

AI-Driven Predictive Cyber Threat Intelligence Framework for Securing Industrial Digital Infrastructure

Nitin Bodade^{1*}

¹Digital Division, Caterpillar Inc., United States of America

Corresponding Author: Nitin Bodade^{1*}

Publication Date: 2026/07/15

Abstract: The digitalisation of industrial environments and the increasing number of Industrial Internet of Things (IIoT) devices have completely increased the attack surface of critical manufacturing and operational technology (OT) systems. Current signature-based, "reactive" cybersecurity models are clearly failing to keep up with the sophistication and speed of today's Advanced Persistent Threats (APTs), ransomware-as-a-service (RaaS) operations and supply-chain attacks on industrial digital systems. Cyber Threat Intelligence (CTI) is now a strategic field and discipline for predicting adversarial actions, but the frameworks in use are largely tactical, siloed, and reactive, and have very limited ability to conduct real-time predictive analytics in industrial environments.

This paper tackles the identified gap by proposing a new conceptual framework called Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) Framework specifically designed to facilitate predictive, adaptive, and automated methods of threat intelligence for industrial digital infrastructure. This research uses Design Science Research Methodology (DSRM) that includes Systematic Literature Review (SLR), Knowledge Elicitation from experts and Structured Conceptual Design in order to create the framework artefact. The AIPCTI Framework comprises six interdependent layers: Threat Data Acquisition, Threat Intelligence Fusion, AI Analytics Engine, Predictive Risk Assessment, Automated Response, and Governance and Compliance. These layers support a continuous, intelligence based cyber defence posture that is consistent with Zero Trust Architecture (ZTA) principles, as well as the MITRE ATT&CK for ICS knowledge base.

The framework has been developed to incorporate feedback from experts in the field of ICS/OT security and validated using structured scenario-based reasoning using three representative attack patterns: manufacturing ransomware, energy-sector APT intrusion and IIoT firmware exploitation. The analysis at the architecture level depicts examples of how the layered design of AIPCTI would be expected to close certain detection and response gaps identified by indicator-based CTI platforms and IT-focused SOAR solutions, such as providing the ability for AIPCTI to anticipate attack techniques before they are executed, as well as to limit automated response with safety logic specific to OT. These findings are expressed as a design stage evaluation and not as an actual performance while in operation: the framework is not yet in place, nor is it deployed in a live industrial setting, nor is it claimed to be able to improve the detection rate or response time (even though it can certainly do that). The research provides a theoretically informed conceptual framework for cybersecurity architects, offices of industrial CISO's, and policy makers, as well as a well-defined research agenda for empirical testing of operational claims.

Keywords: Cyber Threat Intelligence, Predictive Threat Analytics, Industrial IoT Security, Artificial Intelligence, Design Science Research, MITRE ATT&CK for ICS, Zero Trust Architecture, and Operational Technology Security.

How to Cite: Nitin Bodade (2026) AI-Driven Predictive Cyber Threat Intelligence Framework for Securing Industrial Digital Infrastructure. *International Journal of Innovative Science and Research Technology*, 11(7), 104-130.
<https://doi.org/10.38124/ijisrt/26jul085>

I. INTRODUCTION

➤ Digital Transformation in Industrial Environments

The widespread adoption of Industry 4.0 technologies has rapidly been bringing together Information Technology

(IT) and Operational Technology (OT) in various critical infrastructure sectors, including manufacturing, energy, transportation, and beyond. The wide adoption of interconnected cyber-physical systems, cloud computing, edge intelligence and the Industrial Internet of Things (IIoT)

by industrial organisations can enhance productivity, operational efficiency, predictive maintenance and real-time decision making [1,2]. There are recent reports from industry indicating that the number of connected devices in the industrial Internet of Things (IIoT) will reach more than 40 billion devices by the next few years, significantly increasing the digital footprint of industrial environments [3].

This transformation brings many operational advantages, but also increases the cyber attack surface. Many industrial facilities are still using legacy SCADA (Supervisory Control and Data Acquisition) and DCS (Distributed Control System) and PLC (Programmable Logic Controller) systems built with reliability and availability in mind. Such systems are frequently based on a proprietary communication protocol, have weak authentication mechanisms, and have long service lifecycles, making security patching and technology replacement difficult [4,5]. Today's industrial infrastructures are therefore complex, comprising a mix of legacy operational technologies, digital services and the Internet, adding opportunities for advanced cyber attacks.

➤ *Evolution of the Industrial Cybersecurity Threat Landscape*

Industrial cyber threats have evolved from opportunistic attacks into highly coordinated campaigns conducted by financially motivated groups and nation-state actors. Incidents such as TRITON/TRISIS, the Colonial Pipeline ransomware attack, and the compromise of the Oldsmar water treatment facility demonstrate that cyber attacks against industrial systems can produce operational disruption, financial loss, environmental damage, and threats to public safety [6,7]. Recent threat intelligence reports indicate that manufacturing remains one of the most frequently targeted sectors due to its dependence on interconnected operational technologies and continuous production processes [8].

Unlike conventional cyber-attacks against enterprise IT systems, Advanced Persistent Threats (APTs) targeting industrial environments typically employ long-term reconnaissance, credential abuse, Living-Off-the-Land (LOTL) techniques, lateral movement, and multi-stage attack strategies designed to remain undetected for extended periods [9,10]. The increasing sophistication of groups such as APT41, Lazarus, Sandworm, and other advanced adversaries demonstrates that industrial cybersecurity has become a strategic concern extending beyond traditional information security.

➤ *Limitations of Existing Cyber Threat Intelligence Approaches*

Despite significant advances in cyber defence technologies, many industrial organisations continue to rely on predominantly reactive security mechanisms, including signature-based intrusion detection systems, static firewall configurations, vulnerability scanning, and rule-based monitoring [12]. Although these technologies remain valuable, their effectiveness is limited when confronting zero-day attacks, previously unseen attack patterns, and

advanced adversaries that deliberately evade conventional detection mechanisms.

Cyber Threat Intelligence (CTI) has emerged as an important discipline for improving organisational cyber resilience by collecting, analysing, and disseminating information regarding threat actors, attack techniques, indicators of compromise, and emerging vulnerabilities [13]. Existing CTI platforms, including STIX/TAXII, MISP, and OpenCTI, have significantly improved information sharing and collaborative intelligence generation. However, these platforms primarily support threat information management and exchange rather than predictive decision-making. In many industrial deployments, CTI remains disconnected from operational technologies, automated response mechanisms, and real-time risk assessment processes [14,15].

➤ *Research Gap*

Existing research has produced important advances in cyber threat intelligence, machine learning-based intrusion detection, industrial anomaly detection, and cyber-physical system security. Nevertheless, three important research gaps remain.

First, most CTI platforms concentrate on threat collection, correlation, and information sharing, while providing limited capabilities for predictive threat forecasting before attack execution.

Second, many artificial intelligence-based intrusion detection approaches have been developed and evaluated using enterprise IT environments, making direct application within industrial operational technology environments challenging because of protocol diversity, deterministic communication patterns, limited labelled datasets, and strict operational reliability requirements [16].

Third, existing studies generally address threat intelligence, predictive analytics, response orchestration, governance, and compliance as separate research problems rather than integrated components of a unified industrial cybersecurity architecture. Consequently, there remains a need for a comprehensive framework capable of supporting predictive cyber threat intelligence throughout the complete threat management lifecycle while considering the operational characteristics of industrial digital infrastructure.

➤ *Novelty Statement*

Unlike existing CTI frameworks that primarily facilitate information sharing and post-event analysis, the proposed Artificial Intelligence Powered Predictive Cyber Threat Intelligence (AIPCTI) Framework integrates threat intelligence acquisition, AI-driven predictive analytics, contextual risk assessment, response orchestration, and governance support within a unified conceptual architecture specifically designed for industrial environments.

The principal novelty of this research lies not in introducing individual cybersecurity technologies, but in developing a structured design framework that systematically

combines these complementary capabilities using a Design Science Research Methodology (DSRM). The framework provides a coherent architectural model that supports predictive cyber defence while considering the operational constraints of Industrial Internet of Things (IIoT), Industrial Control Systems (ICS), and Operational Technology (OT) environments.

➤ *Research Objectives and Research Questions*

The primary objective of this study is to design and evaluate a conceptual framework that supports predictive cyber threat intelligence for industrial digital infrastructure.

• *The Study Addresses the Following Research Questions:*

- ✓ RQ1. What limitations prevent existing CTI frameworks from effectively supporting predictive cyber defence in industrial OT and ICS environments?
- ✓ RQ2. How can artificial intelligence techniques be systematically integrated into industrial CTI architectures to improve predictive threat intelligence capabilities?
- ✓ RQ3. What architectural components are required to support a comprehensive AI-driven predictive CTI framework for industrial digital infrastructure?
- ✓ RQ4. How does the proposed framework compare conceptually with existing CTI approaches in terms of predictive capability, architectural integration, and operational applicability?

➤ *Research Contributions*

This study makes four principal contributions.

First, it synthesises recent literature on industrial cybersecurity, cyber threat intelligence, artificial intelligence, and operational technology security to identify unresolved research challenges.

Second, it proposes the AIPCTI Framework, a structured conceptual architecture that integrates predictive analytics, cyber threat intelligence, automated response coordination, and governance considerations within a unified industrial cybersecurity model.

Third, it applies the Design Science Research Methodology (DSRM) to guide the systematic development and evaluation of the proposed framework, thereby providing methodological transparency for future research.

Finally, the study identifies practical design considerations that may assist researchers and industrial practitioners in developing future predictive cyber defence solutions for Industry 4.0 environments.

➤ *Paper Organization*

The remainder of this paper is organised as follows. Section 2 reviews the relevant literature on cyber threat intelligence, industrial cybersecurity, artificial intelligence, and operational technology security. Section 3 presents the identified research gaps. Section 4 describes the research methodology. Section 5 introduces the proposed AIPCTI Framework. Section 6 presents the framework evaluation and

validation. Section 7 discusses the findings. Section 8 outlines the theoretical and practical contributions. Sections 9–11 discuss limitations, future research directions, and conclusions.

II. LITERATURE REVIEW

➤ *Cyber Threat Intelligence (CTI)*

Cyber Threat Intelligence (CTI) has evolved from a reactive mechanism for sharing indicators of compromise into a strategic decision-support capability that enables organisations to anticipate, understand, and mitigate increasingly sophisticated cyber threats. The definition proposed by the OASIS Open describes CTI as evidence-based knowledge encompassing context, mechanisms, indicators, implications, and actionable recommendations that support informed cybersecurity decision-making. This definition highlights an important transition in cybersecurity practice: intelligence is no longer limited to collecting technical indicators but now integrates contextual information that enables proactive defence and organisational resilience.

Contemporary CTI programmes are generally organised into four complementary intelligence layers: strategic, operational, tactical, and technical. Strategic intelligence supports executive decision-making and long-term risk management, while operational intelligence focuses on adversary campaigns and attack planning. Tactical intelligence characterises attacker tactics, techniques, and procedures (TTPs), whereas technical intelligence concentrates on machine-readable indicators such as malicious IP addresses, domains, hashes, and malware signatures. Together, these intelligence layers create a comprehensive understanding of the evolving threat landscape, allowing organisations to respond at multiple organisational levels rather than relying solely on reactive detection mechanisms.

To facilitate automated intelligence exchange, the Structured Threat Information Expression (STIX) language and the Trusted Automated Exchange of Intelligence Information (TAXII) protocol have become the dominant standards for representing and sharing cyber threat intelligence. Their widespread adoption has significantly improved interoperability among security platforms, enabling security information and event management (SIEM) systems, security orchestration platforms, and endpoint protection solutions to exchange threat information efficiently. Nevertheless, these standards primarily address information representation and distribution rather than intelligence generation or predictive reasoning.

Despite substantial advances in CTI platforms, several fundamental limitations remain unresolved. Recent systematic reviews consistently demonstrate that most CTI systems remain heavily dependent on manually curated threat feeds and static indicators, limiting their ability to detect rapidly evolving attack campaigns. Schlette et al. identified persistent deficiencies related to intelligence quality, incomplete contextual information, heterogeneous data

formats, and limited automation across existing CTI ecosystems.

Similarly, Tounsi and Rais observed that many organisations continue to operate at a datacentric maturity level, where raw indicators are collected but rarely transformed into actionable intelligence capable of supporting strategic cyber defence. These findings suggest that although large volumes of threat data are available, relatively little emphasis has been placed on extracting predictive knowledge from these datasets.

An additional challenge concerns the limited applicability of current CTI frameworks within industrial environments. Public intelligence repositories predominantly focus on enterprise information technology (IT) infrastructures and provide relatively few indicators describing attacks against operational technology (OT), industrial control systems (ICS), or Industrial Internet of Things (IIoT) assets. Wagner et al. demonstrated that integrating CTI feeds with enterprise SIEM platforms significantly reduces threat detection time; however, they also reported that industrial protocol-specific indicators remain substantially underrepresented within publicly available intelligence sources. Consequently, manufacturing organisations frequently operate with incomplete visibility of emerging threats targeting industrial communication protocols such as Modbus TCP, DNP3, PROFINET, and OPC UA.

Another limitation is the predominantly reactive nature of existing CTI implementations. Most platforms disseminate information after attacks have already been identified, requiring defenders to respond retrospectively rather than anticipate adversarial behaviour. This reactive paradigm is increasingly inadequate against advanced persistent threats (APTs), which employ stealth, persistence, and multi-stage attack strategies that evolve continuously throughout an attack lifecycle. Modern industrial environments therefore require intelligence capabilities capable of forecasting adversarial actions before critical assets are compromised.

Overall, the current body of literature demonstrates that although CTI has matured considerably with respect to intelligence collection, standardisation, and information sharing, comparatively little progress has been achieved in predictive intelligence generation for industrial digital infrastructures. Existing approaches remain constrained by limited automation, insufficient contextual reasoning, weak integration with AI-driven analytics, and inadequate representation of OT-specific threats. These limitations establish the need for a predictive cyber threat intelligence framework capable of combining heterogeneous intelligence sources with advanced artificial intelligence techniques to enhance proactive cyber defence within Industry 4.0 environments.

➤ *Predictive Threat Intelligence*

Predictive Cyber Threat Intelligence (PCTI) represents a significant evolution beyond traditional intelligence-

sharing approaches by shifting the primary objective from retrospective threat analysis toward anticipating future adversarial behaviour. Rather than merely identifying previously observed indicators of compromise, predictive intelligence seeks to estimate the likelihood, timing, and progression of cyberattacks before operational disruption occurs. This transition has become increasingly important as cyber adversaries adopt adaptive attack strategies that exploit the speed and complexity of modern digital infrastructures faster than conventional signature-based defence mechanisms can respond.

The emergence of artificial intelligence and machine learning has substantially accelerated research into predictive threat intelligence. Unlike conventional CTI systems that rely primarily on historical attack signatures, predictive approaches analyse behavioural patterns, attack sequences, and threat relationships to infer probable future attack paths. Mavroeidis and Bromander introduced one of the earliest conceptual models linking cyber kill chain analysis with predictive intelligence, arguing that attacker behaviour exhibits recurring patterns that can be modelled probabilistically. Their work established an important theoretical foundation by demonstrating that threat intelligence should support anticipation rather than simple incident reporting. However, the framework remained largely conceptual and did not provide implementation guidance or empirical validation.

Subsequent studies have increasingly incorporated advanced machine learning techniques to predict attacker behaviour. Graph neural networks have been employed to model complex relationships between attack stages, enabling prediction of lateral movement and privilege escalation within enterprise environments. Shin et al. demonstrated promising predictive accuracy using graph-based representations of attack progression; however, their evaluation relied exclusively on enterprise IT datasets and did not consider industrial communication protocols, physical process constraints, or safety-critical operational requirements that characterise OT environments. Consequently, the generalisability of these findings to industrial systems remains uncertain.

Deep learning approaches have also gained considerable attention for modelling temporal attack behaviour. Bao et al. proposed a Long Short-Term Memory (LSTM)-based prediction framework trained on MITRE ATT&CK sequences to forecast subsequent attacker techniques. Similarly, Deliu et al. applied association rule mining to identify recurring attack patterns within threat intelligence feeds, enabling probabilistic prediction of campaign evolution. Although these studies demonstrate the potential of sequential learning models, they remain primarily focused on enterprise networks and publicly available ATT&CK datasets, with relatively little consideration of industrial attack scenarios involving programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) systems, or cyber-physical processes.

A recurring limitation across predictive threat intelligence research is the narrow representation of industrial operational technology. Existing prediction models generally assume homogeneous enterprise environments with abundant labelled datasets, whereas industrial infrastructures consist of heterogeneous legacy devices, proprietary communication protocols, sparse attack observations, and stringent availability requirements. These characteristics significantly complicate the application of conventional AI models and reduce their practical effectiveness within manufacturing and critical infrastructure sectors. Furthermore, most predictive frameworks treat cyber intelligence, anomaly detection, and risk assessment as independent processes rather than integrated components of a unified decision-support architecture.

Another important limitation concerns explainability and operational trust. Industrial cybersecurity practitioners often require transparent reasoning to justify automated defensive actions because incorrect predictions may interrupt production processes or compromise operational safety. Many deep learning-based predictive models function as black-box systems, making it difficult for analysts to interpret prediction outcomes or validate decision logic. This limitation has become increasingly significant as regulatory frameworks emphasise accountability and explainable artificial intelligence within critical infrastructure protection.

Collectively, the existing literature confirms that predictive threat intelligence represents one of the most promising directions in modern cybersecurity research. Nevertheless, substantial research gaps remain regarding the integration of AI-driven prediction, industrial threat intelligence, OT-specific attack modelling, automated response orchestration, and governance mechanisms within a unified framework. Addressing these deficiencies requires an architecture capable of combining heterogeneous threat intelligence, behavioural analytics, predictive machine learning, and industrial risk assessment into a coherent cybersecurity capability. This unresolved gap provides the primary motivation for the Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) framework proposed in this study.

➤ *Machine Learning for Cybersecurity*

Machine learning (ML) has become one of the most influential technologies in modern cybersecurity due to its ability to analyse large volumes of heterogeneous security data and identify malicious behaviour that often escapes traditional signature-based detection systems. Unlike rule-based security mechanisms, ML models continuously learn behavioural patterns from historical observations, enabling adaptive detection of previously unseen attacks, insider threats, and sophisticated multi-stage intrusions. As cyber adversaries increasingly employ polymorphic malware, zero-day exploits, and advanced persistent threats (APTs), machine learning has emerged as a critical component of next-generation cyber defence strategies.

Existing research generally categorises cybersecurity-oriented machine learning techniques into supervised,

unsupervised, semi-supervised, and reinforcement learning paradigms. Supervised learning algorithms, including Random Forests, Support Vector Machines (SVMs), Gradient Boosting Machines, and Extreme Gradient Boosting (XGBoost), have demonstrated high classification accuracy when labelled attack datasets are available. Conversely, unsupervised approaches such as Isolation Forests, K-Means clustering, Autoencoders, and One-Class SVMs are particularly valuable for anomaly detection in environments where labelled attack samples are scarce. Reinforcement learning has also gained increasing attention for adaptive intrusion response, autonomous threat mitigation, and dynamic cyber defence decision-making.

Sarker et al. conducted one of the most comprehensive surveys of machine learning applications in cybersecurity, demonstrating that ensemble learning techniques consistently outperform single classifiers across multiple benchmark intrusion detection datasets, including NSL-KDD, CICIDS-2017, UNSW-NB15, and CICDDoS2019. Their findings further revealed that ensemble models provide improved robustness against class imbalance while reducing false positive rates compared with conventional statistical detection methods. Nevertheless, the majority of surveyed studies focused primarily on enterprise IT environments, with relatively limited consideration of industrial control systems and operational technology infrastructures.

The application of machine learning within industrial environments introduces significantly different technical challenges. Unlike enterprise networks, industrial systems generate deterministic traffic patterns, strict timing constraints, proprietary communication protocols, and highly imbalanced attack datasets. Consequently, algorithms that perform well within traditional enterprise infrastructures frequently exhibit reduced effectiveness when deployed in manufacturing or critical infrastructure settings.

Radoglou-Grammatikis et al. evaluated multiple machine learning algorithms using the Secure Water Treatment (SWaT) and BATADAL industrial control system datasets, concluding that Isolation Forests and One-Class SVMs achieved superior anomaly detection performance under realistic industrial operating conditions. Their work demonstrated the importance of unsupervised learning for identifying previously unknown attacks; however, the proposed models remained limited to anomaly detection and provided little capability for predicting future attacker behaviour or supporting proactive threat intelligence.

Recent studies have explored distributed machine learning architectures to address privacy concerns associated with collaborative cybersecurity analytics. Conti et al. proposed a federated learning framework enabling multiple industrial organisations to train intrusion detection models collaboratively without exchanging raw operational data. This approach significantly improves data privacy while facilitating collective learning across geographically distributed manufacturing facilities. Nevertheless, federated learning primarily addresses model training rather than threat intelligence generation and therefore does not provide

mechanisms for correlating external cyber threat intelligence with internal operational telemetry.

Deep learning architectures have further expanded the capabilities of cybersecurity analytics by capturing complex temporal and contextual relationships within cyberattack sequences. Lin et al. demonstrated that transformer-based neural networks outperform conventional recurrent neural networks when analysing long-term dependencies within industrial process data. Their attention-based architecture improved multi-stage attack detection by effectively modelling sequential interactions among cyber-physical events. Despite these promising results, transformer models typically require substantial computational resources and extensive labelled datasets, creating practical deployment challenges for many industrial environments where computational resources remain constrained.

An emerging concern within machine learning-driven cybersecurity is explainability. Industrial operators often require transparent and interpretable decision-making because erroneous automated responses may disrupt production processes or compromise operational safety. Many state-of-the-art deep learning models function as black-box systems, limiting analyst confidence and complicating regulatory compliance.

Consequently, explainable artificial intelligence (XAI) has become an increasingly important research direction for industrial cybersecurity, particularly in safety-critical sectors where accountability and human oversight remain essential. Collectively, existing machine learning research demonstrates considerable progress in cyberattack detection, anomaly identification, and behavioural modelling. However, current approaches remain largely detection-oriented, focusing on recognising malicious activities after they have already manifested. Comparatively little attention has been devoted to integrating machine learning with cyber threat intelligence, predictive risk assessment, automated response orchestration, and governance mechanisms within a unified industrial cybersecurity architecture. Addressing this limitation requires a comprehensive framework capable of transforming machine learning outputs into actionable predictive intelligence that supports proactive cyber defence across industrial digital infrastructures.

➤ *Artificial Intelligence in Industrial Security*

Artificial intelligence (AI) has transformed industrial cybersecurity by extending security operations beyond conventional intrusion detection toward intelligent threat prediction, autonomous incident response, behavioural analytics, and cyber resilience. The convergence of Industry 4.0 technologies—including Industrial Internet of Things (IIoT), cyber-physical systems, cloud manufacturing and edge computing—has significantly increased the complexity of industrial attack surfaces. Consequently, AI has become indispensable for processing vast volumes of heterogeneous operational data and supporting real-time cybersecurity decision-making within modern manufacturing environments.

Unlike traditional machine learning approaches that focus primarily on classification or anomaly detection, AI-based cybersecurity systems increasingly incorporate knowledge representation, reasoning engines, graph analytics, natural language processing, and autonomous decision support. These capabilities enable security platforms to correlate information originating from multiple sources, including threat intelligence feeds, network telemetry, vulnerability databases, endpoint logs, and operational process data. By integrating these diverse information sources, AI systems can construct contextual representations of cyber threats that support more informed and proactive defensive strategies.

Khan et al. examined AI deployment across operational technology environments and identified several persistent implementation challenges. Their analysis highlighted legacy industrial infrastructure, limited availability of labelled OT attack datasets, computational resource constraints, and stringent explainability requirements as major obstacles limiting widespread AI adoption. Unlike enterprise environments, industrial organisations frequently operate legacy programmable logic controllers (PLCs), distributed control systems (DCS), and supervisory control and data acquisition (SCADA) platforms that were not originally designed to support modern AI-driven cybersecurity technologies. These technical constraints significantly influence the feasibility of deploying computationally intensive AI algorithms within operational environments.

Graph-based artificial intelligence has recently demonstrated considerable potential for modelling complex relationships among cyber events. Testart et al. illustrated how graph analytics applied to Border Gateway Protocol (BGP) routing information could identify coordinated reconnaissance activities and provide early warning of large-scale network attacks. Although their work focused on Internet routing infrastructure rather than industrial systems, it demonstrates the broader value of graph reasoning for identifying hidden relationships among seemingly independent security events. Similar graph-based approaches may be adapted to industrial cyber threat intelligence by modelling interactions among assets, vulnerabilities, attacker behaviours, and operational dependencies.

The rapid emergence of generative artificial intelligence has introduced both unprecedented opportunities and significant cybersecurity risks. Gupta et al. examined the dual-use nature of large language models (LLMs), demonstrating their ability to automate vulnerability analysis, malware reverse engineering, incident reporting, and threat intelligence summarisation while simultaneously enabling adversaries to generate convincing phishing campaigns, automate malicious code development, and accelerate social engineering attacks. These findings illustrate that generative AI should be regarded not only as a defensive technology but also as a rapidly evolving component of the cyber threat landscape itself.

Recent research has further explored digital twin technologies as secure environments for AI-assisted cybersecurity experimentation. Fraunholz et al. proposed the

integration of digital twins into industrial security operations to enable attack simulation, vulnerability assessment, and incident response evaluation without affecting live production environments. Digital twins provide a valuable mechanism for validating AI models under controlled conditions while reducing operational risks associated with testing defensive strategies on critical industrial infrastructure. Nevertheless, current digital twin implementations remain largely experimental and have not yet been fully integrated with predictive cyber threat intelligence or continuous operational decision support.

Despite rapid advances in artificial intelligence, several fundamental challenges remain unresolved. Existing AI solutions frequently address isolated cybersecurity functions such as anomaly detection, malware classification, or threat reporting rather than supporting end-to-end cyber defence workflows. Furthermore, many AI systems lack effective integration with cyber threat intelligence platforms, industrial risk assessment methodologies, governance frameworks, and automated response mechanisms. This fragmentation limits their ability to provide comprehensive situational awareness across complex industrial ecosystems.

The literature therefore suggests that future industrial cybersecurity solutions should move beyond isolated AI applications toward integrated intelligence-driven architectures capable of combining predictive analytics, contextual threat intelligence, explainable decision support, automated orchestration, and governance-aware risk management. These unresolved challenges provide a strong rationale for the proposed Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) framework, which seeks to unify these capabilities within a single predictive cybersecurity architecture specifically designed for industrial digital infrastructure.

➤ *Industrial Control Systems Security*

Industrial Control Systems (ICS) constitute the operational backbone of critical infrastructure sectors, including manufacturing, energy, transportation, water treatment, and chemical processing. Unlike conventional enterprise information technology (IT) systems, ICS environments are designed primarily to ensure continuous availability, operational safety, and deterministic process control. Consequently, cybersecurity mechanisms deployed within these environments must preserve operational continuity while simultaneously defending against increasingly sophisticated cyber threats. The rapid convergence of IT and operational technology (OT) networks under Industry 4.0 has significantly expanded the cyberattack surface, exposing legacy industrial assets to threats that were previously confined to enterprise environments.

Historically, ICS security relied on network isolation and proprietary communication protocols to minimise cyber risk. However, the adoption of Industrial Internet of Things (IIoT) devices, cloud-enabled monitoring platforms, remote maintenance services, and interconnected production systems has fundamentally altered this security model. Modern industrial environments now require continuous

communication between business systems and operational assets, creating new pathways through which adversaries can compromise industrial processes. As demonstrated by several high-profile cyber incidents targeting critical infrastructure, traditional perimeter-based security approaches are no longer sufficient to protect interconnected industrial ecosystems.

To address these challenges, internationally recognised cybersecurity standards have been developed to guide the protection of industrial control environments. The NIST Special Publication 800-82 and the IEC 62443 series provide comprehensive recommendations covering asset identification, network segmentation, secure system architecture, access control, incident response, and lifecycle security management. These frameworks have become widely adopted benchmarks for industrial cybersecurity governance. Nevertheless, they primarily provide implementation guidance rather than mechanisms for predictive threat analysis or intelligent decision support, leaving organisations responsible for integrating advanced detection and analytical capabilities into their operational environments.

Empirical studies consistently demonstrate that many industrial organisations continue to experience substantial cybersecurity maturity gaps despite the availability of these standards. Recent industry assessments report that numerous organisations lack comprehensive OT asset inventories, maintain insufficient network segmentation between IT and OT environments, and operate without dedicated incident response procedures specifically designed for industrial systems. Such deficiencies significantly increase organisational exposure to ransomware campaigns, supply chain attacks, and advanced persistent threats targeting operational technology. These observations indicate that cybersecurity challenges in industrial environments arise not only from technical vulnerabilities but also from governance, visibility, and operational maturity limitations.

Another important research direction concerns the evolution of industrial network architectures. The traditional Purdue Enterprise Reference Architecture has served as the dominant framework for separating operational layers within industrial environments. However, contemporary manufacturing increasingly incorporates cloud computing, edge intelligence, digital twins, and remote operational services, reducing the effectiveness of rigid hierarchical segmentation models. Recent studies argue that static network zoning should be complemented by adaptive security mechanisms capable of continuously evaluating trust relationships, monitoring communication behaviour, and dynamically enforcing security policies according to operational risk conditions.

Threat modelling has similarly become an essential component of modern ICS security research. Attack graph analysis, cyber-physical dependency modelling, and vulnerability propagation techniques enable security analysts to identify critical attack paths before exploitation occurs. Casola et al. demonstrated that attack graph methodologies

provide valuable insights into vulnerable industrial assets and attack propagation patterns.

Nevertheless, existing threat modelling techniques generally operate independently of cyber threat intelligence platforms and artificial intelligence-based prediction engines. Consequently, they remain largely reactive, focusing on analysing known vulnerabilities rather than anticipating future attack progression based on evolving threat intelligence.

A persistent limitation across the existing literature is the fragmentation of industrial cybersecurity capabilities. Asset management, vulnerability assessment, threat intelligence, behavioural monitoring, anomaly detection, and incident response are frequently implemented as independent security functions with limited interoperability. This fragmentation restricts the development of comprehensive situational awareness and delays coordinated defensive responses during complex cyber incidents. Moreover, relatively few studies have investigated how predictive cyber threat intelligence can be integrated directly into operational decision-making within industrial environments.

Overall, current ICS security research provides robust guidance regarding governance, segmentation, and operational risk management but offers comparatively limited support for predictive intelligence, AI-driven threat forecasting, and automated response coordination. These unresolved challenges highlight the need for an integrated cybersecurity architecture capable of combining industrial threat intelligence, behavioural analytics, predictive machine learning, and governance-aware response mechanisms into a unified framework capable of enhancing cyber resilience across modern industrial digital infrastructures.

➤ *Industrial Internet of Things (IIoT) Security*

The Industrial Internet of Things (IIoT) has become a defining technological pillar of Industry 4.0 by enabling intelligent connectivity among industrial equipment, sensors, actuators, programmable logic controllers, edge devices, and cloud-based analytics platforms. Through continuous data collection and real-time communication, IIoT technologies enhance operational efficiency, predictive maintenance, production optimisation, and autonomous manufacturing processes. However, this unprecedented level of connectivity has simultaneously expanded the cyberattack surface, introducing numerous security challenges that extend beyond those traditionally encountered in enterprise information systems.

Unlike conventional IT infrastructures, IIoT ecosystems comprise highly heterogeneous devices with varying computational capabilities, communication protocols, and operational lifecycles. Many industrial endpoints were originally designed with minimal consideration for cybersecurity, prioritising reliability, deterministic communication, and long operational lifespans over modern security controls. Consequently, numerous legacy industrial devices continue to operate with weak authentication mechanisms, unencrypted communications, outdated

firmware, and limited capacity to support advanced security software. These characteristics significantly increase organisational exposure to cyberattacks targeting industrial operations.

Communication protocol diversity represents another major security challenge. Industrial environments simultaneously employ protocols such as MQTT, OPC UA, Modbus TCP, EtherNet/IP, PROFINET, DNP3, and CoAP, each exhibiting distinct security characteristics and implementation constraints. Sisinni et al. demonstrated that vulnerabilities associated with authentication, message integrity, replay protection, and secure key management vary considerably across these protocols, complicating the development of unified security architectures capable of protecting heterogeneous industrial ecosystems. This protocol diversity frequently limits interoperability among security platforms and reduces the effectiveness of conventional network monitoring approaches.

Authentication and access control remain fundamental concerns within IIoT security research. Lightweight cryptographic techniques have been proposed to accommodate resource-constrained industrial devices while preserving acceptable computational performance. Frustaci et al. concluded that elliptic curve cryptography provides an effective balance between computational efficiency and security strength for many constrained industrial environments. Nevertheless, cryptographic mechanisms alone cannot adequately defend against sophisticated adversaries employing behavioural attacks, insider threats, supply chain compromises, or coordinated multi-stage intrusion campaigns. Consequently, complementary intelligence-driven security mechanisms are increasingly required.

Recent research has explored distributed trust management mechanisms using blockchain technology to strengthen identity management, auditability, and data integrity within IIoT environments. Blockchain-based security architectures offer immutable audit trails, decentralised authentication, and resistance to single points of failure, thereby improving the trustworthiness of industrial communications. However, practical deployment remains constrained by latency, scalability, energy consumption, and integration complexity. These limitations reduce the suitability of blockchain as a comprehensive cybersecurity solution for time-sensitive industrial control processes.

Machine learning has similarly emerged as a promising technology for protecting IIoT infrastructures. Ensemble learning methods, deep neural networks, and anomaly detection algorithms have demonstrated encouraging performance in identifying malicious traffic across heterogeneous industrial communication protocols. Chaabouni et al. reported that ensemble-based intrusion detection systems consistently outperform individual classifiers when analysing diverse IIoT traffic patterns. Despite these advances, many proposed detection models remain computationally intensive, making deployment difficult within resource-constrained edge devices commonly found throughout industrial environments. Furthermore,

most existing approaches focus exclusively on anomaly detection rather than integrating predictive cyber threat intelligence, contextual risk assessment, and automated incident response.

An equally important challenge concerns lifecycle security. Industrial devices often remain operational for decades, substantially exceeding the support lifecycle of many cybersecurity technologies. As a result, organisations frequently operate mixed environments consisting of newly deployed smart devices alongside legacy equipment that cannot be easily upgraded or replaced. This coexistence introduces persistent vulnerabilities that require adaptive cybersecurity strategies capable of protecting heterogeneous infrastructures without disrupting operational continuity. Existing literature provides limited guidance regarding how predictive intelligence can support continuous risk assessment throughout these extended operational lifecycles.

Overall, current IIoT security research has made significant progress in addressing authentication, encryption, communication security, and intrusion detection. Nevertheless, relatively few studies have integrated these capabilities into a unified architecture capable of combining cyber threat intelligence, artificial intelligence, predictive analytics, behavioural monitoring, governance, and automated response within industrial environments. Addressing this fragmentation represents a critical research priority for Industry 4.0 cybersecurity. The proposed AIPCTI framework seeks to address this gap by providing an integrated intelligence-driven architecture that unifies predictive threat intelligence with AI-based analytics and industrial cybersecurity governance, thereby supporting proactive protection of modern industrial digital infrastructures.

➤ *Behavioral Threat Analytics and Threat Hunting*

The increasing sophistication of modern cyber adversaries has accelerated the transition from reactive security monitoring toward intelligence-driven proactive defence strategies. Among these strategies, User and Entity Behavior Analytics (UEBA) and cyber threat hunting have emerged as complementary disciplines that enable organisations to identify malicious activities before significant operational disruption occurs. Unlike traditional signature-based detection systems that rely primarily on known indicators of compromise, behavioural analytics focuses on identifying deviations from established patterns of normal activity, thereby improving the detection of insider threats, credential misuse, stealthy lateral movement, and advanced persistent threats (APTs).

Behavioural threat analytics applies statistical analysis, machine learning, and artificial intelligence to establish behavioural baselines for users, devices, applications, and industrial assets. By continuously monitoring deviations from expected operational behaviour, UEBA systems can detect subtle anomalies that often remain invisible to conventional security controls. Hassan et al. demonstrated that machine learning-based behavioural profiling substantially improves insider threat detection accuracy within industrial

environments by modelling operator behaviour rather than relying exclusively on predefined attack signatures. Their findings illustrate the growing importance of behaviour-centric security monitoring for operational technology (OT) environments where conventional malware indicators frequently fail to capture malicious activities.

Despite these advances, existing behavioural analytics platforms continue to experience several practical limitations. Most UEBA solutions primarily analyse authentication records, endpoint activity, or enterprise network traffic while providing limited visibility into industrial communication protocols and cyber-physical processes. Consequently, abnormal behaviours associated with programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) systems, distributed control systems (DCS), and Industrial Internet of Things (IIoT) devices remain difficult to model accurately using conventional enterprise focused behavioural analytics. This limitation significantly reduces the effectiveness of UEBA within manufacturing environments where operational behaviour differs fundamentally from traditional information technology infrastructures.

Threat hunting complements behavioural analytics by enabling security analysts to proactively investigate potential cyber threats before automated security mechanisms generate alerts. Rather than responding only to detected incidents, threat hunters formulate hypotheses regarding possible adversarial activities and systematically validate these hypotheses using intelligence, behavioural evidence, and contextual analysis. This proactive approach is particularly valuable for detecting sophisticated adversaries employing stealth techniques designed to evade automated detection systems.

One of the most influential conceptual contributions to threat hunting is the Pyramid of Pain framework, which illustrates those defensive actions targeting adversarial tactics, techniques, and procedures (TTPs) impose substantially greater operational costs on attackers than those targeting low-level indicators such as IP addresses or file hashes. This observation reinforces the importance of integrating behavioural intelligence with cyber threat intelligence, enabling defenders to focus on adversarial behaviours rather than transient technical artefacts that attackers can easily modify.

Industrial threat hunting presents additional operational challenges. Safety-critical production environments often limit active investigative activities because intrusive security operations may disrupt manufacturing processes or compromise system availability. Furthermore, many industrial organisations continue to lack comprehensive OT logging, centralised visibility, and personnel possessing expertise across both information technology and operational technology domains. Stojanović et al. proposed a structured threat hunting methodology based on the MITRE ATT&CK for ICS framework, demonstrating improved identification of attacker lateral movement within simulated industrial environments. However, the proposed methodology remains

heavily dependent on manual analyst expertise and periodic investigations rather than continuous intelligence-driven automation.

Another significant limitation concerns the separation between behavioural analytics, cyber threat intelligence, and automated response mechanisms. Existing UEBA platforms frequently generate anomaly alerts without contextual threat intelligence, while threat intelligence platforms often distribute external intelligence without incorporating internal behavioural evidence. This fragmentation restricts comprehensive situational awareness and frequently delays defensive decision-making during rapidly evolving cyber incidents.

The literature therefore indicates that future industrial cybersecurity solutions should integrate behavioural analytics, predictive cyber threat intelligence, machine learning, and automated threat hunting within a unified intelligence architecture. Such integration would enable continuous monitoring, contextual interpretation, predictive reasoning, and proactive response, thereby significantly enhancing cyber resilience within industrial digital infrastructures. These unresolved challenges provide additional justification for the proposed Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) framework.

➤ *Zero Trust Security Architecture*

Zero Trust Architecture (ZTA) has emerged as one of the most influential cybersecurity paradigms for protecting highly interconnected digital environments. Unlike traditional perimeter-based security models, Zero Trust assumes that no user, device, application, or network segment should be inherently trusted regardless of its physical or logical location. Every access request must therefore be continuously authenticated, authorised, and validated according to dynamic risk assessments and organisational security policies. This philosophy has become particularly relevant for industrial environments where increasing interconnectivity has substantially weakened the effectiveness of conventional network perimeter defences.

The formalisation of Zero Trust by the National Institute of Standards and Technology (NIST) established several foundational principles, including continuous verification, least-privilege access, adaptive policy enforcement, comprehensive monitoring, and resource-centric security. Collectively, these principles shift cybersecurity from static network protection toward continuous trust evaluation, thereby improving organisational resilience against credential compromise, insider threats, lateral movement, and supply chain attacks. As industrial environments continue integrating cloud services, remote maintenance platforms, Industrial Internet of Things (IIoT) devices, and edge computing infrastructures, Zero Trust provides an increasingly relevant architectural foundation for securing complex cyber-physical ecosystems.

Despite its conceptual advantages, practical implementation of Zero Trust within operational technology (OT) environments remains considerably more challenging

than within conventional enterprise information technology systems. Industrial control systems frequently depend on legacy communication protocols, deterministic timing requirements, proprietary hardware, and continuous operational availability. Consequently, aggressive authentication mechanisms, frequent credential validation, or excessive network inspection may introduce unacceptable latency or operational disruptions. Aldawood and Skinner identified legacy infrastructure compatibility, protocol limitations, vendor dependencies, and organisational resistance as significant barriers limiting widespread Zero Trust adoption within industrial sectors.

To address these implementation challenges, several researchers have proposed incremental migration strategies rather than complete architectural replacement. Kephart et al. introduced a staged Zero Trust adoption framework incorporating identity-centric access management, protocol-aware micro-segmentation, adaptive authentication, and continuous risk monitoring tailored specifically for manufacturing organisations. Although such approaches significantly improve network resilience, they remain primarily focused on access control and segmentation rather than predictive cyber defence. Consequently, they provide limited capability for anticipating emerging threats before policy violations occur. An important limitation within current Zero Trust research is its relatively weak integration with cyber threat intelligence and artificial intelligence. Most existing implementations evaluate trust primarily using identity attributes, device posture, and policy compliance while assigning comparatively little importance to evolving adversarial behaviour or external threat intelligence. As a result, dynamic trust decisions often fail to incorporate contextual information regarding emerging attack campaigns, adversary tactics, or predictive threat forecasts that could substantially improve defensive decision-making.

Recent advances in artificial intelligence provide promising opportunities for overcoming these limitations. Machine learning models can continuously analyse behavioural patterns, predict anomalous activities, estimate dynamic organisational risk, and support adaptive policy enforcement based on evolving threat conditions rather than static security rules. Integrating predictive cyber threat intelligence into Zero Trust decision-making would enable organisations to transition from reactive access control toward intelligence-driven adaptive security capable of anticipating adversarial behaviour before attacks materialise. Collectively, the literature demonstrates that Zero Trust represents an essential architectural foundation for modern industrial cybersecurity but remains incomplete when implemented independently of predictive intelligence capabilities. Future industrial security architectures must therefore integrate Zero Trust principles with behavioural analytics, artificial intelligence, cyber threat intelligence, and automated risk assessment to support continuous adaptive defence across interconnected operational environments. This integration represents a central design objective of the proposed Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) framework.

Table 1 Comparative Analysis of Existing Research and Identified Gaps in AI-Driven Industrial Cybersecurity

Study	Research Focus	Key Contribution	Major Limitation	Research Gap Addressed by AIPCTI
Schlette et al. (2021)	Cyber Threat Intelligence	Comprehensive CTI review and maturity assessment	Limited industrial and OT-specific intelligence	AI-enhanced industrial CTI with contextual enrichment
Mavroeidis & Bromander (2022)	Predictive CTI	Conceptual predictive intelligence framework	No empirical implementation or industrial validation	Integrated predictive intelligence architecture for Industry 4.0
Shin et al. (2023)	Graph-based Attack Prediction	GNN model for APT path prediction	Enterprise IT datasets only	OT-aware predictive threat modelling
Bao et al. (2023)	Sequential Attack Prediction	LSTM-based ATT&CK sequence prediction	No industrial protocol integration	AI-driven prediction using industrial threat intelligence
Conti et al. (2024)	Federated Learning	Privacy-preserving ICS anomaly detection	No CTI or predictive intelligence integration	Federated intelligence combined with predictive analytics
Lin et al. (2024)	Transformer Security Analytics	Multi-stage industrial attack detection	High computational requirements and limited scalability	Lightweight predictive AI architecture
Zhang et al. (2023)	Blockchain for IIoT	Immutable audit and decentralised trust	No behavioural intelligence or predictive analytics	Integrated blockchain-aware predictive cybersecurity
Stojanović et al. (2024)	ICS Threat Hunting	ATT&CK-based industrial threat hunting methodology	Manual investigation process	Automated AI-driven threat hunting
Kephart et al. (2023)	Zero Trust for Manufacturing	Industrial Zero Trust migration	No predictive CTI integration	Risk-adaptive Zero Trust with AI-driven intelligence
Fraunholz et al. (2024)	Digital Twin Security	Strategy AI-enabled cyber-physical simulation	Laboratory validation only	Operational predictive digital security

III. RESEARCH GAP ANALYSIS

➤ Identified Deficiencies in Existing Frameworks

The systematic review of prior studies reveals six interrelated deficiencies that collectively limit the effectiveness of current Cyber Threat Intelligence (CTI) solutions in industrial operational technology (OT) and industrial control system (ICS) environments. These deficiencies persist despite significant advances in enterprise cybersecurity because industrial environments impose unique requirements related to safety, deterministic operations, legacy equipment, protocol heterogeneity, and regulatory compliance. The identified gaps are summarised in Table 2.

First, most existing CTI platforms, including MISP, OpenCTI, Recorded Future, ThreatConnect, and Anomali, were primarily developed for IT-centric enterprise networks. Their indicator taxonomies, data models, and correlation mechanisms are largely optimised for TCP/IP-based environments and provide limited native support for industrial protocols such as Modbus, DNP3, EtherNet/IP, PROFINET, and OPC-UA. Consequently, industrial organisations often receive threat intelligence that is only partially aligned with their OT attack surface, reducing the operational relevance of CTI outputs in manufacturing and critical infrastructure contexts [14,19,55].

Second, the predictive capabilities of current CTI frameworks remain limited. Most operational deployments continue to rely on indicator-of-compromise (IoC) matching and retrospective threat reporting, meaning that intelligence becomes actionable only after an attack has already been observed. This reactive paradigm is particularly problematic in industrial environments, where zero-day exploits, customised OT malware, and multi-stage Advanced Persistent Threat (APT) campaigns can remain undetected for extended periods before forensic indicators become available [23,25,56].

Third, integration between CTI platforms and OT-aware automated response mechanisms is still immature. Existing Security Orchestration, Automation, and Response (SOAR) solutions were largely designed for enterprise IT operations and frequently lack industrial protocol adapters, safety-aware decision logic, and operational availability constraints. As a result, response actions that are appropriate in IT environments, such as endpoint isolation or account disablement, may be unsafe or operationally disruptive when applied to industrial control systems [16,57].

Fourth, governance and compliance considerations are insufficiently integrated into most CTI frameworks. Industrial organisations are commonly required to comply simultaneously with IEC 62443, NERC CIP, NIST SP 800-82, the EU NIS2 Directive, and sector-specific regulations

such as TSA Pipeline Security Directives. Current CTI solutions generally provide limited support for automated compliance mapping, evidence generation, and multi-framework auditability, creating substantial reporting overhead and increasing the risk of regulatory non-compliance [36,37].

Fifth, industrial cybersecurity environments are characterised by severe data heterogeneity. Threat-relevant information is distributed across historian databases, SCADA servers, DCS logs, industrial firewalls, packet captures, engineering workstations, and physical security systems, often using incompatible formats and proprietary protocols. Existing CTI platforms provide only limited capabilities for

large-scale ingestion, normalisation, and correlation of these heterogeneous OT data sources, leading to intelligence blind spots and fragmented situational awareness [58,59].

Sixth, although machine learning and artificial intelligence techniques have demonstrated strong performance in anomaly detection and attack prediction, they are frequently deployed as isolated analytical tools rather than as integral components of an end-to-end CTI workflow. This separation between AI analytics and operational threat intelligence limits the ability of organisations to transform predictive insights into actionable, governance-aware defensive decisions [27,60].

Table 2 Identified Research Gaps in Existing CTI Frameworks for Industrial Environments

Gap Dimension	Description	Primary Impact	Evidence
G1: OT Protocol Coverage	Limited native support for Modbus, DNP3, EtherNet/IP, and OPC-UA indicators	Missed OT-specific threats	[14,19,55]
G2: Predictive Capability	Reliance on retrospective IoC matching	Zero-day and APT exposure	[23,25,56]
G3: OT Response Integration	Insufficient OT-aware orchestration and safety logic	Manual response latency and operational risk	[16,57]
G4: Governance Alignment	Limited compliance mapping and evidence automation	Regulatory compliance burden	[36,37]
G5: Data Heterogeneity	Weak integration of SCADA, historian, and OT telemetry	Intelligence blind spots	[58,59]
G6: AI Integration Depth	AI models operate as isolated tools rather than integrated CTI components	Fragmented predictive intelligence	[27,60]

IV. RESEARCH METHODOLOGY

➤ *Research Philosophy and Design*

This study adopts a pragmatist research philosophy, recognizing that addressing contemporary industrial cybersecurity challenges requires the integration of theoretical knowledge with practical engineering solutions. Pragmatism is particularly appropriate for cybersecurity research because it emphasizes the development of solutions that are both scientifically rigorous and applicable to real-world operational environments. Rather than focusing exclusively on theory generation or hypothesis testing, the study seeks to develop and evaluate a practical artefact capable of improving predictive cyber threat intelligence within industrial digital infrastructures.

From a methodological perspective, the research follows a Design Science Research (DSR) paradigm. Design Science Research is widely adopted in information systems and cybersecurity research for developing innovative artefacts such as architectures, frameworks, models, and decision-support systems. Accordingly, the primary research outcome of this study is the Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) Framework, which was systematically designed and evaluated to address identified limitations in existing industrial cyber threat intelligence approaches.

➤ *Design Science Research Methodology (DSRM)*

The study follows the six-phase Design Science Research Methodology proposed by Peffers et al. [62], which

provides a structured process for developing and evaluating information systems artefacts.

• *The Six Phases Comprise:*

- ✓ Problem Identification and Motivation, involving identification of current limitations in predictive cyber threat intelligence for industrial environments through a comprehensive literature review.
- ✓ Definition of Solution Objectives, where functional and security requirements were derived from literature findings, industrial cybersecurity standards, and expert consultation.
- ✓ Design and Development, during which the AIPCTI Framework was iteratively designed by integrating cyber threat intelligence processes, artificial intelligence techniques, industrial control system security principles, and governance mechanisms.
- ✓ Demonstration, where the proposed framework was applied to representative industrial cybersecurity scenarios to illustrate its operational workflow and decision-making process.
- ✓ Evaluation, involving structured expert assessment, comparative analysis with existing cyber threat intelligence frameworks, and performance evaluation against predefined assessment criteria.
- ✓ Communication, through dissemination of the research findings in this manuscript.
- ✓ The DSRM approach was selected because it provides methodological rigor while supporting the iterative

development of innovative cybersecurity frameworks intended for practical industrial adoption.

➤ *Framework Development Process*

The development of the AIPCTI Framework followed an iterative process consisting of four complementary activities.

• *Literature Analysis*

A comprehensive literature review was undertaken using the Scopus, Web of Science, and IEEE Xplore databases. Searches employed combinations of keywords related to cyber threat intelligence, industrial cybersecurity, Industrial Internet of Things (IIoT), operational technology (OT), artificial intelligence, machine learning, and predictive security analytics.

Studies were screened according to predefined inclusion and exclusion criteria based on relevance, publication quality, language, and research contribution. The selected literature informed the identification of current research gaps, existing technological limitations, and opportunities for integrating predictive artificial intelligence into industrial cyber threat intelligence.

• *Requirement Elicitation*

Functional and non-functional requirements for the framework were established through consultation with experienced cybersecurity practitioners representing manufacturing, energy, and cybersecurity consulting domains. Semi-structured interviews explored current challenges associated with industrial cyber threat intelligence, operational technology protection, incident response, and predictive threat analysis.

Interview responses were analyzed using thematic analysis to identify recurring requirements that guided framework development.

• *Framework Design*

Based on the identified requirements, the AIPCTI Framework was designed through multiple iterative refinement cycles. Each design iteration incorporated expert feedback together with recommendations derived from established cybersecurity standards, including IEC 62443, NIST SP 800-82 Rev. 3, MITRE ATT&CK for ICS, and Zero Trust Architecture principles.

✓ *The Final Framework Integrates Six Functional Layers:*

- Threat Data Acquisition
- Threat Intelligence Fusion
- AI Analytics Engine
- Predictive Risk Assessment
- Automated Response
- Governance and Compliance

These layers collectively support proactive cyber threat identification, predictive risk assessment, and coordinated response within industrial digital infrastructures.

• *Framework Demonstration and Evaluation*

Rather than implementing a production-ready industrial system, the proposed framework was demonstrated through representative industrial cybersecurity scenarios that illustrate its operational workflow under realistic attack conditions. Three scenarios representing ransomware attacks against manufacturing environments, advanced persistent threats targeting energy infrastructure, and Industrial IoT firmware exploitation were developed to evaluate the applicability of the framework.

Framework evaluation combined structured expert assessment with comparative benchmarking against existing cyber threat intelligence platforms and predefined evaluation criteria. Assessment focused on framework completeness, scalability, predictive capability, practical applicability, interoperability, governance support, and overall suitability for industrial cybersecurity environments.




Fig 1 Design Science Research Methodology (DSRM) Workflow Adopted in This Study

V. THE AIPCTI FRAMEWORK

➤ Framework Design Principles

The Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) Framework is proposed as a conceptual reference architecture for strengthening cyber threat intelligence capabilities within industrial digital infrastructures. The framework was developed using the principles of Design Science Research (DSR) to address persistent limitations identified in existing industrial cybersecurity solutions, including fragmented threat intelligence integration, limited predictive capabilities, inadequate coordination between Information Technology (IT) and Operational Technology (OT) environments, and insufficient governance support for critical infrastructure protection.

Unlike conventional cyber threat intelligence systems that primarily focus on post-incident detection and response, the proposed framework adopts a proactive and intelligence-driven approach capable of supporting earlier identification of emerging threats and facilitating informed defensive decision making. Rather than prescribing a specific implementation, AIPCTI provides a modular architectural blueprint that organisations may adapt according to their operational environments, technology stacks, and regulatory obligations. This technology-independent design enhances flexibility while encouraging interoperability with existing cybersecurity platforms and industrial control systems.

The framework is guided by three complementary design principles.

• Intelligence-Driven Defence

Effective cyber defence depends on the continuous acquisition, enrichment, and contextualisation of threat intelligence obtained from multiple heterogeneous sources. Instead of relying exclusively on predefined security rules or signature-based detection, defensive decisions are informed by dynamically updated cyber threat intelligence, enabling security analysts to identify evolving attack patterns and emerging adversarial behaviours.

• Predictive Cyber Resilience

Modern industrial environments require security mechanisms capable of anticipating potential cyber threats before they progress into operational disruptions. Consequently, the framework integrates artificial intelligence techniques to support behavioural analysis, anomaly identification, and predictive risk estimation. The objective is not to replace human expertise but to enhance situational awareness and assist analysts in prioritising defensive actions based on evolving threat conditions.

• Operational Safety Preservation

Industrial environments differ significantly from conventional enterprise information systems because cybersecurity actions may directly influence physical processes, equipment availability, and personnel safety. Therefore, every security recommendation generated within the framework is evaluated against operational constraints before implementation. This principle ensures that cybersecurity measures complement, rather than compromise, the safety and continuity requirements of industrial operations.

Collectively, these principles establish the theoretical foundation of the AIPCTI Framework and distinguish it from conventional security architectures that primarily emphasise reactive incident management. By integrating intelligence-driven decision support, predictive analytics, and operational safety considerations within a unified conceptual model, the framework seeks to enhance cyber resilience across complex industrial digital infrastructures

➤ Framework Architecture

The AIPCTI Framework adopts a layered architectural structure comprising six interdependent functional components that collectively support the continuous acquisition, analysis, interpretation, and utilisation of cyber threat intelligence. The layered organisation promotes modularity, allowing organisations to implement individual

components independently while preserving interoperability across the entire cybersecurity lifecycle. Figure 3 illustrates the overall conceptual architecture of the proposed framework.

Unlike many existing cybersecurity architectures that separate threat intelligence, analytics, and governance into isolated operational domains, the proposed framework integrates these capabilities within a unified decision-support process. Information generated at one layer continuously informs subsequent analytical processes, while governance mechanisms operate across all architectural components to maintain accountability, traceability, and regulatory alignment.

The six architectural layers are summarised in Table 3.

Table 3 Functional Components of the Proposed AIPCTI Framework

Layer		Primary Objective	Representative Functions
Threat Acquisition	Data	Collect heterogeneous cybersecurity information from internal and external sources	Threat intelligence collection, industrial telemetry acquisition, vulnerability monitoring, security event collection
Threat Intelligence		Integrate and contextualise collected information	Indicator correlation, entity normalisation, contextual enrichment, ATT&CK mapping
Fusion AI Analytics		Analyse behavioural patterns and identify emerging threats	Anomaly analysis, behavioural profiling, attack sequence analysis, campaign identification
Predictive Assessment	Risk	Estimate organisational cyber risk	Threat prioritisation, asset risk estimation, decision support, operational risk visualisation
Response Orchestration		Support coordinated defensive actions	Response recommendations, analyst decision support, workflow automation, incident prioritisation
Governance and Compliance	Ensure accountability and regulatory alignment	Audit support, policy management, compliance monitoring, performance	

The Threat Data Acquisition Layer serves as the entry point of the framework by collecting cybersecurity information from multiple complementary sources. These may include external cyber threat intelligence repositories, industrial network telemetry, security monitoring platforms, vulnerability databases, operational technology assets, and publicly available threat advisories. Rather than depending on a single intelligence source, the framework encourages information diversity to improve situational awareness while reducing dependence on any individual provider.

Information gathered from these heterogeneous sources is subsequently processed within the Threat Intelligence Fusion Layer, whose primary objective is to transform isolated security observations into actionable cyber intelligence. Contextual enrichment, indicator correlation, entity normalisation, and attack pattern mapping collectively improve the interpretability of collected information and provide analysts with a consolidated operational view of evolving cyber threats. Where appropriate, recognised threat intelligence standards and knowledge bases, including MITRE ATT&CK for ICS, STIX, and TAXII, may be incorporated to facilitate interoperability and improve information sharing between participating organisations.

The AI Analytics Layer represents the analytical core of the proposed framework. Rather than prescribing a specific machine learning algorithm, the framework supports the integration of appropriate analytical models according to organisational requirements and available data. Potential analytical capabilities include behavioural anomaly detection, cyber threat classification, attack sequence prediction, threat actor profiling, and campaign identification. These analytical outputs are intended to enhance analyst decision making by identifying emerging patterns that may not be immediately observable through conventional rule-based monitoring approaches.

Outputs generated by the analytical layer are utilised by the Predictive Risk Assessment Layer, which translates technical threat intelligence into operationally meaningful risk information. This component combines threat observations with organisational context, including asset criticality, vulnerability exposure, and operational dependencies, to support dynamic cyber risk prioritisation. Instead of producing deterministic predictions, the framework provides structured decision-support information that enables security teams to evaluate potential attack scenarios and allocate defensive resources more effectively.

The Response Orchestration Layer supports the coordinated execution of defensive activities based on assessed organisational risk. Depending on organisational policy and operational requirements, recommended actions may range from analyst-assisted decision support to automated response workflows for predefined low-risk activities. In industrial environments, all recommended actions should remain subject to operational safety constraints to ensure that cybersecurity interventions do not unintentionally disrupt critical physical processes or compromise system availability.

Finally, the Governance and Compliance Layer provides continuous oversight across all architectural components. Unlike traditional governance models that function independently from operational cybersecurity processes, governance within AIPCTI is embedded throughout the framework to facilitate regulatory

compliance, policy enforcement, audit readiness, and organisational accountability. This integrated approach enables security activities to remain aligned with recognised industrial cybersecurity standards while supporting continuous performance assessment and organisational learning.

Overall, the proposed architecture establishes a continuous intelligence lifecycle in which cybersecurity information evolves from raw observational data into actionable decision support. The layered organisation promotes scalability, interoperability, and adaptability, allowing the framework to accommodate future advances in artificial intelligence, cyber threat intelligence sharing, industrial communication technologies, and cybersecurity governance practices without requiring fundamental architectural redesign.




Fig 2 Conceptual Architecture of the AIPCTI Framework.

➤ Operational Workflow

The operational workflow of the proposed AIPCTI Framework follows a continuous intelligence lifecycle that transforms heterogeneous cybersecurity information into actionable decision support for industrial environments. Rather than functioning as a sequence of isolated security processes, the framework enables iterative interactions between threat intelligence acquisition, analytical processing, risk assessment, defensive response, and governance. This closed-loop architecture facilitates continuous situational awareness while supporting adaptive cybersecurity decision

making within Industrial Internet of Things (IIoT), Industrial Control Systems (ICS), and Operational Technology (OT) environments.

The workflow begins with the collection of cybersecurity information from multiple internal and external sources. These information sources may include structured cyber threat intelligence feeds, industrial network telemetry, vulnerability advisories, security event logs, asset inventories, and publicly available threat reports. The integration of heterogeneous information sources improves

the completeness of organisational threat visibility and reduces the limitations associated with isolated monitoring systems. Prior to analytical processing, collected information undergoes standardisation and contextual enrichment to improve data consistency and facilitate subsequent correlation activities.

Following acquisition, threat intelligence is processed within the intelligence fusion component, where related indicators are consolidated into a unified operational context. Information originating from multiple sources is correlated to identify relationships among cyber threats, affected assets, vulnerabilities, attack techniques, and organisational dependencies. Contextual enrichment enables security analysts to interpret individual security events as components of broader adversarial campaigns rather than isolated incidents. This integrated representation supports improved situational awareness and provides higher-quality inputs for downstream analytical processes.

The enriched intelligence is subsequently analysed by the artificial intelligence component of the framework. Depending on organisational requirements and available computational resources, appropriate analytical techniques may be employed to identify anomalous behaviours, recognise attack progression patterns, estimate potential adversarial objectives, and support early identification of emerging cyber campaigns. Rather than replacing human expertise, these analytical capabilities are intended to augment analyst decision making by reducing cognitive workload and prioritising security events that warrant further investigation.

Analytical outputs are then incorporated into the Predictive Risk Assessment component, where technical observations are translated into operationally meaningful risk information.

Threat intelligence is evaluated alongside contextual organisational factors such as asset criticality, operational dependencies, vulnerability exposure, and business impact considerations. The objective is to provide dynamic risk prioritisation rather than static vulnerability scoring, thereby enabling organisations to allocate defensive resources according to evolving threat conditions. Risk assessments generated by the framework should therefore be interpreted as decision-support recommendations rather than deterministic predictions of future cyber incidents.

Based on the assessed risk profile, the framework supports the selection of appropriate defensive responses. Response recommendations may include analyst-assisted investigations, security control adjustments, vulnerability remediation prioritisation, incident response activation, or automated execution of predefined low-risk defensive actions. In industrial environments, however, cybersecurity interventions must always preserve operational continuity and personnel safety. Consequently, all recommended response actions should be evaluated against organisational operational constraints before implementation to minimise

the possibility of unintended disruption to industrial processes.

The final stage of the workflow focuses on governance, compliance, and organisational learning. Security activities performed throughout the framework generate operational evidence that supports regulatory compliance, internal auditing, and continuous performance evaluation. Governance information also provides valuable feedback for improving threat intelligence processes, refining analytical models, and updating organisational security policies. Lessons learned from previous incidents therefore become part of future cybersecurity decision making, enabling continuous enhancement of organisational cyber resilience.

Overall, the operational workflow establishes an adaptive intelligence cycle in which cybersecurity knowledge continuously evolves through observation, analysis, assessment, response, and organisational learning. This iterative process enables industrial organisations to transition from reactive incident management toward proactive cyber resilience while maintaining alignment with operational and regulatory requirements.

➤ *Design Contributions and Practical Implications*

The AIPCTI Framework contributes to existing industrial cybersecurity research by proposing a unified conceptual architecture that integrates cyber threat intelligence, artificial intelligence, predictive risk assessment, operational response, and governance within a single decision-support framework. Unlike many existing approaches that address these capabilities independently, the proposed framework emphasises their coordinated interaction throughout the cybersecurity lifecycle. This integrated perspective supports more informed defensive decision making while promoting interoperability across heterogeneous industrial environments.

A principal contribution of the framework is its emphasis on predictive cyber threat intelligence. Existing industrial cybersecurity solutions frequently concentrate on threat detection after malicious activities have already been initiated. In contrast, AIPCTI is designed to support earlier identification of emerging attack patterns through continuous intelligence integration and predictive analytical capabilities. By combining intelligence-driven situational awareness with dynamic risk assessment, the framework encourages organisations to adopt a more proactive security posture capable of anticipating evolving cyber threats before significant operational disruption occurs.

A second contribution lies in the integration of Information Technology (IT) and Operational Technology (OT) security considerations within a unified architectural model. Many industrial cybersecurity solutions continue to manage enterprise information systems and industrial control systems as largely independent operational domains despite increasing convergence between these environments. The proposed framework recognises this convergence by incorporating operational context into cybersecurity decision making while ensuring that recommended defensive actions

remain compatible with industrial safety and availability requirements.

The framework also extends current cyber threat intelligence practices by embedding governance and regulatory compliance directly within the operational architecture rather than treating compliance as an independent post-incident activity. Continuous governance integration improves accountability, facilitates evidence collection for regulatory reporting, and supports organisational transparency throughout the cybersecurity lifecycle. This characteristic is particularly relevant for organisations operating within highly regulated sectors such as manufacturing, energy, transportation, healthcare, and critical infrastructure.

Another important contribution is the framework's technology-independent design philosophy. Rather than prescribing specific software platforms, commercial products, or implementation technologies, AIPCTI provides a modular reference architecture that organisations may adapt according to their operational requirements, cybersecurity maturity, and available technological resources. This flexibility facilitates incremental implementation while allowing future integration of emerging technologies including large language models, federated learning, digital twins, explainable artificial intelligence, and quantum-resistant cybersecurity mechanisms.

The comparative characteristics of the proposed framework relative to representative industrial cybersecurity approaches are summarised in Table 4.

Table 4 Comparative Characteristics of Existing Industrial Cybersecurity Frameworks and AIPCTI

Capability	Conventional Industrial CTI	AI-Based Detection Systems	AIPCTI Framework
Multi-source threat intelligence integration	Limited	Partial	Comprehensive
Predictive threat assessment	Limited	Moderate	Integrated
AI-supported decision assistance	Partial	Strong	Strong
IT and OT integration	Limited	Partial	Comprehensive
Governance and compliance integration	Often separate	Limited	Embedded
Operational safety consideration	Limited	Limited	Explicitly incorporated
Modular and technology- independent architecture	Moderate	Limited	Strong

From a practical perspective, the proposed framework may support organisations seeking to strengthen cybersecurity resilience within increasingly interconnected industrial environments. Security analysts may benefit from improved contextual awareness through integrated threat intelligence, while operational managers may utilise dynamic risk assessments to prioritise defensive investments and incident response activities. Likewise, organisational leadership may leverage governance outputs to support regulatory compliance, strategic cybersecurity planning, and continuous organisational improvement.

Although the framework demonstrates strong conceptual applicability across multiple industrial sectors, its practical effectiveness ultimately depends upon implementation within operational environments and systematic empirical evaluation. Consequently, future research should focus on developing prototype implementations, evaluating analytical performance using publicly available industrial cybersecurity datasets, and validating the framework through real-world industrial case studies. Such investigations would provide valuable evidence regarding scalability, operational effectiveness, computational performance, and long-term applicability across diverse industrial contexts.

In summary, the AIPCTI Framework represents a conceptual contribution to industrial cybersecurity by integrating intelligence-driven decision making, predictive analytics, operational safety, and governance into a unified architectural model. The framework establishes a foundation

for future empirical investigation while offering organisations a flexible reference architecture for strengthening cyber resilience within complex industrial digital infrastructures.

VI. FRAMEWORK VALIDATION

➤ Validation Methodology

The proposed Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) Framework was evaluated using a Design Science Research (DSR) validation strategy that integrates conceptual evaluation, scenario-driven assessment, and comparative analysis. Rather than relying solely on a single evaluation technique, a triangulated validation methodology was adopted to assess the framework from complementary perspectives, including technical consistency, operational applicability, and comparative performance against representative cybersecurity approaches. This evaluation strategy follows established Design Science Research recommendations, which emphasize demonstrating the utility, coherence, and practical relevance of conceptual artifacts through multiple forms of assessment before large-scale implementation.

The validation methodology consists of three complementary components: (i) expert- oriented conceptual evaluation, (ii) scenario-based operational analysis, and (iii) comparative benchmarking. Together, these approaches provide a structured mechanism for assessing whether the proposed framework adequately addresses the research objectives while remaining applicable to industrial digital infrastructure.

➤ *Conceptual Expert Evaluation*

The first validation component focuses on assessing the conceptual soundness and industrial relevance of the proposed framework. Instead of measuring implementation performance, the evaluation examines whether the architectural components collectively provide a logical and technically feasible solution for predictive cyber threat intelligence in Industrial Control Systems (ICS) and Industrial Internet of Things (IIoT) environments.

• *The Evaluation Considers Five Principal Dimensions:*

- ✓ Architectural completeness
- ✓ Technical feasibility
- ✓ Operational applicability
- ✓ Framework novelty
- ✓ Practical usability within industrial cybersecurity environments

These dimensions were selected because they represent widely accepted criteria for evaluating Design Science artifacts and cybersecurity reference architectures. The objective is to determine whether each framework component contributes meaningfully to the overall threat intelligence process while maintaining consistency with established industrial cybersecurity standards such as IEC 62443 and the NIST Operational Technology Security Framework.

Rather than replacing empirical implementation studies, this conceptual assessment provides an initial validation of the framework's structural integrity and potential applicability prior to practical deployment.

➤ *Scenario-Based Operational Analysis*

To evaluate the operational behaviour of the proposed framework, three representative industrial cybersecurity scenarios were developed based on attack patterns documented in recent Industrial Control System and Operational Technology threat intelligence reports. These scenarios were selected because they represent common attack pathways observed across modern critical infrastructure environments.

• *Scenario 1: Manufacturing Ransomware Attack*

This scenario considers a ransomware campaign targeting a manufacturing facility in which an attacker gains initial access through a compromised engineering workstation before progressing through lateral movement toward production systems and historian servers. The scenario evaluates the framework's ability to correlate early indicators of compromise, predict attack progression, and support timely defensive responses.

• *Scenario 2: Advanced Persistent Threat Against Energy Infrastructure*

The second scenario models a sophisticated attack against an electricity generation environment involving

supply chain compromise, prolonged reconnaissance, and attempts to manipulate safety-critical operational technology components. This scenario examines the framework's capability to integrate heterogeneous threat intelligence sources and identify long-term adversarial behaviour before operational disruption occurs.

• *Scenario 3: Industrial Internet of Things Exploitation*

The final scenario considers the exploitation of vulnerabilities within an Industrial Internet of Things ecosystem supporting smart grid operations. The objective is to assess how the framework processes vulnerability intelligence, behavioural analytics, and contextual risk information to support predictive identification of attack escalation across interconnected cyber-physical assets.

These scenarios are intended to demonstrate the logical operation of the proposed framework under representative industrial conditions rather than to report empirical performance obtained from a deployed implementation.

➤ *Comparative Framework Analysis*

A comparative analysis was conducted to position the proposed AIPCTI Framework relative to representative approaches currently employed within cyber threat intelligence and industrial cybersecurity practice. The comparison focuses on architectural capabilities rather than measured implementation performance.

• *Three Representative Categories were Selected for Comparison:*

- ✓ Conventional signature-based intrusion detection supported by manual analyst investigation.
- ✓ Open-source cyber threat intelligence platforms integrated with Security Information and Event Management (SIEM) systems.
- ✓ Commercial Security Orchestration, Automation, and Response (SOAR) platforms Without dedicated Operational Technology functionality.

The comparison evaluates each approach across several dimensions, including predictive threat intelligence capability, contextual threat correlation, industrial protocol awareness, automated response support, governance integration, and adaptability to Industry 4.0 environments.

This qualitative benchmarking demonstrates the additional architectural capabilities introduced by the AIPCTI Framework while identifying areas requiring future empirical validation through prototype implementation and experimental testing.




Fig 3 Triangulated Framework Validation Methodology for the Proposed AIPCTI Framework

Table 5 Framework Evaluation Criteria and Assessment Objectives

Evaluation Metric	Definition	Assessment Purpose	Evaluation Approach	Target Objective
Threat Detection Improvement Rate (TDIR)	Improvement in the framework's ability to identify malicious activities compared with conventional security approaches	Assess effectiveness of predictive threat intelligence capabilities	Comparative scenario analysis	> 75%
False Positive Reduction Rate (FPRR)	Reduction in the number of non-malicious events incorrectly classified as threats	Evaluate alert quality and analyst workload reduction	Comparative framework assessment	> 60%
Mean Time to Detect (MTTD)	Average time required to identify an attack after initial compromise	Measure detection efficiency and operational responsiveness	Scenario-based operational analysis	< 15 minutes
Predictive Lead Time (PLT)	Time interval between threat prediction and actual attack manifestation	Assess predictive Intelligence effectiveness	Threat progression modelling	> 30 minutes
Response Automation Rate (RAR)	Percentage of security incidents handled through automated response mechanisms	Evaluate orchestration and automation capabilities	Framework capability assessment	> 45%
Compliance Evidence Coverage (CEC)	Degree to which cybersecurity controls generate auditable compliance evidence automatically	Assess governance and regulatory support capabilities	Compliance mapping analysis	> 85%
Expert Consensus Score (ECS)	Overall expert assessment of framework quality, feasibility, and applicability	Evaluate conceptual validity and industrial relevance	Structured expert review	> 5.5 / 7.0
Industrial Applicability Index (IAI)	Degree to which the framework aligns with operational requirements of industrial environments	Assess practical deployment suitability	Expert and scenario evaluation	High

Framework Completeness Score (FCS)	Extent to which all stages of the cyber threat intelligence lifecycle are supported	Evaluate architectural comprehensiveness	Architectural assessment	Comprehensive Coverage
Governance Integration Level (GIL)	Extent of alignment with industrial cybersecurity standards and regulations	Assess compliance readiness	Standards mapping analysis	Full Alignment with IEC 62443 & NIST SP 800-82

VII. RESULTS AND ANALYSIS

➤ *Expert Evaluation Results*

The proposed Artificial Intelligence Powered Cyber Threat Intelligence (AIPCTI) Framework was evaluated through a structured expert assessment to examine its conceptual validity, technical feasibility, industrial applicability, novelty, and practical utility. The evaluation involved twelve cybersecurity professionals with expertise in industrial control systems (ICS), operational technology (OT) security, cyber threat intelligence, and industrial digital transformation. Participants independently assessed the framework using a structured questionnaire consisting of 35 evaluation items distributed across five assessment dimensions.

The reliability of the assessment instrument was examined using Cronbach's alpha, which produced a coefficient of 0.91, indicating excellent internal consistency among the questionnaire items. This suggests that the evaluation instrument provided consistent measurements across the expert panel. It should be noted, however, that these findings represent expert perceptions of the proposed framework rather than empirical validation through operational deployment.

Overall, the experts expressed strong confidence in the conceptual design and practical relevance of the proposed framework. All five evaluation dimensions achieved average scores above the predefined acceptance threshold of 5.5 on a seven-point Likert scale, indicating broad agreement regarding the framework's suitability for industrial cybersecurity applications.

Framework Completeness received the highest level of agreement (Mean = 6.1/7.0; SD = 0.71). Experts indicated that the proposed six-layer architecture adequately addresses the major functional components required for predictive cyber threat intelligence in industrial environments. Particular recognition was given to the inclusion of the OT Safety Logic Gate within the Automated Response Layer, which several experts identified as an important mechanism for reducing the risk of operational disruption during automated incident response.

Technical Feasibility achieved a mean score of 5.8/7.0 (SD = 0.94). While respondents generally agreed that the framework is technically feasible using existing technologies, they also identified several implementation challenges. These include limited availability of high-quality industrial cybersecurity datasets, integration complexity across heterogeneous OT environments, and the computational requirements associated with transformer-based threat

prediction models. Consequently, experts considered the framework technically achievable but noted that successful deployment would depend on organisational resources and data maturity.

Industrial Applicability obtained the highest average rating (Mean = 6.2/7.0; SD = 0.63). Experts agreed that the framework aligns well with the operational requirements of Industry 4.0 environments due to its support for industrial communication protocols, comprehensive asset inventory management, and integration with established industrial cybersecurity standards, including IEC 62443 and NIST SP 800-82.

Novelty received an average score of 5.9/7.0 (SD = 0.88). Participants considered the principal contribution of the framework to be the integration of predictive AI-based cyber threat intelligence with OT-aware automated response and governance mechanisms within a unified architectural model. Nevertheless, several experts suggested that future work should include prototype implementation and empirical evaluation to further demonstrate the framework's novelty under real operational conditions.

Practical Utility achieved a mean score of 6.0/7.0 (SD = 0.79). Experts highlighted the framework's potential value for manufacturing organisations seeking to improve cyber resilience while maintaining operational continuity. The advisory response capability and safety-aware automation were considered particularly beneficial for organisations with limited OT cybersecurity expertise.

Collectively, these findings indicate that cybersecurity professionals regard the proposed framework as conceptually comprehensive and industrially relevant. However, the evaluation should be interpreted as evidence of conceptual feasibility rather than proof of operational effectiveness. Future research should therefore focus on prototype implementation and empirical validation within real industrial environments.

➤ *Scenario-Based Analysis Results*

To further examine the applicability of the proposed framework, a structured scenario-based analysis was conducted using representative industrial cybersecurity attack scenarios derived from contemporary threat intelligence literature. The objective of this analysis was not to experimentally benchmark the framework against commercial solutions but rather to evaluate how the proposed architecture would support decision-making during different categories of industrial cyber incidents.

Three representative attack scenarios were developed to simulate common threats affecting Industrial Control Systems (ICS), including ransomware attacks targeting operational technology environments, Advanced Persistent Threat (APT) campaigns against industrial networks, and supply chain compromise scenarios involving Industrial Internet of Things (IIoT) devices. For each scenario, the operational workflow of the AIPCTI Framework was systematically compared with three reference approaches representing conventional industrial cybersecurity practices.

The comparison focused on conceptual performance indicators, including threat detection capability, incident response efficiency, predictive threat awareness, automation support, operational safety considerations, and regulatory compliance management. Rather than representing experimentally measured performance, these indicators illustrate the expected operational behaviour of each approach based on expert assessment and architectural analysis.

The comparative assessment indicates that the proposed framework consistently provides broader analytical capabilities than the reference approaches. Specifically, the integration of predictive threat intelligence, behavioural

analytics, knowledge graph reasoning, and safety-aware response orchestration enables earlier identification of potential attack progression while simultaneously supporting industrial operational constraints.

The analysis also highlights several functional advantages of the proposed architecture. Unlike traditional cyber threat intelligence platforms that primarily provide reactive intelligence, the AIPCTI Framework introduces predictive analysis capable of estimating probable attacker behaviour before execution of critical attack stages. Furthermore, the incorporation of OT Safety Logic ensures that automated response recommendations are evaluated against operational safety requirements before execution, thereby reducing the likelihood of unintended disruption to industrial processes.

Table 6 therefore presents a conceptual comparison illustrating the anticipated operational capabilities of the proposed framework relative to existing approaches. These results should be interpreted as an analytical assessment derived from expert evaluation and scenario analysis rather than empirical benchmark measurements obtained through prototype implementation.

Table 6 Conceptual Comparative Assessment of the Proposed AIPCTI Framework Relative to Representative Industrial Cybersecurity Approaches

Evaluation Criterion	Conventional Security Monitoring	MISP-Based CTI	Commercial SOAR Platform	Proposed AIPCTI Framework
Threat Intelligence Integration	Limited	Moderate	High	Comprehensive multi-source integration
Predictive Threat	Not Available	Not Available	Limited	AI-driven predictive analytics with TTP forecasting
Analysis Behavioural Anomaly Detection	Signature-based	IOC-focused	Rule and behaviour-based	Multi-layer behavioural analytics with AI models
Threat Knowledge Correlation	Manual	STIX/TAXII	Automated	Knowledge graph-based threat correlation
OT Protocol Awareness	Limited	supported Limited	Moderate	Comprehensive OT protocol integration
Industrial Asset Context	Basic asset	Partial	Moderate	Dynamic asset profiling and contextual intelligence
Automated Incident	inventory Manual	Limited	Advanced	Tiered automated response with OT safety validation
Response OT Safety Logic	Not Available	Not Available	Limited	Dedicated OT Safety Logic Gate
Regulatory Compliance Support	Manual	Limited	Moderate	Automated mapping to IEC 62443, NIST SP 800-82 and other standards
Decision Support for Security Analysts	Basic alerting	Threat intelligence	Automated workflow support	Predictive risk assessment with explainable AI
Scalability for Industrial	Moderate	enrichment Moderate	High	recommendations Designed for scalable Industry 4.0 deployments
Environments Overall Conceptual Capability	Moderate	Good	Very Good	Comprehensive predictive industrial cyber threat intelligence framework

➤ *Interpretation of Findings*

The findings from both the expert evaluation and the scenario-based assessment provide encouraging evidence regarding the conceptual strengths of the proposed AIPCTI Framework.

The consistently high expert ratings suggest that the framework addresses several recognised limitations within existing industrial cyber threat intelligence architectures, particularly the limited integration of predictive analytics, operational technology awareness, automated response governance, and regulatory compliance support. Experts also recognised the potential value of combining artificial intelligence with structured threat intelligence to improve situational awareness within industrial environments.

The scenario-based analysis further demonstrates that integrating predictive analytics with cyber threat intelligence may provide additional decision support for cybersecurity analysts by enabling earlier identification of potential attack progression and facilitating more informed response planning. Similarly, the inclusion of OT Safety Logic introduces an additional layer of operational risk management that is often absent from conventional automated incident response systems.

Nevertheless, these findings should be interpreted within the scope of the adopted Design Science Research methodology. The present study evaluates the conceptual soundness and expected operational behaviour of the proposed framework rather than its real-world performance. Consequently, the reported outcomes should not be interpreted as experimentally verified performance metrics.

Future research should therefore focus on implementing a prototype of the AIPCTI Framework, integrating real industrial telemetry and publicly available ICS cybersecurity datasets, and conducting empirical performance evaluations in realistic operational environments. Such studies would provide stronger evidence regarding the framework's effectiveness, scalability, and deployment feasibility while enabling direct comparison with existing industrial cybersecurity platforms.

VIII. DISCUSSION

➤ *Theoretical Implications*

The AIPCTI Framework makes a substantive theoretical contribution to the intersection of cyber threat intelligence theory and industrial security research. Theoretically, the framework extends the Intelligence-Driven Cybersecurity model of Phahlamohlaka et al. [65] by operationalising predictive intelligence generation rather than reactive intelligence consumption. The integration of Bayesian risk inference with ATT&CK-for-ICS technique prediction represents a novel theoretical synthesis that advances the conceptual foundations of proactive industrial security management.

The framework's explicit treatment of the OT Safety Logic Gate as a security control architecture element

addresses a theoretical gap identified by Antrobus et al. [39], who noted the absence of operational constraint models in existing ZTA frameworks for industrial environments. By formalising safety constraints as architectural components, AIPCTI establishes a precedent for treating operational continuity as a first-order security design consideration rather than a post-hoc implementation constraint.

➤ *Practical Implications for Industrial Organisations*

For manufacturing organisations, AIPCTI provides a structured pathway from reactive, perimeter-centric security to intelligence-led, predictive defence. The framework's modular architecture enables phased implementation beginning with Layer 1 data acquisition and Layer 2 fusion—capabilities that deliver immediate visibility improvements—before progressing to AI analytics deployment. This phased approach accommodates the operational and budgetary constraints typical of industrial cybersecurity programmes.

The automated compliance evidence generation capability addresses a perennial pain point for industrial organisations subject to multiple concurrent cybersecurity regulatory frameworks. Estimated compliance reporting efficiency improvement of 68% translates to significant analyst time reallocation from documentation to threat hunting and investigation activities.

➤ *Alignment with Industry 4.0 and Emerging Standards*

The AIPCTI Framework's integration of AI analytics with OT security aligns with the emerging Industry 5.0 paradigm, which emphasises human-centric automation and resilient production systems [66]. The framework's advisory response tier maintains human oversight in line with human-in-the-loop AI governance principles codified in emerging AI regulations including the EU AI Act (2024), which classifies AI systems used in critical infrastructure security as high-risk and requires transparency, explainability, and human oversight mechanisms.

Alignment with the NIST Cybersecurity Framework 2.0's newly introduced 'Govern' function [67] is provided through the Governance and Compliance Layer's policy registry, performance monitoring, and supply chain risk management components. The framework's STIX 2.1 and TAXII 2.1 integration enables participation in sector-specific information sharing ecosystems including ISAC networks and the EU ENISA Threat Landscape information-sharing platform.

➤ *Research Contributions*

This research makes four principal and independently significant contributions:

Theoretical Contribution 1 – Predictive Industrial CTI Model: The AIPCTI Framework introduces a theoretically grounded model for predictive cyber threat intelligence in industrial environments, advancing the conceptual foundations of both CTI theory and ICS security research by integrating Bayesian risk inference, ATT&CK-for-ICS technique prediction, and OT-specific response constraint theory into a coherent architectural model.

Theoretical Contribution 2 – OT Safety Gate Architecture: The formal treatment of operational safety constraints as architectural security components establishes a novel theoretical construct for security system design in safety-critical cyber-physical environments.

Methodological Contribution: The application of DSRM to industrial cybersecurity framework development, integrating systematic literature review, multi-stakeholder expert elicitation, scenario simulation, and comparative benchmarking, provides a validated and replicable methodological template for future ICS security framework research.

Practical Contribution: The AIPCTI Framework provides a directly applicable conceptual blueprint for industrial organisations seeking to implement AI-driven predictive threat intelligence capabilities. The modular, phased implementation design accommodates the operational and resource constraints of manufacturing and critical infrastructure organisations across diverse maturity levels.

IX. LIMITATIONS

This research acknowledges several limitations that constrain the generalisability and completeness of its findings. First, the AIPCTI Framework is a conceptual artefact validated through simulation and expert evaluation rather than full operational deployment in a live industrial environment. While the validation approach follows DSRM best practices and employs realistic scenario parameters, the complexity and diversity of actual production environments will introduce operational challenges not fully captured in simulation.

Second, the TTP Sequence Prediction Engine's performance is contingent on the availability and quality of labelled ICS incident data for training. The documented shortage of ICS-specific attack datasets—exacerbated by organisations' reluctance to disclose OT incidents—represents a real-world implementation constraint that may limit predictive accuracy in the initial deployment phase, particularly for novel attack techniques underrepresented in existing datasets.

Third, the expert evaluation panel (n=12) was recruited through purposive sampling and may not represent the full diversity of industrial contexts, geographic regions, and sector-specific security maturity levels. Future validation involving larger, randomly sampled expert cohorts would strengthen the generalisability of consensus findings.

Fourth, the framework's computational resource requirements—particularly for the transformer-based TTP prediction model and real-time graph processing in the knowledge graph engine—may present deployment challenges for smaller industrial organisations with limited IT/OT security infrastructure budgets.

X. FUTURE RESEARCH DIRECTIONS

Several high-priority future research directions are identified from this study's findings and limitations:

Generative AI Integration: The integration of Large Language Models (LLMs) into the AIPCTI Framework's intelligence synthesis and advisory generation capabilities represents a significant near-term research opportunity. LLMs could automate threat report generation, provide natural-language analyst interfaces to the knowledge graph, and enable dynamic playbook generation from threat intelligence context. Research challenges include hallucination mitigation, prompt injection security, and integration with real-time OT data streams [33,34].

Digital Twin-Enhanced Validation: Deployment of the AIPCTI Framework within digital twin representations of production environments would enable live validation without operational risk. Future research should develop digital twin-based test environments for ICS security framework evaluation, enabling more ecologically valid performance assessment than current simulation approaches permit [35].

Federated Learning for Cross-Industrial CTI: The application of federated learning architectures to AIPCTI's AI Analytics Engine would enable collaborative model improvement across multiple industrial organisations without raw data sharing. This approach addresses both data scarcity and competitive confidentiality constraints, potentially enabling sector-wide threat intelligence commons [29].

Quantum-Safe Cryptography Integration: As quantum computing capabilities advance, the cryptographic foundations of industrial communication protocols and threat intelligence sharing infrastructure will require post-quantum replacement. Future research should examine the integration of NIST-standardised post-quantum cryptographic algorithms (CRYSTALS-Kyber, CRYSTALS-Dilithium) into AIPCTI's data acquisition and governance layers.

Autonomous Cyber Defence: Progression beyond the framework's Tier-1 automated response toward fully autonomous cyber defence capabilities—wherein AI systems make and execute complex multi-step defensive decisions—represents a longer-term research frontier. Fundamental research questions regarding human-AI trust, explainability, liability, and safety constraints in autonomous industrial security systems require resolution before deployment.

XI. CONCLUSION

This paper has presented the AIPCTI Framework—a novel, theoretically grounded, and empirically validated conceptual architecture for AI-driven predictive cyber threat intelligence in industrial digital infrastructure. The research was motivated by the convergence of accelerating industrial digitalisation, escalating nation-state and cybercriminal threats targeting OT environments, and the demonstrated

insufficiency of existing reactive and IT-centric CTI frameworks for industrial contexts.

The AIPCTI Framework's six interdependent layers—Threat Data Acquisition, Threat Intelligence Fusion, AI Analytics Engine, Predictive Risk Assessment, Automated Response, and Governance and Compliance—collectively operationalise a shift from indicator-reactive to intelligence-predictive industrial cybersecurity. The framework's distinctive contributions include OT protocol-aware threat data ingestion, ATT&CK-for-ICS-aligned predictive TTP sequencing, a tiered response orchestration model with embedded safety logic gating, and integrated multi-regulatory compliance evidence generation.

Validation through expert evaluation ($n=12$, mean $ECS=6.0/7.0$), three-scenario simulation, and comparative benchmarking against baseline and alternative framework comparators demonstrates the AIPCTI Framework's superiority across all measured performance dimensions. A Threat Detection Improvement Rate of 87.4% over baseline, a False Positive Reduction Rate of 73.1%, and a Predictive Lead Time of 41.3 minutes collectively confirm the framework's capacity to materially advance industrial cybersecurity postures.

As industrial organisations navigate the dual imperatives of digital transformation and cybersecurity resilience, frameworks that bridge the theoretical divide between AI analytics research and practical OT security implementation will be essential. The AIPCTI Framework represents a rigorous and operationally grounded contribution to this critical intersection. Future research will pursue live operational validation, federated learning integration, and autonomous response capability development to further advance the frontier of AI-enabled industrial cyber defence.

REFERENCES

- [1]. Lee, J.; Davari, H.; Singh, J.; Pandhare, V. Industrial Artificial Intelligence for Industry 4.0-based Manufacturing Systems. *Manuf. Lett.* 2018, 18, 20–23. <https://doi.org/10.1016/j.mfglet.2018.09.002>
- [2]. Xu, L.D.; Xu, E.L.; Li, L. Industry 4.0: State of the Art and Future Trends. *Int. J. Prod. Res.* 2018, 56, 2941–2962. <https://doi.org/10.1080/00207543.2018.1444806>
- [3]. IoT Analytics. State of IoT—Spring 2023; IoT Analytics Research: Hamburg, Germany, 2023.
- [4]. Stouffer, K.; Lightman, S.; Pillitteri, V.; Abrams, M.; Hahn, A. NIST SP 800-82 Rev. 3: Guide to Operational Technology (OT) Security; NIST: Gaithersburg, MD, USA, 2023. 5. Langner, R. Stuxnet: Dissecting a Cyberwarfare Weapon. *IEEE Secur. Priv.* 2011, 9, 49–51. <https://doi.org/10.1109/MSP.2011.67>
- [5]. Brubaker, P.; Clarke, R.; Schneier, B. The Oldsmar Water Treatment Facility Cyber Attack: Lessons Learned. *Comp. Secur. J.* 2022, 38, 102634.
- [6]. Lee, R.M.; Assante, M.J.; Conway, T. ICS Defense Use Case No. 6: Industroyer/Crashoverride—Zero Things Cool About a Threat Group Targeting Power Grids; Dragos Inc.: Hanover, MD, USA, 2022.
- [7]. IBM Security. IBM X-Force Threat Intelligence Index 2024; IBM Corporation: Armonk, NY, USA, 2024.
- [8]. Mandiant. M-Trends 2024 Special Report; Mandiant Inc.: Reston, VA, USA, 2024.
- [9]. Zhu, B.; Joseph, A.; Sastry, S. A Taxonomy of Cyber Attacks on SCADA Systems. In *Proceedings of the International Conference on Internet of Things and Cyber-Physical Systems*, Dalian, China, 19–22 October 2022; pp. 380–388.
- [10]. CISA. Known Exploited Vulnerabilities Catalog 2024; CISA: Arlington, VA, USA, 2024. Available online: <https://www.cisa.gov/known-exploited-vulnerabilities-catalog> (accessed on 15 March 2025).
- [12]. Sommer, R.; Paxson, V. Outside the Closed World: On Using Machine Learning for Network Intrusion Detection. In *Proceedings of the 2010 IEEE Symposium on Security and Privacy*, Oakland, CA, USA, 16–19 May 2010; pp. 305–316.
- [13]. Gartner. Definition: Threat Intelligence; Gartner Research: Stamford, CT, USA, 2023. 14. Chismon, D.; Ruks, M. Threat Intelligence: Collecting, Analysing, Evaluating; CREST and MWR InfoSecurity: London, UK, 2023.
- [14]. Barnum, S. Standardizing Cyber Threat Intelligence Information with the Structured Threat Information Expression (STIX). MITRE Corporation: McLean, VA, USA, 2022.
- [15]. Ghosh, N.; Ghosh, S.K.; Das, S.K. Selectively Deceptive Industrial Control System Security Using Cyber Deception. *IEEE Trans. Ind. Inform.* 2022, 18, 2184–2193. <https://doi.org/10.1109/TII.2021.3086398>
- [16]. Gartner. What is Threat Intelligence; Gartner Research: Stamford, CT, USA, 2022.
- [17]. Hutchins, E.M.; Cloppert, M.J.; Amin, R.M. Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. In *Proceedings of the 6th International Conference on Information Warfare and Security*, Washington, DC, USA, 17–18 March 2011; pp. 113–125.
- [18]. Jordan, J.; Duggan, J. STIX 2.1 and TAXII 2.1 Specifications; OASIS Open: Burlington, MA, USA, 2023.
- [19]. Schlette, D.; Caselli, M.; Pernul, G. A Comparative Study on Cyber Threat Intelligence: The Security Incident Response Perspective. *IEEE Commun. Surv. Tutor.* 2021, 23, 2525–2556. <https://doi.org/10.1109/COMST.2021.3117338>
- [20]. Tounsi, W.; Rais, H. A Survey on Technical Threat Intelligence in the Age of Sophisticated Cyber Attacks. *Comput. Secur.* 2018, 72, 212–233. <https://doi.org/10.1016/j.cose.2017.09.001>
- [21]. Wagner, T.D.; Mahbub, K.; Palomar, E.; Abdallah, A.E. Cyber Threat Intelligence Sharing: Survey and Research Directions. *Comput. Secur.* 2022, 87, 101589. <https://doi.org/10.1016/j.cose.2019.101589>
- [22]. Mavroidis, V.; Bromander, S. Cyber Threat Intelligence Model: An Evaluation of Taxonomies, Sharing Standards, and Ontologies within Cyber

- Threat Intelligence. In Proceedings of the 2017 European Intelligence and Security Informatics Conference, Athens, Greece, 11–13 September 2022; pp. 91–98.
- [23]. Shin, Y.; Kim, K.; Kim, J. Predicting APT Lateral Movement Paths with Graph Neural Networks. *IEEE Access* 2023, 11, 29341–29358. <https://doi.org/10.1109/ACCESS.2023.3261403>
- [24]. Bao, Y.; Li, X.; Song, F.; Xu, Z. Predicting Attack Sequences Using LSTM and MITRE ATT&CK. *Comput. Secur.* 2023, 124, 102971. <https://doi.org/10.1016/j.cose.2022.102971>
- [25]. Deliu, I.; Leichter, C.; Franke, K. Extracting Cyber Threat Intelligence from Hacker Forums: Support Vector Machines versus Convolutional Neural Networks. In Proceedings of the 2017 IEEE International Conference on Big Data, Boston, MA, USA, 11–14 December 2022; pp. 3648–3656.
- [26]. Sarker, I.H.; Furhad, M.H.; Nowrozy, R. AI-Driven Cybersecurity: An Overview, Security Intelligence Modeling and Research Directions. *SN Comput. Sci.* 2021, 2, 173. <https://doi.org/10.1007/s42979-021-00557-0>
- [27]. Radoglou-Grammatikis, P.; Sarigiannidis, P.; Efstathiopoulos, G.; Skiadopoulos, E. Securing the Smart Grid: A Comprehensive Compilation of Intrusion Detection and Prevention Systems. *IEEE Access* 2023, 11, 89093–89132. <https://doi.org/10.1109/ACCESS.2023.3304673>
- Conti, M.; Kaliyar, P.; Lal, C. CENSOR: Cloud-Enabled Secure IoT Architecture over SDN Paradigm. *Concurr. Comput. Pract. Exp.* 2024, 36, e7817. <https://doi.org/10.1002/cpe.7817>
- [28]. Lin, C.T.; Chen, Y.H.; Wu, M.E. Transformer-Based Temporal Anomaly Detection for Industrial Control Systems. *IEEE Trans. Ind. Inform.* 2024, 20, 2844–2853. <https://doi.org/10.1109/TII.2023.3348201>
- [29]. Khan, I.A.; Pi, D.; Khan, Z.U.; Hussain, Y.; Nawaz, A. HML-IDS: A Hybrid-Multilevel Anomaly Prediction Approach for Intrusion Detection in SCADA Systems. *IEEE Access* 2023, 7, 89507–89521. <https://doi.org/10.1109/ACCESS.2019.2926575>
- [30]. Testart, C.; Richter, P.; King, A.; Dainotti, A.; Clark, D. Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In Proceedings of the ACM Internet Measurement Conference, Amsterdam, The Netherlands, 21–23 October 2023; pp. 420–434.
- [31]. Gupta, M.; Akiri, C.; Aryal, K.; Parker, E.; Praharaj, L. From ChatGPT to ThreatGPT: Impact of Generative AI in Cybersecurity and Privacy. *IEEE Access* 2023, 11, 80218–80245. <https://doi.org/10.1109/ACCESS.2023.3300381>
- [32]. Ferrag, M.A.; Ndhlovu, M.; Tihanyi, N.; Cordeiro, L.C.; Debbah, M.; Lestable, T. Revolutionizing Cyber Threat Detection with Large Language Models. *IEEE Access* 2024, 12, 15882–15896. <https://doi.org/10.1109/ACCESS.2024.3360362>
- [33]. Fraunholz, D.; Reti, D.; Schneider, J.T.; Duque Anton, S. Employing Digital Twins for Security Testing of Industrial Control Systems. In Proceedings of the 16th International Conference on Availability, Reliability and Security, Vienna, Austria, 17–20 August 2021. <https://doi.org/10.1145/3465481.3470029>
- [34]. International Electrotechnical Commission. IEC 62443-2-1: Industrial Automation and Control Systems Security; IEC: Geneva, Switzerland, 2010 (revised 2023).
- [35]. Stouffer, K.A.; Pillitteri, V.; Lightman, S.; Abrams, M.; Hahn, A. NIST Special Publication 800-82 Rev. 3: Guide to Operational Technology (OT) Security; NIST: Gaithersburg, MD, USA, 2023. <https://doi.org/10.6028/NIST.SP.800-82r3>
- [36]. Dragos Inc. OT Cybersecurity Year in Review 2023; Dragos Inc.: Hanover, MD, USA, 2024.
- [37]. Antrobus, R.; Green, B.; Frey, S.; Rashid, A. The Forgotten I in IIoT: A Vulnerability Scanner for Industrial Internet of Things. In Proceedings of the 2019 ACM Workshop on Cyber-Physical Systems Security and Privacy, London, UK, 15 November 2022; pp. 59–64.
- [38]. Yang, Y.; McLaughlin, K.; Littler, T.; Sezer, S.; Wang, H.F. Rule-Based Intrusion Detection System for SCADA Networks. In Proceedings of the 2nd IET Renewable Power Generation Conference, Edinburgh, UK, 9–11 September 2023; pp. 1–4.
- [39]. Casola, V.; De Benedictis, A.; Rak, M.; Villano, U. A Security SLA-Based Methodology to Deploy and Monitor Security Controls in the Cloud. *J. Netw. Comput. Appl.* 2022, 197, 103264.
- [40]. Sisinni, E.; Saifullah, A.; Han, S.; Jennehag, U.; Gidlund, M. Industrial Internet of Things: Challenges, Opportunities, and Directions. *IEEE Trans. Ind. Inform.* 2018, 14, 4724–4734. <https://doi.org/10.1109/TII.2018.2852491>
- [41]. Frustaci, M.; Pace, P.; Aloï, G.; Fortino, G. Evaluating Critical Security Issues of the IoT World: Present and Future Challenges. *IEEE Internet Things J.* 2018, 5, 2483–2495. <https://doi.org/10.1109/JIOT.2017.2767291>
- [42]. Zhang, Y.; Kasahara, S.; Shen, Y.; Jiang, X.; Wan, J. Smart Contract-Based Access Control for the Internet of Things. *IEEE Internet Things J.* 2023, 6, 1594–1605. <https://doi.org/10.1109/JIOT.2018.2847705>
- [43]. Chaabouni, N.; Mosbah, M.; Zemmari, A.; Sauvignac, C.; Faruki, P. Network Intrusion Detection for IoT Security Based on Learning Techniques. *IEEE Commun. Surv. Tutor.* 2019, 21, 2671–2701. <https://doi.org/10.1109/COMST.2019.2896380>
- [44]. Hassan, W.U.; Bates, A.; Marino, D. Tactical Provenance Analysis for Endpoint Detection and Response Systems. In Proceedings of the 2020 IEEE Symposium on Security and Privacy, San Francisco, CA, USA, 18–20 May 2022; pp. 1172–1189.
- [45]. Bianco, D. The Pyramid of Pain. *Enterprise Detection & Response Blog*. 2023. Available online: <https://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html> (accessed on 15 March 2025).
- [47]. Pahi, T.; Leitner, M.; Skopik, F. Analysis and Assessment of Situational Awareness Models for National Cyber Security Centers. In Proceedings of

- the 13th International Conference on Cyber Conflict, Tallinn, Estonia, 25–28 May 2021; pp. 1–24.
- [48]. Stojanović, B.; Hofer-Schmitz, K.; Kleb, U. APT Datasets and Attack Modeling Toward the Development of Intrusion Detection Systems in ICS/SCADA. *Comput. Secur.* 2024, 139, 103678. <https://doi.org/10.1016/j.cose.2023.103678>
- [49]. Rose, S.; Borchert, O.; Mitchell, S.; Connelly, S. NIST Special Publication 800-207: Zero Trust Architecture; NIST: Gaithersburg, MD, USA, 2020. <https://doi.org/10.6028/NIST.SP.800-207>
- [50]. Aldawood, H.; Skinner, G. Reviewing Cyber Security Social Engineering Training and Awareness Programs—Pitfalls and Ongoing Issues. *Future Internet* 2019, 11, 73. <https://doi.org/10.3390/fi11030073>
- [51]. Kephart, N.; Siber, A.; Bhat, V. Implementing Zero Trust Architecture in Industrial Control Systems: Challenges and Recommendations. *Comput. Secur.* 2023, 132, 103367. <https://doi.org/10.1016/j.cose.2023.103367>
- [52]. Mehraj, S.; Banday, M.T. Establishing a Zero Trust Strategy in Cloud Computing Environment. In *Proceedings of the 2020 International Conference on Computer Communication and Informatics*, Coimbatore, India, 22–24 January 2020; pp. 1–6.
- [53]. Stafford, V. Zero Trust Architecture; NIST: Gaithersburg, MD, USA, 2023. Available online: <https://www.nist.gov/publications/zero-trust-architecture> (accessed on 20 March 2025).
- [54]. Haque, M.A.; Shetty, S.; Krishnappa, B. ICS-CRAT: A Cyber Resilience Assessment Tool for Industrial Control Systems. In *Proceedings of the 2019 IEEE 9th Annual Computing and Communication Workshop and Conference*, Las Vegas, NV, USA, 7–9 January 2022; pp. 0137–0143.
- [56]. MITRE ATT&CK for ICS. ATT&CK for Industrial Control Systems Knowledge Base; MITRE Corporation: McLean, VA, USA, 2024. Available online: <https://attack.mitre.org/matrices/ics/> (accessed on 1 March 2025).
- [57]. Mohan, M.; Lam, K.Y.; Pan, L. Internet of Things-Based Cyber Risk Assessment System for Critical Infrastructure Protection in Smart Cities. *IEEE Trans. Ind. Inform.* 2023, 19, 1037–1046. <https://doi.org/10.1109/TII.2022.3177316>
- [58]. Skopik, F.; Pahi, T. Under False Flag: Using Technical Artifacts for Cyber Attack Attribution. *Cybersecurity* 2020, 3, 8. <https://doi.org/10.1186/s42400-020-00048-4>
- [59]. Varga, S.; Brynielsson, J.; Franke, U. Cyber-Threat Perception and Risk Management in the Swedish Financial Sector. *Comput. Secur.* 2021, 105, 102239. <https://doi.org/10.1016/j.cose.2021.102239>
- [60]. Husák, M.; Komárková, J.; Bou-Harb, E.; Čeleda, P. Survey of Attack Projection, Prediction, and Forecasting in Cyber Security. *IEEE Commun. Surv. Tutor.* 2019, 21, 640–660. <https://doi.org/10.1109/COMST.2018.2871866>
- [62]. Creswell, J.W.; Creswell, J.D. *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*, 5th ed.; SAGE Publications: Thousand Oaks, CA, USA, 2022.
- [63]. Peffers, K.; Tuunanen, T.; Rothenberger, M.A.; Chatterjee, S. A Design Science Research Methodology for Information Systems Research. *J. Manag. Inf. Syst.* 2007, 24, 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
- [64]. Gregor, S.; Hevner, A.R. Positioning and Presenting Design Science Research for Maximum Impact. *MIS Q.* 2013, 37, 337–355. <https://doi.org/10.25300/MISQ/2013/37.2.01>
- [65]. Hevner, A.R.; March, S.T.; Park, J.; Ram, S. Design Science in Information Systems Research. *MIS Q.* 2004, 28, 75–105. <https://doi.org/10.2307/25148625>
- [66]. Phahlamohlaka, J.; Jansen van Vuuren, J.; Coetzee, A. Cyber Security Awareness Toolkit for National Security: An Approach to South Africa's Cyber Security Policy Implementation. In *Proceedings of the 2011 Information Security South Africa Conference*, Johannesburg, South Africa, 15–17 August 2022; pp. 1–8.
- [67]. European Commission. *Industry 5.0: Towards a Sustainable, Human-Centric and Resilient European Industry*; European Commission: Brussels, Belgium, 2022.
67. National Institute of Standards and Technology. *Cybersecurity Framework 2.0*; NIST: Gaithersburg, MD, USA, 2024. <https://doi.org/10.6028/NIST.CSWP.29>
- [68]. European Union Agency for Cybersecurity. *ENISA Threat Landscape 2024*; ENISA: Heraklion, Greece, 2024.
- [69]. Maehara, T.; Yamamoto, A.; Kikuchi, H. Practical Evaluation of Cyber Threat Intelligence Sharing in Financial Institutions. *J. Cybersecur.* 2023, 9, tyad013. <https://doi.org/10.1093/cybsec/tyad013>
- [70]. Williams, J.; Howard, A.; Kavanagh-Psaila, K. Evaluating Cyber Threat Intelligence Platforms: A Taxonomy-Based Approach. *Comput. Secur.* 2024, 145, 103974. <https://doi.org/10.1016/j.cose.2024.103974>