

When AI Comes for the Library

A Vision for Human-Led Knowledge Architecture

Zac Potter | Version 1.0 | September 2026

Artificial intelligence is rapidly becoming one of the primary interfaces through which people search, synthesize, and interact with knowledge. It draws on intellectual infrastructure that universities, libraries, publishers, and researchers have spent generations building.

This paper confronts a simple question:

What if trusted knowledge institutions built their own living knowledge architectures that AI could consult, rather than simply allowing AI to consume their intellectual infrastructure?

An outside AI could bring a question to an institution's digital gateway. Behind that boundary, institutional knowledge, experts, agents, methods, and validation systems could produce a traceable consultation. The AI could receive an answer, supporting evidence, and stated limitations without receiving unrestricted access to the architecture that produced them.

Many of the technical components already exist. At the ISKO 2026 conference in São Paulo, I saw them developing separately: machine-readable knowledge organization systems, ontologies and knowledge graphs, research workflows, and approaches to provenance, transparency, and auditability. The harder challenge may be institutional: whether fragmented universities can align their systems, governance, methods, and people well enough to operate as coherent knowledge architectures.

This is intentionally an optimistic proposal. It imagines universities and libraries remaining active, sovereign participants in an AI-mediated knowledge economy by making human stewardship, institutional trust, and distinctive ways of knowing the center of value.

The Question in the Room

At ISKO 2026 in São Paulo, there was no single attitude toward artificial intelligence. I heard excitement about new possibilities, caution about their limits, ethical concern, and genuine uncertainty about what AI might mean for the professions represented in the room.

The International Society for Knowledge Organization brings together people who study and build the systems through which knowledge is classified, described, related, retrieved, and made useful. Their work includes library classification, metadata, thesauri, taxonomies, ontologies, knowledge graphs, and other forms of knowledge organization.

AI is increasingly operating in this same territory. People already use it to search, retrieve, summarize, compare, connect, and reorganize vast amounts of information. As these systems become more capable, they can extend the work of knowledge organization. They can also change the professional and economic roles of the institutions and people who have traditionally performed and governed that work.

Yet knowledge institutions contain more than collections. They also contain methods, standards, expertise, institutional histories, competing schools of thought, and judgments about what counts as trustworthy knowledge. These are not secondary features. Together, they shape how knowledge becomes credible, navigable, and useful.

During the conference, several variations of the same question seemed to fill the room; something along the lines of:

What happens when increasingly capable AI systems extract the value from these intellectual infrastructures while the very institutions and professions that have created and continue to maintain them are reduced to passive resources rather than remaining active participants in the exchange?

That is the problem this paper engages.

From Repository to Research Engine

The existential problem was certainly in the room and was a recurring topic in roundtables and panels at ISKO. But most of what I encountered was the careful construction of tools and approaches to knowledge organization. There were ontologies, thesauri, controlled vocabularies, knowledge graphs, metadata schemas, research workflows, provenance systems, methods for connecting AI with structured knowledge, and ways of auditing machine activity. I also heard knowledge organization professionals increasingly described as architects.

One of the clearest institutional visions came from Xuemao Wang, Dean of Libraries and Charles Deering McCormick University Librarian at Northwestern University. During the panel *AI and the Future of Research: Strategies and Approaches Through Knowledge Organization and Data Standardization*, I heard Wang describe the academic research library as moving “*from repository to research engine*.”

That phrase was striking to me. A repository preserves and provides access to the products of research. A research engine also helps organize, connect, evaluate, and activate knowledge as part of the research process itself. The library does not abandon its responsibility for stewardship. It makes that stewardship more active, more usable, and more deeply connected to the production of knowledge.

Two presentations helped me imagine how parts of such an engine might operate.

Wiesław Babik of Jagiellonian University presented a model of information logistics in knowledge organization. His central distinction was that organizing knowledge is not enough. A functioning knowledge system must also consider how information is sourced, structured, moved, delivered, and improved through feedback. In this view, knowledge organization gives information structure and meaning, while information logistics helps it reach the right recipient, at the right time, in the right form, and at the right level of quality.

Joseph A. Busch of Taxonomy Strategies approached the problem from another direction. He presented application profiles as documented frameworks that bring together functional requirements, domain models, metadata schemas, controlled vocabularies, usage guidance, and syntax rules for a particular information application. I began to understand such a profile as a kind of institutional method card. It can tell an authorized person, application, or agent what it is allowed to do, which structures and terms it must use, what standards govern its work, and how its output should be validated.

Neither presentation intended to offer a finished university knowledge architecture. However, what stood out to me was the pattern they revealed when placed beside Wang’s institutional vision. Wang described the movement from repository to research engine. Babik described the movement of knowledge through the system. Busch described a documented contract for how particular applications and agents could work within systems.

Repositories hold knowledge. Semantic systems organize it.

Information logistics governs its movement. Application profiles govern how it is used. Provenance and auditing make the process traceable.

I began to see these as pieces of a larger knowledge architecture.

The Pieces Already Exist

This paper’s vision does not require every technical and institutional component to be invented from scratch. Most of the pieces needed to construct coherent knowledge architectures already exist.

- **Institutional knowledge assets:** Universities already maintain special collections, publications, datasets, archives, teaching materials, institutional research, and other distinctive bodies of knowledge.
- **Machine-readable knowledge organization:** Taxonomies, ontologies, thesauri, controlled vocabularies, metadata schemas, knowledge graphs, and standards such as SKOS already give information explicit structure and relationships.
- **Documented methods and workflows:** Research protocols, application profiles, metadata requirements, validation procedures, and rules for human review can guide how people and machines work with institutional knowledge.
- **Provenance and audit:** Systems can record sources, versions, transformations, human actions, machine actions, approvals, and the steps used to produce a result.
- **Governed access:** Authentication, permissions, licensing, APIs, and controlled retrieval can allow knowledge to be consulted without giving unrestricted possession of the underlying collection or architecture.
- **AI-assisted orchestration:** Local agents and retrieval systems can already search structured information, route questions, use tools, follow defined workflows, and return sourced results within established limits.
- **Institutional governance:** Universities already have research standards, ethics processes, access policies, disciplinary methods, and human authorities responsible for maintaining them.

The vision is not for every university or library to build its own frontier AI model. The question is whether these existing capabilities can be assembled into a coherent institutional architecture with enough semantic structure, procedural guidance, technical interoperability, and human authority to interact with outside AI systems on the institution’s own terms.

Universities and libraries also hold forms of value that an AI company cannot simply reproduce: distinctive collections, accumulated expertise, institutional history, established authority, and generations of human stewardship.

The technical gap may be smaller than the institutional one. The larger challenge is synthesis: bringing these pieces together into a functional knowledge architecture.

What Could University Knowledge Architecture Become?

This is the utopian proposal at the center of the paper.

Imagine a university not simply as a collection of departments, libraries, databases, researchers, and publications, but as having a coordinated living knowledge architecture. Its knowledge could remain distributed across real people, disciplines, collections, and systems while the architecture makes those parts coherent enough to work together, govern access, expose their methods, and interact with the outside world as an institution.

There are many ways such an architecture could be designed. After seeing individual pieces of such an architecture developing across ISKO, this is the form I began to imagine.

A Possible Institutional Knowledge Architecture

INSTITUTIONAL CONSTITUTION

The foundation

Who are we, what do we value, and how do we decide?

Mission · values · governance · epistemic commitments · policies · authority

The Institutional Constitution establishes the authority under which the whole architecture operates. It defines what the institution exists to do, which values guide it, who may make consequential decisions, and how its commitments can be questioned or revised.

KNOWLEDGE STRUCTURE

The core

What do we know, and how have we chosen to organize it?

Repositories · collections · vocabularies · ontologies · taxonomies · knowledge graphs · mappings · standards

Each institution's Knowledge Structure is a unique architecture of both the institution's evidence and the systems through which that evidence becomes organized, connected, and retrievable. It does not claim to represent knowledge neutrally or finally. It makes the institution's organizing choices visible enough to use, compare, contest, and revise.

INSTITUTIONAL METHOD

The operating layer

How do we decide how to work with that knowledge?

Meta-methods · methods · application profiles · librarians · domain experts · agents · tools · workflows · validation

Institutional Method governs how work is performed inside the architecture. At its root are meta-methods, which help determine which research method, professional practice, tool, expert, or validation procedure is appropriate for a particular question.

An application profile can function here as an institutional method card. It documents what an authorized person, application, or agent

may do, which structures and vocabularies apply, what standards govern the task, and how the result must be validated.

KNOWLEDGE GATEWAY

The boundary and front door

Who can interact with the institution, and on what terms?

Identity · purpose · permissions · authentication · APIs · licensing · access rules · query handoff

The Knowledge Gateway is the first part encountered by an outside person, organization, or AI system. It establishes identity, purpose, permissions, access level, and terms before a question enters the institution.

The Gateway then attaches the appropriate institutional method card and hands the question into the architecture without granting unrestricted access to everything behind the boundary.

KNOWLEDGE LOGISTICS

Circulation through the system

Where should knowledge move, to whom, when, and in what form?

Sourcing · semantic structuring · routing · distribution · timing · form · quality

Wiesław Babik's work on information logistics helped clarify this element. Organizing knowledge is not enough. A functioning architecture must also govern how knowledge enters the system, how it is structured, where it moves, who should receive it, when it should arrive, what form it should take, and what level of quality the task requires.

Knowledge Logistics is the circulation system of the architecture. It moves questions, evidence, methods, human decisions, agent work, and completed consultations through the appropriate institutional pathways.

TRACEABILITY, AUDIT, AND FEEDBACK

The surrounding accountability and learning layer

What happened, can it be examined, and how does the institution learn?

Sources · versions · methods used · human actions · agent actions · approvals · corrections · outcomes

Traceability records the journey through the architecture. Audit compares what happened with what the Constitution, Gateway, and Institutional Method authorized or required. Feedback allows errors, outcomes, and recurring problems to improve future structures, methods, policies, and logistical pathways.

A university cannot guarantee that every conclusion will be correct. What it can do is make much more of the process visible: where information came from, how it was organized, which methods were used, what humans and machines did, what standards governed the work, and how errors were corrected.

In technical terms, much of this is provenance. In practical terms, it is the ability to show the journey through the institution. Its value comes from the combination of knowledge, structure,

method, human authority, controlled circulation, traceability, and institutional trust.

The architecture makes that trust inspectable, usable, and exchangeable.

AI Comes to the Library

The easiest way to understand the model is to imagine a future request:

Consult MIT

A researcher, company, or individual is already working with an AI system. They reach a question where MIT's knowledge, methods, or expertise might provide additional value. Instead of asking the outside AI to imitate MIT, they ask it to consult MIT.

The AI comes to the library.

- **The user requests the consultation.** "Consult MIT on this problem."
- **The outside AI reaches MIT's Knowledge Gateway.** Identity, purpose, permissions, access level, and terms are established. The consultation might be part of a public commons, an individual service, or an institutional agreement.
- **The institution accepts the handoff.** The appropriate institutional method card is selected for the question. It identifies the structures, methods, tools, evidence standards, validation, and human review that apply. The outside AI does not receive unrestricted access to the internal architecture. In a simple sense, it stays in the lobby.
- **MIT's systems go to work.** Knowledge Logistics routes the question through the appropriate collections, semantic structures, agents, librarians, experts, tools, and validation procedures. Traceability records the sources consulted, methods used, human and machine actions, and required approvals.
- **The institution returns a traceable consultation.** Depending on the service, this might include:
 - a synthesis or answer
 - supporting sources
 - methods used
 - limitations or confidence
 - provenance, approvals, and institutional attribution

The outside AI can then continue serving the user without requiring possession of the underlying library or unrestricted access to its intellectual infrastructure.

The same pattern could apply elsewhere:

Consult Harvard Law.

Consult a leading medical institution.

Consult the Library of Congress.

None of these institutions is infallible. Their value would come from **declared methods, accumulated expertise, traceable sources, institutional governance, and accountable human stewardship.**

Value for Value

Some consultations should remain freely available as part of the public commons. Others could represent substantial professional or institutional value.

A corporation asking an institution to apply its knowledge architecture, methods, agents, and expert systems to an important problem is receiving a real service. Licensing, contracts, usage agreements, attribution, or other forms of compensation could allow some of that value to return to the institutions producing it.

Those resources could support libraries, research infrastructure, knowledge professionals, faculty, public resources, and scholarships.

If AI creates an enormous new demand for trustworthy knowledge, that demand does not have to bypass the institutions that spent generations building it.

AI access would not be synonymous with plunder; AI could come to the library on the institution's terms.

The Hardest Architecture May Be Human

One of the strongest cautions I heard while discussing versions of this idea at ISKO was not that the technology was impossible, but that universities already struggle to coordinate the many people, systems, rules, and authorities among themselves, let alone present a coherent institutional interface to increasingly capable AI systems.

Universities are plural institutions by nature. Libraries, departments, research centers, presses, technology systems, faculty groups, administrators, and outside partners may operate under different standards, incentives, methods, licensing arrangements, and forms of authority.

That plurality is one of the university's greatest intellectual strengths. It is also an architectural challenge.

Before a university can present a coherent knowledge architecture to the outside world, it may first need to establish enough coherence within itself. That means resolving practical questions of governance, ownership, interoperability, licensing, maintenance, authority, and responsibility.

It also requires enough shared institutional will to act together where coordination matters without erasing legitimate differences between disciplines, methods, and schools of thought.

A coherent architecture does not require a uniform institution.

It requires clear boundaries, declared authority, compatible handoffs, and enough common structure for different parts of the institution to work together while preserving meaningful disagreement.

The challenge becomes larger if many institutions eventually want to interact with AI through compatible gateways, shared standards, or recognizable forms of institutional consultation. Each

institution may maintain its own methods and identity while still needing enough interoperability to participate in a wider knowledge ecosystem.

This does not make the architecture impossible. It identifies a central design problem.

Many of the technical pieces can already be assembled. The harder work is creating the human agreements, professional responsibilities, governance structures, and institutional relationships capable of maintaining them over time.

The hardest architecture may actually be human, but it is also the architecture that makes the rest worth trusting.

A Hopeful Frame

One can imagine a darker timeline in which increasingly capable AI systems absorb, reorganize, and reproduce more of the intellectual infrastructure that universities, libraries, researchers, publishers, and knowledge professionals spent generations building. In that future, the institutions that created and maintained this infrastructure risk becoming passive sources for systems operating around them.

This paper offers another possibility.

Universities and libraries could become living knowledge architectures: governed systems with their own collections, structures, methods, expertise, gateways, provenance, feedback, and institutional memory. They would not need to build the largest AI model. They would build the trusted environments that many different models could consult.

The result could be a more plural knowledge economy. Instead of relying on a small number of AI systems with epistemic assumptions and biases that are often difficult for outsiders to inspect, compare, or contest, hundreds of institutions could maintain their own declared methods, commitments, areas of expertise, ethos, and forms of accountability.

AI does not have to inherit the library by absorbing it. It can come to the library, receive a governed consultation, and leave with an answer rather than the architecture itself.

My hope is that our libraries become sovereign architectures of enduring value, and that AI comes to them on their terms.

Conference Sources

Wang, Xuemao. Remarks during the panel AI and the Future of Research: Strategies and Approaches Through Knowledge Organization and Data Standardization. 19th International ISKO Conference, São Paulo, Brazil, 2026.

Babik, Wiesław. Logistyka informacji [Information Logistics]. Kraków: Wydawnictwo Uniwersytetu Jagiellońskiego, 2025. Discussed by the author at the 19th International ISKO Conference, São Paulo, Brazil, 2026.

Busch, Joseph A. "Using Schema and Value Vocabularies to Structure Information Applications." Presentation at the 19th International ISKO Conference, São Paulo, Brazil, 2026.

Share & Discuss

X: @ZacPotter777 • CURRENT VERSION HOSTED @ ZacPotterOfficial.com