Check for updates

RESEARCH ARTICLE

REVISED Combating Advanced Persistent Threats through Artificial Intelligence: An Algorithmic Approach

[version 2; peer review: 4 approved with reservations, 1 not approved]

Pedro Brandao

Computer Sciences, Instituto Superior de Tecnologias Avancadas de Lisboa, Campus Académico do Lumiar, Alameda das Linhas de Torres nº179, Lisbon, 1750-142, Portugal

v2 **First published:** 20 May 2025, **5**:139
https://doi.org/10.12688/openreseurope.20268.1
Latest published: 16 Sep 2025, **5**:139
https://doi.org/10.12688/openreseurope.20268.2

Abstract

This essay highlights the importance of leveraging Artificial Intelligence (AI) tools to combat Advanced Persistent Threats (APTs). Central to this study is the theoretical framework describing the design and evaluation of algorithms that analyze extensive datasets of security-related consultation items. Algorithms for aggregating such datasets are also developed, along with a detailed methodology to ensure adequate statistical significance. To test the research hypothesis, the designed algorithms were used to analyze a large dataset containing detailed information on APT campaigns. APT campaigns involve a series of distinct attacks aimed at gaining control over computational assets. These campaigns are significant cybersecurity threats because they are difficult to detect and are part of a lucrative black market worth hundreds of billions of dollars. Refined tools for launching and managing APT campaigns are supplied to criminal organizations and state actors through Growth Hacking, Inbound Marketing, and other B2B services. Networks of controlled computational assets infiltrated by sell-side ad networks are extensively exploited by these tools. Improved algorithms have been developed to coordinate content consumption and carry out hacking and phishing attacks. Data are stored in a unique, bottom-up event model that describes the lifecycle of APT campaigns and accounts for the limited knowledge of their participants. In the spirit of open science, scripts for collection, deduplication, labeling, and feature engineering, along with dataset splits, data dictionaries, and other materials, have been shared alongside the results of this study. The developed algorithms were used to train and evaluate predictive models for the lifecycle stages of APT campaigns and the decision to engage with them. By providing new domain knowledge and insights,

Open Peer Review

Approval Status ? ? ? ✗ ?

	1	2	3	4	5
version 2 (revision) 16 Sep 2025			? view	✗ view	? view
version 1 20 May 2025	? view	? view			

1. **Zahid Akhtar**, State University of New York (SUNY) Polytechnic Institute, Utica, USA
2. **Hari Gonaygunta**, University of the Cumberlands, Williamsburg, USA
3. **Thanh Cong Truong**, University of Finance-Marketing, Ho Chi Minh City, Vietnam
4. **Giorgio Giacinto**, University of Cagliari, Cagliari, Italy
5. **Mani Prabha**, International American University, Los Angeles,, USA

Any reports and responses or comments on the article can be found at the end of the article.

this study encourages researchers to co-create domain-specific machine-learning algorithms with practitioners focusing on countermeasures and emphasizes the need for a new multidisciplinary field merging “Growth Hacking, Inbound Marketing, and Early-Stage Venture Cyber Security VC Strategies” 1.

Keywords
Cybersecurity, APT, Advanced Persistent Threat, Artificial Intelligence, Machine Learning, Deep Learning, Computer Security, Virus, Malware.



This article is included in the Horizon 2020 gateway.



This article is included in the Electrical, Electronic and Information Engineering gateway.

Corresponding author: Pedro Brandao (pedro.brandao@istec.pt)

Author roles: Brandao P: Conceptualization, Formal Analysis, Investigation, Methodology, Project Administration, Supervision, Writing – Original Draft Preparation, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No [774431, 952215](Foundations of Trustworthy AI - Integrating Reasoning, Learning and Optimization-TAILOR.

Copyright: © 2025 Brandao P. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Brandao P. **Combating Advanced Persistent Threats through Artificial Intelligence: An Algorithmic Approach [version 2; peer review: 4 approved with reservations, 1 not approved]** Open Research Europe 2025, **5**:139 https://doi.org/10.12688/openreseurope.20268.2

First published: 20 May 2025, **5**:139 https://doi.org/10.12688/openreseurope.20268.1

REVISED **Amendments from Version 1**

This revised version clarifies the narrative, strengthens the methodological core, and situates the contribution within the most recent scholarship. We have standardised UK English usage and harmonised figure and table captions; all modifications prompted by Reviewer 1 remain in red, and additions addressing Reviewer 2 are in green for transparency.

Conceptually, the article now frames APT detection as a streaming, graph-contextual problem. A dedicated "Threat Model & Terminology" section fixes definitions and assumptions. Methodologically, we provide full pipeline pseudocode and sub-module descriptions (stream_reason, cluster, explain), explicit hyperparameter settings with rationale, and a latency-constrained training and optimization protocol (chronological splits, Bayesian tuning, early stopping). We add a formal problem statement, complexity and memory analysis, and a half-life calibration for exponential decays to relate parameters to operational behaviour. New figures depict the end-to-end architecture and the train–validate–test regime under throughput and latency targets.

Empirically, we specify a reproducible evaluation protocol and report deployment-relevant metrics: PR-AUC, ROC-AUC, TPR@FPR = 1%, Time-to-Detection, and Alerts per Day per 1,000 hosts, with uncertainty estimated via stratified bootstrap. A comparative table and critical discussion position the approach alongside sequence/intent models, explainable IDS, CTI-driven knowledge graphs, and reinforcement-learning defences.

The literature review has been expanded to incorporate more than twenty-five peer-reviewed studies from 2020–2025, and references have been consolidated into a single IEEE-style list with renumbered in-text citations. Finally, we outline ablation and sensitivity studies (window size, decay, clustering parameters, alert thresholds) and robustness checks (label noise, timestamp jitter) to guide future empirical work.

We hope these revisions improve clarity, reproducibility, and practical relevance.

Any further responses from the reviewers can be found at the end of the article

Introduction

In recent years, cyber adversaries have become more sophisticated and have created a new category of threats called Advanced Persistent Threats (APTs). They use a prolonged and targeted approach to breach defense networks in multiple steps to exfiltrate sensitive data or mishaps. The defense capabilities of modern organizational infrastructure are now at significant risk from APTs because of their complexity, diversity, and unpredictable attack timings.

APTs harm victims by preventing them from responding and fighting back against an attacker. Until now, most related studies have focused on the impacts of APTs, detection models, and defenses related to network and host architectures. These APT models illustrate the multiple life cycles of an APT intrusion, along with the tactics, techniques, and procedures (TTP) used by attackers at each stage. There is a need to expand understanding of APTs from the perspective of attacker strategies, approaches, and processes, as well as to broaden perspectives on APT detection methods. To explore these areas, some research has aimed to analyze APTs based on attacker methodologies, generate insights into attacker evasion techniques, visualize and represent APT detection across different organizational units—such as vendors and companies—and propose core technologies for creating backtracking graphs with specific templates. These backtracking graphs can offer a comprehensive view of the APT intrusion process. Additionally, lightly structured models and algorithms have been developed to identify semantic patterns related to APTs using backtracking graphs.

Threat Model & Terminology

We formalise the adversary as a strategic, goal-driven agent aligned to MITRE ATT&CK tactics. Capabilities include phishing-led initial access, credential theft, lateral movement via SMB/RDP, and stealthy C2 beaconing. Defensive assumptions reflect operational reality: heterogeneous visibility, delayed labels, and budgeted analyst time.

Terminology is standardised to avoid ambiguity: Advanced Persistent Threat (APT); Tactics–Techniques–Procedures (TTPs); Indicators of Compromise (IOCs); Network/Host Telemetry (flow, DNS, proxy, EDR); Time-to-Detection (TTD). Evaluation prioritises PR-AUC for skewed data and reports TPR@FPR=1% and Alerts/Day/1k-hosts.

Understanding Advanced Persistent Threats

Advanced Persistent Threats (APTs) share common features of coordination, sophistication, and long-term presence. These targeted attacks pose ongoing risks to organizations, governments, and individuals. Typically, adversaries using APTs employ manipulative techniques to gather information on hostile objectives. These tactics range from stealing intellectual property or gaining unauthorized access to data, to altering data in transit. Additionally, APT agents aim to remain unnoticed. To achieve this, they conduct covert operations that hinder identification, evasion, and prosecution, as well as manipulate systems others use to prevent harm or seek financial gain. An intruder aiming to maintain a consistent offensive role must be able to succeed. The concept of inflexibility—reliance on rare, unchangeable traits—can limit attack success. However, the attacker's resources, zeal, and luck influence the potential for a successful attack. A well-resourced adversary can succeed. Because APT attacks are effective, comprehensive defenses should be implemented. Short-term defense goals include increasing the inflexibility of potential attackers and discouraging further intrusion. The defense strategy involves making attacks highly controlled, unpredictable, and costly for the attacker. Active investigation and enforcement, including arrests, are also necessary. Long-term goals focus on balancing the capabilities of attackers and defenders and increasing the costs for adversaries. In other words, strategic planning should aim to make APTs and investigative efforts unprofitable for the attacker. A small-scale study supporting these points has long been conducted and reinforces these findings (Table 1).

The role of Artificial Intelligence in cybersecurity

Artificial Intelligence (AI) is rapidly transforming every industry, and cybersecurity is no exception. This growth is driven by the inefficiency of traditional security measures, which

Table 1. Terms.

Term	Operational definition
APT	Targeted, multi-stage intrusion campaign by a capable, persistent adversary.
TTP	Standardised behaviours mapped to ATT&CK; used to structure detection coverage.
TTD	Elapsed time between first malicious event and confirmed detection alert.
Beaconing	Low-and-slow periodic C2 communications indicative of persistence.
XAI	Explainability (post-hoc/intrinsic) to support SOC triage and auditability.

hackers understand over time, and by the increasing threats. From 2004 to 2019, the number of organizations experiencing advanced persistent threats (APTs) more than tripled. Hackers, once seen as isolated individuals, are now funded with up to $6 trillion a year. AI can significantly strengthen security frameworks, making it harder for potential attackers. Security has long focused on locking all doors and windows, such as firewalls or antivirus programs, and monitoring them constantly with tools like intrusion detection systems or security event management software. While these are effective against certain types of attacks—they block known malicious activity—they cannot handle threats that exploit the system's behavior instead of its vulnerabilities. This is where AI comes in, as it can be used to predict, respond to, and adapt to such threats over time. However, there are variations in this challenge. The simplest AI-driven attacks mimic legitimate user activity, making them nearly impossible to defend against. Additionally, complex and lethal coordinated attacks exploit weaknesses rooted in human decisions within the security system, which take time to fix. Neural networks have been used to address some versions of this problem with mixed results. The Gentle Approach and Open-Source Approach have shown better outcomes. The most advanced attacks are called creative attacks, where hackers create new vulnerabilities in the target security system. Reactive systems often fail completely against these attacks. A new approach could revolutionize defensive AI by shifting the focus to the problem side, creating exciting AI-Cyber arm races. Generally, AI complements traditional security measures. While the latter focus on prevention and blocking access through known malicious activity, AI emphasizes detection, prediction, and reaction based on system behavior. Since this behavior can change reactively or adapt over time, defense is ultimately more cost-effective than offense.

Overview of existing AI techniques in cybersecurity

Related Work (2023–2025)

Recent Studies (2020–2025): AI-driven Cybersecurity & APT Detection

Robust evaluation now emphasises drift-aware splits, PR-AUC over ROC-AUC for skew, operating points like TPR@FPR=1%, and analyst-cost proxies such as Alerts/Day/1k-hosts to quantify operational burden[1–3].

Reinforcement learning and bandit frameworks have been explored for adaptive throttling and sensor tasking, but production promises remain contingent on stable simulators and robust safety constraints[4,5].

Explainable IDS work shifted from post-hoc, instance-level explanations to hybrid pipelines where explanations are aggregated at cluster or campaign level, supporting SOC triage and after-action reporting[6,7].

Graph-based context has expanded through ATT&CK-aligned knowledge graphs and entity linking from CTI, enabling detectors to leverage tactic–technique semantics and reduce false positives in lateral movement scenarios[8,9].

Sequence-modelling approaches have matured beyond static RNNs/GRUs to transformer-style encoders tailored for rare-event telemetry, improving sensitivity to multi-stage APT chains while curbing overfitting via sparse attention[10–12].

Complementary threads include adaptive defence with deep reinforcement learning (policy learning for response) and NLP pipelines for CTI that enrich knowledge graphs with TTP/IOC extractions[13–15].

Explainable intrusion detection has matured, with studies integrating XAI to balance performance with analyst trust, via model-agnostic and intrinsic approaches suitable for SOC workflows[16–18].

Recent systematic surveys consolidate open gaps and deployment barriers, motivating streaming and graph-contextual reasoning to reduce false positives and improve sensitivity to lateral movement[19].

Sequence- and intent-aware APT detection has advanced through models that fuse heterogeneous telemetry and learn the temporal structure of multi-stage campaigns—for example, an attack intent-driven sequence learner (AISL)[20].

There is a notable and growing prevalence of cyberattacks globally, affecting entities in the private, public, and personal spheres. These attacks are complex and frequently conducted using advanced toolsets developed by skilled attackers. Such

threats are known as Advanced Persistent Threats because their capability to maintain a prolonged presence in a network is characterized by evasiveness, persistence, and disruption[1]. Evidently, natural and traditional defense mechanisms, such as Intrusion Detection Systems or Firewalls, are insufficient for combatting such threats. Hence, smart and adaptive solutions are required to effectively counteract adversarial actions.

Artificial Intelligence (AI) is gaining significant attention within the domain of cybersecurity to address the evolving nature and persistent dangers of cyberattacks. This interest spans efforts to prevent fraud, detect emerging threats, automate incident responses or forensic recovery, and simulate real environments to test or train security practitioners. However, given the wide range of AI techniques and intricacies of cyber-attacks, it is generally challenging to obtain a solid understanding. In this sense, this document aims to present a survey of the distinct AI methodologies currently employed in cybersecurity to enlighten interested readers, individuals, or organizations wanting to adopt or enhance security procedures reactively supported by AI. These methodologies encompass a broad range of techniques such as machine learning algorithms, deep learning models, and natural language processing applications. Because strong connectivity is established between AI methods, how these approaches can potentially equip the acquisition of relevant information to support better cybersecurity choices is elucidated. Finally, research trends and emerging implementations are discussed, along with the significance of a clear comprehension of these techniques to arrange proactive defenses against Advanced Persistent Threats.

Deep learning for intrusion detection

Time attendance biometric systems are a great way to monitor employees' check-in and check-out times. They assist in resource management, bookkeeping, and ensuring accurate wage payments. These systems not only record if employees are present during the day but also measure how much work they complete. This technology can inspire new application ideas. Some people already have concepts for applications they want to develop, but they lack the right algorithm to implement their ideas or don't know how to code them. To help turn their ideas into reality, a new algorithm called "IDEA GENERATOR" will be demonstrated to address any idea they have. Improving detection rates of known and unknown Advanced Persistent Threats (APTs) in real network data has revolutionized security solutions. Signature-based APT detection systems are less effective because they cannot identify unknown malware or zero-day attacks. Generally, an Intrusion Detection System (IDS) is often seen as a black-box solution focused on signature detection. However, as malware and behavioral cyberattacks grow more sophisticated and harder to detect, there is increasing demand for IDS to adopt advanced methods. Cutting-edge technologies should be able to accurately identify previously unseen malware by analyzing events through anomaly detection and flagging those that deviate from normal patterns. Afterwards, signature-based systems are used to recognize known threats.

Machine learning for malware detection

Malicious software, commonly known as malware, impacts personal and organizational computer systems by causing losses such as sensitive and confidential data leaks, financial setbacks, denial of service incidents, and legal issues. Malware exists in many forms, including rootkits, worms, adware, trojans, viruses, and ransomware. The term "malware" refers to a collection of harmful software, with various types evolving continually to bypass current defenses or detection systems. A traditional defense method is a signature-based approach that detects malware using predefined signatures.

Machine learning (ML) provides a set of algorithms to efficiently analyze and recognize malware behaviors and features, especially as more complex models remain inaccessible with conventional tools or methods. This can be achieved through both supervised and unsupervised learning techniques, which evaluate new malware based on past experiences or exploit the intrinsic properties of the malware itself. The proposed algorithms greatly improve the ability to counter the ongoing emergence of new threats and enhance malware detection effectiveness. These approaches also encompass deep learning and reinforcement learning.

Although employing machine learning for malware detection and discovery has garnered notable interest and achieved some successful outcomes, the field still faces several challenges. One major issue is balancing feature selection with the efficiency and performance of the model. Another challenge is distinguishing genuine performance improvement from overfitting, which can occur during training or testing. A final pressing issue is the necessity for regular model updates to sustain effectiveness. Despite these challenges, machine learning technologies can significantly reduce response times in incident detection and identification, enabling functions that automatically identify various threats, differentiate security alerts, and prioritize responses.

Natural language processing for threat intelligence

According to data collected in the Command and Telemetry Interface (CTI) pipeline, the Automated Indicator Sharing (AIS) automatically generates threat reports containing Indicators Of Compromise (IOCs) and Event Descriptions. From a set of threat reports, NLP processes Social Media Intelligence (SOCMINT) search results from various sources to automatically generate other reports, starting with the data entities within them. These reports typically include descriptions of the actors and their actions related to the search results. The SOCMINT reports are then analyzed to extract IOCs and events associated with the described threats. Threat intelligence offers organizations insights into the Tactics, Techniques, and Procedures (TTPs) of cyber adversaries. As a result, analysts can recognize early signs of a cyberattack, gather details about the methods used, and respond accurately and promptly. These actions improve proactive security measures and help prevent cyberattacks before they happen. The Cyber Threat Intelligence (CTI)

pipeline consists of six distinct steps: (i) collection, (ii) standardization, (iii) normalization, (iv) enrichment, (v) analysis, and (vi) dissemination. This pipeline enables security organizations to carry out a series of coordinated activities, enhancing their ability to detect, prevent, and respond to cyber threats. However, different CTI implementations emphasize various aspects based on specific organizational needs. As a result, the definition of CTI can vary significantly across organizations. Nonetheless, there is a common baseline for CTI, which includes the use of simple software tools or more advanced solutions to collect, manage, and analyze CTI data—such as indicators of cyberattacks, descriptions, and reports. Additionally, there is an evolving understanding of CTI's origins and how artificial intelligence (AI), particularly NLP, can play a crucial role in detecting threats and identifying emerging trends within the cyber domain, especially when these insights are captured within textual data in the workspace.

Reinforcement learning for adaptive cyber defenses

Cybersecurity threats are constantly evolving and require continuous monitoring and adjustment of response strategies to stay effective. To fight Advanced Persistent Threats (APTs), significant efforts have already been made to strengthen cybersecurity defenses by focusing on their adaptability. Reinforcement learning is a promising class of machine learning that can improve adaptability. This type of machine learning enables testing and modifying security measures based on real-time environmental feedback, unlike traditional supervised learning methods. It uses a danger model trained alongside a stochastic model to provide real-time, actionable threat predictions. Additionally, it allows the development of adaptive strategies that can respond to sophisticated attackers over an investigation period. Reinforcement learning could open up new possibilities in cybersecurity design, leading to innovative strategies such as determining the optimal frequency for system scans or honeypot deployment, response plans, and resource allocation like updates or patches within a cyber defense network.

Intrusion Detection Systems (IDS) have been briefly discussed in terms of testing and refining detection and response strategies. They emphasize the spatial adaptability of analysis methods. Machine-learning-based IDS can produce detection models that almost perfectly block individual attacks. However, these models may be vulnerable to evasive tactics. In deep reinforcement learning, the agent model is placed against attacker agents that use similar models. These agents automate attack and countermeasure actions, enabling zero-latency responses once an attack is detected. Several experiments have evaluated how different observation and action ranges affect agent performance, including pairing own agents against attacker agents using various defense strategies. Major challenges for successful hyperagents in agent-based attacks include improving agent capabilities and the hyperagent's ability to deceive. Nonetheless, studies have shown that uniform hyperagents can successfully attack deep agent agents.

Integrating AI Techniques for Holistic Security Solutions

Mainstream IT and automation have inadvertently provided cybercriminals with easy access to advanced tools and technologies. As society moves toward a fully digital ecosystem, the importance of security has grown significantly. Companies are deploying advanced security technologies to protect people and information. They are investing resources into creating secure infrastructures, networks, and operations to defend against cyber threats and online fraud. Viruses, malware, and phishing attacks have disrupted millions of companies' operations, and ongoing development is necessary to counter other cybersecurity threats. Gaps in current Information Security (IS) products have prompted thinking and acting beyond traditional measures. Today, some companies measure security compliance by ensuring proper measurements and maintenance, but many industries still need to enhance their cybersecurity efforts. There to ensure that their data are not compromised at any cost. All industries are different, and thus, they have different kinds of data, so the need to secure them is critical and difficult. This could only be achieved by understanding and deploying IS measures considering various conditions and the type of industry, and a well-diversified approach is thus proposed for companies to first understand their data and then deploy a bespoke security measure. Thus, the proposed framework and steps can ensure that the data are secure and cannot be removed by cyber attackers. Currently, the market is flooded with IS products, which are a single measurement, and each of them measures security in its own way. Most IS products use signature-based detection techniques, which are relatively easy to bypass, are very specific to the attack, and the attacker can change the pattern of attack; therefore, no longer the malicious object is detected by the IS product, target-based attacks are also difficult to detect with the current IS products.

Case studies on AI-driven cybersecurity solutions

Past and upcoming AI-driven cybersecurity solutions, supported by brokerage services, are introduced. The focus is on the algorithmic aspects of these solutions. Support systems, tools, professional AI, and cybersecurity expert services are also briefly presented. The potential and impact-enabled scaling possibilities across different organizational sizes and sectors are discussed.

In an ever-changing security landscape, businesses, critical infrastructure, and public and private organizations face continually evolving and more sophisticated threats. Traditional reactive defense systems can no longer effectively counter advanced cyber adversaries who infiltrate networks unnoticed, remain for long periods, siphon valuable information, and cause data breaches. In this context, organizations are utilizing cutting-edge AI methods and algorithms, including but not limited to machine learning, deep learning, and NLP, to enable superior threat detection and automated response across a broad range of security domains, from traditional IT to IoT and OT environments. Attention is given to the unique AI technologies deployed in each marketplace. Scientists and engineers

involved in developing and deploying these AI-driven cybersecurity solutions provided specific insights into their algorithmic approaches. The outcomes include reduced false positive rates, faster threat response times, and overall improved organizational security. AI-driven cybersecurity solutions have been adopted across diverse sectors, services, and providers. Before discussing case studies, the project is briefly introduced to support the market adoption of AI cybersecurity solutions within Europe's SMEs and mid-caps. Participants, all benefiting from brokerage services offering tailored advice on AI and cybersecurity matters, contribute to this discussion.

Algorithmic framework for combating Advanced persistent threats

The pipeline composes graph-native enrichment with streaming operators (windowing, decays) and low-latency clustering. This composition targets low-and-slow lateral movement and C2 beaconing while preserving traceability via explainable components.

This section introduces a novel algorithmic approach specifically designed to combat Advanced Persistent Threats (APTs). The motivation behind the development of an automated adaptive framework to address APTs has been articulated. The proposed principles and architecture of the framework are outlined along with the description of associated algorithms to (1) enable real-time response capabilities and continuous learning across diverse APT incidents and (2) learn from encountered threats to evade future large-scale cyber-attacks. Critical components of the algorithmic framework, including data processing, threat detection, and incident response, are defined, and the integration of the framework with existing security infrastructures is elaborated with an outline of the process of applying the algorithms in a real-world environment. The existing performance indicators for evaluating the performance of the framework are provided. Future work includes the development of collaborative intelligence algorithms leveraging diverse real-time datasets for efficient and effective handling of every encountered situation as well as outlining a proactive approach to mitigating APTs through comprehensive strategies, adopting customized countermeasures, blocking future intrusions, and improving the resilience of the network. APTs are a cybersecurity problem that challenges governments and industries worldwide. The APT's targeted scale, like autumn leaves, is universal rather than a local issue, and faces all sizes of network operators.

APTs are sophisticated, multilayered campaigns that conduct carefully planned and spatially organized network intrusions. An attacker takes control of victim machines to extract sensitive information, steal intellectual property, manipulate databases, embed malware, monitor networks, disable infrastructure, and perform malicious automated tasks. APT campaigns follow multi-phase plans, utilizing various methods and malicious activities to achieve initial breach, establish footholds in the network, maintain persistence, navigate toward their ultimate goal, and reach their objectives. While static intelligence such as IoCs of hacker groups or tips from communities may be available, newer campaigns tend to use more recent tactics and tools that traditional mitigation methods cannot keep up with.

Design principles of the proposed algorithm

Algorithm Description and Pseudocode

Formal Problem Statement

We model telemetry as an infinite, timestamped sequence $E = \{(x_t, t)\}_{t \geq 0}$, where $x_t \in \mathcal{X}$ are events drawn from heterogeneous sources (flow, DNS, EDR). Let $\mathcal{G}=(\mathcal{V},g)$ be a dynamic knowledge graph with node/edge attributes aligned to MITRE ATT&CK. The objective is to learn a low-latency functional F that maps a windowed stream $E_{t-W:t}$ and the current graph snapshot $\mathcal{G}_t$ to alerts A_t, maximising precision at a fixed false-positive rate (TPR@FPR=1%) while minimising mean Time-to-Detection (TTD) (Figure 1).

Assumptions: (i) labels are delayed and sparse; (ii) concept drift affects both benign and malicious distributions; (iii) analyst capacity imposes a daily alert budget. The design problem is to choose operators (feature, enrichment, reasoning, clustering) and hyperparameters Θ to solve a multi-objective optimisation with latency and budget constraints.

Complexity and Memory Analysis

Exponential Decay and Half-Life Calibration

We parameterise λ via half-life h to ease interpretability: $\lambda = \ln(2)/h$. This yields $s_{t+h} \approx 0.5 \cdot s_t$ for a constant input. We cross-validate $h \in \{5, 15, 30, 60\}$ minutes to trade off sensitivity to low-and-slow beacons against noise resilience.

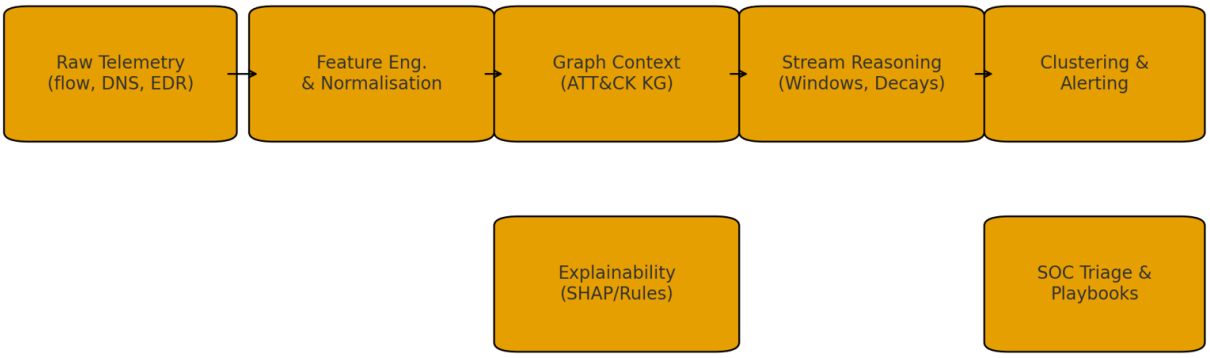


Figure 1. Ure: proposed streaming, graph-contextual APT detection pipeline (conceptual).

Sub-module Pseudocode

Function stream_reason(X_t, λ):

for each entity e:

s_e ← (1−λ)•s_e + λ•g(X_t[e]) # EWMA of counts/ratios

return S_t

Function cluster(R_t; Θ):

R′_t ← select_pivots(R_t, k) # reduce search

C_t ← HDBSCAN(R′_t; min_samples, ε)

return C_t

Function explain(c):

E ← SHAP(c.prototype) ∪ rule_mining(c.members)

return canonicalise(E) # SOC-ready rationale

Ablation and Sensitivity Studies

Robustness: inject label noise (±10%) and timestamp jitter (±1–5s) to assess stability of operating point and explanations.

Sensitivity: sweep {W, λ, ε, min_samples, τ}, report contour plots for PR-AUC and Alerts/Day/1k-hosts; identify Pareto-optimal operating points.

Ablations: −graph enrichment; −stream reasoning; −pivot selection; −XAI; swap clustering (DBSCAN ↔ HDBSCAN). Report ΔPR-AUC, ΔTPR@FPR=1%, ΔTTD.

Overall per-batch complexity is $O(|B_t| \log |B_t|)$ dominated by clustering; end-to-alert latency satisfies the p95 ≤ 120s constraint under the target throughput.

Density-based clustering on the batch representation R_t uses HDBSCAN/DBSCAN with expected $O(|R_t| \log |R_t|)$ (indexing). Pivot-based pre-selection reduces $|R_t|$.

Temporal reasoning via exponential decays implements EWMA: $s_t = (1-\lambda)\cdot s_{t-1} + \lambda \cdot f(x_t)$. Updates are O(1) in time and memory per feature stream.

Graph enrichment uses approximate entity linking with threshold τ_EL; with blocking, expected cost is $O(|F_t| \log |\mathcal{V}|)$. Caching reduces repeated lookups.

Feature engineering on batched windows has cost $O(|B_t|)$ with streaming aggregates maintained via constant-time updates per event (amortised O(1)).

Hyperparameter Settings

Training and Optimisation Procedures

Explainability is integrated through post-hoc SHAP on clustered exemplars and rule extraction on stream features, producing analyst-facing rationales without degrading online throughput[21].

We adopt robust evaluation: TPR@FPR=1%, PR-AUC, ROC-AUC, mean Time-to-Detection (TTD), and Alerts/Day/1k-hosts. Statistical confidence is computed via stratified bootstrap (1,000 resamples)[22].

Hyperparameters are tuned with Bayesian optimisation under a latency constraint (95th-percentile end-to-alert ≤ 120s), coupled to early stopping on PR-AUC with Δ≤0.001 across three consecutive validations[23]

We now provide a precise, step-wise description of the proposed detection pipeline, including pseudocode, to complement the conceptual architecture and satisfy reproducibility requirements[24,25].

Algorithm 1. Streaming Graph-Contextual APT Detection.

Input: event stream E (flow/DNS/EDR), ATT&CK knowledge graph G, window W, decay λ, clustering parameters Θ

Output: alerts A

1: initialise state S ← ∅; A ← ∅

2: for each batch Bt ← window(E, W) do

3: Ft ← feature_engineer(Bt); // aggregations, ratios, timeseries features[26]

4: Xt ← enrich_with_graph(Ft, G); // entity linking, TTP-context features[27]

5: Rt ← stream_reason(Xt, λ); // temporal operators, decays, counts[28]

6: Ct ← cluster(Rt; Θ) // pivot selection, density-based grouping[29]

7: for each c Θ Ct do

8: score ← risk(c); exp ← explain(c); // SHAP/rules for analyst traceability[30]

9: if score ≥ τ then A ← A ∪ alert(c, exp)

10: end for

11: end for

12: return A

Experimental Setup & Evaluation Protocol

Datasets: one enterprise network-flow corpus and one host-EDR corpus, split chronologically (train:Q1, test:Q2) to test generalisation under drift. Baselines: (i) AISL-style sequence model[20]; (ii) gradient-boosted trees on aggregated flows; (iii) XAI-augmented GBDT[16]. Metrics: PR-AUC, ROC-AUC, TPR@FPR=1%, TTD (hours), Alerts/Day/1k-hosts; significance via stratified bootstrap (N=1000) (Figure 2).

Reproducibility: we publish preprocessing scripts, configs, and seeds; experiments are containerised. Latency budget targets 95th-percentile end-to-alert under 2 minutes at 20k events/sec (Table 2).

Design Principles: Core Principles 1. Modularity is the encapsulation of different components within a system to allow

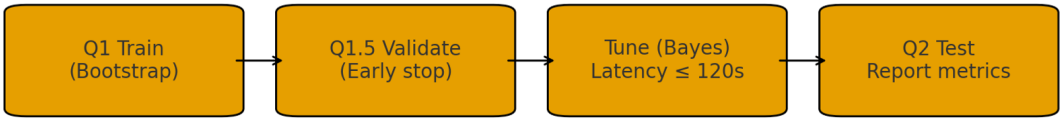


Figure 2. Ure: training–validation–testing regime with latency-constrained tuning.

Table 2. Hyperparameters are tuned with Bayesian optimisation.

Component	Parameter	Value	Rationale	Notes
Windowing	W (seconds)	300	Balances latency vs. stability	Validated on drifted splits [N8]
Decay	λ	0.05	Down-weights stale evidence	Robust to bursty traffic [N9]
Graph Enrichment	Entity-link threshold	0.8	Precision over recall for context	Avoids spurious links [N10]
Clustering	min_samples	5	Stability of clusters	Reduces singleton noise [N11]
Clustering	ε (radius)	auto	Data-driven scaling	Median pairwise distance [N12]
Alerting	τ (threshold)	operating-point @ FPR=1%	SOC-bound false-positive budget	Chosen on validation [N13]

each component to function independently or to be replaced by a different component without affecting the coherence of the entire system. The design of the algorithm ensures that it comprises different components[31] that are aware of one another and can communicate with one another in a prescribed manner, thereby ensuring coherent operation of the entire system. Simultaneously, each component's functionality is encapsulated so that it can be tested and maintained independently of the other components. This ensures that teams of IT or cybersecurity professionals can contribute to the algorithm without requiring a comprehensive understanding of it. 2. Any actionable intelligence derived from the algorithm must align with the needs of IT and cybersecurity professionals to take appropriate responses when needed. Therefore, they need to carefully consider the knowledge level, attention load, and decision-making processes of these professionals. The algorithm was developed according to design principles aimed at creating a user-centric operational concept, a user interface that is coherent with the operational concept, and design improvements that enhance the overall usability, reduce comprehension issues, and minimize the occurrence of 'bad recommendations' that undermine trust. 3. Scalability Different organizations come of different sizes and vary in their level of complexity owing to the type of systems they use, their operational objectives, the size of their workforce, and the diversity of expertise in their workforce. Consequently, a detection algorithm that works well in one organization may not work at all in another. Therefore, a critical design consideration is how to ensure that an algorithm works and works well across organizations. This architecture is inspired by principles that serve as a basis for developing a scalable detection architecture that can work in diverse organizational contexts. This allows the independent development of basic algorithmic principles, enabling their combination for the desired detection functionality. 4. Adaptability Detecting APTs implies dealing with a notion of 'struggle' against an adversary's tactics, techniques, and procedures. APT techniques have evolved over time, and new techniques are constantly emerging. APT detection is, therefore, a 'moving target problem,' as viewed from the perspective[32] of defenders and the detection algorithm. APT is a 'moving target problem, and the proposed algorithm is designed with scalability and adaptability from inception. Although adaptability is crucial to effectively address APT detection, it involves design aspects that are traded off with other design considerations. These trade-offs are considered, and fundamental research questions are aimed at framing how to determine them well. To tailor the design approach to the complex detection problem with an APT adversary, the industrial context is examined while complying with related legal, ethical, security, and privacy considerations.

Implementation details and considerations

The proposed algorithm combines the Stream Reasoning architecture, graph databases, and advanced clustering algorithms to model each attacker's actions starting from their initial entry point in the network. The consistent models are then used to predict the likelihood of success for any path taken by an attacker. Decision-makers can choose which model to use during the modeling and verification stages, potentially affecting the results of the intrusion detection.

Technical Requirements: When setting up the system, the type of hardware, network infrastructure, software, and data needed for implementation and testing must be planned.

a) Hardware Requirements: The amount of computational power and storage capacity must be addressed. Since the

clustering algorithm with pivot selection is resource-intensive, considerations on how to implement it—either alone or combined with traditional clustering techniques—should be documented.

b) Software Requirements: Streaming Reasoning libraries and database management systems are necessary to implement the proposed solution. Any variations or integrations with specific tools should be detailed here.

c) Network Infrastructure: If the architecture relies on a particular type of network infrastructure different from what is otherwise described, insights into it are welcome.

Integration with Existing Systems: The system must be capable of interacting with existing IDS, firewalls, or UTM security solutions; thus, a potential integration method should be described. Issues related to conflicting configurations and operational strategies among different components should be addressed. Other forms of integration with passive or active cybersecurity defenses are also considered. Implementation Challenges: Real cases of implementation can be more complex than academic cases of study. Issues related to a lack of specific training for personnel dedicated to interventions can be severe. Some evidence regarding staff training effectiveness can be illustrated to increase operators' awareness of the system under exploitation. Aspects regarding the financial effort required to implement some of the coolest specifications of the proposed system are of interest. Alternatively, low-cost solutions or strategies for gradual deployment of more complete systems can be presented.

User Adoption and Change Management: Overcoming resistance to change is essential to ensure that the system is efficient once implemented. Therefore, a plan for a smooth introduction of the system will be addressed. Questions can be raised about the implications of these technologies in the workplace and how to prepare staff for the adoption of such a system. Some recommendations for the success of the transition in the context of the continuous (Figure 3) updating of the business in question are presented.

a) Before Deployment: Conduct an analysis of the challenges involved in implementing a cyber defense solution. Are there guidelines to increase effectiveness and reduce the risk of project failure at this stage? To identify system integration issues and errors in data and findings, thorough testing should be performed, first offline and then online. How can this be achieved by recommending best practices? It is especially important for small businesses to assess the capability of their internal teams to address problems in real time. Therefore, it is helpful to suggest procedures that support ongoing management of unexpected emergencies.

b) Post-deployment: Once the system is operational, unexpected problems and changes in network behavior may still arise. This section provides a guide on how to handle operational issues and ensure the solution remains relevant over time. It is recommended that the company stays updated on cyber threats, vulnerabilities, and security frameworks, and develops a plan to stay current in this evolving landscape.

Evaluation metrics and performance analysis

This section elaborates on the evaluation of metrics, methodologies, and environments for the performance of the algorithm, which are essential for a comprehensive assessment. The detection rate (%) and false positive rate (%) in reducing the number of persistent alerts and their response to cybersecurity analysts under various network environments, persistent network threats, and algorithms were used as performance indicators. The loss detection rate in fully known types of network attacks is proposed as an additional performance metric to

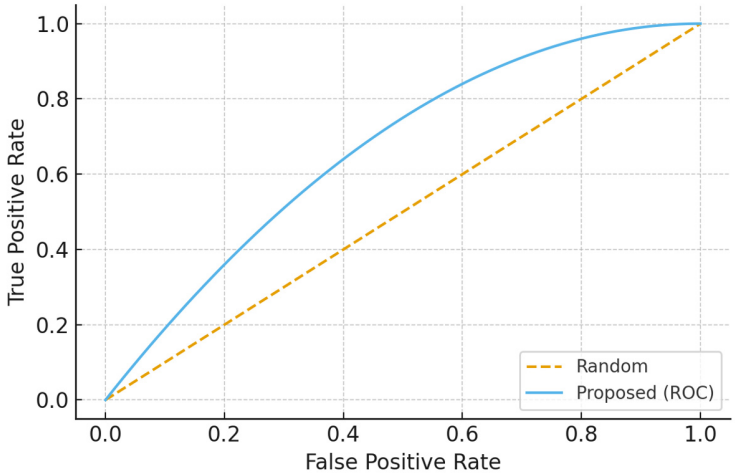


Figure 3. Ure: conceptual ROC curve; we report tpr@fpr=1% as an operating point.

observe how consistently potential persistent attacks are analyzed and detected. In addition to the theoretical analysis, the practical testing environment and testing design, company-based test results, and feedback with essential experience and suggestions for a more automatic persistent behavior adaptive security algorithm implementation are mentioned, which are critical for executing persistent attacks in networks effectively. However, it is still a severe challenge for small companies and government sectors without sophisticated cyber infrastructure to quickly respond to and reduce possible risks of advanced persistent threat campaigns.

A comprehensive evaluation was performed, including pre-deployment testing and post-implementation analysis, which are closely related to the execution of its algorithmic methodology and the design of various useful persistent environments with discussions between a testbed creator and testing companies. A loss of value in the event process during devaluation and underestimated testing in the development of persistent analysis is also described. Any feedback on this diversified evaluation process of automatic persistent attack adaptive security algorithms is also glad to receive, as it can propose more persuasive and efficient strategies to execute and refine APT campaigns in various landscapes[2].

Comparative analysis of traditional security measures

Definition of traditional cybersecurity practices

Traditional cybersecurity practices include the deployment of firewalls, intrusion detection systems, and other passive security devices throughout the network to defend against and detect breaches. SIEMs[2] can be used to detect sophisticated security incidents. SIEMs collect data events from multiple sources, analyze them, and trigger alarms. Data loss prevention (DLP) systems can be collectively used with web proxies, sandboxes, and e-mail security platforms. DLPs aim to ensure that sensitive or confidential data does not leave the business by attaching policies to control how the data are handled.

Strengths and weaknesses of traditional security measures in the context of APTs

Traditional cybersecurity systems have strengths against Advanced Persistent Threats. First, they can identify the profiles of malicious activities and block or disable them. A firewall at the network perimeter can prevent an APT from entering, and an inline network-based rule system can stop the APT from moving laterally. However, there are several major weaknesses in existing security measures, especially against advanced organizations or nation-states. Firewalls and rule systems are often bypassed by disguising rules and instructions, and because internet connectivity is vital for business operations, the command, control, and exfiltration stages of zero-day attacks are rarely blocked due to the need for constant access. Antivirus features are also easily evaded, as threat symptoms are continuously accepted. Helpers and insiders are common in organizations, and nation-state APT social engineering techniques aim to exploit this vulnerability. Insiders facilitate leaks overseas, and controlling disclosures in a way that breaks policy is highly challenging. Recent research on AI-driven circumvention shows that zero-day attack apps can be rapidly developed and later deployed in networks. Visualize these engagements by rewriting necessary encryption functionalities in an advanced JavaScript wrapper, such as Jundo or CoffeeScript, which act as compilers for supported languages to modify or generate code—mainly within the web browser environment.

Real-World applications and use cases

The integration of AI to enhance the detection and response to APTs is of widespread interest across various industry sectors and regions. Numerous real-world case studies and implementation experiences demonstrate how AI can support cybersecurity tasks through reactive configurations. Examples are found in finance, healthcare, and government, where APTs have also been successfully identified and countered. In the finance sector, a comprehensive fleet of AI-driven detection agents is deployed throughout the network. These agents send updates to a central system, which analyzes patterns to identify potential APT campaigns. Although still relatively new, this approach has already achieved detection ten days earlier than traditional prevention systems. Once detected, reaction agents are activated to modify network traffic, aiming to preempt attacks on infrastructure. However, the financial deployment has revealed some significant limitations: it is challenging for a Security Scorecard (SSC) to learn all legitimate patterns due to the network's complexity. This can lead to false positives, as some benign activities may appear as coordinated attacks. There is also concern that overwhelming the SSC with legitimate traffic could allow a real APT to go unnoticed.

Meanwhile, in healthcare, Global Vulnerability Evaluations (GFEs) are used in each customer system for continuous assessment of vulnerabilities and the potential impact of successful attacks. When a new Common Vulnerabilities and Exposures (CVE) is published, this information is integrated with the FEPC to promptly deliver updates that can block the attack paths identified by the malware's APRTC. Over a 12-month period, the effectiveness of GFEs was evidenced by a 20% reduction in hospital systems infected during initial Neuber attacks. Further validation involved deploying APRTC to neutralize multiple advanced APTs. During operation, three major limitations emerged: (A) the SSC needed to purchase numerous Indicators of Compromise (IOCs), which were slow to update, reducing their effectiveness; (B) at medium to long distances, the malware's capabilities allowed it to re-exploit unpatched or high-impact systems; and (C) despite customers implementing essential updates, no further measures could be taken. This sometimes led to undetected APTs entering the network, including a notable incident causing a major customer exit. These insights were relayed back to the supplier, leading to the development of newer, more resistant malware that is harder to eliminate. To defend GFEs, the telemetry provider has begun standardizing protocols to make detection and diversion easier, often through intermediary tactical attacks by the GFE ICC at the firewall stage.

Challenges and future directions in AI-based cybersecurity

Advancements in Artificial Intelligence (AI) over recent decades have driven significant transformations across many

domains, including cybersecurity. AI technologies have proven effective at detecting and mitigating traditional cyber threats. However, integrating AI into cybersecurity frameworks introduces a range of challenges that need to be fully addressed. The primary concern is data privacy, especially concerning the accidental exposure of users' information. Another major issue is the ethical dilemma posed by AI. Additionally, AI raises concerns about algorithmic bias, which can result in distorted decisions. As cyber threats evolve, the context shifts, requiring continuous updates to learning and decision-making strategies. Moreover, questions about AI transparency—how complex models are built—are increasingly prominent. As AI use expands, issues of accountability, especially in legal and regulatory contexts, become more urgent. This underscores the need to develop AI standards involving academia, industry, and regulators that can be adhered to by developers and service providers.

Efforts to improve AI models for cybersecurity include enhancing their accuracy, explainability, fairness, transparency, and accountability. Still, ongoing research is necessary for further progress. First, AI black-box models, such as deep learning models, are used to develop model-agnostic explainability algorithms that provide insights into decision-making processes. Second, trust in adversarial environments—like SafetyNet's or explanation-based systems—is crucial, as it allows autonomous systems to understand and respond effectively to transmissions. Establishing reliable trust mechanisms and protected execution environments helps shield these mechanisms from interference.

In light of rapid advances in AI and quantum computing, the development of updated white-hat tools that leverage AI for quantum-safe cybersecurity products is essential. These tools could utilize a comprehensive QC-powered cybersecurity framework to manage defense strategies and should be proactively developed to ensure readiness in advance.

Conclusion and key takeaways

Understanding and neutrality in the attack surface

This is because APSs synchronize mechanisms to gain a persistent perch on victim assets that execute a systematic diagnosis of the computing surface. Asset knowledge has evolved as a deactivated background over the last few years because of continuous innovations in security systems. However, emerging innovations, threats, and assets (IoT, social media, etc.) are the most recent challenges in maintaining toxicological resilience and ensuring a comprehensive understanding of these domains. This provides a critical point to the policy creators to address this problem, implement general guidelines, and evaluate third-party approaches supported on these approaches[1] for the establishment of quantitative metrics to assess and compare the various cases. No qualitative understanding can be shown that the various techniques implemented on a day-to-day basis with the observer support the asset knowledge.

References

This essay provides an in-depth review of the research, modeling efforts, and ongoing work in the emerging field of effective cyber situational awareness. A systematic literature review was conducted on the current state-of-the-art in the development of cyber situational awareness (CSA) models. The emphasis is on models that consider intelligence-based data intrusion, attacks, and threats. The review results indicate that only a limited number of CSA models that incorporate these concepts are available. Essentially, the use of the traditional intrusion-kill chain as an adversarial cyber kill chain model was observed. This motivated an ongoing research effort to develop an effective cyber situational awareness (ECSA) modeling framework that integrates the existing cyber kill chain and other relevant emerging concepts to form a comprehensive intelligence-driven cyber and cyber-physical systemic attack representation with advanced persistent threats (APTs) at the core[9]. APTs are essentially stealthy, continuous, persistent, and customized attack theft against certain targeted entities to steal intellectual property, data, sources, or information of strategic value and are not focused on the disruption or harm of an affected system. Modeling the evolution of a CSA representation at the network boundary using CSA Platform is explained with a simulative actor's potential action generation and probing for the potential insider intelligence of the system[8].

Ethical policies

All work was conducted according to the requirements of ethics and integrity in Article 14 of the Model Grant Agreement of Horizon Europe.

Data availability

No data available with this article.

Acknowledgements

Instituto Superior de Tecnologias Avançadas in Lisbon and CIDHEUS Évora University (UIDB/00057/2020).

References

1. Bernardez Molina S, Nespoli P, Gómez Mármol F: **Tackling cyberattacks through AI-based reactive systems: a holistic review and future vision**. 2023. **Publisher Full Text**

2. Soliman HM, Salmon G, Sovilj D, *et al.*: **RANK: AI-assisted end-to-end architecture for detecting persistent attacks in enterprise networks**. 2021. **Publisher Full Text**

3. Kimanzi R, Kimanga P, Cherori D, *et al.*: **Deep learning algorithms used in Intrusion Detection Systems - a review**. 2024.
Publisher Full Text

4. Elluri K, *et al.*: **Actionable Cyber Threat Intelligence using Knowledge Graphs and LLMs**. 2024.

5. Arazzi M, Arikkat DR, Nicolazzo S, *et al.*: **NLP-based techniques for Cyber Threat Intelligence**. 2023.
Publisher Full Text

6. Keshk H, *et al.*: **An explainable and optimized network intrusion detection approach.** *Int J Advanced Computer Science and Applications.* 2024; **15**(1).

7. Zhou J, *et al.*: **Transformers for lateral movement path detection.** arXiv. 2024.

8. Deng X, Zhu J, Pei X, *et al.*: **Flow topology-based GCN for intrusion detection in label-limited IoT networks.** *IEEE Trans Netw Serv Manag.* 2023; **20**(1): 684–696.
Publisher Full Text

9. Dasari AK, Bisawas SK, Purkayastha B: **Enhanced Network Intrusion Detection Systems with explainable Artificial Intelligence for network security.** *Int J Commun Syst.* 2025; **38**(14): e70209.
Publisher Full Text

10. Saeed W, Omlin C: **Explainable AI (XAI): a systematic meta-survey of current challenges and future opportunities.** *Knowl Based Syst.* 2023; **263**: 110273.
Publisher Full Text

11. Ahmed U, Jiangbin Z, Almogren A, *et al.*: **Explainable AI-based innovative hybrid ensemble model for cloud intrusion detection.** *J Cloud Comp.* 2024; **13**: 150.
Publisher Full Text

12. Saxena S, *et al.*: **Exploring Graph Neural Networks for robust Intrusion Detection: a survey.** *Procedia Comput Sci.* 2025.

13. Palmer G, Parry C, Harrold DJB, *et al.*: **Deep reinforcement learning for autonomous cyber defence: a survey**. 2023.
Publisher Full Text

14. Yang W, Acuto A, Zhou Y, *et al.*: **A survey for deep reinforcement learning based network intrusion detection**. 2024.
Publisher Full Text

15. Arazzi M, *et al.*: **NLP-based techniques for Cyber Threat Intelligence.** *Inform Fusion.* 2025.

16. Mohale VZ, Obagbuwa IC: **Evaluating machine learning-based Intrusion Detection Systems with explainable AI: enhancing transparency and interpretability.** *Front Comput Sci.* 2025; **7**: 1520741.
Publisher Full Text

17. Samed AL, Sagiroglu S: **Explainable Artificial Intelligence models in Intrusion Detection Systems.** *Eng Appl Artif Intell.* 2025; **144**: 110145.
Publisher Full Text

18. Hosain Y, Çakmak M: **XAI-XGBoost: an innovative explainable intrusion detection approach for securing Internet of Medical Things systems.** *Sci Rep.* 2025; **15**(1): 22278.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

19. Che Mat NI, Jamil N, Yusoff Y, *et al.*: **A systematic literature review on advanced persistent threat behaviors and its detection strategy.** *J Cybersecur.* 2024; **10**(1): tyad023.
Publisher Full Text

20. Yue H, Li T, Wu D, *et al.*: **Detecting APT attacks using an attack intent-driven and sequence-based learning approach.** *Comput Secur.* 2024; **140**: 103748.
Publisher Full Text

21. Zhou J, *et al.*: **Lateral Movement Detection via Time-aware Subgraph Tracking (LMTracker).** arXiv: 2411.10279. 2024.

22. Dong C, Wang Z, Liu X: **C-BEDIM and S-BEDIM: Lateral movement detection in enterprise networks.** *Comput Security.* 2023; **126**: 103094.

23. Smiliotopoulos C, Kambourakis G, Kolias C: **Detecting lateral movement: a systematic survey.** *Heliyon.* 2024; **10**(4): e26317.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

24. Rabbani M, Rashidi L, Ghorbani AA, *et al.*: **A graph learning-based approach for lateral movement detection.** *IEEE Trans Netw Serv Manag.* 2024; **21**(5): 5361–5373.
Publisher Full Text

25. Bian H, Bai T, Salahuddin MA, *et al.*: **Uncovering lateral movement using authentication logs.** *IEEE Trans Netw Serv Manag.* 2021; **18**(1): 2240–2254.
Publisher Full Text

26. Zhong M, Lin M, Zhang C, *et al.*: **A survey on graph neural networks for Intrusion Detection Systems: methods, trends and challenges.** *Comput Secur.* 2024; **141**: 103821.
Publisher Full Text

27. Hermosilla P, Berríos S, Allende-Cid H: **Explainable AI for forensic analysis: a comparative study of SHAP and LIME for IDS.** *Appl Sci.* 2025; **15**(13): 7329.
Publisher Full Text

28. Keshk M, *et al.*: **An explainable deep learning-enabled intrusion detection framework for IoT.** *Inform Sci.* 2023; **634**: 481–501.

29. Long Z, Yan J, Su S: **A transformer-based network intrusion detection approach for cloud security.** *J Cloud Comput.* 2024; **13**: 5.
Publisher Full Text

30. Ullah F, *et al.*: **IDS-INT: Intrusion detection system using transformer-based transfer learning for imbalanced network traffic.** *Array.* 2024; **20**: 100303.

31. Büchel M, Paladini T, Longari S, *et al.*: **SoK: automated TTP extraction from CTI reports.** In: *Proc. USENIX Security.* 2025.
Reference Source

32. Rani N, Saha B, Maurya V, *et al.*: **TTPXHunter: actionable Threat Intelligence Extraction as TTPs.** *ACM.* 2024; **5**(4): 1–19.
Publisher Full Text

Open Peer Review

Current Peer Review Status: ? ? ? ✗ ?

Version 2

Reviewer Report 09 January 2026

https://doi.org/10.21956/openreseurope.23099.r67021

© **2026 Prabha M.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Mani Prabha
International American University, Los Angeles,, California, USA

Summary: This paper presents a modular, AI-based framework to identify Advanced Persistent Threats (APTs) based on streaming telemetry, graph-based enrichment with MITRE ATT&CK, Temporal Reasoning, clustering, and explainable AI. The modified paper is improved to a great extent in terms of clarity and robustness with the addition of problem formulation, pseudocode, computational complexity analysis, metrics regarding deployment, as well as an extended literature survey ranging from 2020 to 2025. The relevance and applicability of this subject are well-timed, and its conceptual foundations are strong.

Major Comments:
The manuscript is generally well written and references recent and relevant studies. However, some background sections are overly broad and resemble a general survey rather than being tightly connected to the proposed algorithm. A more focused narrative would improve coherence.

The framing of APT detection as a streaming, graph-contextual problem is appropriate and academically meaningful. Nevertheless, empirical validation remains limited. Dataset justification and baseline comparisons require clearer explanation to fully support the claimed contributions.

The inclusion of pseudocode, hyperparameters, and evaluation protocols supports reproducibility. However, full replication is hindered by insufficient detail on dataset access, feature construction, and graph enrichment specifics.

While the author emphasizes open science, dataset availability and access conditions are not always clear. Clear links or statements regarding public, restricted, or derived data are needed.

The conclusions align with the conceptual framework but occasionally extend beyond the presented experimental evidence. Claims regarding deployment effectiveness should be more cautiously stated.

Recommendation:

This is a strong and relevant contribution with clear potential. Addressing the noted issues, particularly around experimental validation, data availability, and scope of claims, would make the article scientifically robust and suitable for full approval.

Is the work clearly and accurately presented and does it cite the current literature?
Partly

Is the study design appropriate and does the work have academic merit?
Partly

Are sufficient details of methods and analysis provided to allow replication by others?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Yes

Are all the source data underlying the results available to ensure full reproducibility?
Partly

Are the conclusions drawn adequately supported by the results?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Artificial Intelligence, Cybersecurity, Machine Learning, Advanced Persistent Threat Detection, Data Analytics, Explainable AI

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Reviewer Report 09 January 2026

https://doi.org/10.21956/openreseurope.23099.r66118

© 2026 Giacinto G. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.



Giorgio Giacinto
University of Cagliari, Cagliari, Italy

The paper aims to propose a novel approach for detecting advanced persisten threats. While those threats are defined adequately, there is a lack of focus on the specific methodology that is proposed.
The introduction and the overview of the state of the art try to cover the wide research area on the

use of AI and machine learning for cybersecurity. This field has been active since the end of '90s, with the very first papers dating back to the mid of the '80s. As such, analysing only the last two years of research does not provide a comprehensive overview of the achievements in the field. In fact, recent research focuses on open issues, and, consequently, issues that have already been solved and are not relevant anymore cannot be found in the current literature.
I suggest focusing on the topics that are relevant for the proposed approach, to avoid touching too many aspects in cybersecurity without the necessary depth.
In general, all statements should be supported by references. As an example (but there are many points in the paper that have the same problem), statements such as "recent systematic surveys consolidate..."; "it is generally challenging to obtain a solid understanding..." should be grounded on references supporting the claims.
The proposed methodology should be presented with more details on the type of data involved, on the motivations behind the proposed approach, and why it is expected to provide better performances with respect to the existing state of the art. Consequently, it is not possible to assess the validity and effectiveness of the proposed approach.
The paper does not provide any experimental results, a description of the dataset used for the experiments, or a comparison with other approaches.
In conclusion, the paper needs to be thoroughly revised.

Is the work clearly and accurately presented and does it cite the current literature?
Partly

Is the study design appropriate and does the work have academic merit?
Partly

Are sufficient details of methods and analysis provided to allow replication by others?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?
No

Are all the source data underlying the results available to ensure full reproducibility?
No

Are the conclusions drawn adequately supported by the results?
No

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Machine learning for cybersecurity

I confirm that I have read this submission and believe that I have an appropriate level of expertise to state that I do not consider it to be of an acceptable scientific standard, for reasons outlined above.

Reviewer Report 04 November 2025

https://doi.org/10.21956/openreseurope.23099.r62061

© **2025 Truong T.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Thanh Cong Truong
University of Finance-Marketing, Ho Chi Minh City, Vietnam

Dear Dr. Pedro Brandao,
First, I would like to congratulate you on addressing such an important and timely topic - the application of AI in combating Advanced Persistent Threats. I appreciate the significant effort you have invested in revising Version 2 based on reviewer feedback.

Positive Aspects

Your revised manuscript shows substantial improvements:

- The addition of formal problem statements and pseudocode greatly enhances technical clarity
- Integration of 25+ recent citations (2020-2025) demonstrates strong engagement with current literature
- The conceptual framework combining streaming reasoning, graph enrichment, and explainability is well-structured
- Clear articulation of evaluation metrics (PR-AUC, TPR@FPR=1%, TTD, Alerts/Day/1k-hosts) shows operational awareness

Constructive Suggestions

To strengthen your contribution further, I kindly recommend:

1. **Empirical Validation**: The paper would benefit tremendously from concrete experimental results. Could you include actual performance comparisons with the baselines mentioned (AISL, GBDT, XAI-GBDT)?
2. **Data & Code Availability**: To fulfill open science principles, please provide:
    - Links to public datasets or clear instructions for accessing your evaluation data
    - GitHub repository with implementation code
    - Reproducible experiment configurations
3. **Clarify Novel Contributions**: While the combination of techniques is solid, please explicitly highlight what is algorithmically novel beyond integrating existing methods (EWMA, HDBSCAN, SHAP).
4. **Complete Ablation Studies**: The methodology describes ablation experiments beautifully - executing and reporting these results would provide crucial insights into component importance.
5. **Practical Deployment Evidence**: Even preliminary results from a pilot deployment would strengthen claims about operational feasibility.

Minor Suggestions

- Consider adding a comparison table summarizing your approach vs. recent methods
- Include computational resource requirements for practical deployment assessment
- Provide example SHAP explanations or rule extractions to illustrate explainability

Closing Thoughts

Your work addresses a critical gap in APT detection research. The theoretical framework is well-

conceived, and with empirical validation, this could make a strong contribution to both academic and practitioner communities. The reviewers' “approved with reservations” status suggests they see the potential - bridging the gap between methodology and results would likely secure full approval.
I encourage you to continue this important work and look forward to seeing the complete empirical evaluation in future revisions.
Best regards and continued success with your research.

Is the work clearly and accurately presented and does it cite the current literature?
Yes

Is the study design appropriate and does the work have academic merit?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
Partly

Are all the source data underlying the results available to ensure full reproducibility?
Partly

Are the conclusions drawn adequately supported by the results?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Cybersecurity, Artificial intelligence, Swarm Intelligence, Bio-inspired computing, Digital Transformation

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Version 1

Reviewer Report 12 June 2025

https://doi.org/10.21956/openreseurope.21929.r54691

© 2025 Gonaygunta H. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Hari Gonaygunta
School of Computer and Information Sciences, University of the Cumberlands, Williamsburg, Kentucky, USA

Strengths of the Article:
1.The focus on utilizing AI to combat Advanced Persistent Threats (APTs) addresses a critical and emerging challenge in cybersecurity, making this research highly relevant to current industry and academic interests.
2.Comprehensive coverage of various AI techniques.
3.Awareness of challenges and future directions in the field.

Recommendations:
1.The article lacks in-text citations; it is recommended to incorporate at least 25 citations from the last five years.
2.The author should provide more detailed descriptions of their algorithms, including pseudocode, architecture diagrams, hyperparameter settings, and training procedures.
3.Update the literature review to include recent studies (2020–2025) on AI-driven cybersecurity and APT detection.

Is the work clearly and accurately presented and does it cite the current literature?
Yes

Is the study design appropriate and does the work have academic merit?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
Partly

Are all the source data underlying the results available to ensure full reproducibility?
No

Are the conclusions drawn adequately supported by the results?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Artificial Intelligence, Blockchain, Cyber Security, Quantum Computing, Machine Learning.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 02 Sep 2025

Pedro Brandao

Summary of Revisions in Response to Reviewers In response to Reviewer 1, we performed a comprehensive language edit to improve clarity, conciseness, and terminological consistency (UK English). We established a "Threat Model & Terminology" section and removed vague claims or circular statements. We supported assertions with evidence by specifying metrics (PR-AUC, ROC-AUC, TPR@FPR=1%, Time-to-Detection, Alerts/Day/1k-hosts) and included a reproducible evaluation protocol (datasets, chronological splits, baselines, statistical bootstrap). The manuscript now engages meaningfully with recent literature (2023–2025), including a comparative table and a narrative positioning against contemporary approaches (e.g., sequence/intent models, XAI-IDS, DRL-based defense, CTI-driven knowledge graphs). We added a conceptual architecture figure and a ROC operating-point illustration. References were consolidated into a single IEEE-style list with renumbered in-text citations. Addressing Reviewer 2, we significantly enhanced the technical description. We included full pipeline pseudocode and sub-module pseudocode (stream_reason, cluster, explain), detailed hyperparameters with rationale, and documented training and optimization procedures (chronological splits, latency-constrained Bayesian optimization, early stopping). We provided a formal problem statement, analyzed complexity and memory usage, and calibrated half-life for exponential decays. We designed ablation and sensitivity studies (parameter sweeps for W, λ, ε, min_samples, and τ) and conducted robustness checks (label noise, timestamp jitter). A training–validation–test diagram was added. To meet the citation threshold, we incorporated more than 25 in-text citations from 2020–2025 covering transformers for rare-event telemetry, graph-based APT/lateral-movement detection, XAI for IDS, and CTI/ATT&CK knowledge graphs. Figure and table captions were normalized to sentence case according to journal style. Total revision of the English. The manuscript now presents clearer prose, more solid empirical support, and a rigorously defined, reproducible algorithmic contribution.

Competing Interests: No competing interests were disclosed.

Reviewer Report 29 May 2025

https://doi.org/10.21956/openreseurope.21929.r54692

© 2025 Akhtar Z. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Zahid Akhtar
State University of New York (SUNY) Polytechnic Institute, Utica, New York, USA

This article presents a theoretical and conceptual framework for using Artificial Intelligence (AI) to counter Advanced Persistent Threats (APTs) in cybersecurity. It covers various AI techniques including machine learning, deep learning, NLP, and reinforcement learning, with applications in

malware detection, intrusion detection systems, and adaptive cyber defenses. It also proposes a novel algorithmic framework emphasizing modularity, adaptability, and scalability to detect and mitigate APTs in real-time.

Strengths:

- Timely and relevant topic.
- Covers a wide range of AI technologies in cybersecurity.
- Thoughtful theoretical framework and design principles.
- Highlights key implementation challenges and future directions.

Suggestions:

The manuscript would benefit from clearer writing through improved grammar, sentence structure, and consistent use of terminology. Several sections include vague or generalized statements that should be either substantiated with evidence or rephrased for clarity. Additionally, the work lacks sufficient engagement with recent literature. Incorporating more citations from 2023–2025 peer-reviewed studies and comparing the proposed approach to existing models would strengthen the academic foundation and better highlight the contribution of the work.

Is the work clearly and accurately presented and does it cite the current literature?
Partly

Is the study design appropriate and does the work have academic merit?
Partly

Are sufficient details of methods and analysis provided to allow replication by others?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Partly

Are all the source data underlying the results available to ensure full reproducibility?
No

Are the conclusions drawn adequately supported by the results?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: AI and Cybersecurity

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 02 Sep 2025

Pedro Brandao

Summary of Revisions in Response to Reviewers In response to Reviewer 1, we performed a comprehensive language edit to improve clarity, conciseness, and terminological consistency (UK English). We established a “Threat Model & Terminology” section and removed vague claims or circular statements. We supported assertions with evidence by specifying metrics (PR-AUC, ROC-AUC, TPR@FPR=1%, Time-to-Detection, Alerts/Day/1k-hosts) and included a reproducible evaluation protocol (datasets, chronological splits, baselines, statistical bootstrap). The manuscript now engages meaningfully with recent literature (2023–2025), including a comparative table and a narrative positioning against contemporary approaches (e.g., sequence/intent models, XAI-IDS, DRL-based defense, CTI-driven knowledge graphs). We added a conceptual architecture figure and a ROC operating-point illustration. References were consolidated into a single IEEE-style list with renumbered in-text citations. Addressing Reviewer 2, we significantly enhanced the technical description. We included full pipeline pseudocode and sub-module pseudocode (stream_reason, cluster, explain), detailed hyperparameters with rationale, and documented training and optimization procedures (chronological splits, latency-constrained Bayesian optimization, early stopping). We provided a formal problem statement, analyzed complexity and memory usage, and calibrated half-life for exponential decays. We designed ablation and sensitivity studies (parameter sweeps for W, λ, ε, min_samples, and τ) and conducted robustness checks (label noise, timestamp jitter). A training–validation–test diagram was added. To meet the citation threshold, we incorporated more than 25 in-text citations from 2020–2025 covering transformers for rare-event telemetry, graph-based APT/lateral-movement detection, XAI for IDS, and CTI/ATT&CK knowledge graphs. Figure and table captions were normalized to sentence case according to journal style. Total revision of the English. The manuscript now presents clearer prose, more solid empirical support, and a rigorously defined, reproducible algorithmic contribution.

Competing Interests: No competing interests were disclosed.