

Human-in-the-Loop (HITL) as an Epistemic Imperative in Algorithmic Knowledge Organization

Alireza Noruzi

Department of Information Science and Knowledge Management, Faculty of Public Administration and Organizational Sciences, College of Management, University of Tehran, Tehran, Iran. Email: noruzi@ut.ac.ir

Abstract

The exponential growth of information worldwide has rendered entirely manual cataloging, classification, and metadata creation economically and logistically unsustainable. Conversely, the rush toward fully autonomous, AI-driven knowledge organization systems has proven equally problematic, as these systems are beset by algorithmic hallucinations, contextual blindness, and the implicit encoding of historical biases. Consequently, the discipline of library and information science finds itself trapped in a false dichotomy between the unscalable bottleneck of human labor and the unreliable opacity of machine autonomy. This editorial argues that Human-in-the-Loop (HITL) in knowledge organization is not a concession to technological limitations, but rather a strategic cognitive partnership that leverages both machine capabilities and human nuance. This paradigm transcends a mere quality-control mechanism, emerging instead as a foundational philosophy for knowledge organization in the algorithmic age. Drawing upon empirical evidence from evaluation studies conducted at the Library of Congress, the German National Library of Economics (ZBW), and the Shanghai Library, this study demonstrates that while automated systems perform poorly in subject heading assignment and classification number allocation, with accuracy rates below 50%, HITL workflows can achieve near-perfect accuracy (approaching 100%). The application of this paradigm across five core domains—namely cataloging, classification, indexing, metadata, and ontology—is briefly addressed. The editorial concludes by analyzing the key challenges of implementing HITL in knowledge organization, including intervention timing, training data quality, and organizational adaptation, and argues that the future of the field lies not in human versus machine, but in human-directed AI systems—an approach that guides the enduring profession toward a renewed and empowered purpose in the age of artificial intelligence.

Keywords: Human-in-the-Loop (HITL), knowledge organization, artificial intelligence, epistemic imperative, metadata, cataloging, algorithmic bias

Source: Noruzi, A. (2026). Human-in-the-Loop (HITL) as an Epistemic Imperative in Algorithmic Knowledge Organization. *Librarianship and Information Organization Studies*, 37(2), 7-18. <https://doi.org/10.30484/nastinfo.2026.4173.2414>
https://nastinfo.nlai.ir/article_3384.html

Introduction

The exponential growth of global information has rendered purely manual cataloging, classification, and metadata creation economically and logistically unsustainable. Conversely, the rush toward fully autonomous, AI-driven knowledge organization systems has proven equally flawed, being beset by algorithmic hallucinations, contextual blindness, and the silent encoding of historical biases. The information profession is consequently trapped in a false dichotomy between the unscalable bottleneck of human labor and the unreliable opacity of machine autonomy.

As we navigate the complexities of the modern information landscape, the integration of artificial intelligence into library and information science (LIS) and knowledge management is no longer a futuristic prospect; it is our daily operational reality. Large language models (LLMs) and machine learning algorithms can process, cluster, and tag data at speeds that human beings cannot match. Yet, as we automate the mechanics of information retrieval, we risk forfeiting the interpretive depth that constitutes the very *soul* of information understanding. When multimodal systems can assist special-collections cataloging with reported efficiency gains of 66.7% (Shi et al., 2026), a fundamental question confronts every knowledge organizer: *what role should human beings play in this increasingly automated epistemic ecosystem?*

Today's most advanced AI models, when tasked with the core operations of knowledge organization, remain far from operating independently at a level that can substitute for human judgment. The Library of Congress's [Exploring Computational Description](#) experiment revealed that while AI can extract structured fields such as identifiers (LCCN, ISBN), titles, and authors, with accuracy as high as 85%, its performance on complex semantic tasks, particularly subject heading assignment, genre classification, and dates, drops below 50%. The Library's conclusion was unambiguous: *"Because library catalog records must contain highly reliable and accurate metadata, no current AI tools are currently good enough to run fully automatically for this task."*

The Library of Congress Labs' *Exploring Computational Description* experiment tested five open-source machine-learning models on approximately 23,000 eBooks, measuring performance using the F1 score, a standard classification metric combining precision and recall on a 0-100% scale. Large language models achieved only 26% F1 accuracy for Library of Congress Subject Headings, while Annif software reached 35% for text classification (Potter & Saccucci, 2024). The study established 95% F1 as the quality threshold for production use, which no AI system approached.

Automation is not a panacea. It is particularly fragile when confronted with domain ambiguity, implicit meanings, and incomplete inputs. A single term may carry radically different connotations across disciplinary contexts; an abbreviation may have multiple interpretations within a single organization; the intellectual "intent" of a document frequently cannot be inferred from its text alone. These are commonplace scenarios in knowledge organization, and they remain stubborn hurdles for fully autonomous systems.

It is precisely at these frontiers that the Human-in-the-Loop (HITL) paradigm demonstrates its irreplaceable epistemic value. HITL is not a concession to technological limitations but a strategic cognitive partnership that leverages machine scale alongside human nuance. It represents a transition from a mere quality-control mechanism to a foundational philosophy for responsible knowledge organization in the algorithmic age.

HITL is not merely a linear partitioning of tasks, “machines do this, humans do that.” Rather, it constitutes a deeper form of cognitive collaboration: machines provide scale, speed, and pattern recognition; humans contribute judgment, contextual awareness, and professional accountability.

In cataloging, the “VLM + LLM + HITL” workflow proposed by the Shanghai Library offers a compelling case. The AI’s initial accuracy was 75%, but through structured human–machine collaboration, the final accuracy approached nearly 100%. This is not a one-way “machine generate, human review” pipeline but an iterative feedback loop in which AI proposes candidate records, humans exercise critical adjudication, and the system continuously learns, each human intervention refining the model, and each model output enhancing the efficiency of subsequent human tasks. (Shi et al., 2026).

In classification and subject indexing, the German National Library of Economics (ZBW) exemplifies HITL practice through its AutoSE project. A review panel of seven to eight subject specialists evaluates machine-generated subject headings on approximately 1,000 publications annually. The review does more than ensure quality. It systematically detects algorithmic biases, such as the over-representation of terms like “theory” and “United States,” which arise from skewed training distributions. Such distortions can only be identified and rectified by domain experts who possess the substantive knowledge to recognize aberrant patterns. As the ZBW researchers observe, knowledge organization systems and their inter-system mappings “are usually created and maintained by humans”, not out of methodological conservatism, but out of professional necessity (Kasprzik, 2025).

In metadata management, HITL is equally vital. High-quality metadata requires not only completeness but also accuracy, consistency, and contextual relevance. Automated tools might populate every field, yet they cannot adjudicate what a term genuinely signifies within a specific domain or community of practice. A metadata pipeline without HITL is essentially an unvalidated automation flow, expeditious, but not reliably trustworthy (Dayal, 2026).

Applications of HITL in Knowledge Organization

The HITL paradigm represents the intelligent integration of machine capabilities with the judgment and creativity of human information science professionals. This paradigm reshapes the foundational pillars of our discipline in the following concrete ways.

1. Cataloging: From Data Entry to Semantic Enrichment

In traditional cataloging, the human was the sole creator of the bibliographic record. Today, AI can generate baseline MARC or linked-data records in seconds by extracting data from title pages, tables of contents, and even the full text of the resource itself. Under the HITL approach, the machine acts as the draftsman; the cataloger serves as the editor and semantic enricher. The human-in-the-loop reviews the AI-generated record for accuracy, resolves conflicting or ambiguous data, and adds semantic enrichment, such as relational linkages to other works, creators, and subjects, that requires an understanding of the resource’s broader cultural, academic, and intellectual impact.

Olson (2026) concludes that AI implementation serves vendor interests in labor cost reduction rather than improving information access, while consolidation creates platform monopolies extracting value from professional labor while eliminating the expertise that creates it. Resistance requires collective organizing, alternative infrastructure development, and recognition that current AI implementations conflict with core professional values. The framework extends platform capitalism theory to professional service contexts.

2. Classification: Navigating Ambiguity and Interdisciplinarity

Automated classification systems excel at placing clearly defined, single-subject items into established hierarchies such as the Dewey Decimal Classification (DDC) or the Library of Congress Classification (LCC). However, they frequently fail when confronted with interdisciplinary works, emergent fields, or culturally specific contexts that resist tidy hierarchical placement. In the Library of Congress experiment, the accuracy of subject heading assignment for such works dropped below 50% (Library of Congress, n.d.).

In the HITL model, algorithms suggest classification numbers based on textual analysis and historical usage patterns. The human classifier or cataloger intervenes to render the final judgment, applying professional discernment to determine the “main subject” of a work. The human ensures that the classification reflects the author’s intellectual intent and the specific needs of the user community, rather than merely the statistical probabilities computed by the machine.

3. Indexing: Contextualizing the Aboutness

Indexing requires understanding not only which words appear in a text, but what the text is *about*, its conceptual *aboutness*. AI is highly proficient at keyword extraction and named-entity recognition, yet it struggles with irony, sarcasm, implied meaning, and domain-specific jargon that require inferential reasoning.

Under HITL, the AI generates a broad pool of potential index terms, descriptors, concepts, and named entities. The human indexer curates this candidate list, discarding false positives, combining synonyms, and ensuring that controlled vocabularies (such as LCSH or MeSH) are applied consistently and appropriately. The human guarantees that the index functions as a precise and reliable cartographic guide rather than a noisy, undifferentiated keyword cloud.

4. Metadata: Ensuring Provenance, Ethics, and Accuracy

Metadata is the lifeblood of discoverability, identification, and retrieval. While AI can automatically generate descriptive, structural, and administrative metadata, it cannot verify provenance, assess ethical implications, or understand the sensitivity of the information it processes. HITL is therefore critical for ethical metadata practice. Information professionals must review AI-generated metadata to ensure it does not perpetuate harmful stereotypes, misgender authors, or misrepresent Indigenous knowledge systems. Furthermore, human oversight is essential to establish provenance and rights metadata (e.g., Creative Commons licensing or copyright status), ensuring legal and ethical compliance for which machines cannot be held accountable.

Organizations are generating more metadata than ever before. Document management systems using machine learning extract characteristics of information resources. Large language models now summarize, classify, and tag content at scale. This is precisely where HITL becomes essential. When automated systems fail, humans provide the context, judgment, and accountability that automated systems still struggle to replicate. When metadata must be scalable, interpretable, and reliable, human experts cannot be completely removed from the loop. Thus, HITL approaches remain essential for producing accurate, interpretable, and trustworthy metadata, and the informed oversight of knowledge organization experts strengthens the automated metadata pipeline (Dayal, 2026).

5. Organizing Knowledge: Steering the Epistemology

At the highest level, organizing knowledge is an epistemological act. It dictates how societies categorize truth, history, and scientific inquiry. Because AI models are trained on historical data,

they inevitably reflect the biases, blind spots, and colonial structures of the past. Human experts must therefore remain the ultimate architects of our ontologies and taxonomies. When an AI suggests a new branch in a knowledge graph or a new category in a taxonomy, human knowledge organizers must evaluate that proposal against contemporary ethical standards, diverse cultural perspectives, and evolving academic paradigms. HITL ensures that our knowledge organization systems remain dynamic, inclusive, and reflective of current human values, rather than reifying the limitations embedded in their training corpora.

Ethical Governance and Professional Empowerment

The significance of HITL extends well beyond operational efficiency. In the context of libraries and cultural heritage institutions, it carries profound ethical and governance functions.

The Library of Congress's *Human in the Loop* project explicitly frames HITL as a mechanism to address algorithmic bias, maintain the trustworthiness of machine-generated data, and communicate data provenance and potential inaccuracies to end users. When library patrons trust a catalog, that trust rests on the accuracy and transparency of its metadata, precisely what a purely automated system cannot alone guarantee.

Bias detection and correction represent a unique strength of HITL. As the ZBW case illustrates, expert review uncovered systematic distortions that quantitative evaluations would have entirely missed, distortions that carry substantive consequences for the fairness and representativeness of knowledge organization. Accountability and transparency are equally central: when an AI system produces an erroneous judgment, the HITL workflow provides a clear chain of responsibility attribution and a viable correction mechanism. In data lake and AI governance contexts, HITL is increasingly regarded as a required design element to satisfy regulatory demands for meaningful human oversight, such as those articulated under the GDPR.

Crucially, the adoption of HITL should not be viewed by information professionals as a threat, but as a liberation. By offloading the repetitive, mechanical aspects of data entry and basic extraction to machines, we free human professionals to engage in the higher-order activities for which they are uniquely qualified: critical reasoning, ethical deliberation, and the design of intuitive, user-centric knowledge ecosystems. The future of cataloging, classification, and metadata is not human *versus* machine; it is human *directing* machine. By keeping the human firmly in the loop, we ensure that our systems of organizing knowledge remain not merely efficient, but deeply and inherently human in their orientation.

Conclusion

The core insight of HITL is this: the goal of automation is not to supplant human professional judgment, but to reconfigure its deployment, freeing humans from repetitive chores so they can focus on tasks that demand discernment, contextual understanding, and ethical reflection. In the long arc of knowledge organization, catalogers, classifiers, and indexers have never been mere data movers; they are architects of intellectual order and custodians of cultural memory. AI can accelerate their work, and it can assist it, but it cannot replace the most essential human dimensions: the interpretation of content, the grasp of context, and the commitment to ethical stewardship.

HITL is not a retreat of humanity in the face of automation. Rather, it represents a reconfiguration of the epistemic position of information professionals as stewards of knowledge systems in the age of artificial intelligence. This position ensures that the locus of judgment and decision-making remains concentrated at the highest cognitive and ethical levels. Success in

implementing this paradigm, however, depends heavily on the epistemic robustness of training corpora and the quality of underlying data. As the Library of Congress experiments have demonstrated, without balanced datasets, even the most sophisticated HITL workflows will face performance limitations. Thus, the systematic development of techniques and practices for creating balanced datasets emerges as an ongoing, long-term research endeavor—one whose continued vitality is essential and unavoidable for sustaining the efficacy of this paradigm. *When* to intervene, *how* to intervene, and *who* should intervene, these three questions constitute the core of HITL design. Not every human intervention is equally valuable; not every timing is equally efficacious. Scholars have called for moving beyond generic HITL models toward structured collaboration methodologies that are tailored to the specific demands of distinct knowledge-organization tasks.

References

- Dayal, U. (2026). *Why Human-in-the-Loop Is Critical for High-Quality Metadata?* Digital Divide Data. <https://www.digitaldividedata.com/blog/human-in-the-loop-metadata>
- Kasprzik, A. (2025). Transferring applied machine learning research into subject indexing practice. In: Balnaves, Edmund et al. (Eds.), *New Horizons in Artificial Intelligence in Libraries*. Berlin: De Gruyter Saur, 2025, pp. 199-212, <https://doi.org/10.1515/9783111336435-015>
- Library of Congress. (n.d.). *Exploring Computational Description: Investigating how machine learning can help with cataloging*. The Library of Congress Labs. <https://labs.loc.gov/work/experiments/ECD/>
- Olson, M. (2026). Operating the franchise: vendor consolidation, algorithmic mediation, and the remaking of academic librarians as platform administrators for digital capitalism. *Information Research*, 31(2), 348-367. <https://doi.org/10.47989/ir31262945>
- Potter, A., & Saccucci, C. (2024). Could artificial intelligence help catalog thousands of digital library books? An interview with Abigail Potter and Caroline Saccucci. *The Signal: Digital Happenings at the Library of Congress*. <https://blogs.loc.gov/thesignal/2024/11/could-artificial-intelligence-help-catalog-thousands-of-digital-library-books-an-interview-with-abigail-potter-and-caroline-saccucci/>
- Shi, Z., Zhang, Y., & Wang, P. (2026). Research on a new intelligent cataloging model for special collections based on multimodal large language models. *Journal of Information and Management*, 11(1), 52-63. <https://jim.library.sh.cn/EN/Y2026/V11/I1/52>