

Danus: Orchestrating Mathematical Reasoning Agents with Fact-Graph Memory

Jihao Liu^{1,2,*}, Guoxiong Gao^{1,*}, Zeming Sun^{2,3,*}, Bin Wu^{4,5,*}, Shurui Liu⁶,
Jiedong Jiang⁷, Haocheng Ju¹, Leheng Chen¹, Ronnie Cheng⁶, Xiping
Zhang⁸, Bin Dong^{5,9,10,11}✉

¹School of Mathematical Sciences, Peking University

²Beijing International Center for Mathematical Research, Peking University

³Research Institute for Mathematical Sciences, Kyoto University

⁴School of Mathematics, Tianjin University

⁵Zhongguancun Academy

⁶Department of Mathematics, Stanford University

⁷Westlake Institute for Advanced Study, Westlake University

⁸School of Mathematical Sciences, Key Laboratory of Intelligent Computing and Applications (Ministry of Education), Tongji University

⁹Beijing International Center for Mathematical Research and the New Cornerstone Science Laboratory, Peking University

¹⁰Center for Machine Learning Research, Peking University

¹¹Center for Intelligent Computing, Great Bay Institute for Advanced Study, Great Bay University

Recent LLM-based mathematical reasoning agents have begun to tackle research-level problems and, in several cases, have contributed to the resolution of open problems. However, scaling and orchestrating such agents effectively remains challenging, due to the difficulty of coordinating parallel proof search while keeping intermediate claims organized and reliable. In this paper, we propose Danus, an orchestration system for research-level mathematical reasoning centered on a shared fact graph as a global memory-management mechanism. Danus consists of a main agent that performs planning and coordination, multiple worker agents that carry out proof search in parallel, and a stateless verifier that checks proposed mathematical claims before they are admitted into the fact graph. Each verified fact is stored together with its proof and logical dependencies, allowing the system to build long arguments incrementally while keeping the shared proof state organized. The main agent periodically summarizes the evolving proof state, redirects workers across promising directions, and supports interaction with human mathematicians through progress reports. We evaluate Danus through six research-level case studies in algebraic geometry, singularity theory, and combinatorics, illustrating how the fact-graph memory mechanism enables Danus to construct long, detailed mathematical proofs. Our results suggest that fact-graph-based orchestration provides an effective route toward scaling mathematical reasoning agents for long-horizon research problems. Danus is open source.

*Equal Contribution, ✉Corresponding author

✉ Correspondence: dongbin@math.pku.edu.cn

🔗 Danus Source: <https://github.com/frenzymath/Danus>

1 Introduction

Large language models (LLMs) are increasingly deployed as the reasoning cores of agentic systems that retrieve knowledge, call tools, execute code, interact with external environments, and revise their outputs through feedback. In these systems, performance depends not only on the base model, but also on the harness that governs how information, tools, state, and feedback are exposed to the model. With appropriate harness engineering, LLM-based agents can outperform non-interactive prompting baselines on tasks requiring external information, iterative correction, or long-horizon execution [27, 48].

Several agents have also been proposed for research-level mathematical reasoning. The Aletheia agent [16], built on an advanced version of Gemini Deep Think, consists of a generator, verifier, and reviser, and iterates over these three components. It has resolved several Erdős problems either autonomously or semi-autonomously, and has also been applied to research-level problems in algebraic geometry [41], combinatorics [26], and representation theory [15]. The Rethlas agent [21] is designed to mimic the workflow of human mathematicians. It is equipped with skills and tools tailored to mathematical research, and maintains a working memory of intermediate artifacts generated during the reasoning process, such as constructed examples, counterexamples, and subgoal-decomposition plans. Rethlas has autonomously addressed several open problems in commutative algebra [18], functional analysis [13], and probability [28], and has also assisted in resolving mathematical research problems in algebraic geometry [40]. The QED agent [5] consists of a decomposer, prover, structural verifier, detailed verifier, and regulator, and orchestrates these purpose-specific agents in loops. It has resolved several research-level problems in algebraic geometry, PDEs, probability, and inverse problems. The ProofCouncil system [44] consists of an author agent, an advisory LLM council, a critic agent, and a compute agent, and iterates over these components. It correctly resolved 6 out of 10 problems in the second batch of FirstProof, up to at most minor revisions [1]. The AI co-mathematician [49] is an agentic AI workspace that coordinates multiple specialized agents to decompose problems, explore ideas, search the literature, run computations, draft informal proofs, and iteratively review and revise their outputs. It has helped human mathematicians resolve open problems in group theory in an interactive manner.

Most of the agents mentioned above incorporate, explicitly or implicitly, a generate–verify–revise loop, and the “multi” in existing multi-agent mathematical reasoning systems typically refers to agents with different specialized roles. By contrast, there has been little systematic study of how to scale mathematical reasoning systems built around generate–verify–revise loops by increasing the number of agents directly involved in proof generation. Scaling such agents is not simply a matter of spawning many agents to work on the same problem without further modification. It requires careful memory management and coordination. If the shared memory of multiple agents is not handled appropriately, it can confuse the agents, propagate irrelevant or erroneous intermediate artifacts, and ultimately hurt performance.

To scale and orchestrate mathematical reasoning agents effectively, we propose Danus, an orchestration system centered on a shared fact graph that serves as a global memory-management mechanism. In Danus, a main agent performs planning and coordination, while multiple worker agents carry out proof generation. The shared fact graph keeps intermediate informally verified facts organized, enabling the system to construct long, elaborated proofs.

The interaction between Danus and human mathematicians is also carefully designed. The main agent regularly summarizes the current proof state into a report, allowing human mathematicians to inspect the progress of the proof and provide high-level guidance through the main agent when necessary. The same report can also be used by the main agent to consult advanced systems such as GPT-5.5-pro for additional mathematical guidance. In addition, Danus includes a writing system that turns the completed proof into a paper-style exposition once the target statement has been proved, making the resulting argument more accessible to human readers. We demonstrate the effectiveness of Danus on several challenging research-level problems in algebraic geometry, singularity theory, and combinatorics, and discuss the role of the fact-graph memory mechanism in producing long, detailed, and correct proofs. We also explicitly report the human input provided for each problem and how the mathematicians collaborated with Danus in these examples.

The remainder of the paper is organized as follows. Section 2 presents the design of the Danus system; Section 3, six case studies; Section 4, a discussion of its strengths and limitations; and Section 5 concludes the paper.

2 Methodology

Danus is an automated system for research-level mathematics, built on the worker–verifier core of the previous system, Rethlas [21]. It couples a swarm of proof-search workers and a stateless verifier with a shared memory, all coordinated by a main agent. The design follows a strict separation of powers: the main agent performs the global planning and coordination, the workers carry out the detailed proof search, the verifier is the sole authority on correctness, and a single fact graph holds every verified result and is the system’s only source of truth (Figure 1). The remainder of this section takes up each of these in turn: Section 2.1 describes the workflow; Sections 2.2 and 2.3, the fact graph and the memory beyond it; Sections 2.4 and 2.5, the main agent, the workers, and the verifier; Section 2.6, the skills and tools of each kind of agent; and Section 2.7, summarization and paper writing.

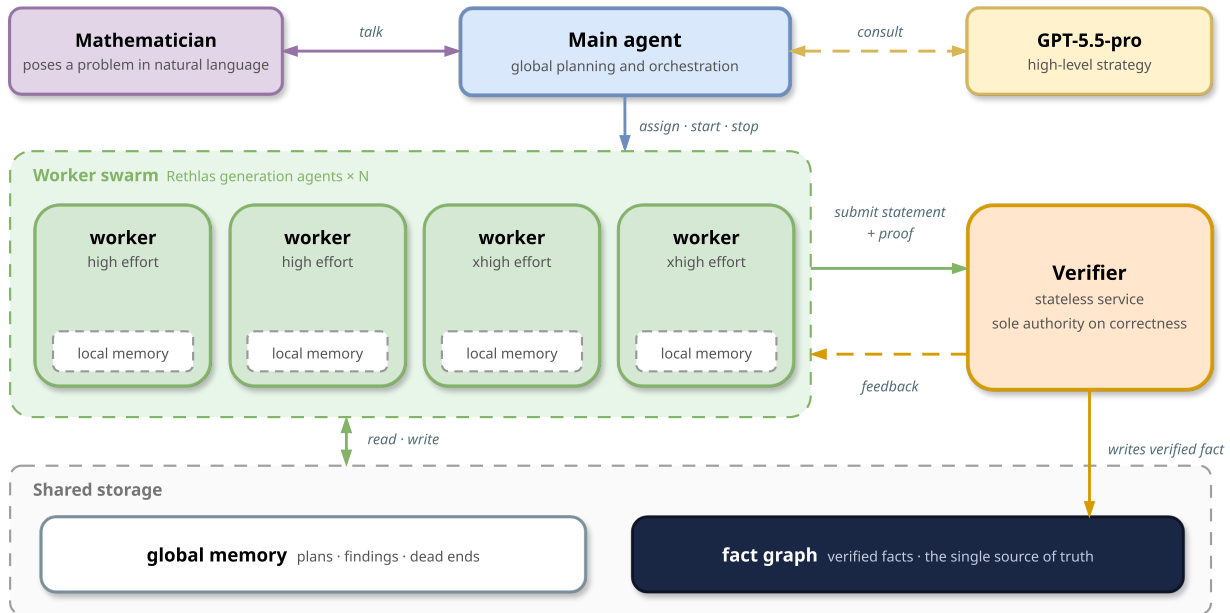


Figure 1 The overall architecture of Danus.

2.1 Workflow overview

These parts work together on a single problem, from the initial statement to a finished paper. A mathematician gives the main agent a problem in natural language. The main agent forms an initial plan, consulting GPT-5.5-pro for high-level mathematical strategy when needed, and assigns directions to the workers. The workers are Rethlas generation agents that run in parallel and explore the problem from several directions, including both constructive and refutational routes. Each worker repeatedly proposes a verifiable statement together with a supporting proof and submits it to the verifier; a statement whose proof passes verification is stored as a *fact*. These facts form a directed acyclic graph (DAG), called the *fact graph*, whose edges record logical dependencies. As the workers proceed, the main agent periodically reads their progress and the fact graph, summarizes the current state, consults GPT-5.5-pro, and re-assigns the workers. The iteration is not stopped after a preset number of rounds; it stops only when the main agent confirms that the target statement, or its refutation, has appeared as a verified fact in the fact graph. The main agent can then turn the fact graph into a paper for human experts to check.

This parallel exploration is the main departure from Rethlas. Where Rethlas pursues one line of reasoning at a time, Danus runs several workers at once, each exploring a different facet of the problem, such as a different lemma to prove, a counterexample to construct, or a toy example to study. This broadens the exploration and makes long, multi-step arguments tractable. It also raises the problem that the fact graph is designed to solve: letting many workers contribute to a single proof without interfering with one another.

2.2 The fact graph




Figure 2 The fact graph behind the matroid tangent-class case study of Section 3.6: 3,157 verified facts, 8,616 dependency edges; nodes darken and grow with dependency depth (up to 54). The clusters are separate lines of attack: bottom, conditional scaffolding that the final proof never cites; left, an independent re-derivation of the Chern-number bound; upper right, the integral lift of the result, whose final route enters the proof.

accumulated proof, whereas the fact graph lets each worker draw on only the facts it needs for its current claim and submit one fact at a time, so the working context stays small even as the proof grows to many pages. Second, it supports parallel work: a single file is hard for many workers to edit at once or to use for separate lines of attack, while the fact graph lets their contributions accumulate into one shared structure.

Revocation. The fact graph also supports revocation: a fact that is later found to be wrong is removed together with every fact that depends on it, directly or transitively. This happens in one of two situations: an error in a cited reference (usually found by the agent during a final review, or pointed out by a human expert) or, more rarely, a conceptual confusion or a flawed proof. In our runs, revocation was needed only very rarely, a point we return to when we discuss the reliability of verification.

2.3 Memory beyond the fact graph

Content that has been verified lives in the fact graph; everything else worth keeping lives in memory. Memory is not part of the truth; it is shared context that helps the workers avoid repeating one another’s failed attempts, and it has two tiers. Each worker keeps a private local memory, a running log of its own activity, used mainly for later analysis of how a line of reasoning developed. Above it is a global memory that every worker and the main agent can read and write. The global memory records the intermediate products of the search, such as plans, promising directions, dead ends, and constructed examples and counterexamples. It also records the prompt and reply of every GPT-5.5-pro consultation. Together, the two tiers keep each worker from repeating its own local exploration and let the main agent see what has happened along each worker’s branch of the search, so that it retains an overall grasp of the reasoning strategy and the plan.

2.4 The main agent

Role. The main agent, the workers, and the verifier are the three kinds of agent that act on the fact graph and the memory; of these, the main agent is the principal addition relative to Rethlas. It is responsible for the global planning and orchestration of the search: it forms and continually revises the overall proof strategy, decides how to decompose the problem and allocate effort across the workers, judges when a line of attack

Structure. The fact graph is the unique source of truth of the whole system and its central design element. It is a directed acyclic graph (DAG) whose nodes are facts and whose edges record logical dependencies. A fact is a mathematical statement together with a verifier-checked proof; an edge from one fact to another means that the proof of the second uses the first. A worker can draw on the facts already in the graph: in principle it has access to the entire graph, and in practice it retrieves the relevant facts by searching it. When a worker submits a new statement and proof, it records the identifiers of the facts that the proof depends on, and these become the new fact’s incoming edges. Every statement that passes verification is added to the graph, which therefore grows until it contains the target theorem (Figure 2).

Why a graph rather than a single blueprint. This design is what lets many workers collaborate. In Rethlas, the whole result was carried by a single Markdown blueprint that held all supporting lemmas, definitions, and the final theorem; a worker repeatedly edited this blueprint and asked the verifier to check it and suggest revisions. Compared with a single blueprint, the fact graph has two advantages. First, it keeps each worker’s context small and focused. Such context management matters because a language-model agent reasons most reliably over a short context holding only relevant material; unrelated content both strains its capacity and interferes with its reasoning. A single blueprint forces every worker to carry the whole

should be abandoned in favor of a new one, and conducts the strategic interaction with the human expert; it also performs delegated tasks such as writing up the results. Its work is confined to this global level: the main agent does not itself carry out the detailed mathematical derivations, which are left to the workers and constitute the only content subject to verification. The system thereby maintains a separation of powers, preventing the agent that steers the search from introducing unverified mathematics into the fact graph. The main agent forms these judgments from its own understanding of the overall structure, its observation of each worker’s progress, and the contents of the global fact graph and memory. For fine points of mathematics it may consult GPT-5.5-pro at low frequency, at most once per hour, as an expert reference. Its strategic decisions therefore rest on a system-wide view of the run, augmented where needed by the mathematical strength of a frontier model.

The strategy loop. In operation, the main agent first interprets the task set by the human expert, consults GPT-5.5-pro for high-level strategic guidance, and assigns initial directions to the workers. At regular intervals (every one to two hours), it re-reads the workers’ logs, the memory, and the fact graph to understand the current state and the points at which the search is stuck, produces a summary of this state, and asks GPT-5.5-pro for the next strategic step. This cycle repeats, and when the main agent confirms that the problem has been solved it notifies the human expert and supports downstream tasks such as paper generation. Throughout, the human can interact with the main agent at any time, to ask about progress, redirect the search, or adjust priorities, without interrupting the running workers. Section 2.7 describes how the progress and the final result are presented to the human.

Reading the global state. A hard part of this role is the periodic summary. The main agent must read through many layered log files and a fact graph that can hold thousands of entries, follow the progress of several workers at once, and combine all of it into a single summary that a human or GPT-5.5-pro can act on. This reading over a large, evolving body of files calls for the capabilities of a coding agent, which is built to navigate and read extensive bodies of code. We evaluated several such agents (Codex, Claude Code, and OpenClaw, loaded with GPT-5.5 or Claude Opus 4.8) and selected Claude Code with Claude Opus 4.8, which handled this reading task best; the main agent is built on it.

2.5 Workers and the verification service

Workers. While the main agent is new, the worker and verifier are inherited, with only minor changes to the agents themselves, from the generation and verification agents of Rethlas; what changes is chiefly how the worker is used. A worker now works under a task assigned by the main agent and typically focuses on one claim at a time, such as a lemma, a counterexample, or a toy example, rather than an entire proof. It repeatedly submits this claim to the verifier and revises it under the verifier’s feedback until it passes, at which point the claim enters the fact graph as a fact. Because the worker retains Rethlas’s core retrieval mechanism, built on the theorem search engine Matlas [20], it can search the mathematical literature precisely for relevant results. In practice we run three to nine workers per project and split them roughly evenly between two reasoning-effort levels, “high” and “xhigh” (the effort setting of the underlying Codex agent). Reserving half of the workers for the lower “high” effort is deliberate: it supplements the search with shallower but useful conclusions and increases the diversity of reasoning across the swarm.

The verifier. The verifier is run as a service that both the workers and the main agent can call, and it is stateless: a fresh instance judges each submission and retains nothing afterwards. Crucially, it is permitted to read the fact graph, so that when a proof depends on other facts it can follow those citations and check the dependency. Owing to the skills and checking procedure carefully engineered in Rethlas, the verifier produced essentially no false positives on the problems we tested; in the rare cases where it did, it accepted a proof containing a small number of skipped steps, or, because it takes cited references to be correct, a proof that relied on an erroneous reference. Both kinds of slip are easy to catch in a final review, which is what makes the fact graph reliable as the system’s source of truth. The full submit–verify–repair cycle is shown in Figure 3.

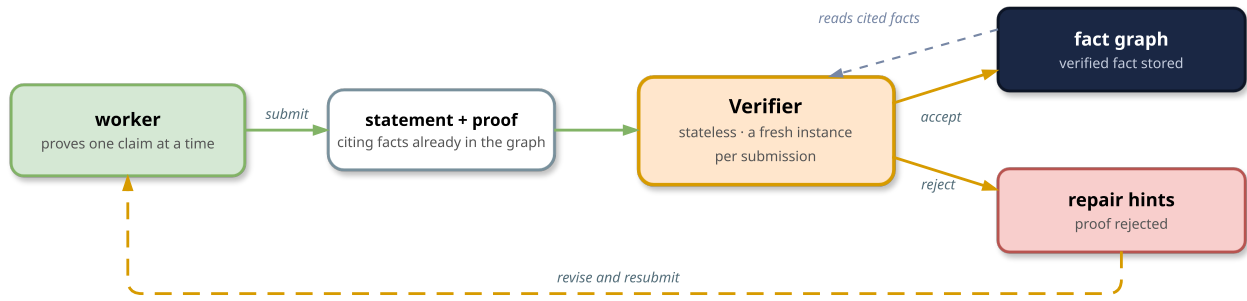


Figure 3 The submit–verify–repair cycle between a worker and the verifier.

2.6 Skills, tools, and interfaces

Each kind of agent carries out its role through its own skills and its own tools (Figure 4). The worker and verifier skills are inherited essentially unchanged from Rethlas: the worker’s are reasoning primitives such as constructing examples and counterexamples, decomposing a goal into subgoals, and direct proving, while the verifier’s are proof-checking primitives such as checking a proof step by step and confirming that cited statements exist and are applicable. The main agent has its own skills for consulting GPT-5.5-pro, producing summaries, and writing. A small command-line interface lets the main agent start, assign work to, monitor, and stop the workers, and a set of tools exposed through the Model Context Protocol (MCP) lets the agents interact with the fact graph and the global memory. Retrieval from the mathematical literature is provided as one such tool, backed by Matlas, and is available to the workers, the main agent, and the verifier alike.

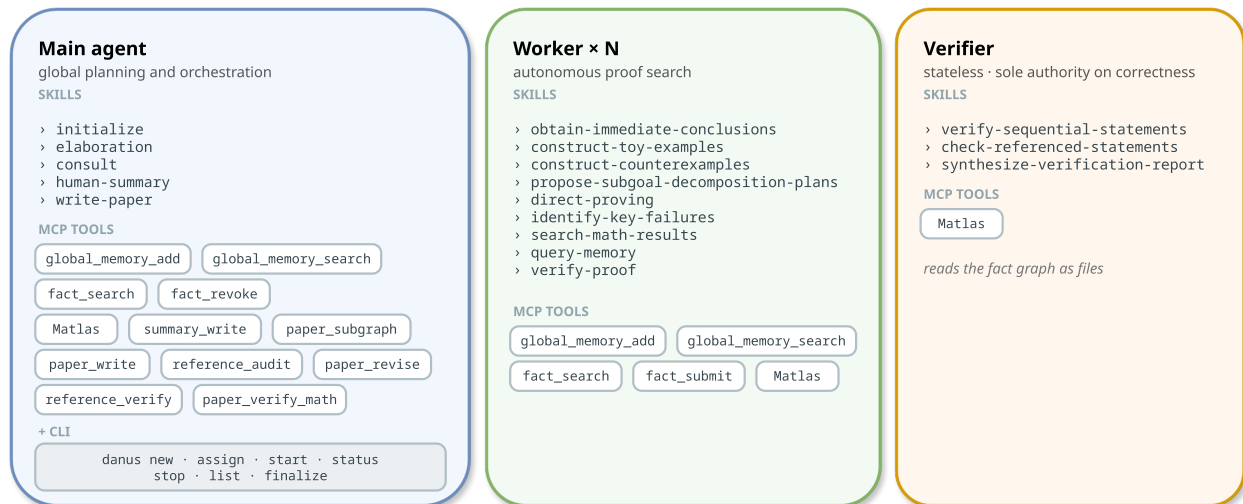


Figure 4 The three kinds of agent, each with its own skills and its own role-gated set of tools.

2.7 Summarization and paper writing

Two further components produce the system’s output for the human experts: summarization while a run is in progress, and paper writing once the proof is complete. The two differ in purpose. Summarization presents the unfinished search through a progress report, so that the experts can follow and steer it; paper writing turns the verified fact graph into a readable manuscript rather than a list of facts, so that they can check the result quickly.

Summarization. The progress report is written by an isolated agent that receives only the problem statement and the verified mathematics, stripped of every reference to the system’s internals, so the report cannot mention what its author never saw. The skill that governs it is a set of prohibitions rather than a guide

to fluent prose: progress is never estimated numerically, and a result is reported as proven only when a verified fact settles it with no hypothesis left to match; the default is to report it as weaker. The prohibitions exist because a model left to itself summarizes optimistically—a conditional argument with one unmatched hypothesis becomes “essentially complete”—and the report is the experts’ window into the run. The same discipline, with one further rule, governs the state summaries that the main agent writes for GPT-5.5-pro (Section 2.4): there, every stalled line must be classified as a failure of the method, which leaves the statement alive, or as evidence against the statement itself, the distinction behind the main agent’s decision to abandon a route.

Paper writing. Danus treats the manuscript as a new mathematical artifact: the completed draft is passed back through the verifier as a whole and revised until it passes as written. Fact-level verification does not transfer, because turning the graph into linear prose creates new mathematics that was never a fact—compressed steps, “it suffices to” reductions, glue between facts whose statements nearly but not exactly meet—so correct facts can be stitched into an incorrect manuscript, with the errors at the seams. When a manuscript is too long for a single pass, the main agent decomposes it into self-contained parts, each culminating in a designated result and carrying the results it relies on as established statements, so that the verifier judges a complete document, never a fragment. Where a verified proof was compressed into an assertion, the original is re-supplied and rendered in full. The case of Section 3.6 records this loop in a full run.

3 Results

To test Danus, we collaborated with a number of mathematicians and conducted a series of experiments in the mode of Danus–human collaboration. Here we summarize a selection of representative papers to illustrate the capabilities of Danus. As a baseline, each of the problems below was also posed, independently of Danus, to the web interface of GPT-5.5-pro; in no case did it produce a meaningful result.

To be precise about the mode of collaboration: these papers were written through interactions in which human experts directed Danus in natural language. In each case, Danus’s work ran through the two stages described in Section 2: the proof search, which ends when the target theorem stands as a verified fact in the fact graph, and the writing stage, in which the system turns the fact graph into a manuscript (Section 2.7); human experts could intervene at either stage. Part of the human input was common to every case: the problem statement itself; operational instructions carrying no mathematical content, such as requests for status summaries and the instruction to produce the manuscript; a final check of the completed manuscript, in which human experts confirmed each proof; and revisions that the human authors made independently of Danus after each manuscript was complete, chiefly fine adjustments to the mathematical notation and the citation format (see Section 4.5). We call this the common input; the “Human input” part of each subsection records only the input that went beyond it.

3.1 Optimal bend-and-break for foliations

The problem. Mori’s bend-and-break is one of the founding tools of birational geometry: on a smooth projective variety X of dimension n with an ample divisor H , through every point of a curve C with $K_X \cdot C < 0$ it produces a rational curve of H -degree at most $(n+1) \frac{H \cdot C}{-K_X \cdot C}$, and the constant $n+1$ is optimal [19, 37, 38]. For a foliation $\mathcal{F} \subset T_X$ of rank r (a structure that decomposes the variety, away from a singular locus, into a family of r -dimensional leaves), the natural analogue asks for the optimal constant for rational curves *tangent* to \mathcal{F} , that is, curves running along the leaves. Explicit but non-optimal constants were known— $2n$ by Shepherd-Barron [45] and $2r$ by Bogomolov–McQuillan [8]—and the optimal constant remained unknown. The paper [35] settles it: the optimal constant is $r+1$.

Theorem 1 (Optimal bend-and-break for foliations [35]). *Let X be a normal projective variety of dimension n , let \mathcal{F} be a foliation on X of rank r , and let H_1, \dots, H_{n-1}, H be ample divisors on X . Let C be a general complete intersection of elements of $|m_i H_i|$ with $m_i \gg 0$, and suppose $K_{\mathcal{F}} \cdot C < 0$. Then through a general point of C there is a rational curve Σ tangent to \mathcal{F} with*

$$H \cdot \Sigma \leq (r + 1) \frac{H \cdot C}{-K_{\mathcal{F}} \cdot C},$$

and the constant $r + 1$ is optimal.

The optimal constant depends only on the rank of the foliation, not on the ambient dimension; it matches the classical bound when $r = n$ and the $r = 1$ case, i.e. the case of rank 1 foliations, and once a cone theorem for foliations is available it yields the optimal length bound $r + 1$ for $K_{\mathcal{F}}$ -negative extremal rays.

How Danus produced it. A human expert gave Danus the problem and directed it to read two references, the optimal bend-and-break bounds of Jovinelly–Lehmann–Riedl [19] and the work of Bogomolov–McQuillan [8]. Danus judged on its own that the notation of [8] was not compatible with that of [19], and replaced it with the work of Kebekus–Solá Conde–Toma [22]. Danus then completed the proof; five workers explored the problem in parallel, and the search recorded 63 verified facts and 239 failed paths on the way. Danus turned the verified fact graph into a manuscript, and on examining it, a human expert suggested that Danus use the algebraicity criterion of Campana–Păun [9] to simplify the proof; Danus carried out the revision and finalized the manuscript.

Human input. Beyond the common input, human input consisted of the two initial references and one suggestion arising from the final check of the manuscript: human experts confirmed the correctness of every detail of the proof but found one section more complicated than necessary, and a human expert suggested the simplification noted above; Danus carried out the revision.

What this case shows. In this case the strategy was supplied by the human experts, and Danus’s contribution was the execution. First, it carried out a sustained synthesis, transporting an intricate technique from ordinary varieties into the foliated setting on its own, together with the connecting arguments this required. Second, when the references it was given proved unsatisfactory, it sought out a better one on its own.

3.2 Shokurov’s global index conjecture for threefold foliations

The problem. Shokurov’s global index conjecture predicts that numerically trivial log Calabi–Yau structures are torsion, with index bounded uniformly in the dimension and the coefficient set (see, e.g., [47]). In plainer terms, the canonical class of such a structure meets every curve in degree zero, and the conjecture predicts that a single positive integer multiple of it is linearly zero, with the integer depending only on the dimension and the coefficient set. Its foliated version, formulated in [33] for foliated log Calabi–Yau triples, is known in dimension two [42] and open in dimension three. The paper [30] resolves the three-dimensional case.

Theorem 2 (Foliated index theorem in dimension three [30]). *Let $\Gamma \subset [0, 1] \cap \mathbb{Q}$ be a set satisfying the descending chain condition. Then there is a positive integer $I = I(\Gamma)$ such that for every log canonical foliated triple (X, \mathcal{F}, B) with $\dim X \leq 3$, coefficients of B in Γ , and $K_{\mathcal{F}} + B \equiv 0$, one has $I(K_{\mathcal{F}} + B) \sim 0$. Moreover, if \mathcal{F} is canonical, not algebraically integrable, and $B = 0$, one may take $I \leq 30$.*

This completely resolves the threefold foliated version of Shokurov’s global index conjecture. Note that Shokurov’s global index conjecture for usual varieties remains open in dimension ≥ 4 , so this is the best foliated result one expects to get now in this direction.

How Danus produced it. A human expert gave Danus the problem. Unprompted, Danus decomposed the problem according to the rank and the algebraic rank of the foliation (the dimension of its leaves, and the dimension of its largest subfoliation with algebraic leaves) into five classes, a classification the human experts considered sound. For three of the five classes, Danus found a Lie-theoretic method and solved them;

the approach was one the human experts had not anticipated. It then attempted a similar method on the remaining two classes, the algebraically integrable ones (those whose leaves are algebraic), without success. In the human experts’ view, these two classes are precisely the ones for which humans already have a clear and simple approach, while the three classes Danus solved are the harder ones. A human expert then suggested handling the remaining two classes with the foliated minimal model program (the foliated version of birational geometry’s standard simplification machinery, recently developed by the human expert and others), and with this hint Danus solved them. The human experts then asked Danus to prove a stronger result—generalizing the index theorem from canonical foliations with $B = 0$ to log canonical foliated triples (X, \mathcal{F}, B) with DCC coefficients—and gave hints toward the ACC for minimal log discrepancies on surfaces [3] and certain results of Spicer [46]. Danus completed the generalization. Across the run, seven workers explored in parallel, the first wave running for about eight hours, and the fact graph grew to 784 verified facts, of which 77 form the supporting closure of the theorem. Danus then turned the verified fact graph into the manuscript.

Human input. Beyond the common input, human input consisted of two interventions, both during the proof search: when Danus stalled on the algebraically integrable classes, a human expert suggested the foliated minimal model program, after which Danus resolved them; and the human experts asked for the stronger result, supplying the hints noted above, upon which Danus completed it.

What this case shows. This time the structure of the attack was Danus’s own. First, it decomposed the problem into an appropriate classification without any human hint. Second, it discovered a method of solution from a direction the human experts had not expected. Third, when it did stall, a brief hint from a human expert was enough: it understood the suggestion and solved the remaining cases.

3.3 Total Cartier indices of rational singularities in families

The problem. The total Cartier index of a normal projective variety is the least common multiple of the Cartier indices of all \mathbb{Q} -Cartier Weil divisors on it—here a divisor is \mathbb{Q} -Cartier if some positive multiple of it is cut out by a single equation near every point, and its Cartier index is the least such multiple. The behavior of this index in families is a basic question about how singularities vary, of the kind the moduli theory of higher-dimensional varieties (in the wake of the KSBA program [4, 24]) constantly needs. Han and Jiang proved boundedness in bounded families of klt-type varieties (a standard class of mild singularities) by running the minimal model program in families, and asked (Problem 4.5 of [17]) whether the much weaker assumption of rational singularities suffices. Rational singularities form a natural frontier here: the class is not stable under the operations of the minimal model program—the tool behind Han and Jiang’s result and behind most boundedness theorems of this kind—and without some hypothesis the index can be infinite (a cone over an elliptic curve already fails). The paper [34] answers the question affirmatively.

Theorem 3 (Total Cartier indices in families [34]). *Let k be an algebraically closed field of characteristic zero and $\pi: X \rightarrow B$ a projective morphism to a finite-type k -scheme. Then there exists a positive integer I such that for every closed point $b \in B$ for which X_b is normal, projective, of pure dimension, and has rational singularities, the total Cartier index of X_b divides I .*

How Danus produced it. A human expert gave Danus the problem, posed at first as a hunt for a counterexample. Seven workers explored refutational and constructive routes in parallel, and the search converged on the affirmative answer: Danus independently discovered a route that synthesizes techniques from markedly different fields: it reduced the algebro-geometric problem to a problem in commutative algebra, and reduced the core of that problem, in turn, to a problem in real algebraic geometry, which it treated with a Hadamard-inequality bound (a classical determinant inequality) among other tools. Along this route Danus also classified the problem, on its own, into the surface (codimension-two) case and the case of dimension at least three (codimension at least three), attacked the two by different methods—controlling the link of the singularity (the boundary of a small neighborhood of the singular point) for surfaces, and using Kollár’s work on local Picard groups [23] in higher dimension—and stitched the two together at the end. Danus then reported that it had completed every case except the codimension three case, which it could not settle. A human expert examined the run and found the cause: unable to obtain the T_EX source of [23], Danus had worked from the PDF and misread the hypothesis “dimension ≥ 3 ” as “dimension > 3 .” Once this was pointed

out, Danus repaired the argument and completed the proof search, and the system turned the fact graph into a manuscript. The first draft was poor: in compressing the verified material into prose, the writing rendered a key lemma in real algebraic geometry incorrectly, while the verified facts beneath it were sound. Submitted back to the verifier (Section 2.7), the draft was rejected; Danus revised it against the verifier’s findings and brought the manuscript to completion.

Human input. Beyond the common input, human input consisted of a single correction: a human expert traced Danus’s failure in the three-dimensional case to the misread hypothesis in [23], and the repair itself was carried out by Danus.

What this case shows. The origination widened here, in both field and scale. First, Danus found its tools far from the problem’s home topic—the minimal model program—and, in the experts’ view, changed their sense of the methodology available for such questions. Second, it decomposed the problem on its own and combined the cases at the end. Third, it carried out a long chain of reasoning whose steps are macro-scale reductions, first to commutative algebra and then to real algebraic geometry, rather than lemma-sized ones. Fourth, because the exploration runs refutational and constructive routes together (Section 2.1), a search launched toward a counterexample could end by settling the question affirmatively. Fifth, when the writing stage compressed verified material incorrectly, the flaw was caught by verification and repaired without human intervention (Section 2.7).

3.4 Factorial asymptotics of the Matryoshka numbers

The problem. The cosmodhedron is a positive geometry—informally, a geometric object generalizing a convex polytope, whose shape encodes quantities of physical interest—introduced by Arkani-Hamed, Figueiredo and Vazão to encode the wavefunction of the universe in a class of cosmological models [7]. Its combinatorics [6] involves the *Matryoshka numbers* (OEIS A177384 [39]), defined by $a_1 = 1$ and $a_n = \sum_{k=1}^{n-1} (k+1)a_k a_{n-k}$. Kotěšovec conjectured, in the OEIS entry, that $a_n \sim c \cdot n! n^4$ with $c \approx 0.0054283$; the conjecture is restated as Conjecture 5.7 of [6], whose authors write that they do not know how to prove it. The conjecture asserts that the numbers grow factorially, and pins the growth law down to the exact multiplicative constant. The note [29] proves it.

Theorem 4 (Matryoshka asymptotics [29]). *There is a real number S with $0 < S < \infty$ such that*

$$\lim_{n \rightarrow \infty} \frac{a_n}{(n+4)!} = \lim_{n \rightarrow \infty} \frac{a_n}{n! n^4} = S, \quad 0.00542831750 \leq S \leq 0.00542831848,$$

an enclosure of width less than 10^{-9} containing the conjectured value.

The proof is elementary and every constant in it is explicit, so the enclosure—an interval proved to contain S —is rigorous rather than numerical; since S has no known closed form, such an enclosure is the strongest meaningful determination of the constant.

How Danus produced it. A human expert gave Danus the problem. Danus produced the complete proof entirely on its own: it proved that the limit exists and derived explicit upper and lower bounds for it, giving the certified enclosure of S stated above. The run went essentially straight to the proof: five workers, about ninety minutes from launch to the verified target theorem, a fact graph of 100 verified facts with dependency chains up to 15 deep, of which 37 form the supporting closure of the theorem, and a global memory recording a single proof attempt and no dead ends. Danus then turned the verified fact graph into the manuscript, likewise without human involvement.

Human input. Human input did not go beyond the common input: no mathematical guidance of any kind was provided.

What this case shows. Here there was no human mathematical contribution: Danus completed the entire pipeline, from the problem statement to the finished manuscript, without any hint. We note also that the problem comes from the combinatorics of mathematical physics: Danus’s range is not confined to a single area of mathematics.

3.5 Weighted homogeneity via logarithmic vector fields

The problem. A classical theorem of Saito [43] characterizes weighted homogeneous isolated hypersurface singularities algebraically ($f \in J(f)$, the Jacobian ideal); a hypersurface singularity is weighted homogeneous if, in suitable coordinates, its defining polynomial becomes homogeneous once the coordinates are given appropriate positive weights. da Silva Machado and Seade [36] conjectured two characterizations of this hidden symmetry that are geometric rather than algebraic, phrased in terms of vector fields and their behavior near the singular point. A pair of companion papers resolves the conjecture affirmatively: one written entirely by hand [32], the other produced by Danus [31].

Theorem 5 (Weighted homogeneity via logarithmic vector fields [31, 32]). *Let $(D, 0) \subset (\mathbb{C}^{n+1}, 0)$ be a reduced isolated hypersurface germ, with $n \geq 2$, or with $n = 1$ and D irreducible. Then the following are equivalent: (i) $(D, 0)$ is weighted homogeneous in suitable coordinates; (ii) in suitable coordinates, there is a holomorphic logarithmic vector field everywhere transverse, in the real-Euclidean sense, to all small links of D ; (iii) there is an ambient holomorphic vector field tangent to D with a nondegenerate isolated singularity at the origin.*

Both new characterizations are geometric, where Saito’s criterion is algebraic.

How Danus produced it. A human expert gave Danus the problem, without suggesting any route. Danus completed a proof and turned the verified fact graph into a manuscript. In parallel, a human expert produced an independent proof of the same conjecture; the two proofs proceed by genuinely different methods, and the expert remarked that Danus’s approach was one they had not anticipated. On examining Danus’s manuscript, the human experts found a gap: the term “nilpotent” is used inconsistently across the literature, the definition in a reference Danus had relied on is itself erroneous, and the flaw had propagated into one step of the argument. After a human expert pointed this out, Danus discarded the affected steps, located a correct reference in place of the flawed one, revised the proof substantially, and completed it correctly. The search ran far wider than the proof it finally kept: seven workers, over two waves, produced 687 verified facts—with 23 more revoked in the repair just described—in dependency chains up to 21 deep, while the supporting closure of the main theorem holds about ten facts; the global memory records 91 proof attempts and 49 counterexamples.

Human input. Beyond the common input, human input consisted of one finding in the final check of the manuscript: the human experts identified the flawed reference noted above and pointed it out to Danus. The subsequent overhaul of the proof was carried out by Danus itself.

What this case shows. This case tested Danus against a defect in the literature itself. First, it produced a proof independently, and by a route the human expert had not anticipated. Second, it repaired its own literature dependency: the facts resting on the flawed reference were revoked through the mechanism of Section 2.2, a correct reference was found, and the proof was rebuilt from the verified remainder.

3.6 Tangent classes of matroids and wonderful compactifications

The problem. When a matroid M of rank $d + 1$ is realized by a complex linear subspace L , the wonderful models of De Concini–Procesi [12] provide smooth projective compactifications $W_{L, \mathcal{G}}$ of dimension d , one for each building set \mathcal{G} of flats of M , and Feichtner–Yuzvinsky [14] presented the Chow ring of $W_{L, \mathcal{G}}$ by generators x_F , indexed by the flats $F \in \mathcal{G}$, and relations that involve only the matroid M . The presentation is intrinsic to the matroid M and therefore defines a graded ring $A_{\mathbb{Q}}(M, \mathcal{G})$, with top-degree map $\deg_{M, \mathcal{G}}$, for every loopless matroid (not necessarily realizable); its maximal-building-set case is the ring in which Adiprasito–Huh–Katz [2] proved log-concavity for matroids; Larson–Li–Payne–Proudfoot [25] constructed the combinatorial K-ring $K(M, \mathcal{G})$, generated by classes τ_F .

One basic piece of the geometry had no counterpart in this dictionary: the tangent bundle of $W_{L, \mathcal{G}}$. The problem is to construct a tangent class $T_{M, \mathcal{G}}$ in the integral combinatorial K-ring of every loopless matroid and every building set containing the top flat E , such that (i) $T_{M, \mathcal{G}}$ specializes to the class of the actual tangent bundle of the wonderful compactification $W_{L, \mathcal{G}}$ for a realizable matroid M realized by L , (ii) its

Hirzebruch–Riemann–Roch numbers recover the graded dimensions of the Chow ring, and (iii) its Chern numbers against the combinatorial hyperplane class are bounded below by $\binom{d+1}{k}$, their value for projective d -space. The construction and key properties of such tangent classes were recently studied by Cheng [10]. The paper of this case study¹ constructs the class first in the rational K-ring and proves the three key properties mentioned above, and then adapts it to the integral version, independently of [10], to which, as described below, it had no access.

Theorem 6 (Integral tangent classes for matroid wonderful models [11]). *Let M be a loopless matroid of rank $d+1$ on a finite nonempty ground set E , and let \mathcal{G} be a Feichtner–Yuzvinsky building set on M containing the top flat E . Then there is a class $Q_{M,\mathcal{G}}^{\mathbb{Z}}$ in the integral combinatorial K-ring $K_{\mathbb{Z}}(M, \mathcal{G})$ such that*

$$T_{M,\mathcal{G}}^{\mathbb{Z}} := \sum_{F \in \mathcal{G} \setminus \{E\}} (1 - \tau_F)^{-1} - Q_{M,\mathcal{G}}^{\mathbb{Z}}$$

is a genuine element of $K_{\mathbb{Z}}(M, \mathcal{G})$, with integer coordinates in the standard τ -monomial basis, and, writing $T_{M,\mathcal{G}}$ for its rationalization: (i) if M is realized by a complex linear subspace L , then $T_{M,\mathcal{G}}^{\mathbb{Z}}$ maps to the class of the tangent bundle of $W_{L,\mathcal{G}}$ under an integral ring isomorphism $K_{\mathbb{Z}}(M, \mathcal{G}) \xrightarrow{\sim} K_0(W_{L,\mathcal{G}})$; (ii) for every $0 \leq i \leq d$,

$$\dim_{\mathbb{Q}} A_{\mathbb{Q}}(M, \mathcal{G})^i = (-1)^i \deg_{M,\mathcal{G}}(\text{ch}(\wedge^i T_{M,\mathcal{G}}^{\vee}) \text{td}(T_{M,\mathcal{G}})),$$

with ch and td taken in the formal splitting-root sense; (iii) for every $0 \leq k \leq d$,

$$\deg_{M,\mathcal{G}}(c_k(T_{M,\mathcal{G}}) \alpha^{d-k}) \geq \binom{d+1}{k}, \quad \alpha := -x_E.$$

How Danus produced it. The problem was posed, in a fixed verbatim prompt and with no suggested route, to three systems: the web interface of GPT-5.5-pro, Rethlas, and Danus. The experiment was run before [10] appeared on arXiv, and that paper was deliberately withheld from all three, so the problem they faced was open. Neither of the first two produced a solution. Rethlas, run on the problem three times, returned nothing that passed the verifier; its attempts foundered precisely on the confluences the problem invites, for instance between a realizable matroid, its wonderful model and the permutohedral toric variety, or between integral and rational K-theory. Danus, which runs the same workers against the same verifier, solved the rational form of the problem and then adapted it to the integral version. Seven workers explored constructive and refutational routes in parallel for about five days. The search was far broader than the proof it left behind: the final fact graph, shown in Figure 2, holds 3,157 verified facts in dependency chains up to 54 facts deep, of which 664 form the supporting closure of the theorem, and the global memory records 636 proof attempts, 151 counterexamples, and 25 dead ends. Danus then turned the verified fact graph into a manuscript; reorganizing the verified material into linear prose introduced errors of its own, so the draft was passed back through the verifier and corrected section by section. After the human experts observed that the completed manuscript solved only the rational version of the problem, whose original statement had asked for an integral class, Danus resumed work and produced a solution to the integral version. Details of the experiment are recorded in [11, Appendix B].

Human input. Beyond the common input, human input consisted of a single mathematical intervention, which came after the system had declared the task finished: the human experts pointed out that the solution addressed the rational rather than the integral version of the problem, upon which Danus produced the integral version. The human authors then checked the manuscript and found it correct, with one local exception: the justification given for [11, Lemma 8.7] (used in the Chern– α lower bound) is incomplete as written, the lemma having been treated as proved without a valid argument. The lemma is nonetheless true and admits a short proof (see [10, Proposition 4.19]), and this small gap affects neither the construction of the tangent class nor the $P^K = \text{Hilb}$ identity.

¹[11] R. Cheng, S. Liu, and G. Gao, *Tangent classes of matroids and wonderful compactifications*. The paper will be posted to arXiv within two days, simultaneously with the present paper; this footnote will then be replaced by the citation.

What this case shows. The final case is the largest of the six searches and the closest our experiments come to a controlled comparison between Danus and its predecessor: Rethlas and Danus ran the same worker and verifier models on the same prompt, so the difference in outcome, three failed runs against a complete verified solution, reflects the orchestration around the models rather than the models themselves. The comparison isolates two capabilities. First, Danus accommodates a much greater depth of reasoning: the proof grew fact by fact into a verified closure of hundreds of statements, and the long-horizon design let the run continue until the theorem was verified rather than for a preset number of rounds. Second, Danus is sensitive to mathematics whose difficulty lies in fine distinctions: where the single-line attempts of Rethlas conflated nearby notions, in Danus several workers attacked the problem simultaneously from different directions, including attempts at refutation, and every claim had to pass the verifier in the exact form in which later facts would use it, so the distinctions the problem turns on stayed intact across hundreds of steps. The class Danus constructed recovers the tangent class and its three key properties of [10], which it had never seen.

4 Discussion

The case studies of Section 3 show what the design of Section 2 delivers: proofs built from hundreds of verified facts, produced in days, with human intervention concentrated at a few decisive points. They show its boundaries just as plainly: the manuscripts still need expert review, and the hardest steps of several searches were unlocked by a human. This section takes up four questions these observations raise: Section 4.1 asks why each model holds its role and what the assembly gains over its parts; Section 4.2, how the writing stage balances fluency against faithfulness; Section 4.3, why reliable verification underlies everything else; and Section 4.4, in what sense the design scales test-time computation and where the scaling stops. Section 4.5 then collects the strengths the experiments exhibited and the limitations they exposed.

4.1 Comparison with GPT-5.5, Claude Opus 4.8, GPT-5.5-pro, and Rethlas

The assignment of models to roles follows their measured strengths. In our experiments GPT-5.5 showed stronger mathematics than Claude Opus 4.8, so the workers and the verifier run on Codex agents with GPT-5.5. Claude Code with Claude Opus 4.8, in turn, was better suited to the main agent’s task of reading a large, evolving body of logs, memory, and a fact graph of thousands of entries (Section 2.4). GPT-5.5-pro has the strongest mathematics of the three but is too costly to call continuously, so it serves as the low-frequency expert reference. We have not yet systematically evaluated the more recent Fable 5 from Anthropic.

Assembled this way, the system is stronger than any of its parts. That the harness can matter more than the raw model was already visible in Rethlas: an agent running on GPT-5.4 and GPT-5.5 solved a research-level problem in algebraic groups on which GPT-5.5-pro, through its web interface, produced a completely wrong proof (Section 5.1.1 of [21]). Moreover, no single model call is designed to return a verified proof of the length our case studies require. Even GPT-5.5-pro offers at best an outline of a direction, often an erroneous one; when posed our problems directly, it produced no meaningful result in any case (Section 3).

Much of the design can therefore be read as machinery for eliciting from each model more than a direct call obtains. The verifier holds the reasoning to mathematical rigor: context isolation and skills that spell out the checking logic largely suppress hallucination and skipped steps, which no direct call achieves, even one explicitly instructed to check its result before answering (Section 2.5). Repeated calls and independent workers absorb the errors that even the strongest models make: a GPT-5.5-pro consultation enters the memory as guidance, never the fact graph as truth; a wrong direction is tried, refuted, and recorded as a dead end, while a right one, even one in ten, is carried through into verified facts. The search thus profits from GPT-5.5-pro’s strength at whatever reliability it offers. Parallelization scales the search with test-time computation, and the fact graph and the memory extend its length: each agent handles a small task while the system advances on the large one (Section 4.4).

The assembled system thus outperforms not only its models called alone but also Rethlas itself, which continues inside Danus essentially as a single worker–verifier loop: on the same problem, with the same models, Rethlas failed three times where Danus produced the verified solution (Section 3.6).

4.2 Balance between writing fluency and truth faithfulness

Faithfulness to the proof record and readability pull in opposite directions, and the tension grows with the size of the fact graph. The most faithful manuscript would simply write out, one after another, every fact the target theorem depends on; it would be correct, since each fact has passed verification, and unreadable, running to hundreds of facts at the scale of our runs. A readable paper instead needs motivation, an order of presentation different from the order of discovery, and the omission of routine detail. But this rewriting is itself mathematical writing, and it introduces errors at a rate far above the graph’s own. The characteristic failure, seen repeatedly in our experiments, is a manuscript whose results, constructions, and route are all correct while the written reasoning goes wrong exactly where several facts were compressed into one (Sections 2.7 and 3.6).

The remedy is to pursue the two goals in sequence: the draft is written for readability, and the completed manuscript must then pass the verifier as written. Until it does, it is revised and resubmitted, the worker’s submit–verify–repair cycle lifted to the whole paper, except that here the main agent faces the verifier and writes no new facts: each revision retrieves existing facts from the graph and renders them more accurately (Sections 2.5 and 2.7). This removes most writing-stage errors while the paper remains a paper, not a fact-by-fact transcript.

What remains is a balance rather than a resolution: the manuscripts can still explain motivation too thinly or compress steps too far, and demanding more rigor pushes the text back toward the transcript. The released versions are the point we chose, and the residual defects are small: the human revisions recorded in Section 3 were chiefly adjustments of notation and citation format.

4.3 Importance of verification

Danus takes its verifier from Rethlas and deliberately narrows its role: an independent, stateless service concerned with nothing but the correctness of the text in front of it (Section 2.5). The isolation is essential: prompting a model to check its own work suppresses hallucination only mildly; the checker must be a separate agent whose context holds the submission, the cited facts, and nothing else, and whose skills prescribe the checking procedure. Engineered this way, the verifier almost never accepts a flawed proof: across our experiments the few slips on record were caught in the final human review (Section 4.5), and revocation, the graph’s repair channel, was needed only once (Section 3.5).

This precision is what makes the fact graph workable. A worker trusts every existing fact without re-deriving it (Section 2.2), so a fact is only as reliable as everything beneath it. The theorem of Section 3.6 rests on 664 supporting facts in chains up to 54 deep, and a structure of this depth stands or falls with the verifier: if verification is imprecise, errors accumulate along the chains, the facts higher up become unreliable, and long-horizon reasoning is impossible; if it is precise, accumulation is safe. The verifier’s precision is thus the load-bearing assumption of the architecture.

Verification also changes the economics of the human–machine loop. In the writing loop of Section 4.2 it acts as a first referee: what it rejects is repaired by the machine and never reaches the expert. This matters because human checking, not machine search, is the bottleneck; results produced in a day or two typically took experts a week or two to check. Verification, more than search speed, is what accelerates the joint production of new mathematics.

4.4 Test-time scaling in width and in depth

The design began with a simple intention: Rethlas spends its test-time computation on one line of reasoning at a time, and we wanted to spend more of it, on many lines at once. Simply adding workers fails: all of them edit the single blueprint carrying the whole proof, and each one’s progress is every other’s interference (Section 2.2). The fact graph makes the parallel work add up: it splits the proof into facts a single worker can own, so contributions accumulate instead of colliding, and the main agent spreads the workers over different lines of attack (Sections 2.2 and 2.4). In the run of Section 3.6, seven workers built 3,157 verified facts in five days; only 664 of them support the theorem, and the rest record how widely the system searched. This is scaling in width: added computation becomes added exploration.

Width alone buys many short arguments rather than one long one; length comes from trusting what is already built and remembering what has already been tried. Verification provides the trust: every fact is checked on admission (Section 4.3), so a worker can stand on a chain 54 facts deep without re-deriving any of it. Memory does the remembering: raw exploration stays in the local logs, distilled lessons rise into the global memory, and the main agent’s summaries compress the state of the run into a page (Sections 2.3 and 2.4). Each level condenses the one below it, so what any agent reads stays bounded while what the system knows grows: a worker on its tenth task inherits the lessons of the first nine from a summary, not five days of logs. This is scaling in depth: an effective reasoning horizon far beyond a single context window.

The machinery behind both axes is modest, a main agent, a fact graph, and the interaction logic between them, yet it separates the case studies of Section 3 from what the same models, or Rethlas itself, produce alone (Section 4.1). The same experiments also mark where the scaling stops. The scaling works where a path to the solution exists and needs only to be found: width locates the entry points, depth carries the argument through. It does not create the path: when the solution demands an idea beyond what any single call proposes, more workers and longer horizons merely let the system circle, accumulating shallow conclusions that never break the problem open. In our experiments such impasses were broken by a human supplying the missing idea (Section 3.2); for the deepest open problems there may be no one to supply it, and attacking them will need agent frameworks more inventive than the present one.

4.5 Strengths and limitations

The preceding subsections each pursued a single question; to close the discussion, we summarize Danus’s capabilities as a whole. Danus showed many strengths in our experiments. First, the quality of the verifier was satisfactory: we observed only a handful of verifier errors in total, most of which stemmed from imprecise or erroneous statements in the cited references (Section 3.5). The fact graph could therefore serve reliably as the system’s source of truth. On this foundation, with the main agent directing the search and the memory carrying its lessons forward, Danus extended its work in depth and in width: it carried an intricate technique into a new setting in one sustained synthesis (Section 3.1), kept a long parallel search running until the target theorem stood verified (Section 3.6), and decomposed problems unprompted into parts worked along constructive and refutational routes (Sections 3.2 and 3.3). In work of this breadth and length, Danus showed notable mathematical capability: at times it completed an entire proof with little or no human assistance, in one case the whole pipeline from problem statement to finished manuscript (Section 3.4); it reached solutions by routes the human experts had not anticipated (Sections 3.2 and 3.5), drew tools from branches of mathematics distant from the problem’s own (Section 3.3), and, when a supplied reference was suboptimal, sought out a better one on its own (Section 3.1). Finally, Danus displayed a measure of writing ability: when a first draft rendered a key lemma incorrectly, the verifier rejected the draft and Danus repaired it automatically (Section 3.3); another manuscript was corrected section by section the same way (Section 3.6).

At the same time, the experiments exposed several shortcomings. There remain cases in which Danus does not find the correct route to a proof on its own, and a hint from a human expert is then needed (Section 3.2). Even when the route is found, Danus’s paper writing is not yet fully satisfactory: citation formatting occasionally strays from the conventions of the field, and its notation, though logically consistent, does not always guard against confusion; these imperfections required manual adjustment by the human experts afterwards. In searching the literature, Danus tends to settle for references that are mathematically sufficient, rather than the ones that would make a proof simplest: the simplifying criterion in the case of Section 3.1 had to be pointed out by a human expert. Finally, Danus relies on the information it extracts from its references, and can inherit an error when a reference is itself flawed; errors of this kind are fully resolved only by expert review: in the case of Section 3.5, an erroneous definition propagated into the argument until the experts caught it and the facts resting on it were revoked (Section 2.2).

5 Conclusion

In this paper, we propose Danus, an orchestration system for research-level mathematical reasoning centered on a shared fact graph as a global memory-management mechanism. Danus separates global planning, parallel proof search, and verification: a main agent coordinates the overall strategy, multiple worker agents explore

and prove local claims in parallel, and a stateless verifier checks proposed claims before they are admitted into the fact graph. By storing each verified fact together with its proof and logical dependencies, the fact graph allows Danus to accumulate long mathematical arguments incrementally while keeping the shared proof state organized and reliable.

We evaluated Danus through six research-level case studies in algebraic geometry, singularity theory, and combinatorics. These examples show that Danus can construct long and detailed proofs, decompose difficult problems into workable subproblems, repair flawed proof dependencies, and turn verified fact graphs into mathematical manuscripts. The amount of human input varied across the case studies, ranging from no mathematical guidance to a small number of high-level hints or corrections, but in all cases Danus carried out substantial proof construction.

Our results suggest that fact-graph-based orchestration provides an effective mechanism for extending LLM-based mathematical reasoning agents to long-horizon research problems. By turning parallel local proof search into an organized and cumulative proof-building process, Danus enables many contributions from different workers to be assembled into coherent mathematical arguments. At the same time, our experiments show that human mathematicians remain essential for choosing meaningful problems, identifying missing ideas, and carrying out final expert review. We therefore view Danus as a step toward collaborative mathematical research systems in which AI agents can contribute substantially to proof construction while remaining integrated with human mathematical judgment.

Acknowledgements

This work is supported in part by the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (JYB2025XDXM113), the National Key R&D Program of China grant 2024YFA1014000, and the New Cornerstone Investigator Program.

References

- [1] Mohammed Abouzaid, Nikhil Srivastava, Rachel Ward, and Lauren Williams. First proof second batch. [arXiv preprint arXiv:2606.18119](#), 2026.
- [2] Karim Adiprasito, June Huh, and Eric Katz. Hodge theory for combinatorial geometries. *Annals of Mathematics*, 188(2):381–452, 2018.
- [3] Valery Alexeev. Two two-dimensional terminations. *Duke Mathematical Journal*, 69(3):527–545, 1993.
- [4] Valery Alexeev. Boundedness and K^2 for log surfaces. *International Journal of Mathematics*, 5(6):779–810, 1994.
- [5] Chenyang An, Qihao Ye, Minghao Pan, and Jiayaun Zhang. Qed: An open-source multi-agent system for generating mathematical proofs on open problems. [arXiv preprint arXiv:2604.24021](#), 2026.
- [6] Federico Ardila-Mantilla, Nima Arkani-Hamed, Carolina Figueiredo, and Francisco Vazão. Combinatorics of the cosmohedron. [arXiv preprint arXiv:2603.03425](#), 2026.
- [7] Nima Arkani-Hamed, Carolina Figueiredo, and Francisco Vazão. Cosmohedra. *Journal of High Energy Physics*, 2025(11):Paper No. 29, 58 pp., 2025.
- [8] Fedor Bogomolov and Michael McQuillan. Rational curves on foliated varieties. In *Foliation theory in algebraic geometry*, pages 21–51. Springer, 2016.
- [9] Frédéric Campana and Mihai Păun. Foliations with positive slopes and birational stability of orbifold cotangent bundles. *Publications mathématiques de l’IHÉS*, 129:1–49, 2019.
- [10] Ronnie Cheng. Tangent classes for matroid building sets. [arXiv preprint arXiv:2606.22650](#), 2026.
- [11] Ronnie Cheng, Shurui Liu, and Guoxiong Gao. Tangent classes of matroids and wonderful compactifications. [arXiv preprint](#), 2026. in preparation.
- [12] Corrado De Concini and Claudio Procesi. Wonderful models of subspace arrangements. *Selecta Mathematica*, 1(3):459–494, 1995.

-
- [13] Xu'an Dou and Zeyu Jin. Degenerate constants in degree inequalities for sobolev circle maps: on some problems posed by brezis. [arXiv preprint arXiv:2605.24626](#), 2026.
- [14] Eva Maria Feichtner and Sergey Yuzvinsky. Chow rings of toric varieties defined by atomic lattices. *Inventiones mathematicae*, 155(3):515–536, 2004.
- [15] Tony Feng. Eigenweights for arithmetic Hirzebruch proportionality. *PNAS Nexus*, 5(5):pgag143, 2026.
- [16] Tony Feng, Trieu H Trinh, Garrett Bingham, Dawsen Hwang, Yuri Chervonyi, Junehyuk Jung, Joonkyung Lee, Carlo Pagano, Sang-hyun Kim, Federico Pasqualotto, et al. Towards autonomous mathematics research. [arXiv preprint arXiv:2602.10177](#), 2026.
- [17] Jingjun Han and Chen Jiang. Total cartier index of a bounded family. *Pure and Applied Mathematics Quarterly*, 22(1):171–179, 2026.
- [18] Jiedong Jiang, Yixiao Li, Zeming Sun, Yuefeng Wang, Liang Xiao, and Jiahong Yu. On some open problems in commutative algebra resolved by rethlas. [arXiv preprint arXiv:2605.25259](#), 2026.
- [19] Eric Jovinelly, Brian Lehmann, and Eric Riedl. Optimal bounds in bend-and-break. *Forum of Mathematics, Pi*, 14:e16, 2026.
- [20] Haocheng Ju, Leheng Chen, Peihao Wu, Bryan Dai, and Bin Dong. Matlas: A semantic search engine for mathematics. [arXiv preprint arXiv:2604.17484](#), 2026.
- [21] Haocheng Ju, Guoxiong Gao, Jiedong Jiang, Bin Wu, Zeming Sun, ShuRui Liu, Leheng Chen, Yutong Wang, Yuefeng Wang, Zichen Wang, Wanyi He, et al. Automated conjecture resolution with formal verification. [arXiv preprint arXiv:2604.03789](#), 2026.
- [22] Stefan Kebekus, Luis Sol'a Conde, and Matei Toma. Rationally connected foliations after bogomolov and mc-quillan. *Journal of Algebraic Geometry*, 16(1):65–81, 2007.
- [23] János Kollár. Maps between local picard groups. *Algebraic Geometry*, 3(4):461–495, 2016.
- [24] János Kollár and Nicholas I Shepherd-Barron. Threefolds and deformations of surface singularities. *Inventiones mathematicae*, 91(2):299–338, 1988.
- [25] Matt Larson, Shiyue Li, Sam Payne, and Nicholas Proudfoot. K-rings of wonderful varieties and matroids. *Advances in Mathematics*, 441:109554, 2024.
- [26] Joonkyung Lee and Jaehyeon Seo. Lower bounds for multivariate independence polynomials and their generalisations. [arXiv preprint arXiv:2602.02450](#), 2026.
- [27] Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. Meta-harness: End-to-end optimization of model harnesses. [arXiv preprint arXiv:2603.28052](#), 2026.
- [28] Zhangsong Li. On injectivity of phase retrieval. [arXiv preprint arXiv:2606.17922](#), 2026.
- [29] Jihao Liu. Factorial asymptotics of the Matryoshka numbers. Note, <https://jihaoliu.org/notes/Matryoshka.pdf>, 2026.
- [30] Jihao Liu and Sheng Qin. Shokurov's global index conjecture for threefold foliations. [arXiv preprint arXiv:2605.22735](#), 2026.
- [31] Jihao Liu and Xiping Zhang. Criteria of isolated weighted homogeneous hypersurface singularities using logarithmic vector fields. [arXiv preprint arXiv:2606.29891](#), 2026.
- [32] Jihao Liu and Xiping Zhang. A criteria of weighted homogeneity via logarithmic vector fields. [arXiv preprint arXiv:2606.29886](#), 2026.
- [33] Jihao Liu, Fanjun Meng, and Lingyao Xie. Complements, index theorem, and minimal log discrepancies of foliated surface singularities. *European Journal of Mathematics*, 10(1):Paper No. 6, 2024.
- [34] Jihao Liu, Ruicheng Hu, and Sheng Qin. Boundedness of total cartier indices for rational singularities in families. [arXiv preprint arXiv:2605.22782](#), 2026.
- [35] Jihao Liu, Zeming Sun, and Jiedong Jiang. Optimal bend-and-break for foliations. [arXiv preprint arXiv:2605.20754](#), 2026.

-
- [36] Diogo da Silva Machado and Jose Seade. Generic vector fields on isolated complex hypersurface germs. arXiv preprint arXiv:2605.09210, 2026.
- [37] Yoichi Miyaoka and Shigefumi Mori. A numerical criterion for uniruledness. Annals of Mathematics, 124(1):65–69, 1986.
- [38] Shigefumi Mori. Threefolds whose canonical bundles are not numerically effective. Annals of Mathematics, 116:133–176, 1982.
- [39] OEIS Foundation Inc. Entry A177384. The On-Line Encyclopedia of Integer Sequences, <https://oeis.org/A177384>, 2010.
- [40] Xiangyu Pan and Jiahong Yu. Lift-independence problem in the p -adic simpson correspondence for curves. arXiv preprint arXiv:2605.29947, 2026.
- [41] Anand Patel. The simplicity of the hodge bundle. Proceedings of the National Academy of Sciences, 123(21):e2610183123, 2026.
- [42] Jorge Vitório Pereira. On the height of foliated surfaces with vanishing kodaira dimension. Publicacions matemàtiques, 49(2):363–373, 2005.
- [43] Kyoji Saito. Quasihomogene isolierte Singularitäten von Hyperflächen. Inventiones mathematicae, 14(2):123–142, 1971.
- [44] Johannes Schmitt, Tim Gehringer, Jasper Dekoninck, Gergely Bérczi, Uri Kreitner, and Liam Price. Proof-Council: An llm agent for solving open mathematical problems. <https://github.com/eth-sri/proof-council>, 2026.
- [45] NI Shepherd-Barron. Miyaoka’s theorems on the generic seminegativity of t_x and on the kodaira dimension of minimal regular threefolds. Flips and abundance for algebraic threefolds, pages 103–114, 1992.
- [46] Calum Spicer. Higher-dimensional foliated mori theory. Compositio Mathematica, 156(1):1–38, 2020.
- [47] Yanning Xu. Complements on log canonical Fano varieties and index conjecture of log Calabi–Yau varieties. PhD thesis, University of Cambridge, 2020.
- [48] John Yang, Carlos Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering. Advances in Neural Information Processing Systems, 37:50528–50652, 2024.
- [49] Daniel Zheng, Ingrid von Glehn, Yori Zwols, Iuliya Beloshapka, Lars Buesing, Daniel M Roy, Martin Wattenberg, Bogdan Georgiev, Tatiana Schmidt, Andrew Cowie, et al. Ai co-mathematician: Accelerating mathematicians with agentic ai. arXiv preprint arXiv:2605.06651, 2026.