

MCP SERVER

NO CODE

CLOUD HOSTED

# Stanford OpenAlex MCP

Map the entire academic knowledge graph.

Stanford OpenAlex gives you open access to the global research catalog. It lets you analyze over 250 million academic works and trace connections between authors, institutions, concepts, and funding bodies worldwide. Check a researcher's full impact profile, map out scientific concept hierarchies, or find every paper funded by a specific organization.

**A+** Quality Score 98.33/100

openalex

academic-research

bibliometrics

institutions

journals

open-access



# 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 OpenAlex MCP

16 tools available

Cloud-hosted on Vinkius

Working with modern academic research means dealing with massive datasets—millions of papers, thousands of institutions, and complex citation graphs. This MCP connects you directly to the OpenAlex API, giving your AI client deep access to this global knowledge catalog. You don't need an API key because OpenAlex is free and open. Instead of manually hopping between Google Scholar, university websites, and journal portals, you ask your agent to do the heavy lifting. You can profile a scholar using metrics like h-index or i10-index; check which organizations are funding specific fields of study; or narrow down papers that have been published under open access licenses. It's the power of having all the academic world's metadata available in one place, accessible through Vinkius and your preferred AI client.

---

## Core Capabilities

### 01 — Assess Academic Impact

Retrieve detailed profiles for authors, including their total citation count, h-index, i10-index, and historical publication trends.

### 03 — Analyze Funding Streams

Determine which specific organizations fund research, allowing you to map funding sources against academic topics.

### 02 — Map Research Connections

Explore how scientific concepts relate to one another by viewing the concept hierarchy or tracing related research areas.

### 04 — Deep Dive into Works

Get comprehensive details on any paper using its DOI, OpenAlex ID, or PubMed ID, including abstracts and citation information.

# One Click on Vinkius — From Prompt to Execution

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

- 01** Subscribe to this MCP within Vinkius. No API key is required because the source data is openly available.
- 02** Connect your AI client (Claude, Cursor, etc.) to the catalog via the MCP connection point.
- 03** Ask your agent a specific question—like 'Show me all papers on deep learning funded by the NSF'—and it executes the necessary queries.

The bottom line is you get structured data about global research fields and people, delivered directly to your AI client for immediate analysis.

---

## Built For

This MCP serves academic libraries, university administrators, bibliometricians, and science policy makers. It's for anyone whose job requires tracing the provenance of ideas—from funding sources back to original authors.

### Research Librarian

Checks journal details using ``get_source`` or filters search results via ``search_sources`` to determine if a required paper is open access.

### Academic Administrator

Benchmarks different institutions by running searches through ``search_institutions`` and comparing citation metrics across multiple schools.

### Policy Analyst

Understands the global flow of money into science using ``get_funder`` or ``search_funders`` to map which organizations are prioritizing specific research areas.

---

## What Changes When You Connect

- 01** Pinpoint scholarly impact immediately. Instead of compiling CVs from multiple sites, use `get_author` to instantly gather a scholar's full citation history and h-index metrics.

- 
- 02 Trace research funding paths. Use `search_funders` or `get_funder` to understand which organizations are investing in specific topics, helping policy makers track global science money.

---

  - 03 Find literature easily. If you only need free papers, use `search_open_access`. This saves time and eliminates the frustrating step of paywall checking before citing a source.

---

  - 04 Understand field relationships. When analyzing a niche topic, run `get_concept` to see its parent fields and related concepts, giving you context beyond just keywords.

---

  - 05 Benchmark institutions accurately. Compare multiple universities' research strength by querying `search_institutions`, letting you compare works counts or citation metrics side-by-side.
- 

---

## Real-World Applications

### Assessing a Collaborator's Standing

A PI needs to prove that a new collaborator has deep expertise in AI. They ask their agent to use `search_authors` with filters like `'works_count:>500'` and then run `get_author_works` to confirm the most cited papers are all within neural network concepts.

### Curating a Review Article

A librarian is writing an overview of CRISPR technology. Instead of searching journal by journal, they use `get_work` multiple times with different IDs and then use `search_sources` to ensure the journals are high-impact before adding them.

### Tracking Policy Shifts

A policy maker needs to know if climate research funding is shifting from government sources to private foundations. They use `search_funders` and then cross-reference the results with papers found using `search_works` filtered by `'climate change'`.

### Mapping a New Field

A PhD student is exploring the boundaries of quantum computing. They ask their agent to run `search_topics`, which maps out the full concept hierarchy, showing related fields like particle physics and high-energy theory.

---

# Patterns to Avoid

---

## Using Google Scholar for everything

### X AVOID

A user copies a list of 50 DOIs from one database and pastes them into another, hoping to get comprehensive metrics like h-index or concept mapping. This fails because different sites use different data models.

### ✓ INSTEAD

You must centralize your queries using ``get_work`` for individual papers, but for bulk analysis, run structured searches via ``search_authors`` and ``search_works``. These tools consolidate the raw metadata you need.

---

## Searching only by keywords

### X AVOID

A user runs a keyword search on one platform and gets 20 results. They then manually try to find out which institutions are involved, requiring dozens of clicks.

### ✓ INSTEAD

Use ``search_works`` with structured filters like `'publication_year:2024' AND 'type:journal-article'`. Then, use the resulting list to check institutional details via ``get_institution`` for a complete picture.

---

## Assuming data is siloed

### X AVOID

A researcher finds an author on one platform and assumes their work history is only visible there. They miss out on funding context.

### ✓ INSTEAD

Always run ``get_author`` first, then immediately use the resulting concepts to cross-reference with ``get_funder``. This connects the person's output directly to who paid for it.

---

## The Right Fit

Use this MCP if your core need is understanding the *structure* of academic knowledge—who published what, where, when, and why (funding). You need to trace connections: author to work, work to concept, or funding body to field. Don't use it if you just need a simple list of papers by title; `search_works` handles that. However, don't rely on this MCP for the most recent news articles, either—it tracks established scholarly literature. If your goal is purely conversational summary generation without needing structural metrics (like h-index or concept hierarchy), then a general LLM might suffice. But if you need verifiable data points—citation counts, specific concepts, funding bodies—this catalog is essential.

---

---

## The Pain of Manual Academic Due Diligence

Right now, proving the impact or lineage of an idea is a nightmare. You start by Googling papers, then jump to CrossRef to get the DOI, then open a university website just to check if that author is still affiliated there, and finally, you cross-reference funding data from another dashboard. It's copy-pasting metadata across five different tabs just to build one citation graph.

With this MCP, your agent does it all in one query. You ask for the full picture—the author metrics, the linked concepts, and the primary source—and you get a single, structured data response. It's instant due diligence.

---

## Using Stanford OpenAlex to Pinpoint Impact

The tedious steps of cross-checking publication years against institutional affiliations or manually verifying the open access status are gone. You run a search and get all that data simultaneously.

Your workflow changes from manual information retrieval to focused analysis. You stop collecting data points, and you start building arguments.

---

# Stanford OpenAlex: 16 Tools for Academic Data

Use these specific tools to query the massive catalog directly, whether you need an author's impact score or a journal's citation count.

#	TOOL	DESCRIPTION
01	<code>get_author</code>	Retrieves a scholar's full profile, including citation count, h-index, i10-index, and publication trends over time.
02	<code>get_author_works</code>	Lists all papers written by a specific author, letting you sort the results by most cited or newest date.
03	<code>get_concept</code>	Provides detailed information on scientific concepts, including their ancestors and related fields of study.
04	<code>get_funder</code>	Shows which organizations fund specific research areas along with statistical details about the funding.
05	<code>get_institution</code>	Gathers detailed metrics for a university or lab, including its overall research output and collaborative partners.
06	<code>get_source</code>	Evaluates the quality of journals or conferences by retrieving their impact metrics and coverage details.
07	<code>get_work</code>	Fetches a paper's title, abstract, authors, and citation count using its DOI, OpenAlex ID, or PubMed ID.
08	<code>search_authors</code>	Searches the database of over 90 million authors by name and filters results based on citations or country code.
09	<code>search_concepts</code>	Finds specific scientific concepts within the knowledge hierarchy, showing associated works counts and citation rates.
10	<code>search_funders</code>	Searches for funding organizations worldwide, detailing their grants count, funded works, and overall citation impact.
11	<code>search_institutions</code>	Finds research institutions across the globe by name or filters them by type, country, or high citation counts.
12	<code>search_open_access</code>	Filters all available works to show only papers that have freely accessible PDFs for reading and meta-analysis.
13	<code>search_publishers</code>	Analyzes the academic publishing market by listing major publishers, their country of origin, and citation counts.

#	TOOL	DESCRIPTION
14	search_sources	Searches over 240,000 journals and conferences, providing details like ISSNs and open access availability.
15	search_topics	Maps the entire landscape of science by searching topic classifications across all domains and parent fields.
16	search_works	Performs a full-text search or structured filter on 250 million works, allowing sorting by relevance, citation count, or date.

---

## See It in Action

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

### **U** Which universities have the highest research output in AI?



I've searched OpenAlex for institutions with the highest works count in artificial intelligence concepts. The top institutions include MIT, Stanford University, Carnegie Mellon, UC Berkeley, and Tsinghua University.

### **U** What are the most cited open access papers on CRISPR?



I've found the most cited open access papers on CRISPR from OpenAlex, sorted by citation count. The foundational papers by Doudna and Charpentier lead the list.

### **U** Show me the concept hierarchy for machine learning



I've retrieved the Machine Learning concept from OpenAlex. It sits at level 1 under Computer Science (level 0), and has child concepts including Deep Learning, Neural Networks, Reinforcement Learning, and Natural Language Processing at finer granularity levels.

---

## Frequently Asked Questions

### **01** How do I use Stanford OpenAlex for citation metrics?

Use the ``get_author`` tool. This retrieves a scholar's full profile, including their total citation count, h-index, and i10-index, giving you immediate impact assessment.

### **02** Can I find out who is funding a specific research topic using Stanford OpenAlex?

Yes, use ``get_funder`` or ``search_funders``. These tools map funding organizations to specific research areas, allowing you to trace the financial backing of science.

---

**03 What is the best way to find open access papers with Stanford OpenAlex?**

The dedicated `search\_open\_access` tool filters all works and returns only those academic papers that have freely available PDFs, saving you time on paywalls.

---

**04 How do I compare multiple universities' research output using Stanford OpenAlex?**

Use `search\_institutions` to find the schools. Then, for a direct comparison of metrics across them, run `get\_institution` on each one you want to benchmark.

---

**05 Does Stanford OpenAlex cover patents or just academic papers?**

It covers works that include papers, books, datasets, and patents. Use the `search\_works` tool with specific filters if you are looking for a particular type of output.







---

# 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-openalex": { "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 OpenAlex 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 OpenAlex. 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 OpenAlex MCP
Server ID	019dea60-e5a2-7074-8986-f7313b31a150
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-openalex](https://vinkius.com/mcp/stanford-openalex).