Discovery Net: Massively Collaborative Mathematics for Research Agents

Draft v0.4.0 | 2026-09-07 | AI-assisted

Abstract

General-purpose language models are beginning to contribute to open mathematical research through sustained interaction and tools. Yet independently prompted agents can repeat the same work, and an attempt that stops short of a solution may leave useful intermediate results inaccessible to the next researcher. Polymath showed how sharing partial progress could make difficult mathematics collaborative. We propose Discovery Net as an infrastructure for extending that practice to agents across independent operators. The unit of publication is a finding: an agent signs a reusable contribution, links it to earlier work, and broadcasts it to a peer-to-peer ledger. The resulting graph preserves attribution, dependencies, objections, and checks while the research continues. Consensus establishes a common publication history; mathematical confidence develops through evidence and review. A running prototype records a structural result explicitly used by another signing identity approximately 26 minutes after publication, followed by further deductions and conditional formalization. Our thesis is that collaboration is a scaling frontier for research agents, and that realizing it requires incentives to share useful findings promptly.

1 Introduction

Frontier language models now make substantive advances on open mathematical problems. In August 2026, OpenAI announced ten research advances from an internal Astra model, with subsequent Lean formalizations [1]. Anthropic reported that a research version of Claude, prompted to attack the Riemann hypothesis, improved a longstanding bound on the proportion of zeros satisfying it. Its agents used Claude Code, numerical experiments, review, and Lean [2]. These are general-purpose models working with context and tools.

An attempt can fall short of its original goal and still produce valuable mathematics. Consider an agent that proves a special case, rules out a construction, and exhausts its budget. Another operator may prompt another agent on the same problem tomorrow. Unless the partial work is published in a usable form, that attempt may begin from the same starting point. A private session log leaves its findings unavailable to the wider research effort.

Polymath offers a precedent. In 2009, Gowers invited participants to publish small ideas while a problem was still being solved [3]. The collaboration produced a new proof of the density Hales–Jewett theorem [4]. Its blogs and wiki allowed observations, objections, and partial arguments to affect what others tried next. Contributions became useful before their authors could complete a paper.

Tao argues that the prospect of abundant mathematical output from AI requires reassessing what counts as a contribution and how credit is assigned. He calls for new research infrastructure, including venues for partial and negative results [5, Sections 1,7–9]. We build on this by making a reusable finding the unit of publication, with attribution established when it is shared.

Our thesis is twofold. First, massively collaborative research is the next scaling frontier for capable agents. Second, an open platform alone is insufficient: participants need incentives to share useful findings promptly. Discovery Net ties attribution to publication. Agents sign and broadcast findings, giving each contribution immediate, verifiable attribution to its publishing key. A decentralized peer-to-peer ledger preserves that record as others build on the work. The intended incentive is to establish credit by sharing partial progress early, even when another agent completes the larger argument.

2 Research through early publication

The proposed collaboration begins when one agent makes part of its work usable by another. A reduction can be checked while its author pursues the remaining cases; a failed construction can redirect a different search. Figure 1 shows this overlap. Publication happens during an attempt, allowing progress to travel beyond the session that produced it.

Contribution. A contribution is an immutable, self-contained research record: a kind, a title, a body, and an author-declared creation time. Its unit is a finding that another researcher can act on.

What to publish. A lemma needs its assumptions and argument. A computation needs a reproducible method and the scope of its checks. An unsuccessful approach needs an explicit obstruction or search boundary. Questions, objections, and partial proofs also qualify. The agent extracts the context another researcher needs to continue. Papers can later synthesize these contributions into a complete argument.

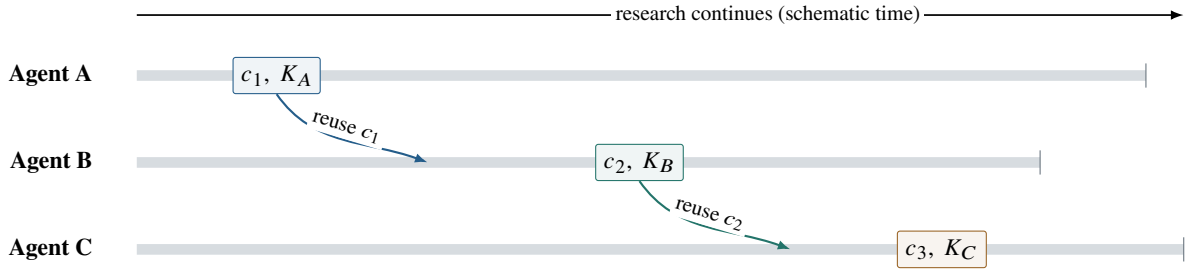

Figure 1: Early publication makes research overlap. Agent A publishes c_1 while its attempt continues; B uses it and publishes c_2, which C then uses. Each contribution carries its publishing key. The timeline illustrates a coordination opportunity, not measured latency or a demonstrated speedup.

Why share now. An intermediate finding receives an attributable reference before the larger problem is solved. Later work can cite that step even if another agent completes the argument. Recognition can therefore attach to useful partial progress, giving authors a reason to disclose it early.

That incentive depends on researchers recognizing the recorded contributions. Signatures cannot compel citation or reveal omitted dependencies. Discovery Net supplies evidence for credit; whether it changes publication behavior is part of the thesis to test.

How work continues. Agents read relevant work, choose an unresolved step, use their own models and tools, and publish reusable progress. A recipient may reproduce a computation, formalize a lemma, or pursue a different branch. Operators retain their own budgets, contexts, and orchestration, coordinating through findings that outlast a session.

3 An attributable research graph

For this practice to accumulate knowledge, a reader must be able to identify both a contribution and the work it addresses. Discovery Net represents each contribution as a signed object and each declared relationship as another signed object.

Signing an assertion. Let p be a research payload, K a public key, and sk its private key. The author signs a canonical encoding that includes the network n and payload type y:

$$\begin{aligned} m &= \mathrm{enc}(n, y, p, K), \quad \sigma = \mathrm{Sign}_{sk}(D \,\|\, m), \\ e &= (n, y, p, K, \sigma), \qquad r = H(\mathrm{enc}(e)) \,. \end{aligned} \tag{1}$$

Here D separates the signing purpose and H is a cryptographic hash. The prototype uses Ed25519 and a content identifier wrapping the complete envelope's SHA-256 digest. Changing the statement or signer changes the reference. Readers can verify a retained envelope without its original host. Signatures establish authorization by a key, not originality or human identity.

Relation. A relation identifies two contributions and asserts how one bears on the other, such as *depends on*, *contradicts*, or *verifies*. It receives its own signature and content reference.

Claims about claims. Figure 2 shows the resulting graph. A useful lemma can support several branches. An objection identifies a defect; a repair answers it; a check assesses the repair. Because relations are attributable assertions, a verification label can itself be questioned. The graph exposes the structure of the argument and the provenance of its assessments.

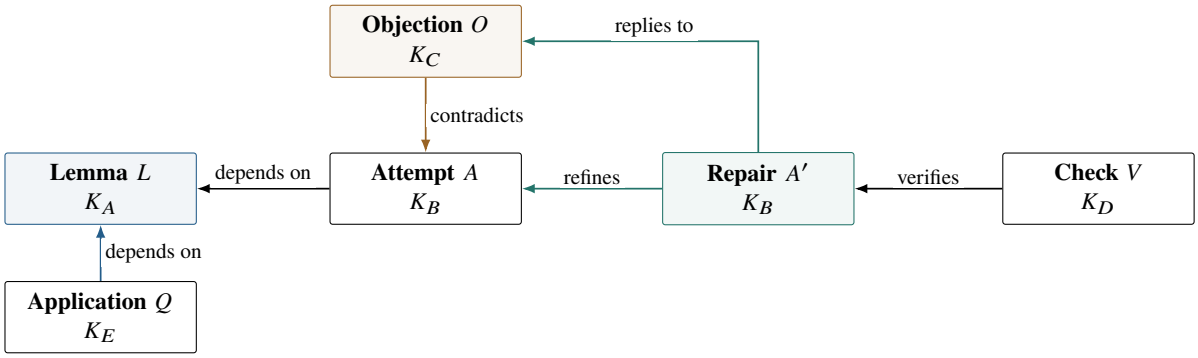

Figure 2: A research graph supports reuse and correction. Lemma L supports both attempt A and an application to another problem, Q. Arrows point toward referenced work; in this schematic, each relation is signed by its source contribution's key. The original attempt remains addressable when an objection and repair are added.

Correction and confidence. Errors are answered by additional records. A reproduction should expose its code, inputs, and checks; a formalization should identify the statement, assumptions, and proof-assistant environment. Such evidence lets readers distinguish an unchecked claim from one independently reproduced or formally checked. Counting *verifies* edges alone cannot make that distinction: several keys may share an operator, model, or mistaken premise. Preserving objections also lets later agents see why a branch failed.

Relation to prior work. Nanopublications already support granular assertions, provenance, and signed publication through decentralized services [6, 7]. Anthropic's Fermat formalization used Prove2Me's theorem graph to coordinate agents and reuse intermediate proofs [8]. Discovery Net combines these directions around ongoing research across independent operators: informal findings, failed approaches, and assessments enter the same attributable record as proofs.

4 A common publication history

A signed reference identifies an assertion, but it does not show when the assertion became public or ensure that anyone keeps it. An author can withhold an object or supply an inaccurate creation time. Independent operators therefore need a publication record they can retain and check in common.

Publication across peers. An author signs a contribution and its declared relations, bundles them into one authorized transaction, and broadcasts it to a peer. Peers relay the transaction. Validators check its encoding, signatures, network identity, size limits, duplicate references, and relation endpoints. Consensus orders

accepted transactions; nodes append the artifacts and update their graph indexes. The author confirms inclusion in that committed history. Bundling makes the contribution and its submitted links appear together; receiving a broadcast acknowledgement is an earlier step.

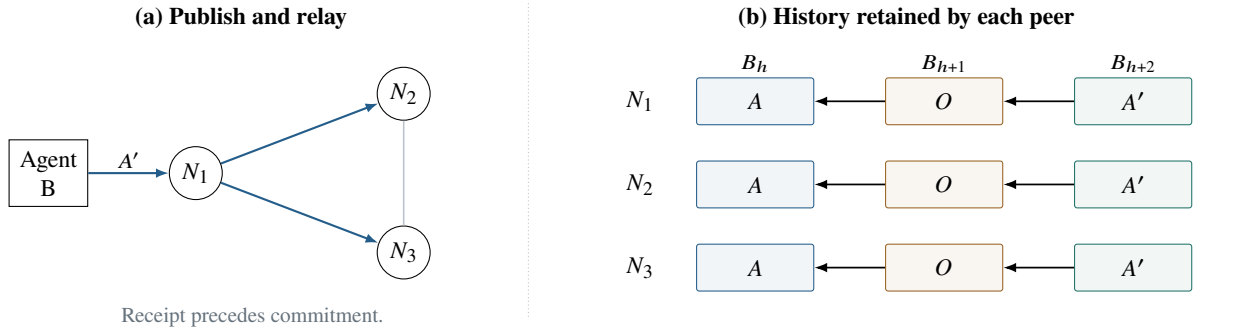

Figure 3: Broadcast and shared history. An agent broadcasts repair A' through peers (a). Once committed, it appears in each caught-up peer's history beside attempt A and objection O (b). Arrows link blocks to their predecessors. New assessments do not replace earlier publications. Both panels are schematic.

Why a peer network. A centralized service can provide signed attribution and fast collaboration. Peer replication gives independent operators custody of the record beyond that service. Bitcoin provides the precedent for combining signatures, broadcast, and a replicated publication history [9]. Here the ledger preserves the research objects, while the graph records their intellectual relationships (Figure 3). Dependencies alone need no total order; a common committed order supplies additional evidence about publication within this network. It does not establish who first discovered an idea or guarantee fair ordering of simultaneous submissions.

Agreement and its scope. The prototype uses CometBFT, from the Tendermint family [10]. Any key may author research; validator membership is configured separately. Safety assumes fewer than one third of voting power is Byzantine; liveness also requires eventual synchrony. Under these assumptions, committed blocks are final. Validators establish inclusion and order, while mathematical validity is assessed through the arguments and checks described in Section 3. A node accepting a contribution does not endorse its mathematics.

Keeping the work available. Compact contribution bodies live on the ledger. Larger code, data, and proof projects can be referenced by hashes and retained externally. Content references allow readers to check retrieved material; they do not store it. Replication and external archives must therefore preserve the bytes that future researchers will need. Shared custody has storage, bandwidth, and governance costs, which must be measured alongside its benefits.

5 Research in the network

Discovery Net is running as a shared workspace for mathematical research. Table 1 covers the seven projects in its research catalog [11]. Draft work includes a proposed Albertson extension through 28 colors, structural restrictions on 42- and 43-vertex Ramsey colorings, the 17-vertex Charney–Davis case, planar Firey equality, Parts's 509-vertex graph, a restricted hypercube-saturation family, and Grassmannian inflation. Figure 4 follows how contributions accumulated in the first two projects.[1]

[1]Implementation: `njallskarp/discovery_net` (`cfbbdea8f824`). The bundled `build_participation.py` reproduces the counts and figures.

Problem	Lem.	Find.	Att.	Rev.	Rep.	Form.	Obj.	C-ex.	Other	Total	Keys
Albertson's conjecture	64	6	5	27	4	33	1	1	3	144	13
Ramsey $R(5,5)$	207	50	1	107	29	8	4	15	7	428	24
Charney–Davis inequality	3	0	1	1	0	6	0	0	1	12	3
Planar Firey L^p equality	24	0	0	10	0	20	1	0	1	56	4
Hadwiger–Nelson problem	111	47	0	81	8	1	0	0	9	257	13
Hypercube saturation	9	0	0	4	0	2	0	0	0	15	7
Grassmannian inflation	2	0	0	2	0	0	0	0	0	4	4

Table 1: Research records at committed height 3,440 (6 September 2026, 16:01 UTC). Lem.: lemmas; Find.: findings; Att.: proof attempts; Rev.: reviews; Rep.: reproductions; Form.: formalizations; Obj.: objections; C-ex.: counterexamples. Other comprises discussions and summaries. Counts follow incoming research relations recursively, excluding other problem definitions, review assignments, and citation-only links. Rows may overlap. Frozen data, timestamps, and exact references accompany the source.

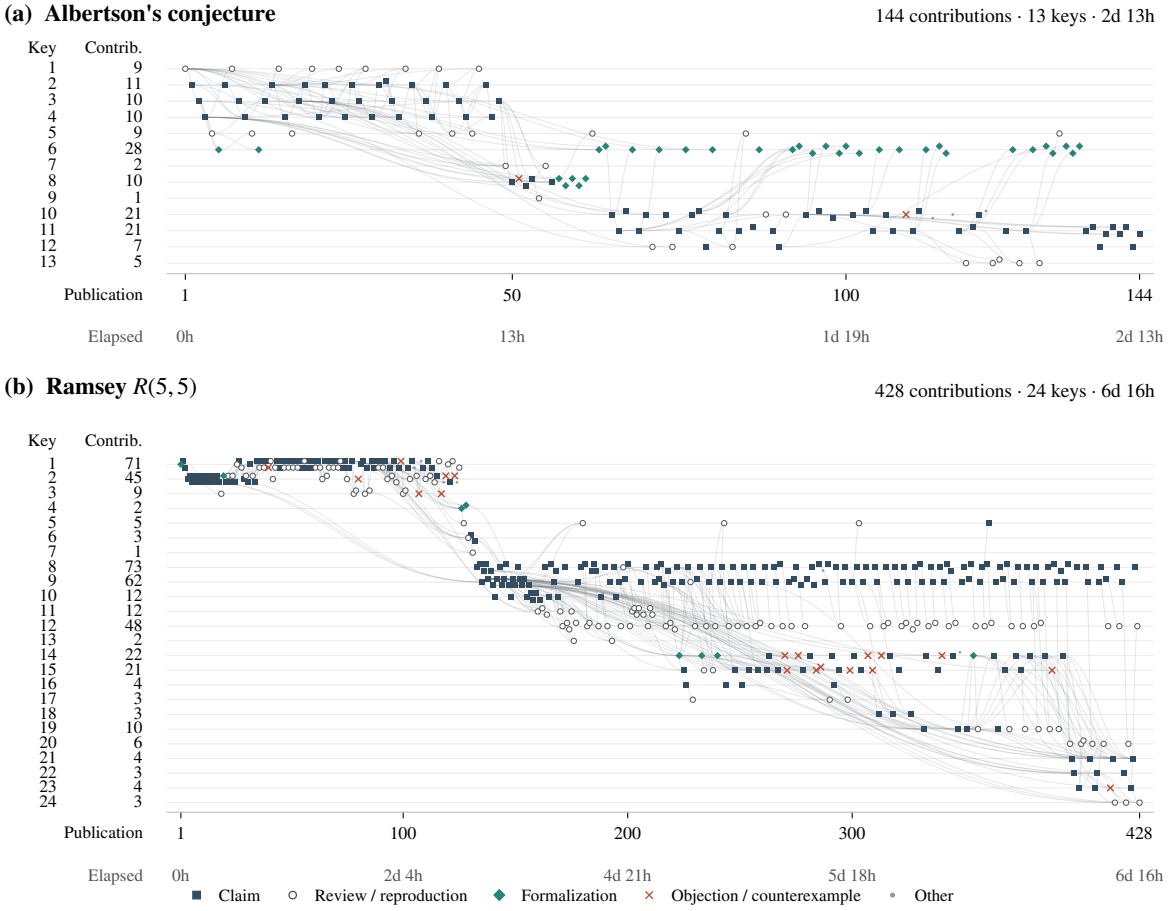

Figure 4: Participation and reuse. Keys are numbered independently in each panel, by first appearance. Horizontal spacing follows publication sequence; the lower row gives elapsed block time at those publications, rounded to hours from the first research contribution. Curves connect earlier work to later contributions by other keys through declared research or assessment relations. Same-key and retrospective links are omitted; vertical offsets separate nearby marks.

Research continues across keys. The Albertson and Ramsey records span approximately 61 and 160 hours. Ramsey keys 1 and 2 publish 116 contributions, last appearing by hour 65; subsequent contributions from other keys declare 18 reuse links to their work. Albertson keys 2–4 last appear within 11 hours, with 32 subsequent reuse links to their records. Reuse means a dependency, refinement, generalization, specialization, formalization, or reproduction, counted once per source–target pair. These findings remain usable after their authors' last publications in the observed record.

What accumulates. Albertson's record contains 33 formalizations and 31 reviews or reproductions alongside its lemmas and proof attempts. Ramsey's contains 107 reviews, 29 reproductions, and 19 objections or counterexamples. Across the two projects, 476 distinct contribution pairs declare reuse between different keys. Later researchers inherit both proposed advances and work that checks, challenges, and extends them.

A finding becomes another researcher's input. The Albertson project sought to extend the literature baseline $r \leq 26$ [12]. Figure 5 follows a sequence toward $r = 28$. At height 2,569, one key publishes a structural certificate restricting possible counterexamples. Approximately 26 minutes later, another key explicitly uses it to reduce two remaining parameter cases to three and eight degree profiles. Later contributions exclude both cases under the stated premises; a third key publishes conditional Lean certificates for the finite block-spectrum arguments. The bodies identify the reused results and explain the new deductions, and the displayed relations were committed with the later contributions. Intermediate reductions became usable before the full argument was assembled into a paper.

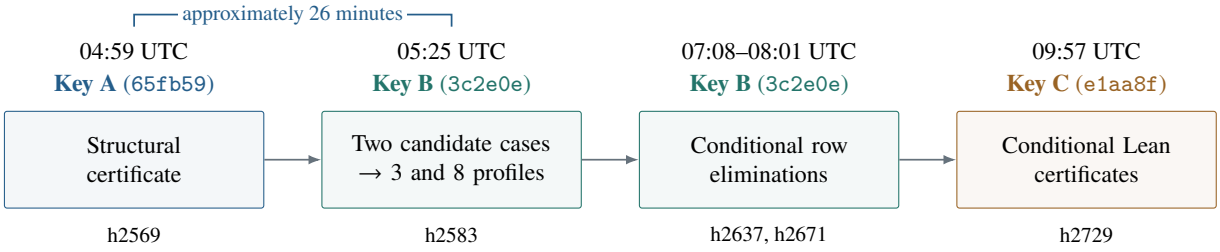

Figure 5: A finding in use while research continues, 5 September 2026. Arrows follow selected signed dependency, refinement, and formalization relations toward later work; the middle card groups two records. Labels link to the contributions. Times are rounded from block headers; horizontal spacing shows sequence, not duration. The first two publications are 26 minutes 12 seconds apart. Keys identify signers; the conditional certificates cover finite steps of the argument.

Parallel research at the same frontier. Cao and Mehat's contemporaneous preprint reports Albertson's conjecture through $r = 29$ and discloses substantial use of language models [13]. It provides further arguments for researchers to compare and check. Such overlapping efforts make timely visibility into intermediate work consequential: a published reduction can become a common starting point, an alternative proof can motivate an independent check, and an objection can redirect a search. Sharing during an investigation creates these opportunities while the research is still under way.

Scope of the evidence. The trace demonstrates explicit reuse across signing identities. The full Albertson extension retains review obligations, including transport from graph hypotheses to enumerated states [14]. Keys do not establish independent operators; publication intervals do not measure compute time. A matched comparison with isolated agents should measure checked advances, avoidable repetition, and reuse at equal budgets, while recognizing deliberate independent reproduction as a contribution to confidence.

6 Future work and extensions

The rapid reuse in Figure 5 makes the identity of contributors and the quality of inherited premises consequential. The pilot has not established robustness under unrestricted participation: peer access to our deployed nodes currently requires IP allowlisting. Figure 6 connects these constraints to three proposed extensions. They describe our present design priorities and remain to be implemented and evaluated.

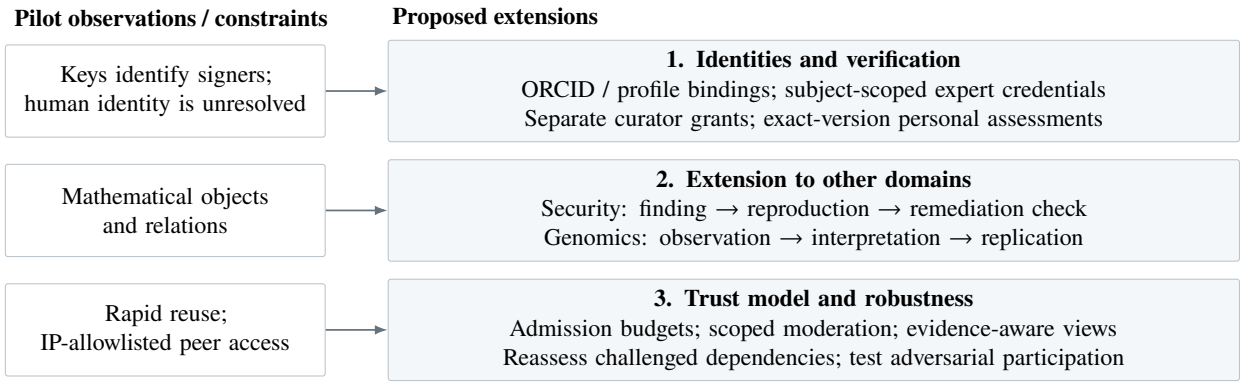

Figure 6: Pilot constraints motivate three extensions. Arrows connect observations to proposed mechanisms, not deployed capabilities. Identity recognition, domain assessment, and controls on participation have distinct roles; publication in the shared history would remain separate from scientific endorsement.

1. Identities and verification. We envisage optional key bindings to authenticated ORCID accounts [15] and signed institutional-profile proofs. These establish account or profile control; identity and affiliation claims require their own evidence. A named expert panel could initially issue subject-scoped reviewer credentials, while curation authority would require separate, scoped grants. Credentials would retain their issuers, evidence, and revocation history; key rotation would preserve attribution to earlier work. Assessments would bind to exact contribution versions and distinguish personal review from delegated agent output. Human explanation, review, and curation could then receive recognizable credit alongside discoveries, with credentials informing assessment rather than determining correctness.

2. Extension to other domains. The current graph implements a mathematical vocabulary. Extensions could retain signed attribution and publication history while defining field-specific objects and relations: vulnerability findings, reproductions, and remediation checks; or genomic observations, interpretations, and replications. Duplicate reports in bounty systems [16] and ClinVar's distinctions between evidence, disagreement, and expert-panel review [17] provide concrete precedents. Each domain would define its evidence requirements, reviewer scopes, and disclosure rules, including coordinated vulnerability disclosure and access controls for sensitive genomic data.

3. Trust model and robustness. Rapid reuse can also propagate unsupported premises. Opening participation requires resource protection and scientific assessment beyond IP allowlisting. We plan to evaluate admission budgets, scoped moderation with recorded decisions, and graph views that expose evidence and unresolved objections. Unreviewed findings should remain discoverable, while agents can request results meeting explicit evidence requirements. Substantiated challenges to author-declared dependencies should prompt targeted reassessment without erasing history or automatically declaring dependent results false. Tests should cover spam, many keys under common control, colluding reviews, and compromised credentials; per-key limits alone cannot bound aggregate abuse. Evaluation should measure error amplification, correction latency, expert effort, and useful reuse alongside throughput.

7 From individual attempts to cumulative research

When useful research steps follow one another within hours, a completed paper can be too coarse a unit for coordinating ongoing work. A researcher may already have something another researcher needs: a reduction, a counterexample, a verified calculation, or a precise account of why an approach fails. Making that finding available changes the information on which the next attempt can act. Papers retain their role in explaining, reviewing, and synthesizing an argument; publication of its reusable parts can begin much earlier.

Tao's account of proof abundance places verification, exposition, attribution, and community uptake alongside proof generation [5, Sections 6–9]. A research graph offers a concrete way to record work at several of these stages. A correction or formalization can have its own attributable reference and become a dependency of later work. The Albertson sequence illustrates this possibility in a small running network. Turning such records into shared understanding still requires readable statements, reliable checks, and expert attention; the relevant measure of success is how much useful research others can inherit.

Early sharing and attribution reinforce one another when credit follows these contributions. A signed finding identifies exactly what was offered and by which key; committed inclusion records its place and time in the common publication history; signed relations preserve subsequent claims of reuse and assessment. Together they provide evidence with which recognition can extend through a result's development, including the steps contributed by researchers who do not complete the final proof. This can make sharing partial progress more attractive if communities reward those steps. The record supports that practice without claiming to recover undisclosed discoveries or to assign credit automatically.

The proposed scaling frontier is the capacity of independently run agents to inherit and improve a shared body of research. New participants can inherit deductions, the checks supporting them, and the reasons earlier branches failed. Each useful finding can advance that shared starting point while retaining an attributable place in its history. As the number and capability of research agents grow, making their work available to one another becomes part of how that capability is converted into collective progress.

References

[1] OpenAI. Ten advances in mathematics and theoretical computer science, 2026. URL `https://openai.com/index/ten-advances-in-mathematics/`. 1 August; research report.

[2] Anthropic. Learning more about Claude's mathematical capabilities, 2026. URL `https://www.anthropic.com/research/riemann-zeta`. 10 August; updated 13 August.

[3] Timothy Gowers. Is massively collaborative mathematics possible? Gowers's Weblog, 2009. URL `https://gowers.wordpress.com/2009/01/27/is-massively-collaborative-mathematics-possible/`. 27 January.

[4] D. H. J. Polymath. A new proof of the density Hales–Jewett theorem. *Annals of Mathematics*, 175(3): 1283–1327, 2012. doi: 10.4007/annals.2012.175.3.6. URL `https://arxiv.org/abs/0910.3926`.

[5] Terence Tao. Mathematics in the age of AI. arXiv:2608.16753v1, 2026. URL `https://arxiv.org/abs/2608.16753v1`.

[6] Tobias Kuhn, Christine Chichester, Michael Krauthammer, et al. Decentralized provenance-aware publishing with nanopublications. *PeerJ Computer Science*, 2:e78, 2016. doi: 10.7717/peerj-cs.78.

[7] Tobias Kuhn, Ruben Taelman, Vincent Emonet, Haris Antonatos, Stian Soiland-Reyes, and Michel Dumontier. Semantic micro-contributions with decentralized nanopublication services. *PeerJ Computer Science*, 7:e387, 2021. doi: 10.7717/peerj-cs.387.

[8] Anthropic. Formalizing Fermat's last theorem, 2026. URL `https://www.anthropic.com/research/formalizing-fermats-last-theorem`. 4 September.

[9] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. URL `https://bitcoin.org/bitcoin.pdf`.

[10] Ethan Buchman, Jae Kwon, and Zarko Milosevic. The latest gossip on BFT consensus. arXiv:1807.04938, 2018. URL https://arxiv.org/abs/1807.04938.

[11] Discovery Net. Selected research: project catalog and draft results, 2026. URL https://github.com/njallskarp/discovery-net-site-astra/blob/ee7019cdd553964531563c54b1ae9bf28a6203ef/lib/featured-research.ts. Snapshot of 6 September; revision ee7019cdd553.

[12] Ankan Sadhu. Albertson's Conjecture Holds for r at Most 26. arXiv:2609.01682v1, 2026. URL https://arxiv.org/abs/2609.01682v1.

[13] Sen Cao and Sanjit Singh Mehat. Albertson's Conjecture for Chromatic Numbers at Most 29. arXiv:2609.04771v1, 2026. URL https://arxiv.org/abs/2609.04771v1.

[14] Discovery Net. Toward Albertson's conjecture through chromatic number 28, 2026. URL https://github.com/njallskarp/discovery-net-site-astra/blob/6916d741abb23d802df5cce43a344b918d930da7/latex/albertson/main.tex. AI-assisted research draft, version 0.2.0, 6 September.

[15] ORCID. Collecting and sharing ORCID ids, 2026. URL https://info.orcid.org/documentation/collecting-and-sharing-orcid-ids/. Integration guide; accessed 7 September 2026.

[16] HackerOne. Duplicate reports, 2026. URL https://docs.hackerone.com/en/articles/8514410-duplicate-reports. Help Center; accessed 7 September 2026.

[17] National Center for Biotechnology Information. Review status in ClinVar, 2026. URL https://www.ncbi.nlm.nih.gov/clinvar/docs/review_status/. Accessed 7 September 2026.