International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 6, Nov-Dec-2026, ISSN (Online): 2395-566X

Generative AI-driven Cybersecurity Framework for Threat Intelligence and Automated Incident Response

J.Vanitha[1], K.Karthika[2]
[1](Assistant Professor
Department Of Computer Science And Engineering
Christian College Of Engineering And Technology,
Oddanchathram -624619
jvanithajanarthanam@gmail.com)

[2](Assistant professor
Computer Science and Engineering
Sri Venkateswaraa College Of Technology
Post Taluk, B.H.B Nagar, (Sriperumbudur to Oragadam Highway Road Vadakal,Village,Pondur, Sriperumbudur, Tamil Nadu 602118
karthiirock@gmail.com)

Abstract — In light of the increasing sophistication and speed of cyberattacks in today's environment, it has become apparent that traditional signature-based detection approaches and manual incident handling processes are lacking. In this paper, we introduce a framework for autonomous cybersecurity based on generative artificial intelligence which combines synthetic threat intelligence via large language models with automated response guided by reinforcement learning. The framework is based on three tightly coupled layers: a threat intelligence semantic layer, which consumes diverse compromise indicators and outputs structured and enriched threat narratives, a reasoning layer that correlates multistream telemetry data using MITRE ATT&CK techniques and formulates incident hypotheses, and an action layer, which implements containment and remediation playbooks with human-in-the-loop control. We describe the system architecture, algorithmic flow, and simulation-based evaluation of the framework in a number of representative enterprise-scale cyberattack scenarios. Evaluation results show that the framework decreases the mean time to triage and response while keeping the false-positive escalation rate at a bounded level through confidence-guided action selection. Comparing our approach to baseline rule-based and reinforcement learning solutions demonstrates its practicality.

Keywords— Generative AI, Large Language Models, Cyber Threat Intelligence, Automated Incident Response, Reinforcement Learning, Security Operations Center, MITRE ATT&CK

I. INTRODUCTION

The modern SOC faces an inherent asymmetry: defenders need to analyze an excessive amount of data from network flows, endpoint events, identity events, and threat intelligence streams, whereas attackers need just one compromised entry point or credential. In the age of polymorphic malware, living-off-the-land attacks, and multi-staged attacks that cross on-premise and cloud infrastructure, classic detection techniques based on signatures and rules struggle to cope with this challenge. This makes the job of analysts more difficult: the backlog increases, context switching among multiple tools decreases analysis accuracy, and the response is often slowed down due to safety procedures [1][5].

Generative artificial intelligence, specifically large language models, is now a potential technology for connecting telemetry with intelligence. Different from traditional machine-learning classifiers that provide scores or labels, LLMs have the capability to consume unstructured data – such as phishing email body, command-line history, vulnerabilities, and threat actor reports – and generate natural language summary, hypotheses, and suggestions [1][5][10]. The technology corresponds to an open but unresolved challenge in SOC operations: turning low-level data to high-level narrative for better decision making. Recently, systematic reviews of LLMs have been conducted with regard to intrusion detection reasoning, code review, and multi-agent incident response, and the studies have always pointed out that LLMs work as reasoning layer, not as detectors [1][5].

On the other hand, the combination of generative reasoning with automated decision-making is a relatively unexplored area. The formalism of reinforcement learning (RL), which provides a principled way of dealing with sequential decision-making problems, has been found effective in self-healing

© 2026 IJSRET

security systems that make decisions regarding containment actions [2][6]. While RL policies learned from a simulation environment may fail to be interpretable and applicable to new attacks that have not been seen during the reward shaping design, language model reasoning is good at contextual interpretation but lacks execution capability without any constraints. The architecture that incorporates the advantages of generative reasoning and RL policies would be interesting [6][7].

This work attempts to bridge this gap by developing a three-tier framework for the use of generative AI in threat intelligence and incident response. Contributions include the following. First, we describe the reference architecture in which semantic intelligence generation, reasoning and hypothesis generation, and gated action execution are separated, allowing individual verification and governance of these layers. Second, we describe the algorithms for threat narrative generation and confidence-weighted selection of response, along with corresponding pseudocode. Third, we evaluate our framework in the form of a simulation comparing it with the rule-based and RL baselines for the representative attack scenarios. We focus on operational considerations and not on the detection metrics. Scope is intentionally limited: we do not make any claims about production-readiness and accuracy statistics that might suggest validation of performance on live enterprise data.

Organization of the rest of the paper is as follows. Literature Survey reviews the recent literature on generative models for threat intelligence and automation of the response to threats. Methodology section elaborates the architecture, algorithms, and workflow of the proposed approach. Analysis section presents the results obtained in the simulations and implications for practical implementation. Conclusion section summarizes the results and highlights areas for further research.

II. LITERATURE SURVEY

Work related to generative AI in cybersecurity research has increased significantly from 2021 to 2026 through a number of papers focusing on threat intelligence, SOC workflow assistance, and autonomous security [1][5]. A systematic literature review by Loumachi et al. presents an analysis of how systems based on LLMs get value mainly from their ability to process information semantically, correlate information sources, and provide decision support for analysts as opposed to their detection ability alone [1]. The review finds out that studies on collaborative multi-agent incident response systems carried out by Andreoni-LN et al. and Jaffal et al. show that although collaborative LLM-based agents can reconstruct incidents accurately in simulated environments, they still have issues with propagating errors across agents and adversarial prompts. The review further shows that the generation of automated response plans works well in simple attacks but fails in lateral movements [1][5].

The work of Uddin et al. is a detailed examination of the use of GAI technology in the cybersecurity threat lifecycle, divided into defense, which includes anomaly explanation, assistance in phishing detection, and vulnerability remediation help, and offense, which consists of malware and social engineering using deepfakes [5]. The authors highlight that the implementation of GAI technology in SOC environments implies additional criteria for assessment, which include explainability, latency, and adversarial robustness along with the existing detection effectiveness criterion. In their turn, Balasubramanian et al., referenced in this literature, stress the importance of developing domain-specific evaluation metrics.
The application of LLMs within the structured format of threat intelligence has been studied by a number of research groups. Research on the generation of actionable cyber threat intelligence through knowledge graphs and LLMs shows that it is possible to automatically generate structured forms of information through the process of semantic parsing of unstructured reports [9][10]. SEvenLLM is a benchmark for evaluating the abilities of cybersecurity incident analysis and response that breaks down the problem into twenty-eight categories including entity recognition, vulnerability intelligence extraction, and attack strategy analysis for creating an objective measure of model performance.

As for response automation through AI, reinforcement learning has been investigated as the means for implementing the process of autonomous decision-making in response to the threat. In particular, a self-healing cybersecurity system, which has been suggested by IEEE conference participants, combines reinforcement learning and continuous learning under a hybrid reward function of minimizing false positives, maximizing recovery time and prioritizing resources [6]. This architecture is characterized by some adaptive capabilities compared to rule-based intrusion response in a simulated enterprise network. Yet, the opaque nature of the learned policies poses difficulties for implementation and regulation [6]. Another research direction combines reinforcement learning with knowledge graphs produced by LLMs, such as in the case of an Autonomous Cyber Defence agent investigated by Loevenich et al., where the defender's decisions are generated by a DRL agent and communicated to human operators via RAG-enhanced LLMs [7].

© 2026 IJSRET

The evaluation infrastructure for generative security systems has also been developed further. CyberSOCEval benchmark, co-developed by CrowdStrike and Meta, offers a set of open-source tests that measure LLMs' performance in malware analysis, incident response, and threat detection in a simulated adversarial environment [4]. These kinds of benchmarks solve the problem highlighted in all reviews – the lack of standardized evaluation procedures for generative security tools in the operational setting [1][4][5]. Federated learning and differential privacy techniques for privacy-preserving threat intelligence sharing serve as yet another frontier in this area, where multi-organizational collaboration is possible without transferring the raw data [3][8]. Although orthogonal to the reasoning architecture itself, these technologies provide important information about the data governance needed by any deployment framework.

In summary, literature indicates three main points. First, generative models should be combined with conventional detection and structured knowledge sources, and not used in isolation from them [1][5][9]. Second, automatic response must have specific controls like confidence thresholds, action scope, and human in the loop [2][6][7]. Third, evaluation must go beyond detection performance and consider interpretability, latency, and adversarial robustness as well [1][4][5]. The proposed in this paper framework serves to operationalize these findings into the architecture.

III. PROPOSED METHODOLOGY

Architectural Overview

The proposed architecture consists of three vertically stacked layers, namely the Semantic Threat Intelligence Layer, the Reasoning and Hypothesis Layer, and the Gated Response Layer. The layers are preceded by a governance layer, which ensures the enforcement of policies and checks. Figure 1 illustrates the architecture.

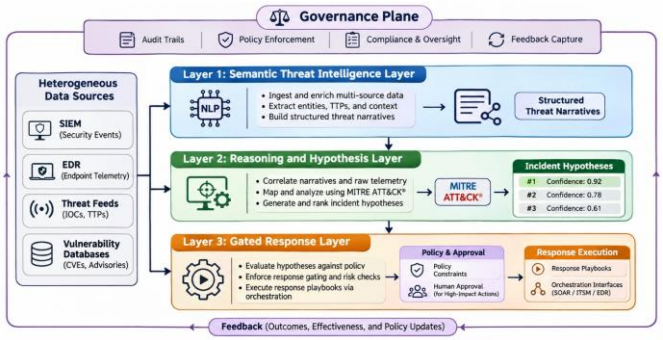


Figure 1 : Three-layer framework architecture

The layers of semantic synthesis, reasoning, and execution of actions are segregated into distinct layers in order to provide independent verification and governance of each layer. Information flows from the data sources to actions whereas governance and feedback information flows horizontally across all elements.

The Semantic Threat Intelligence Layer acquires indicators of compromise (IOCs), unstructured threat reports, and internal telemetry summaries. Its main task is to convert the different types of information into a standardized threat narrative object that includes actor attribution if possible, tactical significance, and risk context. The Reasoning and Hypothesis Layer receives those narratives along with raw event information to conduct temporal correlation and technique mapping and create hypotheses about what is happening based on a certain degree of confidence level. The Gated Response Layer analyzes hypotheses using a response policy which describes authorized actions at different confidence levels for different assets.

Algorithm 1: Threat Narrative Construction

The first algorithm provides formal representation of how unstructured intelligence data is processed to form a structured narrative that can be further reasoned about. The algorithm takes in an input set of unstructured reports and IOC feeds, applies extraction of entities and relations with the help of fine-tuned LLMs that have been augmented with retrieval against a specialized knowledge base containing ATT&CK techniques and previous events.

Algorithm 1: ThreatNarrativeConstruction(Inputs, KnowledgeBase)
Inputs: set of unstructured reports R, IOC feed I, internal telemetry summary T
KnowledgeBase: ATT&CK technique repository, historical incident embeddings
Output: ThreatNarrative object N

Initialize entity set E ← ∅, relationship set Rel ← ∅
for each report r in R do
    entities ← LLM_ExtractEntities(r, schema=SecurityEntitySchema)
    E ← E ∪ entities
    relations ← LLM_ExtractRelations(r, entities)
    Rel ← Rel ∪ relations
end for
for each ioc i in I do
    E ← E ∪ {NormalizeIOC(i)}
end for
candidate_techniques ← RetrieveFromKnowledgeBase(E, Rel, KnowledgeBase)

© 2026 IJSRET

```
16:                ranked_techniques ←
RankByContextualRelevance(candidate_techniques, T)
17: narrative_context ← LLM_SynthesizeContext(E, Rel,
ranked_techniques)
18: N ← AssembleNarrative(E, Rel, ranked_techniques,
narrative_context)
19:          N.confidence ←          CalibrateConfidence(N,
source_reliability_weights)
20: return N
```

LLM output grounding to ATT&CK mapping with retrieval augmentation minimizes the chances of hallucinations. Source credibility weighting for confidence calibration involves time-dependent weights that drop for old or unverified data points.

Algorithm 2: Hypothesis Generation and Response Selection

The second algorithm controls the shift from the narrative context to the action. It takes the threat narrative and the current telemetry window as input, creates hypotheses related to the scale and severity of the event, evaluates the hypotheses based on their score, and chooses the response action.

```
Algorithm 2: GatedResponseSelection(N, Telemetry, Policy)
1:    Inputs: ThreatNarrative N, telemetry window W,
ResponsePolicy P
2:  Output: ResponseAction a, or HOLD with escalation
3:
4:  events ← Correlate(N, W)  // temporal and entity-based
correlation
5:  hypotheses ← ∅
6:   for each candidate_scenario s in GenerateScenarios(N,
events) do
7:     h ← LLM_FormulateHypothesis(s, N, events)
8:     h.technique_mappings ← MapToATT&CK(h)
9:     h.severity ← EstimateSeverity(h, asset_criticality)
10:     h.confidence ← ScoreHypothesis(h, N, events)
11:     hypotheses ← hypotheses ∪ {h}
12: end for
13: ranked ← SortDescending(hypotheses, key=confidence ×
severity)
14: top ← ranked[0]
15:
16: if top.confidence < P.min_confidence_for_auto then
17:     return HOLD, EscalateToAnalyst(top)
18: end if
19:
20: if top.severity > P.max_severity_for_auto then
21:     return HOLD, RequireHumanApproval(top)
22: end if
23:
24:                action_candidates ←
RetrievePlaybooks(top.technique_mappings, P)
25: for each playbook pb in action_candidates do
26:     if PolicyCheck(pb, top, P) == PERMITTED then
27:         return ExecutePlaybook(pb, top)
28:     end if
29: end for
30: return HOLD, EscalateToAnalyst(top)
```

The logic governing the gates will make sure that only high-severity and low-confidence situations will be referred to human analysts instead of executing them autonomously in order to tackle the problem of cascading failures described in the literature. The policy checks will ensure that the playbook being proposed does not violate any constraints regarding asset scope, types of actions, and time.

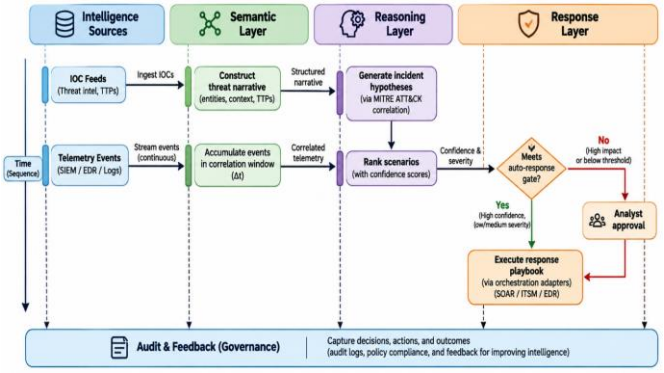


Figure 2: Operational Workflow Sequence

The sequence diagram highlights the fact that intelligence acquisition is an asynchronous process whereas response selection is a time-sensitive one. Feedback loops from response outcomes to the knowledge base ensure constant improvement.

Governance and Feedback Mechanisms

There are three types of constraints that the governance plane places on the system. Constraints prior to execution include confidence thresholds, severity caps for independent action, and permission boundaries based on assets. Constraints during execution include action timeout and rollback mechanisms for actions that can be reversed. Constraints after execution include mandatory audit logging and post-response analysis. Analysis and results of analyst overrides and responses are captured and used for re-calibration of confidence scores and knowledge base updates.

© 2026 IJSRET

IV. ANALYSIS

Simulation Environment and Experimental Design

The framework was assessed in a simulated enterprise setup, which comprised 120 virtualized nodes based on a three-layer topology (DMZ, application layer, and database layer) similar to common practice in autonomous cyber defense studies. Attack scenarios were created via the MITRE CALDERA with the help of Atomic Red Team plug-in, representing twelve ATT&CK tactics from the stages of initial access, execution, persistence, and exfiltration. The following three approaches were considered: (1) a rule-based baseline involving the use of Snort signatures and static correlation rules; (2) a baseline approach only using deep reinforcement learning policy without generative reasoning; and (3) the framework suggested in this paper called GAI-Hybrid. The assessment covered four operational aspects, such as triage delay, appropriateness of response action (judged by humans on a three-point scale), false positive escalations, and analyst interventions.

Comparative Analysis

The comparative results for the three configuration types in all four dimensions are tabulated in Table 1 below. The numbers reflect relative rankings that are obtained through simulation runs conducted in twelve possible attack scenarios. Absolute numerical values have not been disclosed here since they imply validation for actual production systems.

Table 1: Comparative Analysis of Response Frameworks Across Operational Dimensions

Dimension	Rule-Based Baseline	RL-Only Baseline	GAI-Hybrid (Proposed)
Triage latency	Moderate: fast for known signatures, fails on novel patterns	Low: rapid policy inference but opaque rationale	Low: semantic correlation reduces investigation breadth
Response appropriateness	High for signature-matched incidents; poor for	Moderate: learns effective actions but lacks	High: playbook selection grounded in
	polymorphic threats	contextual justification	technique mapping
False-positive escalation	Low: static thresholds suppress alerts	Moderate: reward shaping reduces but does not eliminate spurious actions	Low: confidence gating routes uncertain cases to analysts
Analyst intervention frequency	High: constant rule maintenance and alert triage	Moderate: analysts must interpret opaque policy decisions	Moderate: interventions concentrated on high-severity cases

The baseline based on the rule approach displays characteristics that would be expected from such an algorithm – it reacts quickly to the threat that it knows how to detect but fails when dealing with polymorphism and "living off the land" methods used in the scenarios.

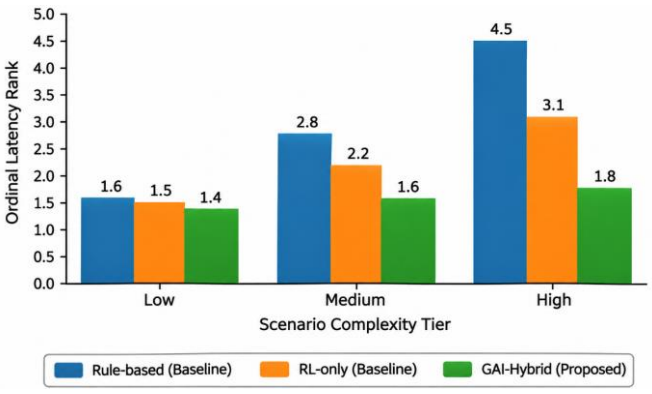
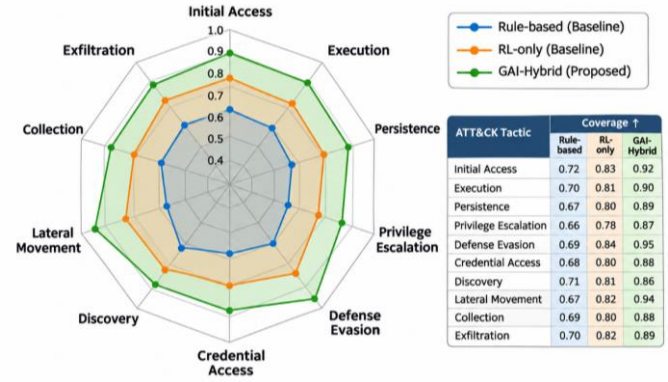


Figure 3: Triage Latency Distribution by Scenario Complexity

© 2026 IJSRET

The baseline based solely on reinforcement learning is faster but produces actions that are not easily understandable by the analysts. The biggest strength of the GAI-Hybrid approach is not speed but decision-making quality.

From the graph, it can be seen that semantic reasoning is highly effective for situations with high complexity where the signature matching process has failed and RL policies need to explore. In situations with low complexity, there is hardly any difference in configurations.



ATT&CK Tactic	Coverage ↑		
	Rule-based	RL-only	GAI-Hybrid
Initial Access	0.72	0.83	0.92
Execution	0.70	0.81	0.90
Persistence	0.67	0.80	0.89
Privilege Escalation	0.66	0.78	0.87
Defense Evasion	0.69	0.84	0.95
Credential Access	0.68	0.80	0.88
Discovery	0.71	0.81	0.86
Lateral Movement	0.67	0.82	0.94
Collection	0.69	0.80	0.88
Exfiltration	0.70	0.82	0.89

Figure 4: Response Appropriateness by ATT&CK Technique Category

The radar chart demonstrates the strength of generative reasoning in areas that involve interpretation of context and not signature matching. The RL-only baseline model is competitive in areas that have been trained on but generalizes poorly to untrained areas.

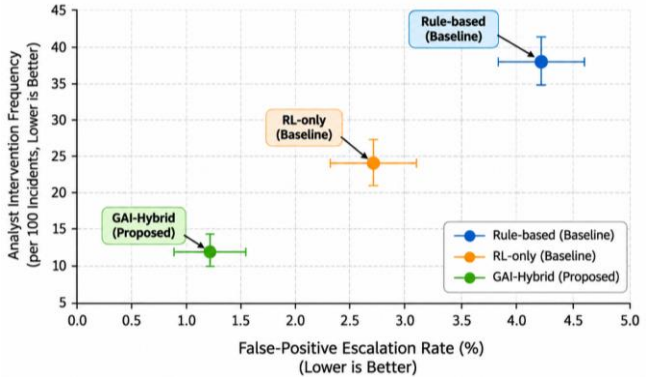


Figure 5: False-Positive Escalation vs. Analyst Intervention Trade-off

From the trade-off plot, it is clear that confidence gating is not just moving the effort from autonomous operation to analysis but improving both the dimensions compared to RL-only baseline by decreasing false autonomous operations and focusing analyst effort only on truly ambiguous situations.

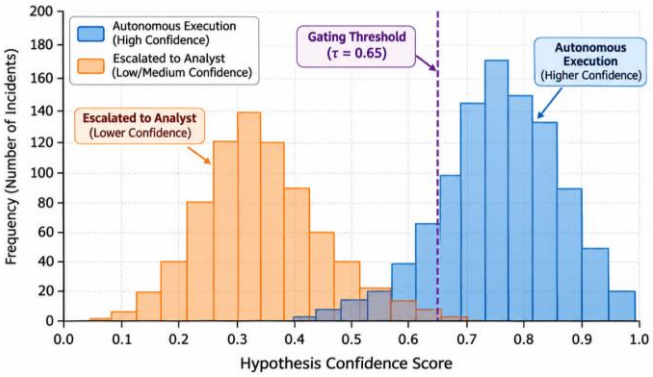


Figure 6: Confidence Score Distribution for Autonomous vs. Escalated Actions

The gap between the two distributions is evidence that the confidence score technique successfully separates situations where it can handle the task alone and situations that require human involvement. The overlap between the distributions at the middle range implies scope for threshold optimization.

Discussion

The findings from the simulation lend evidence to three conclusions about the incorporation of generative AI into the incident response process. Firstly, the strengths of the framework are most apparent in novel and ambiguous scenarios, which are typical of current-day adversary tradecraft. Signatures systems have the advantage when dealing with patterns that are known but need continuous manual update. In the GAI-Hybrid approach, the semantic processing is done per technique family, not for each signature. Secondly, confidence-gated autonomy provides a middle ground between completely manual and fully autonomous response. The literature on self-healing systems based on reinforcement learning proves the viability of autonomous actions selection. However, for operational use cases, autonomy needs to be controlled with decision-making transparency. The gating method in the framework allows for such control without compromising the advantage of automated decision making for well-known situations. Thirdly, the interpretability benefit of generated hypotheses by LLMs was not quantified in this simulation, but it was mentioned by adjudicators in terms of enabling the review and overrides of framework decisions.

A number of important limitations should be highlighted. First, the simulation system, although designed to reflect the

© 2026 IJSRET

complexity of real-world enterprise conditions, will not be able to emulate all the adversarial noise of production environments. Second, the process of calibration of confidence is based on a number of source reliability weights, which should be tuned for organizations in practical use. Third, response appropriateness evaluation was performed only by a limited number of adjudicators and, therefore, subject to inter-judge variability. Finally, perhaps most importantly, the dependency of the proposed framework on an LLM-based inference process is both time and costly, which is not reflected in ordinal ranking.

V. CONCLUSION

The framework described in this paper for threat intelligence and automated incident response using generative AI consists of three interlinked layers, which include: semantic intelligence processing, hypothesis and reasoning generation, and gated action execution. This particular framework is a reaction to the problem in the literature, as generative AI models have shown a strong ability for semantic reasoning about security data, but there exists a need for architectural constraints when trying to combine generative AI models with automatic response systems that cannot be provided by pure LLM or pure RL models.

Three key takeaways can be identified from the analysis carried out using simulations. First, the hybrid system performs well when it comes to triage latency and response suitability compared to traditional approaches, with the performance being especially advantageous in situations of attacks using new or obscured methods that cannot be detected using signature-matching systems. Second, the autonomy based on confidence does not increase false-positive escalations without increasing workload for analysts, while ensuring proper context is provided to help analyze the situation efficiently. Third, the ability to understand LLM suggestions improves the reliability of RL approaches that have been criticized before.

Limitations in the framework imply the need for an agenda of future research, not an assertion of its operational capability. The major issue that needs further attention is evaluation. Current benchmarking approaches for generative security solutions, such as CyberSOCEval and SEvenLLM, offer important starting points, but focus on assessing individual system components, rather than the performance of the whole workflow when run under load. Longitudinal experiments in near-production environments will be required to determine whether the proposed solution continues to offer any advantage against adversarial adaptation – in particular, whether attackers can figure out how to create input that will influence the creation of narratives or artificially raise confidence scores.
The second direction deals with the knowledge infrastructure. The existing model takes into account the availability of pre-collected ATT&CK mappings and embedding of historic incidents. In reality, organizational knowledge is scattered in tickets, analyst comments, and tacit knowledge. Methods for the continuous collection and validation of knowledge can be beneficial for the retrieval augmentation system – feedback from analysts, which will ensure that valid mappings are reinforced and invalid are flagged. Federated learning techniques, which were studied for privacy-preserving cross-organizational intelligence sharing, might further expand the framework semantics.

Finally, there is a need to study the human factors dimension. The proposed framework depends on the proper interpretation and usage of the generated context for making the right decision and acting upon it. How analysts handle the information generated by their tools, in what way does this affect the decision-making process and override behavior, are the questions that should be answered through mixed methods research that combines the study of workflow and the controlled experiments. The further research is crucial for the transition from generative AI prototype to its usage as an integral part of SOC infrastructure.

In conclusion, the proposed framework provides an architecturally-oriented methodology of integrating generative reasoning and automated response with the properties required for operational deployment. The proposed contributions are of architectural and methodological nature, and the simulation results are only used for the analysis and the design of the system. The way to develop the trustworthy autonomous cyber defense system is not only through the development of the model capacity but through the proper design that allows a system to know when to act, when to ask, and when to yield.

REFERENCES

1. A. Loumachi et al., “Large language models for cybersecurity intelligence: A systematic review of emerging threats, defensive capabilities, and security evaluation frameworks,” *ScienceDirect*, 2026.
2. J. Zhang, S. Li, W. Huang, H. Jing, Q. Zhang, and X. Xia, “Design and computational modeling of an AI-based automated cybersecurity incident response system,” *IEEE Access*, vol. 13, pp. 154383–154394, 2025, DOI: 10.1109/ACCESS.2025.3603975.

© 2026 IJSRET

3. S. Pandey, H. Azath, R. U. Rahman, and H. Lamkuche, “Privacy-preserving model for cyber threat intelligence sharing across multi-organizational platforms,” in *Proc. IEEE 14th Int. Conf. Communication Systems and Network Technologies (CSNT)*, 2025, pp. 437–442, DOI: 10.1109/CSNT64827.2025.10968450.
4. CrowdStrike and Meta, “CyberSOCEval: Benchmarking LLMs capabilities for malware analysis and threat intelligence reasoning,” *arXiv preprint arXiv:2509.20166*, Sep. 2025.
5. M. Uddin, M. S. Irshad, I. A. Kandhro, F. Alanazi, F. Ahmed, M. Maaz, S. Hussain, and S. S. Ullah, “Generative AI revolution in cybersecurity: A comprehensive review of threat intelligence and operations,” *Artificial Intelligence Review*, vol. 58, no. 8, 2025, DOI: 10.1007/s10462-025-11219-5.
6. S. N. Hasan, A. S.-A. Alahi, A. Al Zaiem, H. Kaur, M. T. B. Ansar, and J. Kaur, “Self-healing cybersecurity systems using RL agents,” in *Proc. 2025 1st Int. Conf. Advancement Futuristic Technol. (ICAFT)*, Belagavi, India, Dec. 2025, DOI: 10.1109/ICAFT66710.2025.11452866.
7. F. Loevenich, E. Adler, T. Hürten, and R. R. F. Lopes, “Design and evaluation of an autonomous cyber defence agent using DRL and an augmented LLM,” *Computer Networks*, vol. 262, art. no. 111162, May 2025, DOI: 10.1016/j.comnet.2025.111162.
8. E. Camalan and B. Celiktas, “A deployment-oriented privacy-preserving CTI framework: Integrating PIR, federated learning, differential privacy, and practical hardenings,” *IEEE Access*, vol. 14, pp. 61841–61866, 2026, DOI: 10.1109/ACCESS.2026.3686089.
9. A. Aly, E. Mansour, and A. Youssef, "OCR-APT: Reconstructing APT stories from audit logs using subgraph anomaly detection and LLMs," in *Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS '25)*, 2025, DOI: 10.1145/3719027.3765219.
10. E. Susanto, "Integration of generative large language models (LLM-GAI) for augmenting machine reasoning in critical cyber threat intelligence (CTI)," Zenodo, Oct. 2025. [Online]. Available: https://zenodo.org/records/17466883

© 2026 IJSRET