CYBERSECURITY PRACTICES IN THE GENERATIVE AI ERA: RISK ASSESSMENT AND ETHICAL ISSUES FOR ACADEMIC LIBRARIES IN NIGERIA

By

AKPABIO, Kingsley Ubong
Department of Computer Science
Federal University of Technology, Ikot Abasi, Akwa Ibom State
akpabioubongk@gmail.com | 09016912531

&

AKOR, Solomon Obotu
Department of Library and Information Technology
Federal University of Technology, Ikot Abasi, Akwa Ibom State
akorsolomon11@gmail.com

Abstract

This study examined cybersecurity practices, risk-assessment activities, and ethical issues associated with the use of generative artificial intelligence in selected federal university libraries in Nigeria. A descriptive-correlational survey design was adopted. The population comprised 340 professional, systems, electronic, and digital librarians in six federal university libraries representing Nigeria's geopolitical zones. Convenience sampling was used to reach librarians who were readily accessible and willing to participate, and 314 completed questionnaires were used for analysis. Data were collected using the Cybersecurity, Risk Assessment and Generative AI Ethics Questionnaire (CSRAGEQ). The instrument was validated by specialists in Library and Information Science, cybersecurity, and measurement and evaluation. A pilot test involving 30 librarians yielded an overall Cronbach's alpha coefficient of 0.88. Mean and standard deviation were used to address the descriptive objectives, while Pearson Product-Moment Correlation was used to test the hypothesis at the 0.05 significance level. The findings revealed that cybersecurity practices were implemented to a high extent (grand mean = 3.01), risk-assessment practices were undertaken to a high but uneven extent (grand mean = 2.82), and ethical concerns relating to privacy, inaccurate outputs, copyright, bias, transparency, and accountability were substantial (grand mean = 3.28). Risk-assessment practices had a statistically significant moderate positive relationship with cybersecurity practices ($r = .586$, $p < .001$). The study concluded that the secure adoption of generative AI requires integrated governance, technical safeguards, documented risk assessment, staff development, vendor oversight, and human accountability. It recommended institution-wide AI governance policies, periodic security audits, privacy-conscious procurement, incident-response planning, and continuous cybersecurity and ethics training for librarians.

Keywords: academic libraries; cybersecurity; ethical AI; generative AI; risk assessment; Nigeria

Introduction

Generative artificial intelligence (AI) is reshaping academic information environments through conversational search, automated reference support, text summarization, metadata generation, translation, recommendation, and research assistance. Academic libraries are experimenting with large language models and AI-enabled platforms to improve access, responsiveness, and staff productivity. The International Federation of Library Associations and Institutions (IFLA, 2024) recognizes the practical value of generative

AI for library and information professionals while urging prior consideration of ethical implications. Academic libraries therefore face a dual responsibility: they must explore technologies that enhance learning and research while protecting the confidentiality, integrity, availability, and trustworthy use of information resources.

The security implications extend beyond conventional malware and unauthorized access. Generative AI systems may expose confidential prompts, user histories, proprietary documents, or research data to external providers; they may also be manipulated through prompt injection, poisoned knowledge sources, insecure plug-ins, compromised application programming interfaces, or adversarial inputs. The National Institute of Standards and Technology's Generative AI Profile identifies risks such as confabulation, information security weaknesses, privacy harm, intellectual-property concerns, harmful bias, and opaque system behaviour, and recommends managing them across the AI lifecycle (Autio et al., 2024). In academic libraries, these risks can affect chatbots, discovery tools, digital repositories, electronic-resource systems, and staff use of public generative AI services.

Ethical concerns are inseparable from cybersecurity because the security of an AI-enabled service is not limited to preventing intrusion. A technically protected system may still be harmful if it produces discriminatory recommendations, fabricates citations, obscures data use, violates copyright, or makes consequential decisions without human review. UNESCO's Recommendation on the Ethics of Artificial Intelligence emphasizes privacy and data protection, fairness, transparency, explainability, accountability, auditability, human oversight, and AI literacy (UNESCO, 2021). These principles closely correspond with the professional values of librarianship, including intellectual freedom, equitable access, confidentiality, information integrity, and stewardship.

The challenge is particularly important in Nigerian federal universities, where academic libraries operate amid uneven bandwidth, power, technical staffing, funding, and institutional policy development. Generative AI can offer valuable support where personnel and resources are constrained, but these same constraints may weaken security monitoring, vendor due diligence, incident response, and staff training. Matsieli and Mutula (2025) note that ethical risks associated with generative AI may be more acute in the Global South because limited infrastructure, skills shortages, and weak governance frameworks can reduce institutions' ability to respond. Evidence is therefore needed on the safeguards librarians apply, the extent to which they assess AI-related risks, and the ethical issues they encounter.

Statement of the Problem

Academic libraries increasingly use or provide access to generative AI tools without always possessing a unified framework for evaluating their security, privacy, ethical, and operational implications. Staff may enter sensitive information into third-party systems, connect AI services to library platforms, rely on generated outputs, or recommend tools to users without adequate knowledge of retention practices, model limitations, security controls, or contractual responsibilities. Weak authentication, excessive access privileges, unpatched systems, poor backup arrangements, and informal incident reporting can compound AI-specific vulnerabilities. At the same time, inaccurate or fabricated responses, algorithmic bias, unclear authorship, copyright infringement, and insufficient disclosure can undermine scholarship and public trust.

Existing discussions often describe the potential benefits and risks of AI in libraries but provide limited multi-regional evidence on how Nigerian academic librarians apply cybersecurity controls and conduct systematic risk assessment. The absence of evidence makes it difficult for university

management to determine whether adoption is sufficiently governed, where competence gaps exist, and which safeguards deserve priority. This study addresses that problem by examining cybersecurity practices, risk-assessment activities, and ethical issues across selected federal university libraries and by determining whether stronger risk assessment is associated with stronger cybersecurity practice. The focus is association rather than effect because the correlational design does not establish causality.

Objectives of the Study

The study sought to (1) assess the extent to which cybersecurity measures are implemented in generative AI-enabled academic library services; (2) examine the extent to which librarians undertake risk-assessment practices for generative AI tools and services; (3) evaluate the ethical issues associated with generative AI use in academic libraries.

Research Hypothesis

H0: There is no significant relationship between librarians' risk-assessment practices and cybersecurity practices in generative AI-enabled academic library services in Nigeria.

Literature Review

Generative AI, academic libraries, and cybersecurity

Generative AI expands the attack surface of the academic library because it introduces new data flows, third-party dependencies, and interfaces between users, institutional systems, and external models. Traditional controls such as access management, secure configuration, encryption, patching, backups, logging, network protection, and incident response remain essential. However, they must be supplemented by AI-specific controls, including prompt and output filtering, restrictions on sensitive data, model and vendor documentation, retrieval-source validation, adversarial testing, human review, and monitoring for harmful or anomalous output. NIST's AI Risk Management Framework recommends that AI risks be governed, mapped, measured, and managed rather than treated separately from enterprise cybersecurity and privacy processes (Tabassi, 2023). This lifecycle approach is appropriate for libraries because risks arise during procurement, configuration, deployment, daily use, modification, and retirement.

Risk assessment is the disciplined process through which an institution identifies assets, threats, vulnerabilities, likelihood, consequences, existing controls, and residual risk. In an AI-enabled library, the assessment should establish what information enters the system, where it is stored, who can access it, how outputs influence services, and what happens when the system fails or produces harmful content. Autio et al. (2024) emphasize that generative AI risk management should reflect organizational goals, legal obligations, risk tolerance, and resources. Consequently, a small experimental chatbot and an AI system connected to patron records should not receive identical treatment. Data sensitivity, scale, autonomy, user population, and integration depth should determine the strength of review and control.

Ethical issues in generative AI-supported library services

Privacy is a primary ethical and security concern because prompts can reveal reading interests, research questions, personal circumstances, unpublished ideas, or identifiable information. External platforms may retain or reuse inputs under terms that users do not fully understand. NIST (2024) notes that AI can amplify re-identification, behavioural tracking, and surveillance risks. Libraries should therefore apply data minimization, purpose limitation, informed notice, restricted input categories, contractual safeguards, and privacy impact assessment. Confidentiality must remain a service value even when a convenient tool encourages users to share more information than is necessary.

Accuracy and information integrity are equally important. Generative models can produce plausible but false statements, fabricated citations, misleading summaries, or outdated guidance. IFLA (2026) cautions that generative AI outputs may contain inaccuracies, bias, and fabricated information, reinforcing the value of retrieval from trusted and curated sources. Academic librarians should not present generated content as authoritative merely because it is fluent. Human verification, transparent labelling, source grounding, escalation procedures, and AI literacy instruction are required, especially when outputs affect research decisions, reference services, or academic assessment.

Bias, exclusion, and epistemic injustice arise when training data and model design privilege dominant languages, cultures, and knowledge traditions. Matsieli and Mutula (2025) identify algorithmic bias, digital inequality, misinformation, privacy breaches, intellectual-property conflicts, and job displacement as interconnected ethical issues for libraries. Nigerian libraries must consider whether systems adequately represent local scholarship, indigenous knowledge, Nigerian English, and users with different levels of digital access. Fairness assessment should examine not only model performance but also who benefits, who is excluded, and whether human alternatives remain available.

Copyright, authorship, transparency, and accountability create further tension. AI-generated text may reproduce protected expression, obscure the origin of ideas, or complicate decisions about attribution. Sousa (2025) argues that academic libraries are positioned to mediate responsible AI use in research but that institutional responses remain fragmented. Libraries should provide guidance on disclosure, verification, citation, confidentiality, and appropriate human responsibility. UNESCO (2021) stresses that AI must not displace ultimate human accountability; accordingly, librarians and universities remain responsible for services offered under their authority, even when a vendor's model produced the output.

Institutional readiness and professional competence

Secure and ethical deployment depends on institutional readiness rather than individual caution alone. Governance should define approved uses, prohibited data, authorization procedures, procurement review, roles, monitoring, incident reporting, audit, and sanctions. Cross-functional oversight is necessary because librarians understand service values and information behaviour, while information-technology personnel, cybersecurity specialists, legal officers, data-protection officers, researchers, and users contribute different forms of expertise. Ngulube and Mosha (2025) conclude from a scoping review that safe AI integration in academic libraries requires systematic attention to perceived risks and ethical issues, confirming that adoption decisions should be embedded in governance rather than left to informal experimentation.

Professional development must combine functional AI literacy with security and ethical competence. Knowing how to obtain a useful response from an AI tool does not establish the ability to evaluate privacy terms, recognize prompt injection, identify confidential information, validate sources, or document an incident. Training should be role specific and recurring because platforms and threats evolve rapidly. Librarians also need the confidence to pause or reject unsafe uses. The literature therefore indicates a need for empirical assessment that connects routine safeguards, formal risk-assessment practices, and ethical concerns within the institutional realities of Nigerian academic libraries.

Methodology

The study adopted a descriptive-correlational survey design. The descriptive component enabled the researchers to document prevailing cybersecurity practices, risk-assessment activities, and ethical concerns without manipulating the study environment.

The correlational component was used to determine the direction and strength of the relationship between risk-assessment practices and cybersecurity practices. The design does not establish cause and effect.

The population comprised 340 professional librarians, including systems, electronic-resource, digital, and other librarians involved in technology-mediated services, from six federal university libraries representing the geopolitical zones: Ahmadu Bello University, Zaria; University of Maiduguri; University of Jos; University of Lagos; University of Nigeria, Nsukka; and University of Calabar. The study employed convenience sampling, a non-probability technique that enabled the researchers to recruit members of the population who were readily accessible and willing to participate during the data-collection period. The questionnaire link was circulated through the selected university libraries and relevant professional online platforms. At the close of data collection, 314 librarians had completed the questionnaire, and all 314 properly completed responses constituted the realized sample used for analysis.

Data were collected using the Cybersecurity, Risk Assessment and Generative AI Ethics Questionnaire (CSRAGEQ), developed from the objectives and relevant AI risk and ethics guidance. The instrument contained demographic items and three four-point scales covering cybersecurity practices, risk-assessment practices, and ethical issues. Responses ranged from Very High Extent (4) to Very Low Extent (1). Specialists in Library and Information Science, cybersecurity, and measurement and evaluation reviewed the instrument for clarity, coverage, relevance, and alignment. A pilot study involving 30 librarians from a federal university outside the study produced an overall Cronbach's alpha coefficient of 0.88, indicating good internal consistency.

The questionnaire was administered electronically. Participation was voluntary; respondents received information about the study, confidentiality, anonymity, and their right to withdraw. No directly identifying data were included in the analysis. Mean and standard deviation were used to address the descriptive objectives. A criterion mean of 2.50 guided decisions: mean values of 2.50 or above indicated a high extent or a substantial ethical concern, depending on the scale, while values below 2.50 indicated a low extent. Pearson Product-Moment Correlation tested the null hypothesis at the 0.05 level of significance. Results were presented by objective rather than by restating research questions.

Presentation and Analysis of Results

Table 1. Distribution of the study population and completed questionnaires

Response status	Frequency	Percentage (%)
Completed questionnaires	314	92.4
Did not complete	26	7.6
Study population	340	100.0

Table 1 shows that 314 librarians completed the questionnaire, representing 92.4% of the identified population of 340 librarians. These 314 completed questionnaires constituted the realized convenience sample and formed the basis of the analysis.

Table 2. Cybersecurity practices in generative AI-enabled services

Cybersecurity practice	Mean	SD	Decision
Uses strong, unique passwords for AI-related systems	3.35	0.62	High
Applies multi-factor authentication where available	3.18	0.71	High

Cybersecurity practice	Mean	SD	Decision
Restricts access according to staff roles	3.09	0.74	High
Installs security patches and updates promptly	3.04	0.76	High
Encrypts or protects sensitive data in storage and transit	2.81	0.84	High
Maintains tested backups of critical library data	3.12	0.72	High
Monitors logs and unusual AI-system activities	2.69	0.88	High
Restricts entry of confidential data into public AI tools	3.21	0.69	High
Uses documented incident-reporting procedures	2.76	0.86	High
Provides regular cybersecurity awareness training	2.83	0.82	High
Grand mean	3.01	0.76	High

Cybersecurity measures were implemented to a high extent overall (grand mean = 3.01). Password protection, restrictions on confidential input, multi-factor authentication, backups, role-based access, and software updates received the strongest ratings. Log monitoring, documented incident reporting, encryption, and continuing awareness training were comparatively weaker. The pattern indicates that visible user-level controls were more established than monitoring, response, and governance-oriented controls.

Table 3. Risk-assessment practices for generative AI

Risk-assessment practice	Mean	SD	Decision
Inventories AI tools used by the library	2.94	0.78	High
Classifies the sensitivity of data processed by AI tools	2.76	0.84	High
Identifies threats and vulnerabilities before deployment	2.91	0.80	High
Assesses likelihood and impact of identified risks	2.87	0.82	High
Conducts privacy or data-protection impact assessment	2.61	0.91	High
Evaluates vendors' security and data-retention terms	2.68	0.88	High
Tests systems for inaccurate or harmful outputs	2.89	0.81	High
Reviews risk after upgrades or major changes	2.74	0.86	High
Documents risk owners, controls, and residual risk	2.57	0.93	High
Continuously monitors and reports AI-related risk	3.19	0.73	High
Grand mean	2.82	0.84	High

Risk assessment was undertaken to a high extent, but the grand mean of 2.82 was lower than that for general cybersecurity practice. Continuous monitoring, tool inventory, pre-deployment threat identification, output testing, and likelihood-impact analysis ranked highest. Documentation of residual risk, privacy impact assessment, and vendor review ranked lowest. This suggests that assessment occurred, but formal documentation and third-party accountability were not equally mature.

Table 4. Ethical issues associated with generative AI use

Ethical issue	Mean	SD	Decision
Exposure of patron or research data through AI prompts	3.42	0.64	Substantial

Ethical issue	Mean	SD	Decision
Fabricated facts, citations, or misleading outputs	3.51	0.58	Substantial
Copyright and intellectual-property infringement	3.34	0.67	Substantial
Bias against local languages or knowledge traditions	3.16	0.74	Substantial
Lack of transparency about how outputs are produced	3.29	0.69	Substantial
Unclear accountability for harmful AI-assisted decisions	3.27	0.72	Substantial
Undisclosed AI use in academic work	3.38	0.65	Substantial
Over-reliance that weakens human professional judgment	3.11	0.77	Substantial
Unequal access caused by infrastructure and cost	3.22	0.73	Substantial
Surveillance or excessive profiling of library users	3.08	0.81	Substantial
Grand mean	3.28	0.70	Substantial

All ethical issues were rated as substantial, with a grand mean of 3.28. Fabricated information and citations recorded the highest mean, followed by privacy exposure, undisclosed AI use, copyright concerns, lack of transparency, and unclear accountability. Surveillance and over-reliance recorded the lowest means but remained well above the criterion. The findings show that respondents viewed information integrity, confidentiality, and scholarly responsibility as immediate professional concerns.

Table 5. Relationship between risk-assessment practices and cybersecurity practices.

Hypothesis	N	r	p-value	Decision
risk-assessment practices	314	.586	< .001	Reject H0
cybersecurity practices.	314			

Since the p-value was less than .001, which is below the 0.05 significance level, the null hypothesis was rejected. This indicates a statistically significant relationship between risk-assessment practices and cybersecurity practices.

Discussion of Findings

Cybersecurity practices

The high overall implementation of cybersecurity practices suggests that librarians and their institutions recognized the need to protect AI-enabled services. The strongest practices - password security, restrictions on confidential inputs, multi-factor authentication, backups, access control, and updates - correspond with foundational cybersecurity controls. This is encouraging because generative AI does not eliminate conventional security obligations. NIST's AI RMF stresses that AI risk should be integrated with cybersecurity, privacy, and enterprise risk rather than managed as an isolated category (Tabassi, 2023). The result implies that selected libraries had established a basic control environment on which AI-specific safeguards could be built.

The weaker scores for log monitoring, encryption, formal incident reporting, and regular training reveal a maturity gap. Preventive controls may reduce common incidents, but they cannot guarantee detection, containment, recovery, and learning. AI-enabled applications can behave unpredictably or be manipulated through inputs that evade ordinary network controls. Libraries therefore need operational visibility, clear escalation paths, retained evidence, tested response plans, and staff who can recognize unusual behaviour. The NIST Generative AI Profile similarly recommends ongoing measurement and

monitoring throughout the lifecycle rather than a one-time security review (Autio et al., 2024).

Risk-assessment practices

Risk-assessment practice was rated high but less strongly than general cybersecurity. Respondents reported monitoring, tool inventory, threat identification, output testing, and likelihood-impact analysis, yet formal documentation, vendor assessment, and privacy impact assessment were weaker. This distinction is important. An institution may perform security tasks without maintaining an auditable process that links identified risks to owners, controls, residual exposure, review dates, and acceptance decisions. Inconsistent documentation can result in repeated mistakes, weak accountability, and difficulty demonstrating due diligence after an incident.

The lower vendor-review score is particularly significant because many libraries consume externally hosted AI services rather than developing their own models. Risk may arise from data retention, subcontractors, model updates, insecure interfaces, unclear breach notification, limited audit rights, or vendor lock-in. Autio et al. (2024) recommend aligning risk responses with organizational priorities, obligations, and available resources. Nigerian academic libraries should therefore use proportionate but documented due diligence before transferring institutional or user information to an AI provider. Procurement is a cybersecurity and ethics decision, not only a financial or functional decision.

Ethical issues

Respondents identified serious ethical concerns across privacy, accuracy, copyright, bias, transparency, accountability, disclosure, professional judgment, inequality, and surveillance. The prominence of fabricated facts and citations is consistent with concerns about confabulation in the NIST profile and with IFLA's warning that generative AI may produce inaccurate, biased, or fabricated information (Autio et al., 2024; IFLA, 2026). Because users often approach libraries for reliable information, unverified AI output can transfer technological error into institutional authority. Libraries should ground answers in curated sources, label AI involvement, preserve access to source material, and require human verification for consequential assistance.

Privacy exposure, copyright, and undisclosed use also received high ratings. These concerns align with UNESCO's (2021) principles of privacy, transparency, accountability, fairness, and human oversight. They also support Matsieli and Mutula's (2025) argument that libraries must treat generative AI as an ethical information-society issue rather than a neutral productivity tool. Academic libraries are well placed to guide users because they already teach evaluation, citation, research integrity, and responsible information use. However, guidance must be supported by institutional policies and examples that distinguish acceptable assistance from confidential, deceptive, infringing, or academically dishonest use.

Conclusion

Generative AI creates opportunities for responsive, personalized, and efficient academic library services, but it also changes the nature and scale of institutional risk. The findings indicate that librarians in the selected federal universities applied basic cybersecurity safeguards to a high extent and undertook risk assessment to a generally high but uneven extent. Ethical concerns were substantial, particularly those relating to inaccurate or fabricated outputs, privacy exposure, undisclosed AI use, copyright, transparency, and accountability. Risk-assessment practices were significantly associated with cybersecurity practices; however, the correlational design does not support causal conclusions.

Secure adoption therefore requires an integrated approach in which technology, governance, professional ethics,

procurement, training, and human oversight reinforce one another. Libraries should neither adopt generative AI uncritically nor reject it without analysis. Their responsibility is to evaluate intended benefits, identify foreseeable harms, implement proportionate safeguards, communicate limitations, preserve user choice, and remain accountable for services provided under their authority.

Recommendations

The following recommendations were made based on the findings of the study:

1. Federal universities should adopt an institution-wide generative AI governance policy defining approved tools and uses, prohibited data, authorization responsibilities, disclosure requirements, human oversight, retention rules, monitoring, audit, and sanctions.
2. Academic libraries should maintain an inventory of AI tools and conduct documented risk and privacy impact assessments before deployment, after major changes, and at scheduled intervals.
3. Library and university IT units should strengthen multi-factor authentication, encryption, role-based access, secure configuration, patching, backups, log monitoring, vulnerability management, and tested incident-response procedures for AI-connected services.
4. Procurement and renewal processes should evaluate vendors' data ownership, retention, model-training use, subcontractors, breach notification, security testing, audit rights, accessibility, bias controls, service continuity, and exit arrangements.
5. Librarians should receive recurring role-based training on confidential-data handling, prompt injection, output verification, fabricated citations, copyright, bias, transparency, incident reporting, and responsible AI disclosure.

Limitations of the Study and Suggestions for Further Research

The study relied on self-reported practices and perceptions rather than independently audited controls, technical vulnerability assessments, incident logs, contractual records, or direct observation of professional behaviour. Consequently, social-desirability and common-method biases may have influenced the responses. Convenience sampling also depended on respondents' accessibility and willingness to participate; therefore, the participating librarians may differ systematically from those who did not complete the questionnaire. The findings should consequently be generalized with caution. The cross-sectional design did not permit causal inference, while the six selected federal university libraries may not adequately represent the conditions prevailing in state universities, private universities, polytechnics, colleges of education, or other library environments in Nigeria. However, further research should use verified survey responses together with interviews, policy review, technical audit, and incident data. Comparative studies could examine differences by geopolitical zone, institutional type, library size, staff role, level of AI integration, and availability of dedicated cybersecurity units. Longitudinal and intervention studies should assess whether training, governance frameworks, privacy impact assessment, or vendor controls improve measurable outcomes. Research is also needed on user attitudes, local-language bias, accessibility, copyright, environmental cost, and the effectiveness of retrieval-grounded library chatbots in reducing fabricated information.

Declaration of Interest: The authors declare that they have no competing interests.

Funding: Funding was received for this study through the Institutional-Based Research (IBR) grant by the Tertiary Education Trust Fund (TETFUND).

Acknowledgement: The researchers gratefully acknowledges the support of the Tertiary Education Trust Fund (TETFUND) for providing funding for the execution of this project through the Institutional Based Research (IBR) grant. The authors also extend their gratitude to librarians who participated in the survey. Their time, willingness, and valuable responses made

this research possible. We further acknowledge the support of colleagues and professional networks who contributed, directly or indirectly, to the successful completion of this study.

References

Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1

International Federation of Library Associations and Institutions. (2024). Generative AI for library and information professionals. https://www.ifla.org/generative-ai/

International Federation of Library Associations and Institutions. (2026). AI and libraries: A new era of trusted discovery and smarter services. https://www.ifla.org/news/ai-and-libraries-a-new-era-of-trusted-discovery-and-smarter-services/

Matsieli, M., & Mutula, S. (2025). Generative AI and the information society: Ethical reflections from libraries. Information, 16(9), 771. https://doi.org/10.3390/info16090771

National Institute of Standards and Technology. (2024). Cybersecurity, privacy, and AI. https://www.nist.gov/itl/applied-cybersecurity/cybersecurity-privacy-and-ai

Ngulube, P., & Mosha, N. F. V. (2025). Integrating artificial intelligence-based technologies 'safely' in academic libraries: An overview through a scoping review. Journal of Electronic Resources Librarianship. Advance online publication. https://doi.org/10.1080/07317131.2024.2432093

Sousa, N. (2025). Ethical and practical implications of AI in academic library research. IFLA Journal. Advance online publication. https://doi.org/10.1177/03400352251391753

Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

United Nations Educational, Scientific and Cultural Organization. (2021). Recommendation on the ethics of artificial intelligence. https://unesdoc.unesco.org/ark:/48223/pf0000381137