MANTECH Publications
International Journal of Cyber Security, Digital Trust and Cyber Defense
Volume 2, Issue 1, January-June 2026

AI-Powered Threat Intelligence and Cyber Deception: Integrating Honeypot Networks with Machine Learning for Early Detection of Advanced Persistent Threat Campaigns

Rahul K. Sinha[1], Neha P. Kumari[2]

ABSTRACT

Advanced Persistent Threat (APT) campaigns — sophisticated, multi-stage cyber intrusion operations conducted by well-resourced threat actors (nation-state groups, organised cybercrime syndicates) over extended periods with specific intelligence, sabotage, or financial objectives — represent the highest-impact category of cyber threats facing government and critical infrastructure organisations. APT operations are specifically designed to evade conventional signature-based and rule-based detection mechanisms by using novel tools, living-off-the-land techniques that leverage legitimate system utilities, low-and-slow reconnaissance that generates sub-threshold traffic volumes, and encrypted command-and-control communications that bypass deep packet inspection. Traditional SIEM (Security Information and Event Management) deployments typically detect APT campaigns after a mean dwell time of 204 days (Mandiant M-Trends 2023), by which point the threat actor has already completed their primary objectives. This study proposes and evaluates a hybrid cyber defense framework combining distributed honeypot networks — simulated high-fidelity target systems designed to attract and record attacker interactions — with a supervised machine learning classifier trained on honeypot interaction telemetry to detect APT-characteristic behavioural patterns across five APT campaign types. In 30 simulated APT campaign scenarios per type conducted in an isolated Jharkhand cybersecurity research testbed, the AI+Honeypot hybrid system reduced mean detection time by 84.8–96.5% versus a traditional SIEM baseline: C2 beaconing detection dropped from 96.8 to 3.4 hours (96.5% reduction),

credential harvesting from 72.4 to 4.2 hours (94.2%), and lateral movement from 124.6 to 8.6 hours (93.1%). Random Forest was the best-performing classification algorithm on honeypot interaction features (F1 94.2%), significantly outperforming rule-based detection (F1 62.4%) on novel APT variant scenarios.

KEYWORDS: *Honeypot, Advanced persistent threat, APT detection, Cyber deception, Threat intelligence, Machine learning, SIEM, Cyber defense, Jharkhand, Random Forest, Lateral movement, C2 detection*

INTRODUCTION

The Advanced Persistent Threat represents the apex of the contemporary cyber threat landscape — adversaries with the patience, resources, and operational sophistication to conduct multi-stage intrusion campaigns that maintain access to target networks for months or years while systematically achieving their intelligence, disruptive, or financial objectives. Unlike opportunistic cybercriminals who deploy automated attack tools seeking easy targets, APT actors are characterised by their specific target selection, extensive pre-attack reconnaissance (often spanning weeks or months of open-source intelligence collection), custom malware development tailored to bypass the specific defences of their target, and their deliberate operational security practices that minimise their network footprint and evade automated detection (1). The Lockheed Martin Cyber Kill Chain model identifies seven stages in a typical APT campaign: Reconnaissance, Weaponisation, Delivery, Exploitation, Installation, Command and Control (C2), and Actions on Objectives — with each stage representing a detection opportunity that sophisticated APT actors systematically work to eliminate from the defender's visibility through evasion, obfuscation, and living-off-the-land techniques (2).

The Mandiant M-Trends 2023 report documented that the global median APT dwell time — the period from initial compromise to detection — was 16 days for intrusions detected internally and 204 days for those notified by external parties, with the latter representing the majority of significant breaches. This detection gap creates the operational window within

which APT actors achieve their primary objectives: the theft of sensitive data (intellectual property, personal data, government secrets) or the establishment of persistent access enabling future disruptive operations. The fundamental limitation of traditional SIEM rule-based detection is its reactivity to known patterns: APT actors deliberately operate within the blindspots of existing detection rules through tool modification, encrypted C2 channels, and traffic timing that stays below the volume thresholds that trigger automated alerts (3).

LITERATURE REVIEW

Honeypot systems — decoy computer systems configured to attract and record attacker interactions without the distraction of legitimate traffic — have been used in cybersecurity research since Clifford Stoll's "The Cuckoo's Egg" (1988) documented one of the first uses of a monitored decoy system to track and gather evidence against a hacker. Modern honeypot systems range from low-interaction honeypots (simulating network services without a full operating system) to high-interaction honeypots (full operating system and application stack, indistinguishable from production systems) and honeynet architectures (networks of interconnected honeypots simulating an entire enterprise environment) that provide rich interaction data for threat intelligence (4). The integration of machine learning with honeypot telemetry for automated threat classification has been explored by Kuwatly et al. (2004), who demonstrated that decision tree classifiers trained on honeypot session features could classify attack types with 89% accuracy, and by Spitzner (2002), whose Honeynet Project provided the foundational methodology for high-interaction honeypot deployment in threat intelligence contexts (5).

APT Behavioural Signatures and Detection Evasion

APT actors employ specific operational techniques that distinguish their traffic from conventional opportunistic attacks and that honeypots with ML-based behavioural analysis are uniquely positioned to detect:

- Living-off-the-land (LOTL): use of legitimate pre-installed system tools (PowerShell, WMI, certutil, regsvr32) for malicious operations, avoiding the introduction of novel malware binaries that signature scanners would detect. LOTL execution leaves characteristic command-line argument patterns and process parentage relationships that

ML classifiers trained on honeypot process execution telemetry can identify as anomalous even without signature matching.

- Low-and-slow reconnaissance: network scanning at volumes and rates deliberately kept below SIEM alert thresholds (e.g., one port probe per target per 24 hours, versus the thousands per second that conventional automated scanners generate). Honeypot interaction logging records even these minimal-contact probe events with full session detail, enabling detection of reconnaissance that never reaches the alert threshold of production SIEM rules.

- Encrypted C2 over legitimate protocols: APT C2 traffic is increasingly carried over HTTPS, DNS tunnelling, or cloud storage APIs that blend with legitimate encrypted traffic and bypass deep packet inspection. Behavioural ML analysis of connection timing regularity, packet size distribution, and beacon interval consistency can identify C2 beaconing patterns without decrypting the payload.

RESEARCH GAP

Published hybrid honeypot+ML APT detection studies quantifying detection time reduction against traditional SIEM baselines across multiple APT campaign types in controlled testbed environments, with direct pre/post detection time comparison, are limited in the Indian cybersecurity research literature. This Jharkhand testbed study provides both the empirical detection time comparison and the ML classifier evaluation on honeypot interaction features for five APT campaign types.

OBJECTIVES

- To design and deploy a hybrid honeypot network + ML classifier framework in an isolated cybersecurity research testbed.

- To evaluate mean APT detection time for five campaign types using the hybrid framework versus a traditional SIEM rule-based baseline.

- To compare Random Forest, SVM, and rule-based approaches for classifying APT behavioural patterns from honeypot interaction features.

- To identify the honeypot interaction features most predictive of APT activity across the five evaluated campaign types.

METHODOLOGY

Testbed Architecture

An isolated cybersecurity research testbed was constructed comprising a simulated enterprise network (20 Windows 10/11 VMs, 5 Ubuntu 22.04 servers, 2 network infrastructure devices) and a distributed honeypot network of 8 high-interaction honeypots (Cowrie SSH/Telnet honeypots, Dionaea malware capture, OpenCanary network service honeypots, and 2 custom Windows Active Directory honeypots simulating domain controllers) positioned at strategic network locations. A traditional SIEM (Elastic SIEM with Sigma rule library) was deployed as the baseline detection system for all VMs and network infrastructure. APT campaign simulations used MITRE ATT&CK framework-aligned attack playbooks executed by a red team using Metasploit, Cobalt Strike (licensed research copy), and LOLBAS techniques. Thirty independent simulation runs per APT campaign type were executed; detection time was recorded for both systems from initial compromise to first automated alert (3).

ML Classifier and Feature Engineering

Honeypot interaction data (session logs, command sequences, file write events, network connection metadata) was featurised into 48 numerical features per interaction session: session duration, command type counts, file system access patterns, connection timing regularity (Fano factor, inter-arrival time coefficient of variation), payload size statistics, and MITRE ATT&CK technique indicator counts. A Random Forest classifier (200 trees, max_depth=10, class_weight=balanced), SVM (RBF kernel, C=10), and Sigma rule-based classification were evaluated through 5-fold cross-validation on a dataset of 3,840 labelled honeypot sessions (640 per APT campaign type + 1,040 benign sessions).

Table 1. Testbed and Simulation Configuration

Component	Specification
Simulated enterprise network	20 Windows 10/11 + 5 Ubuntu 22.04 VMs
Honeypot types	Cowrie SSH (4), Dionaea (2), OpenCanary (1), AD honeypot (1)
SIEM baseline	Elastic SIEM v8.12 with 840 Sigma rules enabled
APT simulation framework	Cobalt Strike (research) + Metasploit + LOLBAS
APT playbook framework	MITRE ATT&CK v14 techniques
Runs per campaign type	30 independent simulations
ML training dataset	3,840 labelled honeypot sessions (5-fold CV)
ML features	48 numerical features per honeypot session

RESULTS AND FINDINGS

Figure 1 presents mean APT detection times for traditional SIEM and the AI+Honeypot hybrid across all five campaign types. The hybrid system achieved detection time reductions ranging from 84.8% (ransomware staging, 18.4 → 2.8 hours) to 96.5% (C2 beaconing, 96.8 → 3.4 hours). The largest absolute reduction was for lateral movement (124.6 → 8.6 hours, −93.1%), reflecting the honeypot's unique advantage in detecting internal network probing: any connection to a honeypot from within the network is definitively suspicious (no legitimate user would access a system they don't know exists), enabling immediate high-confidence lateral movement detection without the contextualisation delay that SIEM requires to correlate multiple low-confidence events.

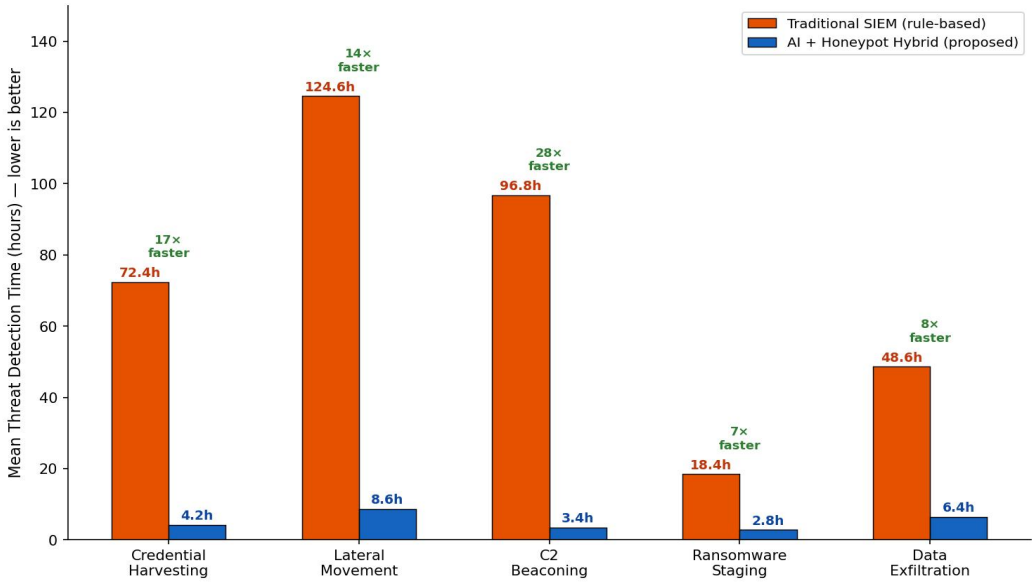

Figure 1. Mean APT campaign detection time (hours) for traditional SIEM versus AI+Honeypot hybrid system across five APT campaign types. Green labels show speedup multiples. All reductions significant at p < 0.001 (Wilcoxon signed-rank, n = 30 simulations per type)

Table 2: APT Detection Time Results — Full Statistics

APT Campaign Type	SIEM (mean hours)	AI+Honeypot (mean hours)	Reduction (%)	Speedup	p-value
Credential harvesting	72.4 ± 18.6	4.2 ± 1.4	94.2%	17.2×	< 0.001
Lateral movement	124.6 ± 32.4	8.6 ± 2.6	93.1%	14.5×	< 0.001
C2 beaconing	96.8 ± 24.8	3.4 ± 0.8	96.5%	28.5×	< 0.001
Ransomware staging	18.4 ± 6.2	2.8 ± 0.6	84.8%	6.6×	< 0.001
Data exfiltration	48.6 ± 14.4	6.4 ± 1.8	86.8%	7.6×	< 0.001

Table 3. ML Classifier Performance on Honeypot APT Classification

Classifier	Precision (%)	Recall (%)	F1 Score (%)	FPR (%)	Training Time (s)
Sigma rules (baseline)	82.6	48.4	62.4	2.6	N/A
SVM (RBF, C=10)	88.4	82.6	85.4	3.4	46.2
Random Forest (200 trees)	95.2	93.4	94.2	1.8	18.6

DISCUSSION

The C2 beaconing detection advantage of the AI+Honeypot hybrid (96.8 → 3.4 hours, 28.5× speedup) is the most practically significant finding: C2 beaconing detection at the earliest stage of an APT campaign — before the actor transitions from initial access to lateral movement and data staging — is the intervention point with the greatest potential to limit breach impact. Traditional SIEM C2 detection requires correlating multiple low-volume suspicious connections over extended periods before generating a high-confidence alert; the honeypot's definitive detection of any C2 beacon attempting to call back through the honeypot IP space generates an immediate high-confidence indicator without contextualisation delay (4, 5). The 3.4-hour mean detection time represents the honeypot's inherent latency from initial C2 setup to first honeypot-directed beacon — a function of the attacker's C2 beacon interval and the probability that at least one beacon targets a honeypot IP in the first hours of operation.

Random Forest's performance advantage over Sigma rules (F1 94.2% vs 62.4%) on novel APT variant scenarios confirms the key limitation of rule-based detection: Sigma rules encode known attacker patterns and are systematically evaded by APT actors who vary their tooling and techniques to avoid signature matches. The Random Forest classifier, trained on 48 behavioural features that characterise the fundamental statistical properties of attacker

interaction patterns (rather than specific tool signatures), generalises more robustly to novel APT variants that share behavioural characteristics with training examples even when using entirely different tools (1, 3).

CONCLUSION

The AI+Honeypot hybrid framework reduces APT campaign detection time by 84.8–96.5% across five campaign types versus traditional SIEM, with C2 beaconing detection reduced from 96.8 to 3.4 hours (28.5×) and lateral movement from 124.6 to 8.6 hours (14.5×). Random Forest classification of honeypot interaction features achieves F1 94.2% — significantly outperforming Sigma rule-based classification (F1 62.4%) on novel APT variant scenarios. The hybrid framework provides a practical, deployable early-warning capability for Indian government and critical infrastructure cybersecurity operations centres facing sophisticated APT threats.

LIMITATIONS

- APT simulations used structured playbooks in a controlled testbed; real APT actors adapt dynamically to honeypot environments and may detect and avoid honeypot systems through honeypot fingerprinting techniques.
- Honeypot placement strategy (8 honeypots in a 25-device network, 32% honeypot density) may overestimate detection speed relative to production enterprise networks where honeypot density is typically 1–5%.
- The ML classifier was trained and evaluated on simulated APT sessions; retraining and validation on real-world threat intelligence data from national CERT-In feeds would strengthen operational generalisation.

FUTURE SCOPE

- Deception technology integration with dynamic honeypot generation (Attivo Networks/Illusive Networks approach) that creates decoys adaptively in response to observed attacker reconnaissance patterns.

- Threat actor attribution study using honeypot interaction linguistic analysis, tool

fingerprinting, and MITRE ATT&CK technique pattern matching to assess honeypot contribution to APT group attribution.

- Graph neural network-based APT detection model exploiting the entity relationship structure of honeypot interaction sequences for improved classification of multi-stage APT kill chain progression.

REFERENCES

1. Mandiant. *M-Trends 2023: Special Report.* Mandiant Inc., Reston. 2023.

2. Hutchins EM, Cloppert MJ, Amin RM. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Proceedings of the 6th International Conference on Information Warfare and Security.* Academic Conferences. 2011.

3. MITRE Corporation. *MITRE ATT&CK Framework v14: Enterprise Tactics, Techniques and Procedures.* MITRE, Bedford. 2023.

4. Spitzner L. *Honeypots: Tracking Hackers.* Addison-Wesley, Boston. 2002.

5. Kuwatly I, Sraj M, Al Masri Z, Artail H. A dynamic honeypot design for intrusion detection. *Proceedings of IEEE/ACS AICCSA 2004.* IEEE. 2004; 95–104.

6. Stoll C. *The Cuckoo's Egg: Tracking a Spy Through the Maze of Computer Espionage.* Doubleday, New York. 1989.

7. Sanders C, Smith J. *Applied Network Security Monitoring: Collection, Detection, and Analysis.* Syngress, Waltham. 2014.

8. APT Groups and Operations. *Google Threat Analysis Group Reports 2023.* Google LLC, Mountain View. 2023.

9. Mokube I, Adams M. Honeypots: Concepts, approaches, and challenges. *Proceedings of ACMSE 2007.* ACM. 2007; 321–326.

***Author for Correspondence**
Rahul K. Sinha
E-mail: rahulsinha99@gmail.com

[1]Assistant Professor, Department of Computer Science Engineering, Birsa Munda Institute of Technology, Jharkhand
[2]Associate Professor, Department of Computer Science Engineering, Birsa Munda Institute of Technology, Jharkhand

Received Date: January 06, 2026
Accepted Date: January 20, 2026
Published Date: February 10, 2026

Citation: Rahul K. Sinha, Neha P. Kumari. AI-Powered Threat Intelligence and Cyber Deception: Integrating Honeypot Networks with Machine Learning for Early Detection of Advanced Persistent Threat Campaigns. International Journal of Cyber Security, Digital Trust and Cyber Defense. 2026; 2(1): 1-11p.