

MCP SERVER

NO CODE

CLOUD HOSTED

# Semantic Scholar MCP

Trace academic influence from single keywords to full citation graphs.

Stanford Semantic Scholar provides AI-powered access to the world's largest academic knowledge graph. Use this MCP to search millions of papers, track citation chains, analyze author impact metrics like the h-index, and discover foundational research related to your topic.

**A+** Quality Score 98.33/100

semantic-scholar

academic-papers

citations

research

literature-review

bibliography



# The connectivity layer between AI and the world's software.



Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

# Your AI Connections Run Through Vinkius Cloud

The world's largest  
managed MCP catalog

Vinkius is the connectivity layer where AI connects to the software your business already runs. We handle the hosting, the security, the credentials, the uptime — you get agents that actually do things.

We operate the world's largest managed MCP catalog. Major SaaS platforms, CRMs, databases, and cloud providers — running, monitored, production-ready. This MCP server is hosted and maintained by the Vinkius Cloud for AI Agents.

*The agent doesn't manage credentials, doesn't manage uptime, doesn't manage security. Vinkius does.*

— Architecture principle

---

## Four Pillars of the Vinkius Runtime

### 01 — Security by design

Credentials stay encrypted at rest via AES-256. The AI agent never touches raw keys — they're injected into a sandboxed V8 isolate at runtime. Actions are logged, and connections have an emergency kill switch.

### 03 — Deterministic observability

Eight immutable metrics per endpoint: request volume, p95 latency, error rate, active connections, cost attribution. A live payload feed logs every tool call with mutation detection.

### 02 — Built on MCP Fusion

This MCP server was built with **MCP Fusion**, the open-source framework (Apache 2.0) that powers the entire Vinkius catalog. Schema-as-firewall strips undeclared fields, compiled PII redaction runs at zero overhead, and cryptographic lockfiles produce git-diffable audit trails.

### 04 — Autonomous operations

Servers are deployed, monitored, and patched autonomously. New capabilities and security patches ship weekly. Zero-downtime deployments ensure continuous availability across all managed MCP servers.

**AES-256**

Encryption at rest

**Ed25519**

PKI vault signatures

**24h TTL**

Ephemeral session keys

**V8 Isolate**

Sandboxed execution

---

## One Token. Instant Access.

Every MCP server on Vinkius is accessed through a **Connection Token**. Tokens are generated in the cloud dashboard and produce a unique MCP endpoint URL. Paste this URL into any MCP-compatible client — no SDK required.

A single token can serve **multiple AI clients simultaneously**, or you can issue separate tokens per client for granular access control. Each token tracks its own request count, last activity timestamp, and can be individually enabled or revoked.

MCP ENDPOINT

`https://edge.vinkius.com/{token}/mcp`

Claude



Cursor



VS Code



Windsurf



Grok



Gemini

---

## Security Is the Architecture

Security in Vinkius is not a feature — it's the foundation of the runtime. The gateway enforces multiple independent protection layers between AI agents and third-party APIs.

### 01 — Ed25519 PKI Vault

Every workspace has an Ed25519 Master Key. Session keys are generated ephemerally (24h TTL) and signed by the Master Key. Credentials never leave the vault boundary.

### 02 — V8 Isolate Sandboxing

Tool code runs inside isolated-vm V8 isolates with 64 MB memory caps and per-request timeouts. No filesystem access, no network access except through the SSRF-guarded fetch bridge.

### 03 — SSRF Guard

All outbound HTTP requests are DNS-resolved and validated before execution. Private IP ranges (10.x, 172.16-31.x, 192.168.x, AWS metadata 169.254.x) are blocked at the network layer.

### 05 — Cryptographic Audit Trail

Every request is signed into a SHA-256 hash chain with Ed25519 signatures. Events form a tamper-proof, SIEM-exportable forensic record.

### 04 — DLP & PII Redaction

A ResponseGuard pipeline intercepts every tool response. Configurable redaction patterns strip sensitive fields (emails, SSNs, card numbers) before data reaches the AI agent.

### 06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot is used outside Vinkius infrastructure, the server is quarantined instantly.

## Emergency Kill Switch

EU AI Act Art. 14(1)  
Compliant

The kill switch is an **emergency halt** mechanism — not a simple toggle. When triggered, it executes three actions atomically:

#### 01 — Server deactivated

The MCP server is immediately taken offline across the entire cluster.

#### 02 — All tokens revoked

Every connection token is invalidated. Total lockout — reconnection blocked until new tokens are issued.

#### 03 — WebSocket connections killed

Active connections terminated via Redis pubsub broadcast. Propagates to every runtime node in the cluster.

## Full Visibility. Zero Guesswork.

The Vinkius cloud dashboard includes a full MCP Governance suite — real-time analytics and security controls for production AI operations.

**Control Plane**

KPI dashboard with request volume, latency, success rate, token consumption, and AI-generated operational briefings.

**FinOps**

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

**Firewall & DLP**

PII redaction activity, sensitive data protection counters, and security event timeline.

**Agent Activity**

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

**Tool Health**

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

**Incident Log**

Error trends, failure rates, status-code breakdowns, and forensic audit trail access.

Get started at [cloud.vinkius.com](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# Stanford Semantic Scholar MCP

16 tools available

Cloud-hosted on Vinkius

Need deep context on a scientific or technical subject? This MCP connects you directly to Semantic Scholar's massive academic database. It lets your agent go beyond simple keyword searches to understand the actual *context* of published work. You can trace how an idea evolved by finding every paper that cited a key source, or conversely, see what foundational papers influenced a modern breakthrough. Need to review dozens of authors? Use this MCP to quickly pull author metrics, seeing their total citation count and h-index without leaving your agent client. When you connect this through Vinkius, your AI can handle the entire literature review process—from identifying key seminal works to building out bibliometric reports on demand. It's how you get deep academic insight into your workflow.

---

## Core Capabilities

### 01 — Search papers by criteria

Find relevant articles using keywords and filtering the results by specific fields, years, or top-tier journals.

### 03 — Trace citation history

Map the intellectual lineage of a paper by finding both its citing works (forward citations) and the papers it references (backward citations).

### 05 — Process batches of metadata

Handle large lists of papers or authors by pulling all necessary metrics in one single request for efficient analysis.

### 02 — Analyze author impact metrics

Retrieve detailed professional profiles for researchers, including their total publication count and h-index score.

### 04 — Get AI-powered recommendations

Discover highly relevant, yet unfamiliar, research using algorithms that analyze content similarity across multiple source papers.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/stanford-semantic-scholar](https://vinkius.com/mcp/stanford-semantic-scholar) — connect your AI agent in three steps.

- 01 Connect your preferred AI client to this MCP via Vinkius.
- 02 Direct your agent to use a search tool, providing the paper title, author name, or specific ID (like a DOI).
- 03 Your agent executes the query and returns structured data: full abstracts, metrics, citation counts, and links to related work.

The bottom line is you get deep academic graph analysis right inside your chat client without needing an API key or running local scripts.

---

## Built For

This MCP is essential for researchers, data scientists doing bibliometric analysis, and R&D teams who need to prove a concept's intellectual foundation. If your job involves proving *\*why\** an idea matters or tracking the evolution of a technology, you need this.

### Academic Researcher

Needs to conduct systematic literature reviews and identify foundational papers for grant proposals.

### Data Scientist (Bibliometrics)

Builds publication analytics, tracking trends in specific fields or measuring the collective impact of a research group.

### R&D Team Lead

Monitors competitors' latest publications and tracks the state-of-the-art across complex technical domains.

---

## What Changes When You Connect

- 01 Stop guessing research gaps. Use `get_multi_recommendations` to find highly relevant, non-obvious papers that build on your existing literature set.

- 
- 02 Build robust bibliometric reports instantly. The `batch_get_authors` tool lets you analyze multiple researcher profiles—including h-index and citation counts—in one shot.

---

  - 03 Understand the full life cycle of an idea. Use `get_paper_references` to find the original foundational work, or use `get_paper_citations` to see how it influenced later research.

---

  - 04 Speed up systematic reviews. The `bulk_search_papers` tool handles massive result sets with continuation tokens, ensuring you never miss a paper in your review set.

---

  - 05 Pinpoint credibility fast. Search by venue lets you filter results only down to top-tier conferences like Nature or Science, guaranteeing high quality.
- 

---

## Real-World Applications

### Proving the scope of a research problem

An R&D team needs to prove that their new AI model addresses an existing gap. They use `get_paper_references` on key papers in the field and then run `search_by_field` (Computer Science) combined with searching by year (2015-2020) to define exactly what research was done before them.

### Building a comprehensive literature review

A student needs to write a survey of transformer architectures. They start with 'Attention Is All You Need' using `get_paper_citations` to find all subsequent work, and then use `get_recommendations` to discover the next key papers they must read.

### Assessing a collaborator's standing

A PI needs to evaluate a potential co-author. They use `get_author` and `search_authors` on the candidate, immediately seeing their total paper count, h-index, and key affiliations before committing to collaboration.

### Comparing multiple related works

A data scientist needs to compare three different models (e.g., BERT, GPT-2, ViT). They feed all three unique identifiers into `batch_get_papers` to pull and analyze the full metadata simultaneously.

---

## Patterns to Avoid

---

### Treating citation counts as proof of quality

#### ✗ AVOID

Pulling every paper found using `search_papers` just because it has a high citation count. This floods your review with irrelevant, foundational work.

#### ✓ INSTEAD

Instead, use `get_paper_citations` on the most relevant papers to track how an idea evolved. Supplement this by filtering results using `search_by_venue` for known top-tier journals.

---

### Relying only on keywords

#### ✗ AVOID

Searching for 'deep learning' and getting thousands of unrelated, general papers that don't address your specific mechanism (e.g., GANs vs VAEs).

#### ✓ INSTEAD

Be precise with your starting point. Use `get_recommendations` from one highly specific seed paper to narrow the focus down to only adjacent, specialized research.

---

### Manually checking multiple authors' records

#### ✗ AVOID

Opening a new tab for every author and running separate queries just to find their total citation count or h-index.

#### ✓ INSTEAD

Use `batch_get_authors`. This single tool allows you to input a list of researchers and get all the key metrics in one efficient call.

---

## The Right Fit

Use this MCP if your research needs deep graph analysis, meaning you need to understand *relationships* between ideas, authors, or papers. You need to know: 'Who influenced this work?' or 'What came next because of this idea?'. If your goal is simply finding general information—like a basic fact about a paper's abstract or title—then the standard search engine is fine. However, if you are doing systematic literature reviews, building bibliometric models, or trying to prove intellectual lineage for a grant, this tool is necessary. Don't use it if all you need is a general definition; that doesn't require academic graph traversal. If you only have titles and no IDs, remember to run `match_paper_title` first to guarantee metadata accuracy before pulling full details with `get_paper`.

---

## The Citation Trail: How hard it is to track academic influence today

Writing a literature review often feels like forensic accounting. You find one key paper, and then you have to manually jump through dozens of links—checking who cited it (forward) and what papers it relied on (backward). This process is tedious: copying DOIs into different search interfaces, opening endless tabs, and painstakingly cross-referencing citation counts just to map the intellectual flow.

With this MCP, your agent handles the entire traversal. You point it at a core paper, tell it to find all citing works using `get_paper_citations`, and then ask for those papers' references using `get_paper_references`. Suddenly, you don't just have data; you have a clear map of academic influence.

---

## Get Author Profiles: Deep Insight with the Stanford Semantic Scholar MCP

Before this, assessing an author's impact meant visiting their personal website or digging through Google Scholar's limited view. You got a snapshot—a handful of papers and maybe one citation number. It was never comprehensive enough for serious analysis.

Now, you can run `get_author` on any researcher. The agent pulls structured data showing the full h-index, total paper count, and all major affiliations in one clean output. This gives you a definitive measure of their academic weight instantly.

---

# Stanford Semantic Scholar with 16 Tools

These tools let your agent search for papers across millions of records, analyze author metrics, trace citation histories, and pull bulk metadata in structured formats.

#	TOOL	DESCRIPTION
01	<code>batch_get_authors</code>	Retrieves multiple author profiles, providing their names, affiliations, paper counts, citation counts, and h-indices at once.
02	<code>batch_get_papers</code>	Accepts lists of IDs (DOIs, ArXiv, PubMed) to retrieve full metadata for multiple papers in a single call.
03	<code>bulk_search_papers</code>	Searches for very large result sets of academic papers and returns continuation tokens so you can process all results systematically.
04	<code>get_author</code>	Pulls a definitive profile for one author, detailing their affiliations, total paper count, citation count, and h-index.
05	<code>get_author_papers</code>	Retrieves every paper by a specific author, listing titles, years, venues, and whether the work is open access.
06	<code>get_multi_recommendations</code>	Generates focused literature suggestions by finding papers similar to a set of positive sources but unlike a set of negative ones.
07	<code>get_paper</code>	Fetches all details for a single paper using multiple identifiers, including DOI, ArXiv ID, or PubMed ID.
08	<code>get_paper_authors</code>	Identifies the contributing authors of a specific article and provides their individual metrics like h-index.
09	<code>get_paper_citations</code>	Finds all follow-up work by listing metadata for papers that cite a given source, showing how an idea was used later.
10	<code>get_paper_references</code>	Determines the intellectual roots of a paper by listing the foundational works it cited when it was written.
11	<code>get_recommendations</code>	Uses content similarity and citation patterns to suggest the most relevant papers you should read next, based on one seed article.
12	<code>match_paper_title</code>	Finds the correct paper metadata when you only have a slightly misspelled or generalized title string.
13	<code>search_authors</code>	Searches across the academic graph to locate researchers by name, providing their full profiles and metrics.

#	TOOL	DESCRIPTION
14	search_by_field	Filters available papers to only include those that fall within a specific discipline like Medicine or Computer Science.
15	search_by_venue	Narrows down the search results to publications from specific, high-impact conferences like Nature or NeurIPS.
16	search_papers	Performs a broad keyword search across 200 million papers, allowing filtering by year, field, and journal.

---

## See It in Action

Real prompts you can use once this MCP is connected to your AI agent through Vinkius Cloud.

### U Find the most cited papers on transformer architectures published since 2020



I've searched Semantic Scholar for "transformer architecture" papers from 2020-2026. The top results include "Attention Is All You Need" (the foundational paper), Vision Transformer (ViT), BERT, GPT-3, and their derivatives, sorted by citation count.

### U What is Geoffrey Hinton's h-index and how many papers has he published?



I've found Geoffrey Hinton's profile on Semantic Scholar. He has published over 400 papers with a combined citation count exceeding 500,000, giving him one of the highest h-indices in computer science.

### U Recommend papers similar to "Attention Is All You Need"



Using the AI recommendation engine with "Attention Is All You Need" as the seed paper, I've found highly relevant papers including BERT, GPT-2, the Universal Transformer, Transformer-XL, and other key works that built upon the attention mechanism paradigm.

---

## Frequently Asked Questions

### 01 How do I find related papers using Semantic Scholar MCP?

Use `get_recommendations` or `get_multi_recommendations`. You feed the tool one or more seed papers, and it analyzes content similarity to suggest relevant literature you might not know exists.

---

**02 Can I search for papers by a specific journal using Semantic Scholar MCP?**

Yes, use `search_by_venue`. You simply name the conference or journal (like Nature or CVPR) and filter all searches to only include articles published there.

---

**03 What if I don't have a DOI for a paper?**

No problem. Try `match_paper_title` first; it uses fuzzy logic to find the correct metadata even if your title is slightly off or incomplete. You can then use `get_paper` with the found ID.

---

**04 How do I compare multiple authors' work?**

The `batch_get_authors` tool is designed for this. Give it a list of names, and you get all their key metrics (h-index, citations) in one request for easy comparison.

---

**05 Is Semantic Scholar MCP limited to Computer Science research?**

Not at all. The tool supports searching across major fields like Medicine, Biology, Physics, and Economics, giving you a massive scope of academic knowledge.







---

# Go Live in 60 Seconds

Get your connection token from [cloud.vinkius.com](https://cloud.vinkius.com), then paste the endpoint URL into any MCP-compatible client.











YOUR MCP ENDPOINT

```
https://edge.vinkius.com/[TOKEN]/mcp
```

CLIENT	WHERE TO CONFIGURE
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"stanford-semantic-scholar": { "url": "..." }</code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	Extensions → Add MCP Server → Paste endpoint URL

## ASK AN AI ABOUT THIS

Let your preferred AI explain this MCP server

-  **Ask ChatGPT** 
-  **Ask Claude** 
-  **Ask Perplexity** 
-  **Ask Gemini** 
-  **Ask Grok** 

READY TO CONNECT

# Stanford Semantic Scholar is live on Vinkius Cloud.

Get your connection token, paste it into your AI agent, and  
start building. No SDK. No deployment. Just results.

[Start at cloud.vinkius.com](https://cloud.vinkius.com) →

[vinkius.com](https://vinkius.com) · [support@vinkius.com](mailto:support@vinkius.com)

### INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by Stanford Semantic Scholar. All third-party trademarks, logos, and brand names are the property of their respective owners. Their use in this document is strictly for informational purposes to identify service compatibility and interoperability.

### DOCUMENT INFORMATION

Generated	June 2026
MCP Server	Stanford Semantic Scholar MCP
Server ID	019dea61-0650-72a8-a7bc-23543ab1a32b
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### LICENSE & USAGE

This document is generated automatically by the Vinkius PDF Engine. Content reflects the MCP server configuration at the time of generation and may change as updates are deployed. For the most current information, visit [vinkius.com/mcp/stanford-semantic-scholar](https://vinkius.com/mcp/stanford-semantic-scholar).