Sciences Methods and Technologies International Journal (SciMeTech)

ISSN: 3085-5284, journal homepage: www.scimetech.com

Explainable AI for Security Operations : Improving Cyber Threat Intelligence Reporting and Visibility Across African Organizations

Blessing N. Ezeobioha
Miva Open University
gabriellaakubueze115@gmail.com

How to cite:
Blessing N. Ezeobioha, "Explainable AI for Security Operations : Improving Cyber Threat Intelligence Reporting and Visibility Across African Organizations", Sciences Methods and Technologies International Journal (SciMeTech), Vol 2, Issue 1, p 207-215

Keywords:
Cyber threat intelligence
Security Operations
SOC automation
Explainable AI
Incident reporting
Africa
Privacy-aware sharing

Abstract

Cyber threat intelligence (CTI) becomes useful only when incidents are reported early, normalized into structured knowledge, and shared in a form that defenders can act upon. Many global CTI systems assume that incident evidence is already available as indicators, reports, or machine-readable objects. In African security operations, however, analysts often face a more basic upstream problem: incidents may be observed, reported informally, or visible on breach monitoring platforms and social media, yet they do not consistently become structured public or sector-level intelligence. This makes monthly TI reporting difficult and weakens regional visibility in global CTI ecosystems. The camera-ready version strengthens the accepted paper by adding practitioner-anchored observations from African security operations, including alleged dark-web exposure claims, cloned public-service portals, financial-sector phishing, payment-platform impersonation, and electricity-utility themed fraud lures. These cases are treated as sanitized reporting-gap observations rather than independently verified public breach claims. The paper proposes a hybrid explainable AI framework that combines rule-guided extraction, lightweight trainable classification, evidence-span tracing, analyst verification, and Privacy-aware STIX-style reporting. To test feasibility, ten analyst-realistic incident narratives were used, including practitioner-informed Nigerian and African SOC scenarios. In a single-analyst feasibility evaluation, the prototype achieved a macro field-extraction accuracy of 0.80 across sector, threat type, indicator, tactic/technique, and privacy sensitivity labels. The failures occurred mainly in higher-level TTP and privacy classification, which supports the need for human verification rather than black-box automation. The paper contributes a regionally grounded CTI reporting problem formulation, an improved hybrid XAI architecture, and a practical validation pathway for African SOCs, CSIRTs, and sectoral information-sharing communities.

I. Introduction

Cyber threat intelligence (CTI) is central to modern security operations because it helps defenders move from isolated incident handling to reusable knowledge about adversary behavior, infrastructure, targets, and techniques. Standards such as STIX/TAXII and platforms such as MISP support machine-readable sharing, while operational frameworks such as the intrusion kill chain and MITRE ATT&CK help analysts describe adversary activity in a common language (Hutchins et al., 2011; MITRE, 2024; OASIS, 2021; Wagner et al., 2016). However, these systems work best when the earlier reporting layer is mature. If incidents are not documented, normalized, de-identified, and routed to trusted communities, CTI sharing becomes an ideal rather than an operational reality.

For African organizations, the problem is not only access to global intelligence feeds. The more difficult issue is often the local production of CTI. A threat intelligence analyst preparing a monthly report may find detailed reporting for North America or Europe, while African incidents are visible only as scattered news, informal WhatsApp or Telegram alerts, X posts, breach-market claims, cloned-domain observations, takedown notes, and private SOC records. In the author's practitioner experience supporting cybersecurity work across multiple African countries, incidents exist across banks, public-sector services, payment channels, utilities, and identity systems, but detailed public CTI reports from affected organizations are often unavailable. This creates a visibility gap: threats exist, but they do not reliably become searchable, structured, and historically reusable intelligence. Recent work on Nigeria and broader African cyber governance also points

to weaknesses in incident reporting, public-private cooperation, and cross-sector sharing (Nainna & Bass, 2024; Sani & Yakub, 2025; Tounsi & Rais, 2018).

This paper addresses that upstream gap. Instead of treating CTI as a downstream sharing problem only, it frames CTI as a reporting and knowledge-generation process. The central argument is that AI-enabled security operations in Africa should not only detect threats; they should help analysts convert imperfect local observations into explainable, privacyaware, and reusable intelligence products. The paper therefore contributes: (1) a refined African CTI reporting-gap model, (2) a hybrid explainable AI architecture for incident-to-CTI conversion, (3) a small practitioner-anchored feasibility evaluation using ten incident narratives, and (4) a concrete camera-ready improvement path based on reviewer recommendations: stronger validation, clearer baselines, and higher technical depth.

II. Methods a. Background and Research Gap

Prior CTI research has produced mature discussions of standards, sharing incentives, technical interoperability, and collective defense. NIST SP 800-150 emphasizes that cyber threat information must be collected, shared, and used within clearly governed processes (Johnson et al., 2016). Skopik et al. (2016) show that collective defense depends not only on technical exchange, but also on trust, incentives, organizational readiness, and governance. Tounsi and Rais (2018) similarly identify quality, timeliness, and usability as persistent CTI challenges in the age of sophisticated attacks.

Technical standards have improved the representation and transport of intelligence. STIX provides structured objects for indicators, malware, vulnerabilities, attack patterns, and relationships, while TAXII supports automated exchange (OASIS, 2021). MISP offers a widely used open-source platform for collaborative threat intelligence and event sharing (Wagner et al., 2016). These systems are important, but they usually assume that organizations already possess structured intelligence or are willing to share it. The African reporting problem often begins earlier: an analyst may have an incident narrative, a suspicious domain, a dark-web listing, or a cloned public-service portal, but not a ready CTI object.

Explainable AI is relevant because SOC analysts must understand and verify system output before acting on it. General XAI work such as LIME and SHAP demonstrates ways to explain predictions by highlighting influential features or local decision logic (Lundberg & Lee, 2017; Ribeiro et al., 2016). In cybersecurity, explainability is especially important because false positives, opaque automation, and unverifiable classifications can harm trust and response quality (Sommer & Paxson, 2010). The gap addressed here is therefore not simply the absence of AI; it is the absence of analystverifiable AI that supports reporting, evidence tracing, privacy labeling, and CTI normalization in African operational contexts.

Recent Africa-focused scholarship also supports the need for localized CTI work. Nainna and Bass (2024) discuss cyber threat intelligence sharing in Nigeria and point to practical and organizational challenges. INTERPOL's African cyberthreat assessments highlight the region's growing cybercrime exposure and the need for stronger coordination and reporting capacity (INTERPOL, 2021; INTERPOL, 2024). Yet there remains limited work that converts these observations into a concrete AI-enabled reporting pipeline for SOCs and CSIRTs.

i. African CTI Visibility and Reporting Gap

In practice, the African CTI gap appears as a visibility problem. Incidents may happen, be noticed by analysts, or surface on third-party platforms, but still fail to become authoritative, searchable, and reusable intelligence. This is different from saying that Africa has no cyber incidents. Rather, the problem is that local incidents are often under-structured, underpublished, or delayed in public reporting. The result is weak historical memory for monthly TI reporting and reduced representation in global intelligence products.

Two categories of practitioner observation motivated this revision. First, alleged dark-web or breach-market claims involving Nigerian public-sector, financial-sector, and service-provider data may circulate before any structured CTI advisory becomes visible to analysts. Second, impersonation and fraud infrastructure can target high-trust Nigerian brands and public services - including immigration-service lookalikes, central-bank themed lures, payment-platform impersonation, bank-themed phishing, and electricity-utility bill/payment scams - while the public record may contain little more than the adversary's own post, a takedown notice, or informal analyst warnings. These observations are not presented as verified breach accusations against any named organization. They are sanitized practitioner cases used to illustrate reporting latency, weak public CTI production, and the difficulty of building African monthly TI reports from authoritative local sources.

Open threat-intelligence and breach-aggregation platforms can also show uneven country visibility. For example, exploratory use of platforms such as Breach.house may return richer country-level visibility for some regions than for Nigeria or other African contexts. This paper treats such observations carefully: absence from a platform is not evidence of absence of incidents. It is evidence that platform visibility, indexing, reporting pathways, language, and disclosure

practices shape what analysts can find. This motivates the need for CTI pipelines that help local analysts transform observations into structured, privacy-aware reports that can be shared with appropriate communities. This is also why personalization strengthens the paper rather than weakens it. The author is not claiming continentwide measurement from private evidence; instead, the author's operational experience is used to ground the research problem, while the technical contribution remains generalizable: any SOC or CSIRT that receives fragmented incident narratives can use the proposed framework to normalize, de-identify, verify, and share CTI more consistently.

b. Proposed Hybrid Explainable AI Framework

The improved framework uses four layers. Layer 1 collects incident narratives from SOC notes, CSIRT emails, takedown reports, phishing submissions, public advisories, and carefully handled platform observations. Layer 2 applies hybrid extraction: deterministic rules identify obvious indicators such as domains, URLs, hashes, IP addresses, and file names, while a lightweight trainable classifier predicts higher-level labels such as sector, threat type, ATT&CK-style behavior, privacy sensitivity, and suggested action. Layer 3 provides explainability by attaching extracted fields to evidence spans from the original narrative. Layer 4 performs privacy-aware CTI packaging using de-identification, Traffic Light

Protocol-style handling, and STIX-style fields. The key technical change from the original version is that the system is no longer described as only rule-based. Rulebased extraction remains useful for transparent IoCs, but the revised design adds trainable classification for ambiguous fields and explanation support through feature contribution and evidencespan tracing. In a future deployment, a small transformer or gradient-boosted text classifier can be trained on anonymized incident narratives, with SHAP- or LIMEstyle explanations used to show which words or phrases influenced a classification. The analyst remains responsible for final approval, because the African reporting context includes legal, reputational, and privacy constraints that automation alone cannot resolve.

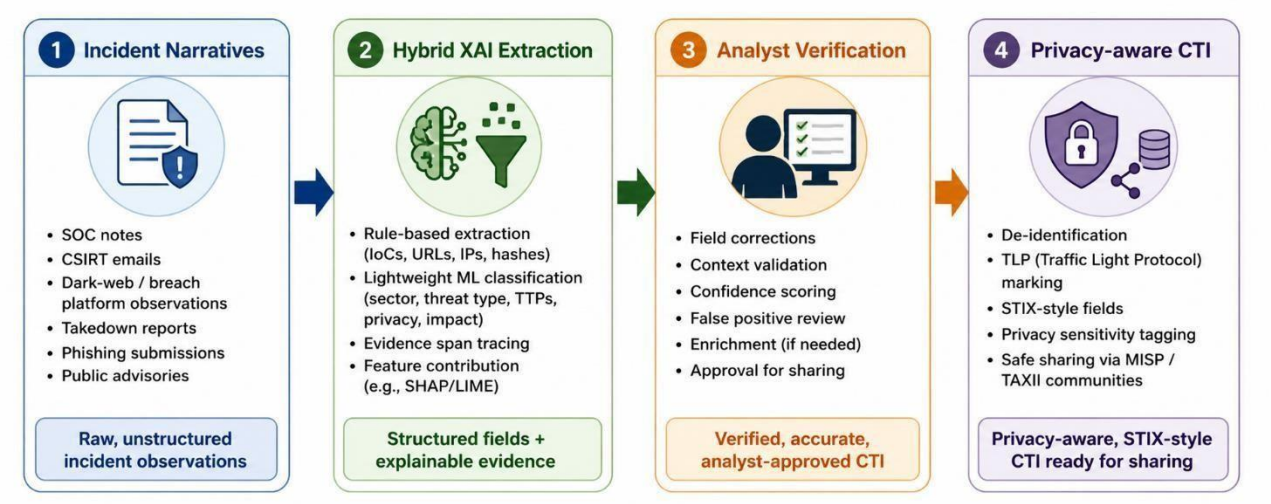

Fig.1 Hybrid explainable AI pipeline for African CTI reporting

c. Feasibility Evaluation Setup

To address reviewer concerns about limited validation, the camera-ready version adds a small feasibility study. Because real African incident reports often contain sensitive information, the evaluation used ten sanitized or synthetic analystrealistic narratives. Six narratives were practitioner-informed Nigerian and African reporting-gap scenarios derived from observed classes of incidents: alleged dark-web exposure, cloned public-service portals, bank-themed phishing, centralbank themed impersonation, electricity-utility payment lures, and payment-platform impersonation. Four narratives were synthetic but designed to reflect common SOC and CSIRT reporting scenarios across African sectors. The aim is not to claim continent-wide generalization or to accuse named organizations of confirmed breaches; the aim is to demonstrate that the reporting pipeline can be implemented, audited, and improved using realistic operational inputs.

The prototype extracted five fields from each narrative: sector, threat type, primary indicator, tactic/technique pattern, and privacy sensitivity. It was compared against manually assigned ground-truth labels. Exact-match field accuracy was used because the immediate reporting task is structured field production rather than free-form summarization. This is a stricter metric than general semantic similarity, but it is appropriate for CTI workflows where fields must be reviewable and machine-readable.

i. Evaluation Narratives

Table 1. Evaluation Narratives

Case	Source	Sector	Threat type	Reporting/visibility issue
C1	sanitized practitioner observation	public sector/maritime	alleged data leak/darkweb sale	Alleged dark-web exposure lacked a readily discoverable structured public CTI advisory.
C2	sanitized practitioner observation	public sector/identity services	phishing/brand impersonation	Cloned immigrationservice style portal had delayed takedown visibility and no clear public CTI bulletin.
C3	practitioner-informed OSINT scenario	financial services	bank-themed phishing or credential harvesting	Bank-related lures may be discussed by analysts, but detailed organization-led CTI reports are not consistently public.
C4	practitioner-informed OSINT scenario	central banking/payment trust ecosystem	CBN-themed impersonation or scam lure	High-trust centralbank branding can be abused while public reporting may remain limited to warnings or informal posts.
C5	practitioner-informed OSINT scenario	electricity utility	payment redirection or brand impersonation	Electricitybill/payment lures require rapid indicator sharing, but structured CTI is often absent or delayed.
C6	practitioner-informed OSINT scenario	fintech/payment platform	Remita-style paymentpage impersonation	Payment-platform impersonation affects users across sectors but may not become
				reusable sector-level intelligence.
C7	synthetic analystrealistic vignette	telecommunications	IoT botnet activity	IoT scanning creates many logs but limited sector-level intelligence.
C8	synthetic analystrealistic vignette	healthcare	ransomware/extortion	Ransomware reporting may omit exfiltration uncertainty and privacy impact.

C9	synthetic analystrealistic vignette	education	phishing/credential theft	Student credential theft often circulates informally before formal reporting.
C10	synthetic analystrealistic vignette	logistics/commerce	infostealer malware	Infostealer activity requires sessiontoken and credentialfocused intelligence.

III. Results and discussion

a. Evaluation Results

Table 2. Evaluation Results

Extracted field	Exact-match accuracy	Interpretation
sector	1.00	Sector labels were identifiable from organization context.
threat_type	0.80	Errors occurred where phishing, impersonation, and malware overlapped.
ioc	0.80	Indicator extraction worked for explicit domains, hashes, and listings, but ambiguous descriptions need analyst review.
ttp	0.70	Higher-level behavior mapping was most difficult and benefits from trainable XAI support.
privacy	0.70	Sensitivity was sometimes underestimated when identity or credential impact was implicit.
Macro average	**0.80**	**Feasible but not sufficient for unsupervised deployment.**

The prototype achieved a macro exact-match accuracy of 0.80. This is intentionally reported as a feasibility result rather than a deployment-ready score. The strongest performance was in sector identification, while TTP and privacy classification were weaker. This pattern is operationally meaningful: the easiest fields are often lexical, while the most important fields for response and sharing require contextual judgment. The result supports the framework's design choice that explainable AI should assist analysts rather than replace them.

The revised version also adds a practical comparison baseline. Manual reporting is flexible but inconsistent and slow when analysts must move from raw narratives to structured monthly CTI. Simple keyword matching is transparent but brittle, especially when the same event contains phishing, malware, privacy, and reputational elements. A black-box classifier may improve labeling but is harder to trust in security operations. The proposed hybrid model occupies a middle ground: it extracts obvious indicators transparently, uses lightweight trainable classification for ambiguous categories, and exposes evidence for analyst verification.

b. Baseline comparison

Table 3. Baseline comparison

Approach	Strength	Weakness	Fit for African CTI reporting

Manual reporting	High human judgment and contextual awareness	Slow, inconsistent, and difficult to scale across sectors	Useful but insufficient for monthly regional visibility
Simple keyword matching	Transparent and easy to audit	Misses context, synonyms, and multi-stage incidents	Useful as a first filter only
Black-box classifier	Can learn complex patterns	Low analyst trust and weak evidence traceability	Risky for legal, public-sector, and privacy-sensitive reporting
Proposed hybrid XAI pipeline	Balances transparency, trainability, evidence spans, and analyst control	Requires curated narratives and feedback loops	Best fit for constrained SOC/CSIRT settings

c. Empirical Observations from African Threat Intelligence Practice

The feasibility evaluation above demonstrates that a hybrid explainable AI pipeline can structure incident narratives into CTI fields, but the stronger motivation for the framework comes from the operational reality of African threat intelligence practice. In many African contexts, the absence of formal incident reports does not mean that cyber incidents are absent. Rather, incidents may be observed by analysts, discussed in private operational channels, appear as alleged claims on breach-monitoring platforms, or surface through phishing and impersonation infrastructure, without becoming structured public or sector-level CTI.

Practitioner experience across African cybersecurity operations shows that analysts frequently encounter suspicious domains, cloned portals, credential-harvesting attempts, financial-sector impersonation, payment-platform abuse, utilitythemed fraud lures, and alleged dark-web exposure claims. However, these observations often remain fragmented. They may appear as screenshots, internal alerts, informal warnings, takedown requests, social-media posts, or adversarycontrolled claims, rather than as standardized CTI reports containing indicators, affected sectors, tactics, techniques, confidence levels, privacy markings, and recommended defensive actions.

This creates a reporting-to-intelligence gap. The problem is not simply that African organizations lack threats to report; it is that observed threats are not consistently converted into reusable intelligence artifacts. This directly affects monthly threat intelligence production. Analysts preparing regional reports may find detailed, searchable, and historically consistent reporting for North America or Europe, while African incidents are represented through scattered fragments that are difficult to verify, normalize, compare, or reuse. As a result, local threat visibility becomes weaker than local threat activity.

Several recurring patterns emerge from analyst-facing African threat monitoring. First, phishing and brand impersonation remain highly visible across public-sector, banking, payment, telecommunications, and utility contexts. Second, website cloning and lookalike portals create persistent risks for identity theft and credential harvesting, particularly where users rely on online public services and payment channels. Third, alleged data exposure and darkweb listing claims may appear before any official advisory or structured disclosure is available. Fourth, sector-specific incidents may be handled internally without producing public CTI artifacts that would help other defenders recognize related infrastructure or tactics.

These observations must be handled carefully. This paper does not present unverified breach-market posts, X posts, or underground claims as confirmed organizational breaches. Nor does it use practitioner observations to accuse named organizations of compromise. Instead, these examples are treated as sanitized evidence of a visibility and reporting problem. From a CTI perspective, even unconfirmed claims can create operational value if they are labeled correctly, assigned confidence levels, de-identified where necessary, and transformed into reviewable intelligence objects for analyst validation.

Open threat-intelligence and breach-aggregation platforms also illustrate the visibility problem. Exploratory use of platforms such as Breach house can show uneven country-level representation, where some regions appear more searchable or more richly indexed than African contexts. Such absence should not be interpreted as absence of incidents. It may reflect differences in disclosure culture, indexing, language, reporting infrastructure, platform coverage, and the maturity of national or sectoral CTI pipelines. This reinforces the argument that African CTI challenges are partly infrastructural and procedural, not merely technical.

The proposed framework directly responds to this gap. Its purpose is not to automate public accusations or replace analyst judgment. Instead, it helps analysts transform fragmented observations into structured, Privacy-aware CTI drafts. A suspicious domain, cloned portal, alleged leak claim, phishing message, or payment impersonation report can be converted into fields such as sector, threat type, indicator, likely tactic or technique, confidence level, privacy sensitivity, and recommended handling. Evidence-span tracing allows the analyst to see why each field was produced, while privacyaware packaging reduces the risk of exposing victims, sensitive infrastructure, or unverified allegations.

This practitioner-grounded evidence strengthens the central argument of the paper: CTI in African security operations should be understood not only as a sharing problem, but also as a reporting and visibility problem. Until local observations become structured and reusable intelligence, African incidents will remain underrepresented in global CTI ecosystems, and analysts will continue to struggle with producing consistent regional intelligence reports. Explainable AI is therefore valuable not because it replaces expertise, but because it supports analysts in turning operational fragments into auditable intelligence.

d. Discussion

The feasibility evaluation and practitioner observations strengthen the original argument. The main contribution is not a claim that a small prototype solves African CTI. The contribution is a practical reporting architecture that recognizes how CTI is actually produced in underreported environments: from messy narratives, uncertain evidence, privacy constraints, platform visibility gaps, and analyst judgment. The practitioner-informed observations show why this matters.A dark-web listing, cloned public-service site, bank-themed phishing campaign, payment-platform impersonation, or utilitypayment scam can be operationally important even before it appears in an official report. Without a structured reporting pathway, such observations remain isolated and disappear from monthly intelligence cycles.

The evaluation also addresses the reviewers' concern about technical novelty. The revised model is not merely a rule based extractor; it is a hybrid XAI workflow. It separates deterministic IoC extraction from trainable classification and adds evidence spans for analyst verification. This makes the contribution suitable for applied AI and cybersecurity: the novelty lies in matching the model architecture to a regional operational bottleneck, not in proposing an abstract new algorithm detached from SOC practice.

Finally, the African focus is not decorative. It changes the design requirements. Reporting systems must protect sensitive identities, avoid unverified public accusations, support de-identification, and work even where formal sectoral sharing structures are immature. The framework therefore emphasizes privacy-aware sharing and confidence labeling. For example, an analyst can publish a de-identified indicator package about a cloned public-service portal without exposing victims or overstating attribution.

i. Limitations and Future Work

This paper remains limited by the small evaluation set and the absence of multi-organization testing. The practitioner observations were sanitized and used as motivating cases, not as independently verified public breach claims. The synthetic narratives improve coverage but cannot replace real anonymized incident reports. Future work should therefore collect deidentified narratives from one or two African SOCs, CSIRTs, or sectoral CERT communities under ethical approval and data-handling agreements.

The paper also deliberately avoids presenting unverified dark-web, X, or breach-platform claims as confirmed organizational breaches. This conservative treatment improves legal and ethical safety, reduces plagiarism and misinformation risk, and aligns with responsible CTI practice. Where named brands appear in the evaluation narratives, they are used only to represent common targeting and impersonation patterns in Nigeria's digital ecosystem, not as verified accusations of compromise.

Future work should also implement the trainable component more fully. A small transformer or gradient-boosted classifier can be trained on anonymized incident narratives, with SHAP or LIME explanations used to show why labels were assigned. The system should then be evaluated using inter-analyst agreement, time-to-report, field completeness, privacy correctness, and downstream utility for monthly TI reporting. These additions would turn the present feasibility study into a stronger empirical system paper.

IV. Conclusion

This paper improves the accepted version by addressing reviewer concerns about real-world validation, technical depth, and African contextual evidence. It argues that CTI in Africa should be treated not only as a sharing problem, but as a reporting and visibility problem. The revised version adds practitioner-informed Nigerian and African scenarios involving

alleged dark-web exposure, cloned public-service portals, bank-related phishing, central-bank themed impersonation, utility-payment fraud, and payment-platform impersonation, while carefully avoiding unverified breach claims. The proposed hybrid explainable AI framework converts incident narratives into structured, Privacy-aware CTI while preserving analyst control. A small feasibility evaluation using ten analyst-realistic narratives achieved a macro exact-match accuracy of 0.80 and revealed where human verification remains necessary. The result is a more practical, transparent, and regionally grounded contribution to AI-enabled security operations.

References

[1] African Union, *African Union Convention on Cyber Security and Personal Data Protection (Malabo Convention)*,2014. Available: https://au.int/en/treaties/african-union-convention-cybersecurity-and-personal-data-protection

[2] S. Barnum, *Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX)*, MITRE Corporation, 2014. Available: https://stixproject.github.io/

[3] Breach.house, *Breach intelligence search platform*, 2026. Accessed: Apr. 29, 2026. Available:https://breach.house/

[4] E. M. Hutchins, M. J. Cloppert, and R. M. Amin, "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains," in *Leading Issues in Information Warfare & Security Research*, vol. 1, pp. 80–106, 2011.

[5] INTERPOL, *African Cyberthreat Assessment Report*, 2021. Available: https://www.interpol.int/

[6] INTERPOL, *African Cyberthreat Assessment Report*, 2024. Available: https://www.interpol.int/

[7] C. Johnson, L. Badger, D. Waltermire, J. Snyder, and C. Skorupka, *Guide to Cyber Threat Information Sharing (NIST SP 800-150)*, National Institute of Standards and Technology, 2016.
doi: 10.6028/NIST.SP.800-150

[8] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proc. 31st Int. Conf. Neural Information Processing Systems (NIPS)*, pp. 4765–4774, 2017.

[9] V. Mavroeidis and S. Bromander, "Cyber threat intelligence model: An evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence," in *Proc. European Intelligence and Security Informatics Conf. (EISIC)*, pp. 91–98, 2017. doi:
10.1109/EISIC.2017.20

[10] MITRE, *MITRE ATT&CK: Design and philosophy*, 2024. Available: https://attack.mitre.org/

[11] M. A. Nainna and J. Bass, "Cyber Threat Intelligence Sharing in Nigeria," *Communications of the IIMA*, vol. 22, no. 1, 2024. Available: https://scholarworks.lib.csusb.edu/ciima/vol22/iss1/1/

[12] OASIS, *STIX Version 2.1 and TAXII Version 2.1 Specifications*, 2021. Available: https://oasisopen.github.io/ctidocumentation/

[13] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," in *Proc. ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, pp. 1135–1144, 2016. doi:10.1145/2939672.2939778

[14] A. I. Sani and I. Yakub, "Evaluating Nigeria's readiness against state-sponsored cyber attacks: A comparative study of cybersecurity policies, incident response, and international cooperation," *Journal of Computational Analysis and Applications*, 2025.

[15] F. Skopik, G. Settanni, and R. Fiedler, "A problem shared is a problem halved: A survey on the dimensions of collective cyber defense through security information sharing," *Computers & Security*, vol. 60, pp. 154–176, 2016. doi: 10.1016/j.cose.2016.04.003

[16] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *Proc. IEEE Symposium on Security and Privacy*, pp. 305–316, 2010. doi: 10.1109/SP.2010.25

[17] W. Tounsi and H. Rais, "A survey on technical threat intelligence in the age of sophisticated cyber attacks," *Computers & Security*, vol. 72, pp. 212–233, 2018. doi:
10.1016/j.cose.2017.09.001

[18] C. Wagner, A. Dulaunoy, G. Wagener, and A. Iklody, "MISP: The design and implementation of a collaborative threat intelligence sharing platform," in *Proc. ACM Workshop on Information Sharing and Collaborative Security*, pp. 49–56, 2016. doi:
10.1145/2994539.2994542

[19] FIRST, *Traffic Light Protocol (TLP) Definitions and Usage Guidance*, 2020. Available: https://www.first.org/tlp/

[20] National Information Technology Development Agency (NITDA), *Nigeria Computer Emergency Readiness and Response Team advisories*, 2022. Available: https://cert.gov.ng/

[21] Nigeria Data Protection Commission (NDPC), *Nigeria Data Protection Act*, 2023. Available: https://ndpc.gov.ng/