

S2N-BIGNUM-BENCH: A PRACTICAL BENCHMARK FOR EVALUATING LOW-LEVEL CODE REASONING OF LLMs

Balaji Rao *

Stevens Institute of Technology
Hoboken, NJ 07030
brao@stevens.edu

John Harrison †

Amazon Web Services
Seattle, WA 98101
jargh@amazon.com

Soonho Kong †

Amazon Web Services
Seattle, WA 98101
soonho@amazon.com

Juneyoung Lee †

Amazon Web Services
Seattle, WA 98101
lebjuney@amazon.com

Carlo Lipizzi

Stevens Institute of Technology
Hoboken, NJ 07030
clipizzi@stevens.edu

ABSTRACT

Neurosymbolic approaches leveraging Large Language Models (LLMs) with formal methods have recently achieved strong results on mathematics-oriented theorem-proving benchmarks. However, success on competition-style mathematics does not by itself demonstrate the ability to construct proofs about real-world implementations. We address this gap with a benchmark derived from an industrial cryptographic library whose assembly routines are already verified in HOL Light. *s2n-bignum* is a library used in AWS for providing fast assembly routines for cryptography and its correctness is established by formal verification. The task of formally verifying this library has been a significant achievement for the Automated Reasoning Group. It involved two tasks: (1) precisely specifying the correct behavior of a program as a mathematical proposition, and (2) proving that the proposition is correct. In the case of *s2n-bignum*, both tasks were carried out by human experts. In *s2n-bignum-bench*, we provide the formal specification and ask the LLM to generate a proof script that is accepted by HOL Light within a fixed proof-check timeout. To our knowledge, *s2n-bignum-bench* is the first public benchmark focused on machine-checkable proof synthesis for industrial low-level cryptographic assembly routines in HOL Light. We believe this benchmark provides a challenging and practically relevant testbed for evaluating LLM-based theorem proving beyond competition mathematics. The code to set-up and use the benchmark are available here: [s2n-bignum-bench](#).

1 INTRODUCTION

Formal theorem proving with Large Language Models (LLMs) and interactive theorem provers has become a central testbed for LLM reasoning, but existing benchmarks emphasize competition-style mathematical problems. Solving complex math problems requires a rigorous framework of steps and logical proofs, and success on such tasks evidences structured reasoning (1). However, excellence on math-centric benchmarks does not automatically transfer to systems with practical engineering consequences. Therefore, the design of diverse and high-quality benchmarks is a key challenge in this research area.

To complement existing benchmarks, we propose *s2n-bignum-bench*, a machine-checkable benchmark distilled from the *s2n-bignum* cryptographic library, focusing on verified low-level code. The benchmark tests whether LLMs can synthesize machine-checkable proofs about real low-level implementations rather than only competition-style mathematics.

*Extended work done after completing internship at AWS; Corresponding author

†Work independent from role at AWS

Our contributions are fourfold. First, we package **2,284**¹ proof obligations from the production-grade *s2n-bignum* cryptography library as isolated *context-query* tasks with stable per-problem identifiers and standalone artifacts. Second, we provide an end-to-end pipeline to build the benchmark, retrieve selected problems, and run fully offline evaluation using the shipped artifacts. Third, we include integrity mechanisms that detect unsound or invalid submissions, including checks for newly introduced axioms, forbidden placeholders such as `CHEAT_TAC`, and parser-level validation of submitted proof expressions. We also provide a contamination-mitigation mechanism based on type-annotation obfuscation. Fourth, because the benchmark is grounded in a deployed cryptographic codebase, it measures ISA-aware, bit-precise reasoning that is closer to real verification workflows than competition mathematics.

2 BACKGROUND AND RELATED WORK

2.1 THEOREM PROVING

Interactive theorem provers (ITPs) following the LCF approach all have their derivations ultimately checked by a small, trusted kernel that produces values of type `thm` (the “theorem” type) (2). Examples of ITPs include Isabelle/HOL, Lean4 and HOL Light. HOL Light is a minimalist proof system designed for higher-order logic (HOL) implemented in OCaml, with a very small trusted kernel and an emphasis on clarity (3). Its proof script is an OCaml program and the proof system uses the OCaml toplevel (REPL) for interactivity. It relies on tactics to discharge goals.

LLMs have become a central part of *neural theorem proving* (NTP), where the model proposes proof steps while an interactive theorem provers (ITP) act as the verifier. Unlike answer-only benchmarks (e.g., `gsm8k` (4), `CruxEval` (5)), NTPs demand *structured, verifiable reasoning*.

2.2 THEOREM PROVING BENCHMARKS

MiniF2F is the most widely adopted cross-system benchmark, with 488 formalized Olympiad-level mathematics problems translated across multiple proof systems including Lean4, Metamath, Isabelle, and HOL Light (6), saturated in 2025 by Seed Prover (7). PutnamBench introduces competition mathematics from the William Lowell Putnam Mathematical Competition, featuring 1,724 hand-constructed formalizations of 672 theorems across Lean4, Isabelle, and Rocq (8), 99.4% solve rate by Aleph (9).

Recent benchmarks have expanded toward verification conditions and repository-scale software verification. In particular, *NTP4VC* studies theorem proving over verification conditions extracted from real systems code, while *VeriSoftBench* evaluates proof synthesis over repository-scale Lean verification tasks(10; 11). `miniCTX` and `VeriBench-FTP` are designed to test the use of context as well as theorem-level, context-level, and project-level generalization across several mathematical, as well as code domains (12; 13). These works represent a paradigm shift toward realistic theorem proving scenarios where models must leverage extensive context from real Lean projects. Other relevant works that have moved toward verification-oriented proving and code-centered formal reasoning, include `miniCodeProps`, `CLEVER`, and `VERINA` (14; 15; 16). Works like `SorryDB` emphasize the need for dynamically-updating benchmarks (17). Our work is complementary to these efforts: we focus specifically on HOL Light proof synthesis for industrial low-level cryptographic assembly with shipped object-code artifacts and trusted ISA semantics.

3 MOTIVATION

A large fraction of neural theorem-proving evaluation has centered on contest or textbook mathematics, with benchmarks such as MiniF2F, and PutnamBench probing abstract reasoning in a formal setting. At the same time, recent work has begun to extend evaluation toward software verification, verification-condition proving, and repository-scale proof synthesis. We aim to supplement these works by proposing an underrepresented ecosystem: mechanized proofs about industrial low-level cryptographic assembly in HOL Light.

¹As of this submission.

Additionally, while some benchmarks have been effective at evaluating reasoning models, they do not test whether models can construct machine-checkable proofs about real implementations. Math-centric proving does not generally involve architectural state, aliasing, or endianness, and competence in abstract mathematics does not, by itself, establish capability for low-level code reasoning.

The correctness of cryptographic libraries and systems code has immediate security and reliability consequences, yet low-level implementation reasoning remains underrepresented in public theorem-proving benchmarks. The *s2n-bignum* library contains hand-tuned big-integer assembly subroutines (x86/ARM) accompanied by HOL Light proofs that the object code meets functional correctness specifications under a trusted ISA model. Building a benchmark from this corpus lets us evaluate an NTP system’s ability to construct a proof that real assembly satisfies its specification.

Benchmark	Formal system	Primary focus	Problems (#)	Task setting
miniF2F (6)	Multiple	Olympiad mathematics (formal)	488	Formal theorem proving
PutnamBench (8)	Multiple	Competition mathematics (formal)	672	Formal theorem proving
NTP4VC (10)	Multiple	Verification conditions from real code	600	Real-world VC proving
VeriBench-FTP (13)	Lean4	Code-verification artifacts	857	Proofs from verification artifacts
miniCTX/miniCTX-v2 (18)	Lean4	Context-dependent proving	762	Context / Project generalization
VeriSoftBench (11)	Lean4	Repository-scale software verification	500	Repository-scale verification proving
s2n-bignum-bench	HOL Light	Verified cryptographic assembly programs	2284	Proof synthesis over machine-code

Table 1: *s2n-bignum-bench* relative to representative theorem-proving and verification benchmarks

4 S2N-BIGNUM-BENCH CONSTRUCTION

Our problems are derived from the open-source *s2n-bignum* repository, an AWS cryptographic library. Each problem in *s2n-bignum-bench* is a HOL Light *context-query* task. We inline the relevant OCaml modules, locate top-level theorem bindings of the form `let THM = prove(goal, proof)`, and extract the goal as the *query*. The accompanying *context* is a self-contained OCaml/HOL Light setup that loads the required definitions, constants, and previously proved results needed to reproduce the original proving environment, while replacing each original proof body with the placeholder `CHEAT_TAC`.² With this, we isolate the task of synthesizing a new, machine-checkable proof under the same interfaces and imports as the source project.

To distinguish between different problems, we introduce the notion of a `Problem identifier`, since the same theorem name may appear across different proof files or multiple times within the same file. A problem identifier has the form `arch.filename.thm.N`, for example: `arm.bignum.montsq_r256.lemma1.0`, where *N* represents an occurrence index of a lemma. In this work, we also include the artifacts to extract selected problems into a directory with (i) `setup.ml` and (ii) `query.txt`. This will allow for reproducible, standalone attempts per problem.

Using lightweight heuristics over theorem names and goal forms, we partition the benchmark into four categories: `Bit-vector lemmas` (**311**), `Program-state lemmas` (**552**), `Functional correctness` (**859**, comprising **437** ARM and **422** x86 problems), and `Generic` (**562**) for auxiliary facts not captured by the preceding categories.

HOL Light proofs are typically developed interactively through the OCaml REPL. Existing tooling, such as `hol_server` and the VSCode extension for HOL Light, can be used to provide an interactive development environment on top of the released benchmark artifacts (19).

5 EVALUATION

5.1 ANSWER SUBMISSION AND GRADING

Challengers submit a proof expression along with the name of the problem attempted. We first perform a syntax and type pre-check by compiling a generated `.synchk.ml` file that pastes the submitted proof expression into the benchmark context. This catches malformed tactic expressions and other immediate parser or type errors before full evaluation.

²`CHEAT_TAC` is a placeholder tactic in HOL Light, analogous to *sorry* in Lean or Isabelle. Any submission that uses it is rejected as cheating.

Each submitted proof attempt yields exactly one verdict per problem: OK, FAIL, CHEATING, TIMEOUT, or ERROR. Results are aggregated into a CSV file, and the primary task metric is binary success at kernel-checked proof completion. To make model comparisons meaningful, an *official* evaluation configuration should fix the proof-check timeout, hardware setting, and submission budget. We also provide support for user-configurable timeouts for exploratory use³. Auxiliary lemmas may be defined inside the submitted proof expression, provided that the overall expression evaluates to a tactic.

The model is not given access to the original proof bodies or the tactics used to prove other theorems. This restriction is important for preserving the validity of the benchmark. It is important to note, however, that the challenger can provide the relevant machine-code context needed to understand the specification being proved to the LLM.

5.2 INITIAL BASELINE EXPERIMENTS

As a preliminary baseline, we evaluate GPT-5.3-Codex (20) through `codex-cli` under the configuration described in Appendix A. The model achieves a binary proof-completion rate of **4.4%** in medium-effort mode and **5.3%** in high-effort mode over the full benchmark. We treat this as an initial baseline rather than an exhaustive estimate of current model capability³.

5.3 INTEGRITY AND CONTAMINATION DEFENSES

To mitigate contamination from memorized theorem statements, we implement an obfuscation mechanism that makes type annotations more explicit. We set the `print_types_of_subterms` of HOL Light to the most verbose mode and reprint the queries. However, this works for only $\approx 70\%$ of the problem set because HOL Light’s printer and parser are not fully Pollack-consistent (21). For such queries, we chose to use their original representations without obfuscation⁴.

To ensure that a challenger did not introduce any forbidden tactics like `CHEAT_TAC` or functions like `new_axiom`, our evaluation checks the output of the `axioms()` function in HOL Light. The function returns a list of theorems that have been axiomatized so far. If the answer from challenger did not use any forbidden functions, the result of `axioms()` must be identical before and after the solution. If they are different, our benchmark script marks the result as “CHEATING”.

A separate class of attacks attempts to introduce a syntactically complex solution that resembles “SQL injection”. To prevent this, we invoke an OCaml parser for each solution and check whether it has one valid expression. Submissions that fail this check are rejected before proof evaluation.

6 CONCLUSION

We introduce *s2n-bignum-bench*, a benchmark for machine-checkable proof synthesis over a deployed corpus of verified low-level cryptographic assembly proofs in HOL Light. This benchmark targets a capability that remains underrepresented in current evaluation: constructing sound proofs about real low-level implementations under trusted ISA semantics. By releasing isolated problem artifacts, an offline evaluation harness, and integrity checks against unsound submissions, we aim to provide a reproducible testbed for future work on theorem proving and verification-oriented reasoning beyond competition mathematics. Although the current release focuses on functional correctness, recent HOL Light developments around *s2n-bignum* also formalize relational properties, including constant-time discipline and equivalence between optimized and verification-friendly routines, suggesting a natural path toward future benchmark extensions beyond extensional correctness (22; 23).

³Detailed explanation in Appendix

⁴We are communicating with HOL Light to fix this

REFERENCES

- [1] T. Achim and V. Tenev, “Harmonic: Building mathematical superintelligence,” <https://harmonic.fun/>, accessed: 2025-09-18.
- [2] J. Harrison, J. Urban, and F. Wiedijk, “History of interactive theorem proving,” in *Handbook of the History of Logic*. Elsevier, 2014, vol. 9, pp. 135–214.
- [3] J. Harrison, “Hol light: An overview,” in *International Conference on Theorem Proving in Higher Order Logics*. Springer, 2009, pp. 60–66.
- [4] K. Cobbe, V. Kosaraju, M. Bavarian, M. Chen, H. Jun, L. Kaiser, M. Plappert, J. Tworek, J. Hilton, R. Nakano *et al.*, “Training verifiers to solve math word problems,” *arXiv preprint arXiv:2110.14168*, 2021.
- [5] A. Gu, B. Rozière, H. Leather, A. Solar-Lezama, G. Synnaeve, and S. I. Wang, “Crux-eval: A benchmark for code reasoning, understanding and execution,” *arXiv preprint arXiv:2401.03065*, 2024.
- [6] K. Zheng, J. M. Han, and S. Polu, “Minif2f: a cross-system benchmark for formal olympiad-level mathematics,” 2022. [Online]. Available: <https://arxiv.org/abs/2109.00110>
- [7] L. Chen, J. Gu, L. Huang, W. Huang, Z. Jiang, A. Jie, X. Jin, X. Jin, C. Li, K. Ma, C. Ren, J. Shen, W. Shi, T. Sun, H. Sun, J. Wang, S. Wang, Z. Wang, C. Wei, S. Wei, Y. Wu, Y. Wu, Y. Xia, H. Xin, F. Yang, H. Ying, H. Yuan, Z. Yuan, T. Zhan, C. Zhang, Y. Zhang, G. Zhang, T. Zhao, J. Zhao, Y. Zhou, and T. H. Zhu, “Seed-prover: Deep and broad reasoning for automated theorem proving,” 2025. [Online]. Available: <https://arxiv.org/abs/2507.23726>
- [8] G. Tsoukalas, J. Lee, J. Jennings, J. Xin, M. Ding, M. Jennings, A. Thakur, and S. Chaudhuri, “Putnambench: Evaluating neural theorem-provers on the putnam mathematical competition,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.11214>
- [9] B. H. Vlad Isenbaev, “Logical intelligence’s aleph solves putnambench.” [Online]. Available: <https://logicalintelligence.com/blog/aleph-solves-putnambench>
- [10] Q. Xu, X. Luan, R. Wang, J. O. J. Leang, P. Wang, H. Li, W. Li, and C. Watt, “Neural theorem proving for verification conditions: A real-world benchmark,” in *The Fourteenth International Conference on Learning Representations*, 2026. [Online]. Available: <https://openreview.net/forum?id=MfDyickxQA>
- [11] Y. Xin, Q. Chen, G. Durrett, and I. Dillig, “Verisoftbench: Repository-scale formal verification benchmarks for lean,” *arXiv preprint arXiv:2602.18307*, 2026.
- [12] J. Hu, T. Zhu, and S. Welleck, “minictx: Neural theorem proving with (long-) contexts,” *arXiv preprint arXiv:2408.03350*, 2024.
- [13] S. Barkallah, S. Daruru, B. Miranda, L. Aniva, A. Nie, and S. Koyejo, “Veribench-ftp: A formal theorem proving benchmark in lean 4 for code verification,” in *The 5th Workshop on Mathematical Reasoning and AI at NeurIPS 2025*.
- [14] E. Lohn and S. Welleck, “minicodeprops: a minimal benchmark for proving code properties,” *arXiv preprint arXiv:2406.11915*, 2024.
- [15] A. Thakur, J. Lee, G. Tsoukalas, M. Sistla, M. Zhao, S. Zetsche, G. Durrett, Y. Yue, and S. Chaudhuri, “Clever: A curated benchmark for formally verified code generation,” *arXiv preprint arXiv:2505.13938*, 2025.
- [16] Z. Ye, Z. Yan, J. He, T. Kasriel, K. Yang, and D. Song, “Verina: Benchmarking verifiable code generation,” *arXiv preprint arXiv:2505.23135*, 2025.
- [17] A. Letson, L. Sarra, A. Poiroux, O. Dressler, P. Lezeau, D. Aranha, F. Pu, A. Hill, M. C. Hidalgo, J. Berman *et al.*, “Sorrydb: Can ai provers complete real-world lean theorems?” *arXiv preprint arXiv:2603.02668*, 2026.

- [18] J. Hu, T. Zhu, and S. Welleck, “minictx: Neural theorem proving with (long-)contexts,” 2025. [Online]. Available: <https://arxiv.org/abs/2408.03350>
- [19] monadius, “Hol light extension for vs code.” [Online]. Available: <https://github.com/monadius/vscode-hol-light>
- [20] OpenAI, “Gpt-5.3-codex system card.” [Online]. Available: <https://deploymentsafety.openai.com/gpt-5-3-codex/gpt-5-3-codex.pdf>
- [21] F. Wiedijk, “Pollack-inconsistency,” *Electronic Notes in Theoretical Computer Science*, vol. 285, pp. 85–100, 2012.
- [22] D. Mazzucato, A. Mohamed, J. Lee, C. Barrett, J. Grundy, J. Harrison, and C. S. Păsăreanu, “Relational hoare logic for realistically modelled machine code,” in *International Conference on Computer Aided Verification*. Springer, 2025, pp. 389–413.
- [23] J. Harrison, “Soundness of s2n-bignum formal verification.” [Online]. Available: https://github.com/awslabs/s2n-bignum/blob/main/doc/s2n_bignum_soundness.md

A A GUIDE FOR CHALLENGERS

The official repository walks a user through setting up the benchmark, and the experimental protocol overview described here is a good place to start to begin exploring this problem set.

A.1 EXPERIMENTAL PROTOCOL OVERVIEW

We conducted preliminary *pass@1* experiments to evaluate whether current language models can synthesize HOL Light tactic proofs for *s2n-bignum-bench*. The experiment(s) follow the benchmark workflow described in the repository documentation: *s2n-bignum-bench*.

Our primary reported baseline uses a closed-source reasoning model accessed through the Codex CLI. The zero-shot prompt template and evaluation scripts are available in official repository.

A.2 BENCHMARK PREPARATION AND PROBLEM RETRIEVAL

Following the repository workflow, we first build the benchmark and extract theorem metadata from the pinned *HOL Light* and *s2n-bignum* sources. This produces a corpus of **2,284** problems. Each problem consists of:

- a HOL Light boolean term stored in `query.txt`, and
- a corresponding `setup.ml` file containing the HOL Light session preamble needed to establish the proof context.

For inference, problems were (and can be) retrieved in two equivalent formats:

- as a flat CSV using `retrieve-problem.py --csv-only`, and
- as a directory tree of `<problem-id>/query.txt` files under a `problems` directory.

A.3 PROMPTING AND INFERENCE PIPELINES

Prompt template. All runs use the same zero-shot prompt template, using the repository prompt:

```
You are an expert in HOL Light. I am going to give a HOL
Light boolean term. Please write a HOL Light proof of it
in a THEN form. Do not use CHEAT_TAC or new_axiom.

<Example>

If the input is

`x * (y + z) = x * z + x * y`

A possible answer is

REWRITE_TAC[LEFT_ADD_DISTRIB] THEN
GEN_REWRITE_TAC LAND_CONV [ADD_SYM] THEN
REFL_TAC

Please include the proof in your proof but not other natural
statements, so that I can easily evaluate your answer.

<Query>

{query.txt}
```

The prompt includes a single worked example and instructs the model to output only a tactic expression, with no surrounding explanation. This matches the expected input format of the evaluation pipeline.

At this point, we have built the benchmark, and we have all the required components to evaluate the answers that are generated by the LLM.

A.4 CODEX CLI BASELINE.

Our main preliminary baseline uses GPT-5.3-Codex through `codex cli`. For each problem, we substitute the goal term into the prompt and invoke the model in a restricted read-only sandbox with shell access, databases, and web search disabled. Responses are written to a CSV together with problem identifiers and auxiliary metadata. We also parse the CLI event stream to detect unexpected tool-related events; these are logged for audit purposes but are not used for evaluation.

A.5 CONSTRUCTION OF THE PROBLEM TIMEOUT MAPPING

The benchmark evaluator executes each submitted proof attempt with a timeout in order to prevent runaway tactics from consuming unbounded compute. We support two timeout mechanisms: category-level defaults in `timeouts.json` and per-problem overrides in `timeout-map.json`. In practice, proof-check times in *s2n-bignum-bench* vary by several orders of magnitude, from a few seconds or milliseconds for small HOL Light lemmas to multiple hours for the largest routine-correctness theorems.

A single blanket timeout therefore creates an undesirable tradeoff: if set too low, it incorrectly penalizes legitimate but expensive proofs; if set too high, failed submissions may waste hours before timing out, this might also lead to OOM issues which could lead to other problems.

To address this, we supply a heuristic supported suggested per-problem timeout map derived from repeated profiling of the ground-truth proofs.

To construct this map, we profile the original human-written proofs using the same benchmark evaluation harness that is later used for submitted answers. During profiling, category-level timeouts are temporarily set to a very large value so that the ground-truth proofs are not prematurely terminated. The evaluator then assembles, compiles, and executes the benchmark proof files exactly as in normal assessment, while recording per-problem wall-clock time for the internal `prove(goal, tactic)` call. We refer to this measurement as `prove_secs`; it excludes compilation overhead and captures only proof execution time inside HOL Light.

Because proof execution times vary across runs due to operating-system scheduling, garbage collection, memory pressure, and parallel execution effects, we repeat this profiling procedure three times over the full benchmark. Across all runs, every ground-truth proof completes successfully. The repeated runs reveal substantial variance for some heavy proofs, including large absolute spreads for the most expensive routine-correctness theorems.

For each problem, we collect the profiled `prove_secs` values from all profiling runs and compute summary statistics including the maximum runtime and an empirical high-percentile runtime. We then divide problems into two broad classes:

- **Light problems:** proofs whose profiled runtimes remain comfortably below the heavy-proof regime.
- **Heavy problems:** proofs whose observed runtimes indicate substantially higher cost or variance.

Light and heavy problems are assigned timeouts using different multiplicative safety margins. Intuitively, heavy problems receive more generous scaling because they exhibit larger absolute runtime variance. In both cases, the assigned timeout is bounded below by a fixed minimum floor (120 seconds) and above by a global cap (10,800 seconds). This yields a timeout map that is conservative enough to accommodate runtime variation while still avoiding the extreme inefficiency of a blanket timeout.

The resulting `timeout-map.json` contains one timeout entry per benchmark problem. During evaluation, the benchmark first checks whether a submitted problem identifier appears in this map. If so, the corresponding per-problem timeout is used. If not, the evaluator falls back to the category-level default from `timeouts.json`. This fallback is primarily intended for exploratory use or for newly added benchmark problems that have not yet been profiled.

A.6 EVALUATION PIPELINE

The inference pipeline is left up-to the challenger to iterate on; with their own models or with API end-point based prompting with more intricate techniques. But ultimately for evaluation, the same artifact expected: a directory of `<problem-id>/answer.txt` files.

Evaluation proceeds in three stages:

(1) Syntax checking and assembly. Each candidate answer is first validated against the benchmark problem identifiers and then syntax-checked by compiling it in context as a HOL Light tactic expression. Concretely, the answer is wrapped into a small OCaml/HOL Light scaffold and checked with the benchmark syntax-checking script. Problems that fail this stage are excluded from further execution.

(2) Proof execution. Answers that pass syntax checking are assembled into benchmark evaluation files and executed in HOL Light using the repository evaluation harness. Each proof attempt is run with:

- a per-problem timeout,
- pre/post axiom-count comparison to detect unsound submissions, and
- standard compile-and-run logging.

(3) Verdict collection. Each problem receives exactly one verdict: OK, FAIL, CHEATING, TIMEOUT, or ERROR.

Our primary reported baselines use Codex-5.3 in medium-effort and high-effort modes under the evaluation configuration described above. Each mode produces one answer per problem.

Out of the full set of **2,284** benchmark problems, the medium-effort run produced **743** answers that passed syntax checking and reached proof execution, while the high-effort run produced **766**; the remaining answers were excluded at the syntax-check stage because they did not compile as valid OCaml/HOL Light tactic expressions.

Across the full benchmark, the medium-effort run solved **101 / 2,284** problems (**4.4%**) and the high-effort run solved **121 / 2,284** (**5.3%**), a net gain of **+20** proofs. Gains concentrate in the `program_state` category (+12 pp) and `generic` (+7 pp). We also host a leader-board for other research groups to make their own attempts on this problem set.

Category	Medium Effort					High Effort				
	Total	Eval	OK	FAIL	OK/Total	Total	Eval	OK	FAIL	OK/Total
generic	562	330	59	239	10.5%	562	349	66	246	11.7%
bit_vector	311	142	26	109	8.4%	311	140	27	107	8.7%
program_state	552	235	16	211	2.9%	552	229	28	193	5.1%
functional_correctness_arm	437	33	0	33	0.0%	437	42	0	42	0.0%
functional_correctness_x86	422	3	0	3	0.0%	422	6	0	6	0.0%
Total	2,284	743	101	595	4.4%	2,284	766	121	594	5.3%