

Scite MCP

A Library Catalog The AI Can Consult, Not A Library It Can Carry Away

The Scite MCP connects AI assistants to Scite's citation database so researchers get real, verifiable papers instead of fabricated references. Search and retrieval only. Paywalled text stays behind the paywall.



How It Works: Five Facts

- 1 **No agreement, no full text.** Unless a publisher has an explicit Scite MCP amendment, no full-text content beyond openly licensed material (CC-BY and equivalent permissive licenses) is ever shown, not to the user and not to the AI tool they're using. Subscription and paywalled text stays behind the paywall, full stop. Even open access articles under restrictive licenses (non-commercial, no-derivatives) are excluded from AI-readable text.
- 2 **Every result carries a DOI.** Each article surfaced to the user includes its DOI and a resolved access link, routing the researcher to the version of record on the publisher's platform, through an institutional subscription, or via a legitimate per-article purchase. Scite sends readers to you. It doesn't replace you.
- 3 **Indexing drives discoverability.** Fully indexed articles (metadata, abstracts, and full-text citation snippets) consistently surface higher in AI-mediated results. Fully indexed content is more likely to appear in a researcher's results and drive link-outs to the publisher's version of record.
- 4 **Usage and denial data support subscription conversations.** Scite provides publishers and institutional customers with the same organization-level and journal-title-level reporting: which journals researchers reached through AI tools, and which they tried to reach but couldn't. Turned-away demand becomes visible evidence for new subscriptions and per-article sales, not lost signal.
- 5 **Everything is logged and read-only.** Access is limited to authenticated Scite subscribers, individually credentialed, usage-capped, and fully auditable.

Common Questions

Does Scite run an AI model on publisher content? No. All language generation happens inside the user's own AI tool. Scite returns structured search results, performs no AI model training of any kind, and publisher content is never used to train AI models, by Scite or by anyone through Scite.

Scite indexes paywalled full text. Isn't that the same as showing it? No, and this is the most common misunderstanding. Indexing lets search look inside the body of an article to locate the right paper. Searching an article is not the same as displaying its text. A paywalled article can be found through full-text search, but its text is never shown to the AI or the user through the MCP.

What does the AI actually receive about an article? Bibliographic metadata (title, authors, journal, DOI), the abstract where the publisher has made one available, citation counts, Smart Citation classifications (whether citing papers support, contrast, or mention the work), and a link to the appropriate access route.

How does a researcher reach paywalled full text? Through the resolved link: the publisher's version of record, an institutional subscription, or a per-article purchase through Article Galaxy, always under existing subscription or purchase terms. Where AI use of an article is specifically needed, Article Galaxy offers a per-article GenAI Rights license from participating publishers, creating a new, rights-respecting revenue channel.



Sharon Mattern Büttiker

Vice President of Publisher Relations & Content Strategy
smattern@researchsolutions.com