Synthesis Report

Standards for Provenance, Attribution, and Usage in the AI Ecosystem

Workshops coordinated by

Workshops hosted by

This report synthesizes key insights and agreements from recent in-person discussion fora concerning the impact of generative and agentic AI systems on scholarly communications. Think of it as an executive summary for those who are too busy to read the full reports.

An unsurprising consensus evolved across the workshops that the infrastructure for content provenance, attribution, and usage measurement needs to adapt.

The Shift to Computational Readership

Traditional scholarly analytics rely on measurable human interactions such as downloads, views, and citations. AI systems are now mediating scholarly discovery by providing synthesized responses and consuming content in fragments (chunks or embeddings), which bypasses conventional readership signals and introduces retrieval noise. This disconnect necessitates developing a new conceptual model for computational readership.

The core challenge lies in the inadequacy of existing formats. The PDF format is insufficient for AI consumption. Full-text XML is more usable, but typically lacks identifiers and other metadata at a sufficiently granular level. Treating components like images, tables, and method sections as first class citizens is necessary, including a mechanism for metadata to consistently travel with these components.

Provenance and Attribution as Supply-Chain Infrastructure

Questions of provenance (where the AI got the content) and attribution (credit for the content's use by the AI tool) are critical for trust and transparency. There is broad agreement that base-level LLM attribution is currently intractable due to the computational scale and proprietary training data. Interventions focusing on the technology layer *above* the base model, such as Retrieval-Augmented Generation (RAG) or multi-agent systems, will be more practicable.

Provenance must be understood and addressed as a supply-chain infrastructure problem. The most actionable approach may be to define and adopt a minimum provenance payload (MPP). This would consist of a defined set of metadata fields that must travel with scholarly content when it is supplied to an AI system. Establishing this infrastructure is likely to be a prerequisite for reliable attribution.

Core Requirements for the Minimum Provenance Payload:

Preliminary technical requirements for this verifiable protocol include component-level tracking, applying provenance not just to whole content items but to individual assets within them. Mooted core fields are:

- Persistent Identifier (PID): DOI or equivalent.
- Content Hash: Crucial for verifying the retrieved object against the Version of Record (VoR). This directly addresses the research integrity issue of the DOI swap problem, in which an AI generates a valid DOI linked to the wrong paper.
- Digital Signature: To establish authenticity and detect modification of the content.
- Source Location and Access Timestamp: To track where and when the content was accessed.
- Version Information and Component Level Identifier: To support nested, traceable components and content granularity.

Adoption will be simplified if publishers take responsibility for wrapping existing or new publications by hashing and digitally signing the content, potentially through controlled API access.

Access Protocols and Economic Sustainability

Current publishing and licensing models, which rely on measurable human engagement, face significant uncertainty as AI-mediated usage obscures the true economic contribution of scholarly materials. Our focus is not on model training (weights), but instead on inference or agent harnesses (real-time access to content).

AI Endpoints and Gateways

AI Endpoints (APIs) may provide a technical and economic solution. These gateways offer several strategic advantages:

- Controlled Access: They standardize entitlement checks and access signaling.
- Monetization and Incentivization: They offer a pathway to monetize open access content and incentivize endpoint usage, which helps reduce DDOS implications.
- Efficiency: They enable token-efficient content delivery, positioning the AI as a librarian providing content via "subscribe to context" relationships.

We need standards around these endpoints, potentially building on existing efforts like Google's WebMCP. The goal is a common web entry point for AI agents.

Evolving COUNTER Usage Metrics

We already know that traditional usage metrics are insufficient for AI-driven usage, and COUNTER has published preliminary best practice to start addressing the issue. The lack of visibility into AI usage is particularly problematic for institutional collection decisions and justifying subscription value.

Our work to date has focused on meaningful use, defined as content that is specifically incorporated into the final AI-generated answer, rather than content merely retrieved or considered. Future metrics are likely to need to measure a broader range of computational activities:

- Component-Level Tracking: Measurement should extend beyond the COUNTER Item (article/chapter) to the component level. We have a mechanism for component reporting, and for rolling this up to Item and indeed Title level. It simply needs to be incorporated into our AI work.
- Holistic Assessment: Developing a metric for all content assessed by an AI system, in addition to the existing meaningful use metrics.

Crucially, the resulting metrics must roll up to the item (article/chapter) level to maintain value for library collection decisions. A top priority for COUNTER is determining how to bring external AI providers and third-party tools into a shared reporting infrastructure.

Key Commitments and Next Steps

The overarching goal is simplicity in standards to encourage adoption by tech companies, who will be incentivized by the need for accuracy and quality in their scholarly outputs.

Immediate Actions:

- Provenance and Attribution: NISO will take the lead on defining required fields for the minimum provenance payload and develop support for nested provenance relationships, specifically for the scholarly ecosystem.
- Usage: COUNTER will take the lead on advancing the AI usage metrics project for the same scholarly space. Interested parties from the workshops have been brought onto the COUNTER working group.
- Engagement: Engage with major technology companies and AI application organizations (e.g., Google Scholar, Consensus) to test and ensure adoption of any proposed standards. This includes defining the behavior of an existing endpoint and exploring the WebMCP initiative.

Caveats

COUNTER and NISO are lean organizations. If the community wants these projects to move quickly, we will require more resource than is currently available. If the feedback is that this work is a priority, Todd and Tasha will put together a business case and request for funding.