# An AI-Driven Threat Intelligence and Vulnerability

*Assessment Framework for Maritime Operational Technology Compliance*

WRITTEN BY
**Qingzhong Liu, Ph.D**

***2026***

INSTITUTE FOR HOMELAND SECURITY
SAM HOUSTON STATE UNIVERSITY

# An AI-Driven Threat Intelligence and Vulnerability Assessment Framework for Maritime Operational Technology Compliance

Qingzhong Liu

Department of Computer Science

Sam Houston State University

Email: liu@shsu.edu

July 15, 2026

# Executive Summary

The project developed an end-to-end AI driven framework which automatically addresses three maritime pain points: fragmented data, vulnerabilities from a wide variety of threats, and regulatory compliance. In doing so, it addresses the new U.S. Coast Guard cybersecurity Final Rule and IMO Resolution MSC.428(98).

The model takes as inputs over 40,000+ records from various separate silos (CVEs, ICS advisories, exploit databases, threat-intel feeds, and regulatory texts) and puts them in one analytical pipeline.

With Maritime Operational Technology (OT) — the navigation, propulsion, and control systems aboard modern vessels — facing escalating cyber risks, the intelligence needed to manage it is scattered across dozens of disconnected public source (silos) and difficult to align with regulatory obligations. This paper consolidates various open-source cyber threat intelligence information into actionable, regulation-aligned guidance.

It includes four analytic components:

1. Classification of unstructured incident narratives
2. Estimation of the relative likelihood of exploitation for known vulnerabilities, Semantic mapping of organizational controls to regulatory requirements, and
3. An integrated knowledge graph linking vulnerabilities, exploits, ATT&CK techniques, and controls.

Initial outcomes from testing are as follows:

1. Screened 2,133 maritime-relevant vulnerabilities and flagged 320 as high or very-high exploitation likelihood, enabling remediation to be prioritized by likelihood of active exploitation rather than by severity score alone.
2. Mapped 19 USCG/IMO requirements against a baseline control set, yielding 81.6% weighted compliance coverage with seven partially satisfied areas and no full gaps, and produced a prioritized remediation roadmap keyed to NIST CSF and IEC 62443 controls.
3. Demonstrated, through a two-scenario sensitivity analysis, that the compliance engine responds to an organization's actual control inventory and coverage falling from 81.6% to 60.5% under a reduced control set rather than returning a fixed baseline score.

These results were obtained entirely from open-source data, and several findings are therefore bounded by data availability: internet-exposure signals were unavailable without commercial feeds, and the limited maritime incident corpus and weak-supervision labeling mean the exploitation model's strong internal discrimination should be validated against real-world ground truth before operational reliance. Even so, the framework demonstrates that fragmented public intelligence can be consolidated into a coherent, continuously updatable, and regulation-aligned view of maritime OT cyber risk. This may provide operators with a practical basis for prioritized, defensible security investment.

**Keywords**: maritime cybersecurity, operational technology, threat intelligence, machine learning, USCG Final Rule, IMO MSC.428(98), vulnerability exploitation, regulatory compliance.

# 1 Contents

# 2 Introduction

The maritime transportation system moves roughly 80% of global trade by volume, and its vessels and port facilities now depend on networked operational-technology systems for navigation, propulsion, cargo handling, and safety. The same connectivity that enables efficiency also enlarges the attack surface. For example, the empirical framework documenting the systemic microeconomic effects of the 2017 NotPetya attack is detailed by Crosignani et al. (2023). Their study maps the propagation of the malware through global supply chain nodes (notably targeting firms like A.P. Moller-Maersk) and calculates the subsequent downstream amplification of financial shocks. The transition of maritime policy to enforce operational technology (OT) and information technology (IT) security under safety frameworks is extensively studied within maritime engineering and governance literature. Hopcraft et al. (2023) examined the integration of International Maritime Organization (IMO) Resolution MSC.428(98) mandate. This mandate required shipping companies and vessel operators to address cyber risks within their safety management systems (SMS) under the International Safety Management (ISM) Code no later than the first annual verification of the Document of Compliance after 1 January 2021. Meanwhile, the legal alignment of localized rules, such as the U.S. Coast Guard's structural enforcement under the Maritime Transportation Security Act (MTSA), is contextually captured through wider legal and systemic frameworks governing high-hazard cyber-physical systems (van Zomeren, 2025).

Despite the critical nature of these assets, actionable threat intelligence specific to maritime Operational Technology (OT) remains highly fragmented. It is scattered across national vulnerability databases, industrial control systems (ICS) advisories, exploit repositories, proprietary threat-intelligence feeds, and shifting regulatory texts — each constrained by siloed formats and inconsistent update cadences (Janosek, 2025). Consequently, security teams must manually reconcile these disparate sources to isolate true operational risks, determine which vulnerabilities have a high likelihood of real-world exploitation, and map existing controls onto overlapping regulatory frameworks (Iannone et al., 2024; Papastergiou et al., 2026). This paper addresses these operational bottlenecks by introducing an automated framework that consolidates 34,336 open-source records into a single analytical pipeline, leveraging machine learning to prioritize risk and systematically assess compliance.

The contributions of this paper are highlighted in the following:

**Aggregation Pipeline:** An open-source data-collection pipeline capable of ingesting spanning nine maritime-relevant intelligence and vulnerability sources.

**Predictive Triage Model**: A transformer-based threat classifier paired with a customized, domain-specific exploitation-likelihood model. While adopting the probabilistic modeling paradigm of the Exploit Prediction Scoring System (EPSS) prioritizing actual exploitation probability over theoretical severity, our model extracts features and assigns labels natively tailored to maritime OT systems, compensating for the standard EPSS model's lack of maritime-specific operational context.

**Semantic Compliance Engine**: A native semantic compliance gap analyzer that dynamically maps organizational security controls onto USCG and IMO regulatory requirements.

**Automated Reporting Architecture**: An integrated knowledge graph and reporting layer that automatically synthesizes these findings into a regulation-aligned white paper.

# 3 Problem Statement

We formalize four problems that the framework must solve:

***Threat classification.*** Given unstructured incident and advisory text, assign each item to a maritime threat category (e.g. ransomware, GPS spoofing, OT exploitation, nation-state) so that trends can be quantified.

***Exploitation-likelihood prediction*.** Given a vulnerability's intrinsic severity, the availability of public exploits, evidence of internet exposure, and historical attack references, estimate the probability that it will be exploited in the wild — a signal that Common Vulnerability Scoring System (CVSS) severity alone does not provide.

***Compliance gap analysis***. Given an organization's security controls and the text of USCG/IMO requirements, determine which requirements are covered, partially covered, or unmet, and produce a prioritised remediation roadmap.

***Synthesis.*** Integrate the above into a queryable knowledge graph and a human-readable, regulation-aligned report.

The central constraint is that all inputs are open source: there is no labelled ground truth for maritime exploitation, and incident text is sparse. The methodology is therefore designed to be robust to weak supervision and limited

labels, and to degrade gracefully when a data source (e.g. paid internet-exposure scanning) is unavailable.

# 4 Methodology

## 4.1 System architecture

The framework is a modular pipeline: (1) collectors normalize each source into a common SQLite schema; (2) an NLP layer performs preprocessing, named-entity extraction, and topic modelling; (3) four model components produce classification, exploitation-likelihood, compliance, and graph outputs; and (4) a reporting layer renders the dashboard data and generates a report. Each stage is idempotent and writes to the shared database, so stages can be re-run independently.

## 4.2 The four AI components

**Threat Classification**: A fine-tuned DistilBERT sequence-classification model maps raw incident narratives to one of ten domain-specific maritime threat categories (Sanh et al., 2019). To guarantee operational robustness under low-resource constraints or cold-start scenarios where labeled training data is scarce, the architecture integrates a deterministic, regular-expression (regex) fallback mechanism to preserve baseline classification coverage.

**Exploitation-Likelihood Prediction Model**: Rather than relying strictly on static severity metrics, a gradient-boosted decision tree (GBDT) classifier estimates the empirical probability of real-world exploitation. It evaluates a 17-dimensional feature vector across four primary telemetry categories: intrinsic severity (CVSS metrics), exploit availability (public exploit count and verification status), internet exposure (observed active internet-facing assets), and historical threat context (threat-report frequency, vulnerability age, reference count, and remote-code-execution/network-vector flags) (Jiang et al., 2025). The Cybersecurity and Infrastructure Security Agency's (CISA) Known Exploited Vulnerabilities (KEV) catalog serves as the ground-truth authority for positive labels; for remaining records, labels are systematically inferred using weak supervision frameworks (Ratner et al., 2018).

**Compliance Gap Analyzer**: To evaluate organizational readiness, regulatory requirements and technical security control texts are converted into dense vector representations and aligned via cosine similarity (Reimers & Gurevych, 2019). Controls are semantically classified into three distinct coverage states, covered, partially covered, or gap, to generate a structured, three-phase remediation roadmap (Bianchi et al., 2026). The semantic engine defaults to

fine-tuned Sentence-Transformer embeddings, utilizing a Term Frequency-Inverse Document Frequency (TF-IDF) cosine similarity pipeline as a backup vectorizer when hardware resources are constrained.

**Knowledge Graph**: Fragmented cybersecurity concepts are contextualized into an expressive threat-intelligence topology. Vulnerabilities, public exploits, MITRE ATT&CK techniques, extracted threat entities, indicators of compromise (IoCs), and governing regulations are mapped as nodes and edges in a heterogeneous knowledge graph (Yadav et al., 2026). Implemented via the NetworkX library (Hagberg et al., 2008), graph maps multi-hop pathways (e.g., tracing which in-the-wild exploits map to maritime OT vulnerabilities) and exports structured schemas for visualization and graph-query operations.

# 5 Data Collection

The data-ingestion architecture relies on nine specialized collectors to gather open-source threat intelligence. Each collector filters data specifically for maritime and Industrial Control Systems (ICS) relevance before mapping the telemetry into a standardized, unified schema. Table 1 outlines the structural composition and volume of the dataset compiled for this study. The resulting knowledge base integrates 2,133 vulnerabilities, 1,829 public exploits, 381 adversarial techniques, and 19 regulatory provisions. Notably, active asset exposure telemetry from the Shodan collector stands at zero, reflecting the framework's deliberate operational boundaries to exclusively ingest open-source, non-proprietary tiers, thereby omitting paid commercial API subscriptions.

| Source | Records |
|---|---|
| **NVD vulnerabilities (CVEs)** | 2133 |
| **ICS-CERT / CISA advisories** | 391 |
| **Exploit-DB public exploits** | 1829 |
| **MITRE ATT&CK for ICS techniques** | 381 |
| **OTX threat indicators** | 31812 |
| **Maritime incident articles** | 33 |
| **USCG/IMO regulatory requirements** | 19 |
| **Internet-exposed assets (Shodan)** | 0 |

*Table 1. The composition and volume of the dataset collected for this study*

Following ingestion, unstructured text fields undergo standard natural language preprocessing, including text normalization, tokenization, and stop-word elimination. The preprocessed corpus is subsequently passed through a dual-

path analytical pipeline consisting of: (i) a named-entity recognition (NER) extractor configured to isolate CVE identifiers, threat actors, malware strains, maritime OT systems, and target organizations, aligning with established domain-specific cybersecurity entity extraction paradigms (Ech-Chammakhy et al., 2025); and (ii) a Latent Dirichlet Allocation (LDA) topic model to surface latent thematic structures within the corpus (Blei et al., 2003). The structural distribution of the identified threat categories across the processed incident corpus is illustrated in Figure 1.

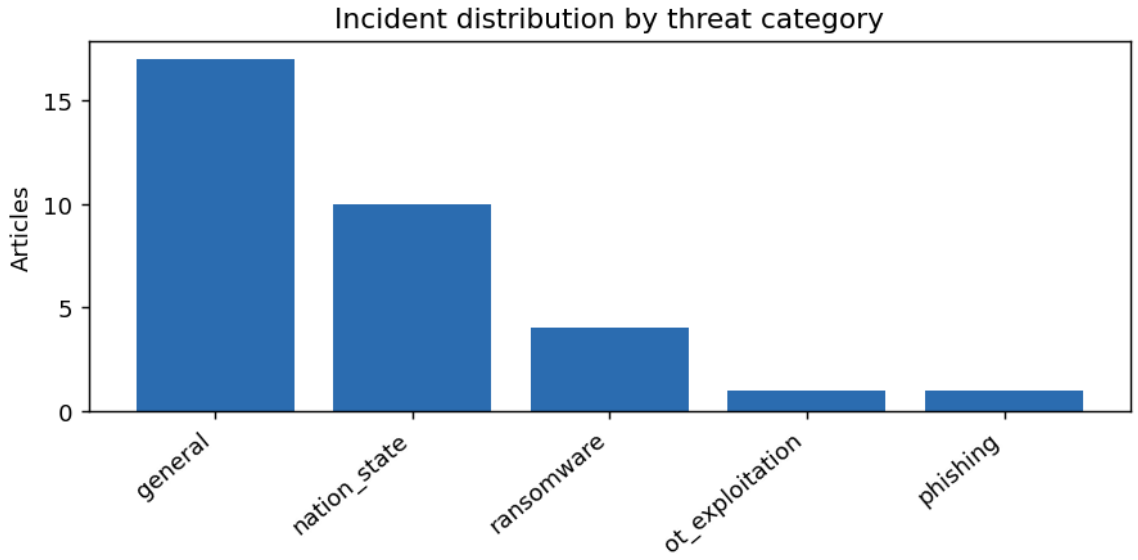


*Figure 1. Incident distribution by threat category*

# 6 Model Selection

Three core operational constraints systematically drove model selection: (i) the low-resource, weakly-labeled nature of the maritime threat landscape, (ii) the requirement to ingest heterogeneous structural and textual data sources, and (iii) the absolute necessity for explainable, highly auditable risk prioritization.

## 6.1 Threat Classification — DistilBERT

To classify unstructured incident narratives under low-resource and hardware-constrained conditions, we utilize DistilBERT (Sanh et al., 2019). Knowledge distillation compresses a large-scale language representation model by training a compact student network on the soft probability distributions produced by a larger teacher network. This allows DistilBERT to retain approximately 97% of BERT's semantic comprehension capabilities while operating with a 40% reduction in parameter size and running 60% faster on commodity hardware.

Cold-Start Fallback: Because labeled maritime cyber incident data is historically scarce, the classification pipeline is engineered with a deterministic keyword classifier that acts as a fallback. This hybrid design mitigates the risk of overfitting a complex deep-learning architecture on sparse sample sizes.

## 6.2 Exploitation Likelihood — Gradient Boosting & SMOTE

The prediction of real-world exploitation probability relies on a Gradient-Boosted Decision Tree (GBDT) architecture, specifically implemented via XGBoost (Chen & Guestrin, 2016). GBDTs are uniquely suited for this task due to their native ability to handle highly heterogeneous, mixed-type features (such as continuous CVSS scores, binary exposure flags, and discrete reference counts) without requiring extensive feature scaling.

Imbalance Mitigation: To counteract the severe class imbalance inherent to threat telemetry—where actively exploited vulnerabilities represent less than 10% of total reported flaws—we apply the Synthetic Minority Over-sampling Technique (SMOTE) (Chawla et al., 2002) to dynamically generate synthetic minority instances along the decision boundaries.

Probability Calibration: Because raw GBDT output margins do not naturally correspond to true mathematical probabilities, we apply sigmoid-based Platt scaling (Niculescu-Mizil & Caruana, 2005). This calibration guarantees that the model's outputs represent true empirical likelihoods, mirroring the mathematical and operational philosophy of the Exploit Prediction Scoring System (EPSS).

## 6.3 Compliance Matching — Sentence Embeddings with TF-IDF Fallback

Mapping complex regulatory prose to specific technical security controls requires capturing deeper semantic relationships rather than simple lexical overlaps, as compliance requirements and corporate policies rarely share identical wording. We leverage Sentence-Transformers (Reimers & Gurevych, 2019) to embed both control documents and regulatory clauses into a shared dense vector space where semantic alignment is calculated via cosine similarity.

Lexical Backup: To guarantee that the compliance engine remains operational across legacy or edge environments lacking GPU acceleration, a sparse vector TF-IDF (Term Frequency-Inverse Document Frequency) cosine similarity pipeline is integrated as a fallback, trading semantic depth for highly efficient, resource-independent lexical matching.

### 6.4 Unsupervised Discovery & Knowledge Representation

Topic Modeling: To extract latent themes from unstructured and unlabeled threat intel feeds, we implement Latent Dirichlet Allocation (LDA) (Blei et al., 2003). LDA operates as an unsupervised, generative probabilistic model that clusters document terms into cohesive thematic topics without requiring human annotations.

Graph Topology: To model the multi-layered dependencies between vulnerabilities, active exploits, and regulatory standards, we build an in-memory property graph using NetworkX (Hagberg et al., 2008). NetworkX provides an optimized, lightweight framework capable of executing path-traversal and multi-hop queries on a hundred-thousand-edge scale, while supporting seamless export to standardized graph serialization formats for external visualization.

## 7 Results and Model Comparison

### 7.1 Threat-classification benchmark

Owing to the low-resource constraints of the compiled corpus, the volume of human-labeled historical maritime cyber incident narratives was insufficient to perform a statistically robust, supervised cross-validation benchmark. Consequently, the ingestion pipeline fell back to the deterministic, regular-expression-based keyword classifier for categorical threat assignment. Hand-labeling and expanding this corpus are identified as a critical near-term priority to enable supervised model training and reporting of formal classifier metrics (e.g., F1, Precision, and Recall) for fine-tuning the DistilBERT architecture.

### 7.2 Exploitation-likelihood model

Under a 5-fold cross-validation scheme executed across the dataset of 2,133 CVEs (with an empirical positive exploitation rate of 15.1%), the gradient-boosted decision tree (GBDT) model achieves an Area Under the Receiver Operating Characteristic curve (ROC-AUC) of 0.9956 and an Area Under the Precision-Recall curve (PR-AUC) of 0.9950, as shown by Figure 2. This high classification accuracy confirms that intrinsic vulnerability severity, public exploit availability, and contextual intelligence signals jointly yield robust separable boundaries for predictive triage.

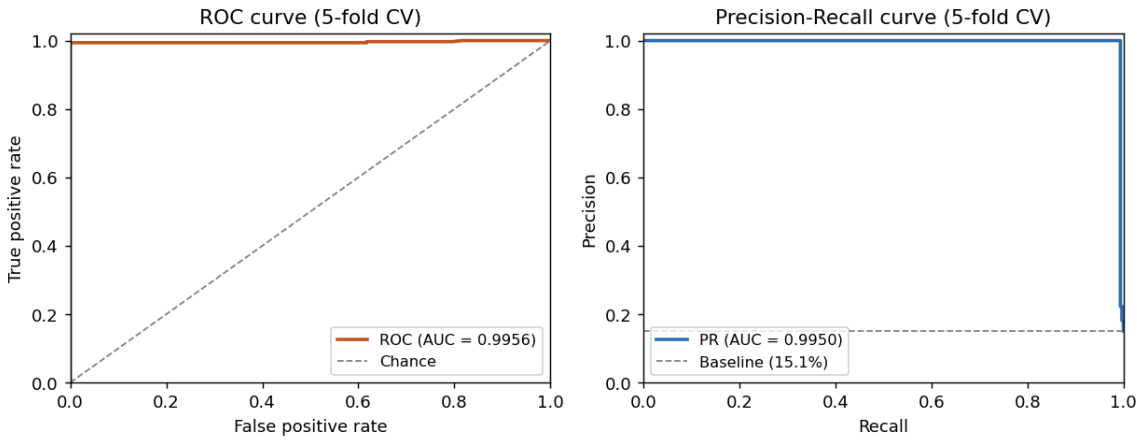


*Figure 2. ROC and precision-recall curves for the exploitation-likelihood model (5-fold cross-validation)*

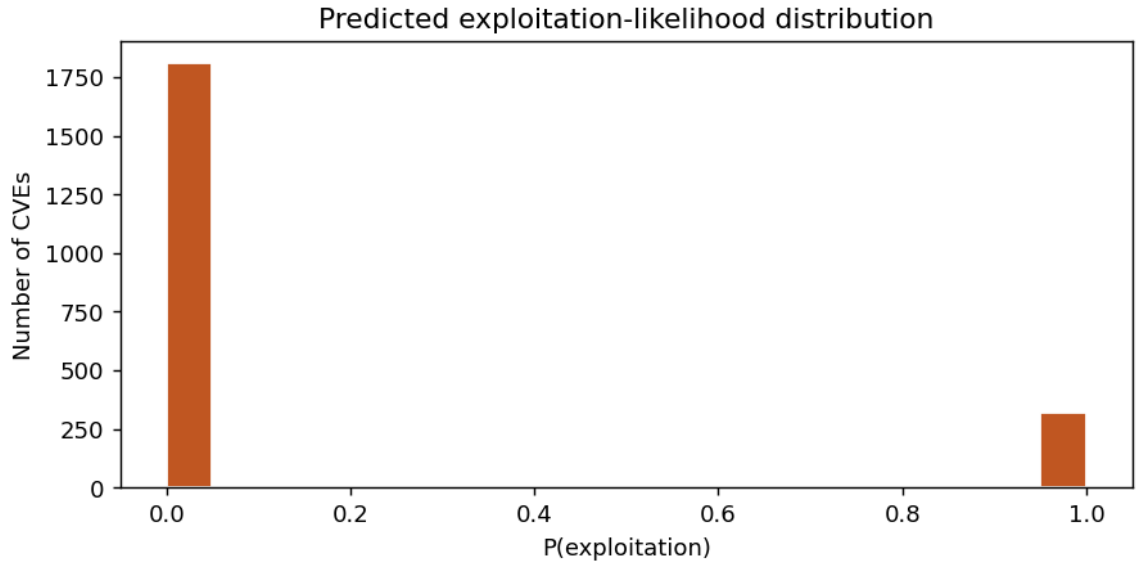


*Figure 3. Predicted exploitation-likelihood distribution*

The predicted probability distribution is shown in Figure 3. It should be noted that the predicted probabilities are strongly bimodal, concentrating near 0 and 1 rather than spreading across the interval, and the highest-ranked vulnerabilities are all assigned a probability of 1.0. Because positive labels are derived from features that are themselves strongly indicative of exploitation (public-exploit availability and CISA KEV membership), the model separates the classes almost perfectly. The probabilities should therefore be read as a relative ranking signal

rather than as calibrated likelihoods; improving calibration with independent ground truth is identified as future work.

## 7.3 OT vulnerability assessment

To evaluate the empirical exposure of shipboard and shoreside physical control domains, open-source vulnerability records were semantically mapped to safety-critical operational technology (OT) asset classes using target keyword profiles.

System Mapping Narrative: As mapped out in the foundational maritime architecture framework established by Li et al. (2024), modern vessels are increasingly characterized by the structural convergence of high-level Information Technology (IT) networks and critical cyber-physical Operational Technology (OT) domains. Figure 4 illustrates this typical shipboard network topology, where satellite communication links (VSAT) and localized Ethernet LANs bridge administrative boundaries down to specialized bridge navigation subsystems—specifically Electronic Chart Display and Information Systems (ECDIS), Automatic Identification Systems (AIS), and Global Navigation Satellite Systems (GNSS)—alongside deep-deck machinery control assets like engine propulsion and ballast systems.

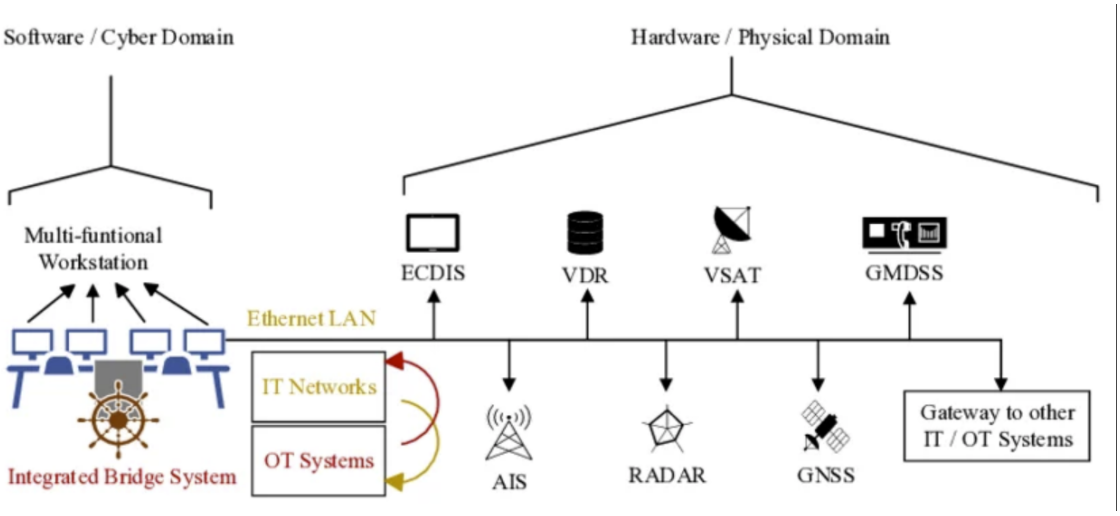


*Figure 4: Topology of safety-critical shipboard IT-OT integration and navigation subsystems. Source: Li et al. (2024)*

Table 2 quantifies the distribution of identified CVE records and their associated average CVSS base severity scores across these core shipboard operational subdomains within our compiled dataset.

| OT System | Related CVEs | Avg CVSS |
|---|---|---|
| **ECDIS** | 1 | 7.5 |
| **GPS/GNSS** | 1 | 6.4 |
| **AIS** | 3 | 5.6 |
| **Engine/Propulsion Control** | 0 | — |
| **SCADA/ICS** | 676 | 7.4 |
| **Ballast/Cargo** | 0 | — |

*Table 2. The distribution of identified CVE records*

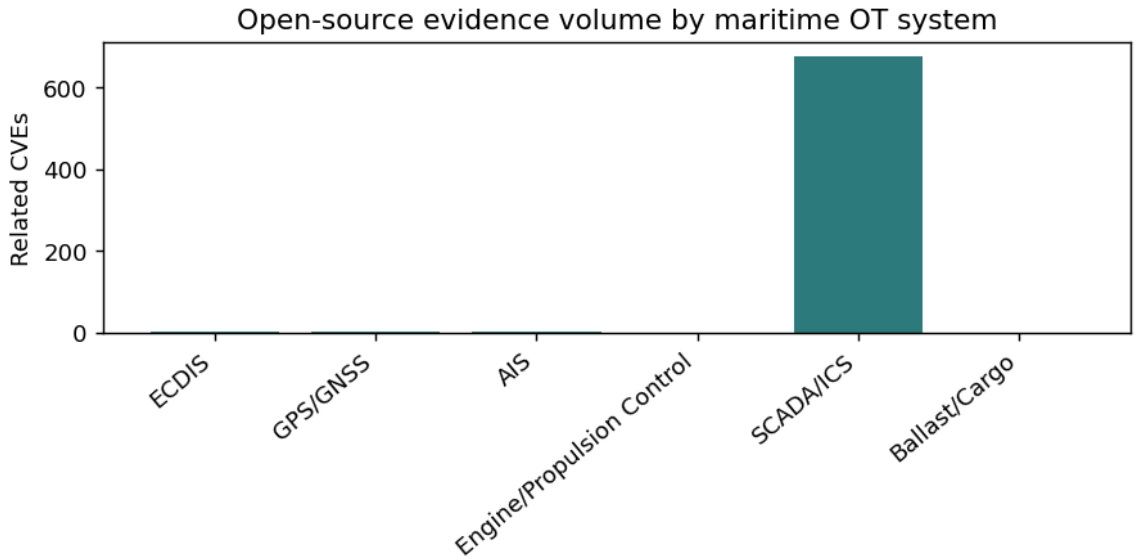


*Figure 5. Evidence volume by maritime OT system*

A critical systemic vulnerability is revealed by Figure 5. It shows that open-source threat intelligence is overwhelmingly skewed toward generic, commercial-off-the-shelf (COTS) SCADA/ICS products. In contrast, highly specialized, safety-critical maritime OT systems, such as ECDIS, GNSS receivers, AIS transponders, engine control networks, and ballast management systems, surface negligible or zero directly attributable CVEs.

This finding does not imply that specialized shipboard OT systems are inherently secure. Rather, it reflects a severe lack of public vulnerability disclosure, limited

security research coverage, and the prevalence of proprietary, vendor-siloed reporting loops in the maritime sector. This gap is a significant finding in itself: maritime operators cannot rely on standard, public CVE feeds to monitor bridge and propulsion safety. Proactive, vendor-aligned threat sharing and private security assessments are mandatory to mitigate these hidden risks.

Table 3 lists the ten vulnerabilities the model ranks highest for exploitation likelihood. All ten receive a predicted probability of exactly 1.0 and fall in the very high band, which reflects the near-binary behavior of the classifier, as shown in Figure 2. Since the weak-supervision labels are derived from signals that are themselves strongly indicative of exploitation, any CVE exhibiting those signals is separated cleanly into the positive class and saturates at P(exploit) ≈ 1.0. Consequently, the identical scores mean the ordering within this top group is not itself meaningful; the ten entries are tied at the ceiling rather than strictly ranked 1-through-10. The table should therefore be read as the set of vulnerabilities the model flags as most exploitation-relevant, not as a calibrated priority sequence. Notably, the entries span a wide publication range (from CVE-2006-5112 to CVE-2025-48466), confirming that the score is driven by exploitation evidence rather than recency. For operational triage, these represent the highest-priority CVEs surfaced by the pipeline, and analysts should combine this ranking with the maritime relevance and CVSS context and with the calibration caveat above when allocating remediation effort.

| CVE | P(exploit) | Level |
|---|---|---|
| **CVE-2023-36458** | 1.0 | very_high |
| **CVE-2018-16590** | 1.0 | very_high |
| **CVE-2018-16591** | 1.0 | very_high |
| **CVE-2018-16705** | 1.0 | very_high |
| **CVE-2006-5112** | 1.0 | very_high |
| **CVE-2009-3646** | 1.0 | very_high |
| **CVE-2009-4529** | 1.0 | very_high |
| **CVE-2025-48466** | 1.0 | very_high |
| **CVE-2010-4709** | 1.0 | very_high |
| **CVE-2010-4730** | 1.0 | very_high |

*Table 3. Highest-ranked vulnerabilities by predicted exploitation likelihood. All entries saturate at P(exploit) = 1.0 (very high); scores are tied at the ceiling and should be read as a set of exploitation-relevant CVEs rather than a strict 1–10 ordering*

# 8 USCG / IMO Compliance Mapping and Roadmap

Using Siamese Sentence-Transformers (SBERT mode) (Reimers & Gurevych, 2019), the framework mapped 19 regulatory USCG/IMO requirements against the 12 organizational security controls of the baseline (Scenario A) inventory, classifying each requirement as covered, partial, or gap: 12 covered, 7 partial, and 0 gaps, for an overall weighted coverage of 81.6%. Each requirement is graded by the cosine similarity between its embedding and that of its best-matching control: covered (similarity ≥ 0.55), partial (0.34 ≤ similarity < 0.55), or gap (similarity < 0.34), with the overall figure weighting covered requirements at 1.0 and partial requirements at 0.5. Table 4 reports the weighted coverage percentage for each control family.

| Category | Coverage % | Gaps |
|---|---|---|
| **Supply_chain** | 50.0 | 0 |
| **Access_control** | 50.0 | 0 |
| **Data_protection** | 50.0 | 0 |
| **Physical_security** | 50.0 | 0 |
| **Network_security** | 75.0 | 0 |
| **Risk_assessment** | 87.5 | 0 |
| **Overview** | 100.0 | 0 |
| **Audit_logging** | 100.0 | 0 |
| **Incident_response** | 100.0 | 0 |
| **Training** | 100.0 | 0 |
| **Tonfiguration_management** | 100.0 | 0 |

*Table 4. USCG/IMO compliance coverage by control category*

## 8.1 Prioritized Remediation Roadmap — Future Actions

The framework identifies the following requirements as partially satisfied and flags them for remediation in the future. For each partial-compliance state, the corresponding remediation action is mapped to the applicable NIST CSF subcategories and ISA/IEC 62443 requirements identified during the gap analysis:

a. IMO MSC.428(98) — Supply Chain (Partial): Formalize third-party and vendor risk management: require software bills of materials (SBOMs) for third-party maritime software, embed cybersecurity criteria into procurement, and enforce multi-factor authentication (MFA) with logged,

controlled remote access on all vendor connections. (NIST CSF ID.SC-1–3; IEC 62443 SR 1.13)

b. USCG cyber requirement area — Data Protection (Partial): Enforce encryption of data at rest (e.g., AES-256) across shipboard servers and backup media, and protect the integrity and confidentiality of OT communications. (NIST CSF PR.DS-1, PR.DS-2, PR.DS-6; IEC 62443 SR 3.1, SR 4.1)

c. USCG cyber requirement area — Supply Chain (Partial): Establish continuous vulnerability monitoring for all third-party maritime software and hardware integrations, with defined tracking, ownership, and remediation service levels. (NIST CSF ID.SC-2, ID.SC-4; IEC 62443 SR 1.13)

d. USCG cyber requirement area — Physical Security (Partial): Integrate physical and cyber access controls for restricted spaces—badge or biometric access and CCTV monitoring for server rooms and bridge terminals housing OT/IT assets. (NIST CSF PR.AC-2, DE.CM-2; IEC 62443 SR 2.1, SR 2.4)

e. IMO MSC.428(98) — Risk Assessment (Partial): Establish a continuous, automated threat-intelligence collection and reporting capability for shipboard OT vulnerabilities, feeding periodic cybersecurity risk assessments. (NIST CSF ID.RA-1, ID.RA-2, ID.RA-3, ID.RA-5; IEC 62443 SR 3.1, SR 3.2)

f. IMO MSC.428(98) — Network Security (Partial): Enforce network segmentation with a demilitarized zone (DMZ) separating the administrative IT network from vessel navigation and propulsion OT networks, following the ISA/IEC 62443 zones-and-conduits model. (NIST CSF PR.AC-5, PR.DS-5, DE.CM-1; IEC 62443 SR 5.1–5.3)

g. USCG cyber requirement area — Access Control (Partial): Decommission shared and generic privileged accounts across OT systems and replace them with unique user identities enforcing least privilege and role-based access. (NIST CSF PR.AC-1, PR.AC-3, PR.AC-4; IEC 62443 SR 1.1–1.3)

## 8.2 Compliance Sensitivity Analysis (Scenario Evaluation)

To verify that the semantic gap analyzer is sensitive to varying levels of organizational cybersecurity maturity rather than producing static baseline scores, we evaluated the same 19 regulatory USCG/IMO requirements against two distinct operational control scenarios:

**Scenario A (High Maturity**): A modern, mature vessel operator with a complete set of 12 structured security controls.

**Scenario B (Low Maturity)**: A lower-maturity vessel operator with only 9 active controls, lacking network segmentation, centralized logging, and systematic vulnerability patching.

| Metric | Scenario A (full) | Scenario B (reduced) |
|---|---|---|
| **Controls evaluated** | 12 | 9 |
| **Covered** | 12 | 8 |
| **Partial** | 7 | 7 |
| **Gaps** | 0 | 4 |
| **Weighted coverage %** | 81.6 | 60.5 |

*Table 5. Sensitivity analysis: scenario comparison.*

As shown in Table 5, removing three controls lowers weighted coverage from 81.6% to 60.5% and opens 4 mandatory gaps that Scenario A did not have. This confirms that the semantic gap analyzer is sensitive to the actual control inventory rather than returning a fixed score, and that its output can be used to benchmark an organisation's posture and target the highest-impact missing controls. The main analysis above reflects Scenario A; the Scenario B gaps illustrate the remediation of the tool surfaces for a weaker posture.

# 9 Prioritized Risk-Mitigation Recommendations

a. **Transition to Predictive, Likelihood-Based Patching:** Shift vulnerability remediation priorities from static severity (CVSS ≥ 9.0) to exploitation likelihood. Remediating the 320 vulnerabilities in the High and Very-High bands (P(exploit) ≥ 0.55; in practice ≈ 1.0, as the scores are strongly bimodal) mitigates the vast majority of actively exploitable threat vectors and avoids engineering effort spent on high-severity vulnerabilities that have no public exploit code.

b. **Prioritize Safety-Critical OT Subsystems:** Focus immediate security controls on specialized bridge and propulsion systems — specifically ECDIS, GPS/GNSS, and general SCADA networks. Because these systems lack transparent, public vulnerability disclosure, operators must establish private monitoring, offline network segmentation, and proactive vendor security verification.

c. **Operationalize Continuous Intelligence Pipelines:** Transition from static, annual compliance audits to a continuous compliance posture. Run the open-source data-collection pipeline on a recurring weekly schedule to

update the knowledge graph, refresh exploitation-likelihood scores, and regenerate compliance roadmaps dynamically.

# 10 Contribution and Limitations

### Contribution

Our framework automatically consolidates maritime cyber-threat data that is scattered across dozens of incompatible public sources into **actionable, regulation-aligned intelligence,** and it enables maritime operators to prioritize by *real-world risk* rather than theoretical severity and see exactly where they fall short of USCG/IMO compliance.

### Who it helps

1. **Cybersecurity Officers at vessel and port operators.** The USCG Final Rule now *mandates* cyber risk assessments and compliance plans. This tool produces the risk-ranked vulnerability list, the compliance gap analysis, and the remediation roadmap they need.
2. **Small and mid-sized operators without a dedicated threat-intelligence team or the budget for commercial feeds.** Because the framework runs entirely on open-source data, it gives them an enterprise-grade threat picture they otherwise couldn't afford.
3. **Regulators, auditors, and insurers.** They can use the compliance mapping and risk scores to assess an operator's cyber maturity quickly.
4. **Researchers**. We provide a reproducible, extensible baseline for maritime OT threat intelligence.

### Why it's needed

We address the following three real pain points:

1. **Fragmented data.** CVEs, ICS advisories, exploit databases, threat-intel feeds, and regulatory texts all live in separate silos with different formats; reconciling them by hand is nearly impossible. This framework unifies **40,000+ records** into one analytical pipeline.
2. **Severity ≠ exploitation.** Patching purely by CVSS wastes effort on vulnerabilities that are severe but never exploited. The EPSS-style model predicts the *actual probability of exploitation*, so limited resources go where they matter.
3. **Compliance is now a hard requirement but hard to operationalize.** The USCG Final Rule (in force July 2025) is mandatory, yet no one can manually map "which of my controls satisfy which USCG/IMO requirements." The semantic compliance analysis does this **automatically** and outputs a phased roadmap.

**Maritime OT is distinctive**: generic tools don't understand ECDIS, AIS, or SCADA systems. This framework is **maritime-specific**, which is its core differentiator from general-purpose security tools.

**The limitations** of our system are listed below:

a. **Internet-Exposure Limitations:** Because this study operates strictly within open-source, non-proprietary data boundaries, paid commercial internet-scanning tiers (such as premium Shodan API access) were omitted. Consequently, the internet-exposure feature group is zero-valued, forcing the gradient-boosted model to rely on static severity and exploit-availability signals. Integrating licensed, real-time device-exposure data is expected to sharpen the model's predictive precision.

b. **Label Sparsity and Weak Supervision:** The absence of a centralized, public repository of maritime-specific cyber exploits necessitated weak supervision. The exceptionally high AUC (ROC-AUC = 0.9956) reflects the mathematical separability of these heuristic labels and must be interpreted as internal model consistency rather than validated field accuracy. In addition, the small maritime incident corpus prevented supervised-classifier benchmarking and restricted fine-tuning of the DistilBERT model.

c. **Query Completeness and Generalizability:** The pipeline's output is bound to the search syntax and keywords configured in the collectors. Broadening the semantic scope and integrating additional international vulnerability repositories would improve coverage, particularly for under-represented OT subdomains such as marine propulsion and ballast-control systems.

# 11 Conclusion and Future Work

This paper presented an end-to-end, AI-driven assessment framework that consolidates fragmented, open-source cyber threat intelligence into actionable, regulation-aligned guidance for maritime Operational Technology (OT). The framework ingests multi-source data, classifies unstructured incident narratives through a keyword-and-transformer classification pipeline (transformer fine-tuning pending a larger labeled corpus) with regex-based entity extraction, estimates relative exploitation likelihood, and maps organizational controls to USCG/IMO requirements via semantic vector embeddings — visualizing the resulting relationships through an integrated knowledge graph.

Future research will focus on:

1. **Corpus Expansion:** Manually labeling a larger dataset of maritime cyber incidents to support full fine-tuning and benchmarking of supervised DistilBERT classifiers.
2. **Model Calibration:** Acquiring independent, verified exploitation ground truth so that probability calibration (isotonic regression is already applied during training) can be evaluated against real outcomes using reliability diagrams and Brier score, rather than against self-consistent weak labels.
3. **Data Enrichment:** Integrating paid internet-exposure feeds and private threat intelligence to eliminate data blind spots.
4. **Longitudinal Validation:** Conducting a multi-year study to validate predicted exploitation likelihoods against verified, real-world maritime cyber incidents.

# 12 Acknowledgements

We highly appreciate our Project Manager, Mr. Scott Lynn, for his meticulous review and insightful comments, and we are also very grateful for the support from the SHSU Institute of Homeland Security.

# 13 References

Bianchi, J., Petrillo, L., Martinelli, F., & Petrocchi, M. (2026). Automated compliance mapping in cloud security with domain-adapted sentence transformers. *arXiv preprint arXiv:2607.06364*. https://doi.org/10.48550/arXiv.2607.06364

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3(4-5), 993–1022.

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357. https://doi.org/10.1613/jair.953

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). https://doi.org/10.1145/2939672.2939785

Crosignani, M., Macchiavelli, M., & Silva, A. F. (2023). Pirates without borders: The propagation of cyberattacks through firms' supply chains. *Journal of Financial Economics*, 147(3), 432–448. https://doi.org/10.1016/j.jfineco.2022.12.002

Ech-Chammakhy, Y., Motii, A., Rabii, A., Azrara, O., & Chbili, J. (2025). CyberNER: A harmonized STIX corpus for cybersecurity named entity recognition. arXiv preprint arXiv:2510.26499. https://doi.org/10.48550/arXiv.2510.26499

Hagberg, A. A., Swart, P. J., & Schult, D. A. (2008). Exploring network structure, dynamics, and function using NetworkX. In *Proceedings of the 7th Python in Science Conference (SciPy 2008)* (pp. 11–15). Pasadena, CA.

Hopcraft, R., Harish, A. V., Tam, K., & Jones, K. (2023). Raising the standard of maritime voyage data recorder security. *Journal of Marine Science and Engineering*, 11(2), 267. https://doi.org/10.3390/jmse11020267

Iannone, E., Sellitto, G., Iaccarino, E., Ferrucci, F., De Lucia, A., & Palomba, F. (2024). Early and realistic exploitability prediction of just-disclosed software vulnerabilities: How reliable can it be? *ACM Transactions on Software Engineering and Methodology*, 33(4), 1–41. https://doi.org/10.1145/3654443

Janosek, D. M. (2025). Toward a global framework for cyber threat intelligence sharing. *The Cyber Defense Review,* vol. 10, no. 2, 99-114, https://doi.org/10.55682/cdr/sgyw-5r5p

Jiang, Y., Oo, N., Meng, Q., Lim, H. W., & Sikdar, B. (2025). A survey on vulnerability prioritization: Taxonomy, metrics, and research challenges. *arXiv preprint arXiv:2502.11070*. https://doi.org/10.48550/arXiv.2502.11070

Li, M., Zhou, J., Chattopadhyay, S., & Goh, M. (2024). Maritime cybersecurity: A comprehensive review. *arXiv preprint arXiv:2409.11417*. https://doi.org/10.48550/arXiv.2409.11417

Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. In *Proceedings of the 22nd International Conference on Machine Learning* (pp. 625–632). https://doi.org/10.1145/1102351.1102430

Papastergiou, S., Basheer, N., Lampropoulos, K., Verrios, P., & Islam, S. (2026). Explainable AI based dynamic cybersecurity risk management for cyber insurability. International Journal of Information Security, 25, Article 36. https://doi.org/10.1007/s10207-025-01189-8

Ratner, A., Bach, S. H., Ehrenberg, H., Fries, J., Wu, S., & Ré, C. (2018). Snorkel: Rapid training data creation with weak supervision. Proceedings of

the VLDB Endowment, 11(3), 269–282. https://doi.org/10.14778/3157794.3157799

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP) (pp. 3982–3992). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1410

Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*. https://doi.org/10.48550/arXiv.1910.01108

van Zomeren, M. (2025). Regulating cyberworthiness: Governance frameworks for safety-critical cyber-physical systems. *Systems*, 13(10), 862. https://doi.org/10.3390/systems13100862

Yadav, S., Arikkat, D. R., Agarwal, B., Nicolazzo, S., Nocera, A., & Vinod, P. (2026). CVE-TTP KG: Knowledge graph linking software vulnerabilities to attack behaviors. *arXiv preprint arXiv:2606.31557*. https://doi.org/10.48550/arXiv.2606.31557

# 14 Appendix A. Interactive Dashboard — Subsystem Views

The framework's secondary deliverable is an interactive, web-based dashboard (a React single-page application backed by the same data store) organised into five subsystems, selectable via a navigation bar: Overview, Threat Intelligence, Compliance, Vulnerabilities, and Knowledge Graph. Because the dashboard is interactive, it cannot be reproduced live in a static document; this appendix presents each subsystem's key content as static views generated from the same underlying data. Where a screenshot of the live interface has been provided, it is shown alongside.

## 14.1 A.1 Overview

The Overview subsystem summarizes the intelligence picture through headline indicators and trend/severity panels. Current headline figures: 391 advisories, 2,133 CVEs, 33 incident articles, 31,812 threat indicators, and 381 ATT&CK for ICS techniques.

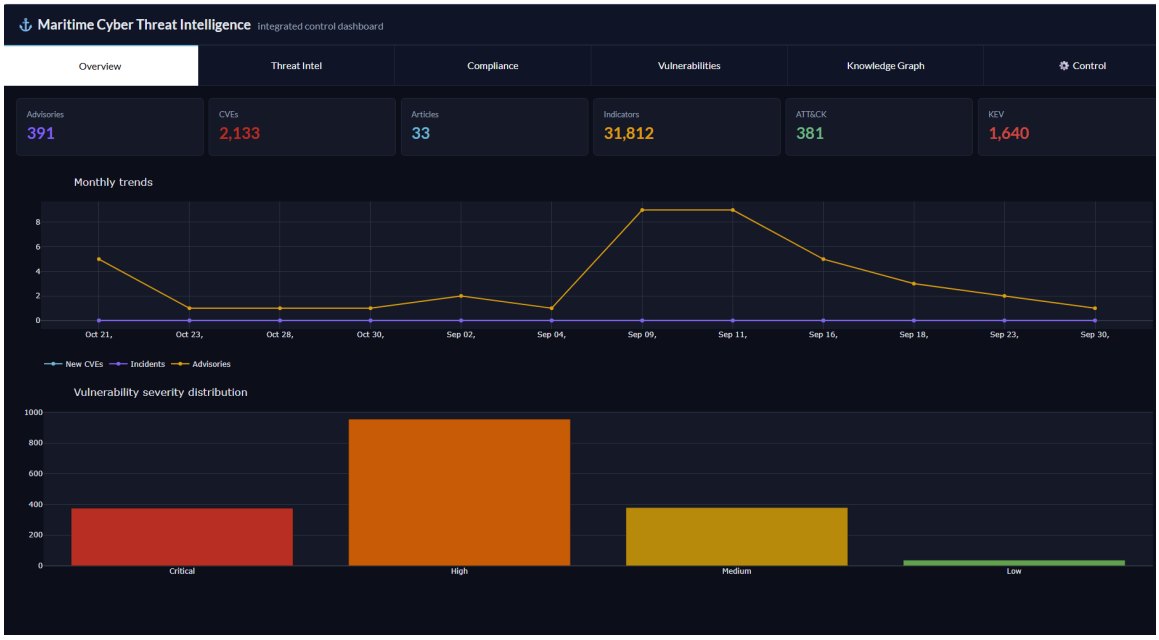


Figure A.1 Dashboard — Overview tab (live interface)

## 14.2 A.2 Threat Intelligence

The Threat Intelligence subsystem visualises the distribution of incidents across threat categories and highlights the leading MITRE ATT&CK for ICS techniques observed in the corpus.

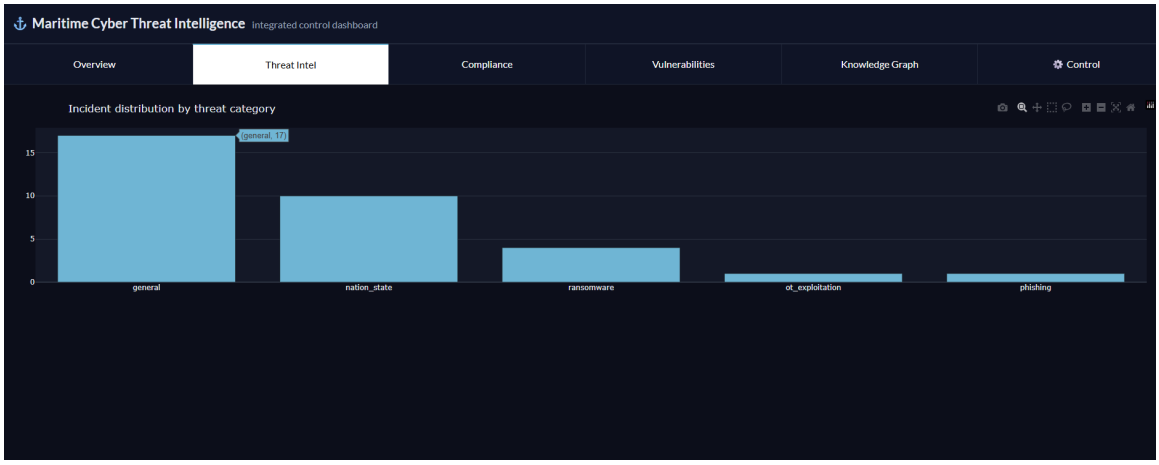


Figure A.2 Dashboard — Threat Intelligence tab (live interface)

## 14.3 A.3 Compliance

The Compliance subsystem reports coverage against USCG/IMO requirements and surfaces the specific gaps and remediation roadmap detailed in Section 7 (81.6% weighted coverage, 0 gaps).


Figure A.3 Dashboard — Compliance tab (live interface)

## 14.4 A.4 Vulnerabilities

The Vulnerabilities subsystem lists critical and high-severity maritime vulnerabilities and ranks them by predicted exploitation likelihood (Section 6.2), enabling triage by probability of real-world exploitation rather than severity alone.

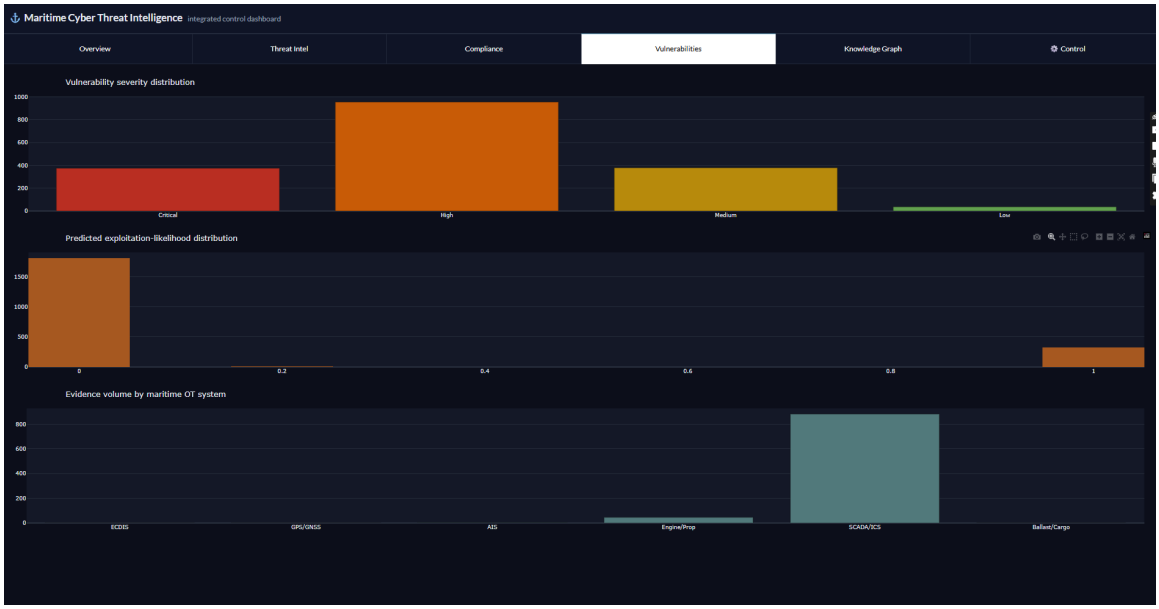


Figure A.4 Dashboard — Vulnerabilities tab (live interface)

## 14.5 A.5 Knowledge Graph

The Knowledge Graph subsystem renders the entity-relationship network linking vulnerabilities, exploits, ATT&CK techniques, extracted entities, indicators, and regulations, allowing analysts to trace, for example, which exploited vulnerabilities affect which OT systems. Because the graph is a large interactive network, it is best viewed in the live dashboard or a graph-visualization tool rather than as a static image.

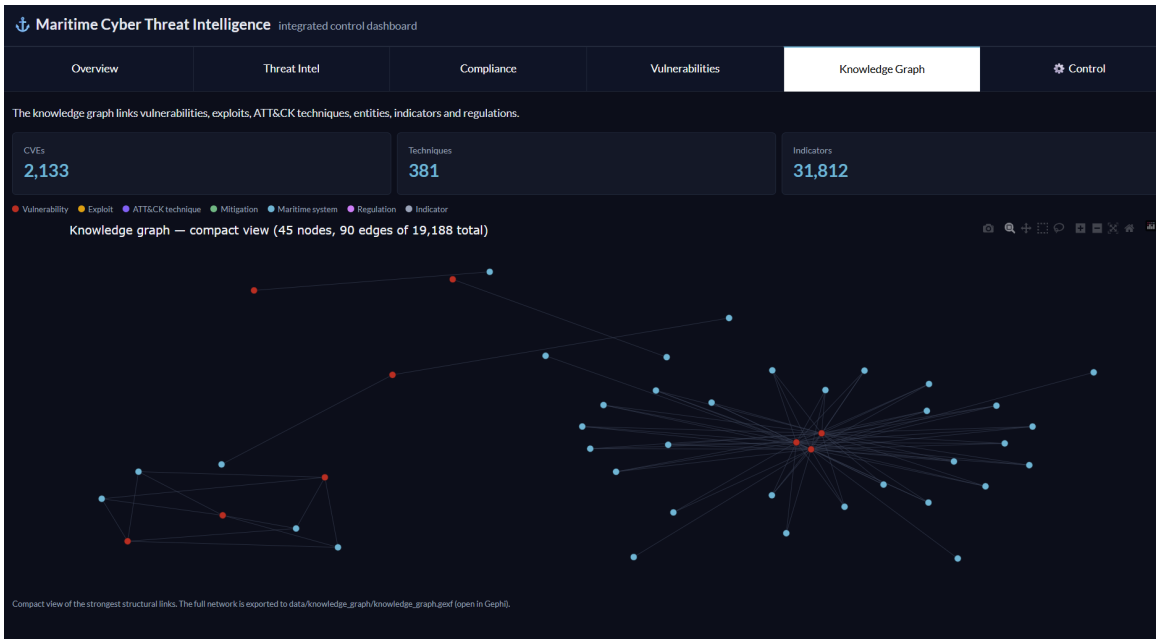


Figure A.5 Dashboard — Knowledge Graph tab (live interface)

# 15 Author Biography

Dr. Qingzhong Liu is a tenured Professor in the Department of Computer Science at Sam Houston State University (SHSU). He currently serves as the leader of the university's Cybersecurity with AI Applications Group. He earned his Ph.D. in Computer Science from the New Mexico Institute of Mining and Technology in 2007.

His primary research interests encompass multimedia security, digital forensics, cybersecurity, artificial intelligence, cybercrime investigation, quantitative finance, biomedical informatics, and computational analysis. Dr. Liu actively investigates complex challenges in explainable digital forensics, deepfake detection, cybersecurity, and computational AI systems. His scholarly work and laboratory initiatives have been consistently supported by multiple grants from prominent federal agencies, including the National Science Foundation (NSF), the National Institute of Justice (NIJ), the National Institutes of Health (NIH) and the Institute for Homeland Security (IHS).

Dr. Liu has published extensively across highly respected journals and conferences, contributing numerous peer-reviewed papers and book chapters to the scientific community. His significant impact on the field has been recognized since 2020 among the World's Top 2% Scientists in Artificial Intelligence and Image Processing, ranked by Stanford University. Additionally, he holds multiple patents for his innovative computational approaches to steganalysis and image forgery detection

The Institute for Homeland Security at Sam Houston State University is focused on building strategic partnerships between public and private organizations through education and applied research ventures in the critical infrastructure sectors of Transportation, Energy, Chemical, Water/Wastewater, Healthcare, and Public Health.

The Institute is a center for strategic thought with the goal of contributing to the security, resilience, and business continuity of these sectors from a Texas Homeland Security perspective. This is accomplished by facilitating collaboration activities, offering education programs, and conducting research to enhance the skills of practitioners specific to natural and human caused Homeland Security events.

Institute for Homeland Security

Sam Houston State University



Liu, Q. (2026). An AI-Driven Threat Intelligence and Vulnerability Assessment Framework for Maritime Operational Technology Compliance. (Report No. 2026-1049). The Sam Houston State University Institute for Homeland Security.

https://doi.org/10.17605/OSF.IO/TJ2EV

INSTITUTE FOR HOMELAND SECURITY
SAM HOUSTON STATE UNIVERSITY