International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 7, Issue 3, 46-51, 2026
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V7I3P106

Original Article

AI-Driven Predictive Cyber Threat Intelligence for Critical Infrastructure Protection

Dr. Naga Venkata Aswini Pavan Kumar Inguva
Department of Information Technology / Cybersecurity Research, USA.

Received On: 26/05/2026 ***Revised On: 24/06/2026*** ***Accepted On: 01/07/2026*** ***Published On: 11/07/2026***

Abstract - *Critical infrastructure systems such as energy, telecommunications, transportation, healthcare, water, and government services increasingly depend on interconnected information technology, operational technology, cloud, and Internet of Things environments. This growing interconnectivity expands the attack surface and increases the consequences of cyber incidents. Traditional defensive approaches remain essential, but many are reactive and depend on known indicators, static rules, or post-event investigation. This paper proposes an AI-driven Predictive Cyber Threat Intelligence Framework for critical infrastructure protection. The framework integrates multi-source cyber threat intelligence, network and endpoint telemetry, vulnerability information, adversary behavior models, machine learning, anomaly detection, risk scoring, explainable analytics, and response orchestration. The proposed model is designed to support proactive threat anticipation, prioritization of high-risk assets, reduction of alert fatigue, and faster response to emerging attacks. The paper also presents a literature-based comparison with existing approaches, a methodology for future experimental validation using public intrusion detection datasets, and implementation considerations for critical infrastructure operators. The contribution of this work is a practical, modular, and standards-aligned framework that connects cyber threat information sharing, AI-based prediction, criticality-aware risk scoring, and operational decision support for high-impact environments.*

Keywords - *Artificial Intelligence, Cyber Threat Intelligence, Critical Infrastructure Protection, Machine Learning, Predictive Analytics, Risk Scoring, Intrusion Detection, MITRE ATT&CK, NIST Cybersecurity Framework.*

1. Introduction

Critical infrastructure provides essential services that support public safety, economic stability, national security, and daily life. Energy generation and distribution, telecommunications, transportation, healthcare, water management, finance, and public-sector services are now highly dependent on digital systems. The convergence of information technology (IT), operational technology (OT), industrial control systems (ICS), cloud platforms, and Internet of Things (IoT) devices has improved efficiency but also introduced new vulnerabilities. A disruption in these environments can have consequences beyond data loss, including service outages, physical damage, safety risks, and cascading failures across dependent sectors.

Cybersecurity guidance for critical infrastructure has increasingly emphasized risk management, governance, information sharing, detection, response, and resilience. NIST Cybersecurity Framework 2.0 provides a broad structure for managing cybersecurity risk across organizations and sectors [1]. CISA's Cross-Sector Cybersecurity Performance Goals provide baseline practices designed for critical infrastructure entities [2]. NIST SP 800-150 highlights the importance of establishing and maintaining cyber threat information sharing relationships [3]. MITRE ATT&CK provides a widely used knowledge base of adversary tactics and techniques for mapping observed behavior to known attack patterns [4]. These resources are valuable, but their effectiveness depends on the ability of organizations to process high-volume security data and operationalize intelligence quickly enough to anticipate attacks.

Many cybersecurity systems still operate primarily as reactive mechanisms. They detect known signatures, match predefined rules, or generate alerts after suspicious activity has already occurred. Such approaches can be effective against known malware, previously observed indicators of compromise, and clearly defined policy violations. However, they may underperform against novel attack chains, living-off-the-land techniques, multi-stage intrusion campaigns, insider threats, and adversaries that deliberately evade detection. Critical infrastructure operators require a shift from reactive monitoring toward predictive cyber threat intelligence that can identify early warning signals and estimate likely attack progression before operational disruption occurs.

Artificial intelligence and machine learning offer capabilities that can support this transition. ML models can detect hidden relationships in large datasets, identify abnormal behaviors, classify attack patterns, and prioritize risks based on contextual factors. Deep learning and sequential models can capture temporal patterns in network traffic and system events. Ensemble learning can improve robustness, and explainable AI can help analysts understand

why a model produced a specific risk score. The challenge is not simply applying AI to logs; rather, it is designing an integrated framework that combines intelligence sources, critical asset context, adversary behavior knowledge, model outputs, and response workflows.

This paper proposes an AI-driven Predictive Cyber Threat Intelligence Framework for critical infrastructure protection. The framework is designed as a modular architecture that can ingest multi-source data, normalize indicators, map events to adversary tactics, predict threat likelihood, assign asset-aware risk scores, and support automated or analyst-guided response decisions. The paper is intended as a framework and methodology paper; it does not claim production deployment results. Instead, it establishes a standards-aligned architecture and an evaluation plan that can be validated using public cybersecurity datasets and sector-specific operational data in future work.

2. Related Work

Research in AI-based cybersecurity has expanded significantly over the past decade. Buczak and Guven surveyed data mining and machine learning methods for cyber intrusion detection and highlighted the importance of training data, feature engineering, and model selection in cyber analytics [5]. Their work remains influential because it positioned ML as a practical approach for detecting misuse and intrusion patterns. Later research has extended ML approaches to deep learning, ensemble learning, behavioral analytics, and IoT intrusion detection.

Intrusion detection datasets have played a central role in cybersecurity experimentation. UNSW-NB15 was created to include modern normal activities and synthetic contemporary attack behaviors, with attack categories such as Fuzzers, Analysis, Backdoors, DoS, Exploits, Generic, Reconnaissance, Shellcode, and Worms [6]. CICIDS2017 was developed to support intrusion traffic characterization and evaluation of intrusion detection approaches [7]. The ToN-IoT dataset provides heterogeneous Industry 4.0, IoT, and IIoT data to evaluate AI-enabled cybersecurity applications [8]. BoT-IoT was developed in a cyber range environment and provides botnet-related attack data in multiple formats [9]. These datasets support model training and benchmarking but do not fully represent the complexity of every critical infrastructure environment.

Critical infrastructure cybersecurity introduces unique requirements that are not always present in enterprise IT systems. Raval et al. surveyed attacks, AI-based defenses, and security risks across critical infrastructure sectors and emphasized that AI mechanisms can support the defense of energy, healthcare, agriculture, finance, and other critical sectors [10]. The U.S. Department of Energy has also described both benefits and risks of AI for the energy sector, including malicious event detection and diagnosis, while warning that AI must be deployed carefully in safety-critical contexts [11]. CSET's workshop report on securing critical infrastructure in the age of AI further emphasized the need for technical mitigations and policy recommendations for AI use in critical infrastructure [12].

Cyber threat intelligence focuses on collecting, analyzing, sharing, and operationalizing information about threats. NIST SP 800-150 provides guidance for establishing sharing relationships, selecting sources, scoping sharing activities, and using shared threat information in cybersecurity practices [3]. However, CTI is often operationalized as a feed of indicators rather than as a predictive intelligence capability. Indicator-based intelligence can be useful, but it may become stale quickly, lacks context, and often does not indicate future attack likelihood. Predictive CTI requires correlating indicators with behavioral patterns, vulnerability exposure, asset criticality, and temporal changes in adversary activity.

Explainability is another important challenge. Rjoub et al. reviewed explainable artificial intelligence for cybersecurity and noted the importance of making AI-based security decisions understandable to analysts and stakeholders [13]. Black-box systems may produce high detection accuracy in controlled environments but face adoption challenges when security teams cannot interpret or validate model outputs. Critical infrastructure operators require explainable predictions because operational decisions can affect availability, public safety, and compliance obligations. Recent public-sector and industry developments also show the increasing relevance of AI in critical infrastructure cybersecurity. DHS formed an AI safety and security advisory board in 2024 to advise on safe AI deployment in critical infrastructure sectors such as energy, transportation, defense, healthcare, and financial services [14]. Reuters also reported that U.S. policy efforts have sought to strengthen the resilience of critical infrastructure against current and future threats, including ransomware and advanced state-linked activity [15]. These developments support the need for AI-enabled security approaches that are not only technically effective but also governable, explainable, and aligned with national resilience objectives.

3. Research Gap and Problem Statement

The literature and current practice reveal several gaps. First, many security architectures remain reactive, with limited capability to forecast attack likelihood before exploitation. Second, CTI platforms often ingest external feeds but do not sufficiently integrate asset criticality, behavioral analytics, vulnerability exposure, and adversary tactics into a unified predictive model. Third, high false positive rates continue to reduce analyst efficiency and can delay response to true incidents. Fourth, many ML intrusion detection studies evaluate models on public datasets but do not address critical infrastructure constraints such as OT availability, legacy systems, safety requirements, and limited tolerance for disruption. Fifth, black-box AI predictions may not be trusted by operators unless accompanied by explainable reasoning.

The problem addressed in this paper is therefore: how can critical infrastructure organizations transform cyber

threat intelligence from a reactive, indicator-driven function into a proactive, AI-enabled predictive capability that supports risk-aware and explainable decision-making?

4. Proposed Framework

The proposed AI-driven Predictive Cyber Threat Intelligence Framework contains seven layers: data acquisition, intelligence fusion, feature engineering, AI prediction, criticality-aware risk scoring, explainability, and response orchestration. Each layer is designed to be modular so that organizations can adopt the framework gradually without replacing existing security operations investments.

The data acquisition layer integrates traditional IT telemetry with OT and critical infrastructure context. Sources may include firewall logs, network flow records, endpoint detection alerts, authentication logs, vulnerability scans, cloud security posture findings, ICS alarms, and sector-specific advisories. The intelligence fusion layer enriches raw events with external context such as known adversary tactics, vulnerability severity, exploited CVEs, malware families, and sector targeting patterns. The feature engineering layer converts events into predictive features. Examples include repeated failed authentication attempts, unusual remote access patterns, sudden changes in traffic volume, privilege escalation attempts, suspicious process execution, communication with known malicious infrastructure, exploitation attempts against exposed services, and MITRE ATT&CK technique mappings. For critical infrastructure, asset context is essential. A low-level alert affecting an internet-facing system that supports grid operations may deserve higher priority than a similar alert on a low-impact development machine.

Table 1: Layered Architecture of the Proposed AI-Driven Predictive Cyber Threat Intelligence Framework

Layer	Function	Example Inputs/Outputs
1. Data Acquisition	Collect raw telemetry and intelligence from heterogeneous sources.	Network flows, SIEM logs, endpoint alerts, OT events, vulnerability scans, OSINT, threat feeds.
2. Intelligence Fusion	Normalize, deduplicate, enrich, and correlate indicators and events.	Mapped indicators, enriched events, confidence scores, sector relevance.
3. Feature Engineering	Transform raw security events into model-ready features.	Temporal patterns, authentication failures, traffic anomalies, ATT&CK technique mappings.
4. AI Prediction Engine	Classify threats, detect anomalies, and forecast attack likelihood.	ML classifications, anomaly scores, predicted attack stage.
5. Criticality-Aware Risk Scoring	Combine prediction output with asset criticality and business impact.	Risk score, priority queue, critical asset exposure estimate.
6. Explainability Layer	Generate analyst-readable explanations for predictions.	Top contributing features, mapped tactics, recommended evidence.
7. Response Orchestration	Recommend or trigger mitigations while maintaining operational continuity.	Analyst alerts, playbook actions, containment recommendations.

The AI prediction engine combines supervised, unsupervised, and sequential models. Supervised models can classify known attack categories, while unsupervised models can identify novel anomalies. Sequential models such as LSTM or Transformer-based architectures can capture attack progression over time. Ensemble approaches can combine model outputs to improve resilience and reduce overdependence on a single algorithm. The risk scoring module converts model output into operational priorities. Risk scoring should incorporate predicted likelihood, impact, confidence, asset criticality, vulnerability exposure, exploit availability, and adversary relevance. The explainability layer then provides analysts with supporting evidence, such as the top features contributing to the score and the relevant ATT&CK tactics. Finally, the response orchestration layer connects predictions to playbooks, ticketing systems, SIEM/SOAR tools, and analyst workflows.

5. Methodology and Evaluation Plan

This paper proposes a methodology for validating the framework in future experimental work. The evaluation can be performed in three phases: baseline model evaluation using public intrusion detection datasets, framework-level evaluation using multi-source simulated CTI feeds, and operational evaluation using organization-specific telemetry subject to privacy and security constraints. In the first phase, public datasets such as UNSW-NB15, CICIDS2017, ToN-IoT, and BoT-IoT can be used to evaluate supervised and unsupervised detection models. The use of multiple datasets is recommended because each dataset represents different environments and attack categories. UNSW-NB15 includes nine attack types and modern traffic conditions [6]. CICIDS2017 supports intrusion traffic characterization [7]. ToN-IoT supports Industry 4.0, IoT, and IIoT scenarios [8]. BoT-IoT provides botnet traffic generated in a cyber range [9].

In the second phase, a synthetic CTI fusion environment can be constructed by combining dataset-derived events with simulated threat intelligence feeds. Indicators can be mapped to ATT&CK tactics and techniques to evaluate whether the framework improves prioritization beyond raw detection accuracy. In the third phase, organizations can validate the framework with internal telemetry in a controlled environment while maintaining operational safety and privacy controls.

Table 2: Proposed Validation Workflow for Framework Evaluation

Step	Description	Expected Output
1. Dataset Selection	Select datasets representing enterprise, IoT, and industrial scenarios.	Training and testing data.
2. Preprocessing	Clean records, normalize fields, encode labels, and handle imbalance.	Model-ready dataset.
3. Feature Engineering	Extract traffic, behavior, temporal, and threat-intelligence features.	Predictive feature matrix.
4. Model Training	Train baseline, ensemble, anomaly, and sequential models.	Candidate models.
5. Risk Scoring	Combine model predictions with simulated asset criticality.	Prioritized threat queue.
6. Explainability	Apply explainability methods to risk decisions.	Analyst-readable rationale.
7. Evaluation	Measure accuracy, precision, recall, F1-score, AUC, false positive rate, and latency.	Performance and usability findings.

6. Comparative Analysis

Table 3 compares the proposed framework with representative categories of existing approaches. The comparison is conceptual and literature-based; it does not claim experimental superiority without future validation.

Table 3: Comparison of Existing Approaches and Proposed Framework

Approach	Strengths	Limitations	Proposed Framework Improvement
Signature-based IDS	Effective for known threats and simple to operationalize.	Weak against zero-day and evolving attack patterns.	Adds ML anomaly detection and predictive risk scoring.
Standalone ML intrusion detection	Can improve classification accuracy on benchmark datasets.	Often lacks CTI, asset criticality, explainability, and response integration.	Integrates prediction with intelligence fusion, criticality, and response orchestration.
Threat intelligence feeds	Provide current indicators and sector-relevant intelligence.	Often reactive, noisy, and not automatically prioritized.	Uses AI and asset context to convert indicators into predictive risk.
SIEM correlation rules	Useful for compliance, alerting, and centralized visibility.	Rule maintenance is difficult and false positives can be high.	Learns dynamic patterns and ranks threats based on likelihood and impact.
SOAR playbooks	Can automate repetitive response actions.	Requires accurate triggers and contextual prioritization.	Feeds SOAR with explainable, risk-ranked predictions.
Proposed framework	Combines ML, CTI, ATT&CK mapping, asset criticality, explainability, and orchestration.	Requires data quality, governance, validation, and careful deployment.	Provides a proactive, modular model for critical infrastructure protection.

7. Discussion

The proposed framework supports a shift from reactive cybersecurity to proactive risk anticipation. Its key advantage is not a single algorithm but the integration of multiple components that are typically separated in operational environments. Security teams often maintain SIEM systems, threat feeds, vulnerability scanners, endpoint tools, and asset inventories as separate workflows. The proposed framework connects these sources into a unified predictive pipeline. This integration is especially important for critical infrastructure because technical alerts must be interpreted in terms of operational impact.

A critical design principle is explainability. In high-impact environments, a model-generated risk score is insufficient unless analysts can understand the basis for the recommendation. Explainable outputs may include the features that influenced the score, mapped ATT&CK techniques, affected assets, relevant vulnerabilities, and recommended mitigation actions. This supports analyst trust and reduces the likelihood of inappropriate automated actions in operational technology environments. Another important consideration is data governance. Critical infrastructure telemetry can be sensitive and may include operational details that cannot be freely shared. Federated learning, privacy-preserving analytics, and controlled information sharing may help organizations collaborate without exposing sensitive operational data. However, such approaches must be carefully evaluated because model updates can also leak information if not properly protected.

The framework should also be deployed with human oversight. Full automation may be suitable for low-risk actions such as enriching alerts, opening tickets, or blocking known malicious domains. High-impact actions such as isolating OT assets, shutting down services, or modifying industrial control configurations should remain analyst-approved unless emergency procedures and safety validation are in place.

8. Implementation Considerations

Implementation should begin with a limited pilot rather than full-scale deployment. Organizations can start by integrating SIEM data, vulnerability information, and a limited set of CTI feeds. Asset inventory and business criticality mapping should be established early because risk scoring depends heavily on asset context. Model governance processes should define training data sources, validation

frequency, false positive review, explainability requirements, and rollback procedures. Security teams should align the framework with NIST CSF 2.0 functions, CISA CPGs, and existing incident response playbooks. The Govern function in CSF 2.0 is particularly relevant because AI-enabled security requires accountability, oversight, and defined risk management practices [1]. The framework should also map predictions to ATT&CK tactics and techniques to support a common language between analysts, executives, and technical teams [4].

9. Limitations

This paper presents a framework and evaluation plan rather than a completed production deployment. Therefore, the paper does not claim measured improvements in detection accuracy, false positive reduction, or response time. Such claims should only be made after experiments are conducted using selected datasets and validated environments. Public datasets may not perfectly represent all critical infrastructure scenarios, and operational telemetry may be difficult to obtain because of confidentiality and safety concerns. AI models may also be vulnerable to adversarial manipulation, data poisoning, model drift, and overfitting. Future work must address these issues through robust validation, continuous monitoring, and governance controls.

10. Future Work

Future work should implement and test the proposed framework using public datasets and, where possible, controlled critical infrastructure testbeds. Experimental evaluation should compare baseline ML models, ensemble models, anomaly detection approaches, and sequential models. Additional research should explore federated learning for multi-operator collaboration, explainable AI for analyst decision support, adversarial resilience, integration with zero trust architectures, and safe automation for OT environments. A future empirical paper can report model performance using accuracy, precision, recall, F1-score, ROC-AUC, false positive rate, mean time to detect, and prioritization effectiveness.

11. Conclusion

Critical infrastructure systems face increasingly sophisticated cyber threats that challenge traditional reactive security approaches. This paper proposed an AI-driven Predictive Cyber Threat Intelligence Framework designed to integrate multi-source intelligence, machine learning, behavioral analytics, criticality-aware risk scoring, explainability, and response orchestration. The framework is aligned with established cybersecurity guidance and addresses the need for proactive threat anticipation in high-impact environments. By connecting predictive analytics with operational context, the proposed model can help organizations prioritize risk, improve situational awareness, reduce alert fatigue, and strengthen resilience. Future experimental validation is required before quantitative performance claims can be made, but the framework provides a practical foundation for advancing predictive cybersecurity in critical infrastructure protection.

Author Biography

Dr. Naga Venkata Aswini Pavan Kumar Inguva is a technology professional and researcher with interests in cybersecurity, artificial intelligence, machine learning, data security, IoT security, and enterprise software systems. His work focuses on AI-driven cybersecurity, cyber threat detection, and secure system development for enterprise and critical infrastructure environments.

References

[1] National Institute of Standards and Technology, The NIST Cybersecurity Framework (CSF) 2.0, NIST Cybersecurity White Paper 29, Feb. 2024.
[2] Cybersecurity and Infrastructure Security Agency, Cross-Sector Cybersecurity Performance Goals, CISA, 2024.
[3] C. Johnson, L. Badger, D. Waltermire, J. Snyder, and C. Skorupka, Guide to Cyber Threat Information Sharing, NIST Special Publication 800-150, Oct. 2016.
[4] MITRE, MITRE ATT&CK Enterprise Matrix, MITRE Corporation, 2026.
[5] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cyber security intrusion detection," IEEE Communications Surveys & Tutorials, vol. 18, no. 2, pp. 1153–1176, 2016.
[6] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection systems," in Proc. Military Communications and Information Systems Conference, 2015.
[7] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in Proc. 4th Int. Conf. Information Systems Security and Privacy, 2018, pp. 108–116.
[8] University of New South Wales Canberra Cyber, The ToN_IoT Datasets, UNSW Canberra, 2021.
[9] University of New South Wales Canberra Cyber, The BoT-IoT Dataset, UNSW Canberra, 2021.
[10] K. J. Raval et al., "A survey on safeguarding critical infrastructures: Attacks, AI-based defenses, and future directions," Computer Science Review, 2024.
[11] U.S. Department of Energy, Potential Benefits and Risks of Artificial Intelligence for Critical Energy Infrastructure, Apr. 2024.
[12] Center for Security and Emerging Technology, Securing Critical Infrastructure in the Age of AI, Georgetown University, 2024.
[13] G. Rjoub et al., "A survey on explainable artificial intelligence for cybersecurity," IEEE Transactions on Network and Service Management, 2023.
[14] Reuters, "US Homeland Security names AI safety, security advisory board," Apr. 2024.
[15] Reuters, "Biden signs new memo to boost security of US critical infrastructure," Apr. 2024.
[16] F. Alwahedi et al., "Machine learning techniques for IoT security: Current research and future vision," Internet of Things and Cyber-Physical Systems, 2024.
[17] A. Arif et al., "An overview of cyber threats generated by AI," International Journal of Modern Data Science and Analytics, 2024.

[18] Microsoft, "Staying ahead of threat actors in the age of AI," Microsoft Security Blog, 2024.
[19] OpenAI, "Disrupting malicious uses of AI by state-affiliated threat actors," OpenAI, 2024.
[20] IEEE Communications Society, "Cybersecurity for Critical Infrastructure Systems," IEEE Communications Magazine Feature Topic Call for Papers, 2026.