International Journal of Modern Computer Science and IT Innovations (IJMCSIT)
**eISSN: 3087-4289**
Volume. 03, Issue. 05, pp. 19-27, May 2026"

# Application of Artificial Intelligence in Digital Risk Protection and External Threat Intelligence

**Kolchin Rustam**
SoftLine PJSC
Almaty, Kazakhstan

Article Received: 12/03/2026, Article Accepted: 16/04/2026, Article Published: 29/05/2026
DOI: https://doi.org/10.55640/ijmcsit-v03i05-04



## ABSTRACT

Organizations now face external digital risks from look-alike domains, leaked credentials, dark web posts, impersonation profiles, and public technical traces that security teams cannot control. The article examines artificial intelligence in Digital Risk Protection and external threat intelligence as an analytical layer for collecting, classifying, prioritizing, and routing external risk signals. The study draws on recent academic publications, systematic reviews, industry guidance, and the MITRE ATT&CK Reconnaissance Framework. Source analysis, comparative analysis, conceptual synthesis, typological classification, and analytical generalization guide the research. The article distinguishes Digital Risk Protection from classical cyber threat intelligence. It explains the use of AI in domain similarity detection, homoglyph analysis, dark web monitoring, language-based extraction, and alert ranking. It replaces a linear workflow figure with a table that connects monitored sources, AI functions, and operational outputs. The proposed model suits enterprise security teams that need defensible prioritization before takedown, identity response, legal review, or SOC escalation.

## KEYWORDS

artificial intelligence, digital risk protection, external threat intelligence, dark web monitoring, phishing detection.

## INTRODUCTION

Corporate security teams now deal with risk signals that form outside the systems they administer. Phishing domains, leaked credentials, spoofed profiles, exposed code fragments, public technical databases, social media posts, dark web forums, paste sites, and attacker-controlled infrastructure can affect an organization before an endpoint alert or network event appears. A security operations center may see the final stage of the attack, while the preparation phase often unfolds in public and semi-closed sources.

Digital Risk Protection (hereafter, DRP) addresses this external layer. The category concerns threats tied to public identity, brands, employees, domains, data exposure, and reputational surfaces. External threat intelligence covers a wider range of adversary behavior, technical indicators, tactics, infrastructure, and campaign context.

The research aim is to define how artificial intelligence can improve DRP and external threat intelligence without turning the process into uncontrolled alert generation. The first objective is to distinguish DRP from classical threat intelligence, brand protection, and external attack surface monitoring. The second objective is to examine AI-based methods for detecting phishing domains, homoglyph attacks, typosquatting, leaked data, and unstructured threat narratives in open, deep, and dark web sources. The third objective is to develop an implementation logic for AI-assisted prioritization, analyst validation, and response routing in enterprise security operations.

The novelty of the article lies in treating DRP as a decision layer between external collection and security response. AI assists that layer through similarity scoring, language extraction, source clustering, entity linking, and explainable prioritization. Analyst judgment remains necessary for evidence review, legal sensitivity, and final escalation. The working hypothesis is that AI delivers practical value in DRP only when teams combine detection models, source governance, confidence scoring, and human validation within a single decision workflow.

## MATERIALS AND METHODOLOGY

The materials consist of ten sources published between 2021 and 2025. The corpus combines peer-reviewed studies on phishing URL detection, homoglyph detection, cybercrime intelligence, dark web monitoring, and LLM-based cyber threat detection with two frameworks and industry documents relevant to reconnaissance and external threat intelligence. The first group covers domain impersonation, phishing websites, lexical URL features, visual similarity, and deep learning classification as studied by Almuhaideb et al. (2022), Haq et al. (2024), and Zieni et al. (2023). The second group examines cybercrime intelligence, surface web, deep web, and dark web collection, anonymity constraints, website features, and business-oriented dark web monitoring in works by Cascavilla et al. (2021) and Dalvi and Bhirud (2024). The third group concerns LLM-assisted threat detection, CTI extraction, model vulnerabilities, and privacy risk mitigation, as addressed in studies by Chen et al. (2024), the European Data Protection Board (2025), and Jaffal et al. (2025). The fourth group frames operational expectations for threat intelligence, DRPS, external attack surface management, MITRE ATT&CK enrichment, and reconnaissance techniques, drawing on MITRE (2025) and Nunez et al. (2024).

The study uses comparative analysis to separate DRP from adjacent security categories. Source analysis extracts recurring methodological positions from the selected literature. Conceptual synthesis builds an AI-assisted DRP workflow. Typological classification groups detection and prioritization scenarios. Analytical generalization supports practical recommendations for enterprise security operations without claiming experimental validation.

## RESULTS AND DISCUSSION

DRP focuses on external risk signals related to the organization's public identity, exposed data, digital assets, and reputation. Classical cyber threat intelligence often begins with adversary behavior, indicators of compromise, malware infrastructure, vulnerability exploitation, or TTP mapping. DRP begins with the attacker's view of the organization: domains that resemble official names, employee identities that can support social engineering, leaked credentials that can open access paths, and public conversations that can trigger reputational or fraud risk.

The Gartner market guide places digital risk protection services, external attack surface management, threat hunting, and threat exposure management near the broader threat intelligence market (Nunez et al., 2024). That placement helps separate DRP from ordinary indicator feeds. A feed can report a suspicious domain or IP address. A DRP process links that signal to a protected brand, an exposed identity, a business process, a legal response path, and an operational owner. The unit of work becomes a signal-to-decision chain: detection, enrichment, confidence assessment, prioritization, escalation, and remediation.

The MITRE ATT&CK reconnaissance tactic gives DRP a defensive monitoring map. MITRE describes reconnaissance as active or passive adversary collection of information that supports targeting, including details about victim organizations, infrastructure, and staff (MITRE, 2025). Security teams can use the same map in reverse. If adversaries search domains, DNS records, email addresses, employee names, code repositories, social media, and closed sources before initial access, DRP teams monitor those surfaces to detect preparation before exploitation.

A systematic multi-vocal review of cybercrime threat intelligence examines the surface, deep, and dark web. It identifies indicators, risk parameters, anonymity degrees, policy constraints, and website features as recurring variables (Cascavilla et al., 2021). A platform that gathers dark web mentions, paste dumps, new domain registrations, and social media references without risk parameters leaves analysts with fragments. Analysts still need to connect fragments with protected entities, plausible attack paths, source credibility, and response ownership.

Dark web monitoring research adds the verification problem. Dalvi and Bhirud (2024) discuss stolen data, zero-day discussions, malware distribution, and hidden

communication channels as sources for business cybersecurity monitoring. Their review points to a practical constraint for DRP: criminal spaces mix actionable evidence with recycled leaks, false claims, scam offers, duplicate dumps, and seller exaggeration. AI can cluster and classify this material, but analysts need source history, sample validation, technical indicators, and organizational relevance before they escalate a finding.

AI-assisted DRP can be grouped into four functional scenarios. The first scenario concerns domain and URL similarity. The second concerns unstructured text analysis in open, deep, and dark web sources. The third concerns entity extraction and enrichment. The fourth concerns prioritization and routing. Each scenario uses a different technical logic, so a single general model cannot handle the full workflow with equal reliability.

Domain and URL detection has the strongest methodological base among these scenarios. Zieni et al. (2023) classify phishing website detection approaches into list-based, similarity-based, visual-similarity, and machine-learning-based methods. Their survey shows that URL structure, page content, third-party information, and visual properties all contribute to phishing detection. In DRP, this matters before a user clicks a link. A newly registered domain that resembles a protected brand, uses a suspicious certificate, copies visible page elements, or appears in a messaging campaign already has a risk profile.

Haq et al. (2024) propose a deep learning system using a one-dimensional convolutional neural network for phishing URL detection. Their study supports the technical feasibility of URL-based classification, whereas DRP introduces an earlier, less explicit problem. External monitoring often sees newly registered domains, inactive domains, parked pages, privacy-protected registrations, and homoglyph variants before malicious content appears. Security teams need URL classification, but they also need similarity scoring, infrastructure enrichment, and watchlist logic for domains prepared for later use.

Homoglyph attacks create a separate detection problem. Visual similarity can hide technical differences when attackers use Unicode substitution, mixed-script characters, or look-alike symbols. Almuhaideb et al. (2022) propose a model that combines a hash function with machine learning to detect homoglyph attacks on forged websites and in phishing scenarios. DRP detection logic benefits from a layered approach: normalized string comparison, Unicode script inspection, edit-distance checks, brand dictionary matching, certificate and registrar enrichment, hosting reputation, and screenshot similarity when content exists.

The comparison of these phishing and domain studies gives a precise boundary for the first objective. Gateway anti-phishing tools protect users at the point of interaction. DRP detects external risk formation before the interaction enters email, browser traffic, or endpoint telemetry. A gateway blocks a URL. A DRP platform detects the domain, enriches it, estimates brand relation, checks infrastructure, and routes takedown or investigation. AI assists at earlier stages by evaluating similarity, clustering variants, and identifying suspicious infrastructure patterns.

The second AI scenario concerns unstructured external data. Dark web forums, paste sites, Telegram channels, social media discussions, breach marketplaces, and threat actor posts do not follow stable schemas. Organization names appear with abbreviations, transliterations, mistakes, screenshots, partial identifiers, and indirect references. Keyword matching misses part of this material and returns many irrelevant matches. NLP and LLM-based pipelines address this problem by extracting entities, relationships, credentials, malware names, actor aliases, infrastructure references, and intent signals from noisy text.

A survey by Chen et al. (2024) reports LLM use in cyber threat detection for threat discovery, anomaly interpretation, detection assistance, and domain-specific security tasks. DRP applies that language capability to uncertain external evidence. A dark web post may contain an organization name, an internal system label, a sample credential, and a price claim. An LLM-assisted pipeline can segment the post, identify entities, link them to an asset inventory, compare them with past incidents, and prepare a provisional analyst brief. The brief reduces time spent on translation, clustering, extraction, and initial narrative reconstruction.

Jaffal et al. (2025) survey LLM applications, vulnerabilities, and defense techniques across several cybersecurity domains, including threat intelligence and social engineering. Their work matters for DRP because platforms that process adversarial text can ingest prompts, poisoned instructions, and deceptive summaries embedded in threat actor posts. A secure

LLM-assisted DRP architecture requires input isolation, retrieval controls, prompt hardening, output validation, and audit logs for model-assisted decisions.

The European Data Protection Board (2025) frames LLM deployment through privacy risk management, data protection by design, data protection by default, and security of processing. External threat intelligence often processes employee names, email addresses, leaked credentials, pseudonymous forum handles, contact details, and screenshots containing personal data. DRP pipelines need data minimization, source classification, role-based access, retention rules, and human review for sensitive findings. These controls belong at ingestion and validation, not after a report circulates across teams.

The third scenario concerns enrichment. A raw external signal rarely carries enough meaning on its own. A domain variant becomes operationally significant after enrichment with registration date, registrar, name servers, certificate metadata, hosting provider, passive DNS, visual screenshot, brand similarity score, and relation to known campaigns. A leaked credential gains priority when analysts link it to a privileged account, active employee, reused password pattern, public executive profile, or a recent phishing campaign. Nunez et al. (2024) describe threat intelligence capabilities such as indicator ratings, enrichment, integration, alert thresholds, contextual dashboards, investigative support, and MITRE ATT&CK mapping. DRP needs these capabilities to move from signal collection to decision support.

The fourth scenario concerns prioritization. External monitoring can produce thousands of weak signals: brand mentions, look-alike domains, expired credential dumps, low-credibility forum posts, suspicious profiles, counterfeit pages, and social media rumors. Ranking them by keyword frequency or source category leaves analysts with unstable queues. A useful prioritization model combines source reliability, confidence, freshness, protected-entity match, asset criticality, exploitability, campaign linkage, possible user impact, and available response path. If a model hides these factors, analysts cannot contest false positives or missed risks.

Table 1 presents an AI-assisted DRP and external threat intelligence processing loop. It adapts threat intelligence use cases described by Nunez et al. (2024) and aligns monitored surfaces with MITRE ATT&CK reconnaissance logic (MITRE, 2025).

**Table 1. AI-assisted DRP and external threat intelligence processing loop, adapted from Nunez et al. (2024) and MITRE (2025)**

| Workflow stage | Source or action layer | AI-assisted function | Operational output |
|---|---|---|---|
| External source monitoring | Open web, social media, messaging channels, code repositories, DNS and certificate data, deep web, dark web, paste sites, leaked credential sources | Source labeling, language detection, entity recognition, and duplicate reduction | Structured external signal set |
| Collection and normalization | Crawlers, APIs, source connectors, timestamped records, entity dictionaries | Text normalization, entity matching, source classification, and noisy record filtering | Comparable and searchable DRP dataset |
| AI-assisted detection | Domain variants, homoglyph patterns, phishing URLs, leaked credentials, threat actor posts, screenshots | Similarity scoring, URL classification, NLP extraction, topic clustering, screenshot comparison | Candidate threats grouped by type and protected entity |
| Enrichment and correlation | WHOIS data, DNS records, certificate metadata, passive DNS, asset inventory, identity records, source reputation | Entity linking, campaign clustering, MITRE ATT&CK reconnaissance mapping, confidence estimation | Contextualized finding with source and asset relevance |

| | | | |
|---|---|---|---|
| Risk scoring and prioritization | Confidence level, freshness, asset criticality, source reliability, campaign linkage, legal sensitivity, response feasibility | Weighted scoring, supervised classification, and analyst queue ordering | Prioritized DRP alert with explainable risk factors |
| Analyst validation | Evidence review, false-positive control, sensitive data handling, escalation assessment | Assisted summarization, evidence grouping, and validation checklist support | Confirmed, rejected, or deferred finding |
| Response routing | Takedown request, credential reset, SOC escalation, legal review, communications response, executive notification | Recommended route selection, owner mapping, and case enrichment | Assigned remediation or investigation action |
| Feedback loop | Confirmed incidents, false positives, failed takedowns, stale sources, validated leaks, analyst notes | Source reputation update, model retraining input, rule refinement | Improved detection logic and calibrated prioritization |

AI produces value in DRP when analysts place it between noisy external sources and structured response decisions. Haq et al. (2024) and Zieni et al. (2023) report phishing URL model performance under defined datasets. Almuhaideb et al. (2022) offer methods for visual deception and Unicode-based impersonation. Cascavilla et al. (2021) and Dalvi and Bhirud (2024) explain why external collection requires verification, source governance, and cautious handling of criminal claims. Chen et al. (2024), the European Data Protection Board (2025), and Jaffal et al. (2025) support the use of language models for cyber threat detection and CTI extraction. In contrast, privacy guidance requires controls for the processing of personal and sensitive data.

A collection engine reports that a domain, credential, mention, or post exists. A decision system explains why the finding concerns a protected entity, what evidence supports the risk level, which owner receives the case, and which action follows. Organizations that deploy DRP as a second alert queue outside the SOC force analysts to manually reconstruct evidence, confidence, and business relevance.

Before monitoring starts, the security team defines official domains, high-value brands, executive names, employee email patterns, product labels, cloud tenant names, public repositories, partner-facing portals, and terms that attackers could reuse in phishing or extortion. Each entity receives a sensitivity profile. A customer authentication domain carries a different profile from a marketing microsite. A generic employee email pattern differs from a privileged administrator account. AI can cluster variants, but people define why the variants matter.

Source governance forms the second layer, open web, social media, dark web, paste sites, messaging channels, code repositories, DNS records, and certificate datasets differ in legality, reliability, volatility, and evidentiary value. A DRP program needs source classes. Public technical databases usually carry lower legal risk and strong technical value. Dark web forums can provide higher intelligence value in selected cases, but teams need controlled access, source reputation tracking, and retention limits. Messaging channels can contain personal data, screenshots, and unverifiable claims. One extraction policy across all sources creates avoidable legal and analytical risk.

DRP architects should use narrow methods when solving the task. Edit distance, Unicode normalization, domain embeddings, and screenshot comparison improve impersonation detection beyond a general LLM alone. LLMs fit linguistically complex input: multilingual posts, fragmented dark web discussions, actor claims, breach advertisements, and long-form external narratives. This division reduces cost, improves explainability, and limits exposure to prompt manipulation.

Table 2 compares the main AI-supported DRP tasks and the operational controls needed for each task. The table separates detection logic from response logic because the

same AI output can have different values depending on source quality, confidence, and the available remediation path.

**Table 2. AI-supported DRP tasks and operational controls**

| DRP task | Suitable AI or analytical method | Main input | Primary output | Required control |
|---|---|---|---|---|
| Look-alike domain detection | Edit distance, Unicode normalization, domain embeddings, classifier models | Domain registrations, passive DNS, certificate logs | Similarity score and suspicious domain cluster | Brand dictionary validation and false-positive review |
| Homoglyph detection | Unicode script analysis, hashing, supervised classification | IDN domains, visual domain variants | Visual deception indicator | Manual confirmation for high-impact assets |
| Phishing URL classification | URL feature extraction, CNN, or transformer-based models | URLs, page metadata, page content | Maliciousness probability | Sandbox isolation and dataset drift monitoring |
| Dark web leak detection | NLP entity extraction, topic clustering, source reputation scoring | Forum posts, paste dumps, breach listings | Matched entity, leak category, confidence level | Sample verification and data minimization |
| Threat narrative summarization | LLM-assisted extraction and summarization | Long-form posts, actor claims, and incident discussions | Analyst brief and extracted indicators | Prompt isolation, source citation, human approval |
| Reputational risk monitoring | Clustering, language detection, cautious sentiment analysis | Social media, forums, news-like sources | Signal cluster and escalation candidate | Communications review and authenticity check |
| Prioritization | Weighted scoring and supervised classification | Enriched findings, asset criticality, source quality | Risk queue and response route | Explainable scoring and periodic calibration |

Table 2 indicates that AI in DRP needs a modular design. A general-purpose model placed over all sources creates fragile automation. A task-specific pipeline produces stronger results because each method has a defined input, output, and control. This modularity strengthens auditability. When legal, SOC, communications, or identity teams receive an alert, they can see which part came from similarity detection, which from source enrichment, and which from analyst validation.

Prioritization requires a scale that combines technical probability and business consequence. A suspicious domain resembling a protected brand may remain low priority while inactive and parked. The same domain becomes a higher priority after DNS activation, a valid certificate, mail exchange records, copied page content, traffic evidence, or appearance in a phishing kit. A credential dump may stay low priority if the accounts are obsolete. It becomes urgent when it contains active privileged identities, recent timestamps, or overlaps with a known campaign. DRP scoring needs source confidence, asset value, freshness, exploit readiness, exposure channel, and response feasibility.

Human validation remains part of the control loop; criminal forums contain false advertising. Threat actors exaggerate access. Paste dumps mix old and new data. Public accusations may reflect ordinary criticism, reputational attacks, or coordinated fraud. AI can prepare

evidence, but analysts check sample authenticity, entity match, source reliability, duplication, legal sensitivity, and operational ownership. Confirmed and rejected findings, then update scoring logic.

Table 3 proposes a response matrix for AI-assisted DRP alerts. The table connects alert categories with evidence thresholds, owners, initial actions, and escalation conditions. This structure prevents DRP from turning into a passive dashboard.

**Table 3. Response matrix for AI-assisted DRP alerts**

| Alert category | Evidence threshold | Primary owner | Initial action | Escalation condition |
|---|---|---|---|---|
| Confirmed phishing domain | Domain similarity plus malicious content, mail records, or campaign evidence | SOC and legal or takedown team | Block, report, initiate takedown, preserve evidence | Active credential collection or customer exposure |
| Suspicious look-alike domain | Similarity score plus registration or certificate signal | Threat intelligence team | Monitor, enrich, watch for content changes | Activation, MX records, hosting change, brand misuse |
| Leaked employee credential | Account match plus sample verification | IAM or identity security team | Reset password, check MFA, review access logs | Privileged account, recent timestamp, repeated exposure |
| Dark web breach claim | Organization match plus source history and partial evidence | Threat intelligence and legal | Validate sample, classify sensitivity, restrict circulation | Verified customer, employee, or regulated data |
| Executive impersonation | Profile similarity plus outreach or fraud evidence | Fraud, legal, communications | Report platform abuse, notify targeted groups | Payment request, credential request, and public reputational impact |
| Source code or secret exposure | Repository or paste match plus secret pattern | AppSec or DevSecOps | Revoke secret, rotate token, remove exposure | Production credential, customer data, active exploitation |
| Reputational campaign | Coordinated cluster plus abnormal amplification | Communications and risk management | Verify source pattern, prepare response option | Media pickup, customer impact, regulatory attention |

The matrix shows that DRP output gains operational value after it is routed into existing security and business processes. A phishing domain without takedown capability remains an observation. A leaked credential without an identity response remains a warning. A reputational cluster without communications ownership remains noise. AI can rank and summarize findings, but the organization needs response channels before automation can reduce risk.

Several implementation principles follow from this model. The DRP program should begin with a limited set of protected entities and high-confidence sources. Broad coverage added too early increases false positives and weakens analyst trust. Every AI-generated risk score should expose its factors. A score without source, freshness, entity match, and confidence cannot support accountable decisions. LLM-assisted summaries should preserve links to source artifacts, extracted fields, and analyst notes. A polished summary without evidence can mislead decision-makers.

Privacy rules need placement at ingestion. Sensitive personal data requires masking, access control, and defined retention. The feedback loop then handles confirmed phishing, failed takedowns, false positives, stale sources, and validated dark web leaks as input for source reputation and detection rules.

The proposed approach has limits. AI-assisted DRP cannot guarantee visibility in closed criminal communities. Access to the dark web depends on the availability of sources, language coverage, forum rules, and operational safety. Domain detection cannot determine intent with certainty before activation. LLM summaries may misread adversarial text or overstate confidence. Legal rules vary across jurisdictions, especially when teams monitor leaked personal data or closed sources. These limits define the boundaries for responsible AI-assisted DRP.

A mature DRP capability has a layered architecture: entity inventory, source governance, task-specific AI models, enrichment, explainable scoring, analyst validation, response routing, and feedback. The architecture supports the article's hypothesis. AI improves DRP when teams place it inside a controlled workflow. Detached from governance and response, it mainly increases the speed at which uncertain external data reaches analysts.

## CONCLUSION

Digital Risk Protection forms a separate external intelligence layer between classical cyber threat intelligence, anti-phishing detection, brand protection, and external attack surface monitoring. Its object is the chain through which an external signal becomes a verified business, technical, legal, or reputational risk. This distinction explains why DRP requires entity modeling, source classification, enrichment, prioritization, and response ownership.

AI supports DRP through domain and URL similarity detection, homoglyph and typosquatting analysis, NLP-based extraction from unstructured open and dark web sources, and risk prioritization across heterogeneous signals. Machine learning and deep learning methods support phishing and domain detection. LLM-focused methods support text extraction, summarization, and CTI assistance. Source validation, privacy controls, prompt isolation, and analyst review determine whether those methods improve security operations.

The proposed implementation logic confirms the working hypothesis. AI strengthens DRP when teams combine confidence scoring, source governance, human validation, and defined response routes. It structures external uncertainty, reduces repetitive triage, and helps security teams move from scattered external observations to defensible decisions.

## References


1. Almuhaideb, A. M., Aslam, N., Alabdullatif, A., Altamimi, S., Alothman, S., Alhussain, A., Aldosari, W., Alsunaidi, S. J., & Alissa, K. A. (2022). Homoglyph attack detection model using machine learning and a hash function. Journal of Sensor and Actuator Networks, 11(3), Article 54. https://doi.org/10.3390/jsan11030054

2. Cascavilla, G., Tamburri, D. A., & Van Den Heuvel, W.-J. (2021). Cybercrime threat intelligence: A systematic multi-vocal literature review. Computers & Security, 105, Article 102258. https://doi.org/10.1016/j.cose.2021.102258

3. Chen, Y., Cui, M., Wang, D., Cao, Y., Yang, P., Jiang, B., Lu, Z., & Liu, B. (2024). A survey of large language models for cyber threat detection. Computers & Security, 145, Article 104016. https://doi.org/10.1016/j.cose.2024.104016

4. Dalvi, A., & Bhirud, S. (2024). Dark web monitoring as an emerging cybersecurity strategy for businesses. International Journal of Information Engineering and Electronic Business, 16(2), 54-67. https://doi.org/10.5815/ijieeb.2024.02.05

5. European Data Protection Board. (2025). AI privacy risks & mitigations: Large language models. https://www.edpb.europa.eu/our-work-tools/our-documents/support-pool-experts-projects/ai-privacy-risks-mitigations-large_en

6. Haq, Q. E. U., Faheem, M. H., & Ahmad, I. (2024). Detecting phishing URLs based on a deep learning approach to prevent cyber-attacks. Applied Sciences, 14(22), Article 10086. https://doi.org/10.3390/app142210086

7. Jaffal, N. O., Alkhanafseh, M., & Mohaisen, D. (2025). Large language models in cybersecurity: A survey of applications, vulnerabilities, and defense techniques. AI, 6(9), Article 216. https://doi.org/10.3390/ai6090216

8. MITRE. (2025). Reconnaissance, tactic TA0043: Enterprise. MITRE ATT&CK. https://attack.mitre.org/tactics/TA0043/

9. Nunez, J., Contu, R., & Schneider, M. (2024). Market guide for security threat intelligence products and services (ID G00794923). Gartner.

**10.** Zieni, R., Massari, L., & Calzarossa, M. C. (2023). Phishing or not phishing? A survey on the detection of phishing websites. IEEE Access, 11, 18499-18519. https://doi.org/10.1109/ACCESS.2023.3247135