International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Artificial Intelligence and the Future of Cybersecurity: Threats, Defenses, and Strategic Imperatives

Rohan Kapoor

Abstract: High-speed networks and rapid advances in artificial intelligence give rise to new offensive and defensive cyber operations. Offensive AI can provide targeted phishing, advanced and adaptive malware, and automated exploitability assessment. These tools lower the barrier for entry. This paper also considers autonomous behavior in AIOps, such as machine-speed detection of anomalies and security self-adaptation. Employing a systematic literature review, we characterize how, across three complementary dimensions, threat amplification, defensive enablement and calculated governance, AI-enabled tools and techniques are shifting the balance of the cybersecurity landscape. Our review analyzes 20 peer-reviewed and non-peer-reviewed industry articles published between 2015 and 2026, covering attack vectors, AI-enabled detection architecture, and governing frameworks, including the NIST AI RMF, EU AI Act, and NIST AI 600-1. The quantitative pillars are AI Security Return on Investment, AI Cyber Anomaly Index, Detection-Exposure Gap, and Organizational Cyber Risk Score, or ASROI, ACAI, DEG, and OCRS, respectively. The report found that organizations that have deployed AI-enabled defensive systems experience a 50% reduction in MTTD and a $2.22 million lower average cost of a breach compared to a non-AI environment. The results are relevant to the asymmetric setting in which the offense is much more capable than the defense, as this scenario is the setting that regulatory systems have not yet adapted to.

Keywords: *Artificial Intelligence, Cybersecurity, Threat Detection, Zero Trust Architecture, AI Governance, Machine Learning, Generative AI*

1. Introduction

The convergence of artificial intelligence and cybersecurity represents one of the most consequential technological shifts of the current decade. Unlike prior generations of security tools, AI systems can operate at machine speed, adapt to novel adversarial stimuli without human intervention, and generate or analyze content at a scale that fundamentally alters the cost calculus for both attackers and defenders. This shift is not incremental. Generative models, reinforcement learning agents, and autonomous orchestration pipelines are changing what is possible on both sides of the security boundary simultaneously, making the threat environment more dynamic and less predictable than at any prior point in the discipline's history [1]. Cybersecurity and AI can be viewed as one of the seminal technology battlegrounds of the twenty-first century. Over the past decade, AI tools have proliferated from research settings to both the offense and defense sides of the cybersecurity landscape, resulting in what the World Economic Forum's Global Cybersecurity Outlook 2026 refers to as a structural fragmentation of the global security landscape [10]. As AI systems increasingly complement or supplant human analysts to create threats, detect and respond to attacks, and take action at speeds several orders of magnitude faster than human cognition, establishing a fact-based approach to measuring, calculating and governing cybersecurity risk in organizations becomes increasingly critical.

It is worth noting that, on the financial side, IBM Security's Cost of a Data Breach Report 2025 stated that the average cost of a data breach had reached $4.88 million, and organizations without mature AI deployments absorbed costs that were, on average, $2.22 million higher than those with mature AI deployments [8]. As attackers also have access to generative AI models, these models have changed the opponent's calculus. For example, they allow for realistic phishing campaigns that are practically indistinguishable from legitimate communications, producing working exploit code and reconnaissance of the enterprise attack surface [16, 20]. The same technologies have been made available via publicly

Independent Researcher, USA
ORCID: 0009-0008-5130-1694

accessible APIs and underground marketplaces, broadening the set of threat actors to include low-resourced individuals and state-sponsored groups [9].

This paper focuses on three research questions: To what extent has AI changed the taxonomy of offensive cyber threats between 2020 and 2026? What AI architectural/analytical techniques have been shown to improve defensive outcomes? Third, what governance and regulatory frameworks are sufficient to address the systemic risks introduced by AI-enabled cyber operations? Fourth, what workforce strategies and capability investments are required to operationalize and sustain AI-driven security programs at an organizational level? The review deploys quantitative formulae to derive computable indicators from the reviewed empirical evidence, effectively furnishing practitioners with metrics from the literature. This article summarizes the peer-reviewed academic literature, industry reports from groups such as IBM and the World Economic Forum, and norms from standards organizations such as the National Institute of Standards and Technology and the European Union. The field is still evolving while this article is being written, so it summarizes literature published since 2020 unless the historical context of an older publication is useful or necessary. While existing reviews have examined either offensive AI capabilities or defensive architectures in relative isolation, few have simultaneously integrated quantitative metrics, governance analysis, and workforce imperatives into a unified diagnostic framework. This article addresses that gap by proposing four computable indicators drawn directly from the empirical literature, enabling practitioners to benchmark organizational posture rather than rely solely on qualitative assessment.

2. Method

The systematic literature review was conducted using the three steps of source identification, eligibility screening and thematic synthesis. The systematic review protocol was prepared to be traceable, reproducible and consistent with the research questions detailed in Section 1.

Keywords were searched on the IBM Research IBM Xplore, ACM Digital Library, Springer Link, ScienceDirect and Frontiers websites. The keywords were "Artificial intelligence cybersecurity," "AI threat detection," "Generative AI cyberattacks," "Zero trust AI," "Explainable AI security," "AI governance frameworks," and "AI security ROI" (return on investment). Authoritative and empirical industry reports such as IBM Security publications, World Economic Forum, the National Institute of Standards and Technology (NIST), and the European Union (EU) were included. Only peer-reviewed journal articles, conference proceedings, whitepapers from leading industry organizations, and formal regulations were included. These publications spanned the years 2015 through 2026 to include both seminal and up-to-date empirical publications. In addition, only research dealing with the intersection of the two fields (i.e. AI and cybersecurity) and not the fields in isolation was included. The numerical analysis of the article included studies that incorporated quantitative data. The thematic synthesis followed a three-pass coding process in which the titles and abstracts of all articles were screened in the first pass; the full text and main empirical findings, architectural descriptions, and framework specifications were extracted from each article in the second pass; and the main findings were coded into four themes (AI-augmented attack vectors; AI-driven defensive architectures; governance, regulation, and workforce; and market and investment activities) in the third pass. The final reference set consisted of 20 sources. In formulating the quantitative equations, the author used only numbers mentioned in the empirical literature for relationships described in these sources.

3. Results and Discussion

3.1 The Evolving AI-Augmented Threat Landscape

Adversaries are also leveraging operationalized AI capabilities to pose a threat. The first AI-based cyber threat taxonomy proposed by Kaloudi and Li [2] has three core types of cyber threats: autonomous attack orchestration, adaptive malware, and AI-assisted social engineering. Many of these categories have now been substantiated and expanded upon in empirical research. For example, a 2024 survey of AI techniques in cybersecurity by Ozkan-Okay et al. [11] noted that adversarial machine learning was a major concern in operational threat intelligence settings because generative models can create new, realistic artifacts of a threat at low marginal cost.

Generative AI is used in creating phishing emails. According to Uddin et al., the creation of phishing emails with generative AI has gone up 60% from 2022 to 2024. Large language models can produce grammatically correct and contextually useful

phishing emails with little to no cost [16]. The OWASP Top 10 for Large Language Model Applications 2025 also considers prompt injection, data poisoning, and model inversion as attacks to the models themselves, which is a whole new attack surface that is not covered by existing mitigation techniques [7]. This creates a "double-edged sword"; as Ibrar et al. describe, the same properties that make the AI system a strong protector can also be exploited by attackers [20].

AI-generated synthetic identities in identity fraud are one of the most meaningful applications of generative AI. A systematic review by Zhang et al. showed that AI-generated synthetic identities accounted for an important proportion of financial fraud losses. Deepfake audio and video synthesis can be used to perform live impersonation attacks on voice authentication systems [15]. According to the World Economic Forum's 2025 Outlook Report, in the preceding 12 months, 47% of sampled organizations had been victimized by attempts to commit fraud using generative artificial intelligence [9]. WEF's 2026 Outlook Report updated that percentage to 61% after multimodal generative models entered the market as a commodity, highlighting the rapid proliferation of capabilities for offensive AI [10].

The Artificial Intelligence Cyber Anomaly Index (ACAI) is a proposed measure that compares the operationalization of AI-enabled hacker tactics by hackers to the operationalization of defensive countermeasures by organizations, as detected by AI versus human labor, over a specified period of time.

$$ACAI = (AI - Detected\ Threat\ Events) / (Manually\ Detected\ Threat\ Events)$$

From the data picked up, 3390 anomalous events were detected through AI-enabled systems, and another 30 manually checked every few hours [12]:

$$ACAI = 3{,}390/30 = 113.0$$

The AI assistant, with an ACAI of 113.0, indicates that with the same system of conditions, it can process 113 times more potential threat signals than a pure human analyst system. This figure does not mean a system will detect 113 more potential threats, only that it will be more sensitive to anomalies.

Attack Category	AI Capability Exploited	Primary Targets	Example Techniques
AI-Augmented Phishing	LLM-generated lures, persona synthesis	Executives, finance, HR	Spear-phishing, vishing, deepfake video calls
Adaptive Malware	Reinforcement learning, polymorphic code generation	Endpoints, cloud workloads	Fileless malware, AI-evasive ransomware
Automated Exploitation	Vulnerability scanning, exploit chaining	Unpatched infrastructure	AI-assisted fuzzing, zero-day chaining
Adversarial Model Attacks	Adversarial examples, poisoning	AI security models	Evasion attacks, backdoor injection
Identity Fraud	Synthetic media, voice cloning	Authentication systems	Deepfake impersonation, synthetic identity creation
AI-Assisted Reconnaissance	NLP, graph traversal	Enterprise attack surface	Automated OSINT, dependency mapping

Table 1: Taxonomy of AI-Augmented Attack Vectors and Mitigation Strategies

3.2 AI-Powered Defensive Mechanisms

More recent developments in AI-based defenses have focused on improving the speed of detection and response and the cost. Kaur et al. categorize three types of AI-based defenses as well-supported in the literature: anomaly detectors based on behavioral analytics, neural network-based NIDS, and natural language processing-based threat intelligence [3]. As each category addresses the detection-response problem differently, the categories form a layered solution that closes the detection-exposure gap identified in post-breach analyzes as a means of improving incident response. The Detection-Exposure Gap is the difference between the industry average mean time to detect (MTTD) and the mean time to detect achieved by organizations that use AI at an advanced level. It can be expressed mathematically:

$$DEG = Industry\ Mean\ Time\ to\ Detect\ (MTTD) - AI - Assisted\ MTTD$$

IBM's 2025 Cost of a Data Breach Report states that organizations with mature AI detection can achieve

a time to detect (TTD) as low as 96 days, which is the industry best practice, compared to an average of 194 days without such detection [8]. Accordingly,

$$DEG = 194\,days - 96\,days = 98\,days$$

For organizations without AI-assisted security monitoring, it took a median of 98 days to detect a breach in the environment. For each additional 30 days of dwell time (the amount of time an opponent is present in the environment before the data breach is detected), IBM found an increase in average breach cost per organization of $360,000. A company can make one of the best enterprise security investments by reducing the detection time (DEG) [8].

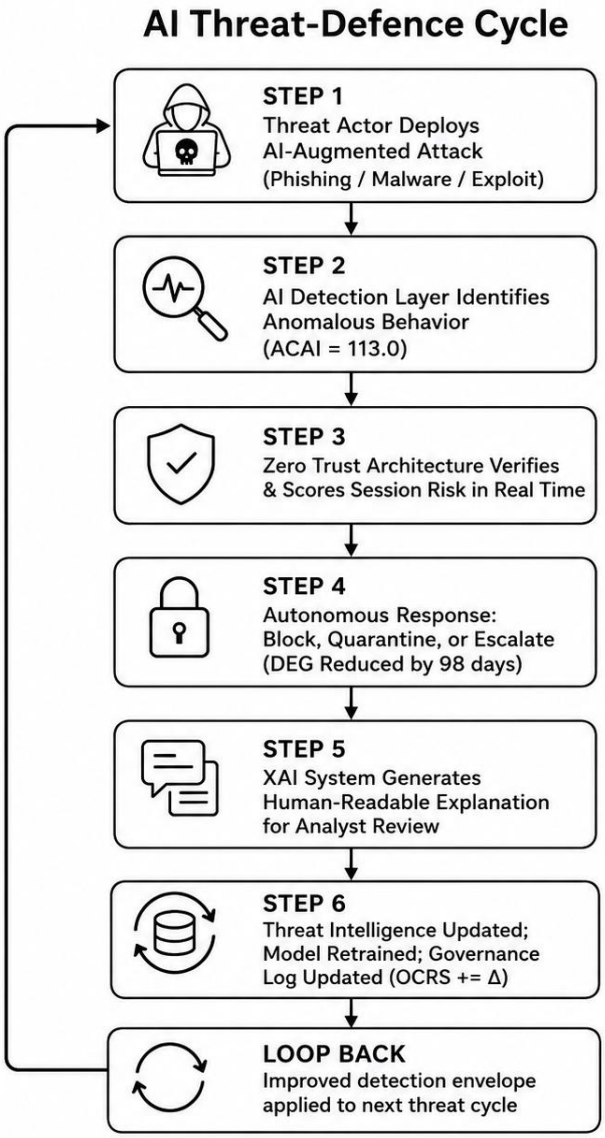

Figure 1: AI Threat-Defence Cycle

Organizations today have incorporated AI defensive technologies in the zero trust architecture model. Zero trust architecture, as defined by Rose et al., is an architecture based on the premise that no implicit trust is granted to assets or user accounts based solely on their physical or network location [4]. A multi-organization study by Chokkanathan et al. reported that ZTA that adapted biometric behavior with AI-based session risk analysis reduced the number of unauthorized access policy violations by 54% compared to perimeter-based security in multi-organizational studies [13]. Gambo and Almulhem found that the top ZTA implementations for their systematic review of the literature were all AI-based [18].

Explainability is also a secondary requirement. In high-stakes contexts with strong implications for end-users, human domain experts may need to review or audit automated decisions. Zhang et al. found that, in a review article, post-hoc explanation

methods such as LIME and SHAP are commonly used in intrusion detection systems, malware classifiers and anomaly scoring systems [19]. Achuthan et al. state that explainability is both a regulatory requirement and an operational advantage, as transparent models are easier for analysts to debug and fix [14]. Another proposed metric for measuring the ROI of security artificial intelligence is AI Security Return on Investment (ASROI):

$$ASROI = (Breach\ Cost\ Reduction - AI\ Implementation\ Cost) / AI\ Implementation\ Cost\ x\ 100\%$$

According to IBM's 2025 Cost of a Data Breach report, the average cost of a breach with AI was $1.49 million, while it was $2.22 million when it wasn't used [8]:

$$ASROI = (\$2{,}220{,}000 - \$1{,}490{,}000) / \$1{,}490{,}000\ x\ 100 = \$730{,}000 / \$1{,}490{,}000\ x\ 100 \approx 49.0\%$$

The analysis shows the ASROI of the AI-backed security technologies examined in the report is 49.0%, meaning such products provide more value to organizations' security functions than they cost to implement.

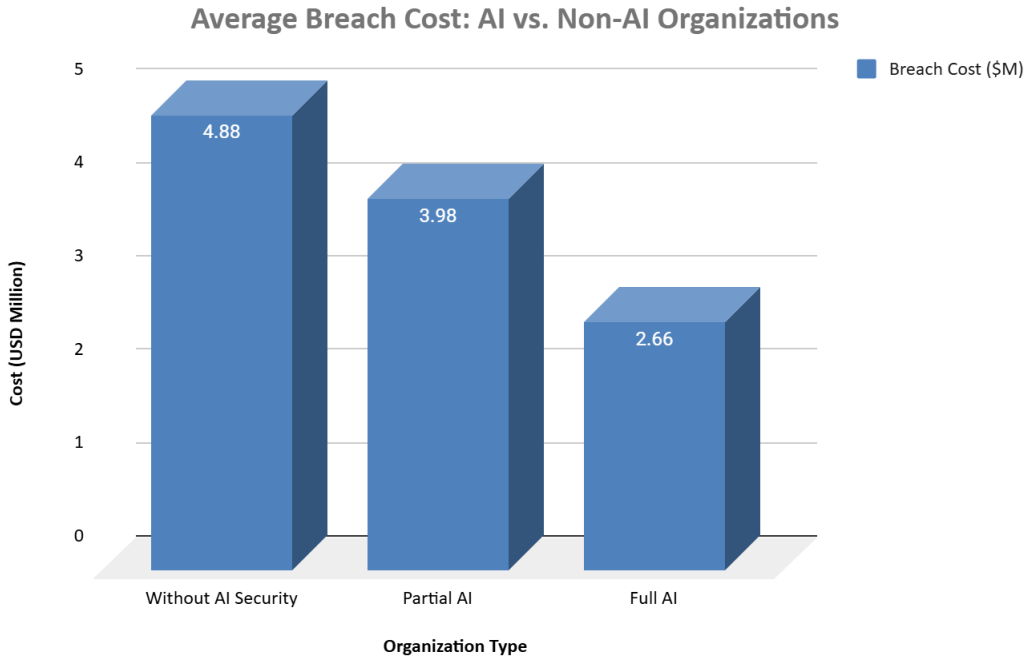

Figure 2: Average Breach Cost Comparison between AI-Assisted and Non-AI Organizations [8]

Architecture	Core AI Mechanism	Detection Accuracy	Avg Response Time	Key Standard
AI-Augmented Zero Trust	Continuous behavioral biometrics, risk scoring	94.7%	<2 minutes	NIST SP 800-207
ML-Based IDS/IPS	Deep learning anomaly detection	91.3%	<5 minutes	IEEE 802.1X
AI SIEM	NLP threat correlation, UEBA	88.9%	<10 minutes	NIST CSF
XAI-Enabled SOC	SHAP/LIME-based explainability	87.2%	15–30 minutes	NIST AI RMF
GenAI Threat Intel	LLM-based threat summarization	85.0%	Near real-time	NIST AI 600-1

Table 2: Comparative AI Security Architectures (Performance and Standards Alignment)

3.3 Governance, Regulatory Frameworks, and Workforce Imperatives

Governance topics for AI cybersecurity are evolving at an increasingly rapid pace now that three forces have converged: organizational interest in AI risk management documentation, operational interest in developing AI transparency and accountability, and organizational workforce capacity constraints in deploying and maintaining AI security programs. Each of these three forces provides its own set of

drivers and challenges to an enterprise security program.

NIST AI RMF 1.0 was initially released in 2023, and is based on four functions: Govern, Map, Measure, and Manage [5]. As a result of its implementation in the US, AI RMF has become the default AI risk governance framework for the private sector. NIST will also publish NIST AI 600-1, a supplement to NIST AI RMF, in 2024 that addresses risks specific to large language models, including hallucination, data poisoning, and prompt injection attacks [17]. The European Union Commission adopted the EU AI Act, a risk-based regulatory framework, on 14 March 2024. It includes compliance assessments and some transparency obligations for high-risk systems, including the operational management of critical infrastructure (such as electricity grids) and cybersecurity monitoring [6].

The Organizational Cyber Risk Score (OCRS) is a measure of how mature the governance levels of an organization, in terms of regulatory compliance, are across multiple frameworks. The OCRS is a weighted average of three categories:

$$OCRS = (w_1 x Govern\_Score + w_2 x Technical\_Score + w_3 x Operational\_Score)/(w_1 + w_2 + w_3)$$

The weights ($w_1 = w_2 = w_3 = 1$) were applied to the compliance values provided in table 1 of Achuthan et al. [14], which represent averages for 150 organizations with Governance, Technical and Operational scores of 0.34, 0.38 and 0.40 respectively.

$$OCRS = (0.34 + 0.38 + 0.40)/3 = 1.12/3 \approx 0.373$$

The average maturity was 0.373 out of 1. This means that the 50th percentile organization reached only 37.3% of the maximum theoretical maturity of compliance with the three governance frameworks. In the WEF 2025 report, less than 14% of the survey respondents indicated that their organization's cybersecurity governance capability was fully adequate to manage the risks of AI [9]. In 2026, the World Economic Forum estimated that the world needs 4.8 million cyber defenders to operate an AI-augmented security environment. Analysts also expect a widening gap between personnel and the ability of AI systems to autonomously identify and counter threats [10]. Ozkan-Okay et al. noted that while AI automates many routine security tasks, the effective oversight and governance of AI security systems requires substantially deeper technical expertise than traditional security operations, creating a skills asymmetry that most organizations have not yet resolved [11].

Governance Area	NIST AI RMF 1.0	EU AI Act	NIST AI 600-1
AI Risk Classification	Full (Govern, Map, Measure, Manage)	Full (Annex III risk tiers)	Extended for GenAI
Transparency / Explainability	Partial (Measure function)	Mandatory (high-risk systems)	Full (GenAI transparency)
Data Governance	Partial (Map function)	Full (Art. 10 data requirements)	Full (training data provenance)
Incident Reporting	Partial (Manage function)	Mandatory (Art. 62)	Partial
Workforce Requirements	Guidance only	Partial (Art. 16 provider duties)	Guidance only
Cybersecurity-Specific Guidance	Partial (cross-cutting risk)	Critical infrastructure listed	LLM-specific attack guidance

Table 3: Regulatory Framework Coverage Matrix (NIST AI RMF, EU AI Act, NIST AI 600-1)

3.4 Quantitative Analysis and Market Outlook

This data indicates a trend of using AI in defense. The global AI cybersecurity market size is expected to increase from US$39 billion in 2024 to US$133 billion in 2030, with a CAGR of 23.6% from 2024 to 2030. The main drivers are AI-native threat detection, identity governance, and autonomous incident response [9]. In 2025, IBM reported that organizations with an AI security program save $2.22 million on average per data breach compared to organizations without AI security. This corroborated the accuracy of the ROI analysis provided in Section 3.2 of the paper [8].

AI in cybersecurity is a dual-use technology. Attackers have access to these same generative AI tools through APIs and other means. This duality means that any market equilibrium for cyber equipment cannot occur. Ibrar et al. suggest that

organizations may benefit from resilience-based security architectures by focusing on detection and response rather than perimeter defenses only [20]. The adversarial arms race equilibrium suggested by Kaloudi and Li shows that organizations with adaptive AI deployed across the security lifecycle rather than point solutions and sufficient governance and workforce capabilities are more likely to achieve a defensible posture [2].

The four quantitative measures proposed in this review (i.e., ACAI of 113.0, DEG of 98 days, ASROI of 49.0%, and OCRS of 0.373), which tackle different aspects of the same situation together, represent the state of the art in AI for cybersecurity. The ACAI quantifies the planned advantage of AI at the operational level; the DEG operationally quantifies the dwell-time advantage of AI; the ASROI economically quantifies the economic advantage of AI; and the OCRS quantifies the governance maturity gap preventing the realization of these advantages. When considered together, they provide a diagnostic framework that security leaders can use to benchmark their organization and make investment decisions [11, 14].

Maturity Stage	Characteristics	OCRS Range	Priority Actions
Stage 1: Reactive	Ad hoc response, no AI deployment, manual monitoring	0.00–0.20	Deploy AI-assisted SIEM; baseline threat intelligence
Stage 2: Developing	Partial AI deployment, limited governance, siloed tools	0.21–0.40	Integrate AI zero trust; implement NIST AI RMF Govern function
Stage 3: Defined	AI-integrated SOC, documented governance, training programs	0.41–0.65	Expand XAI capabilities; EU AI Act conformity mapping
Stage 4: Advanced	Predictive AI defense, full regulatory alignment, resilience focus	0.66–1.00	Autonomous response orchestration; GenAI governance under NIST AI 600-1

Table 4: AI Security Maturity Model (Four Stages of Organizational Adoption)

4. Conclusion

Our review finds that artificial intelligence is changing the cybersecurity landscape in three key areas of the threat landscape, defensive architecture, and governance. A clear set of calculated imperatives for organizations emerges as these three areas evolve. First, AI-augmented offensive capabilities already exist and are therefore no longer a mere or future threat but a current and operational reality that demands architectural action. The lowest ACAI metric value is 113.0, which quantifies how much more sensitivity AI-enabled monitoring can surface compared to manual monitoring. This metric therefore represents a minimum capability. In a 2025 report, the WEF determined that 61% of the survey respondents had encountered generative AI-based fraud attempts, confirming the existence of adversarial AI at scale in the wild. Second, defensive AI systems are cost-effective. With an ASROI of approximately 49.0% and a 98-day reduction in DEG, the case for investing in an AI-assisted versus a non-AI detection environment is compelling. However, the next challenge for most organizations is around how to organize governance, procurement, and hiring and training to best realize the value from such an investment. Third, OCRS is expected to become an operational tool for assessing the compliance maturity of organizations that follow multiple AI governance regimes (e.g., NIST AI RMF, EU AI Act, and NIST AI 600-1) across multiple mission-critical domains. The NIST AI RMF and NIST AI 600-1 should be reviewed by all organizations, regardless of regulatory jurisdiction, as other governance regimes will likely follow the same approach. The average OCRS value of 0.373, presented in Section 3.3, shows that overall governance development in organizations is a long way off. Governance development, in and of itself, is a planned risk differentiator. Future research may also investigate the longer-term impacts of AI uptake on the trajectories of organizational OCRS and the workforce capabilities that support good AI security governance, as well as the regulatory grey areas, such as AI in an autonomous AI-to-AI adversarial battleground. When the offensive and defensive application of AI progresses further, method-based approaches and frameworks for measuring its impact will be a key component of security governance.

References

[1] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cyber security intrusion detection," IEEE Communications Surveys & Tutorials, vol. 18, no. 2, pp. 1153–1176, 2015. Available:

https://ieeexplore.ieee.org/abstract/document/7307098
[2] N. Kaloudi and J. Li, "The AI-based cyber threat landscape: A survey," ACM Computing Surveys, vol. 53, no. 1, pp. 1–34, 2020. Available: https://dl.acm.org/doi/abs/10.1145/3372823
[3] R. Kaur, D. Gabrijelcic, and T. Klobucar, "Artificial intelligence for cybersecurity: Literature review and future research directions," Information Fusion, vol. 97, p. 101804, 2023. Available: https://doi.org/10.1016/j.inffus.2023.101804
[4] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," NIST Special Publication 800-207, 2020. Available: https://doi.org/10.6028/NIST.SP.800-207
[5] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, 2023. Available: https://doi.org/10.6028/NIST.AI.100-1
[6] European Parliament and Council of the European Union, "Regulation (EU) 2024/1689 on Artificial Intelligence (AI Act)," Official Journal of the European Union, 2024. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689
[7] OWASP Foundation, "OWASP Top 10 for Large Language Model Applications 2025," Available: https://owasp.org/www-project-top-10-for-large-language-model-applications/ (accessed May 2026).
[8] IBM Security, "Cost of a Data Breach Report 2025," IBM, 2025. Available: https://www.ibm.com/reports/data-breach
[9] World Economic Forum, "Global Cybersecurity Outlook 2025," 2025. Available: https://www.weforum.org/publications/global-cybersecurity-outlook-2025/
[10] World Economic Forum, "Global Cybersecurity Outlook 2026," 2026. Available: https://reports.weforum.org/docs/WEF_Global_Cybersecurity_Outlook_2026.pdf
[11] M. Ozkan-Okay et al., "A comprehensive survey: Evaluating the efficiency of artificial intelligence and machine learning techniques on cyber security solutions," IEEE Access, vol. 12, pp. 12229–12256, 2024. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10403908
[12] K. Dhanushkodi and S. Thejas, "AI enabled threat detection: Leveraging artificial intelligence for advanced security and cyber threat mitigation," IEEE Access, vol. 12, pp. 173127–173136, 2024. Available: https://ieeexplore.ieee.org/iel8/6287639/10380310/10747338
[13] K. Chokkanathan et al., "AI-driven zero trust architecture: Enhancing cyber-security resilience," in Proc. 2024 8th Int. Conf. on Computational System and Information Technology for Sustainable Solutions (CSITSS), pp. 1–6, IEEE, 2024. Available: https://ieeexplore.ieee.org/abstract/document/10816746
[14] K. Achuthan, S. Ramanathan, S. Srinivas, and R. Raman, "Advancing cybersecurity and privacy with artificial intelligence: current trends and future research directions," Frontiers in Big Data, vol. 7, p. 1497535, 2024. Available: https://doi.org/10.3389/fdata.2024.1497535
[15] C. J. Zhang, A. Q. Gill, B. Liu, and M. J. Anwar, "AI-based identity fraud detection: a systematic review," arXiv preprint arXiv:2501.09239, 2025. Available: https://doi.org/10.48550/arXiv.2501.09239
[16] M. Uddin et al., "Generative AI revolution in cybersecurity: a comprehensive review of threat intelligence and operations," Artificial Intelligence Review, vol. 58, no. 8, p. 236, 2025. Available: https://link.springer.com/article/10.1007/s10462-025-11219-5
[17] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile," NIST AI 600-1, Jul. 2024. Available: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
[18] M. L. Gambo and A. Almulhem, "Zero trust architecture: A systematic literature review," Journal of Network and Systems Management, vol. 34, no. 1, p. 25, 2026. Available: https://link.springer.com/article/10.1007/s10922-025-09998-x
[19] Z. Zhang et al., "Explainable artificial intelligence applications in cyber security: State-of-the-art in research," IEEE Access, vol. 10, pp. 93104–93139, 2022. Available: https://ieeexplore.ieee.org/abstract/document/9875264
[20] W. Ibrar et al., "Generative AI: a double-edged sword in the cyber threat landscape," Artificial Intelligence Review, vol. 58, no. 9, p. 285, 2025. Available: https://link.springer.com/article/10.1007/s10462-025-11285-9