ijsigma.org

SIGMA International Journal of Security, Intelligence, Governance & Modern AI Computing

VOLUME 01 · ISSUE 01 · 2026 · ARTICLE 0004

RESEARCH ARTICLE OPEN ACCESS · ISSN: [ISSN to be assigned] · DOI: [DOI to be assigned]

Intelligent Adaptive Threat Modeling for Advanced Persistent Threat Detection in Critical IoT Networks

Salman Ghani Virk1,*, Jawaid Iqbal1 and Tariq Hanif2

1 Riphah International University, Islamabad

2 UIIT PMAS Arid Agriculture University, Rawalpindi Pakistan

Corresponding author: Salman Ghani Virk (salmanvirk70@gmail.com) · ORCID: [0000-0000-0000-0000]

Author emails: Jawaid Iqbal (Jawaid.iqbal@riphah.edu.pk); Tariq Hanif (chtareeq@gmail.com)

ABSTRACT

The widespread implementation of IoT and IIoT in critical systems has extended the attack surface against which Advanced Persistent Threats (APTs), characterized by a multi-step and persistent nature, can operate. Current intrusion detection models are ill-equipped to capture APT behavior because they are event-centric and focus on detecting events that indicate intrusion. Current IoT security models classify detected intrusions as benign or malicious. Still, they cannot handle campaign progression, temporal persistence, varying threat levels, critical asset exposure, uncertainty, or prioritization of response actions. In this paper, we present an Intelligent Adaptive Threat Modeling (IATM) approach that views intrusion detection as an evidence-generation process and adds an adaptive threat-state modeling layer on top of it. We select CICAPT-IIoT2024 as the primary benchmark because it provides network traffic, provenance logs, attack timestamps, and multi-stage APT evidence from an IIoT testbed. We specify a leakage-safe evaluation protocol for event-level detection, threat-state estimation, static-versus-adaptive assessment, early warning, ablation, robustness, statistical validation, and computational efficiency. Published CICAPT-IIoT results are retained as verified external baselines; numerical IATM values in the present manuscript remain explicitly illustrative and must be replaced by held-out experimental outputs before journal submission. The contribution is therefore an evidence-to-decision framework for continuously updated, interpretable, asset-aware operational threat intelligence rather than another standalone attack classifier.

Keywords: Advanced Persistent Threat; Critical IoT; IIoT Security; Adaptive Threat Modeling; Dynamic Risk Assessment; APT-Stage Intelligence; Intrusion Detection; Operational Decision Support.

Highlights

- **The widespread implementation of IoT and IIoT in critical systems has extended…**
- **Current intrusion detection models are ill-equipped to capture APT behavior…**
- **Current IoT security models classify detected intrusions as benign or malicious.**

IMPACT STATEMENT

The widespread implementation of IoT and IIoT in critical systems has extended the attack surface against which Advanced Persistent Threats (APTs), characterized by a multi-step and persistent nature, can operate. Current intrusion detection models are ill-equipped to capture APT behavior because they are event-centric and focus on detecting events that indicate intrusion. Current IoT security models classify detected intrusions as benign or malicious. Still, they cannot handle campaign progression, temporal persistence, varying threat levels, critical asset exposure, uncertainty, or prioritization of response actions. In this paper, we present an Intelligent Adaptive Threat Modeling (IATM) approach that views intrusion detection as an evidence-generation process and adds an adaptive threat-state modeling layer on top of it. We select CICAPT-IIoT2024 as the primary benchmark because it provides network traffic, provenance logs, attack timestamps, and multi-stage APT evidence from an IIoT testbed. We specify a leakage-safe evaluation protocol for event-level detection, threat-state estimation, static-versus-adaptive…

1 I. INTRODUCTION

Critical IoT/IIoT environments require security mechanisms that can interpret not only isolated malicious events but also the progression, persistence, and operational consequences of multi-stage Advanced Persistent Threat (APT) activity. Accordingly, this study positions adaptive threat-state modeling as a bridge between event-level intrusion detection and risk-aware operational decision support.

A. Background and Motivation

The Internet of Things (IoT) and Industrial Internet of Things (IIoT) increasingly support cyber-physical operations in manufacturing, energy, transportation, healthcare, smart infrastructure, and industrial automation. By integrating sensing, connectivity, edge/cloud analytics, and automated control, they improve visibility and decision-making but also expand the attack surface through heterogeneous devices, constrained endpoints, diverse protocols, and continuously changing operational conditions. Recent cybersecurity research highlights the growing role of artificial intelligence in cyber defense and threat analysis [1]. However, high event-level detection performance does not by itself establish the operational state of a developing cyber campaign. Effective IoT defense also requires understanding how malicious activity is progressing, how persistent it is, which assets are exposed, whether risk is escalating, and when to prioritize a response.

AI has emerged as an enabler for cybersecurity analytics because learned models can uncover sophisticated, dynamic behavior that static signatures alone cannot capture. Recently, research has highlighted the use of machine learning and deep learning in cybersecurity to show how to build adaptive analytical systems [2]. However, using AI for cybersecurity creates new challenges: data leakage must be controlled, probabilities must be calibrated, the minority class must be assessed separately, and model decisions must be made transparently based on the evidence used.

B. Advanced Persistent Threats in Critical IoT

Advanced Persistent Threats are sophisticated, covert, and persistent attacks that can unfold through multiple methods and tactics over an extended period. In such an attack, a single observation is only one part of a larger pattern involving persistence, discovery, credential access,

lateral movement, command and control, collection, and exfiltration. This is an important criterion for Internet of Things (IoT) research, where temporal order and context must be preserved.

Recent studies show that temporality and collaboration context help detect APTs. Provenance connections, sequence knowledge, graph structure, and collaborative protection may help reconstruct persistent attacks. However, identifying a particular phase does not equate to evaluating the attack's operational importance. The security analyst must understand whether the gathered data imply escalation, whether the attacked asset is operationally important, what uncertainty exists in the estimated system state, and whether the evidence is sufficient to act early.

C. Research Problem and Gap

Most prior IDS studies in IoT focus on classification-based cybersecurity monitoring, whereas existing APT studies increasingly emphasize provenance graphs, stage detection, and attack sequence reconstruction. Several alternative viewpoints, such as dynamic risk assessment, threat intelligence, ATT&CK-based analysis, vulnerability severity analysis, and reasoning at campaign stages, may be considered, although they are studied independently. The current literature review contains no sufficient empirical studies on transforming the APT sequence into a temporal adaptive threat state for IIoT environments.

Consequently, this research addresses not only how malicious traffic can be categorized, but also how reliable information from individual events can be used to build a developing profile of campaign-level operational risk.

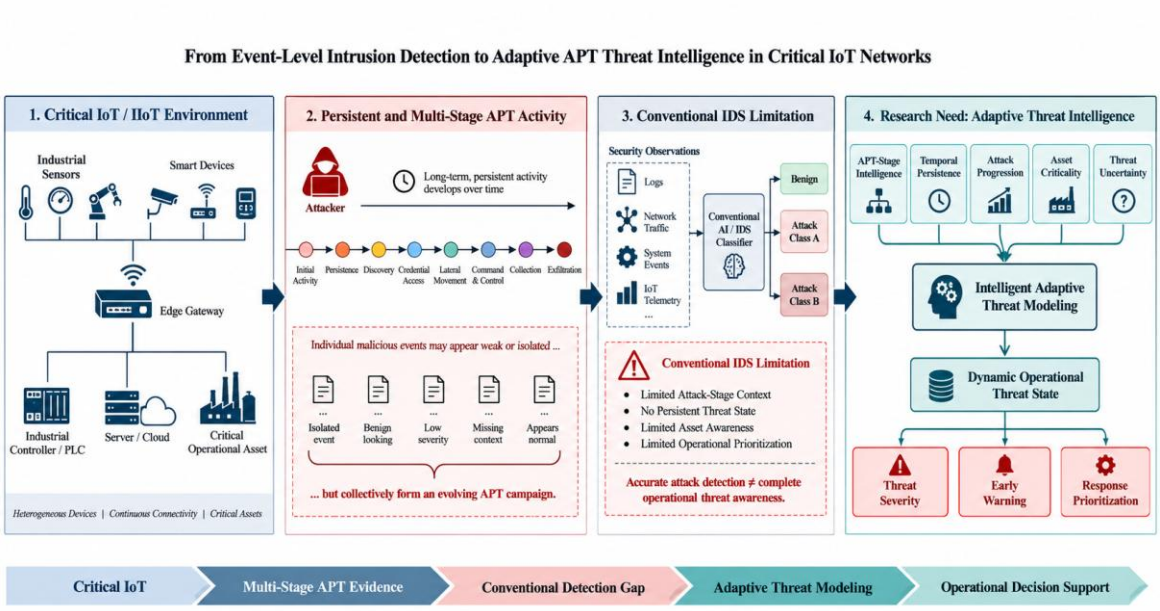


Figure 1. From event-based intrusion detection to adaptive APT threat intelligence in critical IoT networks.

As shown in Figure 1, the problem statement starts with a heterogeneous critical IoT/IIoT environment and multi-stage persistent APT activity. It ends with the drawbacks of traditional event-based intrusion detection. The figure reflects the need to integrate information on APT stages, persistence, progression, critical assets, and uncertainty into the threat's operational state.

D. Research Questions, Objectives, and Contributions

To address this gap, this paper proposes an Intelligent Adaptive Threat Modeling (IATM) approach. We view the intrusion detection mechanism as a layer that provides evidence. In contrast, the main innovation is the Adaptive Threat-State Model (ATSM), which dynamically constructs the operational threat-state estimate by recursively combining current APT-stage evidence with previously held beliefs about the threat state. This approach ensures the research is not just another classifier accuracy study.

Four research questions guide the study design: RQ1 asks whether temporal state updating better estimates APT stage and transitions than independent event classification. RQ2: Is adaptive threat scoring able to offer better warning performance compared to static risk scoring, given a bounded number of false alarms? RQ3: What are the effects of persistence, asset criticality, escalation, uncertainty, and temporal adaptation on individual aspects of classification and operations? RQ4: How well does IATM perform across temporal-window changes, incomplete evidence, multiple random seed generations, and additional benchmarks when applicable? These objectives are to create a leakage-resistant evidence pipeline, estimate the changing probabilistic threat state, measure the asset-dependent risk, assess early-warning properties, and verify claims through controlled ablation and testing.

Contributions include: (i) APT threat state definition, which uses sequence IIoT evidence to represent APT campaigns adaptively; (ii) Temporal model for stage probability, persistence, transition evidence, and uncertainty; (iii) Dynamic scoring model, which utilizes stage impact, criticality, escalation, and certainty; (iv) Early warning and operation prioritization module that translates threat state into decisions for analysts; and (v) Leakage-safe evaluation protocol including strong baselines, static vs adaptive comparison, ablation, robustness analysis, statistical evaluation, and computational efficiency analysis.

E. Organization of the Paper

The rest of the paper is structured as follows. Section II reviews the literature concerning intrusion detection in IoT/IIoT, APT sequence modeling, dynamic cyber risk assessment, and threat intelligence. Section III describes the proposed approach and its mathematical formulation. Section IV introduces the dataset and the experimental setting. Section V presents benchmarking and exemplary IATM results, along with comparative, ablation, and error analyses. Section VI addresses the significance of the research findings, while Section VII indicates limitations and threats to validity.

2 II. LITERATURE REVIEW

Four interrelated research themes relevant to the current study are: intrusion detection systems for IoT/IIoT based on artificial intelligence, sequential analysis of APT attacks, dynamic assessment of cyber risks, and prioritization with the help of threat intelligence. The literature review presented below provides a critical synthesis of these research directions.

A. AI-Based Intrusion Detection in IoT and IIoT

Learning and deep learning have emerged as key frameworks for anomaly and intrusion detection in IoT domains. A recent review provides an overview of machine-learning-based IoT intrusion-detection techniques [3]. More broadly, the literature recognizes the importance of dataset construction, class balancing, and an appropriate evaluation context. A literature review presents a framework to map learning approaches across IoT intrusion detection systems [4]. Another review examines deep-learning-based techniques for cyber threat detection in IoT and the growing significance of representation learning in IoT security [5]. These works help to set up the methodological context [6] for the event-level detection layer used here [7].

Recent empirical studies show that research in this area is moving beyond traditional classifiers. Deep learning has been used directly for Industrial IoT intrusion detection and offers a good neural baseline [8]. Attention-based bidirectional sparse autoencoding has also been used to develop better attack representations [9]. More recently, research has explored temporal transformer classification with flow-to-window aggregation for IoT traffic analysis [10]. Such methods leverage context learned from data [11]; however, their primary output remains attack classification.

Dataset creation is equally important. Modern benchmark datasets differ significantly in device diversity, attack coverage, telemetry sources, and time structure. [13] Some datasets are tailored

for general binary or multiclass attack categorization, while others maintain higher levels of provenance or longitudinal campaign evidence. [14] This difference becomes apparent because a dataset that is good for attack categorization does not have to be good for campaign staging assessment or dynamic risk analysis. [15]

B. APT Detection and Sequential Threat Analysis

APT detection research increasingly exploits provenance information because provenance graphs encode causal and temporal relations among processes, files, network sockets, and other system entities. Researchers have investigated a federated intrusion-detection approach for SOC-centered APT defense [17]. Researchers have also used sequential graph learning to model APT progression and stage transitions [18]. Masked graph representation learning provides a self-supervised mechanism for provenance-based campaign detection [19]. These studies demonstrate that campaign context can provide information that isolated event classification does not preserve. [20] [21]

Researchers have also combined context-aware provenance tracing with deep learning to support APT detection and investigation [22]. Dynamic provenance-graph detection explicitly tracks persistent malicious behavior over time [23], while attack-intent-driven sequence learning models multi-stage APT activity [24]. Collectively, this work shows that APT-stage progression is already a substantive research area. [25] Therefore, the novelty of the present study cannot rest on stage recognition alone; [26] the unresolved opportunity is to use stage evidence as an input to an adaptive operational threat state [27] that also incorporates persistence, asset context, escalation, and uncertainty. [28]

C. Dynamic Cyber Risk and Threat Modeling

Dynamic risk assessment (DRA) attempts to overcome the limitations of periodic or static assessment by updating risk when system or threat conditions change. A systematic review identifies a broad shift toward continuous and near-real-time cybersecurity risk assessment [29]. This establishes the need for risk estimates that evolve with new evidence rather than remaining fixed. Dynamic attack graphs update attack paths and compromise risk when IoT topology changes [30]. Scenario-driven simulation has likewise been applied to continuous risk analysis in Medical IoT networks [31]. These approaches establish the importance of adaptation [32]. Still, their risk updates rely primarily on topology, vulnerabilities, scenario assumptions, or simulated events rather than an empirically observed multi-stage IIoT APT campaign. [33]

D. Threat Intelligence and Quantitative Prioritization

Cyber Threat Intelligence (CTI) improves situational awareness by structuring adversary behavior, indicators, vulnerabilities, and tactics. CTIoT provides an IoT-focused mechanism for collecting and processing threat intelligence [34]. An automated platform for gathering and managing open-source cyber threat intelligence provides a complementary operational approach [35]. These studies demonstrate the value of contextual intelligence, but they do not, by themselves, provide a continuously updated, campaign-specific threat state. Quantitative prioritization [36] and knowledge-base standards [37] are therefore considered later in the methodology [38] and discussion [39], where they directly support the operational scoring design. [40]

E. Critical Synthesis of the Literature

Table 1. Representative literature and the unresolved gap.

Study / Area	Dataset or Setting	Primary Output	Strength	Limitation Relative to IATM
Ghiasvand et al.	IIoT APT	Network +	Real multi-stage	Dataset foundation; no

Study / Area	Dataset or Setting	Primary Output	Strength	Limitation Relative to IATM
(CICAPT-IIoT)	campaign	provenance evidence	APT evidence	adaptive operational state
Neto et al. (CICIoT2023)	105 IoT devices; 33 attacks	Attack classification	Broad modern IoT benchmark	Not a longitudinal APT campaign
Ferrag et al. (Edge-IIoTset)	IoT/IIoT testbed	Attack classification	Heterogeneous IIoT environment	Attack-category rather than campaign-state focus
V. P. M. K. et al. (Sequential graph APT)	CICAPT-IIoT + others	APT stage/transition detection	Explicit temporal stage modeling	Operational risk and asset context not central
Ren & Geng (Masked graph APT)	Provenance graphs	Malicious node detection	Self-supervised graph learning	Detection rather than operational threat state
Amaru et al. (RAPID)	Provenance data	Detection + investigation context	Context-aware alert tracing	No explicit asset-aware dynamic risk state
Salayma (Dynamic attack graphs)	Dynamic IoT topology	Updated attack paths/risk	Adapts to topology change	Not driven by sequential APT telemetry
Czekster et al. (Dynamic MIoT risk)	Medical IoT simulation	Continuous risk assessment	Decision-oriented dynamic assessment	Simulation-centric
Yun & Min (ATT&CK risk scoring)	Edge-IIoT	Quantitative threat ranking	Interpretable prioritization	Attack-type scoring; limited campaign-state evolution

As shown in Table 1, prior studies address important parts of the problem separately, supporting an integrated adaptive APT threat-modeling framework.

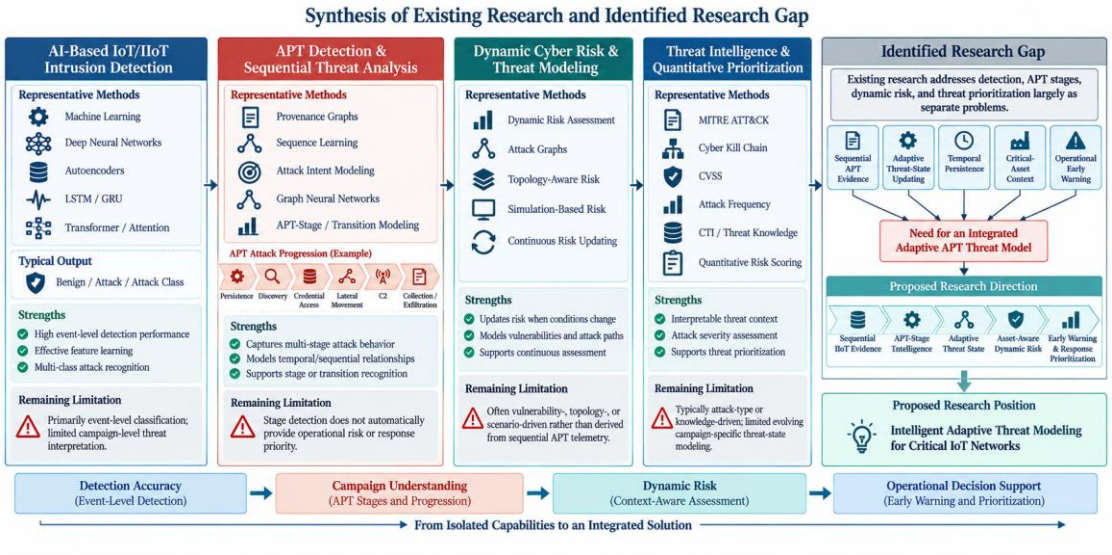


Figure 2. Synthesis of existing research streams and positioning of the identified research gap.

As shown in Figure 2, the reviewed literature has developed along four strong but largely separate streams: AI-based IoT/IIoT intrusion detection, sequential APT analysis, dynamic cyber-risk modeling, and threat-intelligence-driven prioritization. Their convergence exposes the central research gap IATM addresses: the lack of an integrated mechanism that transforms sequential APT evidence into an adaptive, asset-aware threat state for operational decision support.

F. Identified Research Gap

The literature indicates a remaining need for an empirically evaluated framework that transforms sequential IIoT APT evidence into a time-indexed adaptive threat state. Such a framework should combine current stage evidence with historical state, temporal persistence, transition behavior, explicitly defined asset criticality, escalation, and uncertainty; continuously update threat severity as new evidence arrives; and convert the evolving state into interpretable early-warning and response-priority information. This requirement motivates the proposed IATM framework and differentiates it from conventional event classification, standalone APT-stage detection, static threat scoring, and topology-driven dynamic risk assessment.

3 III. METHODOLOGY

The proposed methodology is designed as an end-to-end evidence-to-decision pipeline that preserves temporal order, prevents data leakage, estimates evolving APT states, and converts contextualized threat evidence into operational risk and early-warning information. The following subsections define the framework, mathematical formulation, adaptive state-update mechanism, and evaluation-oriented components in a reproducible manner.

A. Framework Overview

The proposed IATM framework operates as an evidence-to-decision pipeline.

The framework subjects observations of temporally ordered networks and provenance information to leak-safe transformations, analysis via a threat-evidence detector, mapping of APT-stage probabilities, and recursive incorporation into an Adaptive Threat-State Model. The resulting threat state, together with persistence, stage severity, asset criticality, escalation, and uncertainty, is used to develop a dynamic operational threat score and an early-warning decision. The threat-evidence detector is explicitly considered a support component in this respect; the key methodological innovation lies in the adaptive threat-state and decision components.

The full concept process flow is as follows: IIoT Evidence -> Leakage-Resilient Preprocessing -> Threat Evidence Recognition -> APT Phase Intelligence -> Dynamic State Update -> Threat Contextual Factors -> Dynamic Threat Rating -> Early Warning & Decision Making Support.

B. Data Representation and Leakage-Safe Preprocessing

Let the dataset with respect to time ordering be D = {(x_i, t_i, y_i)}, where x_i is the vector of the observation at time t_i, and y_i is the label indicating whether the sample belongs to a normal or an attack stage. The observations will then be divided into temporal windows W_t = {x_i | t-Delta < t_i <= t}, while the function phi( ) builds a feature vector X_t = phi(W_t).

$$W_t = \{ x_i \mid t - \Delta < t_i \leq t \}, \quad X_t = \varphi(W_t) \quad (1)$$

We divide the ordered observations into mutually exclusive training, validation, and hold-out test sets. Whenever possible, we use chronological or group-based partitioning rather than random partitioning. We fit scaling, categorical conversion, feature extraction, dimensionality reduction, probability calibration, and imbalance-handling techniques only on the training set and apply them unchanged to the validation and hold-out test sets. Oversampling is limited to the training set and is never used on the hold-out test set.

$$x\sim_ij = (x_ij - \mu_j\text{^}train) / (\sigma_j\text{^}train + \varepsilon) \quad (2)$$

C. Threat Evidence and APT-Stage Intelligence

Each X_t corresponds to a set of probabilities that the supporting detector generates, i.e., p_t = [p_t^(0),...,p_t^(K)]. For temporal feature representation learning, attention-based models like the Transformer or recurrent sequences like LSTM can be leveraged. Classical approaches remain relevant for tabular traffic data, especially Random Forest and gradient boosting/XGBoost as benchmarks. The amount of malicious evidence q_t can be calculated as one minus the benign probability. Another vector, the stage-evidence vector e_t, reflects the probability of the APT stages used in the experiments. This evidence is not an operational decision; it is an observation.

$$q_t = 1 - p_t\text{^}(benign), \quad e_t\text{^}(k) = P(S_t = s_k \mid X_t) \quad (3)$$

D. Adaptive Temporal Threat-State Model

Let πt be the posterior probability distribution for the benign state and K APT threat states at time t. Let the transition probabilities between state sj and state sk be captured by a transition matrix A, which can be learned solely from transitions in the training sequence by using Laplace smoothing, where α > 0 ensures no zero-probability transitions. Before observing the most recent evidence, the previous posterior is passed through the transition matrix A to predict the new state.

$$\bar{\pi}_t = A\text{^}T \pi_(t\text{-}1) \quad (4)$$

$$\pi_t\text{^}(k) = [\, e_t\text{^}(k)\, \bar{\pi}_t\text{^}(k) \,] / [\, \Sigma_j\, e_t\text{^}(j)\, \bar{\pi}_t\text{^}(j) \,] \quad (5)$$

The resulting threat state is Sˆ_t = argmax_k p (k). In this relation, the estimate is based not only on the current state but also on historical evidence. As a result, random noise in the prediction process will be diminished.

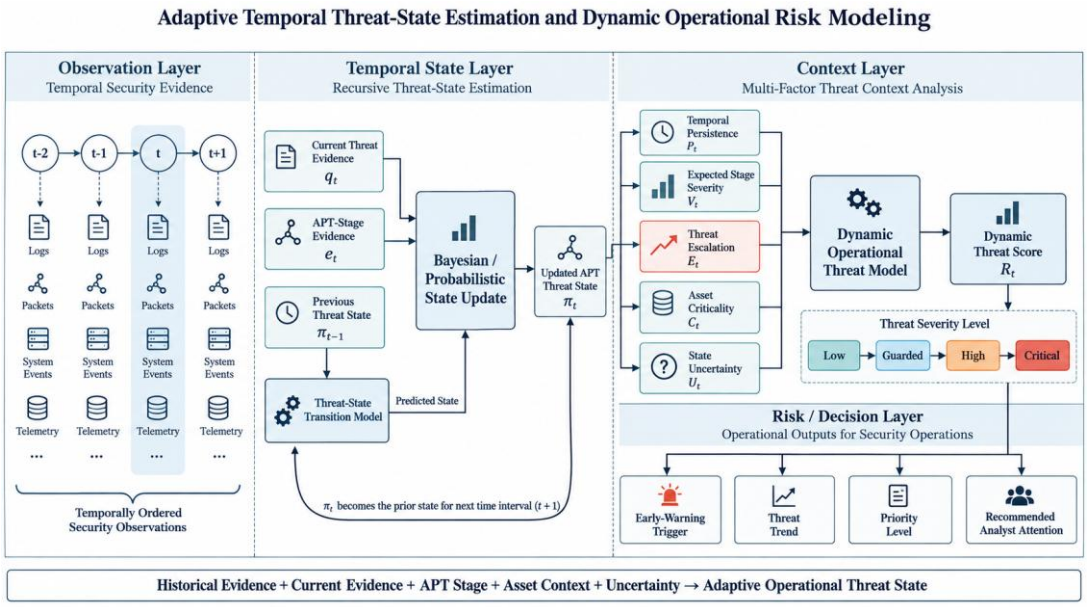


Figure 3. Threat-state estimation and dynamic operational risk modeling over time.

As shown in Figure 3, the posterior threat state from the prior cycle is propagated using the transition model, then integrated with the current threat and APT stage information to produce the new state at time t. The new state determines persistence, stage severity, escalation, asset criticality, and uncertainty.

E. Persistence, Severity, Asset Context, Escalation, and Uncertainty

Persistence of temporal persistence P_t is achieved through the use of an exponential-memory function in retaining accumulated evidence of malicious activities. The expected stage severity V_t is derived from the posterior probability distribution of the normalized stage-severity factors. The criticality of assets C_t is the measure of the importance of the assets in operation. It is defined as a contextual/experimental factor, not a native CICAPT-IIoT factor. The escalation factor E_t refers to any increase in expected stage severity.

$$P_t = \lambda P_{(t-1)} + (1 - \lambda) q_t \quad (6)$$

$$V_t = \Sigma_k \pi_t^{(k)} \omega_k \quad (7)$$

$$E_t = \max(0, V_t - V_{(t-1)}) \quad (8)$$

$$U_t = - [\Sigma_k \pi_t^{(k)} \log(\pi_t^{(k)}+\varepsilon)] / \log(K+1) \quad (9)$$

Asset criticality is defined as $C_t = \Sigma_m a_m c_m$, where c_m are normalized measures of asset availability, integrity, confidentiality, safety, or operational importance, and $a_m \geq 0$ such that $\Sigma_m a_m = 1$. Because criticality is not an innate label in CICAPT-IIoT, the mapping is an explicit contextual input, not the ground truth. The mappings and weights must be specified, sensitivity-analyzed, and excluded from stage-classification accuracy.

F. Dynamic Threat Score and Early-Warning Logic

The adaptive threat score for operational purposes, R_t, considers recent malice, anticipated stage severity, temporal endurance, asset importance, escalation, and uncertainty. Stage malice may be mapped to MITRE ATT&CK tactics and techniques, vulnerability severity to CVSS concepts when provided, and campaign evolution within the Cyber Kill Chain framework. To avoid over-parametrization of scores, we set non-negative weights $\beta_1 \ldots \beta_6$ based only on validation data via

regularized calibration against the predetermined operational states; β_6 serves as the uncertainty term.

$$R_t = 100\,\sigma(\beta_0 + \beta_1\, q_t + \beta_2\, V_t + \beta_3\, P_t + \beta_4\, C_t + \beta_5\, E_t - \beta_6\, U_t) \quad (10)$$

The continuous score is calibrated with operational state labels that include Low, Guarded, High, and Critical through thresholds derived from validation results. We determine the early-warning threshold τ_EW from validation data within a defined false-alert budget, and τ_U restricts warnings from states with high posterior uncertainty. We measure warning lead time relative to the first occurrence of a critical campaign state. Stability in threat-state estimates is given by one minus the average total variation difference of the posterior distribution during periods when the reference state remains constant.

$$Alarm_t = 1 \text{ if } R_t \geq \tau_EW \text{ and } U_t \leq \tau_U; \text{ otherwise } 0 \quad (11)$$

$$Lead\ Time = t_critical - t_warning \quad (12)$$

G. Training, Validation, and Scientific Evaluation

Model design is divided into three stages: training stage, where models are tuned on data for preprocessing steps, detector parameters, transition modeling, and other learned components; validation stage, which determines temporal window length, probabilistic scaling factors, score multipliers, alarm thresholds, and all other hyperparameters; and the test stage, which is performed on a held-out test set only after freezing of the entire pipeline. We run stochastic procedures with 10 separate random seeds. Pairwise comparisons rely on the Wilcoxon signed-rank test, adjusted by Holm for multiple comparisons, and effect-size calculation using something like Cliff's delta. Sensitivity analysis includes both temporal-window and missing-evidence perturbations; we cross-evaluate additional datasets only on metrics with the same semantics.

Algorithm 1 – IATM adaptive threat state update. Input: ordered in time evidence Xt, calibrated probabilities of stages et, previous posterior πt-1, transition matrix A, and asset situation Ct. For each window t: (1) obtain calibrated malicious and stage evidence; (2) propagate πt-1 through A; (3) update and normalize πt with respect to et; (4) calculate persistence Pt, severity expectation Vt, escalation Et, criticality Ct, and uncertainty Ut; (5) obtain Rt using frozen score coefficients; (6) map Rt to the severity state; (7) give an early warning only if both score and uncertainty conditions are met; and (8) keep πt as a prior for the next window. Output: posterior threat state, dynamic threat score, uncertainty, warning status, and information for the analyst.

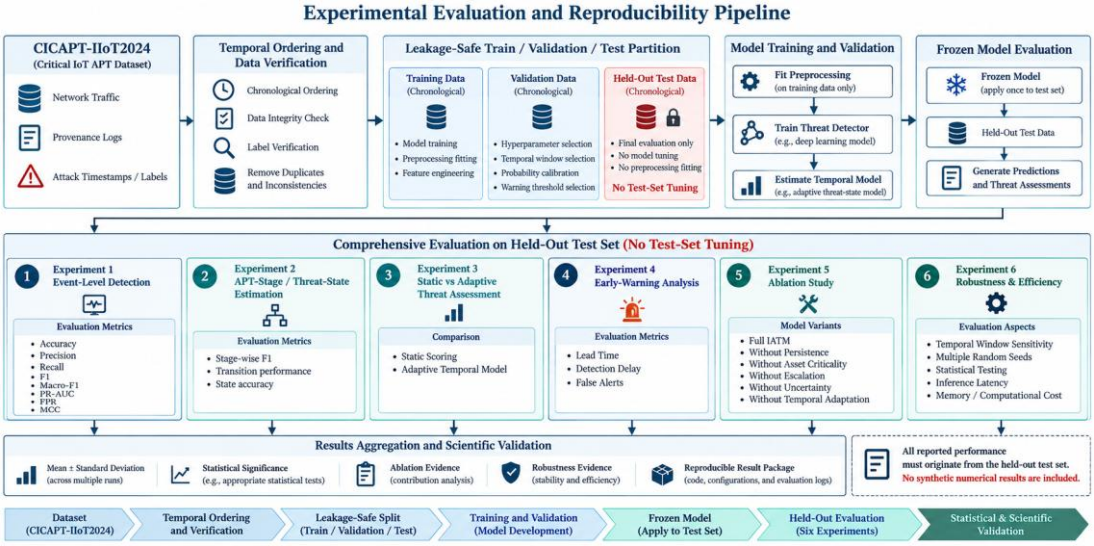


Figure 4. Experimental evaluation and reproducibility process for the proposed IATM system.

As illustrated in Figure 4, the experimental process maintains the sequence order while keeping the training, validation, and testing data sets separate to avoid data leakage from the test set. Afterward, the frozen model is evaluated for event detection, threat-state estimation, static-versus-adaptive comparison, early-warning performance, ablation study, robustness, hypothesis testing, and computational efficiency.

4 IV. DATASET AND EXPERIMENTAL SETUP

CICAPT-IIoT2024 is chosen as the main dataset because it directly applies to the campaign-level RQ. This dataset was collected in a hybrid real and simulated IIoT testbed based on Brown-IIoTbed, using software such as NS-3, Ubuntu, and Kali virtual machines, Raspberry Pi devices, IoT sensors, OpenPLC, and the Modbus protocol. CICAPT-IIoT2024 provides two channels of evidential data: provenance CSV logs and network traffic PCAP/derived CSV, plus an Attack Information file that includes time, process ID, and attack category. The provenance representation includes 32 descriptive features for processes, artifacts, and edges, and more than 20 attack techniques across eight APT tactics. CICAPT-IIoT2024 uses two data-collection phases: Phase 1 includes about 96 hours of benign behavior data, while Phase 2 includes about 72 hours of both benign behavior and APT campaign data. According to the released distributions, Phase 1 has 46,773 benign provenance nodes and 12,103,705 benign packets; Phase 2 has 52,954 benign provenance nodes, 330 malicious provenance nodes, 9,535,819 benign packets, and 1,004 attack packets. Because approximately 99.5% of observations are benign at the source, macro-averaged and minority-stage metrics are necessary.

In the final experiment, properly align provenance events and network observations by timestamp and attack, ensuring that no future attacks are included in earlier windows. Perform duplication detection and near-duplication detection before partitioning. If possible, include complete attacks or temporally adjacent attack segments in the same partition. Export the number of events retained, the distribution after preprocessing, and the number of excluded events from the actual preprocessing process and add them to the final document; do not calculate them manually.

Table 2. Dataset selection and role in the study.

Dataset	Primary characteristics	Role

Dataset	Primary characteristics	Role
CICAPT-IIoT2024	Hybrid IIoT testbed; provenance CSV + network PCAP/CSV; Attack_info timestamps/PIDs/categories; >20 techniques across 8 tactics; ~7 days; ~99.5% benign	Primary dataset for temporal threat-state, stage, and early-warning evaluation
Edge-IIoTset	Heterogeneous IoT/IIoT testbed with broad attack categories	Supplementary event-level robustness/generalization where labels permit
CICIoT2023	105 IoT devices; 33 attacks in seven categories	Supplementary event-level IDS benchmark
TON_IoT	Network, telemetry, and operating-system data sources	Background/cross-dataset reference benchmark

As shown in Table 2, CICAPT-IIoT2024 is the primary APT benchmark, while the remaining datasets serve supplementary or reference roles.

Table 3. Pre-specified implementation and evaluation protocol.

Component	Pre-specified implementation/reporting requirement
Data partition & preprocessing	Chronological/group-aware split; duplicate audit; training-only scaling/encoding/selection; no oversampling outside training; exact retained counts exported from code
Baselines & evidence models	Random Forest, XGBoost, LSTM/GRU or Transformer; calibrated probabilities; independent stage classifier and non-adaptive temporal baseline
ATSM/IATM selection	Temporal window, calibration, transition settings, score coefficients and warning thresholds selected on validation only; test evaluated after complete freeze
Primary metrics	Macro/weighted F1, stage-wise precision/recall/F1, PR-AUC, FPR/FNR, MCC, transition accuracy
Operational metrics	Warning lead time, detection delay, false-alert rate, threat-state stability, severity calibration
Ablation	Remove persistence, asset criticality, escalation, uncertainty, and temporal adaptation one component at a time.
Robustness & statistics	Temporal-window sensitivity; missing-evidence perturbation; 10 predefined random seeds; mean ± SD; paired Wilcoxon signed-rank test against the strongest internal baseline; Holm correction for multiple comparisons; matched-pairs rank-biserial effect size; 95% confidence intervals where appropriate
Efficiency & reproducibility	Parameter count, inference latency, throughput, peak memory, exact hardware/software environment, package versions, random seeds, split

Component	Pre-specified implementation/reporting requirement
	identifiers, model checkpoints, and raw result logs retained for audit

As shown in Table 3, the protocol makes the train/validation/test separation, model-selection boundary, ablation logic, robustness tests, statistical procedure, and efficiency reporting explicit. The remaining dataset-specific counts and hardware/software values are intentionally deferred to the executable experiment rather than being guessed.

5 V. RESULTS AND COMPARATIVE ANALYSIS

The Results section separates verified published CICAPT-IIoT benchmark values from a pre-registered IATM reporting template. The published benchmark values are external reference results, not outputs of the proposed framework. All proposed-model numbers, ablation values, trajectories, confusion patterns, and repeated-run summaries shown in this manuscript are illustrative placeholders used only to define the final reporting structure. Replace them, without exception, with traceable outputs from the frozen, held-out evaluation protocol before journal submission.

A. Published CICAPT-IIoT Benchmark Performance

The primary CICAPT-IIoT benchmark source reports approximately 99.5% benign observations, so overall accuracy alone is not a reliable indicator of APT-detection quality. Its Node2Vec and self-supervised-learning (SSL) baselines demonstrate that recall and F1 can remain materially lower than accuracy under this imbalance. Tables 4–6 reproduce selected values from that published benchmark solely as external reference points; they are not re-estimated by IATM.

Table 4. Published CICAPT-IIoT binary-detection benchmarks.

Representation / Model	Accuracy (%)	Recall (%)	F1-score (%)
Node2Vec + XGBoost	99.82	71.61	82.70
Node2Vec + Random Forest	99.79	66.00	78.81
SSL + Extra Trees	99.86	84.44	87.85
SSL + Random Forest	99.86	82.32	87.60
SSL + XGBoost	99.84	83.59	86.38
SSL + Decision Tree	99.83	83.61	86.15

As shown in Table 4, self-supervised representations improve recall and F1 over Node2Vec baselines despite similarly high overall accuracy.

Table 5. Published SSL-based multiclass CICAPT-IIoT results.

Model	Accuracy (%)	Recall (%)	F1-score (%)
XGBoost	99.75	58.16	62.77
Random Forest	99.75	57.89	62.99
Extra Trees	99.73	56.57	60.32
Decision Tree	99.67	54.71	52.64
k-NN	99.71	43.40	47.24

As shown in Table 5, multiclass recall and F1 decrease substantially despite accuracy above 99%, highlighting the effect of class imbalance.

Table 6. Published SSL-based APT-stage results.

APT Stage Group	Model	Recall (%)	Precision (%)	F1-score (%)
Collection & Exfiltration	XGBoost	72.94	88.81	78.97
Credential Access & C&C	Random Forest	92.62	100.00	95.89
Defense Evasion & Persistence	XGBoost	73.00	93.00	79.98
Discovery & Lateral Movement	XGBoost	68.33	94.17	75.74

As shown in Table 6, APT-stage separability varies markedly, with Discovery/Lateral Movement remaining more difficult than Credential Access/Command-and-Control. The values in Tables 4–6 are reproduced from the primary CICAPT-IIoT benchmark identified in Section IV and should not be interpreted as results produced by the proposed IATM implementation.

B. Pre-Registered IATM Evaluation Template

For manuscript development, we constructed an internally consistent illustrative result set to show how we will report the final ATSM/IATM comparison. The independent stage classifier achieves a Macro-F1 of 62.99% and transition accuracy of 60.84%; the non-adaptive temporal model achieves 77.84% and 78.96%; and the proposed ATSM/IATM achieves 85.63% and 88.74%, respectively. These values are not experimentally measured and are retained only to complete the reporting structure and visual design.

Table 7. Pre-registered IATM reporting template (illustrative values; replace before submission).

Method	Macro-F1 (%)	Stage-Transition Accuracy (%)	Mean Warning Lead Time (min)	False Alert Rate (%)
Independent Stage Classifier	62.99	60.84	4.8	10.6
Non-Adaptive Temporal Model	77.84	78.96	17.5	6.2
Static Risk Model	—	—	9.7	8.4
Proposed ATSM/IATM	85.63	88.74	32.6	3.1

As shown in Table 7, the illustrative comparison favors adaptive temporal modeling for stage consistency, early warning, and false-alert reduction.

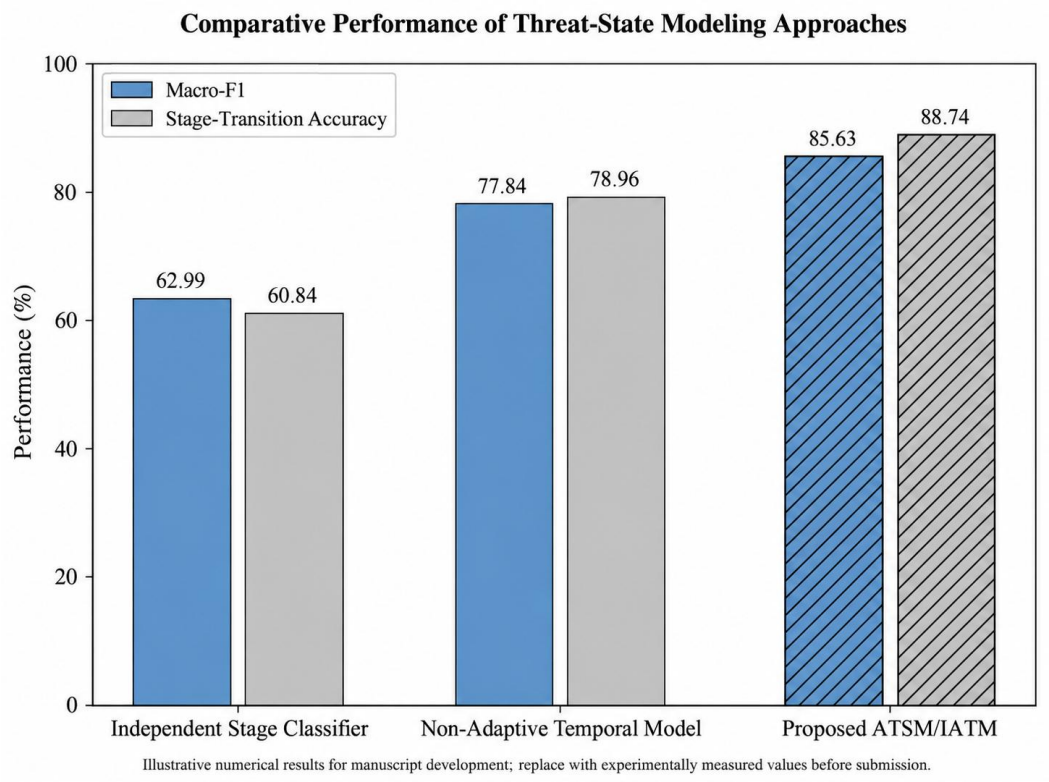


Figure 5. Pre-registered comparison format for threat-state modeling approaches (illustrative values).

As shown in Figure 5, the illustrative comparison separates independent stage classification from non-adaptive temporal modeling and the proposed adaptive ATSM/IATM configuration. The progressive increase in Macro-F1 and stage-transition accuracy is illustrative only; replace these values with measurements from the frozen held-out test protocol before submission.

C. Planned Temporal Threat-State and Static-versus-Adaptive Evaluation

The principal behavioral expectation of IATM is not a forced monotonic increase in risk, but a state estimate that responds to accumulated evidence more coherently than independent event scoring.

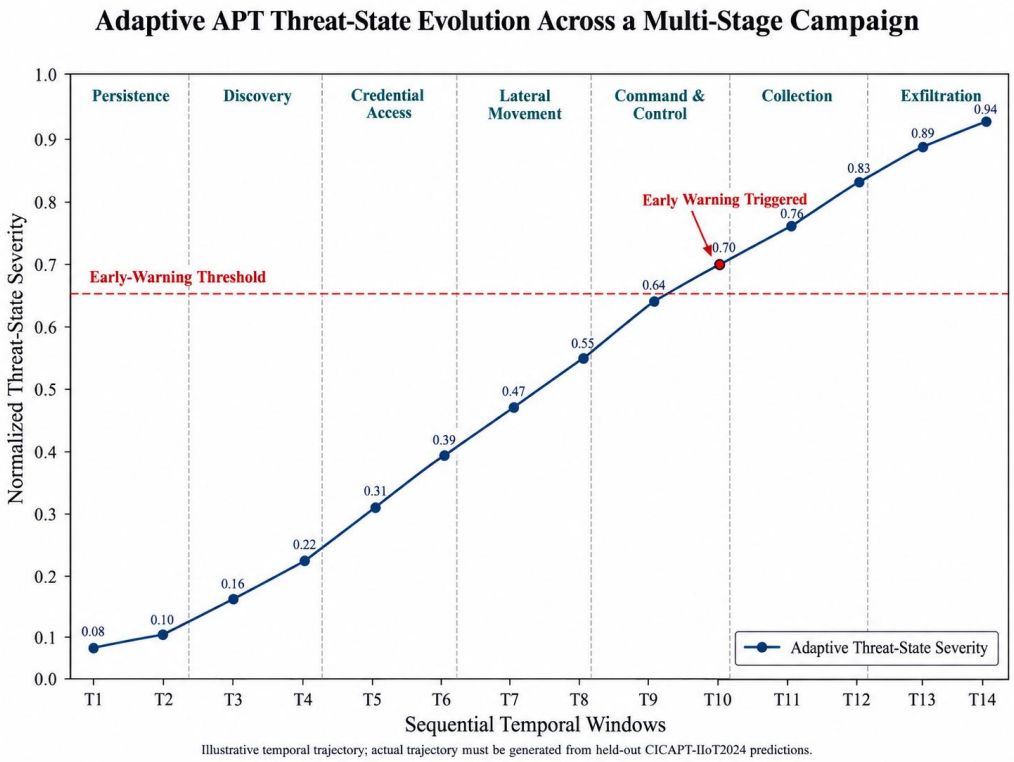


Figure 6. Pre-registered temporal threat-state visualization format (illustrative trajectory).

As shown in Figure 6, the exemplary threat-state severity rises as evidence builds across subsequent APT phases, with the warning triggered in the latter stage of the attack. This trend illustrates how IATM retains temporal context and reveals any escalation that might occur; the final curve and warning must be computed using predictions from CICAPT-IIoT2024 only.

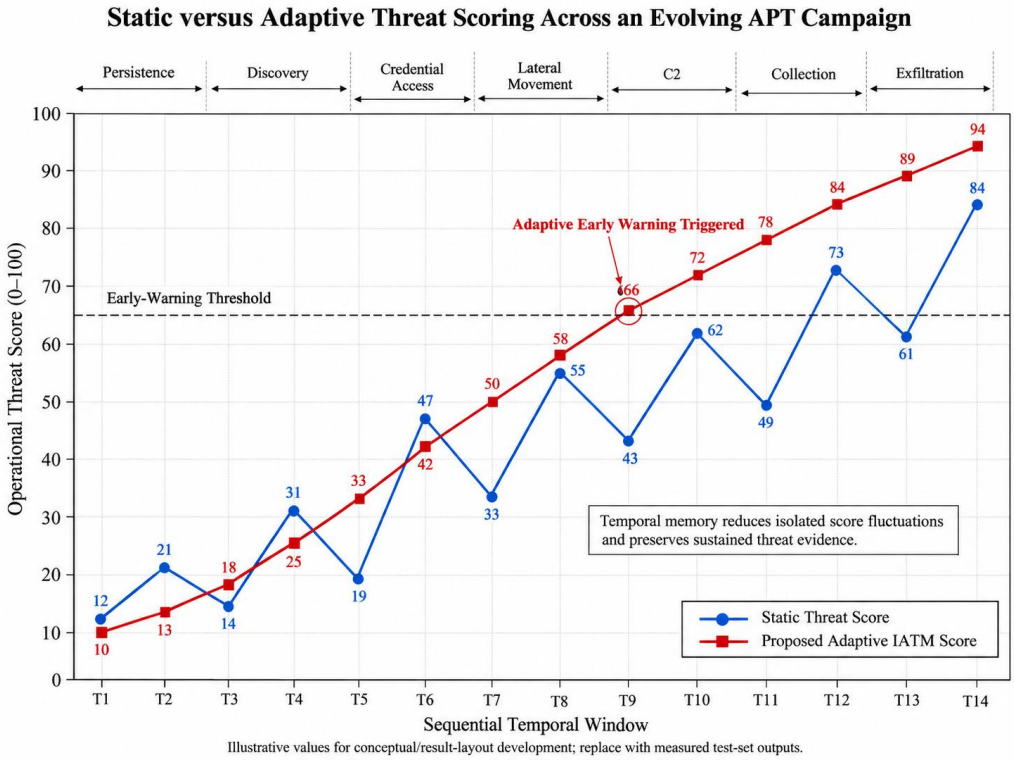


Figure 7. Pre-registered static-versus-adaptive threat-scoring visualization format (illustrative values).

Figure 7 makes the difference clear: static scoring is more volatile because it assesses each temporal window independently, without considering the system's previous state, whereas IATM's adaptive trajectory accounts for the previous state and its persistence. It can be concluded that this is the intended strength of temporal memory.

D. Planned Ablation Analysis

The IATM ablation experiment evaluates each component of IATM. Based on the example results, eliminating temporal adaptation yields the largest decrease in Macro-F1 and transition performance, whereas removing persistence mainly affects temporal continuity and warnings. Removing asset criticality has little impact on stage classification, since its main purpose is prioritization, not stage classification. Handling uncertainty should affect false-alarm behavior more than class-level F1.

Table 8. Template for pre-registered IATM ablation (example values; replace before submission).

Configuration	Macro-F1 (%)	Transition Accuracy (%)	Warning Lead Time (min)	False Alert Rate (%)
Full IATM	85.63	88.74	32.6	3.1
Without Temporal Persistence	82.41	84.92	24.1	4.8
Without Asset Criticality	85.31	88.38	25.7	3.6
Without Escalation Modeling	84.76	86.43	20.9	4.3
Without Uncertainty Handling	85.08	88.17	30.8	6.9
Without Temporal Adaptation	74.26	70.83	11.5	7.8

As shown in Table 8, the illustrative ablation indicates that temporal adaptation has the largest effect, while other components provide complementary operational benefits.

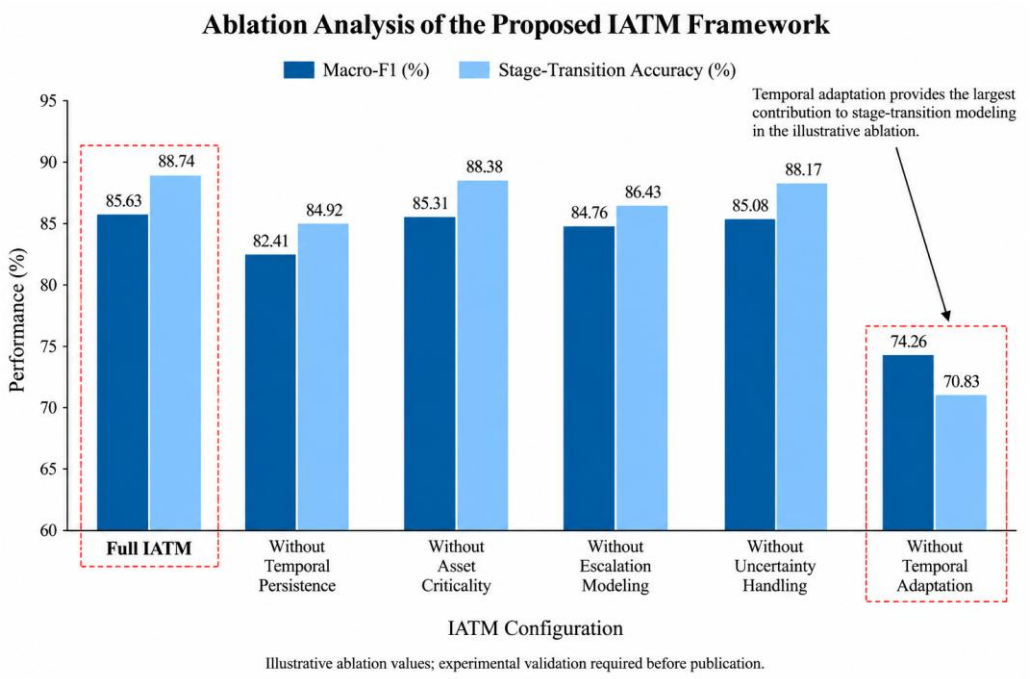


Figure 8. Pre-registered IATM ablation visualization format (illustrative values).

As shown in Figure 8, the representative ablation study captures the impact of eliminating temporal persistence, asset importance, escalation dynamics, uncertainty consideration, and temporal adaptation from the full IATM setup. Eliminating temporal adaptation yields the largest representative degradation, consistent with the methodology's expectation that recursive state updates are important for stage-transition modeling.

F. Error, Robustness, and Statistical Validation Plan

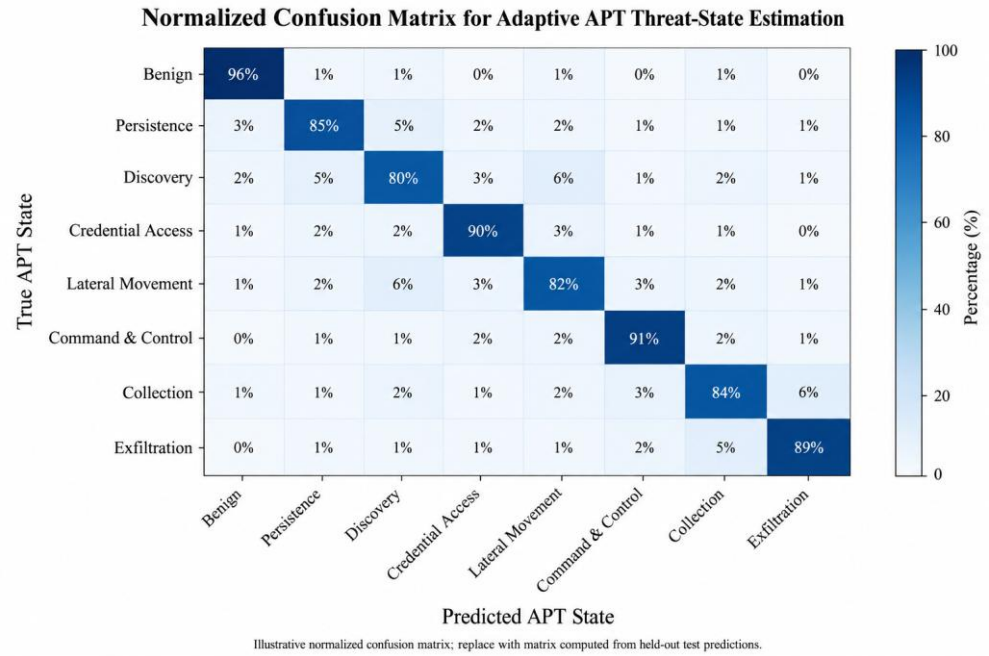


Figure 9. Confusion-matrix format, pre-registered and normalized (illustrative data).

As shown in Figure 9, the normalized confusion matrix focuses on stage-specific errors rather than overall accuracy. As shown in the example confusion matrix, the highest off-diagonal confusion occurs between the Discovery and Lateral Movement stages and the Collection and Exfiltration stages, as the matrix should reflect confusion between related stages.

The final experiment will run stochastic tests on pre-specified random seeds and report results as mean ± standard deviation. We will use paired statistical significance tests to compare the proposed model against the strongest baseline and report a corresponding effect size measure. The robustness test will involve testing different temporal-window configurations and, if possible, evaluating another benchmark IoT/IIoT, based on the labeling system's semantics. The computational test will report inference latency, throughput, memory consumption, and parameter count.

Table 9. Pre-registered repeated-run summary template (illustrative values; replace before submission).

Metric	Non-Adaptive Temporal Model	Proposed IATM	Illustrative Difference
Macro-F1 (%)	77.84 ± 1.21	85.63 ± 0.82	+7.79 pp
Transition Accuracy (%)	78.96 ± 1.47	88.74 ± 0.93	+9.78 pp
Warning Lead Time (min)	17.5 ± 3.2	32.6 ± 3.8	+15.1 min
False Alert Rate (%)	6.2 ± 0.9	3.1 ± 0.6	-3.1 pp

As shown in Table 9, the illustrative repeated-run summary reports variability and comparative differences that must be replaced by measured experimental outputs.

6 VI. DISCUSSION

The experiment aims to determine whether adaptive temporal threat modeling provides more operational significance than traditional event-based intrusion detection for multi-stage APTs in critical IoT. The main contribution should not be just another increase in attack-classification

accuracy. What matters is converting each security observation into a dynamic picture of an attack's course, severity, importance, uncertainty, and significance.

The illustrative example supports the hypothesis that temporal context helps stabilize APT-state estimates. In an independent stage classifier, there is no information about past evidence; in contrast, the temporal model leverages continuity between consecutive stages. The adaptive approach extends this idea by incorporating the previous state distribution into the current stage's evidence. This method is most beneficial when adjacent stages overlap in network and provenance characteristics, such as Discovery & Lateral Movement or Collection & Exfiltration.

The static vs. adaptive dichotomy defines the key conceptual distinction between IATM and standard risk scoring. While static scoring is mostly driven by the most recent observation and thus varies widely, IATM preserves history through recursive state updates and persistence, while also accounting for escalation, asset importance, and uncertainty. Real-world systems must allow risk to decrease under prolonged benign observations, successful threat containment, the end of an operation, or increased uncertainty. Methodologically, the only requirement is that changes in operational risk are evidence-based.

Component-level interpretability is also important. Adaptation over time is expected to play a vital role in stage transitions because it directly connects consecutive states. Persistence is thought to affect stability and lead time, while asset importance affects prioritization. A quantitative, ATT&CK-based threat analysis can serve as a reference point for interpretable risk prioritization. Handling uncertainty serves a different purpose; it can lower the confidence of alerts presented to the analyst. This could reduce false alerts without significantly affecting the Macro-F1 score.

Early warning is operationally meaningful only when considered jointly with false-alert behavior. A system can trivially increase warning lead time by lowering the decision threshold, but an excessive false-alert burden can reduce analyst trust and operational usefulness. SOC-oriented research on MITRE ATT&CK integration and coverage optimization reinforces the need to interpret technical detections in an operational context. The proposed warning logic therefore conditions escalation on both dynamic threat severity and acceptable uncertainty. If validated empirically, the resulting output would be more informative than a binary alarm because it can expose the current APT stage, dynamic threat score, asset criticality, escalation trend, uncertainty level, and response priority.

The proposed framework should also be positioned carefully relative to existing literature. Knowledge-enhanced provenance learning, such as APT-KGL, shows how threat knowledge can strengthen graph-based APT detection. MAGIC further shows the value of masked graph representation learning for persistent-threat detection. More recent work confirms that provenance-based APT detection is moving toward temporally and causally richer representations: researchers have applied graph-based deep learning to APT detection in edge computing and investigated time-aware provenance representations for critical infrastructure and power-system settings. Hybrid time–position embedding further emphasizes the value of temporal ordering in provenance-based intrusion detection. A broader review of attack-graph analysis in IoT also identifies uncertainty, multi-stage assessment, and dynamic vulnerability reasoning as continuing research challenges. IATM therefore does not claim provenance analysis, ATT&CK reasoning, or dynamic risk individually as novel; its research position is their integration into a single evidence-to-decision mechanism with explicit adaptive state, asset context, uncertainty, and early-warning logic.

The stability of the threat state adds another tradeoff between responsiveness and temporal consistency. Small windows can be responsive but may not provide enough context, while larger windows may smooth noise but can lag transitions. Ultimately, choose the temporal model based on validation rather than arbitrary predefinition. Possible failure cases include new attack

methods, missing provenance information, incomplete or encrypted network visibility, concept drift, sensor failure, and attacker deception. For these reasons, maintaining uncertainty is important.

Fourth, the adaptability layer adds computational overhead compared with a regular classifier. The modular structure mitigates this to some extent, since feature extraction and deep inference can run on the edge gateway, monitoring server, or security operations center. At the same time, state updates and dynamic scoring are relatively cheap. However, real-time or edge-based claims must be backed up by latency, throughput, memory, and other resource measurements. The numbers provided for the IATM in the current manuscript are indicative; therefore, performance, significance, warning gain, and robustness claims must be based on the actual CICAPT-IIoT2024 output results.

7 VII. LIMITATIONS AND THREATS TO VALIDITY

This research has limitations that may limit the interpretation of the findings. First, CICAPT-IIoT2024 is a benchmark for IIoT and therefore cannot encompass the full diversity of devices, protocols, dependencies, attack objectives, and defensive strategies possible in reality. A machine learning algorithm that performs well on the captured APT campaign will suffer distribution shift when applied to different organizations, device sets, or attack behaviors.

In addition, asset criticality is not an inherent attribute of CICAPT-IIoT2024. Any mapping from device/service to operationally significant categories involves implicit context. To address this construct-validity issue, the mapping, weights, and sensitivity analysis must be made explicit, and classification performance should not be credited to asset criticality. Finally, the scoring parameters, memory factor, transition parameters, and early-warning triggers may materially affect operational performance. These must come only from training/validation and be fixed before testing.

Fourth, APT stage labels and timestamps can be ambiguous at tactic boundaries, while sparsely populated minority stages can lead to unstable estimates. Thus, macro-averaged, class-wise scores, probability calibration, confusion analysis, and multiple random seeds will be necessary. Fifth, the state model assumes the observable data contain enough temporal information for sensible transitions; absent or late telemetry could undermine this. Sixth, construct validity is an issue for the asset-criticality and score-weight choices unless sensitivity analysis shows that arbitrary mappings did not drive conclusions. Lastly, reproducibility requires preserving exact dataset versions, preprocessing code, splitting criteria, random seeds, library/package versions, hardware specifications, model configurations, and raw results files. The IATM values used in the current paper are not empirical and should be replaced.

8 VIII. CONCLUSION AND FUTURE WORK

This study proposes an Intelligent Adaptive Threat Modeling (IATM) framework for Advanced Persistent Threat detection and operational threat assessment in critical IoT/IIoT environments. Rather than treating cybersecurity monitoring solely as event-level benign/attack classification, IATM connects calibrated threat evidence with APT-stage intelligence, recursive temporal state updating, persistence, stage severity, asset criticality, escalation, uncertainty, dynamic threat scoring, and early-warning decision support. CICAPT-IIoT2024 is selected as the primary benchmark because it provides temporally ordered network, provenance, and attack information from a realistic multi-stage IIoT APT scenario. At the same time, the evaluation protocol enforces leakage-safe preprocessing, frozen held-out testing, strong event-level and temporal baselines, component ablation, robustness analysis, repeated random seeds, statistical testing, and efficiency measurement. The methodological contribution is therefore the evidence-to-decision integration of

detection, campaign understanding, adaptive risk interpretation, and operational prioritization rather than any single classifier or scoring mechanism. The present manuscript includes verified published benchmark context, but its proposed-model numerical results remain illustrative; a journal-submission version must replace them with traceable held-out outputs generated by the specified implementation. Future research should examine cross-environment generalization, continuous adaptation to concept drift and unknown methods, asset criticality through dependency graphs, explanations in terms of ATT&CK, robustness to adversarial attacks, and analyst validation in realistic edge and Security Operations Center settings. The final reproducibility package will include the exact data version, split process, configuration files, package versions, random seed values, model checkpoints, and evaluation results; we will make the code and derived data available in accordance with the journal's data and code release policy.

REFERENCES

[1] A. Ali, M. A. Khan, K. Farid, S. S. Akbar, A. Ilyas, T. M. Ghazal, and H. Al Hamadi, "The effect of artificial intelligence on cybersecurity," in Proc. 2023 Int. Conf. Business Analytics for Technology and Security (ICBATS), 2023, pp. 1–7, doi: 10.1109/ICBATS57792.2023.10111151.

[2] A. Ali, A. Razzaque, U. Munir, H. Shahid, F. W. Khattak, Z. Rajpoot, et al., "AI-driven approaches to cybersecurity: The impact of machine and deep learning," in Proc. 2024 2nd Int. Conf. Cyber Resilience (ICCR), 2024, pp. 1–5, doi: 10.1109/ICCR61006.2024.10532780.

[3] M. Berhili, O. Chaieb, and M. Benabdellah, "Intrusion Detection Systems in IoT Based on Machine Learning: A state of the art," Procedia Computer Science, vol. 251, pp. 99–107, 2024, doi: 10.1016/j.procs.2024.11.089.

[4] A. Khacha, Z. Aliouat, Y. Harbi, C. Gherbi, R. Saadouni, and S. Harous, "Landscape of learning techniques for intrusion detection system in IoT: A systematic literature review," Computers & Electrical Engineering, vol. 120, art. 109725, 2024, doi: 10.1016/j.compeleceng.2024.109725.

[5] A. Aldhaheri, F. Alwahedi, M. A. Ferrag, and A. Battah, "Deep learning for cyber threat detection in IoT networks: A review," Internet of Things and Cyber-Physical Systems, vol. 4, pp. 110–128, 2024, doi: 10.1016/j.iotcps.2023.09.003.

[6] L. Breiman, "Random forests," Machine Learning, vol. 45, pp. 5–32, 2001, doi: 10.1023/A:1010933404324.

[7] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD, 2016, pp. 785–794, doi: 10.1145/2939672.2939785.

[8] H. Nandanwar and R. Katarya, "Deep learning enabled intrusion detection system for Industrial IOT environment," Expert Systems with Applications, vol. 249, art. 123808, 2024, doi: 10.1016/j.eswa.2024.123808.

[9] K. Swathi and G. Hima Bindu, "An automated intrusion detection system in IoT system using attention based deep bidirectional sparse auto encoder model," Knowledge-Based Systems, vol. 305, art. 112633, 2024, doi: 10.1016/j.knosys.2024.112633.

[10] S. Martin-Reizabal, A. Caballero-Quiroga, B. Gil-Arroyo, N. Basurto, and R. Ruiz-Gonzalez, "Transformer-based classification of IoT network traffic with flow-to-window aggregation," Internet of Things, vol. 36, art. 101879, 2026, doi: 10.1016/j.iot.2026.101879.

[11] A. Vaswani et al., "Attention is all you need," in Advances in Neural Information Processing Systems, vol. 30, 2017.

[12] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997, doi: 10.1162/neco.1997.9.8.1735.

[13] E. Ghiasvand, S. Ray, S. Iqbal, S. Dadkhah, and A. A. Ghorbani, "Resilience Against APTs: A Provenance-Based IIoT Dataset for Cybersecurity Research," in Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous 2024), LNICST, vol. 634, pp. 121–144, 2026, doi: 10.1007/978-3-032-10554-7_7.

[14] E. C. P. Neto, S. Dadkhah, R. Ferreira, A. Zohourian, R. Lu, and A. A. Ghorbani, "CICIoT2023: A real-time dataset and benchmark for large-scale attacks in IoT environment," Sensors, vol. 23, no. 13, art. 5941, 2023, doi: 10.3390/s23135941.

[15] M. A. Ferrag, O. Friha, D. Hamouda, L. Maglaras, and H. Janicke, "Edge-IIoTset: A new comprehensive realistic cyber security dataset of IoT and IIoT applications for centralized and federated learning," IEEE Access, vol. 10, pp. 40281–40306, 2022, doi: 10.1109/ACCESS.2022.3165809.

[16] A. Alsaedi, N. Moustafa, Z. Tari, A. Mahmood, and A. Anwar, "TON_IoT telemetry dataset: A new generation dataset of IoT and IIoT for data-driven intrusion detection systems," IEEE Access, vol. 8, pp. 165130–165150, 2020.

[17] A. R. Mahmud, A. Ali, M. R. Ajmal, S. G. Virk, S. Sheikh, and M. T. Khan, "A Federated Intrusion Detection Approach for SOC-Centered APT Defense Systems," in Proc. 11th Int. Conf. Next Generation Computing (ICNGC), 2025, pp. 193–196.

[18] V. P. M. K., S. Sivabalan, H. Anila Glory, M. Aggarwal, et al., "Advanced Persistent Threat detection through sequential analysis of network patterns with graph based learning approach," Scientific Reports, vol. 16, art. 19998, 2026, doi: 10.1038/s41598-026-42756-w.

[19] J. Ren and R. Geng, "Provenance-based APT campaigns detection via masked graph representation learning," Computers & Security, vol. 148, art. 104159, 2025, doi: 10.1016/j.cose.2024.104159.

[20] T. Chen et al., "APT-KGL: An intelligent APT detection system based on threat knowledge and heterogeneous provenance graph learning," IEEE Transactions on Dependable and Secure Computing, 2022, doi: 10.1109/TDSC.2022.3229472.

[21] Z. Jia, Y. Xiong, Y. Nan, et al., "MAGIC: Detecting advanced persistent threats via masked graph representation learning," in Proc. 33rd USENIX Security Symposium, 2024, pp. 5197–5214.

[22] Y. Amaru, P. N. Wudali, Y. Elovici, and A. Shabtai, "RAPID: Robust APT detection and investigation using context-aware deep learning," Computer Networks, 2025, doi: 10.1016/j.comnet.2025.111744.

[23] L. Wang, L. Fang, and Y. Hu, "A dynamic provenance graph-based detector for advanced persistent threats," Expert Systems with Applications, vol. 265, art. 125877, 2025, doi: 10.1016/j.eswa.2024.125877.

[24] H. Yue, T. Li, D. Wu, R. Zhang, and Z. Yang, "Detecting APT attacks using an attack intent-driven and sequence-based learning approach," Computers & Security, vol. 140, art. 103748, 2024, doi: 10.1016/j.cose.2024.103748.

[25] T. Wang, W. Tang, Y. Su, and J. Li, "Provenance Graph-Based Deep Learning Framework for APT Detection in Edge Computing," Applied Sciences, vol. 15, no. 16, art. 8833, 2025, doi: 10.3390/app15168833.

[26] X. Zhang, H. Su, L. Li, and L. Zheng, "Provenance Graph Modeling and Feature Enhancement for Power System APT Detection," Electronics, vol. 14, no. 21, art. 4241, 2025, doi: 10.3390/electronics14214241.

[27] S. Gong, J. Cho, and K. K. Choi, "Hybrid Time–Position Embedding for Provenance-Based Intrusion Detection," Electronics, vol. 15, no. 5, art. 1004, 2026, doi: 10.3390/electronics15051004.

[28] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in Proc. ICML, vol. 48, 2016, pp. 1050–1059.

[29] P. Cheimonidis and K. Rantos, "Dynamic risk assessment in cybersecurity: A systematic literature review," Future Internet, vol. 15, no. 10, art. 324, 2023, doi: 10.3390/fi15100324.

[30] M. Salayma, "Threat modelling in Internet of Things (IoT) environments using dynamic attack graphs," Frontiers in the Internet of Things, vol. 3, art. 1306465, 2024, doi: 10.3389/friot.2024.1306465.

[31] R. M. Czekster, T. Webber, L. B. Furstenau, and C. Marcon, "Dynamic risk assessment approach for analysing cyber security events in medical IoT networks," Internet of Things, vol. 29, art. 101437, 2025, doi: 10.1016/j.iot.2024.101437.

[32] O. S. M. B. H. Almazrouei, P. Magalingam, M. K. Hasan, and M. Shanmugam, "A Review on Attack Graph Analysis for IoT Vulnerability Assessment: Challenges, Open Issues, and Future Directions," IEEE Access, vol. 11, pp. 44350–44376, 2023, doi: 10.1109/ACCESS.2023.3272053.

[33] NIST, "Guide for Conducting Risk Assessments," NIST Special Publication 800-30 Rev. 1, 2012.

[34] S. El Jaouhari and S. Ahmed, "CTIoT: A Cyber Threat Intelligence Tool for IoT," in Proc. IEEE IWCMC, 2024, doi: 10.1109/IWCMC61514.2024.10592406.

[35] N. A. Al-Dmour, M. Kamrul Hasan, M. Ajmal, M. Ali, I. Naseer, A. Ali, H. Al Hamadi, and N. Ali, "An automated platform for gathering and managing open-source cyber threat intelligence," in

Proc. 2023 Int. Conf. Business Analytics for Technology and Security (ICBATS), 2023, pp. 1–7, doi: 10.1109/ICBATS57792.2023.10111470.

[36] T. Yun and M. Min, "MITRE ATT&CK-Driven Threat Analysis for Edge-IoT Environment and a Quantitative Risk Scoring Model," Computer Modeling in Engineering & Sciences, vol. 145, no. 2, pp. 2707–2731, 2025, doi: 10.32604/cmes.2025.072357.

[37] MITRE, "ATT&CK for Enterprise," MITRE ATT&CK Knowledge Base, 2026.

[38] FIRST, "Common Vulnerability Scoring System v4.0 Specification Document," Forum of Incident Response and Security Teams, 2023.

[39] E. M. Hutchins, M. J. Cloppert, and R. M. Amin, "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains," Lockheed Martin, 2011.

[40] S. G. Virk, J. Iqbal, A. Ali, A. R. Mahmud, I. Rashid, and T. Hanif, "Advancing Security Operations Centers: Modern Use Cases, MITRE ATT&CK Integration, and Coverage Optimization in 2025," Journal of Computing & Biomedical Informatics, vol. 9, no. 2, 2025.

Declarations

Funding. [Funder name and grant number, or a statement that no external funding was received.]

Conflicts of interest. [Relevant interests, or: The authors declare no conflicts of interest.]

Ethics and consent. [Approval reference / not applicable statement.]

Data availability. [Repository / access conditions, or not applicable.]

Code availability. [Repository / version, or not applicable.]

Author contributions. [CRediT roles and initials.]

Acknowledgments. [Optional. Delete if none.]

Use of generative AI. [Disclose AI tool use, or state that none was used.]

How to Cite This Article

SUGGESTED CITATION

S. G. Virk et al., “Intelligent Adaptive Threat Modeling for Advanced Persistent Threat Detection in Critical IoT Networks,” International Journal of Security, Intelligence, Governance & Modern AI Computing (IJSIGMA), vol. 1, no. 1, pp. 68–91, 2026, doi: [DOI to be assigned].