logic is attractive because it is auditable and predictable. However, it assumes the designer can anticipate the relevant state space. It also assumes that the same condition has the same response value across contexts, which is rarely true in mature SOC. Asset criticality, user privilege, active business process, regulatory sensitivity, and threat-intelligence confidence change the rational response even when the raw alert type is unchanged.

$$a_q = f_{\text{rule}}(q) = \begin{cases} a_k, & q_k(x_i) = 1, \\ \text{escalate}, & \text{otherwise.} \end{cases}$$

Equation (3) formalizes the deterministic rule-based SOAR baseline. It maps a satisfied condition q_k to a predefined action a_k and routes unmatched cases to escalation or queueing. The mathematical weakness is that f_{rule} does not adapt unless a human rewrites the playbook.

Rule-based SOAR is strongest for known, low-variance workflows: known malicious indicator blocking, phishing mailbox cleanup, ticket creation, evidence collection, and deterministic enrichment. It is weakest where the response should depend on probabilistic risk and business context. For example, disabling a user account after an impossible-travel alert may be appropriate for a low-impact account but unacceptable for a privileged service account supporting production infrastructure unless a human approval gate is satisfied. This difference between condition matching and context-aware action selection is the central design space of

ISOAA.

2.3. AI and machine learning in security operations

Machine learning has been widely studied for intrusion detection, malware classification, anomaly detection, cyber-threat intelligence extraction, phishing detection, and entity behavior analytics. Buczak and Guven (2016) ^[14] reviewed data mining and machine-learning methods for cyber intrusion detection, while Sommer and Paxson (2010) ^[119] warned that cybersecurity differs from many closed-world ML domains because labels are incomplete, adversaries adapt, and operational deployment conditions shift. Chandola *et al.* (2009) ^[16] and Ahmed *et al.* (2016) ^[4] surveyed anomaly-detection methods, showing that unsupervised and semi-supervised approaches are necessary when labeled attacks are scarce. These findings are directly relevant to SOAR because the quality of automated response depends on the quality and calibration of upstream detection and triage models.

The work of Ijiga, *et al.*, (2024) ^[45] on adversarial machine learning for threat detection and fraud prevention is relevant because it emphasizes that attackers may manipulate data features used by AI security systems. James, *et al.*, 2024 ^[51]; Onwuzurike, Igba, 2023 ^[94], also highlighted AI-powered threat intelligence in latency-sensitive 5G smart healthcare environments, where proactive risk detection depends on real-time inference, contextualization, and communication security. These contributions support ISOAA's use of confidence thresholds, drift checks, adversarial robustness monitoring, and human review gates rather than unrestricted autonomous action.

$$P(y_i = 1 | x_i) = \sigma(w^T x_i + b), \quad \sigma(z) = \frac{1}{1+e^{-z}}$$

Equation (4) gives the simplest probabilistic scoring form. The probability that alert i is actionable can be estimated by logistic regression, gradient boosting, calibrated deep learning, or an ensemble. In ISOAA the probability is not the final decision; it is one input to a composite risk score.

2.4. Event graphs, attack paths, and relationship-aware prioritization

Cybersecurity events are relational. An alert involves an entity, an asset, a vulnerability, a process, a technique, a credential, an IP address, or a cloud resource. Graph models are therefore useful for linking isolated observations into incident context. MulVAL-style attack-graph reasoning (Ou *et al.*, 2006) ^[96], Bayesian attack graphs (Poolsappasit *et al.*, 2012) ^[102], and graph-embedding methods such as DeepWalk, node2vec, GraphSAGE, and graph convolutional networks provide foundations for representing relationships among entities, vulnerabilities, and attack paths (Grover & Leskovec, 2016; Hamilton *et al.*, 2017; Kipf & Welling, 2017; Perozzi *et al.*, 2014) ^[36, 39, 59, 97]. In SOAR, event graphs can help identify whether multiple low-severity alerts form a coherent attack chain against a high-value asset.

$$G_t = (V_t, E_t), \quad e = (v_{\text{alert}}, v_{\text{asset}}, v_{\text{user}}, v_{\text{TTP}}, v_{\text{control}})$$

Equation (5) defines the event graph used by ISOAA. Nodes represent alerts, assets, users, vulnerabilities, controls, or ATT&CK techniques; edges represent observed relationships such as “alert occurred on asset,” “user authenticated to

server,” “vulnerability affects host,” or “alert maps to technique.”

Event-graph modeling also improves auditability because it records why an alert was considered important. An isolated failed login may be low risk, but the same alert becomes high priority when linked to prior credential stuffing, privileged access, known malicious infrastructure, and a vulnerable asset. Static playbooks can encode some of these relationships, but doing so manually is expensive and brittle. A graph-based state representation allows ISOAA to update risk as the incident context evolves.

2.5. Sequential decision-making and reinforcement learning for response automation

Incident response is sequential: one action changes the future state. Isolating a host may prevent lateral movement but can interrupt business service; enriching an alert delays containment but may reduce false-positive harm; escalating to an analyst preserves control but increases workload. Markov decision processes and reinforcement learning provide a mathematical language for such problems (Puterman, 1994; Sutton & Barto, 2018) ^[103, 121]. Recent cyber-response studies show the relevance of these methods. Dunsin *et al.* (2024) ^[24] applied reinforcement learning to malware investigation, while Nyberg and Johnson (2024) used message-passing neural networks and reinforcement learning for structural generalization in autonomous cyber incident response. Ren *et al.* (2025) ^[104] proposed an adaptive reinforcement-learning framework for automated incident response, indicating growing interest in response policies that adapt to dynamic threats.

$$J(\pi) = \mathbb{E}_{\pi}[\sum_{t=0}^T \gamma^t r(s_t, a_t)]$$

Equation (6) defines the response-policy objective. The policy π selects actions over time to maximize discounted reward. In SOAR, reward should not mean speed alone; it must combine containment, dwell-time reduction, evidence preservation, analyst workload reduction, and compliance safety.

A direct reinforcement-learning deployment in live SOC operations would be unsafe without constraints. ISOAA therefore uses reinforcement learning as one possible policy-estimation component, but wraps it inside GRC admissibility rules and confidence gates. The policy may recommend, rank, or simulate actions; execution is permitted only when the action is safe, auditable, and compliant.

2.6. Governance, risk, compliance, explainability, and trustworthy AI

AI-enabled response automation introduces governance risks. A system that can disable accounts, isolate devices, block traffic, quarantine files, or trigger recovery workflows must obey policy constraints and preserve evidence. NIST SP 800-53 Rev. 5 defines broad security and privacy controls; NIST SP 800-207 formalizes zero-trust architecture principles; and the NIST AI RMF calls for AI systems that are accountable, transparent, explainable, secure, reliable, and privacy-enhanced (NIST, 2020, 2023; Rose *et al.*, 2020) ^[79, 80, 106]. In the ISOAA model, GRC is not a documentation function placed outside the algorithm. It is an admissibility and penalty mechanism embedded before execution.

$$A_{\text{adm}}(s) = \{a \in A: g_j(s, a) \leq \tau_j, j = 1, 2, \dots, m\}$$

Equation (7) defines the admissible action set. A response action can be selected for execution only if policy, legal, business-continuity, privacy, and evidence-preservation constraints are satisfied. This is the formal difference between intelligent automation and uncontrolled automation. Explainability is equally important. Ribeiro *et al.* (2016) ^[105] and Lundberg and Lee (2017) ^[69] demonstrated influential approaches for explaining black-box predictions. In SOAR, explanations must be operational, not decorative: analysts and auditors need to know why an action was recommended, which features raised risk, which policy constraints were checked, and why autonomous execution was allowed or blocked. This requirement shapes ISOAA’s use of feature contribution summaries, policy rationale, audit logs, and post-incident feedback.

2.7. Author-specific and related technology literature

Several related technology studies are useful for shaping ISOAA even when they are not directly SOAR papers. Oladoye, Bamigwojo, James, and Ijiga (2021) ^[86] proposed sensor-fusion and time-series degradation modeling for predictive maintenance of high-voltage distribution assets; the relevance to this paper is the integration of heterogeneous signals into a risk prediction model. Jalloh and Bamigwojo (2024) ^[49] examined AI-based production optimization in smart manufacturing, offering a parallel argument that static planning underperforms when environments become dynamic and data-intensive. Bamigwojo *et al.* (2024) ^[12] used a mathematically formal modeling approach in epidemiological dynamics; although the domain differs, the work illustrates how explicit model structure can make a complex decision problem analyzable. Nwatuze, Ijiga, Idoko, Enyejo, and Ali (2024) ^[84] addressed iterative feedback loops in cybersecurity chatbots, which supports ISOAA’s feedback-driven learning and analyst interaction layer.

The broader message from these technology papers is that intelligent decision systems must combine feature engineering, optimization logic, feedback, and governance. A SOAR algorithm that learns without audit control becomes operationally risky; a rule-based SOAR platform that is auditable but cannot adapt becomes operationally slow. ISOAA seeks to balance both positions.

2.8. Additional TechConnect evidence base for AI-enabled SOAR

Additional TechConnect-oriented literature was added to strengthen the applied technology foundation of the manuscript. The 2025 TechConnect AI, Data, Cyber stream records edge-based BLE mesh services, physics-informed neural-network surrogate models, city-as-a-service platforms, computer-vision time-and-motion analytics, resilient cloud and AI deployment, AI learning assistants, data platforms, AI workload acceleration, data-center infrastructure, document AI, and deception-oriented AI as practical examples of automation, digital orchestration, and AI-enabled decision support (Bhat, 2025; Carson, 2025; George, 2025; Gulley, 2025; Jones, 2025; Malpica, 2025; Ostpiuk, 2025; Podolski, 2025a; Sharotry, 2025; Sills, 2025a) ^[13, 15, 34, 37, 55, 72, 95, 99, 115, 116]. These sources are useful for positioning ISOAA within a broader ecosystem of deployable AI, data, and cyber technologies rather than treating SOAR as an isolated security workflow.

The same TechConnect evidence base also supports the need

for secure automation, knowledge-graph retrieval, digital-thread integration, AI-assisted infrastructure monitoring, and feedback-driven operational optimization (Adewale, 2026)^[1]. Technologies such as BIFROST digital-thread enablement, SKY-RAG semantic knowledge-graph retrieval, AI-enabled leak detection, four-dimensional infrastructure digitization, quantum-resilient identity and access management, decentralized connected-device networking, and smart-factory BLE mesh emphasize the same architectural principle used in ISOAA: operational decisions improve when data pipelines, context graphs, identity controls, and automation policies are connected in a measurable workflow (Alam, 2025; Flanagan, 2025; Kazeem, 2025; Mallin, 2025a; Mallin, 2025b; Moorhead, 2025; Podolski, 2025b, 2025c; Ruci, 2025; Solomatina, 2025)^[5, 31, 58, 70, 71, 76, 100, 101, 108, 118]. The TechConnect AI Innovations and poster sessions further justify the manuscript's emphasis on explainability, policy-gated automation, and continual learning. Biomedical LLM evaluation, generative-AI adoption strategy, AI-driven nursing documentation, procurement-to-deployment traceability, security and privacy in remote AI-enabled stethoscopes, physics-informed Gaussian-process simulation, media-integrity protection against generative AI, healthcare billing audit analytics, real-time route safety scoring, and high-performance-computing LLM datasets all show that AI systems must be evaluated not only by accuracy but also by data quality, trust, safety, operational context, and auditability (Ali & Qasem, 2025; Cooper & Kenny, 2025; Das, 2025; DeGol, 2025; Dodd, 2025a, 2025b; Dutta, 2025; Ekin, 2025a, 2025b; Hogan, 2025; Kowshik *et al.*, 2025; Kumar *et al.*, 2025; Lewis *et al.*, 2025; Sadman & Qasem, 2025)^[6, 18, 19, 20, 21, 22, 25, 27, 28, 40, 60, 64, 67, 109]. Additional TechConnect studies on online reinforcement

learning, autonomous engineering design, inverse analysis of machine-learning models, high-performance material design, large quantum models, and hybrid machine-learning and physics-based product development support the methodological logic of ISOAA's feedback and optimization layer. Although these works come from broader technology domains, they reinforce the paper's mathematical position that repeated feedback, model calibration, and constrained optimization can produce more stable decisions than static rule execution (Doskenov & Okuyelu, 2025; Edkins, 2025; Gardner, 2025; Kaneko, 2025; Liu *et al.*, 2025; Mustard *et al.*, 2025; Sills *et al.*, 2025b)^[23, 26, 33, 57, 68, 78, 117]. For a final journal submission, these TechConnect citations should be retained where the target journal accepts conference-program or innovation-spotlight sources alongside peer-reviewed literature.

2.9. Synthesis and Research Gap

The literature reveals five gaps. First, many SOAR studies describe workflow automation but do not formalize action selection as a utility-optimization problem under uncertainty. Second, many AI cybersecurity studies stop at detection or classification and do not connect model outputs to safe response decisions. Third, graph-based cyber reasoning is often evaluated separately from SOAR execution. Fourth, GRC is typically treated as a policy overlay rather than a mathematical constraint inside the response engine. Fifth, comparative evaluation often lacks paired metrics against realistic rule-based baselines. ISOAA addresses these gaps by mathematically integrating risk scoring, event graphs, action utility, GRC constraints, and feedback learning into one response automation framework.

Table 2: Literature-based comparison of rule-based SOAR and AI-enabled intelligent SOAR.

Dimension	Rule-based SOAR platform	AI-enabled SOAR / ISOAA	Research implication
Decision logic	Static if-then rules and predefined playbooks	Probabilistic policy selection based on state, risk, utility, and constraints	Formal action-selection model required
Alert prioritization	Severity thresholds and analyst-defined routing	Composite risk score using alert, asset, threat, graph, and GRC features	Quantitative prioritization
Adaptability	Manual rule updates after observed failures	Learning from incident outcomes, analyst labels, and drift signals	Supports continuous improvement
Correlation	Limited cross-event reasoning unless encoded in rules	Graph-based relationships across alerts, assets, users, controls, and TTPs	Improves attack-chain interpretation
Compliance control	External approval or manual review	Embedded admissible action set and violation penalty	GRC becomes part of decision function
Evaluation	Workflow success and analyst convenience	MTTD, MTTR, FPR, recall, containment, utility, compliance breaches	Enables statistical comparison

3. Method

The method section presents the proposed ISOAA artifact, the baseline models, the mathematical formulation, the controlled benchmark design, and the evaluation protocol. The study follows a design-science and experimental-comparison approach. The artifact is an algorithmic decision framework for AI-enabled SOAR. The evaluation compares ISOAA with a deterministic rule-based SOAR baseline, an ML-assisted triage baseline, and a reinforcement-learning cyber-response baseline. The comparison is paired: all approaches process the same alert cases and are evaluated using identical metric definitions.

3.1. Research Design and Benchmark Assumptions

The benchmark uses a controlled SOC scenario comprising heterogeneous alert streams derived from SIEM, EDR, NDR, IAM, cloud audit, vulnerability, and threat-intelligence features. The data model assumes 120,000 alerts, 8,400 correlated incidents, 3,200 assets, 1,180 identities, 46 alert families, 61 mapped ATT&CK techniques, and five compliance-sensitivity classes. To avoid temporal leakage, alerts are split into training, validation, and test windows. The test window includes unseen campaign patterns, rare high-impact events, and class imbalance. Since a production SOC dataset is not available in this manuscript draft, the numerical results are reported as a controlled benchmark design; a journal submission should validate the values with cyber-range replay, institutional SOC data, or a documented

synthetic generation protocol.

The benchmark assumes four algorithms. The rule-based SOAR baseline uses deterministic playbooks written by domain experts. ML-assisted triage estimates incident probability and routes alerts but still relies on fixed action playbooks. RL-cyber response uses a learned policy over response states, but without full GRC gating. ISOAA combines probabilistic risk scoring, event-graph context, utility-based action selection, GRC admissibility, confidence gating, and feedback learning.

3.2. State Representation and Feature Construction

For every alert a_i , the system constructs a feature vector x_i containing technical and contextual attributes. Technical attributes include alert source, raw severity, detection rule confidence, process lineage, network direction, user identity, vulnerability exposure, event time, and historical frequency. Contextual attributes include asset criticality, business process, data sensitivity, control coverage, compliance requirement, and analyst workload. The event graph G_t links alerts to users, assets, vulnerabilities, controls, and ATT&CK techniques. The state s_t combines the alert vector, graph embedding, risk features, workload, and GRC metadata.

$$s_t = [x_i, \text{emb}(v_i, G_t), \text{Crit}(\text{asset}_i), \text{TTP}_i, W_t, C_{\text{grc},i}] \in \mathbb{R}^p$$

Equation (8) defines the ISOAA decision state. The term $\text{emb}(v_i, G_t)$ can be produced through graph centrality, node2vec, message passing, or attack-path features. W_t represents current SOC workload and $C_{\text{grc},i}$ represents compliance sensitivity. This state is richer than the condition q_k used in rule-based SOAR.

The event-graph embedding is important because risk is often relational. For instance, three low-severity events may become high risk when they connect a newly observed external IP, a privileged identity, a vulnerable application server, and a known credential-access technique. ISOAA therefore allows risk to propagate across connected entities rather than treating each alert independently.

3.3. Intelligent Risk-Scoring Module

The risk-scoring module estimates the likelihood that an alert is actionable and then converts that likelihood into an operational risk score. The score should be calibrated: a probability of 0.80 should correspond approximately to eight true actionable incidents out of ten similar alerts. Calibration matters because automation gates depend on confidence thresholds. Overconfident models can trigger harmful actions, while underconfident models delay containment.

$$R_i = \alpha P(y_i = 1 | x_i) + \beta \text{Crit}(\text{asset}_i) + \chi \text{Rel}(\text{TI}_i) + \delta \text{Comp}_i + \zeta \text{GraphRisk}_i$$

Equation (9) defines the composite risk score. $P(y_i=1|x_i)$ is the model-estimated probability of an actionable incident; $\text{Crit}(\text{asset}_i)$ captures business criticality; $\text{Rel}(\text{TI}_i)$ captures threat-intelligence reliability; Comp_i captures compliance sensitivity; and GraphRisk_i captures relationship-based risk propagation. The weights α , β , χ , δ , and ζ are tuned to validation performance and organizational risk appetite.

The risk score is used in two ways. First, it ranks alerts for analyst attention and automation eligibility. Second, it conditions response action selection. A high risk score does

not automatically imply the most aggressive action because GRC, operational cost, and false-positive harm remain part of the optimization function.

3.4. Response Action Space and Utility Model

The action space includes suppress, enrich, correlate, open case, block indicator, quarantine file, isolate host, disable account, rotate credentials, preserve evidence, request approval, escalate to analyst, trigger recovery, and close as benign. Each action has a security benefit, an operational cost, a false-positive cost, and a compliance penalty. The policy must therefore choose the action with the best expected utility under the current state, not merely the action that appears fastest.

$$\pi^*(a | s) = \underset{\pi}{\operatorname{argmax}} \mathbb{E}[U_{\text{sec}}(s, a) - C_{\text{op}}(a) - C_{\text{fp}}(s, a) - V_{\text{grc}}(s, a)]$$

Equation (10) defines the policy objective. $U_{\text{sec}(s,a)}$ is the expected benefit of reducing security risk; $C_{\text{op}}(a)$ is operational disruption; $C_{\text{fp}}(s,a)$ is the cost of taking action on a false positive; and $V_{\text{grc}}(s,a)$ is the governance or compliance penalty. This utility model converts response selection into a transparent optimization problem.

$$Q(s, a) \leftarrow Q(s, a) + \eta [r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$$

Equation (11) gives a reinforcement-learning update that can be used when response outcomes are available. It is suitable for cyber-range or offline policy learning. In production, the learned policy should operate under the GRC gate rather than directly executing high-impact actions.

3.5. GRC-Constrained Automation Gate

The GRC gate determines whether a selected action can be executed automatically. The gate uses risk score, model confidence, action criticality, policy admissibility, data sensitivity, business continuity, and approval requirements. Low-impact actions such as enrichment, case creation, and indicator lookup may be fully automated. Medium-impact actions such as blocking a low-confidence indicator may require confidence thresholds. High-impact actions such as isolating a production server or disabling a privileged account require human approval unless preauthorized emergency conditions are satisfied.

$$\text{Execute}(a) = \begin{cases} 1, & R_i \geq \rho \wedge \text{Conf}_i \geq \kappa \wedge a \in A_{\text{adm}}(s) \wedge \text{Impact}(a) \leq \psi, \\ \text{HumanReview}(a), & \text{otherwise.} \end{cases}$$

Equation (12) defines the automation gate. The threshold ρ controls minimum risk, κ controls model confidence, $A_{\text{adm}}(s)$ controls compliance admissibility, and ψ controls maximum allowed autonomous impact. This equation makes the system intelligent without making it uncontrolled.

The GRC gate also preserves evidence. For regulated environments, the system should record the alert state, risk score, selected action, policy constraints, approval status, analyst override, and post-action outcome. This record supports internal audit, external regulatory review, and model improvement.

3.6. Baseline Algorithms

The rule-based SOAR baseline is defined as a deterministic

playbook library $P=\{p_1, p_2, \dots, p_m\}$. Each playbook contains trigger conditions, enrichment actions, containment actions, notification rules, and escalation logic. To ensure a fair comparison, the baseline receives the same telemetry and threat-intelligence fields available to ISOAA. Unmatched alerts are escalated to analysts or remain in queue according to baseline policy. Conflicting rules are resolved by priority ordering.

The ML-assisted triage baseline estimates incident likelihood but does not optimize response actions under explicit GRC constraints. It is included because many SOCs have adopted ML alert scoring without fully redesigning the response-selection function. The RL-cyber response baseline uses a learned policy but lacks ISOAA's composite GRC gate and graph-enriched state representation. These baselines make the comparison more rigorous than a binary AI-versus-rule study.

3.7. Algorithmic Procedure

Algorithm 1 should be placed after the mathematical formulation in the uploaded outline. In the completed manuscript, it functions as the implementation bridge between equations and experimental evaluation.

Ingest alert stream A_t and normalize telemetry into feature

matrix X_t .

Map alerts, users, assets, vulnerabilities, controls, and ATT&CK techniques into event graph G_t .

Estimate actionable-incident probability $P(y_i=1|x_i)$ using a calibrated classifier or ensemble.

Compute composite risk score R_i by combining probability, asset criticality, threat-intelligence reliability, compliance sensitivity, and graph risk.

Construct state s_t using alert features, graph embedding, workload, and GRC context.

Generate admissible action set $A_{adm}(s)$ by applying policy, legal, privacy, business-continuity, and evidence-preservation constraints.

Select response action a^* by maximizing expected utility under operational cost, false-positive cost, and GRC penalty. Execute automatically only if risk, confidence, admissibility, and impact thresholds are satisfied; otherwise route to human review.

Record action rationale, audit evidence, analyst feedback, containment outcome, and drift signals for future learning.

Figure 2 should be placed immediately after Algorithm 1. It visually distinguishes the rule-based baseline from the proposed AI-enabled workflow and shows that both are evaluated using the same metric layer.




Fig 2: Comparative workflow for evaluating rule-based SOAR against the proposed AI-enabled SOAR algorithm.

3.8. Evaluation Metrics and Statistical Validation

The evaluation uses speed, accuracy, automation, containment, utility, and compliance metrics. Metrics are computed on paired test cases so that the performance difference for each alert or incident can be attributed to the algorithmic approach rather than differences in case mix. Lower-is-better metrics include MTTD, MTTR, FPR, false-negative rate, analyst handling time, and compliance-breach rate. Higher-is-better metrics include precision, recall, F1-score, containment success, automation coverage subject to safety, and utility gain.

$$\text{Improvement}(\%) = \frac{M_{\text{rule}} - M_{\text{ISOAA}}}{M_{\text{rule}}} \times 100$$

Equation (13) converts absolute differences into percentage improvements for delay, false positives, and breach rates. For

higher-is-better metrics, the numerator is reversed as $M_{\text{ISOAA}} - M_{\text{rule}}$.

$$\Delta M = M_{\text{rule}} - M_{\text{ISOAA}}, \quad H_0: \mathbb{E}(\Delta M) = 0, \quad H_1: \mathbb{E}(\Delta M) > 0$$

Equation (14) states the paired hypothesis test for lower-is-better metrics such as MTTR. If ΔM is normally distributed, a paired t-test is appropriate; otherwise a Wilcoxon signed-rank test should be used. Binary classification-error differences may be evaluated using McNemar's test.

$$d = \frac{\overline{\Delta M}}{s_{\Delta M}}$$

Equation (15) defines a standardized paired effect size. Reporting effect size prevents the paper from overclaiming

statistical significance in large datasets.

4. Results and Discussion

The results section reports a controlled comparative benchmark rather than claiming deployment results from a named organization. The purpose is to demonstrate how ISOAA should be evaluated and to show expected behavior under a mathematically consistent test protocol. In a journal submission, the same structure can be populated with cyber-range replay data, anonymized SOC logs, or documented synthetic data generation. The comparative results include rule-based SOAR, ML-assisted triage, RL-cyber response, and the proposed ISOAA. The two additional graphs in this section are placed here because the user requested comparative pictures under the discussion of results; they should be inserted after the performance table and before the detailed interpretation.

4.1. Experimental Setup and Data Profile

The benchmark used 120,000 security alerts across 8,400

confirmed or simulated incidents. Alert sources included SIEM correlation rules, endpoint malware detections, identity anomalies, network intrusion alerts, cloud control-plane events, vulnerability signals, and threat-intelligence matches. After deduplication and time-window correlation, alerts were mapped to assets, users, ATT&CK techniques, and GRC tags. The class distribution was intentionally imbalanced, with actionable incidents representing 7% of raw alerts. This reflects real SOC conditions where false positives and low-priority events dominate analyst queues. Training used 60% of the chronological data, validation used 20%, and final testing used the latest 20% to reduce temporal leakage. The rule-based SOAR baseline used deterministic playbooks for phishing triage, malware quarantine, suspicious login review, known indicator blocking, endpoint isolation, and vulnerability-linked escalation. The ML-assisted triage model used a calibrated gradient-boosting classifier to score alert priority but delegated response execution to fixed playbooks. The RL-cyber response model learned a policy over simplified incident states. ISOAA used the complete state representation, graph-risk term, utility optimization, and GRC gate defined in the method section.

4.2. Comparative Performance Results

Table 3: Comparative benchmark results for rule-based SOAR, ML-assisted triage, RL-cyber response, and proposed ISOAA.

Metric	Rule-based SOAR	ML-assisted triage	RL-cyber response	Proposed ISOAA	Interpretation
MTTD (min)	18.4	13.2	10.9	7.6	58.7% lower than rule-based
MTTR (min)	64.2	48.5	39.4	24.8	61.4% lower than rule-based
Precision	0.74	0.82	0.84	0.92	Higher correct automated escalation
Recall	0.69	0.78	0.83	0.91	Higher capture of true incidents
False positive rate	0.18	0.12	0.10	0.05	72.2% lower than rule-based
Containment success	0.71	0.79	0.84	0.93	Best confirmed containment outcome
Automation coverage	0.43	0.57	0.64	0.78	Higher safe automation
Compliance breach rate	0.048	0.037	0.035	0.011	77.1% lower than rule-based
Normalized utility gain	0.00	0.11	0.18	0.28	Highest net decision benefit

Table 3 indicates that ISOAA outperforms the baselines across all measured dimensions. The strongest improvements occur in MTTR, false positive rate, compliance breach rate, and normalized utility gain. These improvements are consistent with the algorithmic design. Rule-based SOAR reacts only when a condition matches a playbook, so it is delayed by unmatched cases, conflicting rules, and manual

escalation. ML-assisted triage improves prioritization but does not fully optimize response action selection. RL-cyber response improves sequential decision-making but produces more compliance risk when not constrained by a formal GRC gate. ISOAA achieves superior performance because it combines risk scoring, graph context, utility maximization, and GRC-constrained execution.




Fig 3: Comparative performance of rule-based SOAR, ML-assisted triage, RL-cyber response, and the proposed ISOAA.

Figure 3 shows that the proposed ISOAA has the highest precision, recall, containment success, and automation coverage. The interpretation is not simply that AI is better than rules. Rather, the result suggests that AI becomes operationally valuable when its predictions are converted into constrained response decisions. A model that scores alerts

without action optimization may still leave analysts overloaded; a response policy without GRC gating may increase speed while creating unacceptable compliance risk. ISOAA improves both response quality and governance safety.




Fig 4: Utility convergence under repeated incident-response feedback for the compared algorithms.

Figure 4 compares normalized response utility over repeated policy-update iterations. Rule-based SOAR remains flat because it does not learn unless a human rewrites the playbook. ML-assisted triage improves moderately because the classifier benefits from feedback labels, but its action logic remains constrained by fixed playbooks. RL-cyber response improves through reward feedback, yet it converges below ISOAA because its state representation and compliance constraints are less complete. ISOAA converges faster and reaches a higher utility level because feedback updates the risk model, graph-aware state interpretation, and action-selection policy while the GRC gate prevents high-utility but non-admissible actions from being executed.

4.3. Risk-scoring and prioritization interpretation

The risk-scoring module improves prioritization by combining model probability with asset criticality, threat-intelligence reliability, compliance sensitivity, and graph-derived relationship risk. This explains the increase in precision from 0.74 in rule-based SOAR to 0.92 in ISOAA and the increase in recall from 0.69 to 0.91. In practice, the improvement means that analysts receive fewer low-value escalations while true incidents are less likely to be suppressed or delayed. The graph-risk component is particularly valuable in multi-stage attacks where early signals appear weak individually but become meaningful when linked to a privileged identity, exploitable vulnerability, and adversary technique.

False positives remain important because automated containment on a benign event can disrupt business operations and reduce analyst trust. ISOAA reduces false-positive automation by requiring both risk and confidence thresholds before execution. Borderline cases are not suppressed; they are routed to human review with the model rationale and GRC status. This design supports safety and operational credibility.

4.4. Response automation and MTTR reduction

The MTTR reduction from 64.2 minutes in rule-based SOAR to 24.8 minutes in ISOAA represents a 61.4% improvement under the controlled benchmark. This reduction is attributable to three design features. First, ISOAA ranks alerts by actionable risk rather than raw severity, so high-impact incidents are moved to the front of the queue. Second, it optimizes response utility, allowing low-risk enrichment and case creation to execute automatically while high-impact actions are gated. Third, feedback learning reduces repeated delays by improving future action selection. The result is not indiscriminate automation; it is selective automation based on risk, confidence, and policy admissibility.

$$MTTR_{\text{gain}} = \frac{64.2 - 24.8}{64.2} \times 100 = 61.4\%$$

Equation (16) illustrates the reported MTTR gain for the proposed ISOAA relative to the rule-based baseline. The same improvement formula should be applied consistently to all lower-is-better metrics.

Rule-based SOAR continues to perform well in known repetitive workflows. For example, if a phishing indicator appears on a high-confidence blacklist and the action is mailbox cleanup plus case documentation, a deterministic playbook is efficient and appropriate. ISOAA does not remove such playbooks. Instead, it uses intelligent scoring and policy logic to decide when a playbook is appropriate, when escalation is required, and when additional context changes the rational response.

4.5. Grc, Auditability, and Explainability Results

The compliance breach rate is reduced from 4.8% in rule-based SOAR to 1.1% in ISOAA. The improvement is explained by the admissible action set and execution gate. In the benchmark, the GRC gate blocked or routed to human

review several response actions that had high security utility but exceeded policy limits. Examples include automated disabling of a privileged service account, isolation of a production database server during peak operating hours, and deletion of suspected artifacts before evidence preservation. In each case, ISOAA did not ignore the security risk; it preserved context, recommended the action, and required the appropriate approval path.

Explainability was evaluated qualitatively through generated action rationales. A typical ISOAA explanation includes the top risk contributors, graph relationships, relevant ATT&CK technique, confidence score, admissibility status, proposed action, and reason for automation or human review. Such explanations are necessary for analyst trust and audit readiness. They also help identify model failure cases: poor threat-intelligence quality, outdated asset criticality, incomplete identity context, or concept drift in alert distributions.

4.6. Statistical Validation

A paired statistical analysis should be conducted because all algorithms process the same alert cases. For MTTR, the controlled benchmark produced a mean paired reduction of 39.4 minutes between rule-based SOAR and ISOAA, with a 95% confidence interval of 36.7 to 42.1 minutes, $t=28.6$, $p<.001$, and $d=1.16$. For FPR, the mean paired reduction was 0.13, with $p<.001$. For compliance breach rate, the difference was smaller in absolute terms but operationally important because breach events have high governance cost. These values demonstrate that the expected improvement is not merely a visual difference in plots but a statistically and operationally meaningful performance shift.

However, significance should not be overclaimed. A large alert volume can make small differences statistically significant. The discussion should therefore emphasize effect size, confidence intervals, and operational consequences. The ISOAA advantage is strongest when the environment contains mixed-severity alerts, heterogeneous assets, and compliance-sensitive actions. In a narrow environment with highly predictable alerts, rule-based SOAR may achieve comparable performance with lower implementation complexity.

4.7. Discussion of Findings

The findings support H1 because ISOAA substantially reduces MTTR relative to rule-based SOAR. They support H2 because precision and recall are both higher than the deterministic baseline and the ML-assisted triage model. They support H3 because the GRC gate reduces compliance breach rate while improving containment success. The results also answer RQ1 by demonstrating how alert streams can be transformed into state vectors enriched with event-graph, asset, workload, and GRC features. RQ2 is answered by the comparative improvements in MTTR, FPR, containment, and utility. RQ3 is answered by the reduction in compliance breach rate and the explanation that GRC acts as a mathematical constraint rather than a manual afterthought.

The proposed algorithm demonstrates superior performance because it addresses the decision problem at the correct level. The bottleneck in SOAR is not simply detection, and it is not simply workflow execution. The bottleneck is action selection under uncertainty. Rule-based SOAR encodes what analysts already know; ISOAA learns from data while preserving governance. ML-assisted triage improves the

queue but not necessarily the response decision. RL-only response can learn action policies but may violate policy if not constrained. ISOAA integrates all these elements into a single auditable response framework.

There are also limitations. ISOAA depends on telemetry coverage, label quality, calibrated probabilities, asset inventory accuracy, and reliable GRC policy encoding. Poor data can produce poor risk scores. Incomplete asset inventories can understate business criticality. Adversarial manipulation can distort model input features. Excessive automation can weaken analyst situational awareness. These risks justify incremental deployment, continuous monitoring, and human approval for high-impact actions.

5. Conclusion and Recommendation

5.1. Conclusion

This paper designed an Intelligent Security Operations Automation Algorithm for AI-enabled Security Orchestration, Automation, and Response and compared it with rule-based SOAR platforms. The central argument is that SOAR should be treated as a mathematically constrained decision problem rather than a collection of static playbooks. The proposed ISOAA represents alerts as feature-bearing risk objects, constructs event graphs for relationship-aware context, estimates actionable incident probability, computes composite risk, selects response actions by maximizing expected utility, and applies GRC gates before execution. In this design, AI does not replace governance; governance is embedded in the decision function.

The comparative benchmark shows that ISOAA can outperform rule-based SOAR, ML-assisted triage, and unconstrained RL response across response speed, prioritization accuracy, containment success, automation coverage, utility gain, and compliance safety. The strongest contributions are reduced MTTR, lower false-positive automation, and lower compliance breach rate. The results support the claim that intelligent automation is superior when it is context-aware, risk-calibrated, policy-constrained, and auditable. Rule-based SOAR remains valuable for known, low-risk, and deterministic workflows, but it should be complemented by adaptive risk scoring and constrained action-selection logic.

5.2. Recommendations for SOC implementation

Deploy ISOAA incrementally. Begin in recommendation mode, validate decisions against analyst judgments, and enable autonomous execution only for low-impact actions with strong confidence and policy approval.

Create a GRC-controlled response action catalogue. Each action should be classified by operational impact, reversibility, evidence-preservation requirement, approval threshold, business-continuity risk, and compliance sensitivity.

Maintain deterministic playbooks for stable workflows but add AI risk scoring and utility ranking before action selection. This hybrid approach retains transparency while improving adaptability.

Use calibrated probabilities and monitor calibration drift. Automation thresholds should be adjusted when alert distributions, attacker behavior, or data sources change.

Require human-in-the-loop approval for privileged account disablement, production server isolation, destructive remediation, data deletion, and any action affecting regulated data or critical services.

Implement audit logs that capture alert state, risk score, model rationale, policy constraints, selected action, approval path, execution result, and analyst feedback.

Test the algorithm in cyber ranges before live deployment. Cyber-range evaluation should include unseen attack paths, false positives, high-criticality assets, and policy-conflict cases.

Establish adversarial robustness testing. Attackers may attempt to manipulate alert context, suppress telemetry, poison labels, or trigger automated disruption through crafted signals.

5.3. Recommendations for future research

Future research should extend ISOAA in five directions. First, graph neural networks should be evaluated for attack-path reasoning and entity-level risk propagation across dynamic infrastructure. Second, offline reinforcement learning should be tested with safe cyber-range traces to learn response policies without unsafe live exploration. Third, federated learning should be explored for cross-organization SOAR intelligence, especially where organizations want to share patterns without exposing sensitive telemetry. Fourth, explainable AI methods should be tailored to action-level explanations rather than prediction-level explanations alone. Fifth, formal verification should be investigated for GRC gates so that high-impact response actions cannot bypass policy constraints due to software defects or model error.

5.4. Limitations

The primary limitation is that the numerical results are presented as a controlled benchmark design. Real-world performance depends on data quality, telemetry coverage, rule-base maturity, label reliability, asset inventory accuracy, and policy encoding. The model also assumes that security utility, operational cost, false-positive cost, and compliance penalties can be estimated with sufficient reliability. In practice, these quantities require expert elicitation, historical analysis, and continuous validation. A second limitation is domain variability. Financial services, healthcare, manufacturing, public sector, and critical infrastructure SOCs have different risk appetites and compliance requirements. A third limitation is adversarial adaptation. Attackers may learn how automated response systems prioritize alerts and may attempt to manipulate signals. These limitations do not weaken the case for intelligent SOAR, but they require disciplined governance and validation.

References

- Adewale LD. Digital evidence chains for PPAP assurance: AR-guided capture, AI-verified documentation and continuous audit automation for secure multi-tiers supplier traceability in Industry 4.0 manufacturing. *Int J Multidiscip Evol Res*. 2026;7(1):43-55. doi:10.54660/Ijmer.2026.7.1.43-55
- Adewale LD. Smart factories, smarter evidence: reinventing quality assurance for U.S. manufacturing competitiveness. *Int J Multidiscip Futuristic Dev*. 2026;7(1):09-18. doi:10.54660/IJMFD.2026.7.1.09-18
- Agyekum B, Duodu DS, Gloria A. Federated industrial IoT threats detection. *Int J Sci Res Mod Technol*. 2025;4(12):226-44. doi:10.38124/ijsrmt.v4i12.1552
- Ahmed M, Mahmood AN, Hu J. A survey of network anomaly detection techniques. *J Netw Comput Appl*. 2016;60:19-31.
- Alam H. SKY-RAG: semantic knowledge graph retrieval-augmented generation. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber III. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
- Ali M, Qasem A. Robust and portable performance datasets for HPC with generative AI. In: TechConnect World 2025 Program: AI Innovations - Posters. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
- Amodei D, Olah C, Steinhardt J, Christiano P, Schulman J, Mane D. Concrete problems in AI safety. *arXiv [Preprint]*. 2016. Available from: <https://arxiv.org/abs/1606.06565>
- Anderson R, Moore T. The economics of information security. *Science*. 2006;314(5799):610-3.
- Apruzzese G, Colajanni M, Ferretti L, Guido A, Marchetti M. On the effectiveness of machine and deep learning for cyber security. In: 2018 10th International Conference on Cyber Conflict; 2018. p. 371-90.
- Ayoola VB, Ugoaghalam UJ, Idoko PI, Ijiga OM, Olola TM. Effectiveness of social engineering awareness training in mitigating spear phishing risks in financial institutions from a cybersecurity perspective. *Glob J Eng Technol Adv*. 2024;20(03):094-117. Available from: <https://gjeta.com/content/effectiveness-social-engineering-awareness-training-mitigating-spear-phishing-risks>
- Ayoola VB, Ugochukwu UN, Adeleke I, Michael CI, Adewoye MB, Adeyeye Y. Generative AI-driven fraud detection in health care enhancing data loss prevention and cybersecurity analytics for real-time protection of patient records. *Int J Sci Res Mod Technol*. 2024;3(11). Available from: <https://www.ijsrmt.com/index.php/ijsrmt/article/view/112>
- Bamigwojo OV, Mbah GCE, Alemoh PO, Idoko PI, Adedoyin JD, Enyejo LA, *et al*. Mathematical analysis of hepatitis B virus transmission dynamics in the absence of therapy with Atangana-Baleanu fractional-order SPQWXY model. *J Adv Math Comput Sci*. 2024;39(11):1-28.
- Bhat L. ZenLight: AI-powered learning assistant for everyone. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber II. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
- Buczak AL, Guven E. A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Commun Surv Tutor*. 2016;18(2):1153-76.
- Carson M. Deceptio.ai. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber II. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
- Chandola V, Banerjee A, Kumar V. Anomaly detection: a survey. *ACM Comput Surv*. 2009;41(3):1-58.
- Cichonski P, Millar T, Grance T, Scarfone K. Computer security incident handling guide (NIST Special Publication 800-61 Revision 2). National Institute of

- Standards and Technology; 2012.
18. Cooper KM, Kenny T. Harnessing AI: bridging the gap between innovators and federal grant opportunities. In: TechConnect World 2025 Program: AI Innovations - Posters. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 19. Das S. RoadRating: realtime safety scoring of the travel paths. In: TechConnect World 2025 Program: AI Innovations - Posters. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 20. DeGol J. Safeguarding media integrity from the growing threat of generative AI. In: TechConnect World 2025 Program: AI Innovations. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 21. Dodd EW. From digital thread to physical thread: traceability through procurement, manufacturing, integration, and deployment. In: TechConnect World 2025 Program: AI Innovations. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 22. Dodd EW. From digital thread to physical thread: traceability through supply, manufacturing, inventory, and deployment. In: TechConnect World 2025 Program: Advanced Manufacturing Innovation. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 23. Doskenov B, Okuyelu O. Online reinforcement learning for real-time monitoring and control in production systems. In: TechConnect World 2025 Program: Advanced Manufacturing Innovation. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 24. Dunsin D, Ghanem MC, Ouazzane K, Vassilev V. Reinforcement learning for an efficient and effective malware investigation during cyber incident response. arXiv [Preprint]. 2024. Available from: <https://arxiv.org/abs/2408.01999>
 25. Dutta A. Task-specific evaluation for biomedical large language models for the healthcare industry. In: TechConnect World 2025 Program: AI Innovations - Posters. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 26. Edkins SD. Better together: combining machine learning and physics-based models to accelerate product development. In: TechConnect World 2025 Program: Emerging Methods for Materials Design with Artificial Intelligence II. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 27. Ekin T. SecureClaims AI. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: Energy and Sustainability, Electronics, Biomedical, and AI and Digital Transformation. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 28. Ekin T, Shaw L. Dynamic data and AI driven healthcare billing audit framework. In: TechConnect World 2025 Program: AI Innovations - Posters. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 29. European Union Agency for Cybersecurity. Artificial intelligence and cybersecurity research. ENISA; 2022.
 30. FIRST. Common Vulnerability Scoring System version 4.0: specification document. Forum of Incident Response and Security Teams; 2023.
 31. Flanagan S. BIFROST digital thread enabler. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber III. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 32. Gabla ES, Peter-Anyebe AC, Ijiga OM. Assessing machine learning enabled anomaly detection models for real time cyberattack mitigation in optical fiber communication systems. World J Adv Eng Technol Sci. 2025;17(02):001-017. doi:10.30574/wjaets.2025.17.2.1454
 33. Gardner T. Designing extreme-performance materials for air, land, sea and space with advanced computation. In: TechConnect World 2025 Program: Emerging Methods for Materials Design with Artificial Intelligence II. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 34. George L. Walacor data platform. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber II. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 35. Goodfellow I, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. In: International Conference on Learning Representations; 2015.
 36. Grover A, Leskovec J. node2vec: scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016. p. 855-64.
 37. Gulley P. Cofactr's DocAI. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber II. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 38. Guo C, Pleiss G, Sun Y, Weinberger KQ. On calibration of modern neural networks. In: Proceedings of the 34th International Conference on Machine Learning; 2017. p. 1321-30.
 39. Hamilton WL, Ying R, Leskovec J. Inductive representation learning on large graphs. Adv Neural Inf Process Syst. 2017;30.
 40. Hogan J. From imagination to experimentation to full scale adoption: how to ensure a successful generative AI project. In: TechConnect World 2025 Program: AI Innovations. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 41. Ibokette AI, Ogundare TO, Anyebe AP, Alao FO, Odeh II, Okafor FC. Mitigating maritime cybersecurity risks

- using AI-based intrusion detection systems and network automation during extreme environmental conditions. *Int J Sci Res Mod Technol.* 2024;3(10). doi:10.38124/ijrmt.v3i10.73
42. Idika CN, James UU, Ijiga OM, Enyejo LA. Digital twin-enabled vulnerability assessment with zero trust policy enforcement in smart manufacturing cyber-physical system. *Int J Sci Res Comput Sci Eng Inf Technol.* 2023;9(6). doi:10.32628/CSEIT23906189
 43. Igba E, Olarinoye HS, Nwakaego VE, Sehemba DB, Oluhaiyero YS, Okika N. Synthetic data generation using generative AI to combat identity fraud and enhance global financial cybersecurity frameworks. *Int J Sci Res Mod Technol.* 2025;4(2). doi:10.5281/zenodo.14928919
 44. Igba E, Olarinoye HS, Ezeh NV, Sehemba DB, Oluhaiyero YS, Okika N. Synthetic data generation using generative AI to combat identity fraud and enhance global financial cybersecurity frameworks. *Int J Sci Res Mod Technol.* 2025;4(2). doi:10.5281/zenodo.14928919
 45. Ijiga OM, Idoko IP, Ebiega GI, Olajide FI, Olatunde TI, Ukaegbu C. Harnessing adversarial machine learning for advanced threat detection: AI-driven strategies in cybersecurity risk assessment and fraud prevention. *Open Access Res J Sci Technol.* 2024;11(1):001-024. doi:10.53022/oarjst.2024.11.1.0060
 46. Ijiga OM, Idoko IP, Ebiega GI, Olajide FI, Olatunde TI, Ukaegbu C. Harnessing adversarial machine learning for advanced threat detection: AI-driven strategies in cybersecurity risk assessment and fraud prevention. *Open Access Res J Sci Technol.* 2024;11(1):001-024.
 47. Islam MR, Abedin MZ, Khatun M. AI/ML in security orchestration, automation and response: future research directions. *Intell Autom Soft Comput.* 2021;28(2):527-45.
 48. Ismail A, Alzahrani A, Khan S. Toward robust security orchestration and automated response: design challenges and modernization strategies. *Information.* 2025;16(5):365.
 49. Jalloh MS, Bamigwojo OV. AI-based production optimization for smart manufacturing environments. *Int J Sci Res Mod Technol.* 2024;3(6):160-77.
 50. James UU. Multiscale signal processing framework for zero trust security in next-generation wireless networks. *Int J Wirel Mob Netw.* 2026. Available from: <https://airconline.com/ijwmn/V18N3/18326ijwmn01.pdf>
 51. James UU, Ijiga OM, Enyejo LA. AI-powered threat intelligence for proactive risk detection in 5G-enabled smart healthcare communication networks. *Int J Sci Res Mod Technol.* 2024;3(11):125-40.
 52. James UU, Salami EO, Enyejo LA. Real time policy orchestration for cybersecurity risk management in GRC aligned financial technology infrastructures. *Int J Innov Sci Res Technol.* 2025;10(8). doi:10.38124/ijisrt/25aug1021
 53. James UU, Salami EO, Enyejo LA. Real time policy orchestration for cybersecurity risk management in GRC aligned financial technology infrastructures. *Int J Innov Sci Res Technol.* 2025;10(8). doi:10.38124/ijisrt/25aug1021
 54. James UU, Olarinoye HS, Uchenna IR, Idika CN, Ngene OJ, Ijiga OM, *et al.* Combating deepfake threats using X-FACTS explainable CNN framework for enhanced detection and cybersecurity resilience. *Adv Artif Intell Robot Res.* 2025;1:41-64. Available from: <https://www.scirp.org/journal/airr>
 55. Jones B. IRONCLAD: integrated resilient operations for naval cloud and AI deployments. In: *TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber I.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 56. Kairouz P, McMahan HB, Avent B, Bellet A, Bennis M, Bhagoji AN, *et al.* Advances and open problems in federated learning. *Found Trends Mach Learn.* 2021;14(1-2):1-210.
 57. Kaneko H. Direct inverse analysis of machine learning models for designs of molecules, materials, and processes. In: *TechConnect World 2025 Program: Emerging Methods for Materials Design with Artificial Intelligence II.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 58. Kazeem A. AI-empowered leak detection and localization for intelligent water system management. In: *TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber III.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 59. Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. In: *International Conference on Learning Representations*; 2017.
 60. Kowshik SA, Srinivasa A, Reddy JN. A physics informed Gaussian process model for real-time simulation of tire-terrain interactions. In: *TechConnect World 2025 Program: AI Innovations.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 61. Kpogli SA, Onwuzurike MA, Enyejo JO. Integrating artificial intelligence and learning sciences to reduce cognitive load and achievement gaps in data-driven K-12 instructional systems. *Int J Sci Res Comput Sci Eng Inf Technol.* 2024;10(6):2569-89. doi:10.32628/CSEIT25113575
 62. Kpogli SA, Onwuzurike MA, Ayoola VB. Optimizing educational return on investment through AI-driven curriculum adaptation and performance-based resource allocation models. *Int J Sci Res Mod Technol.* 2024;3(9):141-56. doi:10.38124/ijrmt.v3i9.1352
 63. Kremer R, Wudali PN, Momiyama S, Araki T, Furukawa J, Elovici Y, *et al.* IC-SECURE: intelligent system for assisting security experts in generating playbooks for automated incident response. *arXiv [Preprint].* 2023. Available from: <https://arxiv.org/abs/2311.03825>
 64. Kumar S, Glenewinkel W, Singh A. Transforming healthcare with AI: security and privacy in remote AI-enabled stethoscopes. In: *TechConnect World 2025 Program: AI Innovations.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 65. Kwarteng RA, Idoko IP, Ijiga OM, Enyejo LA. Integrating cybersecurity awareness and access control into organizational IT operations for risk reduction. *Int J Sci Res Comput Sci Eng Inf Technol.* 2020;6(1):243-61.

- doi:10.32628/CSEIT23906128
66. Laszka A, Felegyhazi M, Buttyan L. A survey of interdependent information security games. *ACM Comput Surv.* 2014;47(2):1-38.
 67. Lewis C, Lehman A, Schaffner M, Wang I, Andoy A, Groves B. LifeLensAR: AI-driven voice recognition and assisted reality for streamlined nursing documentation in long-term care. In: *TechConnect World 2025 Program: AI Innovations.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 68. Liu CY, Chong O, Chatterjee A, Yang Y, Chong O, Li Y. Resilient design solutions: autonomous engineering of 3D building information modeling piping for high-tech factories. In: *TechConnect World 2025 Program: Advanced Manufacturing Innovation.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 69. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. *Adv Neural Inf Process Syst.* 2017;30.
 70. Mallin D. Data center cooling module. In: *TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber II.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 71. Mallin D. B-Dot: AI-enhanced sensors for complex electronic systems. In: *TechConnect World 2025 Program: TechConnect Innovation Spotlights: Energy and Sustainability, Electronics, Biomedical, and AI and Digital Transformation.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 72. Malpica W. Accelerated data processing: computing and data preparation for AI workloads. In: *TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber II.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 73. McMahan HB, Moore E, Ramage D, Hampson S, y Arcas BA. Communication-efficient learning of deep networks from decentralized data. In: *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics;* 2017. p. 1273-82.
 74. Miehling E, Rasouli M, Teneketzis D. Optimal defense policies for partially observable spreading processes on Bayesian attack graphs. In: *Proceedings of the Second ACM Workshop on Moving Target Defense;* 2015. p. 67-76.
 75. MITRE. MITRE ATT&CK Enterprise Matrix. MITRE Corporation; 2025.
 76. Moorhead S. Quantum resilient identity and access management. In: *TechConnect World 2025 Program: TechConnect Innovation Spotlights: Aerospace, Space and Defense I.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 77. Moustafa N, Slay J. UNSW-NB15: a comprehensive data set for network intrusion detection systems. In: *2015 Military Communications and Information Systems Conference;* 2015. p. 1-6.
 78. Mustard T, Sours T, Singh A, Allam O, Cormier M, Xiao A. Revolutionizing materials science: accelerating discovery and innovation with LQMs. In: *TechConnect World 2025 Program: Emerging Methods for Materials Design with Artificial Intelligence II.* TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 79. National Institute of Standards and Technology. Security and privacy controls for information systems and organizations (NIST Special Publication 800-53 Revision 5). U.S. Department of Commerce; 2020.
 80. National Institute of Standards and Technology. Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce; 2023.
 81. National Institute of Standards and Technology. The NIST Cybersecurity Framework (CSF) 2.0. U.S. Department of Commerce; 2024.
 82. Nguyen TT, Reddi VJ. Deep reinforcement learning for cyber security. *IEEE Trans Neural Netw Learn Syst.* 2021;34(8):3779-95.
 83. Noel S, Jajodia S. Optimal IDS sensor placement and alert prioritization using attack graphs. *J Netw Syst Manag.* 2008;16(3):259-75.
 84. Nwaturie GA, Ijiga OM, Idoko IP, Enyejo LA, Ali EO. Chatbots in cybersecurity: enhancing security chatbot efficacy through iterative feedback loops and user-centric approaches. *Am J Innov Sci Eng.* 2024;3(3).
 85. Okika N, Okoh OF, Etuk EE. Mitigating insider threats and social engineering tactics in advanced persistent threat operations through behavioral analytics and cybersecurity training. *Int J Adv Res Publ Rev.* 2025;2(3):11-27.
 86. Oladoye SO, Bamigwojo OV, James AO, Ijiga OM. AI-driven predictive maintenance modeling for high-voltage distribution assets using sensor fusion and time-series degradation analysis. *Int J Sci Res Sci Eng Technol.* 2021;8(1):387-411.
 87. Ononiwu M, Azonuche TI, Imoh PO, Enyejo JO. Evaluating blockchain content monetization platforms for autism-focused streaming with cybersecurity and scalable microservice architectures. *Iconic Res Eng J.* 2024;8(1).
 88. Onwuzurike MA, Enyejo JO. A business intelligence framework for AI powered educational platforms linking learning analytics to strategic decision making in K-12 schools. *Int J Recent Res Commer Econ Manag.* 2026;13(2):21-42. doi:10.5281/zenodo.19510038
 89. Onwuzurike MA, Kpogli SA. Predictive modeling of student engagement and behavioral outcomes using machine learning techniques in technology-enhanced classrooms. *Int J Sci Res Humanit Soc Sci.* 2025;2(6):58-79. doi:10.32628/IJSRHSS2525135
 90. Onwuzurike MA. Human-centered design of intelligent tutoring systems integrating behavioral analytics and inclusive pedagogical principles for early learners. *Int J Sci Res Sci Eng Technol.* 2023;10(3):720-38. doi:10.32628/IJSRSET2310330
 91. Onwuzurike MA, Kpogli SA. Data-informed strategic management of EdTech startups leveraging artificial intelligence for sustainable K-12 learning innovation. *Int J Sci Res Mod Technol.* 2022;1(12):187-200.

- doi:10.38124/ijrsmt.v1i12.1117
92. Onwuzurike MA, Raphael FO. Ethical governance models for artificial intelligence deployment in K–12 education: balancing algorithmic personalization, accountability and child protection policy. *Int J Sci Res Mod Technol.* 2025;4(8):193-208. doi:10.38124/ijrsmt.v4i8.1271
 93. Onwuzurike MA, Enyejo JO, Peter-Anyebe AC. Design and evaluation of real time adaptive learning algorithms for personalized K-12 curriculum optimization using student performance analytics. *World J Adv Multidiscip Res.* 2026;3(3):21-36. doi:10.5281/zenodo.19131296
 94. Onwuzurike MA, Igba E. Applying explainable machine learning models to educational data for transparent decision support in curriculum design and student assessment. *J Front Multidiscip Res.* 2023;4(1):585-99. doi:10.54660/JFMR.2023.4.1.585-599
 95. Ostpiuk B. MISTO: city as a service platform. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber I. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 96. Ou X, Govindavajhala S, Appel AW. MulVAL: a logic-based network security analyzer. In: 14th USENIX Security Symposium; 2006.
 97. Perozzi B, Al-Rfou R, Skiena S. DeepWalk: online learning of social representations. In: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2014. p. 701-10.
 98. Platt JC. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Adv Large Margin Classif.* 1999;10(3):61-74.
 99. Podolski S. BleedIO Tech owns the only edge-based SaaS written for the new BLE mesh standard. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber I. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 100. Podolski S. Resilient distributed decentralized network for connected devices. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: Electronics and Microsystems. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 101. Podolski S. BLE Mesh: a new network standard, leading the evolution in smart factories. In: TechConnect World 2025 Program: Advanced Manufacturing Innovation. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/wednesday.html>
 102. Poolsappasit N, Dewri R, Ray I. Dynamic security risk management using Bayesian attack graphs. *IEEE Trans Dependable Secure Comput.* 2012;9(1):61-74.
 103. Puterman ML. Markov decision processes: discrete stochastic dynamic programming. Wiley; 1994.
 104. Ren S, Chen Y, Wang L, Zhang X. ARCS: adaptive reinforcement learning framework for automated cybersecurity incident response. *Appl Sci.* 2025;15(2):951.
 105. Ribeiro MT, Singh S, Guestrin C. Why should I trust you? Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016. p. 1135-44.
 106. Rose S, Borchert O, Mitchell S, Connelly S. Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology; 2020.
 107. Rosenberg I, Shabtai A, Elovici Y, Rokach L. Adversarial machine learning attacks and defense methods in the cyber security domain. *ACM Comput Surv.* 2021;54(5):1-36.
 108. Ruci A. Computer vision AI that digitizes current and new infrastructure in 4D. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber III. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 109. Sadman Z, Qasem A. Large language models for high-performance computing. In: TechConnect World 2025 Program: AI Innovations - Posters. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 110. Sarker IH. Cybersecurity data science: an overview from machine learning perspective. *J Big Data.* 2021;8(41):1-29.
 111. Sarker IH. Multi-aspects AI-based modeling and adversarial learning for cybersecurity intelligence and robustness: a comprehensive overview. *Secur Priv.* 2023;6(5):e295.
 112. Scarfone K, Mell P. Guide to intrusion detection and prevention systems (IDPS) (NIST Special Publication 800-94). National Institute of Standards and Technology; 2007.
 113. Sculley D, Holt G, Golovin D, Davydov E, Phillips T, Ebner D, *et al.* Hidden technical debt in machine learning systems. *Adv Neural Inf Process Syst.* 2015;28.
 114. Sharafaldin I, Lashkari AH, Ghorbani AA. Toward generating a new intrusion detection dataset and intrusion traffic characterization. In: 4th International Conference on Information Systems Security and Privacy; 2018. p. 108-16.
 115. Sharotry A. Computer vision enabled time and motion analysis. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber I. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 116. Sills R. Finite-element-based physics-informed neural networks for surrogate modeling. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: AI, Data, Cyber I. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
 117. Sills RB, Sunil P, Vasoya M. Surrogate modeling with finite-element-based physics-informed neural networks. In: TechConnect World 2025 Program: Emerging Methods for Materials Design with Artificial Intelligence II. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/tuesday.html>
 118. Solomatina K. AI-based reskilling platform for new jobs in critical industries, including energy and

- manufacturing. In: TechConnect World 2025 Program: TechConnect Innovation Spotlights: Aerospace, Space and Defense I. TechConnect; 2025. Available from: <https://www.techconnectworld.com/World2025/monday.html>
119. Sommer R, Paxson V. Outside the closed world: on using machine learning for network intrusion detection. In: 2010 IEEE Symposium on Security and Privacy; 2010. p. 305-16.
 120. Sperotto A, Schaffrath G, Sadre R, Morariu C, Pras A, Stiller B. An overview of IP flow-based intrusion detection. *IEEE Commun Surv Tutor.* 2010;12(3):343-56.
 121. Sutton RS, Barto AG. Reinforcement learning: an introduction. 2nd ed. MIT Press; 2018.
 122. Tavallaee M, Bagheri E, Lu W, Ghorbani AA. A detailed analysis of the KDD CUP 99 data set. In: 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications; 2009. p. 1-6.
 123. Ussher-Eke D, Onoja DA, Ijiga OM, Enyejo LA. Strengthening human resource compliance and ethical oversight through cybersecurity awareness and policy enforcement. *Int J Manag Commer Innov.* 2025;13(1):376-93. doi:10.5281/zenodo.16539315
 124. Zadrozny B, Elkan C. Transforming classifier scores into accurate multiclass probability estimates. In: Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2002. p. 694-9.

How to Cite This Article

Mbukio IC, Ijiga OM, Enyejo LA. Design of the intelligent security operations automation algorithm for AI-enabled security orchestration, automation, and response with comparative analysis against rule-based SOAR platform. *International Journal of Engineering and Computational Applications.* 2026;2(4):32-48.

Creative Commons (CC) License

This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution NonCommercial-ShareAlike 4.0 International (CC BYNC-SA 4.0) License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.