



AI INTEROPERABILITY FOR SCHOLARLY PUBLISHING


Todd A. Carpenter
Executive Director, NISO

JST Webinar
February 19, 2026


AI in Scholarly Publishing

Scholarly publishers are rapidly adopting and applying artificial intelligence tools in their workflows and business processes.

What are publishers doing with AI?
What can we standardize?
Issues as we move forward.



AI OPPORTUNITIES ABOUND



There are many applications of AI in research and scholarly publishing, such as research support, knowledge discovery, author recruitment or product development, but today we're primarily going to focus on interoperability processes and community awareness

AI Applications in STM Publishing Processes

Peer Review

Editorial support

Metadata creation

Accessibility, alt-text generation

Integrity & plagiarism detection

Marketing

Search & discovery

Research assessment



Standards Needs for AI Applications in Scholarly Publishing



Standards Needs for AI and STM Publishing

Consistent Terminology

Disclosure of AI Usage

Ingestion Signaling

Model License Language

Reference Linking in AI Output

Agentic Access Control

Usage Tracking and Reporting



Maintain versioned glossaries governed by a neutral standards body to prevent semantic drift as technologies evolve.




Disclosure of AI Usage

Standardize user-facing disclosures indicating when AI systems are used to generate, summarize, translate, or recommend content.

Define thresholds for disclosure based on materiality (e.g., fully generated vs. AI-assisted editorial workflows).

Enable machine-readable disclosure flags to support downstream transparency, audits, and compliance checks.



Ingestion Signaling

Support controls on content crawling and ingestion.

Create standardized metadata fields signaling permissions or restrictions for AI training, fine-tuning, and inference-time retrieval.

Align signals with existing publishing metadata frameworks (e.g., rights, licenses, embargoes) to avoid duplication.

Support both human-readable policy statements and automated enforcement by crawlers and ingestion pipelines.

Create enforcement mechanisms



Model License Language

Reduce legal friction and cost of negotiating with AI tool providers

Develop modular, machine-interpretable license clauses that specify allowed AI uses (training, RAG, summarization, embeddings).

Define clear attribution, compensation, retention, and revocation terms to reduce negotiation friction.

Enable composability so licenses can scale from single works to large, heterogeneous corpora



Reference Linking in AI Output

Standardize citation formats for AI-generated outputs, including persistent identifiers (DOIs, URLs, work IDs) .

Define confidence or provenance indicators distinguishing direct retrieval, paraphrase, and inferred synthesis.

Support bidirectional linking so outputs can be traced back to source content for verification and credit.



Agentic Access Control

Support subscription access as well as manage open access traffic loads on systems.

Define permission models governing what agents can read, write, execute, or delegate across systems.

Standardize scoped credentials and consent mechanisms for multi-step, tool-using agents.

Enable policy-based constraints to limit autonomous behavior by context, risk level, or domain.



Usage Tracking and Reporting

Differentiating between human and machine-based usage of content

Establish common metrics for tracking content access, transformation, and reuse by AI systems.

Define reporting formats that support auditing, royalty calculation, and regulatory review.

Enable privacy-preserving aggregation so usage can be measured without exposing proprietary or personal data.



Implementing Standards for AI Tool in Scholarly Publishing (One Example)




Model Context Protocol

Open, standardized interface that allows large language models (LLMs) to securely and predictably access external context—such as tools, databases, files, APIs, and services—at inference time

Protocol specifies how context is described, requested, authenticated, and returned, ensuring that models receive machine-readable, scoped context rather than unstructured prompts

Context is provided on demand, often with explicit permissions, schemas, and lifecycle controls, allowing enterprises to manage access, logging, and compliance

Separates the model's reasoning engine from the data and tools it relies on, improving modularity and governance



Concerns About MCP

Added architectural complexity – MCP adds a client-server layer, schemas, and permissions that can add system loads to simple prompt-based or inline tool integrations.

Latency and performance overhead – Repeated context discovery and tool calls over MCP can increase round-trip latency, that is problematic for real-time or high-throughput applications.

Fragmented ecosystem risk – While MCP aims to standardize, competing protocols and partial implementations can dilute interoperability, leading to “MCP-like” silos rather than true plug-and-play integration.

Operational burden on developers – Designing clean context schemas, managing versioning, and reasoning about what context to expose or withhold requires additional expertise and ongoing maintenance.

Limited guarantees on model behavior – MCP standardizes access to context, not how well models use it; models may still misuse, ignore, or hallucinate around provided context, limiting end-to-end reliability.



Open standards versus open protocols

De facto control vs. formal openness – Even if MCP is published as “open,” Anthropic’s role as primary designer, maintainer, and reference implementer gives it disproportionate influence over the protocol’s evolution, priorities, and interpretation.

Vendor lock-in by gravity, not mandate – Tooling, examples, and best practices may naturally optimize for Anthropic’s models and workflows, making MCP technically open but practically aligned with one vendor’s ecosystem.

Standards-capture risk – Other model providers, tool vendors, or enterprises may worry that design decisions subtly encode Anthropic’s assumptions about safety, alignment, or agent behavior, limiting neutral, multi-stakeholder governance.

Strategic asymmetry – Anthropic gains early visibility into how developers structure tools, data access, and agent architectures via MCP, potentially creating informational advantages over competitors.

Long-term governance uncertainty – Without an independent standards body, clear IP commitments, and transparent change control, adopters face risk that MCP’s direction could shift in ways that prioritize Anthropic’s business or safety philosophy over broader industry needs.




Be prepared to fix today's AI problems in the future.

Rapid adoption of early-stage technologies does come with risk.

It is likely that the worst AI tools you will ever use in the future are the one's you are using right now.

How many digitization projects had to be redone as the imaging technologies and storage capacity improved?

Don't necessarily hesitate, but plan to redo and improve in the future.



Thank you

Todd A. Carpenter

Executive Director

National Information Standards Organization (NISO)

tcarpenter@niso.org

ORCID: [0000-0002-8320-0491](https://orcid.org/0000-0002-8320-0491)

