

# Artificial Intelligence and Machine Learning in Indian Academic Libraries: Toward a Trustworthy Sociotechnical Framework for Classification, Cataloguing, and User Services

Haridas Karmakar

M.A, M.Lis. , Librarian Pragati College, Agomani, Dhubri, Assam

**Abstract-** Artificial intelligence (AI) and machine learning (ML) are increasingly shaping the way libraries organize knowledge, describe resources, and interact with users. In academic libraries, these technologies promise faster classification, automated metadata enrichment, improved discovery, and more responsive user services. Yet their adoption also raises serious questions about bias, transparency, privacy, professional accountability, and growing dependence on vendor-controlled systems. This conceptual paper examines the application of AI and ML in three core areas of library practice: classification, cataloguing, and user services, with particular attention to the Indian academic library context. Drawing on library and information science literature, together with policy and ethical guidance, the paper argues that AI in libraries should be understood not as a neutral automation tool, but as a sociotechnical intervention that affects representation, access, and trust. The paper proposes a trustworthy sociotechnical framework based on six dimensions: task-risk alignment, data and metadata stewardship, human accountability, transparency and contestability, user rights and inclusion, and organizational readiness. It concludes that AI can strengthen library work only when it remains aligned with professional judgment, user welfare, and the public mission of librarianship.

**Keywords:** artificial intelligence, machine learning, academic libraries, cataloguing, user services, knowledge organization.

## I. INTRODUCTION

Artificial intelligence has become a serious subject of discussion in library and information science rather than a distant technological possibility. Libraries across the world have been experimenting with machine learning for metadata generation, discovery enhancement, chatbot-based assistance, collection analysis, and workflow automation. In academic libraries, interest in AI has grown alongside the expansion of digital repositories, linked data, large-scale digitization, and user expectations for faster and more personalized access to information (Asemi et al., 2021; Wheatley & Hervieux, 2019). These developments suggest that AI and ML are no longer external to library work. They are beginning to reshape how information is organized, located, and mediated.

This shift is especially significant for Indian academic libraries. The Indian higher education system is marked by linguistic diversity, uneven infrastructure, expanding digital learning ecosystems, and strong pressure to improve access, quality, and research visibility. National Education Policy 2020 places emphasis on digital transformation, multilingualism, equitable access, and the strengthening of knowledge institutions (Ministry of Education, 2020). In that setting, AI appears attractive because it promises efficiency in repetitive technical services, improved search and retrieval, better support for digital collections, and more responsive user-facing services. Yet the same technologies may also intensify structural problems if they are adopted without critical scrutiny. Automated classification may reinforce biased vocabularies, chatbots may provide fluent but unreliable responses, and proprietary discovery tools may reduce institutional

control over data, privacy, and professional standards.

The central problem addressed in this paper is that AI adoption in libraries has often been discussed in terms of innovation and efficiency, while questions of governance remain underdeveloped. Libraries do not merely store information. They create descriptive systems, define pathways of access, and mediate between users and knowledge. Classification, cataloguing, and user services are therefore not purely technical functions. They are interpretive and ethical practices that shape visibility, exclusion, and trust. This point has long been clear in critical work on classification and representation (Bowker & Star, 1999; Olson, 2002), and it becomes even more urgent when machine-generated outputs begin to influence how resources are named, grouped, recommended, or explained.

This paper argues that AI in libraries should be understood through a sociotechnical lens. The relevant question is not simply whether AI can perform a task, but under what conditions its use improves knowledge organization and access without weakening equity, accountability, or user trust. The paper focuses on three interrelated domains: classification, cataloguing, and user services. These were selected because they represent core library functions in which AI is both practically attractive and normatively consequential. Classification affects how knowledge is arranged and retrieved. Cataloguing affects bibliographic identity, metadata quality, and interoperability. User services affect how people encounter the library, seek help, and make decisions about information.

The paper is conceptual in nature. It does not report new survey or case-study data. Instead, it synthesizes scholarship, policy guidance, and critical debates in order to develop a framework for trustworthy adoption. The contribution is twofold. First, it brings technical services and user services into a single analytical conversation, rather than treating them as separate arenas of AI deployment. Second, it proposes a trustworthy sociotechnical framework tailored to the realities of academic libraries, with special relevance to the Indian context.

The framework aims to help library professionals, researchers, and institutional decision-makers evaluate AI not only by speed or novelty, but by whether it supports the broader mission of librarianship. In the Indian tradition of library science, this concern resonates strongly with Ranganathan's insistence that library systems must save the time of the reader while remaining fundamentally oriented toward access, use, and service (Ranganathan, 1931).

## II. LITERATURE REVIEW

The literature shows that AI in libraries has already been discussed across several connected areas. Asemi et al. (2021) present a broad review of intelligent libraries, showing that expert systems, robotics, and AI applications have been explored in relation to service delivery, information retrieval, management, and user support. Wheatley and Hervieux (2019) similarly identify AI as an emerging field in academic libraries, with implications for discovery, research support, automation, and professional change.

In classification and cataloguing, the most relevant literature focuses on automated subject assignment, metadata generation, and linked bibliographic infrastructures. Kragelj and Kljajic Borstnar (2021) demonstrate that machine learning can support the automatic classification of older electronic texts into the Universal Decimal Classification. Their study is important because it shows the practical value of AI for large-scale knowledge organization. At the same time, it also reveals the dependence of such systems on existing corpora and prior classification decisions.

Cordell (2020), in a report for LC Labs, argues that machine learning offers significant opportunities for access and description but requires careful human interpretation and institutional reflection. This connects directly with ongoing developments in linked data and BIBFRAME, where bibliographic description is increasingly conceived as part of a networked web environment rather than as a static local record (Gonzales, 2014; Library of Congress, 2016). Under such conditions, errors or biases in machine-generated metadata are not isolated problems. They can circulate across systems.

The literature on user services focuses heavily on conversational systems and virtual reference. McNeal and Newyear (2013) provide an early account of chatbots in libraries, arguing that automated agents can handle routine queries and reduce staff burden. Mckie and Narayan (2019) further explore the implications of chatbots for the academic library experience, emphasizing the need for thoughtful design and a clear understanding of user expectations. Kane (2019) adds an important operational perspective by showing that chatbot deployment requires continuous management, transcript analysis, and refinement rather than one-time implementation. Rodriguez and Mune (2022) illustrate how a university library can implement a virtual reference chatbot with limited coding demands, but their work also suggests that such systems need clear boundaries and ongoing oversight.

A more critical strand of literature complicates any simple celebration of AI in libraries. Bowker and Star (1999) showed that classification systems are not neutral containers of knowledge; they are social arrangements shaped by power, omission, and institutional priorities. Olson (2002) made a similar argument in relation to subject representation, highlighting how library systems can marginalize or flatten certain forms of knowledge. Noble (2018) extended this critique into algorithmic environments by demonstrating how search systems can reinforce structural bias under the appearance of technical objectivity. These insights are highly relevant to AI in library classification and discovery because machine learning systems are trained on historical data and existing descriptive regimes. They may therefore reproduce past exclusions at a greater scale and speed.

Despite this growing body of work, an important gap remains. Much of the literature describes applications, reports environmental scans, or discusses individual tools such as chatbots and machine learning classifiers. Less attention has been given to a unified conceptual framework that connects classification, cataloguing, and user services under a common logic of trustworthy adoption, particularly in the context of Indian

academic libraries. Given India's multilingual knowledge environment, uneven technological capacity, and the strong public-interest role of universities and libraries, this gap is significant. A more integrated framework is needed.

### III. METHODOLOGY

This study is a conceptual paper based on critical integrative analysis. It does not involve surveys, interviews, experiments, or case-study fieldwork. Instead, it draws together scholarship from library and information science, knowledge organization, AI ethics, and library policy in order to develop an interpretive framework for analyzing AI and ML in academic libraries.

The conceptual approach proceeds in three stages. First, the paper identifies the major library functions in which AI and ML are being applied, focusing on classification, cataloguing, and user services. These domains were selected because they are central to the library's role as an institution of knowledge organization and access. Second, the paper synthesizes recurring claims in the literature about the benefits of AI, including efficiency, scale, metadata enhancement, and user responsiveness. Third, it examines the recurring concerns that accompany those claims, especially bias, opacity, privacy risks, deskilling, and vendor dependence.

The aim is not to construct a universal technical standard. Rather, it is to build a framework that can help academic libraries, particularly in India, think more clearly about what trustworthy AI would require in practice. The analysis is sociotechnical because it assumes that AI systems are shaped not only by algorithms, but also by data quality, metadata standards, institutional capacity, procurement decisions, professional values, and user rights. From this perspective, trustworthiness is not a feature built into software alone. It is a property that emerges from the relationship between technology, governance, and library practice.

#### IV. RESULTS OR ANALYSIS

The analysis suggests that AI and ML alter library work in three broad ways. First, they shift some activities from explicit rule-based processing to probabilistic inference. Second, they change professional labor from sole production to oversight, validation, correction, and governance of machine-generated outputs. Third, they alter the interface between the library and the user by introducing systems that can recommend, summarize, and converse rather than simply retrieve. These changes create clear operational opportunities, but they also raise deeper questions about authority, representation, and trust.

In classification, AI offers scale and speed. Large collections, backlogs of digitized material, and multilingual corpora make manual subject assignment difficult. Machine learning can identify patterns across records and propose class numbers or subject groupings with greater speed than manual workflows allow. Kragelj and Kljajic Borstnar (2021) show that this is especially useful for older electronic texts that would otherwise remain poorly organized. For Indian academic libraries, where retrospective conversion, institutional repository growth, and multilingual collections remain important concerns, such systems can be attractive.

Yet classification is never only a technical sorting exercise. It is a knowledge claim. It determines what is placed together, what is separated, and what becomes discoverable through particular conceptual pathways. If AI systems learn from uneven or biased descriptive histories, then they may amplify those histories rather than correct them. This is particularly important in the Indian context, where local knowledge traditions, regional languages, and interdisciplinary subjects may not fit neatly into inherited or globally dominant classification patterns.

Cataloguing and metadata work reveal a similar tension. AI can support automatic keyword extraction, entity recognition, duplicate detection, authority matching, and metadata enhancement. In practical terms, this may reduce repetitive effort and

help libraries process large volumes of material more efficiently. It may also improve discoverability by enriching sparse records. However, the movement toward linked bibliographic infrastructures means that metadata errors are no longer contained within a single catalog. They may travel across systems and platforms. This is one reason why human review remains essential. Bibliographic description affects scholarly visibility, retrieval accuracy, and interoperability. The relevant question is therefore not whether AI can generate metadata, but whether the institution can govern the quality, provenance, and revision of those outputs. In Indian academic libraries, where metadata quality often varies across institutions and platforms, AI could either support standardization or deepen inconsistency, depending on how it is implemented.

User services are the most visible area of AI adoption because they directly affect the user experience. Chatbots, recommendation tools, and automated query support are appealing because they offer immediacy, continuous availability, and low-friction assistance. Routine questions about opening hours, borrowing rules, resource access, or basic navigation can often be handled by conversational systems. This can save staff time and improve responsiveness. For large universities with growing online and hybrid user communities, such tools may be especially useful. Yet the risk is that the fluency of AI systems may be mistaken for reliability. A chatbot can generate a plausible response even when the answer is incomplete, misleading, or contextually inappropriate. In libraries, where users often require trustworthy guidance rather than generic interaction, this distinction matters greatly. Automated systems are most defensible when they are clearly bounded, transparent about their role, and able to transfer complex or uncertain questions to human staff.

These observations lead to the central proposal of this paper: a trustworthy sociotechnical framework for AI in academic libraries. The framework is built around six dimensions.

The first dimension is task-risk alignment. Libraries should not treat all AI use cases as equally suitable for automation. Some tasks are low risk, such as draft

suggestions for internal processing or basic triage of routine service requests. Others are medium risk, including subject suggestion, recommendation, and first-line user support. Higher-risk tasks include outputs that substantially affect representation, retrieval, or user decision-making without meaningful review. In such cases, the potential harm of error is much greater. Trustworthy adoption therefore requires that the level of automation be proportionate to the stakes of the task.

The second dimension is data and metadata stewardship. AI systems are shaped by their training data, descriptive vocabularies, and metadata environments. Libraries must therefore ask what data are being used, whose knowledge is visible in those data, which languages are represented, and how the provenance of machine-generated output can be traced. This is especially important in India, where linguistic diversity and knowledge heterogeneity are major realities. A system trained primarily on English-language records or dominant disciplinary materials may perform poorly in regional language contexts or on local knowledge collections. Data quality is not a purely technical concern. It is central to fair and meaningful access.

The third dimension is human accountability. AI in libraries should augment professional work rather than displace responsibility. Human accountability means more than a nominal review step. It requires identifiable roles for validation, correction, policy setting, and exception handling. If a machine-generated subject term is misleading or a chatbot provides incorrect guidance, responsibility for correction and system improvement must remain clearly located within the institution. This is essential to professional integrity and to user trust.

The fourth dimension is transparency and contestability. Libraries must be open about when AI is being used, what kind of outputs it produces, and how those outputs can be questioned or corrected. In technical services, transparency may include marking machine-assisted fields or documenting workflow decisions. In user services, it may involve clear disclosure that a user is interacting with an automated system, together with visible routes to

human support. Contestability is equally important. Users and staff need mechanisms to challenge outputs and seek correction. Without such mechanisms, AI systems become difficult to govern and easy to normalize.

The fifth dimension is user rights and inclusion. Library values require sustained attention to privacy, accessibility, fairness, and equitable treatment. AI systems that depend on extensive data tracking may conflict with long-established library commitments to reader privacy. Recommendation systems may narrow exposure if they over-personalize. Conversational tools may exclude users if they are not designed for multilingual and accessible interaction. In Indian academic libraries, inclusion also requires sensitivity to regional languages, first-generation learners, and unequal levels of digital literacy. Trustworthy AI must therefore be designed with user diversity in mind rather than assuming a uniform user profile.

The sixth dimension is organizational readiness. Even promising tools can fail when institutions lack policy, training, documentation, or infrastructural stability. Ajani et al. (2022) show that awareness of AI does not necessarily translate into readiness. This point is highly relevant in India, where academic libraries vary significantly in funding, technical expertise, staffing patterns, and digital maturity. A trustworthy framework must therefore include staff development, procurement literacy, institutional policy, and regular evaluation. Adoption without readiness can create superficial modernization while weakening actual service quality and governance.

## V. DISCUSSION

The proposed framework has several implications for academic librarianship in India. The first is that AI should not be approached as a shortcut to modernization. It is often presented as a way to make libraries "smart," but that language can obscure the real issue. The core question is not whether a library uses AI, but whether its use improves access and service in a manner consistent with professional values. In this sense, trustworthy adoption is more important than rapid adoption.

The second implication is that classification and cataloguing should remain central to the AI debate. Much public discussion of AI in libraries focuses on chatbots and user-facing tools because they are easy to demonstrate. However, the deeper transformation may occur in the technical infrastructure of knowledge organization. Machine-generated metadata, automated subject assignment, and AI-assisted discovery can quietly reshape access long before users realize it. For that reason, technical services should not be treated as a secondary concern in AI policy. They are among the main sites where trust is built or lost.

The third implication concerns the professional role of librarians. The literature reviewed in this paper does not support the idea that AI will make librarians redundant. Instead, it suggests a shift in expertise. Librarians may spend less time on some repetitive tasks and more time on validation, system interpretation, policy formation, user education, and ethical oversight. This is not a minor adjustment. It implies the need for new forms of professional development, especially in metadata auditing, conversational system evaluation, procurement review, and AI literacy. In India, where capacity-building is often uneven, this need is especially pressing.

The fourth implication is that Indian academic libraries require context-sensitive frameworks rather than imported technological models. Many AI tools are designed in environments with different assumptions about language, infrastructure, user behavior, and regulatory norms. A system that works adequately in a narrowly English-language, vendor-rich environment may not perform well in multilingual university settings with diverse disciplinary collections and variable digital access. Local adaptation is therefore not optional. It is a condition of trustworthy use.

Finally, the framework highlights a deeper continuity between AI governance and the ethical foundations of library science. Ranganathan's laws remain surprisingly relevant in this discussion. If books are for use, then AI should improve meaningful use rather than merely automate internal processes. If

every reader has his or her book and every book its reader, then classification and discovery systems must avoid reinforcing exclusion. If the library is a growing organism, then technological change must be guided by values that evolve without abandoning the public mission of the institution. Seen in this way, trustworthy AI is not an external add-on to librarianship. It is an extension of core library commitments into a new technical environment.

## **VI. CONCLUSION AND FUTURE WORK**

This conceptual paper has examined the application of AI and ML in library classification, cataloguing, and user services, with a particular focus on Indian academic libraries. It has argued that these technologies should not be treated merely as tools for efficiency, but as sociotechnical systems that shape representation, retrieval, mediation, and trust. The literature demonstrates clear opportunities in automated subject assignment, metadata enrichment, and virtual user support. At the same time, it reveals substantial concerns related to bias, opacity, privacy, weakened accountability, and uneven organizational readiness.

To address these tensions, the paper proposes a trustworthy sociotechnical framework based on six dimensions: task-risk alignment, data and metadata stewardship, human accountability, transparency and contestability, user rights and inclusion, and organizational readiness. The framework is intended to guide conceptual and practical evaluation rather than serve as a rigid checklist. Its main contribution is to connect technical innovation with library values in a way that is especially relevant to the Indian context.

Future work should test and refine this framework through empirical studies in Indian universities and research institutions. Comparative case studies could examine AI-assisted cataloguing in multilingual settings, chatbot deployment in academic reference services, and the effects of recommendation systems on discovery behavior. Further research is also needed on procurement ethics, institutional policy, and the role of librarians in AI literacy initiatives. As academic libraries continue to evolve, AI is likely to

become more visible in both infrastructure and service design. The challenge is not whether libraries will engage with AI. The challenge is whether they will do so in ways that remain trustworthy, inclusive, and professionally accountable.

## REFERENCES

1. Ajani, Y. A., Tella, A., Salawu, K. Y., & Abdullahi, F. (2022). Perspectives of librarians on awareness and readiness of academic libraries to integrate artificial intelligence for library operations and services in Nigeria. *Internet Reference Services Quarterly*, 26(4), 213-230.
2. Asemi, A., Ko, A., & Nowkarizi, M. (2021). Intelligent libraries: A review on expert systems, artificial intelligence, and robot. *Library Hi Tech*, 39(2), 412-434.
3. Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.
4. Bradley, F. (2022). Representation of libraries in artificial intelligence regulations and implications for ethics and practice. *Journal of the Australian Library and Information Association*, 71(3), 189-200.
5. Cordell, R. C. (2020). *Machine learning + libraries: A report on the state of the field*. LC Labs, Library of Congress.
6. Cox, A. (2022). The ethics of AI for information professionals: Eight scenarios. *Journal of the Australian Library and Information Association*, 71(3), 201-214.
7. Gonzales, B. M. (2014). Linking libraries to the web: Linked data and the future of the bibliographic record. *Information Technology and Libraries*, 33(4), 10-22.
8. Hervieux, S., & Wheatley, A. (Eds.). (2022). *The rise of AI: Implications and applications of artificial intelligence in academic libraries*. Association of College and Research Libraries.
9. International Federation of Library Associations and Institutions. (2020). *IFLA statement on libraries and artificial intelligence*. IFLA.
10. Kane, D. (2019). Creating, managing and analyzing an academic library chatbot. *BiD: Textos Universitaris de Biblioteconomia i Documentacio*, 43.
11. Kragelj, M., & Kljajic Borstnar, M. (2021). Automatic classification of older electronic texts into the Universal Decimal Classification-UDC. *Journal of Documentation*, 77(3), 755-776.
12. Library of Congress. (2016). *Overview of the BIBFRAME 2.0 model*. Library of Congress.
13. McNeal, M. L., & Newyear, D. (2013). Introducing chatbots in libraries. *Library Technology Reports*, 49(8), 5-10.
14. Mckie, I. A. S., & Narayan, B. (2019). Enhancing the academic library experience with chatbots: An exploration of research and implications for practice. *Journal of the Australian Library and Information Association*, 68(3), 268-277.
15. Ministry of Education. (2020). *National education policy 2020*. Government of India.
16. Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. New York University Press.
17. Olson, H. A. (2002). *The power to name: Locating the limits of subject representation in libraries*. Kluwer Academic.
18. Padilla, T. (2019). *Responsible operations: Data science, machine learning, and AI in libraries*. OCLC Research.
19. Ranganathan, S. R. (1931). *The five laws of library science*. Madras Library Association.
20. Rodriguez, S., & Mune, C. (2022). Uncoding library chatbots: Deploying a new virtual reference tool at the San Jose State University Library. *Reference Services Review*, 50(3/4), 392-405.
21. UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. UNESCO.