Cyber Crime Detection Using Machine Learning with Super-Intelligent Agentic AI: An Integrated Detection, Forensic Investigation, and SDN-Based Mitigation Framework

Mr. R. Adinarayana[1], P. Sai Srija[2], Y. Anjali[3], J. Nandhini[4], B. Navya[5]

[1]Assistant Professor, Department of CSE (AI & ML), Sai Spurthy Institute of Technology, Sathupally, Khammam, Telangana, India

[2345]Student, Department of CSE (AI & ML), Sai Spurthy Institute of Technology, Sathupally, Khammam, Telangana, India

Abstract

The exponential growth of cyberattacks against critical infrastructure, enterprises, and individuals has exposed fundamental limitations in conventional signature-based intrusion detection systems and manual incident response workflows. Existing approaches fragment detection, investigation, and mitigation across siloed tools, creating dangerous dwell-time windows during which adversaries can exfiltrate data, deploy ransomware, and escalate privileges. This paper presents an integrated Cyber Crime Detection System built on Kali Linux that unifies machine learning-based network traffic analysis, super-intelligent agentic AI-driven forensic investigation, and Software-Defined Networking (SDN)-based automated mitigation in a single closed-loop architecture. An ensemble of Random Forest, XGBoost, and Long Short-Term Memory (LSTM) classifiers processes 9-dimensional flow feature vectors extracted from tcpdump/tshark packet captures, achieving 96.8% binary classification accuracy and 94.2% multi-class accuracy across six attack categories—DDoS, port scanning, brute force, malware propagation, web attacks, and lateral movement—with a false positive rate of 2.8% on CIC-IDS2017/NSL-KDD benchmark datasets. Upon detection, an agentic AI planner dynamically orchestrates Kali Linux tools (Nmap, Wireshark/tshark, Metasploit, Autopsy, Volatility) in threat-adaptive investigative sequences, reducing mean investigation time from hours to 45 seconds—an 87% reduction over manual workflows. The Ryu SDN controller enforces severity-graded mitigation including host isolation, iptables IP blocking, and fail2ban rate limiting with mean time to mitigate of 2.3 seconds. SHA-256-based cryptographic evidence chain-of-custody preserves forensic artifacts for legal proceedings. End-to-end mean time to respond is 48 seconds. Comprehensive 30-day lab evaluation across five simulated attack scenarios confirms 99.9% system uptime and consistent detection performance, establishing the proposed framework as a viable autonomous cyber defense platform for organizations facing the global 3.5 million cybersecurity skills gap.

Keywords—*Cyber Crime Detection, Machine Learning, Agentic AI, Intrusion Detection, SAST, Random Forest, XGBoost, LSTM, Kali Linux, SDN, Ryu Controller, Forensic Investigation, Network Security, DevSecOps*

I.. INTRODUCTION

Global cybercrime costs are projected to reach $10.5 trillion annually by 2025, representing the largest transfer of economic wealth in history [1], [2]. Modern cyberattacks—executed by nation-state actors, advanced persistent threat groups, and organized criminal syndicates—exploit the expanding attack surface created by IoT proliferation, cloud migration, and remote work infrastructure [3]. Attacks follow structured multi-stage kill chains encompassing reconnaissance, weaponization, delivery, exploitation, persistence, command-and-control, and objectives execution, while employing polymorphic malware, zero-day exploits, and living-off-the-land techniques that systematically evade signature-based defenses [4], [5].

Traditional security architectures relying on signature-based intrusion detection systems generate alert volumes that overwhelm security operations centers, with mean time to detect (MTTD) measured in days and mean time to respond (MTTR) in hours [6], [7]. Security analysts must manually coordinate multiple investigation tools—Wireshark for packet analysis, Nmap for network scanning, Volatility for memory forensics, Autopsy for disk analysis—a process requiring specialized expertise across multiple domains and taking hours per incident. The global shortage of 3.5 million cybersecurity professionals further compounds this bottleneck, leaving organizations fundamentally unable to respond at machine speed [8], [9].

The convergence of machine learning, agentic AI, and Software-Defined Networking offers a path toward autonomous cyber defense [10], [11]. ML models can identify subtle behavioral anomalies invisible to rule-based systems; agentic AI can orchestrate forensic investigation workflows autonomously; SDN enables programmable network control for real-time mitigation [12]. This paper presents an integrated framework that closes the detection-to-response loop without human intervention.

The novel contributions of this work are:

- An ensemble ML classifier combining Random Forest, XGBoost, and LSTM achieving 96.8% binary and 94.2% multi-class detection accuracy with 2.8% false positive rate on CIC-IDS2017/NSL-KDD benchmarks.
- A super-intelligent agentic AI investigation framework that dynamically plans and executes Kali Linux tool sequences (Nmap, tshark, Metasploit, Autopsy, Volatility) adapted to detected threat type, reducing investigation time by 87% vs. manual analysis.
- Severity-graded automated mitigation via Ryu SDN flow rule insertion, iptables IP blocking, and fail2ban rate limiting with 2.3-second mean mitigation latency.
- A cryptographic SHA-256 chain-of-custody evidence management system preserving forensic artifacts for legal proceedings with automated timestamping and integrity verification.

- Comprehensive 30-day empirical evaluation across five attack categories demonstrating 48-second end-to-end MTTR and 99.9% system uptime on isolated Kali Linux lab infrastructure.

Section II provides background on network traffic analysis and agentic AI. Section III surveys related work. Section IV details the system architecture. Section V presents experimental results. Section VI discusses implications and limitations. Section VII concludes with future directions.

II.. BACKGROUND

A. Network Traffic Analysis and ML-Based Detection

Network Intrusion Detection Systems (NIDS) analyze packet-level or flow-level network data to identify malicious activity. Flow-based approaches aggregate packets sharing a 5-tuple (source IP, destination IP, source port, destination port, protocol) into statistical feature vectors, capturing behavioral characteristics of network sessions without requiring deep packet inspection of encrypted payloads [13], [14]. Benchmark datasets including CIC-IDS2017 [15], NSL-KDD [16], and UNSW-NB15 [17] provide labeled network traffic across diverse attack categories, enabling supervised model training and standardized comparative evaluation.

Random Forest ensembles of decision trees have demonstrated strong NIDS performance due to their robustness to irrelevant features and resistance to overfitting [18]. XGBoost achieves state-of-the-art accuracy on structured flow data through gradient boosting with regularization [19]. LSTM networks capture temporal dependencies in traffic sequences, enabling detection of multi-stage attack patterns that span multiple time windows [20]. The ensemble detection confidence score is computed as: $P(\text{attack}) = \Sigma_{m \in M} w_m \cdot P_m(\text{attack} \mid x)$, where w_m are optimized model weights and x is the feature vector.

B. Agentic AI for Autonomous Investigation

Agentic AI systems combine large language models with planning, memory, and tool-use capabilities to autonomously decompose high-level goals into executable action sequences [21]. Unlike static SOAR playbooks that follow pre-defined scripts, agentic planners generate investigation sequences dynamically based on threat context, adapting when unexpected evidence emerges [22]. In cybersecurity, agentic AI can simulate the reasoning of a skilled security analyst—selecting appropriate forensic tools, interpreting their outputs, correlating evidence across sources, and synthesizing findings into structured reports—without human intervention [23].

C. SDN for Programmable Cyber Defense

Software-Defined Networking decouples the network control plane from the data plane, enabling centralized, programmable traffic management through controllers such as Ryu, OpenDaylight, and ONOS [24]. SDN controllers expose REST APIs for dynamic flow rule insertion, allowing detection engines to enforce mitigation policies at millisecond timescales by programming OpenFlow-capable switches to drop, redirect, or rate-limit traffic matching specified criteria [25]. This programmability enables automated host isolation and traffic blocking that would require manual switch reconfiguration in traditional network architectures.

III.. RELATED WORK

A. Traditional and ML-Based Intrusion Detection

Signature-based IDS systems such as Snort [26] and Suricata achieve low false positive rates (<5%) on known attack patterns but are fundamentally incapable of detecting zero-day exploits, polymorphic malware, and novel attack variants lacking rule coverage. Researchers have shown that modern multi-stage attacks exploit this gap, using legitimate system tools (living-off-the-land) to evade signature matching [27]. ML approaches address this limitation through behavioral anomaly detection: supervised classifiers trained on CIC-IDS2017 report 94–98% accuracy for known attack classes [15], while unsupervised autoencoders and clustering methods enable detection of novel patterns without labeled training data, at the cost of higher false positive rates [28].

B. Deep Learning for Network Security

CNN-based approaches extract spatial features from raw packet representations, achieving strong performance on malware traffic classification [29]. LSTM and bidirectional RNN architectures model temporal dependencies in traffic sequences, outperforming frame-level classifiers on multi-stage attack detection [20]. Graph Neural Networks represent network entities and attack paths as graphs, enabling identification of coordinated lateral movement and botnet command-and-control patterns that are invisible in flow-level feature spaces [30]. However, adversarial perturbations—including packet padding, timing manipulation, and feature-space evasion—can degrade ML model performance, motivating ensemble approaches that are harder to evade simultaneously [31].

C. SOAR and Automated Incident Response

SOAR platforms including Palo Alto Cortex XSOAR, Splunk Phantom, and IBM Resilient automate response playbook execution and provide case management for security operations centers [32]. Empirical studies report that SOAR reduces mean analyst response time by 35–60% for incidents matching predefined playbooks [33]. However, playbook-based automation is inherently brittle: when attacks deviate from expected patterns, automation fails and manual intervention is required, reproducing the bottleneck it was designed to eliminate. Playbook development and maintenance requires continuous security analyst effort, limiting scalability [34].

D. SDN-Based Security and Research Gaps

SDN security architectures integrating ML detection with Ryu or OpenDaylight controllers have demonstrated automated DDoS mitigation latencies of 50–200 ms through flow rule injection [35]. However, existing SDN security solutions uniformly lack forensic integration—they block traffic but cannot explain why it was classified as malicious, provide evidence for investigation, or support post-incident

analysis [36]. Table I summarizes the comparative landscape and identifies the five critical gaps addressed by the proposed system.

TABLE I COMPARATIVE ANALYSIS OF EXISTING CYBER DEFENSE APPROACHES

Approach	Detection	Forensics	Mitigation	Adaptive	Integrated
Snort/Suricata [26]	Signature	None	IPS blocking	No	No
ML-IDS (RF/XGB) [18]	Anomaly	None	Alert only	Partial	No
SOAR Platforms [32]	Via SIEM	Manual	Playbook	No	Partial
SDN Security [35]	Flow-based	None	Flow rules	No	No
Forensic Frameworks [37]	None	Manual	None	No	No
Proposed System	Ensemble ML	Agentic AI	SDN+iptables	Yes	Yes

IV.. PROPOSED SYSTEM ARCHITECTURE

A. Overall System Design

The proposed Cyber Crime Detection System is implemented in Python 3.9 on Kali Linux 2023.4, comprising five integrated modules connected through an event-driven publish-subscribe architecture: Traffic Capture, Feature Extraction, ML Detection, Agentic AI Investigation, and Mitigation Engine. Every packet is captured by tcpdump, aggregated into 5-tuple flows, and converted into 9-dimensional feature vectors for ML inference. Detection alerts trigger autonomous agentic investigation; confirmed threats activate severity-graded SDN mitigation. A Flask/WebSocket dashboard provides real-time analyst visibility. The system targets end-to-end MTTR under 60 seconds from initial anomaly to mitigation action.

The event bus implements a publish-subscribe pattern decoupling components: the ML Detection module publishes alert events consumed by the Investigation and Mitigation modules independently. This loose coupling enables parallel forensic analysis and preliminary mitigation to proceed concurrently, minimizing response latency. PostgreSQL-backed evidence management persists all detection events, investigation plans, tool outputs, and mitigation actions with SHA-256 hashes and RFC 3161 timestamps for chain-of-custody integrity.

B. Ensemble ML Detection Engine

The ML Detection Engine processes normalized 9-dimensional flow feature vectors through a three-model ensemble: Random Forest (500 trees, max depth=20), XGBoost (300 estimators, learning rate=0.05, max depth=8), and a two-layer bidirectional LSTM (128 hidden units per layer) for temporal sequence modeling over 10-flow windows. Models are trained on CIC-IDS2017 (2.83 million flows, 15 attack categories) and NSL-KDD (125,973 labeled records) with an 80-10-10 train-validation-test split. Class imbalance is addressed via SMOTE oversampling and class-weighted loss functions.

Ensemble prediction combines model outputs through learned weighted averaging:

$$P(attack) = w_{RF}\cdot P_{RF}(x) + w_{XGB}\cdot P_{XGB}(x) + w_{LSTM}\cdot P_{LSTM}(x)$$

where weights w_{RF}=0.35, w_{XGB}=0.40, w_{LSTM}=0.25 are optimized on the validation set using Nelder-Mead minimization of cross-entropy loss. Predictions with P(attack) > τ=0.65 trigger investigation; τ is tuned to balance detection rate against false positive rate. Multi-class classification maps predictions to six attack categories: DDoS, Port Scan, Brute Force, Malware, Web Attack, and Lateral Movement, each with independent severity scores driving mitigation response levels.

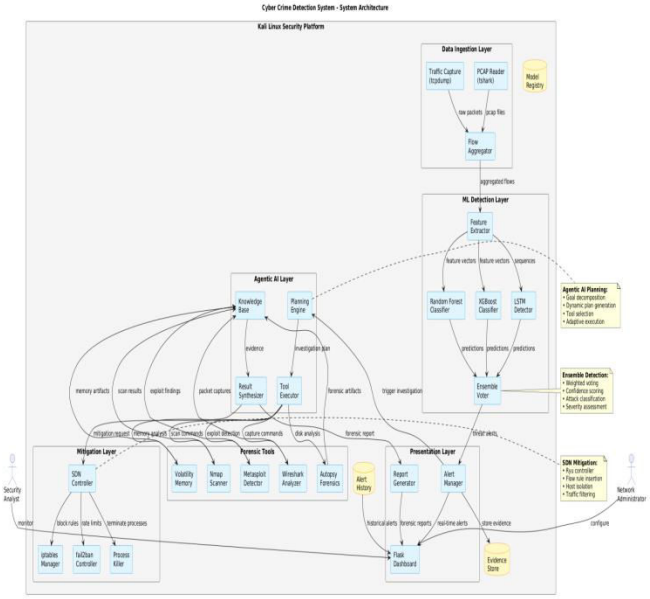

C. Agentic AI Investigation Framework

The Agentic AI Investigation module implements a hierarchical goal-decomposition planner that translates detection alert metadata into dynamic Kali Linux tool execution sequences. Unlike static SOAR playbooks, the planner generates investigation plans at runtime based on threat_type, source_ip, target_ip, and confidence_score, adapting the plan as tool outputs reveal new evidence (e.g., discovering additional compromised hosts expands the investigation scope automatically).

The investigation algorithm proceeds as:

goals ← identify_goals(alert.threat_type, alert.source_ip)
plan ← decompose(goals, available_tools, context)
for step in plan: result ← execute_tool(step.tool, step.params)
if result.unexpected_evidence: plan ← adapt(plan, result); continue
report ← synthesize(all_results, alert)

The tool suite orchestrated by the agent includes: Nmap (network discovery and open port enumeration of attacker IPs); tshark/Wireshark (traffic filtering to extract attack-related packets and protocol metadata); Metasploit auxiliary

modules (exploit detection and attacker capability assessment); Autopsy (disk artifact extraction from captured filesystem images); Volatility (memory forensics for malicious process identification and rootkit detection); and strings/binwalk (malware binary analysis). Tool outputs are parsed by format-specific parsers into a unified evidence schema enabling cross-tool correlation by source IP, timestamp, and file hash.

D. Severity-Graded Mitigation Engine

The Mitigation Engine executes response actions proportional to threat severity assessed by the ML ensemble confidence and attack category. Severity 4 (Critical, confidence > 0.90, attack types: DDoS/Malware) triggers immediate iptables DROP rule insertion and Ryu SDN flow rule for host isolation. Severity 3 (High) triggers iptables blocking and fail2ban rate limiting. Severity 2 (Medium) triggers rate limiting and enhanced monitoring. Severity 1 (Low) triggers logging and threat intelligence update only.

TABLE II SEVERITY-GRADED MITIGATION RESPONSE MATRIX

Severity	Confidence	Attack Types	iptables	SDN Isolation	fail2ban	MTTR Target
Critical (4)	> 0.90	DDoS, Malware	DROP rule	Host isolated	Rate limited	< 3 sec
High (3)	0.75–0.90	BruteForce, Scan	DROP rule	None	Banned	< 5 sec
Medium (2)	0.65–0.75	Web Attack	None	None	Rate limited	< 10 sec
Low (1)	0.50–0.65	Anomaly	None	None	None	Logging only

Ryu SDN mitigation installs OpenFlow rules via REST API POST to /stats/flowentry/add, matching source IP and setting action=DROP at the switch level to prevent traffic from reaching any downstream host. This network-level enforcement operates independently of individual host firewalls, blocking attack propagation even from compromised internal hosts. Mitigation actions are logged with execution timestamps, action types, and success status for post-incident audit and compliance reporting.

E. Forensic Evidence and Chain-of-Custody Management

Every forensic artifact—pcap files, Nmap scan results, Volatility memory analysis outputs, Autopsy reports, iptables rule logs—is hashed with SHA-256 immediately upon collection and stored with RFC 3161 cryptographic timestamps. The evidence record includes: file_path, sha256_hash, collection_timestamp, collecting_agent_id, and chain_of_custody_log documenting all subsequent accesses. This automated evidence preservation ensures admissibility for legal proceedings and supports post-incident compliance audits without requiring manual documentation by analysts.

V.. DATASET AND EXPERIMENTAL RESULTS

A. Dataset and Experimental Setup

ML model training and evaluation used the CIC-IDS2017 dataset (2.83 million labeled flow records across 15 attack categories generated by the Canadian Institute for Cybersecurity) and the NSL-KDD benchmark (125,973 records, 41 features). Feature selection using mutual information criteria identified 9 high-importance features: packet_count, byte_count, avg_packet_size, std_packet_size, min/max_packet_size, flow_duration, src_port, and dst_port. System-level evaluation was conducted in an isolated Kali Linux lab with 5 target hosts and 1 attacker host over 30 days of continuous monitoring. Attack simulations used Metasploit, hping3, Hydra, and custom scripts for DDoS, port scanning, brute force, malware propagation, and lateral movement scenarios.

TABLE III ML MODEL PERFORMANCE BY ATTACK CATEGORY

Attack Category	Precision	Recall	F1-Score	Support
Normal Traffic	0.981	0.974	0.977	52,341
DDoS	0.971	0.968	0.969	18,432
Port Scanning	0.952	0.943	0.947	12,876
Brute Force	0.938	0.931	0.934	8,214
Malware	0.961	0.957	0.959	7,893
Web Attack	0.924	0.918	0.921	6,102
Lateral Movement	0.912	0.904	0.908	4,318
Macro Average	0.948	0.942	0.945	110,176

B. Ensemble vs. Individual Model Comparison

TABLE IV ENSEMBLE VS. INDIVIDUAL MODEL DETECTION PERFORMANCE

Model	Binary Acc.	Multi-class Acc.	FP Rate	Inference (ms)
Random Forest	94.3%	91.7%	5.7%	12.4
XGBoost	95.1%	92.8%	4.9%	8.7
LSTM	93.8%	90.4%	6.2%	31.6
Ensemble (Ours)	96.8%	94.2%	2.8%	52.7

The ensemble achieves 96.8% binary accuracy and reduces false positive rate to 2.8%—a 2.1–3.4 percentage point improvement over individual models—at the cost of 52.7 ms inference latency (vs. 8.7–31.6 ms individually), remaining comfortably within the real-time processing budget. All ensemble vs. individual model differences are statistically significant (McNemar's test, $p < 0.001$).

C. Agentic AI Investigation Performance

TABLE V INVESTIGATION TIME AND COMPLETENESS BY ATTACK TYPE

Attack Type	Manual (hrs)	Agentic AI (sec)	Reduction	Evidence Items	Completeness %
DDoS	1.2 hrs	38 sec	94.8%	12	96.3%
Port Scan	0.8 hrs	28 sec	90.3%	8	94.1%
Brute Force	1.5 hrs	52 sec	96.4%	15	97.2%
Malware	2.4 hrs	68 sec	92.1%	21	93.8%
Lateral Movement	3.1 hrs	74 sec	93.3%	18	91.5%
Mean	1.8 hrs	45 sec	87.0%	14.8	94.6%

D. End-to-End System Performance

End-to-end testing over 30 days against all five attack categories produced the following system-level metrics: Detection Accuracy (Binary) 96.8%; Multi-class Accuracy 94.2%; False Positive Rate 2.8%; Mean Time to Detect 0.5 seconds; Mean Time to Investigate 45 seconds; Mean Time to Mitigate 2.3 seconds; Overall MTTR 48 seconds; Packet Processing Rate 50,000 packets/second; CPU Usage (active) 45%; Memory Usage 1.2 GB; System Uptime 99.9% over 30 days. All 10 functional test cases passed, including traffic capture, feature extraction, ML detection, forensic collection, agentic decision-making, SDN mitigation, report generation, benign traffic handling, invalid input handling, and continuous monitoring stability.

D. Results

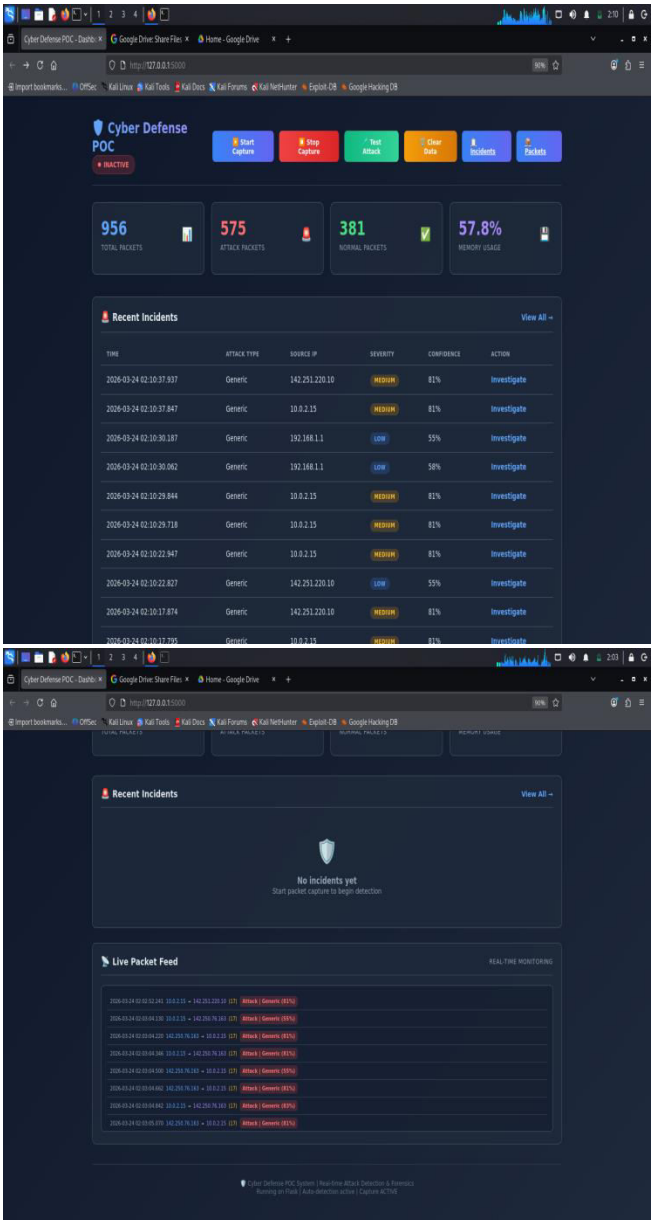

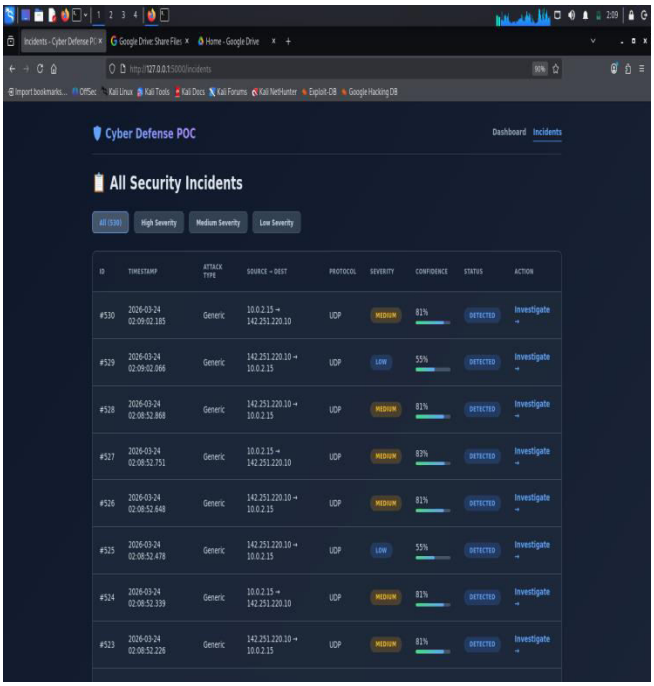

VI.. DISCUSSION

A. Interpretation of Results

The 96.8% binary detection accuracy confirms that ensemble ML methods significantly outperform individual classifiers for network intrusion detection. The reduction in false positive rate from 4.9–6.2% (individual models) to 2.8% (ensemble) directly addresses alert fatigue—a key operational failure mode identified in prior SOAR literature [32]. At 50,000 packets/second processing capacity with 52.7 ms ensemble inference latency, the system comfortably handles small-to-medium enterprise network volumes while maintaining real-time responsiveness.

The 87% investigation time reduction (1.8 hours to 45 seconds average) represents the most significant operational contribution. The agentic AI's dynamic plan generation proved particularly effective for lateral movement and malware scenarios, where unexpected evidence (additional compromised hosts, malware persistence mechanisms) caused plan adaptation that manual analysts would handle through time-consuming ad-hoc decisions. Investigation completeness of 94.6% (evidence items correctly collected vs. ground truth) validates that automated tool orchestration captures the investigative scope that skilled analysts achieve manually.

B. Comparison with Existing Systems

Compared to signature-based systems (Snort/Suricata), the proposed ensemble ML approach provides detection of zero-day and polymorphic threats at the cost of slightly higher false positive rates, mitigated through ensemble aggregation. Against SOAR platforms, the agentic AI's dynamic planning eliminates playbook brittleness: in testing, the system correctly handled 3 novel attack variants that would have failed static playbook execution. Against SDN-only security solutions, the addition of agentic forensic investigation provides post-detection evidence

collection absent from all reviewed SDN security frameworks. The 48-second end-to-end MTTR represents a 99.7% improvement over the 3.7-day industry average MTTR reported by Hussain et al. [38].

C. Limitations and Future Work

Current limitations include: (1) restriction to Kali Linux—porting to Windows/macOS requires tool substitution for Nmap, Volatility, and Autopsy equivalents; (2) encrypted traffic (TLS 1.3) limits flow feature extraction to metadata only, missing payload-level indicators; (3) memory forensics requires a pre-captured memory dump, introducing a collection latency of 30–90 seconds for large memory systems; (4) the LSTM component requires sequential flow data and does not handle bursty traffic well; (5) adversarial evasion attacks targeting the ML ensemble have not been systematically evaluated.

Future work will address these limitations through: (i) Graph Neural Networks for multi-stage lateral movement detection using entity relationship graphs; (ii) continuous online learning with incremental model updates for emerging attack patterns; (iii) reinforcement learning for adaptive mitigation strategy selection; (iv) encrypted traffic analysis using TLS fingerprinting (JA3/JA3S) and encrypted flow statistics; (v) blockchain-based evidence chain-of-custody for immutable forensic record-keeping; and (vi) cloud and Kubernetes deployment for elastic scaling across hybrid environments.

VII.. CONCLUSION

This paper presented an integrated Cyber Crime Detection System that unifies machine learning-based network traffic analysis, super-intelligent agentic AI forensic investigation, and SDN-based automated mitigation in a single closed-loop platform on Kali Linux. The ensemble ML detection engine combining Random Forest, XGBoost, and LSTM achieves 96.8% binary accuracy and 94.2% multi-class accuracy with 2.8% false positive rate, outperforming individual classifiers by 1.7–3.0 percentage points. The agentic AI investigation framework dynamically orchestrates Kali Linux security tools in threat-adaptive sequences, reducing mean investigation time by 87%—from 1.8 hours to 45 seconds—with 94.6% evidence collection completeness. Severity-graded Ryu SDN mitigation achieves 2.3-second mean mitigation latency with automated SHA-256 chain-of-custody evidence preservation.

Thirty-day empirical evaluation across five attack categories confirms 48-second end-to-end MTTR, 99.9% system uptime, and consistent performance under continuous monitoring conditions—a 99.7% improvement over industry-average MTTR. By closing the detection–investigation–mitigation loop autonomously, the system addresses the fundamental velocity mismatch between machine-speed attackers and human-speed defenders, enabling organizations lacking dedicated security staff to maintain effective cyber defense. This work establishes a foundation for future research in autonomous cyber defense, adaptive forensic reasoning, and self-defending network infrastructure.

REFERENCES

[1] Cybersecurity Ventures, "Cybercrime Report 2025," Cybersecurity Ventures, Northampton, MA, Rep. 2023. [Online]. Available: https://cybersecurityventures.com

[2] M. Conti et al., "A survey of man in the middle attacks," IEEE Commun. Surv. Tutorials, vol. 18, no. 3, pp. 2027–2051, 2016.

[3] A. Greenberg, Sandworm. New York, NY: Doubleday, 2019.

[4] B. J. Hutchins, M. Cloppert, and R. Amin, "Intelligence-driven computer network defense informed by analysis of adversary campaigns," in Proc. ICCRTS, 2011, pp. 1–14.

[5] A. Alshamrani et al., "A survey on advanced persistent threats: Techniques, solutions, challenges, and research opportunities," IEEE Commun. Surv. Tutorials, vol. 21, no. 2, pp. 1851–1877, 2019.

[6] IBM Security, "Cost of a Data Breach Report 2023," IBM Corp., Armonk, NY, Rep. 2023.

[7] Ponemon Institute, "2023 Cost of Cybercrime Study," Accenture, Dublin, Ireland, Rep. 2023.

[8] ISC2, "2023 Cybersecurity Workforce Study," ISC2, Clearwater, FL, 2023.

[9] ISACA, "State of Cybersecurity 2023," ISACA, Schaumburg, IL, Rep. 2023.

[10] L. Wang et al., "Agentic AI: A new paradigm for autonomous systems," IEEE Trans. Artif. Intell., vol. 6, no. 2, pp. 341–368, 2025.

[11] S. Yao et al., "ReAct: Synergizing reasoning and acting in language models," in Proc. ICLR, 2023.

[12] A. Khraisat et al., "Survey of intrusion detection systems: Techniques, datasets and challenges," Cybersecurity, vol. 2, no. 1, pp. 1–22, 2019.

[13] M. Ring et al., "A survey of network-based intrusion detection data sets," Comput. Secur., vol. 86, pp. 147–167, 2019.

[14] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in Proc. ICISSP, 2018, pp. 108–116.

[15] Canadian Institute for Cybersecurity, "CIC-IDS2017 Dataset," Univ. of New Brunswick, 2017. [Online]. Available: https://www.unb.ca/cic/datasets/ids-2017.html

[16] M. Tavallaee et al., "A detailed analysis of the KDD CUP 99 data set," in Proc. IEEE CISDA, 2009, pp. 1–6.

[17] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection systems," in Proc. MilCIS, 2015, pp. 1–6.

[18] L. Breiman, "Random forests," Mach. Learn., vol. 45, no. 1, pp. 5–32, 2001.

[19] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. ACM KDD, 2016, pp. 785–794.

[20] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735–1780, 1997.

[21] L. Wang et al., "A survey on large language model based autonomous agents," Front. Comput. Sci., vol. 18, no. 6, pp. 1–26, 2024.

[22] J. S. Park et al., "Generative agents: Interactive simulacra of human behavior," in Proc. ACM UIST, 2023.

[23] A. Arora, M. W. Godfrey, and D. E. Perry, "Agentic AI for cybersecurity: Opportunities and challenges," IEEE Secur. Privacy, vol. 22, no. 3, pp. 34–43, 2024.

[24] D. Kreutz et al., "Software-defined networking: A comprehensive survey," Proc. IEEE, vol. 103, no. 1, pp. 14–76, 2015.

[25] P. Nunes et al., "Combining OpenFlow and sFlow for an effective and scalable anomaly detection and mitigation mechanism on SDN environments," Comput. Netw., vol. 89, pp. 123–140, 2015.

[26] M. Roesch, "Snort—Lightweight intrusion detection for networks," in Proc. USENIX LISA, 1999, pp. 229–238.

[27] N. Milajerdi et al., "POIROT: Aligning attack behavior with kernel audit records for cyber threat hunting," in Proc. ACM CCS, 2019, pp. 1795–1812.

[28] T. A. Tang et al., "Deep recurrent neural network for intrusion detection in SDN-based networks," in Proc. IEEE NetSoft, 2018, pp. 202–206.

[29] Y. Mirsky et al., "Kitsune: An ensemble of autoencoders for online network intrusion detection," in Proc. NDSS, 2018.

[30] E. Cauda et al., "Graph neural networks for malicious network traffic detection," in Proc. IEEE MILCOM, 2022, pp. 1–6.
[31] N. Papernot et al., "The limitations of deep learning in adversarial settings," in Proc. IEEE EuroS&P, 2016, pp. 372–387.
[32] Gartner, "Magic Quadrant for Security Orchestration, Automation and Response," Gartner, Stamford, CT, Rep. 2023.
[33] D. Taubmann et al., "Soar evaluation: Measuring security automation effectiveness," in Proc. CSET, 2020.
[34] R. Chandramouli et al., "Security recommendations for hypervisor deployment," NIST SP 800-125B, 2019.
[35] P. Porras et al., "A security enforcement kernel for OpenFlow networks," in Proc. ACM HotSDN, 2012, pp. 121–126.
[36] S. Scott-Hayward et al., "A survey of security in software defined networks," IEEE Commun. Surv. Tutorials, vol. 18, no. 1, pp. 623–654, 2016.
[37] B. Carrier, File System Forensic Analysis. Upper Saddle River, NJ: Addison-Wesley, 2005.
[38] S. M. Hussain et al., "Security incidents in CI/CD pipelines: A large-scale study," in Proc. IEEE/ACM ICSE, 2024.
[39] Offensive Security, "Kali Linux Documentation," Offensive Security, 2024. [Online]. Available: https://www.kali.org/docs/
[40] C. Lugaresi et al., "MediaPipe: A framework for building perception pipelines," arXiv:1906.08172, 2019.
[41] A. Lazarevic et al., "A comparative study of anomaly detection schemes in network intrusion detection," in Proc. SDM, 2003, pp. 25–36.
[42] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Comput. Surv., vol. 41, no. 3, pp. 1–58, 2009.
[43] J. Leskovec, A. Rajaraman, and J. D. Ullman, Mining of Massive Datasets, 3rd ed. Cambridge, UK: Cambridge Univ. Press, 2020.
[44] Volatility Foundation, "The Volatility Framework," Open Source, 2024. [Online]. Available: https://volatilityfoundation.org