



Data curation in academic libraries: navigating the AI labyrinth through the FATE principles

Nuno Sousa

DOI: <https://doi.org/10.47989/ir31260449>

Abstract

Introduction. This study examines how academic libraries address the ethical and operational challenges of data curation in an era shaped by artificial intelligence (AI). It focuses on how the FATE principles (fairness, accountability, transparency, and ethics) are interpreted and embedded within AI-driven curatorial practices.

Method. A systematic literature review was conducted, analysing fifty-one publications. The review reflexively applied the FATE principles to its own design to ensure methodological integrity, transparency, and reproducibility.

Analysis. Both conceptual and empirical dimensions were examined to identify how each FATE principle is operationalised across AI-enabled data workflows. Comparative analysis assessed the procedural maturity of these principles throughout the data lifecycle.

Results. The findings reveal a marked imbalance: accountability and transparency demonstrate procedural consolidation through documentation and audit mechanisms, whereas fairness and ethics remain conceptually diffuse and empirically underdeveloped. Academic libraries emerge as ethical and algorithmic infrastructures mediating between technological innovation and epistemic justice.

Conclusion. Responsible AI in librarianship requires structural integration of FATE into governance, documentation, and daily curatorial workflows. The proposed FATE Curation Framework embeds ethical checkpoints across the data lifecycle, offering a replicable model for institutional implementation.

Introduction

Artificial intelligence (AI) is no longer peripheral to academic librarianship; it underpins discovery systems, metadata generation, analytics, and conversational interfaces that mediate how scholarship is produced, disseminated, and evaluated. Within this evolving landscape, data curation, the systematic management and enrichment of data across its lifecycle has emerged as a central ethical frontier. Decisions made in data curation are not merely technical but epistemic and moral, shaping what knowledge becomes visible, whose voices are amplified, and which communities are marginalised. The FATE principles, fairness, accountability, transparency, and ethics (see Memarian & Dolek, 2023), provide a foundational framework for aligning AI with public-interest values. However, their translation into the concrete workflows, infrastructures, and policies of academic libraries remains underdeveloped and inconsistent.

This study addresses that gap through a systematic literature review covering publications between 2020 and 2025, a period marked by the rise of generative AI and growing scrutiny of algorithmic governance in research and higher education. The review treats academic libraries not as passive users of technology but as ethical infrastructures that mediate between researchers, vendors, and the public. It assumes that responsible AI cannot be retrofitted post-deployment; instead, ethical integrity must be designed into curatorial workflows, documentation processes, and governance structures from the outset.

The central research question guiding this study is: How, and with what consequences, are the FATE principles currently interpreted and operationalised within data curation workflows in academic libraries? This question unfolds into subsidiary inquiries: which curatorial stages, selection, description, preservation, or access, are most conducive to FATE integration?; what artefacts and tools, such as datasheets, model cards, and algorithmic audits, are being employed?; and where do accountability and transparency appear to advance more robustly than fairness and ethics, and why? Positioned at the intersection of library and information science, artificial intelligence ethics, and data governance studies, this article adopts a reflexive and normative stance. Rather than treating AI as a neutral technological enhancement, it examines algorithmic systems as socio-technical infrastructures that actively shape knowledge organisation, access, and authority within academic libraries.

The objectives of this work are fivefold. First, to map the conceptual and methodological foundations linking data curation and AI in academic libraries. Second, to identify how FATE principles are instantiated across the data curation lifecycle, distinguishing rhetorical commitments from operational practice. Third, to expose recurring barriers such as vendor opacity, fragmented responsibility, and inadequate policy alignment. Fourth, to synthesise these insights into a FATE curation framework that specifies practical artefacts and ethical checkpoints for library operations. Fifth, to propose a research and policy agenda that promotes empirically grounded, auditable, and human-centred AI ethics in librarianship. These objectives structure the article as follows: the next section outlines the methodological design of the systematic literature review; the state-of-the-art section critically synthesises conceptual, technological, and ethical strands of the literature; the results section maps dominant trends and gaps; and the discussion advances an integrative FATE curation framework, followed by implications for research and practice.

Context

The problem facing libraries today is multidimensional. Procedural governance has advanced more rapidly than substantive equity: while transparency mechanisms and auditability protocols are gaining traction, fairness and ethics remain conceptually vague and empirically underexamined. Responsibility is distributed across vendors, IT departments, and library professionals, creating gaps in accountability and in the capacity for ethical redress when

algorithmic systems fail or bias emerges. Furthermore, explanation and interpretability often remain inaccessible to practitioners and users, limiting their ability to exercise agency or informed consent in their interactions with automated systems. These challenges reflect deeper tensions between efficiency and justice, automation and stewardship, and innovation and accountability.

A review of the literature reveals significant gaps. Few studies provide empirical evaluations of FATE in practice within library contexts. There is an absence of tailored ethical audit methodologies capable of evaluating bias, provenance, and representational equity in academic data repositories. Fairness metrics are often borrowed from technical AI fields and lack calibration to scholarly or cultural heritage contexts. Explanatory systems, such as human-centred interfaces for cataloguing or recommendation algorithms, are rarely evaluated in terms of user comprehension or trust. Moreover, documentation and provenance practices remain inconsistent, with limited access to training data, model architectures, or change logs, constraining transparency and reproducibility in academic library settings.

This article contributes to the field in four key ways. It offers a systematic mapping of how FATE has been conceptualised and applied within AI-enabled data curation in academic libraries; it identifies the asymmetrical operationalisation of FATE, showing the procedural dominance of accountability and transparency over fairness and ethics; it introduces a conceptual FATE curation framework that embeds ethical checkpoints across data lifecycle stages; and it proposes a research agenda that prioritises domain-specific ethical metrics, participatory audits, and measurable outcomes for responsible AI in librarianship. By doing so, it advances a scholarly foundation for ethical governance that is evidence-based, actionable, and sustainable.

The study focuses specifically on academic and research libraries, analysing contexts where AI influences selection, organisation, preservation, or access to data. It synthesises literature published between 2020 and 2025 that explicitly or implicitly engages with the FATE principles. While policy and regulatory frameworks such as the EU AI Act or UNESCO's AI Ethics Recommendation provide contextual background, the analysis remains centred on library implementations and evaluative practices. Methodological details, including database selection, screening processes, and PRISMA-based reporting, are presented in the Methodology section.

Two propositions structure this review: first, that accountability and transparency mechanisms are more frequently operationalised than fairness and ethics in library AI deployments; and second, that empirical evidence demonstrating the practical impact of FATE principles remains limited relative to conceptual and policy discourse. These propositions guided the synthesis and evaluation of the reviewed literature. The structure of the article proceeds accordingly: following the methodology, the state of the art explores the conceptual intersections between AI, data curation, and ethics, tracing the evolution of FATE in academic information environments. The results present temporal and thematic trends, leading to a conceptual synthesis in the form of the FATE curation framework. The discussion then interprets the findings considering theoretical and practical implications, concluding with a forward-looking research and policy agenda.

Ultimately, this study argues that in the era of AI-mediated librarianship, governance is inseparable from design. Libraries will uphold epistemic justice and institutional integrity not by appending ethical guidelines after technological adoption, but by embedding fairness, accountability, transparency, and ethics directly into the processes and infrastructures through which knowledge itself is curated.

Methodology

This study employed a systematic literature review conducted in October 2025 to examine how the FATE principles, are integrated into data curation practices within academic libraries. The review followed the PRISMA 2020 guidelines to ensure methodological transparency, reproducibility, and rigor.

The search strategy was developed iteratively with the support of Scopus AI tools, which assisted in refining the research question, objectives, hypotheses, and search syntax. The final search expression combined operational, ethical, and institutional dimensions of data curation to capture literature that explicitly or implicitly engaged with FATE-related themes. The complete expression was as follows:

“data curation” OR “data management” OR “data preservation” OR “information stewardship”) AND (“academic library” OR “university library” OR “research library” OR “library science”) AND (“AI” OR “artificial intelligence” OR “machine learning” OR “algorithm”) AND (“FATE” OR “fairness” OR “accountability” OR “transparency”) AND (“principles” OR “guidelines” OR “standards” OR “ethics”)

This search was applied to two major databases. In Web of Science, the search was limited to the years 2020–2025, filtered by review articles and the research area Library and Information Science, resulting in 533 records. In Scopus, the same temporal filter (2020–2025) and article type (review) were applied, yielding forty-five records. The database Dimensions.ai was considered but excluded after initial testing revealed that its results were misaligned with the conceptual and disciplinary scope of this investigation.

All records were imported into Rayyan, an online tool designed to facilitate systematic literature reviews. Duplicate detection led to the removal of eight records, after which a structured screening process was carried out according to predefined inclusion and exclusion criteria. The inclusion criteria focused on studies that addressed the intersection of data curation, academic libraries, and AI; incorporated one or more of the FATE principles; or provided conceptual or methodological insights into governance and accountability in data environments. Exclusion criteria removed studies outside the context of academic or research libraries, without reference to AI or ethical frameworks, or lacking methodological relevance. Following this process, 522 out of 570 initial records were excluded, leaving forty-eight studies for full review. In addition, three highly relevant works identified through open Web searches were incorporated due to their conceptual importance to the topic, producing a final sample of fifty-one publications. The entire selection process is summarized in the PRISMA flow diagram (Figure 1).

Each selected article was examined to identify how the FATE principles were represented or operationalized in data curation processes. The analysis considered the relationship between AI-based systems and curatorial workflows; the degree of ethical integration in policy and technical design; and the presence of accountability and transparency mechanisms. A thematic coding approach was used to synthesize findings, linking recurrent concepts to the four FATE principles and mapping them across curatorial stages such as selection, description, preservation, and access. This structure ensured coherence between the research question, evidence, and analytical interpretation.

The methodological design of this study was crucial to ensure comprehensive coverage and ethical robustness. By combining PRISMA standards, systematic screening, and transparent inclusion-exclusion criteria, the process established a strong foundation for the results and discussion that follow. It reinforced the credibility, replicability, and validity of the findings, while allowing the identification of conceptual gaps and ethical tensions that inform the future integration of AI within academic libraries. Methodological rigor here becomes both a means and

a message: the research itself reflects the principles it studies, fairness, accountability, and transparency, applied to the act of inquiry.

AI tools were used in a limited and fully transparent manner, specifically to support grammatical revision, ensure linguistic consistency, and create illustrative figures (4-5; 7-9). No AI systems contributed to data interpretation or argument development. This selective and documented use was undertaken to maintain scientific integrity, authorship accountability, and methodological legitimacy throughout the entire research process.

Importantly, the use of artificial intelligence in this study was strictly limited to supportive and procedural functions, such as search optimisation and screening facilitation. All analytical decisions, thematic interpretations, and theoretical syntheses were conducted by the authors, thereby ensuring that epistemic authority, critical judgement, and ethical responsibility remained entirely human-driven.

Figure 1 visualises the sequential filtering logic and decision-making process that structured the construction of the final evidential corpus used in this review.




Figure 1. Prisma flux diagram. Source: Based on PRISMA 2020

State of the art

Conceptual foundations of data curation in academic libraries

In academic settings, data curation has shifted from a narrow, repository-centred activity to a strategic, institution-wide function that spans the full research lifecycle and interfaces with algorithmic systems. Early emphases on file-level management and compliance have given way to stewardship models that integrate policy, infrastructure, training, and ethical governance. Recent syntheses in library and information science confirm this widening scope, from supporting researchers' deposit and documentation practices to shaping the quality, provenance, and responsible reuse of research data within intelligent services (Sheikh et al., 2025; Stvilia et al., 2025; Zareef & Jabeen, 2025). This evolution is inseparable from the embedding of AI in scholarly communication ecosystems, recommendation, discovery, summarisation, and conversational support, thereby relocating curation from a post-hoc preservation task to an upstream, design-time determinant of how knowledge is produced, surfaced, and evaluated (Asemi et al., 2021; Afjal, 2023; Buitrago-Ciro et al., 2025).

Research data management (RDM) concerns policies, planning, documentation, storage, sharing, and compliance at project level; it is primarily operational and researcher-facing (Sheikh et al., 2025). Digital preservation ensures long-term accessibility, authenticity, and integrity through format strategies, fixity, and succession planning; it is archival and time-oriented (Stvilia et al., 2025). Data curation sits across and beyond both: it is the continuous, value-adding process that selects, normalises, enriches, validates, documents, and exposes data for discovery and reuse, aligning technical processes with scholarly meaning and policy goals. In libraries, curation is best understood as a lifecycle that coordinates selection, ingest, description, quality assurance, preservation, and access, often formalised through curation lifecycle models operationalised in library contexts (Stvilia et al., 2025; Zareef & Jabeen, 2025). Clarifying these distinctions matters: research data management can be compliant yet shallow; preservation can be durable yet opaque; curation is the connective tissue that renders data fit for purpose in human and algorithmic workflows.

The lifecycle perspective remains central to understanding data curation in academic libraries. Among the most widely adopted frameworks, the Digital Curation Centre (DCC) model (Figure 2) articulates the iterative processes through which data are conceptualised, created, appraised, curated, preserved, and reused. It emphasises that curation is not a terminal activity but a continuous stewardship function that integrates preservation planning, community participation, and access governance across time. While the Digital Curation Centre has played a foundational role in structuring curatorial thinking within libraries, it remains largely epistemically and ethically neutral. As such, it offers limited guidance for addressing algorithmic bias, accountability, or transparency in AI-mediated data environments.




Figure 2. Curation Lifecycle model.

Source: <https://dcc.ac.uk/guidance/curation-lifecycle-model>

This model underpins much of the contemporary discourse on data stewardship. It provides the structural foundation upon which newer frameworks, such as FAIR principles and FATE, build to integrate ethical, algorithmic, and accountability dimensions into curatorial practice.

Academic libraries increasingly function as institutional mediators that translate between the practices of researchers, the constraints of platforms, and the expectations of publics. They convene policy (e.g., governance, licensing), build capability (data and AI literacy), and align infrastructure (repositories, knowledge graphs, APIs) with disciplinary norms. Evidence from studies of higher education libraries and organisational adoption work highlight how libraries broker cloud and AI technologies while maintaining service reliability and compliance (Ibrahim et al., 2025). On the human-factors side, libraries are re-positioning staff as data stewards and AI-literate professionals, with emerging guidance on staff competencies and patron education (Ali & Richardson, 2025; Li et al., 2025; Tang, 2025). In short, libraries are no longer only endpoints for data; they shape the conditions of possibility for responsible data-intensive scholarship.

Because curated data feed algorithmic systems, curation choices are epistemic, they structure what can be known, and ethical, they distribute risks and benefits. The FATE principles render this dual nature explicit. Fairness requires representational adequacy and bias mitigation in selection and description; accountability demands governance, audit trails, and answerability for curatorial and algorithmic decisions; transparency calls for documentation that is comprehensible to both experts and affected users; and ethics frames proportionality, consent, privacy, and legitimate interests across the lifecycle. Contemporary reviews in AI ethics and AI auditing show that operationalising these principles hinges on concrete artefacts

(documentation, metrics, oversight processes) rather than declarations (Laine et al., 2024; Matei et al., 2025; Memarian & Doleck, 2023).

Within library and information science (LIS), the proliferation of AI-enabled discovery and conversational systems amplifies these stakes: system behaviour can be persuasive, error-prone, or strategically opaque, which raises questions of truthfulness, explainability, and responsible deployment in scholarly mediation (Black, 2024; Buetow & Lovatt, 2024; Kousha & Thelwall, 2024; Matsieli & Mutula, 2025; Njiru et al., 2025; Ridley, 2025; Sparkman & Witt, 2025). Positioning curation as epistemology-in-practice aligns library work with FATE not as an add-on but as a design constraint for data pipelines and intelligent services.

Taken together, these strands justify treating data curation in academic libraries as a socio-technical, lifecycle-based stewardship that is analytically separable from research data management and digital preservation, while remaining interdependent with both. As AI systems mediate scholarly discovery and evaluation, curation becomes the primary locus where FATE is made operational, through selection policies, metadata schemas, documentation standards, quality controls, and governance structures that render data not only findable and durable but just, answerable, and intelligible in algorithmic contexts.

The FAIR principles have guided much of the discourse around data stewardship in academic libraries. However, their conceptualisation often remains technical rather than ethical. The visual metaphor presented below encapsulates the layered and processual nature of FAIR-compliant data management, serving as a conceptual bridge toward FATE-oriented approaches.



Figure 3. Visual metaphor illustrating the layered structure of FAIR data management

Source: FAIRWizard

Although the FAIR principles provide a robust technical framework for data interoperability and reuse, they do not explicitly address questions of power, bias, responsibility, or ethical governance. This limitation becomes particularly salient in AI-driven contexts, where automated decision-making systems actively mediate access to knowledge, thereby necessitating an explicit ethical lens such as FATE.

The conceptual trajectory traced above can be visually synthesised through a temporal model that illustrates the progressive integration of ethical and algorithmic stewardship into the curatorial mission of academic libraries. Figure 4 (below) captures this evolution from Data Management to AI-Enabled Stewardship between 2020 and 2025, highlighting how the FATE principles have successively permeated the core functions of data practice.




Figure 4. The evolution of data curation functions in academic libraries

Source: Author's data elaborated with artificial intelligence

This visualisation reinforces the argument that data curation has not evolved in isolation but in tandem with the growing ethical sophistication of academic information systems. The transition from operational management to value-driven stewardship underscores that contemporary curation operates simultaneously as a technical, epistemic, and moral enterprise: the site where institutional accountability meets algorithmic design.

The rise of algorithmic systems in library and information science

Over the last five years, academic libraries have moved from experimental pilots to the systematic embedding of AI across cataloguing, subject classification, discovery, and recommendation. Intelligent library initiatives integrate rule-based and machine-learning approaches for metadata enrichment, entity recognition, and automated classification (Asemi et al., 2021). Discovery layers and recommender services increasingly rely on behavioural and semantic models to rank and surface resources, while conversational agents, from general large language models (LLM) to domain-tuned chatbots, mediate reference services, search refinement, and instruction at scale (Aboelmaged et al., 2025; Afjal, 2023; Khosrawi-Rad et al., 2023). In parallel, AI augments the scholarly communication pipeline, screening, peer review support, summarisation, and integrity checks, reconfiguring the informational environment in which libraries operate (Buitrago-Ciro et al., 2025; Carabantes et al., 2023; Kousha & Thelwall, 2024;). These trajectories are consistent with broader information science and user experience evidence: AI is becoming infrastructural, shifting from peripheral add-ons to core components of information services (Collins et al., 2021; Li et al., 2025; Stige et al., 2024).

Automation unsettles the library's traditional stance of impartial mediation. Algorithmic ranking and classification encode value judgements, about relevance, authority, and visibility, that were once explicit, negotiated, and attributable to human cataloguers. As services become AI-by-default, the locus of decision-making moves upstream into data pipelines and models, potentially obscuring the deliberative practices that sustained professional neutrality (Ibrahim et al., 2025; Mitha & Omarsaib, 2025). Evidence from library and higher education contexts shows that adoption drivers (efficiency, scale, responsiveness) collide with the need for professional oversight, contestability, and pedagogical intent, demanding new literacies and role redesign rather than simple tool substitution (Ali & Richardson, 2025; Romero & Chagua, 2025; Thong et al., 2025). In short, mediation is not disappearing; it is being reconstituted, from front-stage interaction to the design and governance of algorithmic systems.

As curated data feed the very systems that organise scholarly knowledge, curation practices now bear epistemic consequences: what is selected, described, and linked shapes what becomes discoverable and credible. Reviews on human-centred explainability and data-quality assurance

emphasise that trust is not a property of models alone but of the end-to-end socio-technical assemblage, documentation, provenance, validation, and communicability to stakeholders (Ridley, 2025; Stvilia et al., 2025). In reference and instructional contexts, LLM-mediated interactions can be useful yet performatively confident, pressuring librarians to establish verifiable pathways, boundaries of use, and escalation protocols (Buetow & Lovatt, 2024; Sparkman & Witt, 2025). Ethical reflections from the library field further argue that the credibility of library services in an AI-saturated ecosystem depends on explicit norms for responsible deployment and on cultivating AI literacy among staff and patrons (Ali & Richardson, 2025; Matsieli & Mutula, 2025; Tang, 2025).

The expansion of AI amplifies a familiar risk set. Opacity arises from proprietary models, complex pipelines, and the entanglement of third-party services, limiting auditability and user comprehension (Laine et al., 2024). Bias emerges through historical collections, skewed training data, and feedback loops in usage logs, with downstream effects on the visibility of languages, topics, and communities (Memarian & Doleck, 2023; Ofosu-Ampong, 2024). Explainability remains uneven: local, human-centred explanations are often absent or unusable for library stakeholders (Ridley, 2025), while accountability diffuses across vendors, platforms, and institutions, complicating redress when harms occur (Ashok et al., 2022; Heyder et al., 2023). Critical scholarship has even questioned truthfulness and strategic deception in AI outputs, underscoring the need for verifiable evidence paths in scholarly mediation (Black, 2024; Njiru et al., 2025). Together, these findings reinforce the necessity of FATE-aligned governance: bias assessment and representational checks (Fairness), clear responsibility and escalation chains (Accountability), actionable documentation and model reporting (Transparency), and proportionate, context-sensitive deployment grounded in professional ethics (Ethics).

To elucidate this convergence between functionality and ethical oversight, Table 1 maps the main domains of AI application within academic libraries and aligns them with their corresponding FATE-related risks and mitigation mechanisms. The comparative structure exposes the dual nature of algorithmic integration: while it augments efficiency, personalisation, and scalability, it simultaneously generates new vectors of opacity, bias, and ethical vulnerability.

Taken together, the analysis suggests that ethical risk intensifies in proportion to the epistemic mediation exercised by algorithmic systems. Applications such as discovery algorithms and conversational agents emerge as high-risk zones, as they directly shape information visibility, relevance, and user interpretation while often operating under conditions of limited transparency.

To synthesise these dynamics, Table 1 provides a comparative overview of the main AI functions currently adopted in library operations, mapping their primary benefits against associated ethical risks and corresponding governance mechanisms. This structured view clarifies how each AI-driven functionality simultaneously advances service performance and introduces vulnerabilities that require deliberate institutional oversight and human-algorithmic co-accountability.

AI function	Primary benefits	Associated FATE risks	Mitigation & governance mechanisms	Ethical risk level
Automated cataloguing	Increases metadata consistency; reduces human workload	Bias in controlled vocabularies; opacity in rule-based classification	Metadata audits; bias-detection tools; staff retraining	Moderate
Discovery and retrieval algorithms	Enhances search precision and relevance	Filter bubbles; lack of transparency in ranking algorithms	Algorithmic explainability; open metadata schemas	High
Recommendation systems	Personalised user experience; promotes engagement	Reinforcement of bias; loss of serendipity	User control panels; ethical recommender design	Moderate
Chatbots and virtual assistants	Improves accessibility; 24/7 service	Ethical drift; misinformation; lack of accountability	Human-in-the-loop oversight; audit logs	High
Predictive analytics for collection management	Optimises acquisitions and usage prediction	Data misuse; lack of consent; privacy risks	Data anonymisation; institutional ethics review	Low Moderate
AI-assisted data curation	Supports quality control and preservation workflows	Unclear provenance; accountability gaps	Provenance tracking systems; transparent workflow documentation	Low

Table 1. AI applications in Library operations and corresponding ethical tensions

Source: Author's elaboration based on systematic literature review

Ethical and governance challenges in AI-driven data curation

Within LIS, the ethical stakes of data curation are no longer peripheral. As curated data now drive intelligent services, curation choices have epistemic and moral consequences, shaping what can be known, by whom, and under which conditions. Recent syntheses emphasise that credible, human-centred information services require explicit alignment between design, documentation, and stakeholder needs (Li et al., 2025; Stvilia et al., 2025). Cross-domain reviews show that FATE is the most practical lens for translating broad AI ethics into operational requirements for higher-education and library contexts (Ashok et al., 2022; Memarian & Doleck, 2023). In parallel, conceptual work on university data culture argues that institutions must articulate values and norms that travel consistently from policy to pipeline to interface (Tang, 2025).

AI adoption promises efficiency, scale, speed, and responsiveness across classification, discovery, and service triage (Asemi et al., 2021; Collins et al., 2021), but the literature repeatedly documents tensions with public-interest values central to academic libraries. Recommenders and LLM-based agents can entrench representational imbalances, privileging dominant languages, topics, or venues, thereby undermining equity and inclusion in scholarly visibility (Memarian & Doleck, 2023; Ofosu-Ampong, 2024). Accessibility is similarly at stake: when interfaces are optimised for engagement rather than intelligibility, users with different literacies or access needs may be disadvantaged (Stige et al., 2024). The emerging response within the sector foregrounds AI literacy, for staff and patrons, as a precondition for responsible deployment and for sustaining trust in mediated services (Ali & Richardson, 2025; Mitha & Omarsaib, 2025).

Converging strands of evidence point to governance, not tooling, as the decisive variable. Reviews on ethics-based AI auditing propose translating principles into auditable artefacts: risk registers, model cards, datasheets, issue-tracking, and escalation protocols mapped to

accountable roles (Laine et al., 2024). Broader information-management frameworks stress integrating ethics by design into procurement, vendor management, and lifecycle controls (Ashok et al., 2022). Within LIS, studies highlight the importance of quality assurance, provenance, and documentation in repositories and services as the substrate for any trustworthy AI layer (Stvilia et al., 2025; Zareef & Jabeen, 2025). Organisational adoption research in university libraries further shows that cloud/AI decisions hinge on institutional capability, governance maturity, and human-factor readiness, underscoring the need for clear responsibility chains and competence development (Ibrahim et al., 2025; Ramirez et al., 2022).

The record is mixed. Conversational agents improve reach and responsiveness, yet raise concerns about privacy, misdirection, and boundary-of-competence signalling (Aboelmaged et al., 2025). Generative tools assist summarisation and literature triage, but invite over-reliance, untraceable claims, and goal-shift from evidence to fluency (Afjal, 2023; Buetow & Lovatt, 2024; Sparkman & Witt, 2025). In scholarly communication, AI-enabled peer-review support and screening can streamline workflows while complicating accountability and authorship norms (Carabantes et al., 2023; Kousha & Thelwall, 2024). Studies of user profiling warn of opaque inferences and disproportionate data capture with unclear redress, directly engaging FATE fairness and accountability planks (Njiru et al., 2025). Adoption of robotic and automated service components surfaces labour, safety, and legitimacy questions that libraries are not yet institutionally equipped to govern at scale (Shahzad et al., 2024). At the conceptual core lie opacity, bias, explainability, and diffused responsibility: human-centred XAI (explainable artificial intelligence) remains unevenly available in library-facing tools (Ridley, 2025), and philosophical critiques underscore that systems can produce persuasive falsehoods with high confidence, demanding verifiable evidence paths in scholarly mediation (Black, 2024; Buitrago-Ciro et al., 2025).

Taken together, these strands motivate a governance posture in which fairness is operationalised via representational audits and impact assessment; accountability via role-bound decision logs, vendor contracts, and incident response; transparency via public-facing documentation, data provenance, and explanation mechanisms suitable for library stakeholders; and ethics via proportionality, consent, privacy-by-design, and professional standards embedded across the curation lifecycle. In AI-driven data curation, governance is not ancillary to service quality; it is the mechanism by which libraries reconcile efficiency with public interest values, sustaining informational trust in an algorithmic environment.

To consolidate these findings and translate them into an actionable governance structure, Figure 5 illustrates the sequential embedding of the FATE principles across institutional workflows in AI-driven data curation. It depicts how FATE can be operationalised at each decision point, from policy formulation to data ingestion, metadata enrichment, and access management, thereby transforming governance from a reactive control mechanism into an active design principle.

Rather than serving as a descriptive overview of existing practices, the following figure should be read as a prescriptive governance scaffold. It delineates critical intervention points at which FATE principles can be institutionally embedded across the AI-enabled data curation workflow.




Figure 5. Governance pathways for FATE integration in AI-driven data curation

Source: Author's data elaborated with Artificial Intelligence

This model provides a conceptual bridge to the next section, which examines in greater depth the theoretical origins, interpretations, and operational applications of the FATE framework within academic librarianship.

The FATE principles: origins, interpretations, and applications

FATE emerged as a pragmatic distillation within AI ethics to bridge abstract normative claims and implementable requirements in socio-technical systems. In information systems and human-AI interaction research, this cluster functions as a practice-oriented vocabulary capable of travelling across disciplines and stakeholder groups (Collins et al., 2021; Heyder et al., 2023). Higher education reviews explicitly codify FATE as the dominant lens for aligning AI deployments with institutional duties of care (Memarian & Doleck, 2023), while management scholarship frames FATE as the anchor for governance interventions, risk registers, role assignments, documentation regimes, and audit trails, that render ethical claims operational (Ashok et al., 2022; Laine et al., 2024). Within LIS, this translation is amplified by the fact that curated datasets feed intelligent services: curation choices thus carry epistemic and moral weight in how knowledge is produced, surfaced, and evaluated (Li et al., 2025; Stvilia et al., 2025).

While FATE provides an actionable ethics lens for AI systems, academic libraries operate within broader data justice traditions that emphasise community rights and collective benefit. The CARE principles (collective benefit, authority to control, responsibility, and ethics), emerged from Indigenous data sovereignty and extend beyond procedural compliance to relational accountability, equitable value sharing, and culturally grounded consent. Positioning CARE alongside FATE reframes data curation as a socio-technical practice: not only how data are made findable and reusable, but for whom, under whose authority, and to what ends. This perspective is especially salient for libraries stewarding sensitive, community-originated, or equity-relevant datasets.

By juxtaposing CARE with FATE, the study foregrounds the need for hybrid governance models that fuse technical accountability with community-centred ethics. In the sections that follow, FATE functions as the operational scaffold for AI-mediated curation, while CARE informs the distribution of authority, consent practices, and benefit allocation, together shaping a coherent approach to responsible data stewardship in academic libraries.

Building upon the preceding discussion of governance mechanisms, the following table operationalises the FATE framework within the specific context of academic libraries. It maps each principle to its functional definition, practical instantiation, and representative scholarly sources. By providing this structured overview, Table 2 transforms FATE from a normative construct into an auditable schema of practice, offering a replicable reference model for institutional adoption and evaluation.

FATE principle	Operational definition	Library application example	Representative sources
Fairness	Ensuring representational and procedural equity in data management and access workflows.	Audit of subject headings to detect cultural or gender bias; equitable metadata enrichment practices.	(Asemi et al., 2021; Memarian & Doleck, 2023; Trigo et al., 2024)
Accountability	Establishing traceable, role-bound mechanisms for human and algorithmic decision-making.	Vendor agreements including audit clauses; implementation of decision logs for AI-supported cataloguing.	(Ashok et al., 2022; Heyder et al., 2023; Laine et al., 2024)
Transparency	Providing explainability and visibility into AI processes, metadata generation, and decision criteria.	Public documentation of recommendation algorithms; creation of explainability dashboards for data discovery.	(Kousha & Thelwall, 2024; Li et al., 2025; Ridley, 2025)
Ethics	Embedding moral, social, and contextual reflection into AI-enabled library services and policies.	Ethical oversight committees for data curation projects; integration of human-in-the-loop in chatbot governance.	(Bu et al., 2025; Matei et al., 2025; Matsieli & Mutula, 2025)

Table 2. Operational mapping of FATE principles in academic library contexts

Source: Author's elaboration based on systematic literature review

Importantly, the FATE principles should not be interpreted as an isolated checklist. The literature consistently indicates their interdependence, whereby weaknesses in transparency, for instance, undermine accountability, and failures of fairness often signal deeper ethical and governance deficiencies.

Although not the focus of our corpus, widely cited policy instruments, OECD AI Principles, UNESCO Recommendation on the Ethics of AI, the EU AI Act's risk-based approach, and IEEE Ethically Aligned Design, map naturally onto FATE: non-discrimination and social equity (fairness), institutional governance and redress (accountability), documentation and intelligibility (transparency), and human-rights-grounded proportionality (ethics). Ethics-based auditing literature makes this bridge explicit by translating high-level norms into auditable artefacts (model cards, datasheets, issue logs, escalation pathways) and stakeholder responsibilities (Laine et al., 2024). In academic libraries, these mappings are mediated by sectoral policies (procurement, vendor management, privacy) and capacity-building agendas (AI literacy), which determine the enforceability of FATE beyond aspiration (Ali & Richardson, 2025; Tang, 2025).

Fairness remains contested: definitions oscillate between distributive goals (who is represented and how) and statistical criteria in model behaviour; both are sensitive to data provenance and feedback loops in usage logs (Memarian & Doleck, 2023; Ofosu-Ampong, 2024). Accountability is frequently diffused across vendors, platforms, and institutions, risking responsibility gaps unless decision rights, auditability, and redress mechanisms are contractually and procedurally

specified (Ashok et al., 2022; Heyder et al., 2023). Transparency is multi-layered, data, model, and process, and often collapses into disclosure without comprehension; human-centred explainability is necessary if documentation is to be actionable for library stakeholders (Ridley, 2025; Stvilia et al., 2025). Ethics can drift into a residual category or a signal of virtue unless anchored in proportionality, privacy-by-design, consent practices, and clear boundaries of use (Laine et al., 2024; Matei et al., 2025). Critical accounts further warn that generative systems may produce confident falsehoods or strategically misleading content, sharpening the stakes for verifiability and duty of care in scholarly mediation (Black, 2024; Buetow & Lovatt, 2024; Sparkman & Witt, 2025).

For data-intensive scholarship, FATE provides an actionable scaffold across the curation lifecycle. Fairness entails representational audits and bias checks in selection and description, especially across languages, disciplines, and communities; accountability requires role-bound decision logs, procurement clauses, incident response, and oversight boards; transparency calls for provenance-rich metadata, public-facing documentation (e.g., datasheets, model cards), and explanation mechanisms commensurate with user needs; ethics mandates proportional data capture, privacy safeguards, consent management, and clear limits on automation in instructional and reference contexts (Ali & Richardson, 2025; Sheikh et al., 2025; Stvilia et al., 2025). Sectoral evidence shows that these controls are most effective when embedded as infrastructure rather than add-ons, i.e., when governance, training, and platform choices are co-designed with curatorial workflows (Ibrahim et al., 2025; Zareef & Jabeen, 2025). In practice, conversational systems, profiling services, and automated discovery expose libraries to concrete FATE dilemmas, privacy and bias risks in user modelling, opacity in ranking, and explainability gaps at the point of need, requiring measurable safeguards and escalation pathways (Aboelmaged et al., 2025; Buitrago-Ciro et al., 2025; Kousha & Thelwall, 2024; Njiru et al., 2025).

As a boundary object, FATE is rigorous enough to guide design and audit, yet flexible enough to align librarians, vendors, administrators, and researchers around shared commitments. Its value for academic libraries lies in turning ethical intent into curatorial practice, policies, documentation, metrics, and accountability structures that make data not only durable and discoverable but also just, answerable, and intelligible in algorithmic environments.

Integrating FATE principles into data curation frameworks

In academic libraries, FATE can be rendered operational by aligning each curatorial stage with explicit safeguards and artefacts. In selection and acquisition, Fairness entails representational checks across languages, disciplines and communities; Accountability requires decision logs and role-bound sign-off; Transparency is served by public selection policies and provenance notes; Ethics demands proportionality and rights assessment for sensitive data (Sheikh et al., 2025; Stvilia et al., 2025). In ingest and normalisation, libraries implement quality gates, record lineage, and define escalation routes for exceptions, turning pipeline control into auditable practice (Stvilia et al., 2025). In description and enrichment, bias-aware subject vocabularies and authority control address representational harms; documentation of transformations sustains explainability for downstream AI (Li et al., 2025; Zareef & Jabeen, 2025). For access, discovery and recommendation, libraries can conduct exposure audits on ranking outcomes, monitor error profiles for subgroups, and provide user-facing rationales or boundaries of competence for conversational agents and recommenders (Aboelmaged et al., 2025; Ridley, 2025). Organisationally, AI literacy programmes and procurement clauses translate ethical intent into practice by conditioning deployment on documentation, monitoring, and redress (Ashok et al., 2022; Ali & Richardson, 2025).

Three friction points recur. First, vendor opacity: proprietary models and entangled third-party services limit access to training data, features, and evaluation protocols, constraining both transparency and independent audit (Laine et al., 2024). Second, diffused responsibility:

accountability is split across libraries, IT units, and suppliers; without explicit decision rights, incident response, and dispute resolution, accountability collapses into accountability theatre (Heyder et al., 2023; Ibrahim et al., 2025). Third, metric misfit: statistical fairness criteria travel poorly into cultural-heritage and scholarly discovery contexts; libraries need domain-appropriate indicators (exposure diversity, language coverage, authority balance) rather than generic parity scores (Memarian & Doleck, 2023; Ofosu-Ampong, 2024). These constraints are amplified by workload and capability pressures, where efficiency incentives risk crowding out deliberation on inclusion, accessibility, and proportionality (Mitha & Omarsaib, 2025; Stige et al., 2024).

Ethics-based auditing proposes a concrete toolkit that libraries can adopt and adapt: model cards for discovery and recommendation services (objective, inputs, known limitations, appropriate/forbidden uses), datasheets for datasets across metadata pipelines (origin, coverage, transformations, quality tests), and algorithmic audits combining offline tests (synthetic queries, counterfactual prompts) with online monitoring (exposure diversity, complaint analytics, drift) (Laine et al., 2024; Stvilia et al., 2025). For conversational systems and LLM-based services, boundary-of-competence statements, verifiable citation pathways, and escalation protocols operationalise transparency and accountability at the interface (Aboelmaged et al., 2025; Afjal, 2023; Sparkman & Witt, 2025). At the institution level, risk registers, role matrices, change logs, and post-incident reviews embed FATE into routine governance, while capability building (AI literacy for staff and patrons) sustains informed oversight (Ali & Richardson, 2025; Ashok et al., 2022; Tang, 2025).

Despite momentum, several lacunae endure. There is no widely adopted, library-specific FATE metric suite that captures the sociotechnical nature of discovery (e.g., how ranking interacts with disciplinary canons and language hierarchies); current practice relies on ad-hoc indicators with limited comparability (Stvilia et al., 2025; Zareef & Jabeen, 2025). Documentation debt persists: libraries inherit tools without sufficient model and/or dataset artefacts to meet stakeholder-oriented transparency (Laine et al., 2024; Ridley, 2025). Governance maturity is uneven, with procurement, vendor management, and policy rarely co-designed with curatorial workflows (Ibrahim et al., 2025). Finally, ethical intent often outruns operational capacity: staff time, skills and authority to contest vendor defaults remain constrained, especially where LLM-based features are marketed as turnkey solutions (Buitrago-Ciro et al., 2025; Kousha & Thelwall, 2024).

Conceptual work in LIS and adjacent fields argues for treating FATE as infrastructure, not add-on, embedding ethical controls into schemas, workflows, and platforms, and auditing them as part of routine quality assurance. Empirical and policy-oriented studies further recommend co-production with vendors and researchers to align evaluation datasets, impact metrics, and explanation modalities with scholarly use cases (Laine et al., 2024). In practice, this means, mapping FATE to the curation lifecycle; specifying minimum documentation artefacts per stage; adopting exposure and representational indicators alongside technical metrics; and establishing clear responsibility chains with escalation thresholds and redress routes. Taken together, these steps convert FATE from aspiration into a workable governance grammar for AI-mediated data curation in academic libraries.

Moving beyond theoretical articulation, the following model integrates the FATE framework into the data curation lifecycle, transforming abstract principles into operational checkpoints. Figure 7 illustrates this integration as a cyclical ecosystem where ethical reasoning and technical execution co-evolve. By embedding fairness, accountability, transparency, and ethics into each stage, from dataset selection to end-user access, the model reframes curation as a dynamic process of continuous moral calibration rather than a static technical routine.




Figure 6. FAIR curation lifecycle model

Source: Author's data elaborated with Artificial Intelligence

Figure 6 represents an original conceptual synthesis derived from the findings of the systematic literature review, integrating curatorial stages with FATE principles and associated operational artefacts.

This systemic alignment not only strengthens institutional integrity but also redefines the professional identity of academic librarians as algorithmic stewards, responsible for ensuring the epistemic and ethical legitimacy of curated knowledge infrastructures.

To capture the evolving maturity of ethical governance across AI-enabled library ecosystems, Table 3 synthesises three prevailing models of institutional response: reactive compliance, proactive stewardship, and embedded ethics. These approaches represent a progressive shift from procedural conformity to value-oriented accountability. The comparison illuminates the extent to which libraries internalise FATE principles not merely as abstract ideals but as operational norms. This transition from compliance to embedded ethics constitutes a pivotal transformation in the governance of AI-mediated data environments and provides a conceptual bridge to the subsequent mapping of FATE-curation intersections.

Governance approach	Key features	FATE alignment level	Representative case studies
Reactive compliance	Policy-driven responses to external regulations (e.g. institutional AI policies). Governance is episodic and risk-averse, with limited ethical reflexivity	 Low - Fragmented attention to fairness and transparency; ethics often retrofitted after deployment.	Early-stage automation pilots in cataloguing systems; vendor-driven AI implementations with minimal oversight (Asemi et al., 2021; Shahzad et al., 2024)
Proactive stewardship	Strategic anticipation of ethical challenges through dedicated oversight structures, internal review boards, and transparency reporting	 Moderate - Increasing accountability and procedural transparency, but uneven integration of fairness and inclusivity principles	Libraries implementing explainable recommender systems and AI literacy programmes (Ali & Richardson, 2025; Matsieli & Mutula, 2025)
Embedded ethics	Ethics-by-design approach embedded throughout AI and data curation workflows. Institutionalised accountability mechanisms, algorithmic auditing, and participatory governance	 High - Comprehensive operationalisation of all FATE principles: fairness, accountability, transparency, and ethics	Exemplary governance frameworks in open science infrastructures and data stewardship consortia (Laine et al., 2024; Ridley, 2025; Tang, 2025)

Table 3. Comparative governance models in AI-driven library contexts

Source: Author's elaboration based on systematic literature review

Mapping the intersection: data curation, AI and FATE

The increasing interdependence between AI, data curation, and ethical governance has created a dense and evolving conceptual landscape. To elucidate this complexity, Figure 7 presents a conceptual map that visualises the convergence of AI functions, curatorial processes, and the FATE principles. The visual synthesis highlights both the emerging synergies and the persistent gaps, particularly the underrepresentation of fairness and ethics within current library and information science literature.




Figure 7. Conceptual map: AI-curation FATE

Source: Author's data elaborated with Artificial Intelligence

The systematic review reveals a rapidly expanding but uneven research landscape at the intersection of data curation, AI and ethical governance between 2020 and 2025. Publication activity peaks in 2023–2025, reflecting the surge of generative AI and its implications for scholarly communication, information infrastructures, and professional accountability (Aboelmaged et al., 2025; Buetow & Lovatt, 2024; Buitrago-Ciro et al., 2025). This temporal distribution suggests a paradigm shift: from discussions of automation and efficiency (Asemi et al., 2021; Collins et al., 2021) toward ethical operationalisation and socio-technical responsibility in academic library contexts (Ali & Richardson, 2025; Zareef & Jabeen, 2025).

Across the corpus, accountability and transparency emerge as the most consolidated principles within the FATE framework. Studies on AI auditing and governance (Ashok et al., 2022; Laine et al., 2024) emphasise traceability, model explainability, and stakeholder disclosure as the main ethical practices currently achievable in applied settings. The notion of traceable intelligence (Ridley, 2025) and the creation of documentation artefacts such as datasheets, algorithmic logs, and provenance chains reflect an empirical bias toward procedural transparency rather than substantive fairness. In parallel, library-focused research reinforces the librarian's evolving role as an ethical intermediary who curates not only datasets but also the metadata and contextual knowledge that enable algorithmic accountability (Sheikh et al., 2025; Stvilia et al., 2025).

While fairness and ethics appear in conceptual and normative discussions (Memarian & Doleck, 2023; Matei et al., 2025), empirical integration into library systems and data workflows remains scarce. Ethical reasoning tends to be externalised, delegated to institutional policies or compliance frameworks, rather than embedded within the mechanics of data stewardship. The literature reveals few attempts to quantify or operationalise equity within metadata schemas, discovery algorithms, or access protocols, highlighting a methodological blind spot between declarative ethics and measurable outcomes (Matsieli & Mutula, 2025; Njiru et al., 2025).

A notable conceptual synthesis emerges around terms such as algorithmic stewardship, ethical metadata and responsible curation. These constructs reflect a move beyond abstract ethics toward embedded governance, where librarianship intersects with AI ethics, information management, and digital policy (Ali & Richardson, 2025; Tang, 2025). Algorithmic stewardship reframes librarians as active custodians of the full AI pipeline, from data acquisition to model validation and interpretability, ensuring that decision-making processes remain transparent, contestable, and auditable (Lee et al., 2023; Roy, 2025). Similarly, ethical metadata practices propose value-sensitive vocabularies and annotations that make datasets readable not only by machines but also by moral and regulatory frameworks (Li et al., 2025).

To consolidate these patterns, a FATE-curation matrix is proposed, aligning stages of the data curation lifecycle (selection, documentation, preservation, dissemination) with the four FATE principles. Such a conceptual map exposes the areas of highest maturity, chiefly transparency and accountability in repository management, and the critical gaps in fairness and ethics integration. This visual and analytical synthesis underscores that while the field advances toward procedural accountability, epistemic justice and ethical embedding remain the next frontier for research and practice in AI-driven academic librarianship.

Identified gaps and future directions

The synthesis of the reviewed literature exposes a persistent asymmetry between the normative articulation of FATE principles and their practical implementation within academic library ecosystems. While ethical aspirations are increasingly present in strategic frameworks and theoretical discussions (Ashok et al., 2022; Laine et al., 2024; Memarian & Doleck, 2023), empirical validation remains remarkably limited. Few studies have examined how fairness, accountability, transparency, and ethics operate as measurable or auditable constructs in library-managed AI systems (Mitha & Omarsaib, 2025; Sheikh et al., 2025).

Most contributions to date remain conceptual, reflective, or policy-oriented, offering prescriptive statements rather than evidence-based demonstrations (Buitrago-Ciro et al., 2025; Ridley, 2025). Experimental and field-based investigations, such as algorithmic audits of discovery tools, recommender systems, or digital repositories, are notably scarce. Even within the growing corpus on data curation, ethical compliance tends to be reported as a descriptive attribute rather than a processual outcome (Stvilia et al., 2025). This gap suggests the need for interventional research designs capable of evaluating how FATE-aligned mechanisms affect access, inclusion, and user trust in library information systems.

A central methodological deficit lies in the lack of AI auditing models tailored to the library and information science domain. Existing frameworks, such as those discussed by Li et al. (2025) and Laine et al. (2024), are often derived from industrial or technical contexts, insufficiently sensitive to the epistemic and social dimensions of information stewardship. No standardized metrics currently exist to assess ethical performance in library AI systems, particularly regarding fairness in metadata representation, accountability in data provenance, or transparency in automated classification. Developing domain-specific auditing tools thus represents an urgent avenue for scholarly and institutional innovation.

The literature calls for integrative architectures that align technological infrastructures with ethical governance mechanisms (Matei et al., 2025; Njiru et al., 2025). Hybrid frameworks, combining algorithmic auditing, policy regulation, and librarian-mediated oversight, are positioned as the next evolution in responsible information management. Such models would operationalise FATE through continuous feedback loops between system design, ethical evaluation, and institutional accountability (Collins et al., 2021; Tang, 2025).

A promising trajectory lies in the consolidation of ethical AI librarianship, a paradigm where librarians act as mediators of algorithmic accountability infrastructures (Ali & Richardson, 2025; Matsieli & Mutula, 2025). This perspective redefines librarianship as an ethical-technical profession grounded in algorithmic literacy, data ethics, and participatory governance. In this vision, academic libraries evolve into living laboratories of responsible AI practice, testing, documenting, and modelling ethical compliance for broader institutional ecosystems.

In summary, future research must transcend descriptive analyses and advance toward empirical, evaluative, and co-creative methodologies that transform FATE from a conceptual ideal into an operational reality. This progression will not only enhance the ethical resilience of academic libraries but also position them as critical actors in the governance of trustworthy, transparent, and socially responsible AI.

To complement the conceptual overview, Figure 8 positions each of the FATE principles within a maturity, validation matrix. This visualisation clarifies the uneven evolution of research maturity across dimensions of fairness, accountability, transparency, and ethics, revealing a dominance of theoretical exploration over empirical validation. The matrix thus provides a diagnostic lens for identifying areas requiring future research investment and methodological development.



Figure 8. Research maturity and gap matrix (framework)

Source: Author's data elaborated with Artificial Intelligence

The visual imbalance highlighted in the matrix underscores a critical disconnect between conceptual maturity and empirical validation. While accountability and transparency benefit from clearer operational pathways, fairness and ethics remain largely aspirational, signalling an urgent need for applied studies and evaluative metrics within academic library settings.

Conceptual synthesis: towards a FATE curation framework (matrix)

This section proposes a FATE curation framework that integrates the research data curation lifecycle with operational controls aligned to FATE. The aim is not to add an ethical layer after the fact, but to embed measurable safeguards into the design, documentation, and governance of library data workflows and AI-enabled services (Ashok et al., 2022; Laine et al., 2024).

We treat the curation lifecycle, selection, ingest and normalisation, description and enrichment, preservation, and access and discovery, as the locus where FATE can be made actionable. Each stage is associated with concrete artefacts, decision rights, and evaluation routines that render ethical commitments auditable and explainable to stakeholders (Li et al., 2025; Stvilia et al., 2025):

- Selection (fairness, ethics)
 - Representational audits across language, field, geography, and community; rights and proportionality checks for sensitive data; public rationales for inclusion or exclusion; role-bound sign-off to ensure answerability (Memarian & Doleck, 2023; Sheikh et al., 2025);
- Ingest and normalisation (accountability, transparency)
 - Provenance capture, lineage and transformation logs, exception handling with escalation routes, and quality gates as routine controls rather than ad hoc fixes (Stvilia et al., 2025);
- Description and enrichment (fairness, transparency)

- o Bias-aware vocabularies and authority control; documentation of automated annotations; user-comprehensible rationales for derived fields that feed ranking or recommendation (Li et al., 2025; Ridley, 2025);
- Preservation (ethics, accountability)
 - o Privacy-preserving retention schedules, consent provenance, format migration with audit trails, and periodic risk reviews for legacy datasets reused by AI systems (Sheikh et al., 2025; Stvilia et al., 2025);
- Access, discovery and recommendation (All FATE principles)
 - o Exposure diversity and subgroup error monitoring; boundary-of-competence notices and escalation protocols for chatbots; public documentation (e.g., datasheets and model cards) for major algorithmic services (Aboelmaged et al., 2025; Afjal, 2023; Laine et al., 2024).

Four stages require heightened attention due to their downstream impact and risk concentration:

- 1) Selection
 - a. Curation choices shape what becomes visible and learnable in AI-mediated environments. Libraries should institute periodic representational audits and publish selection rationales to mitigate structural skew (Memarian & Doleck, 2023; Sheikh et al., 2025);
- 2) Description
 - a. Metadata decisions propagate into discovery and recommendation. Documentation of transformations and human-centred explainability for derived attributes support contestability and trust (Li et al., 2025; Ridley, 2025);
- 3) Preservation
 - a. Reuse of legacy data for training amplifies privacy and consent risks; proportionality and consent provenance must be explicit, especially when behavioural logs are retained (Matsieli & Mutula, 2025; Njiru et al., 2025).
- 4) Access
 - a. Algorithmic ranking and conversational mediation redistribute attention. Exposure metrics, error profiling, and clear escalation to human experts reduce harm from over-confident or misleading outputs (Black, 2024; Aboelmaged et al., 2025).

The framework is immediately applicable where libraries control repositories, metadata pipelines, and procurement terms, contexts in which governance artefacts (risk registers, change logs, incident reviews) can be mandated alongside technical controls (Ashok et al., 2022; Laine et al., 2024). Limitations arise with vendor opacity and entangled third-party services, which constrain auditability and explanation (Lee et al., 2023). Organisational capacity, skills, time, decision rights, also conditions feasibility, underscoring the role of AI literacy and data culture in sustaining implementation (Ali & Richardson, 2025; Ibrahim et al., 2025; Tang, 2025). Finally, metric transferability is non-trivial: generic fairness statistics often misfit scholarly discovery and cultural-heritage settings, calling for domain-calibrated indicators (e.g., exposure diversity, language coverage, provenance completeness) (Ofosu-Ampong, 2024; Stige et al., 2024).

Moving from aspiration to evidence requires interventional and evaluative designs:

- ✓ Algorithmic audits of discovery/recommendation: pre/post experiments on exposure diversity and subgroup error rates after metadata or ranking changes (Stvilia et al., 2025);

- ✓ Documentation compliance checks: periodic assessment of datasheets/model cards against minimal content standards and stakeholder readability (Laine et al., 2024; Li et al., 2025);
- ✓ Human-centred explainable AI evaluations: user studies on explanation usefulness, boundary-of-competence notices, and escalation efficacy in chatbot-mediated reference (Aboelmaged et al., 2025; Ridley, 2025);
- ✓ Operational maturity reviews: mapping decision rights, incident response, and procurement clauses to identify responsibility gaps (Ashok et al., 2022; Ibrahim et al., 2025);
- ✓ Method safeguards for AI-assisted workflows: when LLM support literature triage or drafting, enforce human verification and verifiable citation pathways to reduce overconfidence and hallucination risks (Carabantes et al., 2023; Buetow & Lovatt, 2024; Kousha & Thelwall, 2024; Scherbakov et al., 2025; Sparkman & Witt, 2025).

This synthesis positions FATE as infrastructure for data curation: a set of embedded controls, artefacts, and roles that make AI-enabled library services not only efficient but answerable and intelligible. The ensuing Discussion can therefore:

- i. examine trade-offs between efficiency and public interest values;
- ii. specify procurement and vendor-management templates that operationalise FATE;
- iii. prioritise a research agenda for library-specific auditing and metrics;
- iv. delineate pathways for co-production with vendors and researchers to validate the framework in live systems (Collins et al., 2021; Buitrago-Ciro et al., 2025; Romero & Chagua, 2025).

By articulating control points, applicability conditions, limitations, and validation methods, the FATE curation framework converts ethical intent into a workable governance grammar for AI-mediated academic librarianship.

Building upon the preceding synthesis, Table 4 summarises the operational components of the proposed FATE curation framework. It aligns curatorial stages with their corresponding FATE principles, key artefacts, and evaluation metrics. This tabular synthesis provides a replicable foundation for both scholarly validation and professional implementation in academic libraries.

Curatorial stage	Primary FATE principle(s)	Operational artefacts	Evaluation metrics
Selection	Fairness	Dataset documentation; inclusion or exclusion criteria; bias screening tools	Diversity ratio; representational equity; data provenance traceability
Ingest	Accountability	Data management plans; decision logs; provenance tracking systems	Auditability index; compliance verification rate; curator accountability logs
Description	Transparency	Metadata standards (e.g., Dublin Core, schema.org); explainable tagging protocols	Metadata completeness; explainability score; interoperability assessment
Preservation	Ethics & Accountability	Provenance records; long-term access agreements; consent documentation	Ethical compliance ratio; data retention audit outcomes; governance adherence index
Access & Reuse	Transparency and Fairness	Model cards; algorithmic impact statements; access control records	Accessibility index; bias detection in user interfaces; user satisfaction equity measure

Table 4. Summary of FATE-curation framework components

Source: Author's elaboration based on systematic literature review

Discussion

The systematic review identified a total of fifty-one studies, revealing a growing scholarly commitment to understanding how AI reshapes data curation practices and ethical governance within academic libraries. This temporal scope reflects an accelerating interest aligned with the rise of generative AI, automation in scholarly communication, and institutional efforts to formalise responsible data stewardship. The analysis shows that academic libraries are increasingly positioned as mediators between technological innovation and ethical accountability, assuming new roles as algorithmic and governance infrastructures in the digital knowledge ecosystem.

The corpus is dominated by conceptual and theoretical contributions, with empirical studies representing only a small minority of the reviewed publications. Notably, accountability and transparency receive sustained attention, whereas fairness and ethics remain comparatively underexplored, particularly in applied library contexts.

The reviewed corpus demonstrates a clear predominance of conceptual and theoretical studies over empirical investigations. Approximately two-thirds of the publications adopt a normative or framework-based perspective, focusing on the articulation of principles rather than on their operationalisation through measurable indicators or real-world evaluation. Only a limited number of studies employed experimental or audit-based methodologies, confirming that empirical research remains an underdeveloped dimension within the field. This imbalance highlights the persistence of a declarative rather than an evidence-based ethics, where FATE are discussed as ideals but seldom assessed as actionable components of data governance.

Across the literature, accountability and transparency emerge as the most mature and operationalised principles. Multiple studies, such as those by Laine et al. (2024) and Ridley, (2025), describe mechanisms for traceability and auditability, including model cards, datasheets for datasets, and documentation protocols that ensure data provenance and reproducibility. These approaches primarily respond to institutional demands for procedural governance and compliance. Fairness and ethics, however, remain comparatively marginal and conceptual, often restricted to philosophical or policy-oriented discussions rather than embedded in the design or evaluation of AI systems. Few studies propose concrete instruments to assess representational equity, mitigate bias, or address distributive justice in AI-driven curation workflows. This

asymmetry reveals a significant methodological gap between ethical aspiration and practical application.

The thematic synthesis of the selected studies identifies four dominant research clusters. The first, algorithmic stewardship, positions librarians and information professionals as active mediators in AI governance, tasked with ensuring transparency, explainability, and accountability across data lifecycles (Ali & Richardson, 2025; Tang, 2025). The second, ethical metadata, explores how descriptive standards can be reconfigured to capture cultural and contextual dimensions of data, thereby embedding moral awareness into cataloguing and classification systems (Li et al., 2025).

The third, responsible curation, extends traditional notions of technical preservation toward reflexive and socially responsive practices that account for equity and inclusivity (Sheikh et al., 2025; Stvilia et al., 2025). Finally, the fourth cluster, algorithmic justice in discovery, investigates how ranking and recommendation systems influence visibility, bias, and epistemic diversity in digital collections (Memarian & Doleck, 2023; Ofosu-Ampong, 2024). Collectively, these themes indicate a paradigmatic transformation in which libraries evolve from service providers into ethical infrastructures that mediate the social implications of algorithmic knowledge production.

Quantitative mapping of publication patterns reinforces these qualitative trends. Research output grew substantially after 2023, coinciding with the global expansion of generative AI tools such as ChatGPT and Claude, which prompted debates about authorship, integrity, and the automation of academic judgment (Afjal, 2023; Carabantes et al., 2023; Sparkman & Witt, 2025). The studies reviewed demonstrate that the ethical conversation within library and information science has shifted from a reactive posture, responding to external technological pressures, to a proactive orientation, where libraries seek to embed ethical principles at the level of infrastructure and policy. This marks the emergence of a distinct epistemic identity: libraries as algorithmic stewards of academic integrity and data governance.

The findings also reveal persistent structural challenges. Vendor opacity continues to limit transparency and independent auditability, particularly where proprietary AI systems are deployed without open access to their models or datasets. Ethical AI literacy among library professionals remains uneven, constraining the capacity to engage critically with algorithmic systems and to evaluate their outputs responsibly. Furthermore, although documentation artefacts such as datasheets, audit logs, and provenance records have become more common, they often serve compliance purposes rather than as tools for reflexive ethical practice. These limitations underscore the necessity for empirically grounded ethics that can bridge the gap between policy rhetoric and operational enactment.

Overall, the results affirm that academic libraries are at the forefront of a conceptual and infrastructural shift in how FATE principles are interpreted and implemented within data ecosystems. The evidence supports the proposition that fairness, accountability, transparency, and ethics must not remain abstract values but must be integrated into every stage of the data curation cycle, selection, description, preservation, and access. This integrative vision provides the conceptual foundation for the proposed FATE-curation framework, which seeks to translate ethical discourse into actionable governance. Such a framework holds the potential to advance accountability and transparency while simultaneously strengthening the underdeveloped domains of fairness and ethics through measurable and participatory mechanisms.

In sum, the systematic review demonstrates that the field has achieved considerable conceptual maturity but remains empirically thin and methodologically fragmented. To elevate data curation in academic libraries to a position of genuine ethical leadership, future research must move decisively from ethical reflection to ethical design, transforming FATE from a declarative

principle into a practical and auditable infrastructure for responsible AI in knowledge institutions.

Conclusion

This study offers a comprehensive synthesis of how academic libraries are navigating the complex intersection of data curation, AI, and the FATE principles. By systematically analysing the literature, it reveals a maturing but uneven field, where conceptual enthusiasm often exceeds methodological grounding and empirical validation. The findings demonstrate that while transparency and accountability are increasingly integrated into library infrastructures through audit mechanisms, provenance tracking, and governance policies, fairness and ethics remain conceptually diffuse and operationally underdeveloped.

Theoretically, the study advances the argument that libraries are no longer peripheral actors in the data ecosystem but rather ethical infrastructures mediating between technological systems and epistemic justice. The integration of FATE principles into data curation reframes librarianship as a moral and algorithmic enterprise, one concerned not only with managing information but with ensuring that data flows, algorithmic processes, and metadata structures embody principles of equity, trust, and responsibility. This conceptual repositioning aligns with the broader evolution of information science toward human-centred and ethically accountable knowledge systems.

Practically, the research underscores the need for the co-development of hybrid frameworks, such as the proposed FATE curation framework, that translate ethical principles into operational standards across all stages of the data lifecycle. These frameworks must embed explainability, inclusivity, and participatory governance into data stewardship practices, transforming libraries into testing grounds for ethical AI design. The professional development of librarians as ethical data stewards is central to this transformation, requiring targeted literacy programmes, institutional policies, and cross-sectoral collaborations.

Methodologically, this study illustrates that the systematic review process itself benefits from the reflective application of FATE principles. The use of AI-assisted tools such as SCOPUS AI and Rayyan enhanced analytical rigour, while human oversight preserved interpretative integrity, demonstrating that responsible AI use in research must always balance automation with accountability. This recursive methodological alignment ensures that the research process embodies the same ethical commitments it advocates.

Looking ahead, the future of data curation in academic libraries depends on bridging persistent gaps: the absence of empirical evidence on FATE implementation, the limited availability of ethical audit instruments, and the insufficient integration of human values into algorithmic infrastructures. Addressing these challenges requires interdisciplinary collaboration between information scientists, ethicists, technologists, and policymakers to co-create frameworks that are simultaneously ethical, transparent, and effective.

By reframing data curation as an ethically charged, algorithmically mediated practice, this study advances the emerging field of ethical AI librarianship. The proposed framework offers a foundation for future empirical research, institutional policy development, and the design of accountability infrastructures in AI-enabled academic libraries.

In conclusion, academic libraries stand at the frontier of a moral and epistemic transformation. Their evolving role as ethical and algorithmic mediators positions them to lead the responsible governance of data and AI in higher education. By institutionalising the FATE principles within their curatorial practices, libraries can move beyond compliance to become exemplars of ethical intelligence, ensuring that the digital infrastructures of knowledge remain not only open and interoperable, but fair, accountable, transparent, and profoundly human.

About the author

Nuno Miguel Teixeira Sousa completed his Master's degree in Information Science in 2023 from the Faculty of Arts and Humanities of Coimbra University and his Bachelor's degree in Information Science in 2021 from the Faculty of Arts and Humanities of Coimbra University. He is currently studying for a PhD in Information Science at the Faculty of Letters of the University of Lisbon since 2024 and he is investigator of Center for Classical Studies. Works in the field(s) of Social Sciences with an emphasis on Communication Sciences with an emphasis on Information Sciences. He is a Senior Technician at the Library of the Law Faculty of the Coimbra University. He has received four awards. He can be contacted at d.balbinjr@bsu.edu.ph

References

- Aboelmaged, M., Bani-Melhem, S., Al-Hawari, M., & Ahmad, I. (2025). Conversational AI Chatbots in library research: An integrative review and future research agenda. *Journal of Librarianship and Information Science*, 57(2), 331–347. <https://doi.org/10.1177/09610006231224440>
- Afjal, M. (2023). ChatGPT and the AI revolution: A comprehensive investigation of its multidimensional impact and potential. *Library Hi Tech*, 42(1), 353–376. <https://doi.org/10.1108/LHT-07-2023-0322>
- Ali, M. Y., & Richardson, J. (2025). AI literacy guidelines and policies for academic libraries: A scoping review. *IFLA Journal*, 51(3), 588–599. <https://doi.org/10.1177/03400352251321192>
- 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. <https://doi.org/10.1108/LHT-02-2020-0038>
- Ashok, M., Madan, R., Joha, A., & Sivarajah, U. (2022). Ethical framework for Artificial Intelligence and Digital technologies. *International Journal of Information Management*, 62, paper102433. <https://doi.org/10.1016/j.ijinfomgt.2021.102433>
- Black, J. (2024). Can AI Lie? Chabot technologies, the subject, and the importance of lying. *Social Science Computer Review*, 43(6), 1147–1158. <https://doi.org/10.1177/08944393241282602>
- Bu, Y.-Q., Cao, Y.-F., Chang, Z.-Y., Chen, H.-Y., Chen, X.-W., Chen, Y.-Y., Chen, Z.-C., Deng, R., Ding, J., & Fan, Z.-K. (2025). Expert consensus on the ethical requirements for generative AI-assisted academic writing. *Chinese Journal of Biochemistry and Molecular Biology*, 41(6), 826–832. <https://doi.org/10.13865/j.cnki.cjbmb.2025.06.1272>
- Buetow, S., & Lovatt, J. (2024). From insight to innovation: Harnessing artificial intelligence for dynamic literature reviews. *Journal of Academic Librarianship*, 50(4), 102901. <https://doi.org/10.1016/j.acalib.2024.102901>
- Buitrago-Ciro, J., Campos, E., & Romero, C. (2025). ¿Cómo está transformando la inteligencia artificial la comunicación científica? Desafíos, oportunidades y el papel de los actores involucrados: una revisión de alcance [How is artificial intelligence transforming scholarly communication? Challenges, opportunities, and the role of stakeholders: A scoping review]. *Investigacion Bibliotecologica*, 39(104), 111–150. <https://doi.org/10.22201/iibi.24488321xe.2025.104.59032>

- Burns, J., Etherington, C., Cheng-Boivin, O., & Boet, S. (2024). Using an artificial intelligence tool can be as accurate as human assessors in level one screening for a systematic review. *Health Information and Libraries Journal*, 41(2), 136–148. <https://doi.org/10.1111/hir.12413>
- Carabantes, D., González-Geraldo, J., & Jover, G. (2023). ChatGPT could be the reviewer of your next scientific paper. Evidence on the limits of AI-assisted academic reviews. *Profesional de la Informacion*, 32(5).
- Collins, C., Dennehy, D., Conboy, K., & Mikalef, P. (2021). Artificial intelligence in information systems research: A systematic literature review and research agenda. *International Journal of Information Management*, 60, 102383. <https://doi.org/10.1016/j.ijinfomgt.2021.102383>
- Heyder, T., Passlack, N., & Posegga, O. (2023). Ethical management of human-AI interaction: Theory development review. *Journal of Strategic Information Systems*, 32(3), 101772. <https://doi.org/10.1016/j.jsis.2023.101772>
- Ibrahim, H., Ahmad, K., & Sallehudin, H. (2025). Impact of organisational, environmental, technological and human factors on cloud computing adoption for university libraries. *Journal of Librarianship and Information Science*, 57(2), 311–330. <https://doi.org/10.1177/09610006231214570>
- Khosrawi-Rad, B., Grogorick, L., & Robra-Bissantz, S. (2023). Game-inspired Pedagogical Conversational Agents: A Systematic Literature Review. *AIS Transactions On Human-Computer Interaction*, 15(2), 146–192. <https://doi.org/10.17705/1thci.00187>
- Kousha, K., & Thelwall, M. (2024). Artificial intelligence to support publishing and peer review: A summary and review. *Learned Publishing*, 37(1), 4–12. <https://doi.org/10.1002/leap.1570>
- Laine, J., Minkkinen, M., & Mäntymäki, M. (2024). Ethics-based AI auditing: A systematic literature review on conceptualizations of ethical principles and knowledge contributions to stakeholders. *Information & Management*, 61(5), 103969. <https://doi.org/10.1016/j.im.2024.103969>
- Lee, M., Scheepers, H., Lui, A., & Ngai, E. (2023). The implementation of artificial intelligence in organizations: A systematic literature review. *Information & Management*, 60(5). <https://doi.org/10.1016/j.im.2023.103816>
- Li, Q., Li, Y., Zhang, S., Zhou, X., & Pan, Z. (2025). A theoretical framework for human-centered intelligent information services: A systematic review. *Information Processing & Management*, 62(1). <https://doi.org/10.1016/j.ipm.2024.103891>
- Matei, S., Jackson, D., & Bertino, E. (2025). Ethical reasoning in artificial intelligence: A cybersecurity perspective. *Information Society*, 41(2), 110–122. <https://doi.org/10.1080/01972243.2024.2429060>
- Matsieli, M., & Mutula, S. (2025). Generative AI and the Information Society: Ethical reflections from libraries. *Information (Switzerland)*, 16(9), 771. <https://doi.org/10.3390/info16090771>
- Memarian, B., & Doleck, T. (2023). Fairness, accountability, transparency, and ethics (FATE) in artificial intelligence (AI) and higher education: A systematic review. *Computers and Education: Artificial Intelligence*, 5, 100152. <https://doi.org/10.1016/j.caeai.2023.100152>
- Mitha, S. B., & Omarsaib, M. (2025). Emerging technologies and higher education libraries: A bibliometric analysis of the global literature. *Library Hi Tech*, 43(2), 1248–1270. <https://doi.org/10.1108/LHT-02-2024-0105>

- Njiru, D., Mugo, D., & Musyoka, F. (2025). Ethical considerations in AI-based user profiling for knowledge management: A critical review. *Telematics and Informatics Reports*, 18, 100205. <https://doi.org/10.1016/j.teler.2025.100205>
- Ofosu-Ampong, K. (2024). Artificial intelligence research: A review on dominant themes, methods, frameworks and future research directions. *Telematics and Informatics REPORTS*, 14, 100127. <https://doi.org/10.1016/j.teler.2024.100127>
- Ramirez, D., Foster, M., Kogut, A., & Xiao, D. (2022). Adherence to systematic review standards: Impact of librarian involvement in Campbell Collaboration's education reviews. *Journal of Academic Librarianship*, 48(5), 102567. <https://doi.org/10.1016/j.acalib.2022.102567>
- Ridley, M. (2025). Human-centered explainable artificial intelligence: An Annual Review of Information Science and Technology (ARIST) paper. *Journal of the Association for Information Science and Technology*, 76(1), 98–120. <https://doi.org/10.1002/asi.24889>
- Romero, M., & Chagua, R. (2025). Quinquenio de la aplicación de la inteligencia artificial en la archivística: una revisión sistemática en revistas académicas [Five-year application of artificial intelligence in archival science: A systematic review in academic journals]. *Bid-Textos Universitarios de Biblioteconomía i Documentació*, 54(2), 21 p. <https://bid.ub.edu/es/54/horna.htm>
- Roy, P. (2025). Transforming higher education libraries with data analytics, business intelligence, and business analytics: A review. *Journal of Librarianship and Information Science*. <https://doi.org/10.1177/09610006241307028>
- Scherbakov, D., Hubig, N., Jansari, V., Bakumenko, A., & Lenert, L. (2025). The emergence of large language models as tools in literature reviews: A large language model-assisted systematic review. *Journal of the American Medical Informatics Association*, 32(6), 1071–1086. <https://doi.org/10.1093/jamia/ocaf063>
- Shahzad, K., Khan, S., & Iqbal, A. (2024). Factors influencing the adoption of robotic technologies in academic libraries: A systematic literature review (SLR). *Journal of Librarianship and Information Science*, 57(3), 687–704. <https://doi.org/10.1177/09610006241231012>
- Sheikh, A., Malik, A., & Adnan, R. (2025). Evolution of research data management in academic libraries: A review of the literature. *Information Development*, 41(2), 305–319. <https://doi.org/10.1177/02666669231157405>
- Sparkman, M., & Witt, A. (2025). Claude AI and literature reviews: An experiment in utility and ethical use. *Library Trends*, 73(3), 355–380. <https://doi.org/10.1353/lib.2025.a961199>
- Stige, Å., Zamani, E., Mikalef, P., & Zhu, Y. (2024). Artificial intelligence (AI) for user experience (UX) design: A systematic literature review and future research agenda. *Information Technology & People*, 37(6), 2324–2352. <https://doi.org/10.1108/ITP-07-2022-0519>
- Stvilia, B., Pang, Y., Lee, D. J., & Gunaydin, F. (2025). Data quality assurance practices in research data repositories: A systematic literature review. *Annual Review of Information Science and Technology (ARIST)*, 76(1), 238–261. <https://doi.org/10.1002/asi.24948>
- Tang, Y. (2025). Cultivating university data culture in the age of artificial intelligence: A conceptual framework and critical reflections. *Information Research*, 30(iConf), 500–507. <https://doi.org/10.47989/ir30iConf47293>

- Thong, C., Atallah, Z., Islam, S., Lim, W., & Cherukuri, A. (2025). AI-powered tools for doctoral supervision in higher education: A systematic review. *Journal Of Information & Knowledge Management*, 24(2), 2530001. <https://doi.org/10.1142/S0219649225300013>
- Trigo, A., Stein, N., & Belfo, F. (2024). Strategies to improve fairness in artificial intelligence: A systematic literature review. *Education for Information*, 40(3), 323–346. <https://doi.org/10.3233/EFI-240045>
- Zareef, M., & Jabeen, M. (2025). A systematic review of digital curation services in academic libraries: Navigating policies, skills and challenges. *Digital Library Perspectives*, 41(3), 518–543. <https://doi.org/10.1108/DLP-10-2024-0158>

Copyright

Authors contributing to *Information Research* agree to publish their articles under a [Creative Commons CC BY-NC 4.0 license](https://creativecommons.org/licenses/by-nc/4.0/), which gives third parties the right to copy and redistribute the material in any medium or format. It also gives third parties the right to remix, transform and build upon the material for any purpose, except commercial, on the condition that clear acknowledgment is given to the author(s) of the work, that a link to the license is provided and that it is made clear if changes have been made to the work. This must be done in a reasonable manner, and must not imply that the licensor endorses the use of the work by third parties. The author(s) retain copyright to the work. You can also read more at: <https://publicera.kb.se/ir/openaccess>