AI and Data Privacy in Academic Libraries: Emerging Issues, Risks and Ethical Considerations

DipRanjan[1], Salman Khan[2]
Dept Library and Information Science
[1] Research Scholar, YBN University Ranchi
[2] YBN University Ranchi

Abstract- *Academic libraries are integrating AI in discovery, reference/recommendation, metadata, research support and user analytics. Those tools can help to accelerate retrieval and extend the staff. They also rely on information about which libraries have always had a view of privacy – search history, borrowing history, reference questions, chatbot transcripts, research interests and usage logs. This paper explores the impact of AI on data privacy in academic libraries and the need for integrating AI ethics and Responsible AI and Data governance with traditional Library and Information Science concerns of confidentiality, intellectual freedom and trust. The study is descriptive and analytical literature-based design. It references peer-reviewed literature on LIS from 2018 to 2026, guidance from professional bodies, and legislative and policy documents, such as the Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025, the GDPR, the EU AI Act, UNESCO Recommendation on the Ethics of Artificial Intelligence, the OECD AI Principles and the NIST AI Risk Management Framework. It is seen in the literature that there is unequal strategic adoption, increasing use of generative tools and sometimes ethics discussion does not match the implementation. Data files of circulation are no longer the only source of privacy risk. It now incorporates inference, profiling, prompt leakage and vendor side training. An original conceptual model called the Responsible AI and Data Governance Framework for Academic Libraries (RAIDG-AL) is presented. It is a synthesis, and is not yet empirically tested.*

Keywords- Academic Libraries, Artificial Intelligence, AI Ethics, Data Governance, Data Privacy

I. INTRODUCTION

For two decades, academic libraries have transitioned from their local catalogues to suite networks, licensed databases, institutional repositories and research-data services. The newest addition to that stack is AI (Artificial Intelligence). It has been used in ranking algorithms, recommender modules, chat-based reference, automated subject assignment, text and image analysis, preservation workflows, and preservation dashboards that are supposed to describe user behavior [1]– [3]. Huang and Cox and Cox discovered that there remained a lack of explicit AI language in universities and libraries, and that there were stronger vision statements in some Chinese universities than in many plans in the UK, despite some isolated applications already in place [4] The guidance for AI in the education sector has since morphed into a call to action for libraries to move literacy, staff engagement and local norms and rules from experiments to leadership [5].

Libraries do not collect names of a sudden. Since then, they have always maintained membership files. The challenge is that AI systems are being fed data that historically resided in ‘data silos’ and that is used to learn from, combine and infer across. These pictures, which contain a search string, a downloaded article, a late-night question asked by a chatbot, and a repository deposited picture, can be combined to create a working picture of a student's or scholar's intellectual life. Prior to the current generative-AI trend, Jones and Salo wrote about the difficulties that librarians have been confronting in relation to learning analytics, which already conflict with librarians' responsibilities for intellectual freedom and confidentiality [6]. Zimmer's previous study on attitudes towards librarian privacy revealed the same conflict in this Web 2.0 context [7]. Generative tools intensify it because prompts may contain unpublished research, medical or legal worry, or examination material, and because many tools process that text on third-party servers [8], [9].

Traditional LIS ethics did not vanish when the interface changed. The ALA Code of Ethics still requires protection of each user’s privacy and confidentiality with respect to information sought or received and resources consulted, borrowed, acquired or transmitted [10], [27]. IFLA treats privacy as a condition of freedom of access and expression [11]. NISO (National Information Standard Organization) assigns shared privacy duties to libraries, publishers and software providers [12]. ALA’s later guidance on AI in libraries returns to the same core values—public good, intellectual freedom, privacy, sustainability, diversity, equity, inclusion and accessibility—and tells staff to review

vendor tools for privacy, security and equitable access before adoption [13]. Therefore, AI governance in an academic library is not a computer-science side project. It is technology work, management work, privacy work, ethics work, information-governance work, legal work and professional work at once.

A research gap remains. Reviews of library AI still find ethics treated thinly, often reduced to accuracy or a human-in-the-loop slogan, with weaker attention to consent, vendor training data, surveillance and intellectual freedom [1], [2], [14]. Safe-integration research is described as early [15]. Strategy studies show low prioritization even where tools are in use [4]. Indian academic libraries now sit under the Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025, and under the constitutional privacy holding in Justice K. S. Puttaswamy (Retd.) v. Union of India [16]– [18]. Internationally, GDPR, the EU AI Act, UNESCO's 2021 Recommendation, the OECD AI Principles and NIST's AI Risk Management Framework offer overlapping but distinct layers of law, ethics and technical guidance [19]–[23]. This study combines six core concepts, namely, Artificial Intelligence, Data Privacy, Academic Libraries, AI Ethics, Responsible AI and Data Governance, in one argument, which few papers have tried to corroborate in the field of LIS.

The basic issue is thus practical. What are the possibilities of applying AI in academic libraries to enhance discovery, reference, research support and internal workflows without compromising user privacy, confidentiality, intellectual freedom, data rights, autonomy, and trust? This paper aims to provide answers to this question by literature synthesis and proposed framework. Does not produce the new survey data.

II. RESEARCH OBJECTIVES AND QUESTIONS

The study will be guided by seven questions. RQ1. How is AI already impacting academic-library services and strategies? RQ2. Which types of users and research/operational information are processed by AI systems in libraries? RQ3. What privacy risks are posed, or exacerbated, by those systems when they infer, profile, store or send to vendors? RQ4. What ethics are related to algorithmic ranking, recommendations, generation of metadata and generative assistance? RQ5. How do Responsible AI principles translate to library practice and leave behind being just "buzz words"? RQ6. What data-governance arrangements are required if librarians, administrators, IT units, data-protection officers and vendors share the same data? RQ7. How can academic libraries create a conceptual framework to govern AI without forfeiting innovation in service?

The objectives are: to study and analyze the adoption of AI in academic libraries; to identify AI-enabled services that process data and analyze their emerging privacy risks; to address the issue of AI ethics in LIS context and introduce Responsible AI as a library management challenge; to explore data governance and vendor controls; to review Indian and international legal and professional instruments; and to propose the concept of RAIDG-AL as a model derived from literature for planning.

III. RESEARCH METHODOLOGY

The design is descriptive, analytical and literature-based. No questionnaire, interview sample, experiment or statistical test was conducted for this paper, and none is invented. The problem is conceptual, fast-moving and multi-jurisdictional. Peer-reviewed LIS work, official reviews and public legal texts already support a critical synthesis.

Sources were located through Google Scholar and publisher platforms associated with Elsevier, Emerald, SAGE, Taylor & Francis, Springer, ALA/ACRL and IFLA, together with official sites for UNESCO, OECD, NIST, the European Union and the Government of India. Search strings combined academic library, artificial intelligence, generative AI, chatbot, recommendation, privacy, confidentiality, AI ethics, Responsible AI, data governance, intellectual freedom, DPDP, GDPR and AI Act. Preference was given to 2018–2026 literature, with emphasis on 2023–2026 generative-AI and governance writing. Foundational privacy texts were kept where they still define professional duty.

A source was included only when bibliographic details could be verified and the work spoke to libraries, archives, scholarly communication or a directly applicable governance instrument. Unverified items, generic blogs and guessed DOIs were excluded. Thematic analysis organized the material around adoption, data types, privacy, ethics, Responsible AI, governance, bias and surveillance, generative AI, vendors, law and professional standards. Findings from earlier studies are distinguished below from the authors' synthesis and from the unvalidated RAIDG-AL proposal.

IV. CONCEPTUAL FOUNDATION

Artificial intelligence, in the sense used here, covers machine-learning and related systems that produce predictions, rankings, classifications, recommendations, generated text or decisions with some degree of autonomy

[23]. In an academic library that may mean a discovery ranker, a recommender, a reference chatbot, an automated subject string or a preservation classifier. Data privacy is the user's ability to inquire without unnecessary identification, tracking or disclosure of intellectual activity [7], [10], [11]. Academic libraries are teaching and research infrastructures whose users are students, faculty, researchers and staff, and whose services include discovery, reference, collections, repositories and research-data support.

AI ethics asks whether those systems are fair, transparent, accountable, explainable, inclusive and subject to human oversight, and whether they respect dignity and rights [20], [21]. Responsible AI is the organizational attempt to make those principles operational through design, procurement, monitoring and refusal where a tool cannot be justified [1], [13], [23]. Data governance is the set of rules and roles for ownership, classification, quality, minimization, purpose limitation, access, retention, deletion, sharing and audit. Related ideas—algorithmic bias, transparency, accountability, explainability, user autonomy, intellectual freedom and confidentiality—are not ornaments. They determine if a suggestion is a convenience, or a quiet steer, if a chatbot is a reference desk, or an unlogged interview.

V. REVIEW OF LITERATURE

A. AI Adoption in Academic Libraries

Huang, Cox and Cox demonstrated that when the websites reveal pilots, the strategy documents still don't necessarily say AI, just as the websites still don't necessarily say AI when strategy documents reveal pilots. national context shapes how boldly universities speak about the technology [4]. Danquah and colleagues reviewed This was to highlight aspects of implementation literature pertinent to Ghanaian academic libraries and emphasized skills and staged adoption as opposed to fashion [24]. Cox's editorial for a Journal of Academic Librarianship collection referred to generative tools as a change in user behaviour, but also of the role discovery is expected to play [3]. More recently, Fan and Hu consider other user needs, such as security and ethics and information quality, as fundamental requirements of generative-AI knowledge services [25] rather than as add-ons. The pattern is "take, but do not speak about it in governance language.

B. Privacy, Ethics and Responsible Practice

Responsible practice of AI in libraries and archives was explored by Mannheimer and co-authors. AI work was concentrated in big academic libraries, primarily on digitized text and images, and was most frequently for metadata or reference. Only about one-third went beyond the accuracy and human-in-the-loop (HITL) [1] and just over half referred to ethics or values. The ethical and fair approaches have been discussed in Hodonu-Wusu and it was stated that power must be accompanied by responsibility towards library users in AI [14]. In the scoping review of 'safe' integration, Ngulube and Mosha discovered that the literature was very young and that "anxiety of losing their jobs dominated other concerns, with privacy violation emerging as a real threat due to the role that libraries play in preserving seeking behaviour" [15]. As outlined by Sousa in his review of AI in academic-library research in the IFLA Journal, issues of privacy, bias, integrity, and diminishing human agency remain to be addressed [2].

C. Generative AI, Integrity and Refusal

Vogus summarized early discussion on ChatGPT as a binary option between adopting or rejecting, and noted a leaning towards informed use [26]. In Library Trends, Slater proposed 'critical refusal' when there is conflict between generative systems and antiracism, sustainability and intellectual freedom and privacy [27]. Matsieli and Mutula embedded generative-AI ethics in the context of the information-society, and emphasized the need to articulate harms, especially in environments where there is a lack of policy and underdeveloped infrastructure [8]. The contrast helps out. It illustrates that AI literacy can refer to the way that a tool is prompted or the determination that a tool is not applied to a reference interview.

D. Governance Instruments and Research Gaps

Professional and public instruments now sit beside the LIS literature: ALA AI guidance [13], IFLA and NISO privacy texts [11], [12], UNESCO [20], OECD [21], NIST AI RMF [23], GDPR [19], the EU AI Act [22], [28], and India's DPDP In accordance with the framework and Puttaswamy judgment [16]–[18]. Sousa has been laying out the AI Act's risk classification, and correlating it with library classification and procurement practice, and suggesting that ethical language continues to outpace viable regulation [28]. This paper aims to fill this missing link between those instruments and everyday academic-library data activities.

VI. ARTIFICIAL INTELLIGENCE IN ACADEMIC LIBRARIES

Already, AI is linked to routine library tasks. Discovery systems rank results by importance. Recommenders suggest titles. Chatbots respond to queries about hours and are now able to respond to questions about topics. Subjects and abstracts suggested by metadata generators. Theses or Heritage scans are sorted by classifiers. Research Assistants produce paper summaries. Risk is identified using preservation tools. Collection dashboards predict needs. Who used what is described by analytics. Entities are extracted

from the deposits using Repositories. Units of writing assistants are tested for scholarly-communication. There are benefits and data shadows associated with each use, as summarized in Table I.

TABLE I. AI APPLICATIONS IN ACADEMIC LIBRARIES AND PRIVACY IMPLICATIONS

AI application	Library function	Data used	Privacy risk	Ethical concern	Mitigation
Discovery ranking	Search and retrieval	Queries, clicks, holdings	Query logs reveal topics	Opaque ranking of scholarship	Short logs; explain ranking rules
Recommendation	Personalised suggestion	Loans, views, profiles	Reading profiles	Filter bubbles	Opt-in only; aggregate models
Reference chatbot	Enquiry and triage	Questions, chat text	Prompt leakage	Unreviewed advice	Local or contracted models; human referral
Metadata generation	Cataloguing	Full text, images	Sensitive corpus content	Biased description	Documented models, cataloguer review
User analytics	Service planning	Use, device, identity	Re-identification	Surveillance feel	Anonymise; purpose limitation
Generative research help	Research support	Prompts, uploads	Unpublished work to vendor	Hallucination; integrity	Do an audit and ban sensitive uploads;check

VII. AI AND DATA PRIVACY

There are four ways AI alters privacy. Firstly, it is about collection - models require history, not just the current transaction. Second, aggregation is the connection of library systems with campus identity and publisher platforms and cloud AI. Third, inferences generate new facts that no one typed: probable discipline, probable distress, probable exam timing. Fourthly, automated decisions are based on those inferences and are made by ranking, recommending or routing a user [6] and [15]. The chain is called personalization because it's friendly. A full behavioural dossier on the library does not have to be presented to a reader before he/she is told of the existence of a database. Most service value could be carried by opt-in profiles, and session-only signals and anonymized counts.

VIII. EMERGING PRIVACY RISKS

The live hazards include excessive collecting, inadequate consent, too long storage periods, secondary use, re-identification of ‘anonymous' logs, instant exposure of prompt data on public generative tools, cross-border processing, and the inaudible sharing between discovery, learning-analytics and AI layers [1, 8, 9, 15]. Don't need to use cameras here for surveillance. A repeatable individualized graph, certainly suffice to cool a search on sexuality, opposition politics or a supervisor conflict. Causes and responses are presented in Table II.

TABLE II. AI PRIVACY RISKS IN ACADEMIC LIBRARIES

Risk	Cause	Affected data	Potential impact	Recommended response
Prompt leakage	Public or vendor LLMs	Reference and research text	Confidentiality loss	In enterprise level terms, without any sensitive prompts.
Profiling	Cross-system analytics	Search, loans, identity	Chilling of inquiry	Minimise joins, personalisation opt-in.
Vendor training reuse	Unclear contracts	Chats, deposits, queries	Secondary exploitation	No training with library data should be done.
Re-identification	Rich ‘anonymous’ logs	Device + topic trails	Broken anonymisation	Coarser aggregates, deletion clocks
Opaque ranking	Black-box discovery	Queries and clicks	Differential visibility of work	Record features; periodic audit.
Weak consent	Bundled campus login	Credentials and traces	No meaningful choice	Separate artificial intelligence notices; Easy refusal.

IX. AI ETHICS IN ACADEMIC LIBRARIES

Demographic parity is not the only path to fairness in a library. Whether it is a ranker hiding working titles from the Global South, or a recommender nudging first-year students into a narrow canon, or a chatbot answering some accents worse than others [2,8,14]. Transparency involves user awareness of dealing with a machine [21]. Accountability is the person or persons (library officer named) with the authority to turn a tool off. Explainability refers to the possibility of seeing the reasons for the proposal of a hearing on the part of a cataloguer. Autonomy: a researcher does not have to be trained to a profile to be able to search. Confidentiality and intellectual freedom continue to be the focus of the LIS tests: if a system creates a question which is unsafe, it is unsafe professionally [10], [11], [13]. Inclusiveness and accessibility are important because the AI service that doesn't support screen-readers or regional languages isn't a neutral improvement. Human oversight is practical non-maleficence: the human librarian retains the reference interview.

TABLE III. AI ETHICS PRINCIPLES AND LIBRARY APPLICATIONS

Ethical principle	Meaning	Library application	Potential risk	Recommended practice
Fairness	Avoid unjustified disadvantage	Ranking and recommendation	Canon bias	Diverse evaluation sets
Transparency	Users understand the use of AI.	Chatbots and discovery	Hidden automation	Clear interface labels
Accountability	Named responsibility	Procurement and incidents	Vendor blame-shifting	On every tool, library owner.
Explainability	Reasons can be looked at.	Metadata suggestions	Unreviewable output	Keep cataloguer in the loop
Autonomy	User control of data	Personalisation	Forced profiles	Default off; easy export/delete
Intellectual freedom	Inquiry without fear	Search analytics	Chilling effect	No secondary use of queries

X. RESPONSIBLE AI IN ACADEMIC LIBRARIES

Responsible AI, as far as a librarian's point of view is concerned, is not a hype. It's a decision rule. A tool may be used if used by a human that can be supervised, if the purpose is stated, if data are minimized, if no bias can be detected, if appropriate security measures are available and if the library does not lose the collection in selecting to use the vendor. Reliability is important as a misquoted (hallucinated) book on a reference desk is a professional error, not some "quirky feature" [2, 26]. Inclusiveness is important as students that don't have their own devices will first encounter AI on devices in public libraries. In the context of low-level tasks: sustainability should be part of the discussion when energy-intensive models are performed. Those principles become library requirements in Table IV.

TABLE IV. RESPONSIBLE AI REQUIREMENTS FOR ACADEMIC LIBRARIES

Responsible AI principle	Library requirement	Implementation strategy	Expected benefit
Human oversight	No unsupported high-stake tips.	Review staff and set referral thresholds.	Fewer harmful answers
Privacy	Least data for the task	Local or contracted processing	Lower leak surface
Fairness	Documented bias checks	Periodic ranking audits	More equitable discovery
Transparency	Public AI use register	Notices on each tool	Informed users
Security	Same bar as ILS data	Access control; encryption	Reduced breach impact
Accountability	Named AI owner	Incident playbook	Faster response

XI. DATA GOVERNANCE FOR AI-ENABLED ACADEMIC LIBRARIES

Data governance is a response to the questions: Who can collect what, why, how long and with whom? When a sentence is touched by a campus ID provider, a library discovery layer and a cloud model provider, that ownership is not clear. Map stewardship to the library for patron inquiry data, research-data offices for research data deposits, and IT for infrastructure logs, under a written map [12], [16]. Classification should distinguish between the public metadata, identifiable use data and research content. Quality is important, because "biased training or dirty metadata is biased output" [1]. Minimization and purpose limitation are the same grammar of the DPDP as of the GDPR, namely the one of the old circulation ethics: no retaining the retainer's traces after the service need has ended [16], [19]. The paperwork that makes that ethic enforceable is access, retention, deletion, audit and data-sharing agreements. The Privacy impact assessment should be carried out before a new model; it should not come in response to a complaint.

TABLE V. DATA GOVERNANCE REQUIREMENTS

Governance area	Requirement	Responsible stakeholder	Key action
Ownership	Named data steward	Library administration	Write a data map
Minimisation	Collect only what the service needs	Service owners	Strip fields at design
Retention	Deletion with precise time of identifiable AI logs	IT and library	Automated purge jobs
Vendor sharing	No silent training use	Procurement	Contract clauses and audits
Access control	Least privilege	IT	Role-based accounts
PIA	Assess before launch	DPO and librarians	Record residual risk

XII. ALGORITHMIC BIAS, PROFILING AND SURVEILLANCE

This may be introduced via training corpora, via click logs that award virility to already-visible titles, via classification schemes that carry the seeds of colonialism, or via recommender loops that restrict a syllabus [2], [8], [14]. The patterns are used to create a “standing picture” of the person, through the process of profiling. That picture is then analyzed as a forecaster and becomes predictive. This is not something naturally found in all systems and to state that every ranker is biased in the same manner is sloppy. The evidence-based claim is narrower: library AI often lacks published evaluation against fairness and privacy, and personalization can be built without asking whether the library needed identifiable history at all [1], [6], [15]. Intellectual freedom is affected when a student avoids a topic because the system is known to remember. Academic freedom is affected when a scholar’s exploratory reading becomes a management statistic.

XIII. GENERATIVE AI AND ACADEMIC LIBRARIES

Generative system now draft answers, summaries articles, propose metadata, translate abstracts and tutor information-literacy sessions [3], [25], [26]. They help when staff time is short. They harm when a prompt containing a thesis chapter leaves the institution; when a fluent paragraph cites a paper that does not exist; when copyrighted collections are used as informal training stock; or when students treat library chat as a substitute for citing sources [2], [8], [9]. Academic integrity is therefore a library issue, not only an examinations-office issue, because libraries teach how evidence is found and checked. Responsible adoption means enterprise or locally hosted models for confidential work, documented limits, human review of public-facing answers, and literacy that includes refusal skills as well as prompt skills [5], [13].

XIV. LEGAL, ETHICAL AND PROFESSIONAL FRAMEWORK

These layers should not be collapsed. Law states duties. In India, Puttaswamy recognized privacy as a fundamental right [18]. The DPDP Act, 2023 and the DPDP Rules, 2025 regulate digital personal data, including notice, consent or other lawful bases, reasonable security safeguards and breach intimation [16], [17]. Where applicable, there is a lack of reporting obligations or duties for intermediaries under the Information Technology Act, 2000 and the CERT-In directions. Many international collaborations are subject to GDPR for processing of personal data [19]. The EU AI Act (Regulation (EU) 2024/1689) introduces a risk-based framework for AI systems, and includes obligations for literacy and transparency that libraries as deployers of AI systems might be able to fulfil even if they are not a high-risk provider [22][28].

Ethics is the law in action plus, or the law that goes above law. The 2021 Recommendation of UNESCO is rooted in the principles of human rights, proportionality, fairness, transparency and human oversight [20]. Inclusive growth, human-centered values, transparency, robustness and accountability are the OECD AI Principles 2024 [21].

Requirements of librarianship: Professional standards, including ALA ethics and AI guidance [10]– [13] and IFLA privacy. Institutions can use a framework like NIST AI RMF 1.0 (Govern, Map, Measure, Manage) or NIST CSF 2.0 as a structure for the risk work without assuming they are legislation [23] [31]. NITI Aayog is aiming for 'AI for All' in its Policy Horizon Document, the National Strategy for Artificial Intelligence, and not a library code [32].

XV. CHALLENGES IN IMPLEMENTING RESPONSIBLE AI

Libraries have no personnel, who can read a model card and a license with equal confidence [5] [24]. Budget comes before governance time. Legacy discovery contracts come with AI capabilities, though they don't necessarily receive a new privacy assessment [28]. Algorithms stay opaque. Policies refer to privacy and subsequently install a public chatbot on the home page. Unclear legal framework in regards to training data and cross-border transfers postpones decisions or makes for reckless decisions. All data quality within catalogues is carried through into models. Institutional resistance is manifested either in uncritical enthusiasm or in uncritical ban. Both skip evaluation. Meanwhile, users are coming with consumer-AI habits, and no notion that the library search ever was confidential.

XVI. PROPOSED RESPONSIBLE AI AND DATA GOVERNANCE FRAMEWORK FOR ACADEMIC LIBRARIES

The Responsible AI and Data Governance Framework for Academic Libraries (RAIDG-AL) is suggested in this section. It is a conceptual model that draws from and combines the LIS literature, the AI-ethics tools, Responsible AI guidelines, data-governance practice, privacy principles and the laws and guidance referenced above. It has never been field tested; is not an empirically validated instrument.

The fourteen components are operated as a cycle. AI governance identifies owners and assigns tools to risks. Hidden logs are made visible using data inventory and classification. Privacy and risk assessment: What is the inference that a system can draw? Data minimization and purpose limitation reduces the feed. Implementing responsible AI principles translates to local rules. Human oversight provides a line for a librarian to cross. Transparency/Explainability informs users what is automated. Bias Fairness audit of ranking and metadata. If there are any leftover data, they are protected by security and privacy controls. Vendors hold processors accountable through vendor governance. The rules become usable through staff literacy with AI. Real choices through user awareness and control. Monitoring and audit are search for drift. Incident response and Improvement close the loop. The sequence is illustrated in Fig.1.

This is written version of the figure 1. The RAIDG-AL cycle is a single descending sequence that starts at the top and then comes back to governance after incidents:

AI Governance → Data Inventory and Classification → Privacy and Risk Assessment → Data minimization and Purpose Limitation → Responsible AI Principles → Human Oversight → Transparency and Explainability → Fairness and Bias Auditing → Security and Privacy Controls → Vendor and Third-Party Governance → Staff AI Literacy → User Awareness and Control → Continuous Monitoring and Auditing → Incident Response and Continuous Improvement → (return to AI Governance).

The LIS values the framework aims to safeguard, as well as external reference points, discussed above, sit around such spine: user confidentiality, intellectual freedom, trust, equitable access, and academic integrity. Control is each step down. The arrow back is the audit and incident lesson that should impact the next governance decision. Purpose, actions and owners are provided for each step in Table VI.

Fig. 1. RAIDG-AL framework sequence (written representation).

TABLE VI. RAIDG-AL FRAMEWORK

Framework component	Purpose	Key actions	Responsible stakeholder
1. AI governance	Assign authority	Sign up AI tools.	Library leadership
2. Data inventory	Know the data	Catalogue logs and vendors	Systems librarians
The privacy and risk assessment are incorporated.	Anticipate harm	PIA before launch	DPO and librarians
4. Minimisation	Cut unnecessary data	Default anonymous search	Service owners
5. Responsible AI principles	Local ethical rules	Publish a short code	Ethics or AI group
6. Human oversight	Keep professional judgement	Referral rules for chat	Reference staff
7. Transparency	No hidden AI	Label automated services	Web and service teams
8. Fairness auditing	Find systematic skew	Sample ranking outputs	Assessment librarians
9. Security controls	Protect residual data	Encryption; access control	IT
10. Vendor governance	Bind processors	No-training clauses; audit	Procurement
11. Staff AI literacy	Capable judgement	CPD on privacy and bias	LIS educators / HR
12. User control	Meaningful choice	Opt-in personalisation	Public services
13. Monitoring	Detect drift	Log review; complaints	IT and library
14. Improvement	Learn from failure	Revise policy after incidents	Administration

Stakeholders are not all working together equally. Confidentiality still remains at the desk with the librarians, and a chatbot now has to be asked what he or she stores. The vendor feature is enabled or disabled at the administrator's discretion. Technical locks are put in place by IT personnel. Library Traces will be calculated in campus analytics if it is desired by university management. If not contracted otherwise, the models are written by the vendors. There is a need for students and researchers of. real off switch for personalization, not a buried clause. Responsible AI fails when any one of those groups is treated as a spectator.

XVII. DISCUSSION

It is coherent, but the way (pathway) you see in the literature is not fate. The more the innovation with AI, the more processing power is required. The more it's processed, the more privacy is compromised. Without ethics and governance, exposure turns to be an institutional risk. If risk comes to pass, it's an expense to the user's trust. Academic libraries operate based on the asset of trust for inquiry. On one hand, personalization, and on the other hand, privacy. Efficiency fights humans! minimization is at odds with data utilization. Automation vs. expert judgement. Policy is laggard, innovation is quick. The javelin-like name for that refusal is Responsible AI.

Librarians continue as privacy advocates and are now evaluators of AI. Policy and money belong to Administration. IT owns controls. Campus Analytics may "devour" library logs, owned by university management. Vendor owns code (unless contracts otherwise state). Users own their questions and should be able to retain them. If the other is imaginary, none of those works.

XVIII. RECOMMENDATIONS

Academic librarians should presume query/chat/consultation text is confidential and refuse non-essential profiling, and should offer AI literacy with verification and refusal. Library administrators should maintain an up-to-date publicly available register of AI tools, budget to conduct impact assessment, and designate an owner for each AI system. Failing to have a clearly documented justification, library logs should not be folded into learning analytics by university management [6]. IT departments need to provide isolation for AI workloads, encryption of identifiable AI stores and logging of administrator access. DPDPs and, if applicable, GDPRs should be mapped onto chatbot and discovery data by DPOs [16, 19]. Librarians should be part of AI governance committees, beyond central IT. LIS teachers need to provide an algorithm awareness lesson, vendor reading, privacy design lesson. Vendors should provide no training contracts, deletion, audits and plain documentation [12, 13]. Policymakers need to give memory institutions sector notes, not let libraries figure out the application of the AI statutes [28]. Any frameworks like RAIDG-AL should be tested in real libraries and researchers should not postulate success measures in advance.

XIX. FUTURE DIRECTIONS

Unanswered yet speculative questions that are coming up are around agentic options that work across accounts, multimodal options that can handle scans and speech, privacy preserving approaches like on device or federated learning, synthetic data for dashboard testing, increased explainability for ranking, and stricter control of general-purpose models. Pressure will also be placed on libraries to allow commercial crawlers to use library collections for training. But every one of those developments is to be regarded as a procurement and ethics event, rather than an inevitable upgrade [8, 9, 13, 25].

XX. LIMITATIONS

Design of this paper has its limitations. Nor does it provide a survey of a named set of libraries, nor does it measure user confidence empirically, nor does it validate RAIDG-AL in practice. There are variations in coverage of language/region because the searchable peer-reviewed record is not similar across language/region. Information on legal instruments on the level required for LIS argument, not formal advice. The limits are reasons for empirical follow-up, not reasons for a sample invention.

CONCLUSION

It is clear from the literature that academic libraries already experience AI, frequently even before their strategy papers, and even before their ethics papers [1], [4], [15]. At the heart of it, privacy is prominent, since the data that AI desires is the data that was once a professional secret of the circulation file. AI ethics is important because ranking and recommendation and answers generated now become part of what a student can "find", in practice. Librarians have always held roles that are key to responsible use of AI and data governance: confidentiality, intellectual freedom, equitable access and care with other people's questions [10, 11, 13]. Human oversight is not “something from the past.” It's what makes a reference a reference service. There is nothing user trust can handle more than a season of leaked prompts and mysterious profiles. These tools can be used by future academic libraries. The tools can't be emptied of moral content.

REFERENCES

[1] S. Mannheimer et al., “Responsible AI practice in libraries and archives: A review of the literature,” Information Technology and Libraries, vol. 43, no. 3, 2024, doi: 10.5860/ital.v43i3.17245.

[2] N. Sousa, “Ethical and practical implications of AI in academic library research,” IFLA Journal, 2025, doi: 10.1177/03400352251391753.

[3] A. M. Cox, "Artificial intelligence and the academic library," Journal of Academic Librarianship, vol. 50, art. 102965, Nov. 2024.

[4] Y. Huang, A. M. Cox, and J. Cox, "Artificial intelligence in academic library strategy in the United Kingdom and the Mainland of China," Journal of Academic Librarianship, vol. 49, no. 6, art. 102772, 2023, doi: 10.1016/j.acalib.2023.102772.

[5] L. S. Lo, "An initial interpretation of the U.S. Department of Education's AI report: Implications and recommendations for academic libraries," Journal of Academic Librarianship, vol. 49, no. 5, art. 102761, 2023, doi: 10.1016/j.acalib.2023.102761.

[6] K. M. L. Jones and D. Salo, "Learning analytics and the academic library: Professional ethics commitments at a crossroads," College & Research Libraries, vol. 79, no. 3, pp. 304–323, 2018, doi: 10.5860/crl.79.3.304.

[7] M. Zimmer, "Librarians' attitudes regarding information and Internet privacy," Library Quarterly, vol. 84, no. 2, pp. 123–151, 2014, doi: 10.1086/675329.

[8] M. Matsieli and S. Mutula, "Generative AI and the information society: Ethical reflections from libraries," Information, vol. 16, no. 9, art. 771, 2025, doi: 10.3390/info16090771.

[9] K. Slater, "Against AI: Critical refusal in the library," Library Trends, vol. 73, no. 4, pp. 588–608, May 2025, doi: 10.1353/lib.2025.a968497.

[10] American Library Association, "Code of Ethics of the American Library Association." [Online]. Available: https://www.ala.org/tools/ethics

[11] IFLA, "IFLA Statement on Privacy in the Library Environment," 14 Aug. 2015. [Online]. Available: https://www.ifla.org/publications/ifla-statement-on-privacy-in-the-library-environment/

[12] NISO, NISO Consensus Principles on Users' Digital Privacy in Library, Publisher, and Software-Provider Systems. Baltimore, MD, USA: NISO, 2015.

[13] American Library Association, "Guidance on the Use of Artificial Intelligence in Libraries." [Online]. Available: https://www.ala.org/tools/standards-and-guidelines/guidance-use-artificial-intelligence-libraries

[14] J. O. Hodonu-Wusu, "The rise of artificial intelligence in libraries: The ethical and equitable methodologies, and prospects for empowering library users," AI and Ethics, vol. 5, pp. 755–765, 2025, doi: 10.1007/s43681-024-00432-7.

[15] P. Ngulube and N. F. V. Mosha, "Integrating artificial intelligence-based technologies 'safely' in academic libraries: An overview through a scoping review," Technical Services Quarterly, vol. 42, no. 1, pp. 46–67, 2025, doi: 10.1080/07317131.2024.2432093.

[16] Government of India, The Digital Personal Data Protection Act, 2023, Act No. 22 of 2023.

[17] Ministry of Electronics and Information Technology, Government of India, Digital Personal Data Protection Rules, 2025, G.S.R. 846(E), 13 Nov. 2025.

[18] Justice K. S. Puttaswamy (Retd.) and Anr. v. Union of India and Ors., (2017) 10 SCC 1.

[19] Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation), 2016.

[20] UNESCO, Recommendation on the Ethics of Artificial Intelligence. Paris, France: UNESCO, 2021.

[21] OECD, OECD AI Principles, adopted 2019, updated 2024. [Online]. Available: https://www.oecd.org/en/topics/ai-principles.html

[22] Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), 2024.

[23] E. Tabassi, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, Jan. 2023, doi: 10.6028/NIST.AI.100-1.

[24] M. M. Danquah, P. S. Dadzie, K. Gyesi, F. Yeboah, and C. Y. Nyarko, "Artificial intelligence implementation strategies for Ghanaian academic libraries: A scoping review," Journal of Academic Librarianship, vol. 50, no. 6, art. 102975, 2024, doi: 10.1016/j.acalib.2024.102975.

[25] S. Fan and Z. Hu, "Repositioning academic libraries with generative artificial intelligence: Identifying and prioritizing the evolving needs of knowledge services," Journal of Academic Librarianship, art. 103288, 2026, doi: 10.1016/j.acalib.2026.103288.

[26] B. Vogus, "Generative AI and ChatGPT: Friend or foe for academic libraries?" Public Services Quarterly, vol. 19, no. 4, pp. 309–312, 2023, doi: 10.1080/15228959.2023.2266358.

[27] American Library Association, "Privacy: An Interpretation of the Library Bill of Rights." [Online]. Available: https://www.ala.org/advocacy/intfreedom/librarybill/interpretations/privacy

[28] N. M. T. Sousa, "Classifying classification: Risks of the AI Act and the (re)organization of artificial intelligence in academic libraries," Alexandria, 2025, doi: 10.1177/09557490251406841.

[29] Government of India, The Information Technology Act, 2000, Act No. 21 of 2000.

[30] Indian Computer Emergency Response Team, Directions under sub-section (6) of section 70B of the Information Technology Act, 2000, 28 Apr. 2022.

[31] National Institute of Standards and Technology, The NIST Cybersecurity Framework (CSF) 2.0, NIST CSWP 29, 26 Feb. 2024.

[32] NITI Aayog, National Strategy for Artificial Intelligence. New Delhi, India: Government of India, 2018.