



The Document Editor with a *Logic Validation Engine*

AI writes fast. Then breaks its own logic. We catch it.

Columbia SEAS '28 · Austin Wijaya, Akito Yamauchi, Hayden Leung

• Real-Time Logic Linting

• Knowledge Graph

• AI Verification



01 · THE PROBLEM

AI-assisted research is everywhere. *So are its artifacts.*

In February 2024, a peer-reviewed *Frontiers* paper on stem-cell biology made it to print with Midjourney-generated figures, including a rat with anatomically impossible anatomy and gibberish labels, and survived peer review for three days before retraction.^[1] Authors disclosed the AI use. A reviewer flagged concerns. No tool caught it.

768

AI ARTIFACTS IN PRINT

Published papers flagged for chatbot tells like “regenerate response”, likely a small fraction of the true total.^[2]

19.9%

FABRICATED CITATIONS

When researchers checked 176 citations GPT-4o generated for mental-health literature reviews, 19.9% were fabricated, the cited papers didn't exist.^[3]

1 of 12

A CORE GAI RISK

NIST, the US federal standards body, names “confabulation” as one of 12 core risks in its 2024 Generative AI Profile. The exact failure mode we address.^[4]

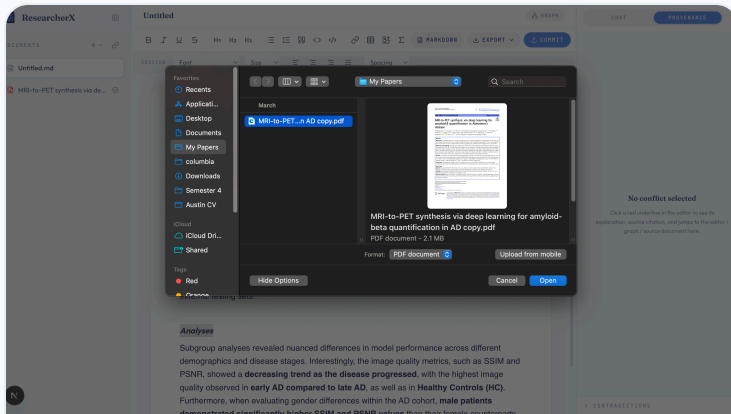
[1] Frontiers retraction, 13–16 Feb 2024 · [2] Glynn, *Suspected Undeclared Use of AI in the Academic Literature*, arXiv:2411.15218, Nov 2024. N=768 cases ·

[3] Linardon et al., *JMIR Mental Health* 2025, N=176 citations · [4] NIST AI 600-1, Jul 2024. doi.org/10.6028/NIST.AI.600-1

02 · THE SOLUTION

Your editor, with a *Logic Validation Engine*

We didn't build another AI writer. We built the thing that checks it.



1

Import & Write

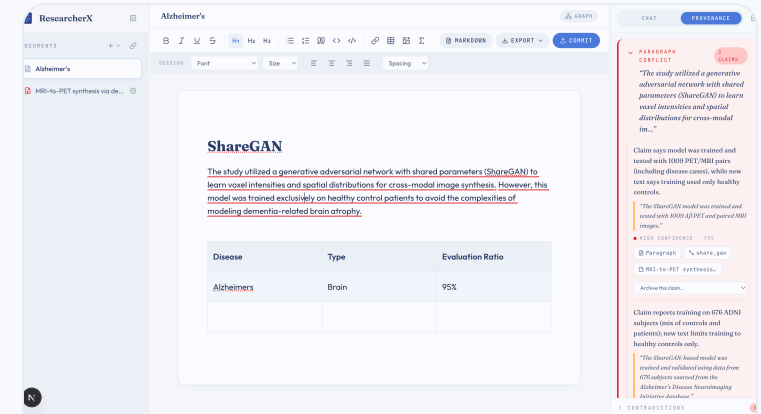
Upload PDFs, paste URLs, or draft directly. Sources are parsed and semantically chunked.



2

Graph Builds As You Work

Imported PDFs extract automatically. Live drafts build paragraph by paragraph, typed entities, natural-language relationships.



3

Conflicts Flagged Instantly

Red underline for high-confidence conflicts, orange for likely ones. Click for source, explanation, and full provenance.

INPUT
Type paragraph



EXTRACT
Claims extracted



RETRIEVE
Graph queried



COMPARE
Conflicts surfaced



OUTPUT
Red underline

04 · COMPETITIVE LANDSCAPE

Why *we win*

No existing product verifies the facts in a document, against the document itself or its sources.

We occupy a new category.

CAPABILITY	ELICIT (\$10-65/MO)	COCOUNSEL (\$110+/MO)	NOTION AI (\$10+/MO)	RESEARCHERX
Cross-doc contradiction detection	×	×	×	✓ Core
Real-time verification	×	×	×	✓ Core
Graph-anchored provenance	Partial	✓ (Westlaw)	×	✓ (Neo4j)
Integrated drafting editor	×	Partial	✓	✓
Works with your own sources	Partial	✓	✓	✓

Elicit finds papers. CoCounsel researches case law. Notion AI writes prose. **None of them verify the facts.**

07 · THE BIGGER PICTURE

WITH FUNDING · Q4 2026

The Foundational *Truth Graph*

We'll extract **billions of arXiv and PubMed papers** into a structured graph of what science actually claims, and serve it as an **MCP server**, the grounding layer frontier AI plugs into.

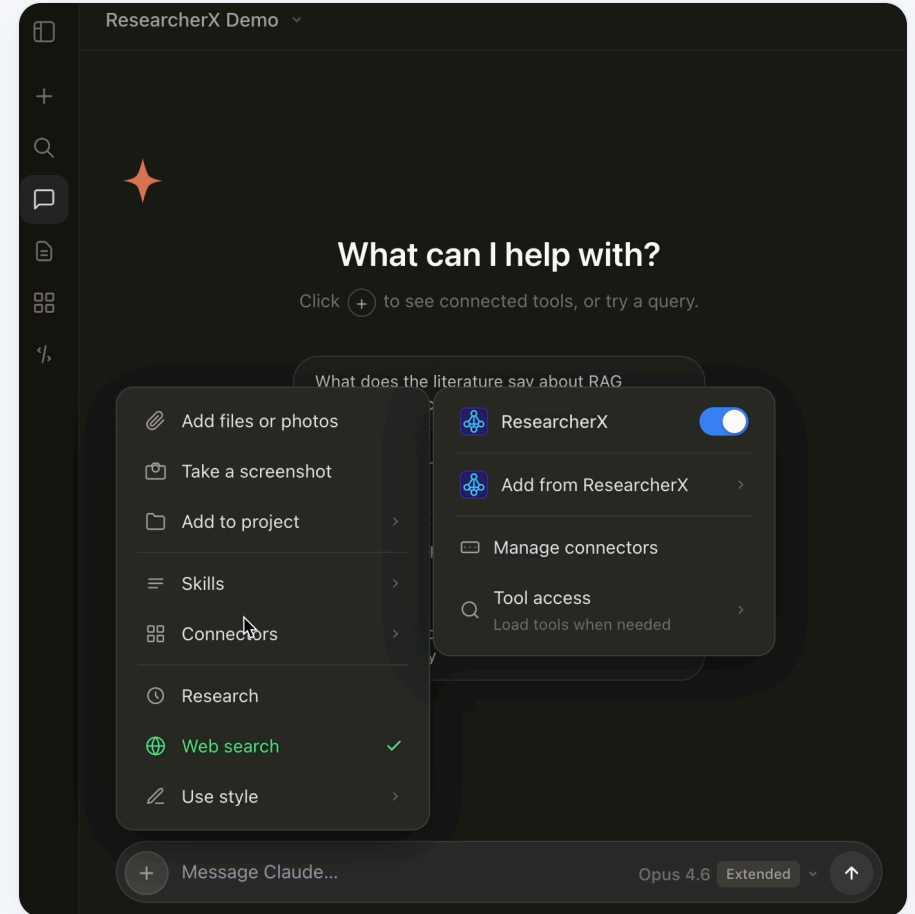
No one else is building this.

• arXiv

• PubMed

• MCP Server

• Data Licensing



05 · MARKET OPPORTUNITY

Adoption is running ahead of *trust*

66% use AI intentionally · only 46% trust it. ^[8] Adoption is outpacing verification, we're the layer that makes AI safe to ship.

TAM · TOTAL ADDRESSABLE

\$9.6B

~40M PEOPLE · @ \$240/YR

Global fact-sensitive knowledge workers: academics, lawyers, journalists, medical writers, analysts.

SAM · SERVICEABLE ADDRESSABLE

\$768M

3.2M PEOPLE · @ \$240/YR

US STEM researchers (2M) + non-STEM (200k) + legal professionals (1M).

SOM · OBTAINABLE

\$38M

~157K PEOPLE · @ \$240/YR

Columbia research-active (~7k) + US conference-reachable researchers (~150k).

[8] Gillespie et al., *Trust, Attitudes and Use of AI: A Global Study 2025*. KPMG × Univ. of Melbourne. N=48,340 across 47 countries.

06 · BUSINESS MODEL

Tiered SaaS with *clean margins*

Free

\$0 / month

200k cloud tokens (~\$1 of compute). After exhaustion, local inference via Ollama keeps the editor functional at zero cost.

Pro

\$20 / month

Unlimited documents, premium cloud LLMs, full knowledge graph + chat, LaTeX support, multi-user collaboration.

Enterprise

Custom pricing

Shared team knowledge graphs, SSO, on-prem deployment, bulk corpus ingestion, API access to your team's graph.

PRO TIER · MONTHLY COST / USER

// Illustrative: Claude Haiku^[9]
// \$1.00 in / \$5.00 out per 1M tok

KB INGESTION

50 PDFs/mo: **\$2.25**

LINTING

~500 lints/mo: **\$1.25**

CHAT

~200 queries/mo: **\$0.90**

Embeddings (local): **\$0.00**

Neo4j + LanceDB: **\$0.00**

COGS / user / mo: **~\$5.00**

Margin @ \$20/mo: **~75%**

// Provider flexibility: Claude · Gemini · Ollama

[9] Anthropic model pricing (Haiku: \$1.00 input / \$5.00 output per 1M tokens). Vendor pricing page, snapshot 2026-04.

From alpha to *scale*

NOW (2026)

Q2 TO Q3 2026

Q4 2026

Q1 2027

Q3 2027

Functional Alpha

3D graph view, GraphRAG chat, LaTeX + Markdown editors, confidence-scored linting, all live.

Columbia Pilot

Feedback loop with 5+ Columbia labs across disciplines.

Conference Outreach

US conference circuit, NeurIPS, CVPR, ACL, AMIA, plus foreign-language support for global academic rollout.

Truth Graph + Shared Graphs

Extraction pipeline on open-access academic corpora begins. Labs work on one shared graph, multi-user, multi-document.

Scale & MCP Server

Truth Graph serves frontier AI models. Enterprise deployments + data licensing.

11 · THE TEAM

The people behind the *logic engine*

Columbia SEAS '28. Combining expertise in AI agents, medical deep learning, and knowledge graphs.



GROWTH LEAD

Austin Wijaya


asw2215@columbia.edu

AI Agents & Data Engineering at the Columbia Data Agents and Processes Lab

Large-scale data lake search

Data pipelines & analytics

60k+ social followers



RESEARCH LEAD

Akito Yamauchi


ay2668@columbia.edu

Medical AI & Deep Learning across two Columbia labs (LIINC & Guo Lab). DSI Scholar.

Two first-author ISMRM abstracts

Collaborating w/ **Google DeepMind**

Medical foundation model



STRATEGY LEAD

Hayden Leung

cl4627@columbia.edu

Knowledge Graphs & large-scale data structures.
Drives ResearcherX product strategy.

Lead-author at **2023 IEEE IRI**

Automated graph frameworks

LLMs on genomic data

12 · THE ASK

\$25K to turn pilot into *platform*

Columbia Pilot

\$7,000

FROM ALPHA TO BATTLE-TESTED

Subsidize premium API compute for ~200 researchers across 5+ Columbia labs. Direct feedback loop with lab leads on STEM, medical, and legal edge cases.

Engineering for Scale

\$8,000

BRIDGE TO THE TRUTH GRAPH

Optimize Neo4j + LanceDB for millions of graph-anchored claims. Finalize MCP server so third-party LLMs plug into each user's verified knowledge base.

Benchmarking

\$2,000

PUBLISHED, MEASURED RESULTS

Test extraction algorithms, frontier models, and context-engineering strategies against SciFact-Open, DocNLI, HoVer, and RAGuard.

Market Entry & Onboarding

\$8,000

PILOT → US RESEARCH COMMUNITY

2-week Pro trials for ~3,000 power users at ~\$5 COGS/user. Present pilot results at academic and AI conferences to drive organic network effects.

\$25K TOTAL ASK

75% GROSS MARGIN

\$768M SAM



Deterministic verification *for every document an AI touches.*

A real-time linter for facts and hallucinations.

researcherx.columbia@gmail.com

austin-senna.github.io/researcherx-pitch

• Working Alpha

• 75% Margins

• \$768M SAM

• Columbia SEAS '28