Cyber-Physical Incident Attribution in UAV/Rail Attacks

Isaac Ojeh[1], Xavier Palmer[2] and Lucas Potter[2]
[1]MorphHats InfoSecure, Waterloo, Canada
[2]BiosView Labs, Dayton, Ohio

morpheus@morphhats.com
Biosview1@proton.me

Abstract: As unmanned aerial vehicles (UAVs) and smart rail systems become increasingly integrated into critical logistics infrastructure, they also present new surfaces for hybrid cyber-physical attacks. Coordinated adversarial actions such as cyber intrusions that manipulate physical trajectories, sensor spoofing, or disruptions to control systems pose significant challenges for real-time detection and post-incident analysis. Effective attribution of such incidents is crucial not only for identifying responsible parties but also for enhancing resilience and enabling coordinated defense responses across infrastructure operators, governments, and private stakeholders. This paper examines the problem of cyber-physical incident attribution in the context of combined UAV/rail attacks, where attack vectors may span networked systems, edge devices, and physical actuators. We propose a layered attribution framework that fuses telemetry from cyber logs, UAV flight data, rail signaling systems, and environmental sensors to reconstruct the sequence and origin of coordinated attacks. The system leverages graph-based causality analysis, trust scoring mechanisms, and cross-domain forensic correlation to associate anomalies with likely sources and attack pathways. Our approach combines both deterministic rules and machine learning models trained on simulated and real-world incident data to balance explainability and adaptive intelligence. This paper, overall, adopts an exploratory perspective, examining foundational challenges and design trade-offs involved in attributing cyber-physical incidents within a multimodal UAV/rail logistics environment. Rather than proposing a finalized solution, the work seeks to identify key data-fusion requirements, threat-modeling gaps, and policy implications to inform future technical and legal frameworks for attribution. Preliminary results from simulated hybrid attacks using a digital twin environment show promising attribution accuracy, particularly when incorporating temporal patterns and system interdependencies. However, limitations in sensor coverage and adversarial evasion tactics underscore the need for multi-source trust validation and international collaboration in standardizing attribution protocols. Ultimately, this research aims to lay groundwork for a scalable, context-aware attribution system that can support accountability, deterrence, and rapid response in the evolving landscape of autonomous transportation.

Keywords: Cyber physical security, UAV attacks, Rail infrastructure, Incident attribution, Forensic analysis, Hybrid threats, Digital twins, Critical infrastructure defense

1. Introduction

Some modern logistics can benefit from the integration of cyber-physical systems such as UAVs and smart railway networks; this can enhance efficiency but can also carry the risk of expanding attack surfaces. Coordinated cyber-physical incidents, such as a cyber intrusion into rail signaling combined with drone hijacking, are hard to detect and attribute (Longo et al. 2025). Attribution, or identifying responsibility, can be made more complex in these multi-domain scenarios. Traditional cybersecurity already faces hurdles, and hybrid UAV-rail attacks amplify this due to disparate systems (IT, wireless, mechanical) and unclear intrusion signs (Hartmann and Giles, 2016). Effective attribution is crucial for accountability, deterrence and improved incident response (Rid and Buchanan 2015). Despite its importance, cyber-physical incident attribution is underdeveloped. Existing railway cybersecurity and UAV security research focus on prevention, not sufficiently on post-incident attribution (Ibadah et al. 2024; Soderi et al. 2023). This gap stems from siloed monitoring, incompatible data, and a lack of cross-domain forensic standards. This paper aims to address that through proposing a layered attribution framework for UAV-rail incidents. It fuses data from various sources (network logs, UAV telemetry, sensors, surveillance) to build an event timeline, inferring attack pathways and entities through cause-and-effect analysis. Techniques include graph-based causality modeling and machine learning for pattern recognition (Qiu et al. 2024). Our exploratory work outlines the framework's architecture, discusses challenges (data integrity, real-time analysis, explainability), and presents a preliminary evaluation using a digital twin. Results show that correlating cyber and physical telemetry significantly improves attribution accuracy, identifying malicious events across domains where isolated analysis failed. Limitations include sensor blind spots and adversarial data injection. This research aims to build foundational, context-aware attribution systems for critical autonomous transportation. As railways integrate IoT and drones become ubiquitous, the security community must preempt attackers exploiting these convergences. Lessons learned will inform technical designs and foster collaborations for collective defense. The following section reviews related work and threats in UAV/rail environments.

1.1 Definition of Cyber-Physical Incident Attribution

In this paper, cyber-physical incident attribution is defined as the structured process of identifying the likely origin, pathway, and responsible entities behind a coordinated cyber-physical attack by correlating technical evidence across digital and physical domains. This definition is intentionally scoped to technical and operational attribution rather than legal adjudication. Attribution in cybersecurity literature often spans multiple contexts. In technical cybersecurity, attribution refers to identifying the infrastructure, tools, and tactics used in an attack through forensic analysis and threat intelligence (Rid and Buchanan, 2015). In criminal law, attribution establishes an individual or group's responsibility for prosecution. In contrast, international law determines whether a state is accountable for wrongful acts committed directly or through proxies. While these legal interpretations are critical, they fall outside the scope of this study. This work focuses specifically on technical attribution within cyber-physical systems, aiming to reconstruct attack sequences, correlate cross-domain evidence, and support investigative decision-making. Legal responsibility and political assessment are treated as downstream processes that may build upon the technical attribution outputs generated by the proposed framework.

2. Background and Related Work

Cyber-physical attacks, such as Stuxnet (2010), show that digital actions can have physical consequences. In rail, incidents range from malware to ransomware, exposing vulnerabilities (Karnouskos, 2011). UAVs pose new threats, with drones interfering with airports or being hijacked. Both rail and UAV domains face compounding cyber-physical risks. Modern rail systems are industrial control systems (ICS). Securing rail ICS is crucial due to safety and economic implications. While standards exist, gaps remain in covering legacy equipment and integrating physical safety with cyber threat detection. Traditional safety systems assume benign failures, necessitating new approaches that combine operational and cybersecurity data for incident analysis. UAVs are flying cyber-physical systems. UAV platforms share architectural similarities with industrial control systems, making them susceptible to comparable cyber threats. Research has identified threats such as GPS spoofing, command-link hijacking, and sensor manipulation in UAV platforms (Kim et al. 2021). Many UAV platforms share ICS characteristics, enabling the adaptation of industrial cybersecurity techniques. However, UAVs introduce unique challenges as they are mobile, operate in unregulated airspace, and can reach sensitive locations. Attribution in cyberattacks is a studied problem that involves tracing activity through digital forensics. Cyber-physical incident attribution is a newer approach that extends digital forensics to include physical causal reasoning. Frameworks are emerging that correlate cyber and physical event patterns to distinguish equipment failures from cyberattacks. Cross-domain correlation techniques have been shown to distinguish cyberattacks from physical faults in power systems. Cross-domain correlation of cyber and physical events has been shown to improve the distinction between system faults and deliberate cyberattacks (Semertzis et al. 2024). Digital forensics readiness is critical, and methods such as digital twins improve accident investigation accuracy. Maintaining extensive logs and synchronized, trustworthy data is essential for post-incident analysis. The convergence of UAV and rail systems is practical, with drones used for inspection and security. For instance, Sensonic uses fiber optic sensors to detect anomalies and dispatch drones for visual confirmation. This integration, while improving response, also creates potential attack vectors that adversaries could exploit to manipulate sensors and hijack drones. Commercial rail operators are already integrating drones with track monitoring systems, increasing both resilience and attack surface. Rail operators have begun integrating drones with track monitoring and security systems, expanding both operational resilience and the attack surface (Pyke 2025). While no sufficient incidents of this nature have been reported, the individual components are vulnerable, suggesting multi-vector attacks are a plausible future threat. In conclusion, cyber and physical aspects of critical infrastructure are increasingly interdependent. Our work aims to bridge the gap in investigative and attribution methods for UAV/rail contexts, drawing on insights from ICS security, drone forensics, and multi-stage attack analysis.

2.1 Indicators of Compromise and Indicators of Attack

Cyber-physical attribution relies on multiple classes of indicators. Traditional cyber forensics emphasize Indicators of Compromise (IoCs), which are technical artifacts such as malicious IP addresses, file hashes, or anomalous log entries that indicate a system has been compromised (Haber and Rolls, 2019; Haber and Rolls, 2024). IoCs are typically reactive and focus on identifying what was affected. In contrast, Indicators of Attack (IoAs) are behavioral and contextual indicators that describe how and why an attack is unfolding (Hwang and Kim, 2023; Yi and Kim, 2024). IoAs capture patterns such as coordinated timing across systems, repeated manipulation of control logic, or consistent sequencing of cyber and physical actions. These indicators are

particularly relevant for attribution, as they provide insight into adversary intent and operational design rather than isolated technical artifacts. The framework proposed in this paper integrates both IoCs and IoAs. IoCs support detection and forensic validation, while IoAs inform causal graph construction and hypothesis generation, enabling attribution of coordinated cyber-physical attacks beyond simple compromise detection.

2.2 Hybrid Cyber-Physical Attacks in UAV and Rail Systems

Hybrid cyber-physical attacks are defined here as coordinated adversarial actions in which cyber operations are deliberately used to induce, amplify, or conceal physical effects (Simons et al, 2020; Lian et al, 2024). In UAV and rail environments, such attacks often involve multiple stages, including digital intrusion, control system manipulation, sensor spoofing, and physical actuation. Examples include unauthorized modification of rail signaling logic combined with UAV-based reconnaissance or obstruction, GPS spoofing used to redirect drones while masking physical intrusions, or denial-of-service attacks that delay operator response during a physical sabotage event. These attacks differ from purely cyber or purely physical incidents in that their effects emerge from the interaction between digital commands and physical system behavior. Understanding these hybrid attack structures is essential for attribution, as isolated analysis of either cyber logs or physical anomalies is often insufficient to reconstruct the whole attack sequence.

3. Layered Attribution Framework

3.1 Design Objectives and Overview

The primary goal of the proposed framework is to reconstruct the sequence of events in a cyber-physical attack and identify the responsible entities (be it specific attackers, malware, or compromised devices) with high confidence. To achieve this, the framework is designed with several objectives in mind. To start, a comprehensive framework is essential, capturing data from all relevant layers, including network traffic, system logs, application-level events, physical sensor outputs, and operator actions, as no single vantage point suffices for complex incidents. For real-time or near-real-time responses, timeliness is crucial, requiring the framework to process and correlate events as they happen and to raise attribution hypotheses even before an incident entirely unfolds; however, offline forensic analysis is also supported for thorough post-incident investigations. The framework must also assess the trustworthiness and integrity of data sources, recognize when an attacker might have tampered with logs or sensor readings, and incorporate redundant or independent data sources to help validate events. Furthermore, the results of the analysis of the attribution should be explainable and interpretable by humans, especially in critical infrastructure, where accountability is key. They should provide a chain of supporting evidence or reasoning rather than just an actor's attribution. These principles lay the groundwork for a system that can effectively adapt to evolving threats.

Table 1: Layered Attribution Framework for UAV and Rail Incidents

Layer	Purpose	Example Data Sources
Data Collection	Capture cyber and physical telemetry	Rail signaling logs, UAV flight data, network flow logs
Event Processing	Detect anomalies and normalize data	Safety rule violations, ML-based anomaly alerts
Causality Analysis	Correlating events across domains	Temporal and causal event graphs
Attribution Inference	Generate and rank attack hypotheses	IoC and IoA correlation, trust-weighted reasoning
Investigation Interface	Support analyst review and reporting	Timelines, evidence chains, hypothesis summaries

Our proposed layered architecture, depicted in Figure 1, includes several modules designed to achieve our objectives. At the foundational data collection layer, interfaces gather information from various sources. These sources include UAV telemetry streams, such as flight path, controller inputs, GPS data, and payload status, as well as rail system logs, which encompass signaling commands, track sensor triggers, and train control system data. Additionally, network monitors collect traffic from relevant networks, potentially including attacker command-and-control communications. Physical surveillance, including CCTV footage and distributed-acoustic-sensing fiber alerts, also contributes to the data. Finally, environmental data, such as GPS timing signals and weather conditions, is crucial for ruling out natural causes, such as solar storms or GPS errors. The system processes data in several layers. First, Event Processing and Anomaly Detection analyze data streams using rule-based and algorithmic methods to flag unusual events, applying domain-specific knowledge (Ibadah et al. 2024). Next, the Causality Graph Construction module links timestamped events from various sources into a graph where nodes are events/states and edges represent causal/temporal relationships. This approach, inspired by attack graph methodologies, establishes edges through deterministic logic or learned patterns (Qiu et al. 2024).

Each graph element carries a trust score that reflects sensor reliability. If a sensor is compromised, its data receives lower confidence unless corroborated. The Attribution Inference Engine then analyzes event clusters and matches them to known attack signatures or threat actors. This involves hypothesis generation, potentially using knowledge bases such as MITRE ATT&CK to label tactics, techniques, and procedures (TTPs), which can hint at attribution (Strom et al., 2018). Finally, an Investigation Interface provides human analysts with visualizations, causal graphs, and narrative explanations, linking evidence to conclusions. It also allows analysts to verify data for legal defensibility.

3.2 Graph-Based Causality Analysis

Our framework's foundation lies in the construction of an event-causality graph, an approach supported by previous research that has used graph models to untangle intricate incident chains (Casey et al. 2015; Qiu et al. 2024). Within this graph, we delineate several node types. These include Cyber Event Nodes, such as "Firewall log entry: port scan detected" or "PLC command issued: switch track," and Physical Event Nodes, such as "Drone altitude change" or "Train emergency brake activated." Additionally, Alert/Anomaly Nodes represent triggers from detection algorithms, like an "Intrusion Detection System alert." At the same time, State Nodes capture relevant system states or conditions, for instance, "Train A was in section X at time T." Understanding these node types is crucial for accurately mapping the sequence of events during a cyber-physical incident. Edges show causal or temporal links. Causal links indicate direct precipitation, while temporal edges suggest ordering and potential influence. Machine learning also identifies correlated events, like network latency anomalies linked to sensor faults. To manage incident complexity, we layer graphs, maintaining subgraphs per system (e.g., UAV, rail) with cross-links. This allows analysts to view simplified slices of an incident (e.g., drone vs. train events) and their convergence in the larger picture, like network attack graph layers. Graph-based analysis systematically evaluates attack paths and assesses the likelihood of each. We use metrics like out-degree and betweenness to identify critical pivot points, similar to Qiu et al. (2024). A single malicious packet causing downstream failures would appear as a high-influence node, highlighting the attacker's pivotal action.

3.3 Trust Scoring Mechanism

Our framework's reliance on fusing data from multiple sources necessitates considering the potential compromise or unreliability of some sources. For instance, if an attacker gains administrative access to a control server, the server's logs could be manipulated to conceal their activities. Unquestioning trust in all logs could lead to erroneous conclusions, such as attributing a malfunction when, in fact, a malicious command was deliberately deleted from the log. We propose a straightforward trust-scoring mechanism: each data source, and, by extension, every event originating from it, is assigned a numerical trust level. This assignment can be static, based on prior assessments of the source's security hardening, or dynamic, depending on whether the system exhibits normal behavior during an incident. For example, an independent GPS sensor on a drone might be deemed more trustworthy than the drone's internal GPS if GPS spoofing is suspected, as a tamper-proof independent sensor could detect discrepancies. Similarly, if a rail control center is isolated and highly secured, its logs might be assigned to high trust; however, if abnormal behavior or an intrusion alert is triggered during an incident, confidence in its data would consequently diminish. In practice, trust scores influence analysis by weighing the edges in the causality graph. Trust-weighted sensor fusion has been proposed as a mechanism for improving reliability in cyber-physical systems. Trust-weighted sensor data fusion improves reliability in cyber-physical attribution by accounting for compromised or unreliable data sources (Yadav et al. 2023). An event chain dependent on low-trust data will be regarded with skepticism unless corroborated. The inference engine can present alternative hypotheses: one assuming data accuracy, another assuming data fabrication. By comparing these, investigators can determine whether an attacker may have concealed their tracks. In essence, our framework endeavors to achieve forensic robustness by avoiding exclusive reliance on a single data source.

3.4 Machine Learning Components

While rule-based detection and expert logic are fundamental for explainability, machine learning (ML) offers the adaptability and pattern recognition necessary to address emerging threats. We apply ML in two key areas: anomaly detection and attack classification/attribution. For anomaly detection, we deploy unsupervised or semi-supervised learning algorithms, such as clustering, autoencoders, or one-class SVMs, to model the normal behavior of UAV operations and rail systems, subsequently flagging outliers. For instance, a deep learning model could learn the typical correlation between train speed profiles and signaling commands, where any deviation might suggest interference. Regarding attack classification/attribution, we train supervised models using simulated incident data to categorize attack types and even link them to known attacker profiles. For example, if a particular threat actor consistently employs a specific sequence of actions (e.g., jamming GPS followed by

injecting malware), a classifier can identify this pattern in incoming data and label it accordingly (e.g., "Pattern resembles Attack Group A's known operations"). This information can then contribute to attribution hypotheses. It is crucial to thoroughly test these ML models in the simulation environment to prevent overfitting to expected patterns and to ensure they can handle real-world noise and variance. Furthermore, by integrating deterministic rules with ML outputs, we aim to prevent "black-box" decision-making. ML suggestions are evaluated alongside logical rules (e.g., "if a valid operator command was never received, the train should not have accelerated; if it did, that's a strong indication of malicious control"). Essentially, we do not permit the system to make an attribution claim based solely on an ML output; ML assists in prioritizing and guiding the search for supporting evidence, but final conclusions require corroboration through the causal graph.

3.5 Data Fusion and Cross-Domain Correlation

Our framework's core is data fusion, correlating events across domains to uncover insights that single-domain analysis misses. For example, simultaneous rail signal errors and drone connection loss, while seemingly isolated, strongly suggest a linked cause like electromagnetic interference or a cyberattack. Our system automatically identifies these temporal and contextual alignments. To achieve this, all data is time-stamped with a common reference clock (NTP/GPS for UAV and rail systems) and normalized into a unified schema (time, location, severity, etc.). Distributed system techniques handle clock drift. Cross-domain correlation also identifies higher-level patterns. Multiple anomalies across subsystems in a short period often indicate a broader attack. Our heuristic: "if X or more subsystems report significant anomalies within Y minutes, suspect a coordinated attack," mirrors observations in past power grid incidents. In summary, this layered framework systematically transforms raw data into a coherent attack narrative and attribution.

4. Challenges and Trade-Offs in Attribution

Designing an attribution system for cyber-physical incidents presents several inherent challenges, including managing high data volumes while ensuring real-time analysis. Our framework addresses this by filtering relevant or abnormal data from trains and UAVs to reduce load; however, this risks false negatives, which we mitigate by sampling normal data and updating ML models. Another challenge involves balancing deterministic rules and machine learning, where we prioritize ML's strong indication of an attack over rules, accepting more false positives for critical infrastructure safety. To provide human-understandable explanations for ML outputs despite their complexity, we utilize causal graphs to link anomalies and attack classifications to evidence, such as sensor deviations or code signatures, ensuring ML informs without replacing evidence. Furthermore, we account for sensor and coverage gaps by simulating data gaps and flagging "insufficient data" in the analysis, acknowledging uncertainty while planning future mitigation through redundant sensors and data sharing. Although false flags and deception are not entirely solvable, multi-source correlation helps expose inconsistencies in fabricated attack scenarios, and we err on the side of caution, suspecting attacks if other anomalies are detected concurrently. Finally, privacy and legal considerations are addressed by ensuring data integrity through timestamping, cryptographic signing, and tamper-evident storage, and by filtering or anonymizing personally identifiable information; compliance also necessitates transparent governance on data usage. Despite these challenges, integrated attribution demonstrates promising benefits, as detailed in our simulation-based evaluation.

5. Evaluation in a Simulated Environment

We evaluated our framework using a digital twin simulation of a logistics corridor with a railway and UAVs. Digital twin approaches have also been applied to UAV forensic analysis, improving post-incident reconstruction accuracy. The framework was evaluated through a digital twin simulation that combined a rail control system and UAV operations (Almusayli et al. 2024). The simulation combined an open-source train control simulator with ROS 2 and Gazebo for UAVs, mimicking realistic signaling and drone operations, including communication networks. We designed attack scenarios based on threat models. One, "Coordinated Derailment Attempt," involved an attacker penetrating the rail network, changing a track switch to derail a train, jamming a drone's GPS, hijacking a surveillance drone to drop a payload onto the tracks, and launching a denial-of-service attack. The attacker also attempted to erase system logs. Our attribution framework, connected in real time, ingested data such as train speed, signal states, drone telemetry, and network logs. It successfully detected anomalies, correlating the unauthorized track switch command with erratic drone GPS and flight deviation, flagging a "Possible coordinated attack: rail switch manipulation and drone deviation concurrently observed. Likely sabotage attempt." It identified log tampering and, using its knowledge base, categorized the incident as a deliberate cyber-physical attack. In 9 out of 10 runs, the framework correctly associated the rail and UAV

anomalies. Another "False Flag" scenario tested the framework's ability to identify technical events while providing two possible hypotheses (malware signature vs. behavior pattern) to the analyst, aligning with our design principle of aiding human judgment. Performance metrics showed the system processed around 500 events per second, sufficient for simulated scenarios. Limitations included the variability of real-world systems and the challenge of tuning anomaly detectors. Cascade effects from a train derailment also temporarily overwhelmed the causal graph, highlighting the need for better abstraction. Nonetheless, the simulation confirmed the approach's viability, significantly improving situational awareness by integrating data that would traditionally be siloed. The use of a digital twin environment enabled controlled experimentation with hybrid attack scenarios, allowing repeatable testing of attribution accuracy under varying conditions without risk to operational infrastructure.

6. Discussion, Limitations, and Future Works

Our framework shows promise for cyber-physical attribution in UAV and rail incidents. Robust attribution requires multidisciplinary collaboration, including IT security, rail engineers, drone operators, and forensic analysts. Organizations must foster cross-domain teams that integrate IT, OT, and UAV experts, especially as they confront an increasingly multidisciplinary cybersecurity landscape (Abdo and Hossain, 2025). Standardized data formats and sharing protocols are crucial for scalable attribution. Policymakers should mandate greater data exchanges, addressing liability and ownership issues. International cooperation, through bodies such as ICAO, is vital for addressing cross-border threats (Minnaar, 2022; Weber, 2025). Adversaries can be expected to evade attribution, necessitating continuous refinement, adversarial testing, and integration of real-time threat intelligence. Our methodology, which fuses multi-source telemetry and causal graphs, applies to diverse cyber-physical systems, such as autonomous trucks and maritime drones. Ethical and privacy concerns demand strict policies and oversight to prevent misuse of surveillance data. Transparency and independent audits remain essential. Investigations are ethically imperative, supported by our framework's data fidelity and analytical transparency. A limitation is the need for more straightforward explanations for non-technical stakeholders. A user-friendly explanation module could generate plain-language incident reports, translating technical insights into actionable intelligence for legal and policy purposes. Implementing such a system requires significant investment, but it is essential. Prevention is not foolproof; however, attribution aids response, deter attackers, and strengthens defenses. A phased rollout focused on critical operations is advisable. Attribution must integrate seamlessly with incident response, providing intelligence for containment and remediation and supporting law enforcement. While not a complete solution, our research is a crucial step towards demystifying cyber-physical attacks. By clarifying the unfolding of attacks, in their many forms and combinations, we empower defenders and responders, making sophisticated attribution a cornerstone of future safety and security. Future work involves real-world testing of the framework with industry partners using benign scenarios or historical data, alongside enhanced automation to streamline investigations by automatically querying related systems for corroborating data. Within this span of efforts, a focus on protecting infrastructure and workflows more closely tied to the bioeconomy across land and water is of interest due to logistical necessities with some biological handling and processing that have increasingly involved digital means (Potter and Palmer, 2023; Barnett et al., 2024). Furthermore, integrating with legal frameworks through collaboration with policymakers on data sharing and liability for cyber-physical incidents is crucial, as is developing predictive defenses by adapting elements for early warning or prediction through the detection of attack precursors. Finally, understanding human factors, including user interaction, interface design, and training, is essential for effective use during incidents.

7. Conclusion

Autonomous systems like UAVs and smart railways offer immense benefits but introduce new security vulnerabilities spanning cyber and physical domains. This paper addressed the attribution of incidents across these domains, which is crucial for accountability and defense in logistics and transportation. Using graph-based causality, it linked malicious cyber actions (e.g., network intrusion) to physical consequences (e.g., train malfunction). Trust scoring and combined deterministic and machine-learning methods handle uncertainty. Our simulated UAV/rail implementation showed the system could identify coordinated malicious actions that would otherwise appear benign. Key contributions included highlighting data-fusion needs for cross-domain attribution, identifying system gaps (e.g., the lack of rail control/UAV telemetry integration), and providing initial evidence that unified analysis can detect complex attack sequences. We also discussed design trade-offs and made recommendations. In summary, attributing cyber-physical incidents in multimodal environments is challenging but essential. It requires breaking down silos between engineering domains and technical/legal

processes. Our framework and insights lay the groundwork for more resilient autonomous infrastructure, enabling stakeholders to respond effectively and deter future attacks, safeguarding trust and safety in increasingly connected digital and physical systems.

Ethics declaration: Ethical clearance for the research referred to in this paper was not required.

AI declaration: Generative AI in the form of assistive tools within Google Docs was primarily used to reorganize this manuscript. Originally, it was more than twice the length of this current draft. Grammarly was used in the assistance of smoothing writing styles and addressing grammar.

References

Abdo, J.B. and Hossain, L., 2025. Interdisciplinarity of Cybersecurity: Towards Convergence Science. Authorea Preprints.

Almusayli, A., Zia, T. and Qazi, E.-U.-H. (2024) 'Drone Forensics: An Innovative Approach to the Forensic Investigation of Drone Accidents Based on Digital Twin Technology', Technologies, 12(1), p. 11. doi: 10.3390/technologies12010011.

Barnett, M., Samori, I., Griffin, B., Palmer, X.L. and Potter, L., 2023, June. A Commentary and Exploration of Maritime Applications of Biosecurity and Cybersecurity Intersections. In ECCWS 2023 22nd European Conference on Cyber Warfare and Security (No. 1). Academic Conferences and publishing limited.

Casey, E., Barnum, S., Griffith, R., Snyder, J. and van Beek, H. (2015) 'Advancing coordinated cyber-investigations and tool interoperability using a community-developed specification language', Digital Investigation, 12, pp. S102–S110. doi: 10.1016/j.diin.2015.01.013.

Haber, M.J. and Rolls, D., 2024. Indicators of Compromise. In Identity Attack Vectors: Strategically Designing and Implementing Identity Security, Second Edition (pp. 87-107). Berkeley, CA: Apress.

Hartmann, K. and Giles, K. (2016) 'UAV Exploitation: A New Domain for Cyber Power', in Proceedings of the 8th International Conference on Cyber Conflict (CyCon 2016). Tallinn, Estonia: NATO CCD COE Publications, pp. 205–221. Haber, M.J. and Rolls, D., 2019. Indicators of compromise. In *Identity Attack Vectors: Implementing an Effective Identity and Access Management Solution* (pp. 103-105). Berkeley, CA: Apress.

Hwang, S. and Kim, T.S., 2023. An exploratory study on artifacts for cyber attack attribution considering false flag: using Delphi and AHP methods. IEEE Access, 11, pp.74533-74544.

Ibadah, N., Benavente-Peces, C. and Pahl, M.-O. (2024) 'Securing the Future of Railway Systems: A Comprehensive Cybersecurity Strategy for Critical On-Board and Track-Side Infrastructure', Sensors, 24(24), p. 8218. doi: 10.3390/s24248218.

Karnouskos, S., 2011, November. Stuxnet worm impact on industrial cyber-physical system security. In IECON 2011-37th Annual Conference of the IEEE Industrial Electronics Society (pp. 4490-4494). IEEE.

Kim, B., Lee, S., Park, J. and Kim, H. (2021) 'Cybersecurity for unmanned aerial vehicles', IEEE Communications Magazine, 59(8), pp. 34–40. doi: 10.1109/MCOM.001.2000919.

Lian, Z., Shi, P. and Chen, M., 2024. A Survey on Cyber-Attacks for Cyber-Physical Systems: Modeling, Defense and Design. IEEE Internet of Things Journal.

Longo, A., Nicoletti, M., Milazzo, G. and Mendola, S. (2025) 'Cyber–Physical Resilience: Evolution of Concept, Indicators, and Legal Frameworks', Electronics, 14(8), p. 1684. doi: 10.3390/electronics14081684.

Minnaar, A., 2022. Border security: An essential but effective tool in combatting cross-border crime. In The Handbook of Security (pp. 357-378). Cham: Springer International Publishing.

Potter, L. and Palmer, X.L., 2023. Mission-aware differences in cyberbiosecurity and biocybersecurity policies: Prevention, detection, and elimination. In Cyberbiosecurity: A new field to deal with emerging threats (pp. 37-69). Cham: Springer International Publishing.

Pyke, D. (2025) 'Future Rail Security: How DAS and Drones Redefine Infrastructure Protection', Sensonic Blog, 22 July. Available at: https://www.sensonic.com/en/start/future-rail-security-how-das-and-drones-redefine-infrastructure-protection--3286/ (Accessed: 8 September 2025).

Qiu, S., Shao, Z., Wang, J., Xu, S. and Fei, J. (2024) 'Research on Power Cyber-Physical Cross-Domain Attack Paths Based on Graph Knowledge', Applied Sciences, 14(14), p. 6189. doi: 10.3390/app14146189.

Rid, T. and Buchanan, B. (2015) 'Attributing cyber attacks', Journal of Strategic Studies, 38(1–2), pp. 4–37. doi: 10.1080/01402390.2014.977382.

Semertzis, I., Goyel, H. and Stefanov, A. (2024) 'Towards Real-Time Distinction of Power System Faults and Cyber Attacks on Digital Substations Using Cyber-Physical Event Correlation', in Proceedings of the 12th International Workshop on Modeling and Simulation of Cyber-Physical Energy Systems (MSCPES 2024). Los Alamitos, CA: IEEE.

Simons, G., Danyk, Y. and Maliarchuk, T., 2020. Hybrid war and cyber-attacks: creating legal and operational dilemmas. Global Change, Peace & Security, 32(3), pp.337-342.

Soderi, S., Masti, D. and Lun, Y.Z. (2023) 'Railway cyber-security in the era of interconnected systems: A survey', IEEE Transactions on Intelligent Transportation Systems, 24(7), pp. 6764–6779. doi: 10.1109/TITS.2023.3254442.

Strom, B.E., Applebaum, A., Miller, D.P., Nickels, K.C., Pennington, A.G. and Thomas, C.B., 2018. Mitre att&ck: Design and philosophy. In Technical report. The MITRE Corporation.

Weber, L., 2025. ICAO in the 21st century. In The Elgar Companion to the Law and Practice of the International Civil Aviation Organization (pp. 677-692). Edward Elgar Publishing.

Yadav, A., Diedrich, H. and Chowdhury, O. (2023) 'Secure and Trustworthy Sensor Data Fusion for Cyber-Physical Systems', IEEE Internet of Things Journal, 10(15), pp. 12544–12557. doi: 10.1109/JIOT.2023.3244775.

Yi, C.G. and Kim, Y.G., 2024. Hypothesis Generation Model for Cyber Threat Hunting. IEEE Communications Magazine, 62(10), pp.110-116.