REUF2F: BENCHMARKING CODING AGENTS ON FORMALIZED UNDERGRADUATE RESEARCH PROBLEMS

Anonymous authors
Paper under double-blind review

ABSTRACT

As AI systems begin to resolve open research problems, reliable verification becomes increasingly important. Formal mathematics provides a trusted framework for mathematical reasoning, enabling agents to produce machine-checkable proofs while conducting mathematical exploration. Existing formal mathematics benchmarks often focus on known results or a small set of prominent open problems, which can become difficult to evaluate reliably once solutions are discovered and widely disseminated. Existing evaluation protocols are also not always aligned with how modern coding agents operate. We therefore introduce REUF2F, a Lean 4 benchmark of 244 open problems drawn from Research Experiences for Undergraduates (REU), spanning entry-level open research problems across a broad range of mathematical subfields. We expect these less prominent problems to attract less concentrated attention and their eventual solutions to circulate less widely, reducing the risk of future solution leakage. We adopt a standard coding-agent evaluation setting, placing each problem and its logical negation into a Lean project in a Docker container. We also build Leanimum-agent on mini-SWE-agent to provide a unified harness for fair comparisons across models. Our evaluation finds that even the best-performing model resolves only ~20% of the problems, leaving substantial room for improvement on formalized open research problems.

1 INTRODUCTION

Large language models (LLMs) have made remarkable progress over the past few years, particularly in mathematics and coding. In both domains, solutions can be difficult to construct but comparatively easy to verify. This makes them well suited to reinforcement learning with verifiable rewards (RLVR), as demonstrated by (Guo et al., 2025; Yue et al., 2025). In mathematics, progress has extended from solving competition problems to resolving open research problems that had resisted expert efforts for decades (Bukh et al., 2025; Zheng et al., 2026). Notable examples include the disproof of the unit distance conjecture (OpenAI, 2026d), a complex structure on the six-sphere (Alpöge, 2026), the ten advances in mathematics and theoretical computer science reported by OpenAI (OpenAI, 2026e), and the recent finite-time blowup results for the Navier–Stokes equations (OpenAI, 2026b). For human experts, even verifying such results can be a substantial undertaking. Rapid progress in mathematical reasoning therefore creates a corresponding challenge for evaluation, as most research problems are proof-based, and the correctness of a natural-language proof is difficult to assess reliably at scale.

Formal proof assistants such as Lean 4 provide a natural connection between mathematics and coding. Through the Curry–Howard correspondence, Lean represents mathematical statements as types and proofs as terms whose correctness can be checked by its kernel (Moura & Ullrich, 2021). This allows an argument to be verified without first requiring a human reader to develop a detailed understanding of the argument. Progress in coding agents also makes turning mathematical proofs into formal code increasingly practical. Recent projects tackle Fermat's Last Theorem, three-dimensional Kakeya, and sphere packing (Anthropic, 2026; Bailey et al., 2026; The Sphere Packing in Lean Contributors, 2024). Anthropic reports that its agent-assisted formalization of Fermat's Last Theorem took 11 days (Anthropic, 2026), illustrating how agents can make Lean implementation less of an obstacle to mathematical evaluation. Formal proofs are also becoming part of the presentation of new mathematics, as illustrated by the Polynomial Freiman–Ruzsa project (The PFR Formalization Contributors, 2026), work on diagonal Ramsey bounds (Mehta, 2023), and Lean certificates ac-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Table 1: Comparison with representative mathematics benchmarks.

Benchmark	# Tasks	Setting	Formalized?	Open?	Diverse subfields?	Lower leakage risk?
AIME	30	Competition	✗	✗	✗	✗
IMO-Bench (Luong et al., 2025)	400/60/1,000/60	Competition	△	✗	✗	✗
MathArena Apex (Dekoninck et al., 2026)	12	Competition	✗	✗	✗	✗
ArXivMath (Dekoninck et al., 2026)	30/48	Research	✗	✗	✓	✗
First Proof (Abouzaid et al., 2026a;b)	10/10	Research	✗	△	✓	✓
Soohak (Son et al., 2026)	439	Research	✗	△	✓	✓
FrontierMath (Glazer et al., 2024)	350	Research	✗	△	✓	✓
miniF2F (Zheng et al., 2021)	244	Competition	✓	✗	✗	✗
PutnamBench (Tsoukalas et al., 2024)	640	Competition	✓	✗	✗	✗
FATE (Jiang et al., 2026)	150/100/100	Research	✓	✗	✗	✗
ProofNet (Azerbayev et al., 2023)	371	Undergraduate	✓	✗	✓	✗
FormalProofBench (Ravi et al., 2026)	100	Research	✓	✗	✓	✗
ArXivLean (Dekoninck et al., 2026)	41/46	Research	✓	✗	✓	✗
Lean-eval (Lean FRO, 2026b)	128	Research	✓	✗	✓	✗
FrontierMath Erdős (Adamczewski & Bloom, 2026)	68	Research	✓	✓	✗	✗
OEIS Open (Adamczewski, 2026)	492	Research	✓	✓	✗	✓
Formal Conjectures (Firsching et al., 2026)	2,615	Research	✓	△	✓	✗
REUF2F (ours)	244	Research	✓	✓	✓	✓

✓ yes; ✗ no; △ mixed. *Research* includes advanced graduate and research-facing tasks, not necessarily unsolved problems. *Open?* means open at collection. *Diverse subfields* means substantial coverage beyond one specialized area. *Lower leakage risk* is a qualitative assessment of whether solutions, once available, are less likely to be readily retrievable online during the benchmark's evaluation lifecycle. Counts: First Proof first/second batch; IMO-Bench Answer/Proof/Grading/LeanProof; ArXivMath and ArXivLean Mar./Jun. 2026; FrontierMath original Tiers 1–4 total (300 core + 50 Tier 4); miniF2F test; FATE M/H/X; Lean-eval v1.

companying recent Navier–Stokes and Euler results (OpenAI, 2026c). This development has made Lean-based open-problem benchmarks an increasingly practical setting for evaluating the mathematical reasoning capabilities of language models. While Lean enables verification of candidate solutions, many collections such as FrontierMath Erdős and Formal Conjectures include research questions of substantial community interest (Adamczewski & Bloom, 2026; Firsching et al., 2026). Once one of these problems is solved, its proof is often published or widely shared online. Since mathematical research naturally involves searching the literature, an agent may retrieve that solution during a later evaluation. The same dissemination that benefits mathematical research can thus undermine a problem's value as a held-out test of new problem solving.

To address this problem, we introduce REUF2F, a Lean 4 benchmark of 244 open conjectures from undergraduate research, primarily Research Experiences for Undergraduates (REU) programs. REUs give students the opportunity to conduct original research with a mentor. Their projects are chosen both for their mathematical interest and for their value in training young researchers, and often leave precise conjectures for further investigation. Compared with prominent conjectures that have attracted sustained attention across a field, we expect these questions to attract less concentrated attention and their solutions to circulate less widely. They therefore offer a source of genuine research problems with a lower expected risk of future solution leakage. The mathematicians and formalization experts on our team selected and formalized 244 such conjectures across a broad range of mathematical areas. Table 1 summarizes the features of REUF2F and compares it with existing mathematics benchmarks.

We evaluate REUF2F using the task setup and submission workflow of modern coding-agent benchmarks. Each task is placed in an independent sandbox where the agent develops its solution. Because an open conjecture can be false, each task provides both the original conjecture and its full logical negation. The agent is free to pursue either direction during its attempt. After the agent submits its solution, we then check the submitted code in a separate evaluation container. We use Comparator to check that the submission establishes the original statement or its negation without changing the underlying definitions or introducing unproved assumptions. A problem counts as solved when a proof of either direction passes both Lean's kernel and the independent nanoda checker. Both the problem and its accepted solution are thus entirely formal. Following miniF2F (Zheng et al., 2021), the F2F in REUF2F refers to this formal-to-formal setup.

For fair comparisons across models, we also build Leanimum-agent, a lightweight adaptation of mini-SWE-agent (Yang et al., 2024) for formal theorem proving. It provides every model with the same task instructions, shell interface, context handling, and interaction loop. Through Leanimum-agent, models can read and edit files, search Mathlib, run computations, and revise proofs using

Lean kernel feedback. This setup retains the capabilities of a modern coding-agent harness while minimizing performance differences attributable to model-specific tools or agent designs.

We evaluate frontier models on REUF2F. Under the tested budgets, the best-performing model, GPT-6 Astra, resolves 19.67% of the problems on average over two runs. These results leave substantial room for improvement and show that REUF2F remains challenging for current coding agents. Progress on REUF2F would therefore provide a concrete indicator of agents' ability to resolve open mathematical questions with machine-checkable proofs.

2 BACKGROUND

2.1 AI FOR MATHEMATICS AND EVALUATION

Over the past couple of years, LLMs' mathematical capabilities have advanced rapidly from solving competition problems (Guo et al., 2025; Hubert et al., 2026; Chen et al., 2025b; Google DeepMind, 2025; Chen et al., 2025a), to tackling genuine research questions. Examples include resolutions of Erdős problems (Tsoukalas et al., 2026), the unit distance conjecture (OpenAI, 2026d), the cycle double cover problem (OpenAI, 2026a; Oum, 2026), a complex structure on the six-sphere (Alpöge, 2026), OpenAI's ten-result collection (OpenAI, 2026e), and finite-time blowup for Navier–Stokes (OpenAI, 2026b). These selected successes, while impressive, do not establish a systematic comparison of LLMs' overall mathematical capabilities.

Mathematical benchmarks have become more challenging in response, progressing from educational and competition problems to advanced and research-level tasks (Cobbe et al., 2021; Hendrycks et al., 2021; He et al., 2024; Gao et al., 2025; Glazer et al., 2024; Schmitt et al., 2025). Yet evaluation methods have struggled to keep pace as model capabilities have rapidly improved. Performance on most established benchmarks has approached saturation, leaving little room to distinguish stronger models. Moreover, most research questions are proof-based, and checking a proposed solution may require specialist knowledge and close reading, making it difficult to obtain evaluations that are both reliable and scalable.

2.2 FORMAL THEOREM PROVING

Formal theorem proving offers a solution to this verification problem. Proof assistants such as Isabelle, Coq, and Lean encode mathematical statements and arguments in formal languages, allowing proofs to be checked mechanically by a small trusted kernel and providing both reliable correctness guarantees and a scalable alternative to expert evaluation.

Among these systems, Lean 4 has seen particularly broad adoption in current AI-driven formal mathematics. Mathlib, Lean's centralized library, contains more than two million lines of code and covers definitions and results across a wide range of mathematical areas (Moura & Ullrich, 2021; The mathlib Community, 2026). This shared foundation supports formalization projects for highly complex mathematical developments, including sphere packing, three-dimensional Kakeya, and Fermat's Last Theorem (The Sphere Packing in Lean Contributors, 2024; Bailey et al., 2026; Anthropic, 2026), with some projects extending to millions of lines of code. This demonstrates the breadth and depth of mathematics that can be built on Mathlib, thereby illustrating the maturity and expressive power of Lean as a language for formalizing research mathematics.

Proof assistants support research on neural theorem proving and autoformalization, including Draft, Sketch, and Prove (Jiang et al., 2022). In Lean, representative examples include miniCTX (Hu et al., 2025), Lean-STaR (Lin et al., 2025), and BFS-Prover-V2 (Xin et al., 2025). Increasingly, solutions to open mathematical problems are released together with formal proofs that provide directly checkable evidence of correctness, with notable examples including the Polynomial Freiman–Ruzsa conjecture and the recent resolution of Navier–Stokes (The PFR Formalization Contributors, 2026; OpenAI, 2026c). These emerging practices point toward a new paradigm of mathematical verification driven by machine-checkable formal proofs.

2.3 CODING AGENTS

Coding is another major application of LLM agents and an important driver of their recent development. Models have progressed from assisting with short code completions to independently editing repositories, implementing features, and carrying out substantial development tasks (Chen et al., 2021; Jimenez et al., 2024; Yang et al., 2026; Huang et al., 2026; Shepard & Salimans, 2026; Merrill et al., 2026). Formal theorem proving lets us turn a mathematical problem into a coding task. We can therefore reuse the mature training and evaluation infrastructure developed for coding agents by replacing unit-test-based correctness checks with formal proof verification. This provides a natural way to evaluate mathematical reasoning through standard coding-agent interfaces.

2.4 RESEARCH EXPERIENCES FOR UNDERGRADUATES

Research Experiences for Undergraduates (REU) programs give students research experience through work with faculty mentors (National Science Foundation, 2026). Well-known mathematics programs include Duluth REU, Williams SMALL, Stanford SURIM, CMU SURF, and MIT SPUR. Their projects cover all major branches of mathematical research. Although the choice of a project takes students' mathematical backgrounds into account, the work combines research training with genuine mathematical investigation. The resulting papers report new results and often propose conjectures that remain unresolved beyond the project itself.

These conjectures provide a useful source for evaluation. They arise from genuine mathematical research, but typically receive less broad attention than famous conjectures. We expect this lower visibility to make their eventual solutions less likely to become widely disseminated reference answers, thus reducing the risk of future solution leakage.

3 REUF2F

In this section, we introduce REUF2F, a benchmark consisting of 244 formalized open problems drawn from undergraduate research projects. We first formulate the task, then describe how the problems were collected and formalized. The evaluation protocol and implementation details are deferred to Section 4.

3.1 TASK FORMULATION

For each problem P_i in the benchmark, let φ_i denote its formal Lean statement. An open conjecture may ultimately be resolved either by a proof or by a refutation, so we consider both φ_i and its logical negation $\neg\varphi_i$. Let $s_i^+ \in \{0,1\}$ indicate whether a proof of φ_i is accepted by our evaluation protocol, and let $s_i^- \in \{0,1\}$ analogously indicate whether a proof of $\neg\varphi_i$ is accepted. Since both statements cannot admit accepted proofs in a consistent proof system, $s_i^+ s_i^- = 0$. We therefore define the resolution indicator for P_i as $r_i = s_i^+ + s_i^- \in \{0,1\}$, where $r_i = 1$ indicates that the problem has been resolved.

3.2 PROBLEM SELECTION

We collect publicly available research papers, preprints, and project reports from established REU programs and other undergraduate research initiatives. Within each paper, we identify the central research questions that the authors leave unresolved. We verify that each selected question was first proposed in the paper, is central to its research contribution, and remains unresolved. This focuses the benchmark on genuinely open questions arising from the undergraduate research itself, rather than on known results or external conjectures mentioned only for context.

We organize the collected problems into five broad areas: discrete mathematics and combinatorics; number theory and arithmetic; algebra and algebraic geometry; geometry and topology; and analysis, dynamics, probability, and partial differential equations. We include 45 problems from each of these five categories to maintain a balanced disciplinary distribution. The remaining 19 problems, which do not fit naturally into any of these categories, are grouped under Others.

3.3 Formalization and Quality Control

The formalization work is distributed among contributors on our team, including several active Mathlib contributors. Formalizers work primarily on problems within their areas of expertise and submit completed formalizations for review. Before a problem is admitted to the benchmark, we conduct additional independent checks using both Codex and Claude Code to assess whether the problem was originally proposed in the source work, remains novel and unresolved, is central to the research contribution, and has been faithfully formalized in Lean. Any issues identified during these checks are returned to the formalizer for revision or replacement. Problems that pass these checks undergo a final review by a Mathlib contributor before being included in the benchmark.

4 Evaluation Protocol

Traditional formal theorem-proving benchmarks are typically distributed as Lean repositories (Zheng et al., 2021; Tsoukalas et al., 2024; Firsching et al., 2026). Each target theorem contains a `sorry`, Lean's placeholder for a missing proof, which must be replaced to solve the task. How these proof holes are filled is left to the evaluation system: a step-level prover may perform tree search over tactics, a whole-proof model may generate a complete proof script, and an agent may interact iteratively with Lean feedback to construct and refine a proof. The final evaluation signal is often determined by whether the resulting development compiles after the required proof holes have been filled. While flexible, this setup leaves four fundamental issues that motivate the design of our evaluation protocol.

1. Proof-only evaluation. A proof-only task assumes that the supplied proposition is provable. This is appropriate for known theorems with trusted formalizations, but not for open conjectures. If a conjecture is false, failures due to its falsity are difficult to distinguish from failures due to mathematical difficulty. A meaningful evaluation of open problems should therefore accept either a valid proof or a valid refutation.

2. Trusted verification. Lean permits metaprogramming and code execution during proof development. This creates ways for a submission to compile without providing the intended proof, often by interfering with the environment or exploiting implementation bugs, making compilation alone an imperfect correctness signal. This motivates additional independent verification beyond Lean's standard elaboration and kernel checking during `lake build`. (Lean FRO, 2026a; Moura, 2026).

3. Harness confounding. Evaluation results obtained with different harnesses measure the combined system, making the underlying models harder to compare. A leaderboard that includes both heavy multi-agent proving systems and individual agent sessions does not effectively isolate differences in model capability. Vals AI's ProofBench addresses this problem by providing a common, lightweight harness with Mathlib search and Lean execution feedback (Ravi et al., 2026). However, this specialized scaffolding does not give the model the general filesystem access used in modern coding agent designs.

4. Environment dependence. Modern agents necessarily interact with the filesystem. Without a clean and consistently configured environment, performance can depend on hardware, accessible files, installed dependencies, and even the order in which tasks are run if changes persist between attempts. A clean, consistently configured environment is therefore necessary to compare runs reliably.

To address these four issues, we adopt a modern coding-agent evaluation setting for REUF2F, augmented with Lean-specific mechanisms for open-problem evaluation and trusted proof verification. Section 4.1 introduces our two-way protocol and independent verification with Comparator, addressing the first two issues. Section 4.2 then presents our sandboxed evaluation environment and Leanimum-agent, a formal theorem-proving adaptation of mini-SWE-agent, which provide a standardized harness and reproducible execution environment for addressing the remaining two issues.

4.1 Kernel-Trusted Two-Way Evaluation

Since the problems selected for the benchmark are open, their truth values are not known in advance. We therefore allow models to prove either direction. Specifically, for each formalized theorem `foo`

Table 2: Verification layers and the invalid submissions they are designed to detect

Method	What is checked	Errors addressed
`lake build`	Elaboration and compilation	Syntax, elaboration, and type errors
+ Axiom checking	Dependency	Unauthorized axioms
+ SafeVerify	Specification	Type manipulation
+ Comparator	Sandboxed export and kernel replay	`.olean`-level attacks
+ nanoda	Independent kernel checking	Lean kernel bugs

with full Lean type T, we use `Export.lean`, included in Appendix A, to obtain its declaration name and universe parameters. Task preparation then appends a theorem `foo_neg` whose type is `¬ type_of% @foo`; Lean elaborates this to $\mathrm{Not}\ T$, including all of T's binders. We leave it to the agent to decide which direction to pursue and count a problem as resolved when a proof of either direction passes our Lean verification.

The second issue is that "passing Lean verification" can refer to different levels of checking. The Lean developers have progressively strengthened their verification tools in response to adversarial testing, including AI-assisted attempts to find vulnerabilities (Moura, 2026). A successful `lake build` establishes that a project passes the compilation process in its current environment. It does not, by itself, establish that an untrusted submission proves the unchanged target using only the intended assumptions. This distinction matters when the model can execute code and modify the environment in which it develops its proof.

We instead use Comparator (Lean FRO, 2026a), the "Gold Standard" proof validation method recommended by Lean official. It takes a trusted challenge and an untrusted solution, builds and exports them separately in sandboxed environments, and compares their formal specifications. For each target theorem, Comparator checks that the submitted theorem proves the same statement as the trusted challenge and that all declarations transitively referenced by the statement agree between the two environments. It also checks the proof's transitive axiom dependencies and replays the exported solution through Lean's kernel. These checks effectively prevent a submission from changing the target statement or its dependencies, assuming the desired conclusion through an additional axiom, or otherwise bypassing the kernel check. Our configuration permits only `propext`, `Quot.sound`, and `Classical.choice`, so a proof depending on any additional axiom is rejected.

Comparator's Lean-kernel check still relies on the correctness of that kernel implementation. AI-assisted investigations have also exposed bugs at this level, including a flaw in the handling of nested inductive types that allowed an invalid proof of `False` (Moura, 2026). To avoid reliance on any single checker implementation, Comparator supports nanoda (Bailey, 2026), an independent type checker for Lean implemented in Rust. We enable this additional check and require a submission to be accepted by both kernels. Table 2 summarizes the layers of verification and the classes of invalid submissions they address. Our final acceptance criterion is a successful Comparator invocation with nanoda enabled in a fresh evaluation container.

4.2 Sandboxed Environments and Leanimum-agent

The third and fourth issues are not specific to Lean. Coding-agent benchmarks already provide established approaches to harness consistency and environment isolation, which we adopt for formal theorem proving tasks.

The SWE-bench and ProgramBench evaluation use mini-SWE-agent as a common harness, while Terminal-Bench provides the Terminus harness (Yang et al., 2024; Merrill et al., 2026). These shared harnesses standardize the inference loop, including tool protocols, context management, and the agent lifecycle, allowing comparisons to focus more directly on differences between models. We use the same principle to compare models on REUF2F.

For the execution environment, modern coding-agent evaluations typically place each task in its own sandbox container. This provides a reproducible starting state and prevents performance from being affected by differences in filesystem state. We follow the same design for REUF2F. Each task is instantiated as a self-contained Lean project containing the two proof targets described in

Section 4.1, together with a fixed Lean and Mathlib environment. The agent interacts only with this isolated workspace during its attempt, ensuring that every run begins from the same task state.

Each task runs in a fresh Docker container built on codex-universal, with Lean, Mathlib and its cache, Comparator, and the required dependencies preinstalled. We initialize the project, run `lake build`, and create a Git checkpoint before starting the agent. The agent can then inspect the workspace, search Mathlib, edit code, and test proofs. At the end of the attempt, its changes are submitted as a Git patch. We replay this patch in a separate evaluation container and run the Comparator verification described above. This follows the familiar pattern of evaluating a coding agent's patch against tests in a clean environment, with formal proof verification serving as the success criterion instead of conventional unit tests.

Our common harness, Leanimum-agent, is adapted from mini-SWE-agent, whose simple shell-based design has proved effective as a common harness for SWE-bench and ProgramBench (Yang et al., 2024; 2026). Despite its lightweight implementation, it provides the main capabilities needed by a modern coding agent: filesystem access, command execution, persistent file edits, and iterative feedback. We preserve this design and add instructions specific to Lean. The prompt explains the proof-or-refutation objective, names the two target declarations, shows how to search Mathlib with `rg`, and asks the agent to check its changes with `lake build` before submitting a source patch. Appendix B includes the full prompt, and Appendix C gives the verification configuration. All evaluated models use this common harness, so the comparison does not confound model capability with differences in agent scaffolding.

5 EXPERIMENTS

We evaluate frontier models on REUF2F with Leanimum-agent. Our main results are shown in Table 3. Even under a relatively generous inference budget, the best-performing model, GPT-6 Astra, resolves only 19.67% of the problems, and most other models resolve much fewer. These results indicate that REUF2F is a challenging and timely benchmark for current models.

5.1 EVALUATION CONFIGURATION

All experiments use Lean and Mathlib `v4.34.0`. We evaluate GPT-6 Astra (xhigh), GPT-6 Sol (xhigh), Grok 4.7 (xhigh), Muse Spark 1.3 (xhigh), Qwen3.8-Max (max), Kimi K3 (max), and Gemini 3.8 Flash (high), where the values in the parentheses represent the reasoning effort setting. Every model uses the same Leanimum-agent harness described in Section 4.2.

Each agent container and each evaluation container is allocated 20 CPU cores and 60 GB of RAM with Internet access, so that Lean does not become a bottleneck itself. We cap each attempt at 1,000 agent steps and 6 hours of wall-clock time, with a timeout of 1,200 seconds for each individual shell command. No cost limit is imposed. Following ProgramBench (Yang et al., 2026), the agent is reminded to finalize and submit its patch once fewer than 20 steps or 20 minutes remain. Each submitted patch is graded in a fresh evaluation container with a timeout of 1,800 seconds. We choose these limits to be generous, so that the results approximate each model's capability ceiling under the common harness.

We run each model $n = 2$ times independently on every problem. We report pass@1, the average resolution rate over the two runs, and pass@2, the fraction of problems resolved in at least one of the two runs.

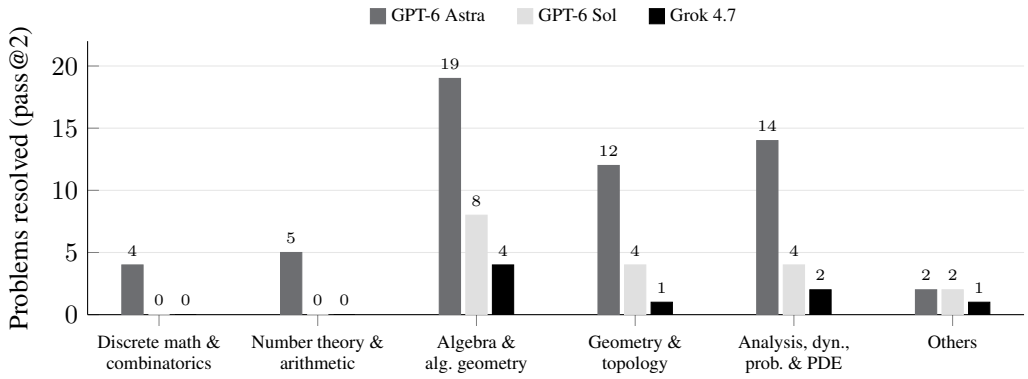


Figure 1: Number of problems resolved under pass@2 in each mathematical area by the three strongest models. Each of the five main areas contains 45 problems, and Others contains 19.

5.2 Main Results

Table 3: Main results on REUF2F ($N = 244$). Tokens and step counts are averaged over all attempts.

Model	Effort	pass@1 (%)	pass@2 (%)	proved@2	refuted@2	Avg. out tokens	Avg. steps
GPT-6 Astra	xhigh	19.67	22.95	7.38	15.57	17.08K	22.0
GPT-6 Sol	xhigh	6.56	7.38	1.64	5.74	45.64K	61.6
Grok 4.7	xhigh	3.07	3.28	0	3.28	159.49K	170.9
Gemini 3.8 Flash	high	2.46	2.87	0	2.87	169.79K	135.18
Kimi K3	max	1.64	2.05	0	2.05	199.76K	60.7
Qwen3.8-Max	max	1.23	2.05	0	2.05	301.65K	110.9
Muse Spark 1.3	xhigh	1.23	1.23	0	1.23	113.26K	98.8

REUF2F is far from saturated. The strongest model, GPT-6 Astra, reaches a pass@1 of 19.67% and a pass@2 of 22.95%, followed by GPT-6 Sol at 7.38% pass@2, Grok 4.7 at 3.28% pass@2, and Gemini 3.8 Flash at 2.87% pass@2. The remaining models resolve 2.05% or fewer problems under pass@2.

Of the 56 problems resolved by at least one model, 18 were resolved by proving the original conjecture and 38 by proving its negation. Interestingly, GPT-6 Astra resolves all 56 of these problems, so every problem resolved by another model is also covered by GPT-6 Astra. Moreover, GPT-6 Astra and GPT-6 Sol are the only models that prove conjectures, while every resolution by the other five models is a refutation. This suggests that finding and certifying a counterexample is considerably more accessible to current agents than constructing a proof.

5.3 Performance across Mathematical Areas

Figure 1 breaks down the problems resolved under pass@2 by GPT-6 Astra, GPT-6 Sol, and Grok 4.7 across the six problem categories. Performance is highly uneven across areas. Most resolutions come from algebra and algebraic geometry, analysis, and geometry and topology, where all three models resolve at least one problem. Discrete mathematics and combinatorics, and number theory and arithmetic, are resolved almost exclusively by GPT-6 Astra, with only 4 and 5 of 45 problems, respectively. The problems resolved by weaker models are largely a subset of those resolved by stronger models.

We suspect that this gap is partly due to formalization bias rather than to mathematical difficulty alone (see the discussion of formalizability in Section 6). Combinatorial arguments often rely on intuitive reasoning, explicit constructions, or case analyses that are short on paper but laborious to formalize in Lean. Many algebraic and analytic arguments, by contrast, can be built more directly on abstractions that Mathlib already provides.

6 DISCUSSION AND LIMITATIONS

A benchmark of open problems is consumed as it is solved. Once a problem is resolved by a model or a human mathematician, a solution may become public, and the problem loses its value as a held-out test. We will maintain REUF2F as a versioned benchmark. We will periodically recheck the status of each conjecture and replace resolved problems with newly collected and formalized undergraduate-level research problems.

Below we discuss a few design choices and limitations of our benchmark and evaluation that should be kept in mind when interpreting the results.

Model coverage and sampling. Due to limited time and funding, we evaluate a restricted set of models, run each model twice, and do not run every model at its maximum reasoning effort. Two runs provide only a coarse estimate of each model's resolution rate, and our pass@k estimates are noisy for the small number of resolved problems. As time and resources permit, we will add more models, more runs per problem, and additional inference settings.

Common harness. Leanimum-agent separates model capability from differences in agent scaffolding, but no single harness is optimal for every model. Some models may perform better in the native agents they were trained with, or in heavier systems that use parallel sampling, retrieval, or multiple agents. Our results should therefore be read as a comparison under a common, lightweight coding-agent harness, rather than as an upper bound on what each model can achieve.

Formalization faithfulness. Despite expert formalization, independent AI-assisted checks, and a final review by Mathlib contributors, some Lean statements may still differ subtly from the intended conjectures. Such errors can make a problem either easier or harder. However, we argue that they still provide meaningful evaluation signals. A slightly misformalized statement is still a precise and nontrivial mathematical question, and resolving it still requires a correct proof or refutation that passes our verification. Moreover, discrepancies that make a problem easier are most likely to surface exactly when a model resolves it. At that point, we can compare the accepted proof with the source conjecture and correct the formalization in a later version of the benchmark.

Bias from formalizability. We include only conjectures that can be stated faithfully in Lean with Mathlib. Conjectures whose statements require substantial missing infrastructure are therefore underrepresented, even though we try to balance the benchmark across five mathematical areas. As a result, the distribution of REUF2F follows the current coverage of Mathlib as well as the distribution of REU research. Formalization in Lean can also shift the relative difficulty of problems across areas. An agent may struggle on a problem not because its mathematics is hard, but because its solution is hard to formalize. This is especially common for arguments based on intuitive reasoning. Resolution rates across areas therefore reflect both mathematical difficulty and formalization difficulty.

Leakage risk. Our argument that REU conjectures carry a lower leakage risk is qualitative. The source papers themselves are public and may appear in training data, and a conjecture may be resolved and published by humans at any time.

6.1 CONCLUSION

We introduced REUF2F, a benchmark of 244 open conjectures drawn from undergraduate mathematics research. The problems span multiple mathematical areas and are supplied with complete Lean definitions and statements. To evaluate modern coding agents, we place each task in a sandboxed container and allow either a proof or a refutation. Submitted patches are checked independently by Comparator with both official Lean kernel and nanoda, while Leanimum-agent provides a common harness for model comparisons. The strongest evaluated model resolves only approximately 20% of the problems under the tested budgets, leaving substantial room for progress. We hope REUF2F will provide a reproducible foundation for evaluating coding agents on open mathematical research and tracking their ability to produce new, formally verified results.

AI DISCLOSURE

In this work, we used generative AI tools to implement methods and to clean and reformat datasets. We have not used generative AI tools for the rest of the required disclosure tasks are not applicable to this work. Additionally, we used generative AI tools to create or modify scientific figures, create or edit software code, source and search for information, format references, and edit the paper to improve readability. We have reviewed all AI-assisted work. AI-assisted dataset checks were used only to flag issues, and every problem was reviewed by a Mathlib contributor before inclusion. AI-assisted code was reviewed and tested by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Mohammed Abouzaid, Andrew J Blumberg, Martin Hairer, Joe Kileel, Tamara G Kolda, Paul D Nelson, Daniel Spielman, Nikhil Srivastava, Rachel Ward, Shmuel Weinberger, et al. First proof. *arXiv preprint arXiv:2602.05192*, 2026a.

Mohammed Abouzaid, Nikhil Srivastava, Rachel Ward, and Lauren Williams. First proof second batch, 2026b. URL https://arxiv.org/abs/2606.18119.

Tom Adamczewski. Oeis open: How many conjectures can language models turn into theorems? *arXiv preprint arXiv:2608.11941*, 2026.

Tom Adamczewski and Thomas F Bloom. Frontiermath erd\h {o} s. *arXiv preprint arXiv:2609.25050*, 2026.

Levent Alpöge. The $(3, 4, \infty)$ modular family of 2-tori, completed at its three special points, is a complex structure on s^6. https://alpo.ge/s6.pdf, 2026.

Anthropic. Formalizing fermat's last theorem. https://www.anthropic.com/research/formalizing-fermats-last-theorem, 2026.

Zhangir Azerbayev, Bartosz Piotrowski, Hailey Schoelkopf, Edward W Ayers, Dragomir Radev, and Jeremy Avigad. Proofnet: Autoformalizing and formally proving undergraduate-level mathematics. *arXiv preprint arXiv:2302.12433*, 2023.

Bolton Bailey, Jia Li, Jiaqi Wang, Junqi Liu, Justin Asher, Pietro Pellegrino, Ruichen Qiu, Wenda Li, Xiao Ma, Yichuan Cao, Yihe Huang, Yunfei Li, Yunzhou Xie, Zekai Zhu, Zihao Zhou, and Weikun He. Kakeya in dimension three. https://github.com/project-numina/kakeya-3d, 2026.

Chris Bailey. nanoda: An external type checker for lean 4. https://github.com/ammkrn/nanoda_lib, 2026.

Boris Bukh, Ting-Wei Chao, and Zeyu Zheng. The oddtown problem modulo a composite number. *arXiv preprint arXiv:2509.00586*, 2025.

Jiangjie Chen, Wenxiang Chen, Jiacheng Du, Jinyi Hu, Zhicheng Jiang, Allan Jie, Xiaoran Jin, Xing Jin, Chenggang Li, Wenlei Shi, et al. Seed-prover 1.5: Mastering undergraduate-level theorem proving via learning from experience. *arXiv preprint arXiv:2512.17260*, 2025a.

Luoxin Chen, Jinming Gu, Liankai Huang, Wenhao Huang, Zhicheng Jiang, Allan Jie, Xiaoran Jin, Xing Jin, Chenggang Li, Kaijing Ma, et al. Seed-prover: Deep and broad reasoning for automated theorem proving. *arXiv preprint arXiv:2507.23726*, 2025b.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems, 2021. *URL https://arxiv. org/abs/2110.14168*, 9, 2021.

Jasper Dekoninck, Nikola Jovanović, Tim Gehrunger, Kári Rögnvaldsson, Ivo Petrov, Chenhao Sun, and Martin Vechev. Beyond benchmarks: Matharena as an evaluation platform for mathematics with llms. *arXiv preprint arXiv:2605.00674*, 2026.

Moritz Firsching, Paul Lezeau, Salvatore Mercuri, Miklós Z Horváth, Yaël Dillies, Calle Sönne, Eric Wieser, Fred Zhang, Thomas Hubert, Pushmeet Kohli, et al. Formal conjectures: An open and evolving benchmark for verified discovery in mathematics. *arXiv preprint arXiv:2605.13171*, 2026.

Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma, Liang Chen, Zhengyang Tang, et al. Omni-math: A universal olympiad level mathematic benchmark for large language models. In *International Conference on Learning Representations*, volume 2025, pp. 100540–100569, 2025.

Elliot Glazer, Ege Erdil, Tamay Besiroglu, Diego Chicharro, Evan Chen, Alex Gunning, Caroline Falkman Olsson, Jean-Stanislas Denain, Anson Ho, Emily de Oliveira Santos, et al. Frontiermath: A benchmark for evaluating advanced mathematical reasoning in ai. *arXiv preprint arXiv:2411.04872*, 2024.

Google DeepMind. Advanced version of gemini with deep think officially achieves gold-medal standard at the international mathematical olympiad. `https://deepmind.google/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-international-mathematical-olympiad/`, 2025.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3828–3850, 2024.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.

Jiewen Hu, Thomas Zhu, and Sean Welleck. minictx: Neural theorem proving with (long-) contexts. In *International Conference on Learning Representations*, volume 2025, pp. 10106–10130, 2025.

Wenqi Huang, Charley Lee, Leonard Tng, and Serena Ge. Deepswe: Measuring frontier coding agents on original, long-horizon engineering tasks. *arXiv preprint arXiv:2607.07946*, 2026.

Thomas Hubert, Rishi Mehta, Laurent Sartran, Miklós Z Horváth, Goran Žužić, Eric Wieser, Aja Huang, Julian Schrittwieser, Yannick Schroecker, Hussain Masoom, et al. Olympiad-level formal mathematical reasoning with reinforcement learning. *Nature*, 651(8106):607–613, 2026.

Albert Q Jiang, Sean Welleck, Jin Peng Zhou, Wenda Li, Jiacheng Liu, Mateja Jamnik, Timothée Lacroix, Yuhuai Wu, and Guillaume Lample. Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. *arXiv preprint arXiv:2210.12283*, 2022.

Jiedong Jiang, Wanyi He, Wang Yuefeng, Guoxiong Gao, Yongle Hu, Jingting Wang, Nailin Guan, Peihao Wu, Bryan Dai, Liang Xiao, et al. Fate: A formal benchmark series for frontier algebra of multiple difficulty levels. In *International Conference on Learning Representations*, volume 2026, pp. 10543–10580, 2026.

Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *International Conference on Learning Representations*, volume 2024, pp. 54107–54157, 2024.

Lean FRO. Comparator: A trustworthy judge for lean proofs. https://github.com/leanprover/comparator, 2026a.

Lean FRO. Lean ai formalization leaderboard. https://lean-lang.org/eval/, 2026b.

Haohan Lin, Zhiqing Sun, Sean Welleck, and Yiming Yang. Lean-star: Learning to interleave thinking and proving. In *International Conference on Learning Representations*, volume 2025, pp. 66041–66062, 2025.

Minh-Thang Luong, Dawsen Hwang, Hoang H Nguyen, Golnaz Ghiasi, Yuri Chervonyi, Insuk Seo, Junsu Kim, Garrett Bingham, Jonathan Lee, Swaroop Mishra, et al. Towards robust mathematical reasoning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 35406–35430, 2025.

Bhavik Mehta. Exponentialramsey: Formalization of an exponential improvement for diagonal ramsey. https://github.com/b-mehta/ExponentialRamsey, 2023.

Mike Merrill, Alexander Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Shin, Thomas Walshe, E Kelly Buchanan, et al. Terminal-bench: Benchmarking agents on hard, realistic tasks in command line interfaces. In *International Conference on Learning Representations*, volume 2026, pp. 40903–40986, 2026.

Leonardo de Moura. Postmortem for kernel soundness bug #14576. https://leodemoura.github.io/blog/2026-8-1-postmortem-for-kernel-soundness-bug-14576/, 2026.

Leonardo de Moura and Sebastian Ullrich. The lean 4 theorem prover and programming language. In *International Conference on Automated Deduction*, pp. 625–635. Springer, 2021.

National Science Foundation. Research experiences for undergraduates (reu). https://www.nsf.gov/funding/initiatives/reu, 2026.

OpenAI. Cdc lean formalization. https://github.com/openai/cdc-lean, 2026a.

OpenAI. Finite time blowup for navier–stokes. https://cdn.openai.com/pdf/32d9f210-8b73-45e0-91bc-82a30aef8a9a/navier-stokes.pdf, 2026b.

OpenAI. Navierstokesandeuler: Lean formalizations of finite time blowup for navier–stokes and euler equations. https://github.com/openai/NavierStokesAndEuler, 2026c.

OpenAI. Planar point sets with many unit distances. https://cdn.openai.com/pdf/74c24085-19b0-4534-9c90-465b8e29ad73/unit-distance-proof.pdf, 2026d.

OpenAI. Ten advances in mathematics and theoretical computer science. https://cdn.openai.com/pdf/ten-proofs-oai.pdf, 2026e.

Sang-il Oum. A proof of the cycle double cover conjecture by openai: An exposition. *arXiv preprint arXiv:2607.16356*, 2026.

Nikil Ravi, Kexing Ying, Vasilii Nesterov, Rayan Krishnan, Elif Uskuplu, Bingyu Xia, Janitha Aswedige, and Langston Nashold. Formalproofbench: Can models write graduate level math proofs that are formally verified? *arXiv preprint arXiv:2603.26996*, 2026.

Johannes Schmitt, Gergely Bérczi, Jasper Dekoninck, Jeremy Feusi, Tim Gehrunger, Raphael Appenzeller, Pieter Belmans, Alessio Bottini, Jim Bryan, João Camarneiro, et al. Improofbench: Benchmarking ai on research-level mathematical proof generation. *arXiv preprint arXiv:2509.26076*, 2025.

Daniel Shepard and Robin Salimans. Automationbench. *arXiv preprint arXiv:2604.18934*, 2026.

Guijin Son, Seungone Kim, Catherine Arnett, Hyunwoo Ko, Hyein Lee, Hyeonah Kang, Jiang Longxi, Jin Yun, JungYup Lee, Kyungmin Lee, et al. Soohak: A mathematician-curated benchmark for evaluating research-level math capabilities of llms. *arXiv preprint arXiv:2605.09063*, 2026.

The mathlib Community. Mathlib4: The lean 4 mathematical library. https://github.com/leanprover-community/mathlib4, 2026.

The PFR Formalization Contributors. The polynomial freiman–ruzsa conjecture: Lean 4 formalization. https://github.com/teorth/pfr, 2026.

The Sphere Packing in Lean Contributors. Sphere packing in lean. https://github.com/thefundamentaltheor3m/Sphere-Packing-Lean, 2024.

George Tsoukalas, Jasper Lee, John Jennings, Jimmy Xin, Michelle Ding, Michael Jennings, Amitayush Thakur, and Swarat Chaudhuri. Putnambench: Evaluating neural theorem-provers on the putnam mathematical competition. *Advances in Neural Information Processing Systems*, 37: 11545–11569, 2024.

George Tsoukalas, Anton Kovsharov, Sergey Shirobokov, Anja Surina, Moritz Firsching, Gergely Bérczi, Francisco JR Ruiz, Arun Suggala, Adam Zsolt Wagner, Eric Wieser, et al. Advancing mathematics research with ai-driven formal proof search. *arXiv preprint arXiv:2605.22763*, 2026.

Ran Xin, Zeyu Zheng, Yanchen Nie, Kun Yuan, and Xia Xiao. Scaling up multi-turn off-policy rl and multi-agent tree search for llm step-provers. *arXiv preprint arXiv:2509.06493*, 2025.

John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik R Narasimhan, and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://arxiv.org/abs/2405.15793.

John Yang, Kilian Lieret, Jeffrey Ma, Parth Thakkar, Dmitrii Pedchenko, Sten Sootla, Emily McMilin, Pengcheng Yin, Rui Hou, Gabriel Synnaeve, et al. Programbench: Can language models rebuild programs from scratch? *arXiv preprint arXiv:2605.03546*, 2026.

Yu Yue, Yufeng Yuan, Qiying Yu, Xiaochen Zuo, Ruofei Zhu, Wenyuan Xu, Jiaze Chen, Chengyi Wang, TianTian Fan, Zhengyin Du, et al. Vapo: Efficient and reliable reinforcement learning for advanced reasoning tasks. *arXiv preprint arXiv:2504.05118*, 2025.

Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. Minif2f: a cross-system benchmark for formal olympiad-level mathematics. *arXiv preprint arXiv:2109.00110*, 2021.

Zeyu Zheng, Shengtong Zhang, Jeremy Avigad, Prasad Tetali, and Sean Welleck. The problem is the problem: Towards scalable mathematical discovery. *arXiv preprint arXiv:2608.16977*, 2026.

A TWO-WAY TASK EXPORT

`Export.lean` checks that each target declaration is a theorem and exports its name and universe parameters. Task preparation preserves the original declaration and its definitions, then appends a `_neg` theorem with type `¬ type_of% @target` (including explicit universe parameters when needed). Lean elaborates `type_of% @target` to the target's complete type, so the negation includes all binders without printing and reparsing the statement.

Listing 1: Lean code for exporting target declaration metadata.

```
import Lean

open Lean Elab Command

namespace Export

/-- Export each target's universe parameters, named `u0, u1, ...`, so
    the task can
    state its refutation as `Not type_of% @target.{u0, ...}`. -/
elab "#reuf2f_export_types " manifestPath:str " to " outputPath:str :
    command => do
  let text <- liftIO <| IO.FS.readFile manifestPath.getString
  let .ok manifest := Json.parse text | throwError "Invalid task
    manifest"
  let .ok entries := manifest.getArr? | throwError "Task manifest must
    be an array"
  let mut results : Array Json := #[]
  for entry in entries do
    let .ok name := entry.getObjValAs? String "declaration"
      | throwError "Task has no declaration"
    let info <- getConstInfo name.toName
    unless info.isTheorem do throwError "Not a theorem: {name}"
    let levels := info.levelParams.zipIdx.map fun (_, i) => s!"u{i}"
    results := results.push <| Json.mkObj
      [("declaration", toJson name), ("universes", toJson levels)]
  liftIO <| IO.FS.writeFile outputPath.getString (Json.arr results).
    pretty

end Export
```

B LEANIMUM-AGENT PROMPT

All models share the system prompt "You are a helpful assistant that can interact with a computer shell to solve Lean tasks." and the instance template below. The template specifies the two-way task, gives instructions for searching Mathlib and checking Lean code, and defines the patch-submission procedure. The `{{task}}` variable is replaced with the task's problem statement, "In `<file>.lean`, prove either `<declaration>` or refute it by proving its logical negation `<declaration>_neg`."

Listing 2: The Leanimum-agent task prompt.

```
<problem_description>
Consider the following Lean problem:
{{task}}
</problem_description>

<instructions>
# Task Instructions

## Overview

You're a mathematician interacting continuously with a computer by
    submitting commands.
```

```
You'll be helping prove or refute the Lean statement described above.
Your task is specifically to make changes to files in the current
    directory in order to formally prove or refute the conjecture
    described in the problem description, keeping the target
    declarations and relevant dependencies unchanged.
<IMPORTANT>This is an interactive process where you will think and issue
     AT LEAST ONE command, see the result, then think and issue your
    next command(s).</important>

For each response:

1. Include a THOUGHT section explaining your reasoning and what you're
    trying to accomplish
2. Provide one or more bash tool calls to execute

## Important Boundaries

- MODIFY: Regular Lean files in /testbed (this is the working directory
    for all your subsequent commands)
- DO NOT MODIFY: Target statements and their dependencies, the `lean-
    toolchain` version; do not introduce `lakefile.lean`

## Recommended Workflow

1. Analyze the codebase by finding and reading relevant files
2. Search the project and Mathlib for relevant definitions and lemmas
3. Edit the task file to prove the original theorem or its full logical
    negation (`_neg`)
4. Check your changes with `lake build`
5. Check that your proof does not depend on introduced axioms or sorries

## Command Execution Rules

You are operating in an environment where

1. You issue at least one command
2. The system executes the command(s) in a subshell
3. You see the result(s)
4. You write your next command(s)

Each response should include:

1. **Reasoning text** where you explain your analysis and plan
2. At least one tool call with your command

**CRITICAL REQUIREMENTS:**

- Your response SHOULD include reasoning text explaining what you're
    doing
- Your response MUST include AT LEAST ONE bash tool call. You can make
    MULTIPLE tool calls in a single response when the commands are
    independent (e.g., searching multiple files, reading different parts
     of the codebase).
- Directory or environment variable changes are not persistent. Every
    action is executed in a new subshell.
- However, you can prefix any action with `MY_ENV_VAR=MY_VALUE cd /path/
    to/working/dir && ...` or write/load environment variables from
    files

Example of a CORRECT response:
<example_response>
I need to understand the target statement. Let me find relevant
    definitions and check the project structure.
```

```
[Makes multiple bash tool calls: {"command": "ls -la"}, {"command": "rg
    -n 'theorem|def' -g '*.lean'"}, {"command": "cat lean-toolchain"}]
</example_response>

## Environment Details

- You have a full Linux shell environment
- Always use non-interactive flags (-y, -f) for commands
- Avoid interactive tools like vi, nano, or any that require user input
- You can use bash commands or invoke any tool that is available in the
    environment
- You can also create new tools or scripts to help you with the task
- If a tool isn't available, you can also install it
- Search Mathlib with `rg 'PATTERN' .lake/packages/mathlib/Mathlib`.
- Register new helper libraries with `[[lean_lib]]` entries in `lakefile
    .toml`, then import the needed modules in the task file.

## Submission

When you've completed your work, you MUST submit your changes as a git
    patch.
Follow these steps IN ORDER, with SEPARATE commands:

Step 1: Create the patch file
Run `git add path/to/file1 path/to/file2 && git diff --cached > patch.
    txt` listing exactly the source files you modified.
Do NOT commit your changes.

<IMPORTANT>
The patch must only contain changes to the specific source files you
    modified to prove or refute the theorem.
Do not use any of the following to complete a proof:

- `sorry` or `admit` in the chosen proof or any of its dependencies
- new axioms or changes to the target's definitions
- changes to the verifier or its configuration (lakefile.toml library
    layout may change)
- binary or compiled files
</IMPORTANT>

Step 2: Verify your patch
Inspect patch.txt to confirm it only contains your intended changes and
    headers show `--- a/` and `+++ b/` paths.

Step 3: Submit (EXACT command required)
You MUST use this EXACT command to submit:

```bash
echo COMPLETE_TASK_AND_SUBMIT_FINAL_OUTPUT && cat patch.txt
```

If the command fails (nonzero exit status), it will not submit.

<CRITICAL>
- Creating/viewing the patch and submitting it MUST be separate commands
     (not combined with &&).
- If you modify patch.txt after verifying, you SHOULD verify again
    before submitting.
- You CANNOT continue working (reading, editing, testing) in any way on
    this task after submitting.
</CRITICAL>
</instructions>
```

C VERIFICATION CONFIGURATION

Comparator compares the trusted challenge with the submitted solution. The configuration below checks the original theorem; a second invocation checks the theorem with the suffix `_neg`. The example shows the task `achenjang2019gaps`. For other tasks, the solution module and theorem name are replaced with the task's module and target declaration.

```
{
  "challenge_module": "Challenge",
  "solution_module": "achenjang2019gaps",
  "theorem_names": ["achenjang2019gaps.conjecture_7_2"],
  "permitted_axioms": [
    "propext", "Quot.sound", "Classical.choice"
  ],
  "external_kernels": {
    "nanoda": ["/usr/local/bin/nanoda_lib/target/release/nanoda_bin"]
  }
}
```

The `external_kernels` entry registers nanoda as an additional checker, so the proof must be accepted by both the official Lean kernel and nanoda. The evaluator uses a clean container with the same pinned Lean and Mathlib dependencies as the task environment.