MOSAIC: **M**odular **O**pen **S**cience, **AI** & **C**ollective Intelligence

Unlocking large-scale networked research to tackle complex scientific challenges

Ronen Tamari (ronen@cosmik.network) | Francisco Carvalho (francisco.l.p.carvalho@gmail.com)

Developed through the BiTS program, a part of Renaissance Philanthropy

1. What are you trying to do?

The scientific paper has been the unit of science for centuries, and it is breaking. The bundle of functions it supports — disseminating results, communicating ideas, evaluating quality, assigning credit — is coming apart, and AI is now pushing submission volumes past what the system can absorb. MOSAIC is a program to architect the next medium for science, combining AI, open social networks, and modular research tools (representing contribution units such as questions, hypotheses, critiques, evidence) into a networked research stack designed for the age of the internet and powerful AI.

If we succeed, science becomes modular and composable like Wikipedia and open-source software, enabling orders of magnitude more people to do real research as contributors whose situated knowledge shapes the agenda. Science gets its “Claude Code” moment: just as agents ship software because code lives in structured units, modular research lets agents become collaborators alongside humans. Scientific knowledge becomes a living commons that informs intelligent funding allocation; feedback loops shrink from years to days. Wicked problems (climate, social, epistemic) begin to yield to networked research that can match their complexity.

2. How is it done today, and what are the limits of current practice?

The paper remains the only officially recognized unit of scientific contribution. No other form of work (such as the modular contributions discussed above) receives formal credit. This means the vast majority of intellectual labor that actually moves science forward is invisible to the system that allocates funding, hiring, and reputation.

This also has two important implications for AI for science: **(1) Poor training data.** Papers are neither machine-readable nor faithful records of how research actually happens; they bury the reasoning behind a polished narrative, so AI trained on them inherits a distorted view of science. **(2) AI "slop science".** Papers are cheap for agents to generate, and with no smaller, verifiable unit of contribution to anchor review, the system cannot tell real progress from plausible noise.

Meanwhile, researchers already do valuable real-time thinking on social platforms, but too often this happens on closed networks where data is silored and engagement is optimized over epistemic quality. And whether on open networks or not, this work is not recognized as a scientific contribution (e.g., as modular contributions like questions, critique, discussions, etc).

Existing institutions won't solve this. Academia lacks the coordination and engineering capacity for integrated infrastructure. Startups need near-term returns, limiting their capacity for collaboration. Both startups and incumbent platforms depend on moats and data enclosure. Open alternatives that resist this pattern exist (open protocols like ATProtocol or modular research tools) but operate on volunteer labor or small budgets that can't reach critical mass on their own.

3. What is new in your approach and why do you think it will be successful?

Surprisingly, science publishing and communication has never been treated as a key R&D priority. We have seen the potential of large-scale networked research through "accidental" emergence in platforms that were never designed for the purpose (e.g., Science Twitter, see below). We also have a good theoretical understanding of the principles driving large-scale collaborative efforts, like open-source code and Wikipedia (e.g., Michael Nielsen's *Reinventing Discovery*). These have never been applied to science. MOSAIC's Big-if-True bet is that catalytic funding and intentional engineering can deliver for science what has only emerged by accident on platforms built for other purposes. And like Wikipedia or open source, an open-protocol substrate is generative: researchers, labs, and AI builders can extend it themselves, growing the system beyond any fixed feature set.

Core hypothesis

Three mutually reinforcing developments are converging to create a unique opportunity. (1) **Modular research tools** (Discourse Graphs, Nanopublications) are making it possible to represent the actual process of research (hypotheses, questions, critiques, evidence, etc) as reusable, machine-readable units rather than burying them inside monolithic papers. (2) **Open social network protocols**, particularly Bluesky/ATProtocol, now offer a credible alternative to closed platforms and are attracting large science communities (see Appendix A). These spaces are already where many scientific microcontributions happen (discussion, critique, data sharing) but the infrastructure to recognize them as actual research contributions doesn't yet exist. (3) **AI capabilities** have reached the point where it can meaningfully contribute across the full research lifecycle but lack the structured scaffolding needed to deliver — exactly what modular research tools and open protocols provide (see Appendix C).

None of these alone is enough: modular units without network adoption, social networks without recognition, AI without scaffolding. Together, they form the substrate for a fundamentally different architecture of science.

Anchor evidence

Three lines of evidence anchor this thesis (details in Appendix A). First, **networked research has already produced transformative outcomes** on platforms not designed for it, such as COVID Moonshot's preclinical antiviral candidate (a 150-researcher, “Twitter-fueled” drug-discovery collaboration), informal peer review that resulted in consequential retractions, and scores of researchers for whom Science Twitter profoundly shaped their intellectual and professional development. Second, the **ATProto Science ecosystem**, which we founded in December 2025, has seen more than 20 independent projects emerge in ~6 months — most volunteer or bootstrapped — showing that researchers and developers will build research infrastructure on shared open protocols once they exist. Third, **Discourse Graphs are actively deployed** across 10+ research labs (100+ weekly active users, 6,000+ structured contributions in a year), with adopters reporting substantial improvements to the quality and confidence of their scientific thinking. Modular research is emerging as the medium of choice to scaffold agentic AI for science contributions (see Appendix C).

Prior initiatives

Prior efforts’ impact has been limited by multiple factors (see Appendix B for full comparison): closed platforms (ResearchGate), premature financialization (blockchain DeSci), AI-for-science tools built on paper-based substrates, and modular research initiatives without reward alignment. MOSAIC's power lies in novel integration: every component is designed to deliver tangible benefit to the working researcher — better sensemaking in their own work, clearer credit for the full range of their contributions, and productive collaboration with AI and larger teams — while also generating the structured signals needed for funding allocation and credible research assessment. In particular, MOSAIC connects modular contribution data directly to a research funding platform (Catalyze), closing the loop between scientific contribution and capital allocation in a way no prior effort has achieved. Adoption follows when contributing this way is itself the reward.

4. Who cares? If you are successful, what difference will it make?

We build proof-of-viability with early adopter research communities, with broader institutional buy-in following initial success.

The early adopters phase focuses on (1) research communities who adopt the MOSAIC stack because the tools improve their day-to-day work (as described in Sec. 3), and (2) individual and institutional backers using funding platforms like Catalyze who need better signals to allocate capital to individuals and small teams. MOSAIC gives them a living, verifiable view of where promising work is happening. Together they form a self-reinforcing loop: better signals attract better-allocated funding, which rewards researchers using the tools, which draws in more adopters.

The second phase broadens adoption through research-lifecycle integrations (ORCID, journals, preprint servers, scientific authoring tools). MOSAIC integrations are backwards-compatible: they enrich existing publications with structured signals, turning the literature into a connected substrate for both researchers and AI. As adoption compounds, the gains scale: mainstream researchers join with less risk once early adopters have demonstrated the rewards; MOSAIC data provides a real-time progress map that mainstream funders and institutions rely on for allocation and evaluation; and social networks enable broad, active public participation.

5. What are the risks?

- **Adoption:** Network effects keep researchers on closed platforms. Even where MOSAIC succeeds at the individual level, new contribution and reputation signals may take longer than the program timeline to earn trust from traditional institutions and funders.
- **Sustainability:** Usage may not reach the scale needed to generate self-sustaining revenue, and the underlying economics remain unproven.
- **Programmatic:** Many interdependent developments — attribution, standards, novel publishing tools — must each work and integrate with the others. Bottlenecks in one area can cascade across the system.
- **AI displacement:** If AI capabilities advance faster than expected, infrastructure designed for human-AI collaboration may become redundant in domains where AI can operate autonomously. *Mitigation:* AI agents need a structured substrate to contribute into just as humans do, so MOSAIC's value persists across a wide range of AI futures, even if the balance of human and AI contribution shifts.

6. How much will it cost?

The 5-year program is budgeted at $30M, structured as a thesis-driven fund with two phases, Phase 2 conditional on Phase 1 success. **Matching funds** are a co-investment mechanism: MOSAIC contributes alongside community-driven crowdfunding (e.g., Catalyze), encouraging adoption of the MOSAIC stack and amplifying funding for projects the community itself signals confidence in.

Phase 1: Viability blitz (Years 1–2, $10M + $1M matching funds). Funds a core team from Discourse Graphs and ATProto Science, with subgrants to ATProto Science ecosystem projects, integrations with major preprint servers and journals (e.g., openRxiv, eLife), and community pilots with established research communities (e.g., Allen Institute). Detailed allocations in Appendix D.

Phase 2: Lifecycle integration and sustainability (Years 3–5, $15M + $4M matching funds). Conditional on Phase 1 milestones. Funds broader research-lifecycle integrations (ORCID, publishing and authoring systems), a larger crowdfunding matching pool, and paid institutional subscriptions that begin to reduce grant dependence.

7. How long will it take?

Years 1-2 build the core MOSAIC stack and seed key ATProto Science apps. Years 2-4 focus on research-lifecycle integrations to broaden adoption. Years 3-5 demonstrate sustainability: paid institutional subscriptions, a larger matching pool, and reinvested revenue together reduce the program's dependence on grants.

8. What are the mid-term and final "exams" to check for success?

Phase 1 (~2 years): Active deployments in ~20 pilot communities with sustained use. 20k monthly active researchers on the MOSAIC stack. At least $2M total flowing through Catalyze ($1M community contributions matched 1:1 by MOSAIC's pool) to researchers and infrastructure in the MOSAIC ecosystem. Initial integrations with journals and preprint servers (e.g., eLife, bioRxiv, medRxiv). Retrofitting of existing literature at scale. Demonstrated reuse of modular contributions across multiple research domains.

Phase 2 (~5 years): Over 200k monthly active researchers on the MOSAIC stack. Between $10M-20M in funding allocated via Catalyze to researchers and infrastructure in the MOSAIC ecosystem (with $4M from MOSAIC's matching pool). Annual recurring revenue of $3-5M from institutional partnerships, demonstrating a credible trajectory toward self-sustainability. Substantial deepening of integrations from Phase 1.

Appendix A: Anchor Evidence for Networked Research

Networked research outcomes on platforms not designed for it.

- **COVID Moonshot** (von Delft et. al, 2021)**.** In March 2020, a researcher posted a viral crystal structure on Twitter. Within weeks, 150+ scientists had self-organized into a global, open drug-discovery collaboration that produced a preclinical antiviral candidate. The collaboration was an unprecedented mode of scientific work, running on social-media coordination and open data sharing without an institutional sponsor. As the researchers observed: "It is a way of working that none of us realized was possible."
- **The Polymath Project.** Beginning in 2009, mathematicians Timothy Gowers, Terence Tao, and others demonstrated that difficult mathematical problems could be solved through massively collaborative online effort. Gowers and Michael Nielsen published in *Nature* on the lessons for collaborative science. In 2023, Polymath contributors translated a proof of the Polynomial Freiman-Ruzsa conjecture into Lean, where machine-verification caught a minor error humans had missed.
- **Science Twitter.** In 2022, researcher Andy Matuschak wrote that "Twitter radically improved the quality of [his] creative and intellectual conversations" — and asked, "What happens if you build a pivot Twitter along this axis?" (https://x.com/andy_matuschak/status/1518738028377755650) Many researchers have reported similar experiences - for a broader review, see *Twitter changed science — what happens now it's in turmoil?*
- **OpenPath** (Huang et. al, 2023) demonstrates the AI-for-modular-science thesis directly: a Stanford team observed that pathologists were already producing microcontributions in everything but name on medical Twitter: image-plus-diagnosis annotations. They scraped roughly 200,000 of them, trained a pathology foundation model, and beat the prior state-of-the-art.
- **Informal peer review on social media.** A longitudinal digital ethnography by Patel & Chan (2026) documented 100+ cases of "informal peer review": substantive evaluation of published research occurring on Bluesky, Twitter/X, LinkedIn, PubPeer, and personal blogs. The study covers a diverse ensemble of credentialed scholars, hobbyists, and methodologists who would often be passed over for formal peer review. Three cases the authors studied in depth show this work producing real institutional consequences: a viral BMJ Nutrition paper on apple cider vinegar was retracted after informal reviewers exposed statistical flaws; a high-profile metascience study in Nature Human Behavior was retracted following critique from informal reviewers like Joe Bak-Coleman; and informal reviewers' fraud allegations against Cassava Sciences triggered FDA reports and university investigations. The authors estimate that 0.35% of all Bluesky posts in 2025 commented on specific published research — roughly 58 times the rate observed on Twitter/X — and that scientific commenting on Bluesky concentrates in dedicated feeds (e.g., #SciComment) that route conversations to PubPeer and Retraction Watch. The substantive work of peer review is already happening on open social platforms; what's missing is the infrastructure to recognize, credit, and accumulate it.

EleutherAI. Adjacent open-source AI research shows the same dynamics at more developed scale. Founded as a Discord server in July 2020 to build open-source alternatives to GPT-3, EleutherAI's volunteer community produced The Pile (an 886 GB training dataset later adopted by Microsoft for Megatron-Turing) and the GPT-Neo/GPT-J/GPT-NeoX models (the largest open-source GPT-3-style models in the world at release). In early 2023, it incorporated as a non-profit research institute backed by Stability AI, Hugging Face, and Canva. The trajectory – open chat-platform community → significant scientific outputs → institutional formalization with mainstream funder backing – is the arc MOSAIC envisions for networked research at large. Even this case emerged largely by accident: through exceptional persistence, ad-hoc compute arrangements, and years of volunteer labor. With intentional infrastructure, this trajectory could become routine rather than rare.

Scientists migrating to Bluesky/ATProtocol. Following changes to X, researchers have been migrating to Bluesky. Researchers describe it as "a place of joy" offering more control over content and engagement (Mallapaty, Nov 2024). A population-scale study tracking 276,431 academics on Twitter/X (January 2023 to December 2024) found that "one in five academics adopted Bluesky," roughly 50,000 researchers (arXiv:2505.24801). A separate 2025 analysis found over 2.6 million Bluesky posts referencing 532,302 scholarly articles from January 2023 to July 2025, with a sharp increase in scholarly activity from November 2024 to January 2025 (Zheng et al., arXiv:2507.18840). In the months following October 2024, Bluesky reached 39% of academic mentions for Nature and Science, 48% for PNAS, and exceeded X in some scholarly communities — 65% for JASIST and 56% for Quantitative Science Studies (Arroyo-Machado et al., 2025). Zheng et al. conclude that Bluesky has emerged as a credible platform for science communication.

The ATProto Science ecosystem. Co-founded by Ronen in 2025, the ecosystem now contains more than 20 independent, interoperating research projects, most built by volunteer or bootstrapped teams. Ronen also co-founded the Cosmik R&D lab, which is building Semble, a social research curation tool on ATProto. Other projects include annotation tools, research feeds, a science crowdfunding platform (Catalyze), independent research portals (research.space), and a researcher-focused social client (Lea). See https://atproto.science/projects/ for a longer list. The pace shows that researchers and developers will build on shared open protocols once they exist.

Discourse Graphs deployments. Active deployments across 10+ research labs (see list below), with 100+ weekly active users generating 6,000+ structured discourse nodes (claims, evidence, questions, hypotheses) in approximately one year. From one pilot user: *"Even in the limited time I've used Discourse Graphs, I've found huge improvements in my thinking and doing of science."* The same user said using discourse graphs is "*more aligned with what the actual process [of doing science] is like*" than the current publication system. A key design insight from these deployments: when tools scaffold how research actually works rather than how it's presented after the fact, adoption becomes natural rather than burdensome.

Discourse Graph pilot research labs:

- Allen Institute for Cell Science
- McGill University
- NYU
- Quantum Biology Institute
- Stowers Institute
- University of Toronto
- University of Washington
- Washington University in St. Louis
- University of Maryland

Appendix B: Prior Art and Differentiation

Prior approach	Limitations
Mainstream science social networks (ResearchGate, Academia.edu)	Closed platforms that resist community-driven evolution; researchers can't extend them or build on their data
Blockchain-based decentralized science (DeSci)	Built tokens and tools before communities; financialization friction; adoption stalled
AI-for-science tools (literature search, summarization, paper-writing assistants)	Inherit the limitations of papers as primary data source; amplify the existing system rather than rebuilding it
Prior modular research initiatives (nanopublications, semantic web)	Researchers are not rewarded for modular formats, so adoption stalls regardless of technical merit. Uptake also limited by lack of broad adoption surfaces (integration with social networks and research lifecycle apps and services)

Appendix C: Technical Scaffolding for AI in Science

The case for modular, machine-readable research has a long lineage, including the FAIR (Findable, Accessible, Interoperable, Reusable) data initiative and the Nanopublications project. Progress in AI-for-science capabilities is leading to a resurgence of interest in these ideas, with recent initiatives more explicitly geared towards supporting powerful agentic AI. Examples include Open eXchange Architecture, OpenEval and the Agent-Native Research Artifact (ARA)

In particular, as Rowan Cockett argues in *Standards Matter More Than Ever in the Age of AI*, AI coding tools work because code lives in a richly structured environment: version control, dependency graphs, type systems, tests, executable artifacts, and decades of accumulated tooling. "If we had jumped directly from Windows Notepad to chat-based code generation — skipping IDEs, compilers, and tooling — the result would have been stunted." AI for science is currently being asked to reason over PDFs: the equivalent of building modern software tooling on plain-text files alone.

Appendix D: Pilot Partners and Candidate Allocations

Phase 1 community pilots (in conversation):

- Allen Institute
- Quantum Biology Institute
- McGill University
- NYU
- Weizmann Institute
- Metagov
- Additional pilots TBD based on community fit

Candidate Phase 1 grants:

ATProto Science / Modular Research

- **Lea**: researcher-focused social client for Bluesky/ATProto
- **research.space**: independent research portal providing ATProto identity without requiring a Bluesky account
- **Cosmik:** Building Semble, a social research curation tool on ATProto
- **OXA (Open eXchange Architecture) <> ATProto integration.**
- **Discourse Graphs**

AI for Modular Science

- **Entailment.ai** a framework that breaks scientific ideas into structured, evaluable parts
- Claim Trees: Framework for extracting Discourse Graphs from existing literature

Phase 1 integrations (in conversation):

- **openRxiv (bioRxiv / medRxiv)**: preprint server integrations
- **eLife**: journal integration with modular contribution support

Candidate phase 2 lifecycle integrations:

- **ORCID**: researcher identity and credit attribution