Available online www.ejaet.com

European Journal of Advances in Engineering and Technology, 2024, 11(2):69-76

Research Article ISSN: 2394 - 658X

Proactive Threat Hunting: Leveraging AI for Early Detection of Advanced Persistent Threats

Kumrashan Indranil Iyer

Email: indranil.iyer@gmail.com

ABSTRACT

Advanced Persistent Threats (APTs) represent one of the most sophisticated and long-term cybersecurity threats faced by enterprises, government agencies, and critical infrastructure operators. APTs frequently employ stealthy tactics to maintain persistent access, evade detection, and methodically exfiltrate valuable data. Traditional reactive security measures (e.g., signature-based antivirus tools and basic intrusion detection) often fail to detect APTs in their early stages.

This paper presents an overview of proactive threat hunting strategies that harness artificial intelligence (AI) to identify and neutralize APTs before they escalate. We discuss the role of machine learning (ML) in detecting anomalous behaviors within vast telemetry datasets, leveraging techniques such as unsupervised clustering, graph-based anomaly detection, and behavioral analytics. For instance, AI-powered models can analyze deviations in user authentication patterns, lateral movement within a network, and subtle exfiltration techniques (key indicators of an ongoing APT attack).

We further highlight key challenges in AI-driven threat hunting, such as adversarial evasion tactics, data quality constraints, and false positive reduction. Finally, we propose a conceptual model that integrates AI analytics with human expertise for timely, accurate APT detection. We conclude with recommendations for practical AI-based threat-hunting deployments and future research directions in proactive cyber defense.

Keywords: Advanced Persistent Threats (APTs), Artificial Intelligence (AI), Machine Learning (ML), Threat Hunting, Intrusion Detection Systems, Cybersecurity, Proactive Defense, Data Exfiltration, Security Operations, AI Analytics.

INTRODUCTION

As the digital ecosystem grows more complex and interconnected, it enables unprecedented levels of global collaboration but also exposes organizations to increasingly sophisticated cyber threats. Among these, Advanced Persistent Threats (APTs) stand out due to their stealth, persistence, and highly targeted nature. APT operations often unfold over weeks, months, or even years, allowing attackers to gather intelligence, escalate privileges, and exfiltrate sensitive data with minimal risk of detection. For example, the 2010 Stuxnet attack, which targeted industrial control systems, is a classic case of an APT that went undetected for a significant period [1].

Traditional security measures, such as signature-based intrusion detection systems (IDS) and antivirus software, typically rely on identifying known attack patterns. While effective against commodity malware, these methods fall short when faced with novel, polymorphic malware or sophisticated Tactics, Techniques, and Procedures (TTPs) employed by APT groups. For instance, signature-based IDS failed to detect the 2014 Sony Pictures attack due to the use of customized malware and advanced evasion techniques.

To counter this evolving threat landscape, proactive threat hunting has emerged as a more aggressive approach, where defenders actively search for hidden threats or anomalies before they trigger alerts. AI-driven threat hunting, which leverages big data analytics, machine learning, and sometimes User and Entity Behavior Analytics (UEBA), plays a crucial role in uncovering subtle signs of compromise and infiltration. For example, machine learning models can identify unusual network traffic patterns that may indicate an APT, such as data exfiltration attempts or lateral movement within the network.

Research Objectives

1. Examine the key elements of proactive threat hunting, emphasizing AI's capacity to process large volumes of heterogeneous security data.

2. Evaluate AI and ML techniques that can detect APT behaviors early in the kill chain.

3. Propose a conceptual framework for integrating AI-driven threat hunting with human expertise, balancing automation and analyst-driven inquiry.

4. Identify challenges in scaling and adopting AI-enabled threat hunting tools across diverse organizational contexts.

LITERATURE REVIEW

Advanced Persistent Threats (APTs)

Advanced Persistent Threats (APTs) typically follow a structured attack lifecycle consisting of reconnaissance, initial compromise, persistence, lateral movement, and data exfiltration. These threats are highly targeted, stealthy, and sustained over long periods, often lasting months or even years. APT actors, frequently state-sponsored or part of well-organized criminal groups, are known for their resourcefulness, employing sophisticated techniques such as zero-day exploits, custom malware, and advanced evasion tactics to bypass traditional detection systems [2].

In the reconnaissance phase, attackers gather intelligence on the target organization's network and personnel. The initial compromise usually occurs through spear-phishing or exploiting a vulnerability in the network. Once access is gained, attackers establish persistence by embedding themselves within the system, often utilizing backdoors, rootkits, or custom malware. Lateral movement across the network follows, with the goal of escalating privileges and gaining further access to critical systems. Finally, data exfiltration occurs, where sensitive or valuable information is extracted stealthily, often over extended periods to avoid detection [3].

Reactive vs. Proactive Defense

Traditional security operations are heavily reliant on reactive defense mechanisms, where security teams respond to alerts generated by systems like Endpoint Detection and Response (EDR) and Intrusion Detection Systems (IDS). These systems typically detect suspicious behavior after it has occurred, often based on known attack signatures. While incident response is crucial, it can be insufficient when dealing with Advanced Persistent Threats (APTs), which are designed to evade detection and may remain undetected for extended periods, sometimes even months or years.

In contrast, proactive defense strategies aim to prevent APTs from reaching critical stages of their attack lifecycle. These strategies involve actively hunting for anomalies that deviate from normal behavior, using advanced techniques such as behavior analytics, threat intelligence, and machine learning to identify suspicious patterns before they manifest as security breaches. Proactive approaches, therefore, shift the focus from responding to threats to detecting and neutralizing them before they can cause harm, enabling a more robust defense posture.

AI in Threat Hunting

Artificial Intelligence (AI) has demonstrated significant potential in enhancing the effectiveness of threat hunting, particularly when dealing with vast data volumes from network logs, endpoint telemetry, and user activities. Machine learning (ML) and Deep learning (DL) models are good at uncovering subtle outliers or relationships that may evade traditional, rule-based systems. Below are some key AI techniques used in threat hunting:

- **Clustering Algorithms:** These algorithms, such as k-means or DBSCAN, group similar behaviors together, enabling the identification of outliers that may signify malicious activity. For example, clustering techniques can detect irregular patterns in user login times or unusual access to critical assets, flagging them for further investigation.
- **Anomaly Detection Models:** These models, which use statistical learning methods like Gaussian Mixture Models (GMM) or Isolation Forests, establish a baseline of "normal" behavior for network traffic or user actions. Any significant deviation from this baseline can trigger an alert for potential intrusions. For instance, anomaly detection can identify abnormal data flows across a network, potentially signaling a data exfiltration attempt by an APT group [4].
- **Neural Networks (e.g., Autoencoders):** Autoencoders (a type of neural network) can detect anomalies in high-dimensional data by learning an efficient encoding of normal behavior. When they encounter unfamiliar patterns, they can raise alarms. For example, autoencoders can be used to spot new forms of malware that do not match existing signatures, providing early warning signs before the attack escalates.

These AI-powered techniques enhance the accuracy and speed of threat detection, enabling security teams to respond proactively to potential APTs.

CORE COMPONENTS OF PROACTIVE THREAT HUNTING

Data Collection and Normalization

Threat hunters rely on comprehensive, high-quality data from various sources, including:

● **Endpoint Telemetry:** Process creation logs, registry changes, file integrity monitoring, kernel events, and command-line activity. For instance, detecting a PowerShell script execution with unusual obfuscation techniques may indicate malicious activity [5].
● **Network Traffic:** Packet captures, flow data (e.g., NetFlow, sFlow), deep packet inspection (DPI), domain name system (DNS) queries, and encrypted traffic analysis. For example, an unusually high volume of outbound connections to a previously unknown domain could suggest a command-and-control (C2) channel [4].
● **Identity and Access Logs:** Authentication attempts, privilege escalation events, changes in role-based access control (RBAC), and failed login attempts across multiple geolocations. Anomalous authentication patterns (e.g., a login from an unusual IP address followed by a privilege escalation) can indicate credential theft.
● **Threat Intelligence Feeds:** Indicators of compromise (IoCs), known malicious IP addresses, domain reputation data, tactics, techniques, and procedures (TTPs) from frameworks like MITRE ATT&CK. Integrating real-time threat intelligence helps correlate ongoing activity with known attack patterns.
● **Application and Cloud Logs:** API call logs, container activity, cloud workload behaviors (e.g., AWS CloudTrail, Microsoft Defender for Cloud). Monitoring unusual API requests, such as excessive file downloads from a cloud storage bucket, can flag potential data exfiltration [6].
● **Deception Technology and Honeytokens:** Deploying decoy accounts, credentials, or services can help lure attackers and detect lateral movement. If an attacker interacts with a decoy, an alert can be triggered early in the intrusion lifecycle [7].
Ensuring data consistency, integrity, and timeliness is crucial, as stale or incomplete data may obscure critical signs of intrusion. Data normalization (harmonizing log formats, timestamp conventions, and enrichment with contextual metadata) facilitates seamless analytics across multiple data sources, enabling more effective anomaly detection.

Machine Learning–Enhanced Analysis

As cyber threats become more sophisticated and attack surfaces expand, conventional correlation rules often fail to detect subtle attack patterns. Machine learning (ML) enables security teams to analyze large-scale data, uncover hidden anomalies, and correlate events beyond human capability. Key ML-enhanced techniques include:

1. Unsupervised Learning for Anomaly Detection

Unsupervised techniques are particularly useful for identifying deviations from normal behavior without requiring labeled attack data.
● **Isolation Forest:** This model isolates anomalies by randomly selecting features and partitioning data. Anomalous instances require fewer splits to be isolated. For example, an employee accessing a sensitive database for the first time at an unusual hour could be flagged for review [8].
● **Autoencoders:** These neural networks learn compact representations of normal behavior. When encountering anomalous data, they produce higher reconstruction errors. For instance, an autoencoder trained on normal network traffic can detect deviations indicative of command-and-control (C2) activity [9].

2. Behavioral Profiling

By establishing dynamic baselines, behavioral profiling techniques identify abnormal user or system activity.
● **User & Entity Behavior Analytics (UEBA):** ML models create personalized profiles of users and entities, tracking deviations. For example, if a user suddenly downloads large volumes of data from a confidential repository, UEBA can flag it as potential insider threat activity.
● **Graph-Based Analytics:** Security graphs model interactions between users, devices, and services to identify abnormal communication patterns. A sudden surge in connections from an endpoint to multiple external servers, bypassing usual routes, might indicate lateral movement [10].

3. Deep Learning for Advanced Correlation

Deep learning models enhance detection by recognizing complex patterns across time-series data.
● **Long Short-Term Memory (LSTM):** This recurrent neural network (RNN) variant is well-suited for analyzing time-dependent security events. LSTM can detect unusual login sequences that resemble brute-force attacks or credential stuffing [11].
● **Transformers:** Advanced transformer models, like those used in natural language processing (NLP), can analyze security logs as sequential data. By learning event correlations over long timeframes, transformers can help detect multi-stage APT campaigns that evade rule-based systems [12].

Continuous Hunting Cycles

Proactive threat hunting is an iterative process that operates in structured cycles, ensuring continuous refinement of detection capabilities. Each cycle consists of four key phases:

1. Hypothesis Formulation

● Analysts develop threat hypotheses based on intelligence sources, industry attack trends, and organization-specific risks.
● Example: If recent threat intelligence reports suggest an APT group is exploiting PowerShell for fileless malware attacks, analysts may hypothesize that such activity exists within their environment [13].

2. Data Analysis & ML Insights

● AI-driven tools analyze vast datasets to identify anomalies, suspicious patterns, or deviations from historical baselines.

● Example: An unsupervised ML model detects an increase in anomalous remote login attempts outside business hours, indicating a possible compromised account.

3. Investigative Drilling

● Analysts perform pivoting, i.e., investigating flagged entities across multiple datasets including network traffic, endpoint logs, and access patterns.

● Example: If an anomaly is detected in Windows Event Logs, analysts trace the associated user session to identify command execution history and lateral movement attempts [14].

4. Verification & Response

● onfirmed threats trigger containment actions, such as isolating affected endpoints, revoking compromised credentials, or blocking malicious IPs.

● False positives or benign anomalies are used to fine-tune ML models, reducing future noise.

● Example: If an alert turns out to be a legitimate admin task, it is marked as a known behavior, reducing false positives for similar future actions [5].

This cyclic approach ensures continuous learning and adaptation, strengthening defenses against evolving threats.

Human-AI Collaboration

While AI-driven tools provide automation, scalability, and rapid anomaly detection, human expertise remains critical for interpreting ambiguous signals, understanding adversary intent, and adapting to evolving threats. The synergy between AI and human analysts enhances multiple aspects of proactive threat hunting:

1. Validating Alerts & Contextual Awareness

● AI detects statistical anomalies but lacks real-world context, leading to false positives. Human analysts validate alerts to differentiate between actual threats and benign anomalies.

● Example: An AI model flags an executive's late-night login from a new location as suspicious. However, an analyst checks corporate travel records and confirms the login is legitimate [5].

2. Root Cause & Attribution Analysis

● Security teams analyze flagged behaviors to trace attack origins, identify exploited vulnerabilities, and link activities to specific adversaries.

● Example: An anomaly detection model identifies unusual outbound data transfers. Analysts correlate it with threat intelligence feeds and determine it aligns with a known APT exfiltration technique [13].

3. Fine-Tuning AI Models for Continuous Adaptation

● Threat landscapes evolve, requiring constant retraining of AI models. Human analysts refine models by adjusting detection thresholds and updating behavioral baselines.

● Example: AI frequently misclassifies security researchers' penetration testing activity as a real attack. Analysts fine-tune the model to recognize red team exercises without triggering unnecessary alerts.

4. Hypothesis-Driven Hunting & Intuition-Based Analysis

● Unlike rule-based alerts, hypothesis-driven threat hunting involves making educated assumptions about potential attack vectors and actively searching for evidence.

● Example: Analysts hypothesize that an APT targeting financial institutions may exploit Active Directory misconfigurations. They proactively query AI-driven behavioral analytics to detect irregular privilege escalation events.

5. Investigating Subtle Tactics & Living-off-the-Land (LotL) Attacks

● Many adversaries avoid malware and use built-in tools (e.g., PowerShell, WMI, RDP) to evade detection. AI may miss these tactics, but human hunters recognize low-and-slow activity indicative of stealthy attacks.

● Example: AI overlooks an attacker using normal IT admin scripts for lateral movement. A human analyst notices that the admin account being used does not match its usual behavior profile [4].

6. Adversary Emulation & Purple Teaming

● Human-led red teams simulate real-world APT attacks, helping AI systems learn how adversaries operate and improving detection accuracy.

● Example: Analysts simulate ransomware propagation in a controlled environment. AI refines its detection model based on observed TTPs from the test.

7. AI Bias Mitigation & Explainability

● AI models can inherit biases from training data, leading to skewed detections. Analysts interpret model outputs, explain decisions, and mitigate unintended biases.

● Example: An AI system disproportionately flags third-party contractor logins as risky. Analysts adjust the model to account for remote access norms.

By integrating AI-driven automation with expert human analysis, organizations achieve a resilient, adaptive, and intelligence-driven approach to threat hunting.

PROPOSED FRAMEWORK FOR AI-DRIVEN THREAT HUNTING

To counter evolving cyber threats, an AI-driven threat hunting framework should integrate real-time data processing, machine learning analytics, human expertise, and automated response mechanisms. Below is a conceptual model that balances automation and human oversight to enhance threat detection, investigation, and response.

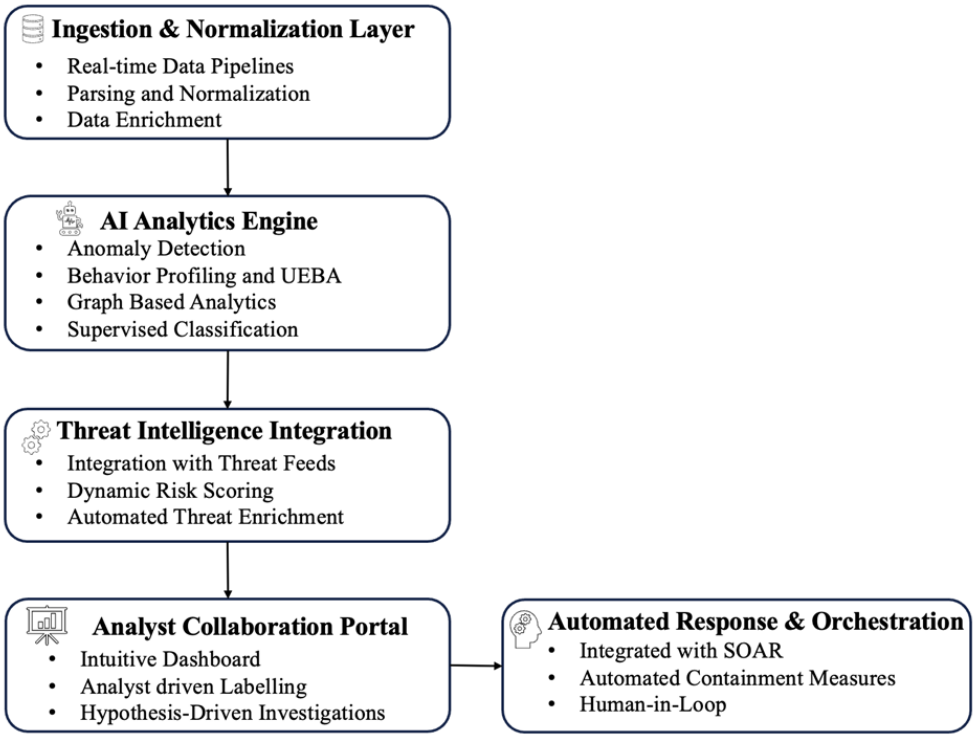

Figure 1: Framework of AI-Driven Threat Hunting
Source: Owner's Own Processing

1. Ingestion & Normalization Layer

Purpose: Ensures that raw security data from multiple sources is collected, standardized, and made available for analysis.

Key Components:

● Real-time Data Pipelines: Utilizes technologies like Apache Kafka, Fluentd, or Logstash to aggregate telemetry from various endpoints, servers, network devices, and cloud services.

● Parsing & Normalization: Converts logs into a unified format, ensuring timestamps, IP addresses, and event codes are standardized across different security tools.

● Data Enrichment: Adds contextual information such as geolocation data, threat intelligence indicators, and historical patterns.

2. AI Analytics Engine

Purpose: Leverages machine learning and advanced analytics to detect patterns, anomalies, and potential threats.

Key Techniques & Examples:

● Anomaly Detection:

O Isolation Forests: Identify deviations from normal network behavior, such as an internal machine suddenly connecting to a foreign IP address.

O Autoencoders: Detect new forms of malware by learning normal system behavior and flagging deviations.

● Behavioral Profiling & UEBA (User and Entity Behavior Analytics):

O Example: If an employee logs in from an unusual location at an odd hour and downloads sensitive data, behavioral models can flag this as a possible insider threat.

● Graph-Based Analytics for Lateral Movement Detection:

O Example: An attacker compromises a user's credentials and moves laterally across multiple devices. AI models can map relationships between users and devices to detect unauthorized privilege escalation.

● Supervised Classification:

O Uses pre-labeled data (known Indicators of Compromise - IoCs) to classify threats.

O Example: A deep learning model trained on malware signatures and attack behaviors can quickly classify new network traffic as benign or malicious.

3. Threat Intelligence Integration

Purpose: Enhances AI-generated detections by correlating them with external intelligence on known threats.

Key Functions:

● Integration with Threat Feeds: Sources include:

O Commercial Feeds: CrowdStrike, FireEye, IBM X-Force

o Open-Source Intelligence (OSINT): AlienVault OTX, MITRE ATT&CK, VirusTotal

● Dynamic Risk Scoring: Assigns risk levels to threats based on IoC reputation and context.

● Automated Threat Enrichment: Matches alerts with known attack techniques and previously seen adversary behaviors.

4. Analyst Collaboration Portal

Purpose: Serves as the interface between AI-driven analytics and human investigators, allowing analysts to validate and refine threat detections.

Key Features:

● Intuitive Dashboards: Visualizes high-risk anomalies and attack patterns.

● Analyst-Driven Labeling: Enables security teams to tag false positives, refine AI models, and improve hunting queries.

● Hypothesis-Driven Investigations: Analysts can create and test hunting hypotheses (e.g., "Are there any unusual admin credential logins outside working hours?").

5. Automated Response & Orchestration

Purpose: Reduces mean time to respond (MTTR) by enabling automatic remediation actions while allowing analysts to retain control over critical decisions.

Key Functions:

● Integration with SOAR (Security Orchestration, Automation, and Response):

O Automates predefined incident response playbooks.

O Example: If an endpoint is detected communicating with a known malicious IP, the SOAR system can automatically isolate it from the network.

● Automated Containment Measures:

O Blocking suspicious IPs, disabling compromised user accounts, or quarantining infected files.

● Human-In-The-Loop Review for Critical Incidents:

O Ensures automated actions for high-impact threats (e.g., mass account lockouts) require human analyst approval.

CHALLENGES AND CONSIDERATIONS

1. Data Quality & Volume

● **Challenge:** AI engines rely on complete and consistent data for effective analysis. Missing log sources or disruptions in data pipelines can degrade the accuracy and reliability of machine learning (ML) insights.

● **Consideration:** As environments scale, the volume of logs increases significantly, posing a challenge for data processing systems. To manage this, organizations must adopt distributed architectures and efficient ML algorithms capable of handling large-scale data flows.

2. False Positives & Analyst Workload

● **Challenge:** AI-driven anomaly detection systems often generate a high number of false positives, leading to alert fatigue among security analysts. This can overwhelm teams, increasing response time and diminishing the effectiveness of threat detection efforts.

● **Consideration:** Striking a balance between sensitivity (capturing all potential threats) and specificity (minimizing false positives) is critical. Ongoing fine-tuning and validation with real-world data are necessary to optimize performance and ensure that only the most relevant anomalies are flagged.

3. Model Drift & Evolving Threats

● **Challenge:** Advanced persistent threats (APTs) continually adapt their tactics, techniques, and procedures (TTPs) to bypass known detection models. Over time, even legitimate environmental changes (e.g., remote work adoption, system updates) can alter what is considered "normal" behavior, causing models to drift.

● **Consideration:** To combat this, continuous retraining of ML models is necessary, alongside regular iterative threat hunting cycles. This approach helps the organization stay ahead of evolving threats and ensures that detection systems remain effective.

4. Privacy & Compliance

● **Challenge:** The collection of detailed user activity logs raises significant privacy concerns, particularly under stringent data protection regulations such as GDPR and CCPA. Organizations must balance security needs with privacy rights.

● **Consideration:** Implementing privacy-preserving analytics or employing data minimization strategies is essential to comply with regulatory requirements while still maintaining strong visibility into potential threats. Techniques such as data anonymization and user consent management can help organizations navigate this delicate balance.

5. Skill Gaps

● **Challenge:** The integration of AI in cybersecurity requires a specialized workforce capable of interpreting AI-driven insights, refining ML models, and understanding sophisticated APT tactics. There is often a shortage of skilled professionals with the necessary expertise.

● **Consideration:** Organizations must invest in training programs and encourage cross-functional collaboration between data scientists and security analysts. This ensures that AI models are aligned with the practical needs of security operations and that teams are equipped to interpret and act on findings effectively.

FUTURE DIRECTIONS

As AI-driven cybersecurity evolves and advanced persistent threats (APTs) grow more sophisticated, proactive threat hunting must continue to adapt. Future research areas include:

1. Explainable AI (XAI) for Security

O Enhancing interpretability of machine learning decisions to improve analyst trust in AI-driven alerts.

O Developing visualization tools and explainable models to clarify anomaly root causes.

2. Federated Threat Hunting

O Enabling organizations to collaborate by sharing anonymized AI models or aggregated anomaly patterns.

O Leveraging collective intelligence while preserving data privacy.

3. AI-Enhanced Deception Technologies

O Using AI to create realistic honeypots and decoys that mimic legitimate environments.

O Tricking APT adversaries into revealing their tactics, techniques, and procedures (TTPs).

4. Integration with Zero Trust Architectures

O Merging continuous trust verification with AI-driven anomaly detection.

O Strengthening multi-layered security against stealthy attackers.

5. Adaptive Defense Against Adversarial AI

O Countering adversarial techniques such as data poisoning and adversarial inputs that attempt to mislead AI detectors.

O Researching robust machine learning methodologies that resist manipulation while maintaining detection accuracy.

By advancing these areas, UEBA and AI-driven security analytics can stay ahead of emerging cyber threats, ensuring resilient and adaptive defense mechanisms.

CONCLUSION

Advanced Persistent Threats constitute a formidable challenge to modern organizations, necessitating proactive methods that go beyond traditional reactive measures. **AI-driven threat hunting** offers a promising solution by continuously scanning vast data sets, identifying hidden anomalies, and accelerating the detection of early-stage intrusions. However, maximizing the effectiveness of these systems requires a careful balance between AI-driven automation and human analytical expertise.

Challenges around data quality, false positives, evolving attacker tactics, and regulatory constraints underscore the need for robust frameworks and well-trained security staff. Moving forward, research on explainable AI, adversarial resilience, and cross-organizational collaboration will be critical in refining proactive threat-hunting capabilities. By embracing these advances, security teams can better anticipate and counter emerging threats, reducing the time adversaries spend undetected and strengthening overall cyber resilience.

REFERENCES

[1]. U.S. National Security Archive, “Shadows of Stuxnet: Recommendations for U.S. policy on critical infrastructures,” National Security Archive, 2010. https://nsarchive.gwu.edu/sites/default/files/documents/3173339/Document-07.pdf.

[2]. P. Chen, L. Desmet, and C. Huygens, “A study on advanced persistent threats,” *Informatics*, vol. 1, no. 3, pp. 1–22, 2014. : https://inria.hal.science/hal-01404186/document.

[3]. Mandiant, “APT1: Exposing one of China’s cyber espionage units,” *Mandiant Report*, 2013. https://www.mandiant.com/sites/default/files/2021-09/mandiant-apt1-report.pdf

[4]. Mandiant, *M-Trends: APT Threat Hunting and Detection*, Mandiant Research Report, 2020.

[5]. R. Sommer and V. Paxson, “Outside the closed world: On using machine learning for network intrusion detection,” in *IEEE Symposium on Security & Privacy*, 2010.

[6]. AWS Security Hub, *Best Practices for Threat Detection in Cloud Environments*, Amazon Web Services, 2022.

[7]. J. Yuill, M. Zappe, and D. Denning, “Honeyfiles: Deception in the file system,” *Journal of Computer Security*, vol. 14, no. 6, pp. 549–573, 2006.
[8]. F. T. Liu, K. M. Ting, and Z. H. Zhou, “Isolation forest,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 23, no. 6, pp. 713–725, 2008.
[9]. M. Sakurada and T. Yairi, “Anomaly detection using autoencoders with nonlinear dimensionality reduction,” in *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2014.
[10]. L. Akoglu, H. Tong, and D. Koutra, “Graph-based anomaly detection and description: A survey,” *Data Mining and Knowledge Discovery*, vol. 29, no. 3, pp. 626–688, 2015.
[11]. S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
[12]. A. Vaswani *et al.*, “Attention is all you need,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
[13]. MITRE, *MITRE ATT&CK Framework: Enterprise Tactics, Techniques, and Procedures*, 2023. : https://attack.mitre.org
[14]. E. M. Hutchins, M. J. Cloppert, and R. M. Amin, “Intelligence-driven computer network defense informed by the kill chain methodology,” *Lockheed Martin White Paper*, 2011.