BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

AI/ML-Based Multi-Stage Detection and Automated Defence Against Advanced Persistent Threats: A Review and Conceptual Framework

[1]Sajid Tadvi, [2]Prof. Geeta Kodabagi,

[1]Student (T.E. Undergraduate), [2]Mentor,
[1]Department of Artificial Intelligence and Data Science,
[1]Ajeenkya DY Patil School of Engineering, Pune, India
[1]Email: sajid.tadvi@dypic.in [2]Email: geetakodabagi@dypic.in

Abstract

Advanced Persistent Threats (APTs) are targeted, long-term cyberattacks that unfold in several stages and are deliberately designed to stay hidden, so detecting them from isolated alerts is unreliable. Machine learning (ML), behavioural analysis, event correlation, provenance analysis and security information and event management (SIEM) are increasingly applied to this problem. This paper reviews fifteen studies published between 2018 and 2026 that cover APT tactics, techniques and procedures (TTPs), supervised, deep, hybrid and evolutionary detectors, provenance-graph learning, attack-stage prediction and uncertainty-aware classification. A criterion-based comparison shows that none of the reviewed studies combines full-lifecycle coverage, multi-source data fusion, prediction of the attacker's next action, handling of uncertainty and drift, and validation in a realistic environment. To address this need, we propose a conceptual framework for multi-source data fusion, multi-stage detection, stage and next-action prediction, uncertainty-aware prioritisation and automated response, and define an evaluation protocol for it. The framework is proposed rather than implemented; open challenges and future directions are discussed.

Keywords: Advanced Persistent Threats, Cybersecurity, Machine Learning, Multi-Stage Detection, Provenance Graphs, SIEM, Automated Response.

I. Introduction

Advanced Persistent Threats (APTs) are targeted, sophisticated and long-term cyberattacks in which attackers gain unauthorised access to a system and then remain undetected for an extended period [1, 2, 10]. An APT campaign is generally multi-stage: it involves initial access, persistence, lateral movement, command and control, and data exfiltration. Attackers often rely on legitimate tools, customised techniques and stealthy methods, which allows them to evade conventional security mechanisms [1].

Traditional security systems may analyse individual alerts separately, which makes it difficult to understand the complete behaviour of an attack [5]. For this reason, recent research increasingly applies AI/ML, behavioural analysis, event correlation, provenance analysis and SIEM to APT detection [4, 5, 7, 12, 15]. Models of this kind have reported very high accuracy on some datasets; for example, Random Forest and XGBoost reached 99.6% accuracy in identifying APT alerts [5].

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Nevertheless, four limitations remain for practical use. First, datasets that represent all stages of an APT attack are scarce, are often synthetic or derived from other data, and suffer from severe class imbalance [3, 7, 14]. Second, many studies address only one or a few stages instead of the complete lifecycle, and they analyse network, host, log and vulnerability information independently. Third, most systems decide whether an activity is malicious but do not estimate what the attacker is likely to do next. Fourth, even high-performing models need validation in realistic environments, because changing attacker behaviour can reduce their performance against unseen campaigns.

This paper makes the following contributions.

- A review and outline of recent research on ML-based APT detection, provenance-graph learning, stage prediction and automated response.
- A conceptual, end-to-end framework that brings these themes together, and a proposed evaluation protocol.
- A discussion of applications, limitations and future directions.

The rest of the paper is organised as follows. The background is presented in Section II, the literature is reviewed in Section III, research gaps are discussed in Section IV, the framework is presented in Section V, and applications, limitations, future work and conclusions are presented in Sections VI–IX.

II. BACKGROUND

A. APT Lifecycle and TTPs

An APT is characterised by the tactics, techniques and procedures (TTPs) of the attacker across a campaign. The typical stages are initial access, persistence, lateral movement, command and control (C2), and data exfiltration. Each stage leaves different evidence: network traffic, host activity, security alerts and system logs. Detection research is therefore moving from signature-based methods to behaviour, then context, and ultimately attacker intent. Network-level behaviour remains an important source of evidence [1].

B. ML-Based Detection and Prediction

Supervised classifiers such as Random Forest and XGBoost have been used to detect APT-related events and to reduce alerts [5]. Hybrid ensembles and CNN-LSTM models have been applied to an APT-aware dataset derived from UNSW-NB15 [7], and tree-based genetic programming has been applied to the DAPT-2020 dataset [11]. For stage prediction, a sequence of observations $o_{1:T}$ can be modelled with a Hidden Markov Model (HMM), whose most likely stage sequence is obtained by Viterbi decoding [3, 18]:

$$\hat{s}_{1:T} = \text{argmax}\, P(s_1)P(o_1|s_1) \prod_{t=2..T} P(s_t|s_{t-1})P(o_t|s_t) \quad (1)$$

C. Provenance-Graph Learning

A provenance graph $G = (V, E)$ represents system entities, such as processes, files and network sockets, as nodes V and the events between them as edges E. Graph neural networks can learn node representations directly from this structure. In GraphSAGE [16], which is used in [4, 13], the hidden state of a node is updated from its neighbours:

$$h_v^{(l+1)} = \sigma(\, W \cdot [\, h_v^{(l)} \,\|\, \text{AGG}(\{\, h_u^{(l)} : u \in N(v) \,\}) \,] \,) \quad (2)$$

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

where σ is a non-linearity, W a learned weight matrix and AGG a neighbourhood aggregator. The resulting embeddings are then scored by anomaly detectors or clustered to find malicious nodes or subgraphs [4, 12, 13].

D. Uncertainty and SIEM-Based Response

Evidential deep learning lets a classifier output the parameters α_k of a Dirichlet distribution over K classes, from which the uncertainty $u = K / \Sigma_k \alpha_k$ follows [17]. Applied to attack-stage classification, this gives calibrated confidence and can flag out-of-distribution inputs that may signal changed tactics [8]. On the operational side, SIEM-based pipelines combine event-detection modules with ML to reduce alerts and support automated post-alert decisions [5]. An analyst-support pipeline can be summarised as detection, classification, prioritisation, decision and response.

III. LITERATURE REVIEW

The methodology, main finding and limitations reported by the authors of the fifteen primary studies were analysed. Table I summarises this information.

Surveys and landscape studies. Alageel and Maffeis [1] investigate network-based APT TTPs and note that better multi-protocol behavioural detection is required. Yuldoshkhujaev *et al.* [2] give a decade-long view of APT actors, techniques and trends from threat-intelligence reports, but this historical knowledge must still be translated into real-time adaptive detection and prediction. The surveys in [9, 10, 14] cover APT behaviours, the lifecycle and IoT environments, and point to the need for better integration of attack characteristics, to evolving techniques, and to limited datasets and class imbalance.

ML and ensemble detectors. Ali *et al.* [5] combine Random Forest, XGBoost and a SIEM with six event-detection modules and 14 malicious-behaviour types, reporting 99.6% accuracy, although dataset limitations and generalisation remain open. Tijjani *et al.* [7] show that hybrid ensembles outperform single classifiers and that CNN-LSTM surpasses traditional ML, but only on a derived dataset. Mamun *et al.* [11] reach 79.75% average balanced accuracy with genetic programming and a very fast inference time of 0.0003 ms, at the cost of a 5.97% false-positive rate in binary mode and roughly 16–17 minutes of training. MLAPT [15] predicts APT attacks in their early steps with 84.8% accuracy, but dates from 2018. Accuracy figures from different datasets are not directly comparable.

Provenance-graph methods. APT-MCL [4] learns structural and behavioural views of a provenance graph without supervision and reaches a macro-F1 of up to 0.999 on CADETS, although supervised models still win on some datasets and training cost is higher. RAPID [12] obtains graph-level precision and recall of 1.00 on THEIA, CADETS and Public Arena at about 3.6×10^4 logs/s, but node-level precision is only 0.55–0.74 and periodic updates are needed for concept drift. ProcSAGE [13] reaches 98% accuracy and F1 on StreamSpot with 69% less time and 78% less memory than StreamSpot, but was tested on one dataset and assumes attack-free training data.

Stage prediction, uncertainty and response. E-HiDNet [3] couples a CNN-RNN with an HMM and modified Viterbi decoding to predict APT stages with 98.8–100% accuracy, but only on a synthetic dataset. Gaudenzi *et al.* [8] use evidential deep learning for uncertainty-aware stage classification, in preliminary simulated experiments without reported metrics. Gilbert *et al.* [6] show that advanced detection combined with a structured response plan can improve APT defence, though real environments add complexity that simulations miss.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

TABLE I

SUMMARY OF THE REVIEWED STUDIES

Ref.	Approach and main finding	Reported limitations
[1] 2026	Investigation of network-based APT TTPs, attack techniques and evasion; highlights the importance of network-level attacker behaviour.	Better multi-protocol behavioural detection is required.
[2] 2026	Longitudinal analysis of APT campaigns using CTI and public reports; gives a decade-long view of actors, techniques and global trends.	Historical CTI must be translated into real-time adaptive detection and prediction.
[3] 2026	E-HiDNet: CNN-RNN with a Hidden Markov Model and modified Viterbi decoding; predicts APT stages with 98.8–100% accuracy on a synthetic dataset.	Evaluated only on synthetic data; real-world performance is not demonstrated.
[4] 2026	APT-MCL: unsupervised multi-view provenance-graph learning (GraphSAGE, iForest, co-training); macro-F1 up to 0.999 on CADETS.	Supervised models still win on some datasets; higher training cost; threshold-sensitive.
[5] 2025	Random Forest, XGBoost and SIEM with six event-detection modules and 14 behaviour types; ML-based alert reduction and automated post-alert decisions; 99.6% accuracy.	Dataset limitations; real-world generalisation.
[6] 2025	Analysis of APT characteristics, anomaly detection, behavioural analytics, threat intelligence and incident response.	Real environments may add complexities not captured by simulations.
[7] 2025	Multi-stage hybrid ensembles and CNN-LSTM on an APT-aware dataset derived from UNSW-NB15; hybrids beat single classifiers.	Simulation results on a derived dataset, not real APT traffic.
[8] 2025	Evidential deep learning (Dirichlet output) for uncertainty-aware stage classification; flags out-of-distribution inputs.	Preliminary simulated experiments; no reported metrics.
[9] 2024	PRISMA-based systematic review of APT behaviours and detection strategies.	Better integration of different attack characteristics is required.
[10] 2024	Survey of the APT lifecycle and detection techniques.	Evolving attack techniques make reliable detection difficult.
[11] 2024	Tree-based genetic programming on DAPT-2020; 79.75% average balanced accuracy, 0.0003 ms inference.	5.97% false-positive rate (binary mode); about 16–17 min training.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Ref.	Approach and main finding	Reported limitations
[12] 2024	RAPID: self-supervised Bi-LSTM on provenance-log embeddings with graph tracing and DBSCAN; graph-level precision and recall of 1.00.	Needs updates for concept drift; node-level precision 0.55–0.74.
[13] 2024	ProcSAGE: process-centric provenance graphs with GraphSAGE and clustering; 98% accuracy and F1 on StreamSpot.	One dataset; assumes attack-free training data; no concept-drift handling.
[14] 2022	Review of ML-based IDS and APT challenges in IoT environments.	Limited comprehensive datasets; severe class imbalance.
[15] 2018	MLAPT: eight detection modules, alert correlation and an ML module that predicts APT progression; 84.8% accuracy in early steps.	Dates from 2018; may not reflect current techniques and tooling.

IV. RESEARCH GAP ANALYSIS

All reviewed studies were marked against five criteria that we regard as the most important research challenges: lifecycle coverage, multi-source fusion, next-step prediction, uncertainty and drift handling, and real-world validation. The rating in Fig. 1 is the authors' qualitative reading of the information summarised in Table I; "not addressed" means that the criterion is not reported there. It is not a quantitative benchmark.

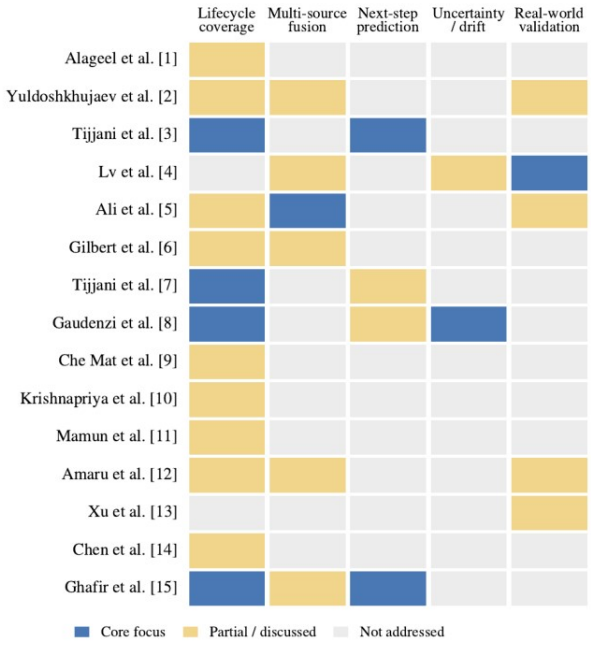


Fig. 1. Coverage of five design criteria by the reviewed studies (qualitative rating).

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

The following can be seen from the matrix.

- **G1 – Lifecycle fragmentation.** Most studies address only one or a few stages. Datasets that represent every stage are scarce, so connections between attack activities can be missed.
- **G2 – Siloed data.** Network, host, log, vulnerability and behavioural information are rarely combined; the SIEM-based work in [5] comes closest.
- **G3 – Reactive detection.** Most systems decide whether an activity is malicious; only [3] and [15] also predict how an attack will progress.
- **G4 – Drift and uncertainty.** Only [8] models uncertainty explicitly, and it is preliminary; changing attacker behaviour and concept drift are mostly left open.
- **G5 – Validation.** Several studies use synthetic or derived data [3, 7, 11], and few evaluate in realistic environments.

V. THE ENVISIONED CONCEPTUAL FRAMEWORK

To address G1–G5 we sketch the pipeline in Fig. 2. It is a proposal: no implementation or experimental result is presented in this paper.

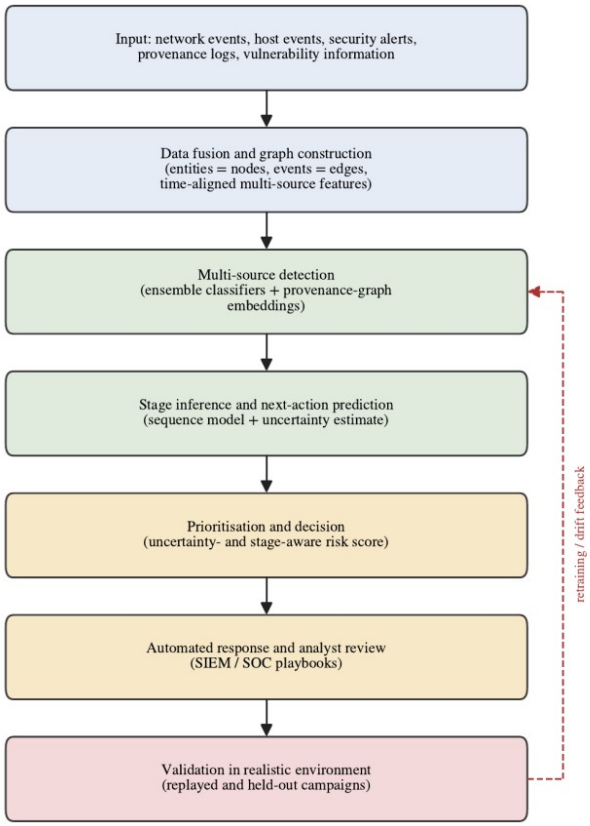


Fig. 2. Proposed multi-stage detection and automated defence pipeline.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

A. Data Collection and Graph Construction

The input combines network events, host events, security alerts, provenance logs and vulnerability information. Public datasets used in the reviewed work, such as the UNSW-NB15-derived dataset [7], DAPT-2020 [11], CADETS and THEIA [4, 12] and StreamSpot [13], can be complemented by in-house data where available. Events are time-aligned, and host activity is turned into provenance graphs as described in Section II-C. Complete campaigns are set aside as a held-out group that is never used for training.

B. Multi-Source Detection

Ensemble classifiers such as Random Forest and XGBoost [5] process alert and network features, while provenance-graph embeddings [4, 13] describe host behaviour. Their outputs are fused so that evidence from several sources contributes to one decision, which targets G1 and G2.

C. Stage Inference and Next-Action Prediction

A sequence model, for example a CNN-RNN with HMM decoding [3], infers the current stage and the distribution over likely next stages from the alert sequence, which targets G3. An evidential output [8] provides an uncertainty estimate, and inputs with high uncertainty are flagged as possible new tactics, which targets G4.

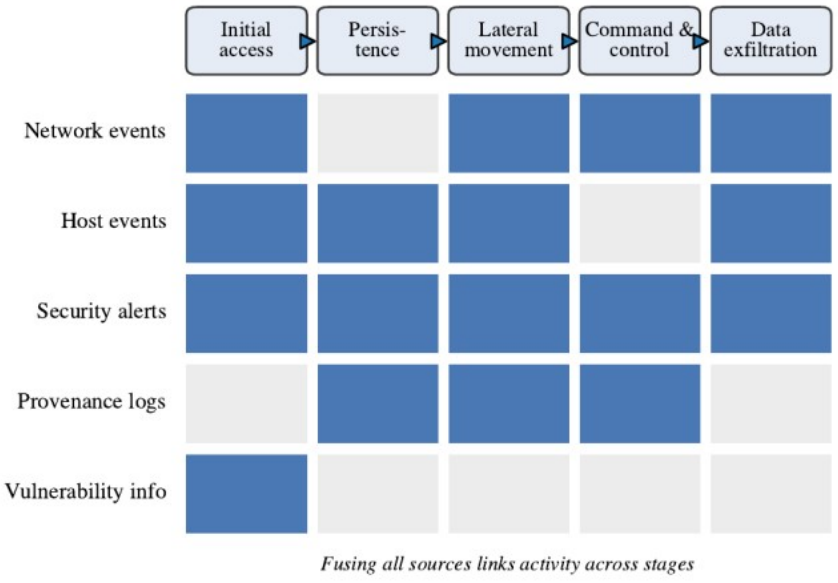


Fig. 3. Conceptual illustration of APT lifecycle stages and data sources; filled cells mark sources that may give evidence at a stage (schematic, not experimental data).

D. Prioritisation, Decision and Automated Response

Detection, stage and uncertainty are combined into a risk score that prioritises alerts for the SOC. Low-uncertainty, high-risk cases can trigger SIEM response playbooks [5, 6], whereas high-uncertainty cases are passed to an analyst. Candidate techniques for each component are listed in Table II.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

TABLE II

PIPELINE COMPONENTS AND CANDIDATE TECHNIQUES

Component	Candidate technique	Basis
Alert detection	Random Forest, XGBoost with SIEM event-detection modules	[5]
Host detection	GraphSAGE embeddings, iForest, clustering	[4, 12, 13]
Stage prediction	CNN-RNN with HMM/Viterbi; CNN-LSTM	[3, 7]
Uncertainty	Evidential deep learning (Dirichlet output)	[8]
Response	SIEM-driven post-alert decisions	[5, 6]

E. Validation and Feedback

The framework should be tested in a realistic environment, for example by replaying real campaigns, rather than only on synthetic data, which targets G5. Analyst-confirmed outcomes are fed back as new labels, and the models are updated periodically to follow concept drift [12].

TABLE III

PROPOSED EVALUATION METRICS

Component	Metrics
Detection	Precision, recall, F1/macro-F1, false-positive rate; node-level and graph-level
Stage prediction	Stage accuracy; next-stage accuracy
Uncertainty	Calibration error; out-of-distribution detection rate
Generalisation	Performance drop on held-out campaigns
Efficiency	Throughput (logs/s), training time, memory
Operations	Alert reduction; time to detect and respond; analyst workload

F. Evaluation Protocol

The framework should be compared against appropriate baselines, such as a single-source, single-stage classifier on alerts and a provenance-only detector. Because a model can look strong on a familiar dataset yet fail on a new campaign, evaluation should include held-out campaigns, and results should be reported on more than one dataset. The measures are listed in Table III.

VI. APPLICATIONS

When realised, the approach can support security operations centres (SOCs) in alert triage, prioritisation and decision-making, reducing manual workload and response time. It can be used for threat hunting and investigation on hosts using provenance graphs, for monitoring enterprise and IoT environments, and for early warning, since predicting the next attacker action leaves time to respond. Because the output includes stage and uncertainty, analysts can see why an alert was raised and where human review is needed.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

VII. LIMITATIONS AND CHALLENGES

The review covers only fifteen primary studies, so conclusions about the literature as a whole should be drawn with care, and the ratings in Fig. 1 are qualitative. Reported results come from different datasets and metrics and cannot be compared directly. The proposed framework has not been implemented, so its benefit is not demonstrated. Significant challenges include the availability of realistic multi-stage datasets, class imbalance, the cost of training and of building provenance graphs at scale, false positives, concept drift and distribution shift between training and deployment, and the risk of automated responses acting on wrong predictions. Any predicted detection still requires validation in a realistic environment.

VIII. FUTURE DIRECTIONS

Natural extensions include building datasets that cover the complete APT lifecycle, fusing network, host, log and vulnerability information, and predicting attacker intent and next actions rather than only detecting malicious activity. Uncertainty-aware and drift-aware models, with online updating, should be explored so that new campaigns are handled more reliably. Tighter coupling of detection with automated response, with analysts in the loop, is another direction. Implementing and benchmarking the proposed framework on public APT datasets and in realistic environments is the natural next step.

IX. CONCLUSION

Recent work applies ML, behavioural analysis, provenance-graph learning and SIEM to APT detection and reports strong results, but no reviewed study combines full-lifecycle coverage, multi-source fusion, next-action prediction, uncertainty handling and realistic validation. We proposed a conceptual framework and an evaluation protocol that bring these elements together. Such an approach may move defence from raising alerts towards AI-assisted decision-making and automated response, provided it is confirmed by experiments in realistic environments.

REFERENCES

[1] A. Alageel and S. Maffeis, “Investigation of advanced persistent threats network-based tactics, techniques and procedures,” *Comput., Mater. Continua*, 2026.

[2] S. Yuldoshkhujaev, M. Jeon, D. Kim, N. Nikiforakis, and H. Koo, “A decade-long landscape of advanced persistent threats: Longitudinal analysis and global trends,” 2026.

[3] S. I. Tijjani, B. Ghita, N. Clarke, and M. Craven, “Deep recurrent hidden Markov learning framework for multi-stage advanced persistent threat prediction,” arXiv:2601.06734, 2026.

[4] M. Lv, H. Liu, S. Zhang, T. Chen, and T. Zhu, “APT-MCL: An adaptive APT detection system based on multi-view collaborative provenance graph learning,” arXiv:2601.08328, 2026.

[5] G. Ali, S. Shah, and M. ElAffendi, “Enhancing cybersecurity incident response: AI-driven optimization for strengthened advanced persistent threat detection,” 2025.

[6] C. Gilbert, M. A. Gilbert, and M. Dorgbefu Jnr, “Detection and response strategies for advanced persistent threats (APTs),” 2025.

[7] S. I. Tijjani, B. Ghita, N. Clarke, and M. Craven, “Hybrid ensemble and deep learning architectures for advanced persistent threat detection,” in *Proc. IEEE IEMCON*, 2025, pp. 34–40.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

[8] A. Gaudenzi, L. Nodari, L. Kaplan, A. Russo, M. Sensoy, and F. Cerutti, "Preliminary investigation into uncertainty-aware attack stage classification," arXiv:2508.00368, 2025.

[9] N. I. Che Mat, N. Jamil, Y. Yusoff, and M. L. Mat Kiah, "A systematic literature review on advanced persistent threat behaviors and its detection strategy," 2024.

[10] S. Krishnapriya and S. Singh, "A comprehensive survey on advanced persistent threat (APT) detection techniques," 2024.

[11] A. A. Mamun, H. Al-Sahaf, I. Welch, and S. Camtepe, "Advanced persistent threat detection using genetic programming," *Appl. Soft Comput.*, vol. 166, Nov. 2024.

[12] Y. Amaru, P. Wudali, Y. Elovici, and A. Shabtai, "RAPID: Robust APT detection and investigation using context-aware deep learning," arXiv:2406.05362, 2024.

[13] B. Xu, Y. Gong, X. Geng *et al.*, "ProcSAGE: An efficient host threat detection method based on graph representation learning," *Cybersecurity*, vol. 7, Art. no. 51, 2024.

[14] Z. Chen *et al.*, "Machine learning-enabled IoT security: Open issues and challenges under advanced persistent threats," *ACM Comput. Surv.*, vol. 55, no. 5, Art. no. 105, 2022.

[15] I. Ghafir *et al.*, "Detection of advanced persistent threat using machine-learning correlation analysis," *Future Gener. Comput. Syst.*, vol. 89, pp. 349–359, 2018.

[16] W. L. Hamilton, R. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 30, 2017.

[17] M. Sensoy, L. Kaplan, and M. Kandemir, "Evidential deep learning to quantify classification uncertainty," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 31, 2018.

[18] A. J. Viterbi, "Error bounds for convolutional codes and an asymptotically optimum decoding algorithm," *IEEE Trans. Inf. Theory*, vol. 13, no. 2, pp. 260–269, Apr. 1967.