AUXSPEC: LEARNING TO AUTOFORMALIZE PROGRAM SPECIFICATIONS VIA EXECUTABLE POSTCONDITIONS

Haocheng Wang[2,1*] **Weixuan Yuan**[1*] **Florian Hübler**[1] **Yinya Huang**[1,4]
Zhijiang Guo[2,3] **Sorrachai Yingchareonthawornchai**[5] **Rasmus Kyng**[1†]
[1]Department of Computer Science, ETH Zurich [2]HKUST(GZ) [3]HKUST
[4]ETH AI Center [5]Institute for Theoretical Studies, ETH Zurich

arXiv:2611.21421v1 [cs.SE] 16 Nov 2026

ABSTRACT

Formal specifications define intended program behavior and underpin verification, but writing them in theorem provers such as Lean is slow and expertise-intensive, making specification a major bottleneck to scaling formally correct code. Autoformalization from natural-language problems is promising, yet training specification models requires scalable feedback on whether generated specifications capture intended behavior. Prior approaches obtain this feedback through explicit proofs or separately implemented code with formal correspondence proofs, both expensive. We introduce AUXSPEC, a Lean framework that pairs declarative postconditions with executable counterparts, obtained through type class resolution when possible and otherwise defined separately with an equivalence proof. Test cases can then evaluate declarative specifications by execution, letting us use agreement as a reinforcement-learning reward. Using AUXSPEC, we synthesize 3,037 LeetCode problems and expand training to 19.5K samples through expert iteration on Codeforces. Experiments show that training Qwen3.5-9B and Qwen3.8-27B with the execution-based reward improves SpecGen pass@4 on AUXSPEC-Test 0.7% $\rightarrow$ 82.7% and 14.3% $\rightarrow$ 84.3%, and the 27B model outperforms all frontier models. On the out-of-domain Verina benchmark with its original prompts, both models improve 7.0% $\rightarrow$ 90.5% and 21.9% $\rightarrow$ 90.5% and surpass VeriSpecGen-30B-A3B, the prior state of the art among open-source specification models, on pass@1 (theirs 47.6% vs. ours 73.7% and 84.5%) and pass@4 (theirs 63.9% at pass@10 vs. ours 90.5%). We release the weights of the trained 9B and 27B models and the 3,037 problems with test cases, including the 300-problem AUXSPEC-Test benchmark with expert review and correctness proofs for future evaluation.

Models and data: huggingface.co/collections/eth-algopt/auxspec
Website: auxspec.org

1 INTRODUCTION

Formal program specifications define the correct behavior of a program over all valid inputs, providing the reference against which implementations can be proved correct or tested. Yet writing such specifications in an interactive theorem prover requires substantial expertise and effort, making specification a fundamental barrier to scaling formal code verification and generation. As AI models make rapid progress in generating code and proving its correctness, this barrier becomes increasingly important: realizing the promise of formally correct code requires automating specification itself. Autoformalization from natural-language programming problems offers a path forward, but immediately raises a central question: *when does a generated formal specification faithfully capture the intended problem?*

This question is particularly important in modern code-generation pipelines. Large language models (LLMs) can generate solutions to programming problems (e.g., LeetCode) from natural-language descriptions. These solutions can be evaluated using input-output test cases, which provide scalable

*Equal contribution.
†Corresponding author: kyng@inf.ethz.ch.

Table 1. Comparison of AUXSPEC with closely related work on specification generation. Specification validation checks whether a specification matches the source problem, and Expert spec. indicates a reference specification written by experts. [1] $\vee$ indicates that the check falls back to the second method when the first is inconclusive, and $\rightarrow$ that the second method follows the first. PBT indicates property-based testing. [2]The postcondition equivalence proof of AUXSPEC shows that the executable postcondition agrees with the postcondition on all valid inputs and outputs. Correctness indicates whether every implementation is proved to satisfy the specification (✓), only some are (Partial), or the work has no implementation task (–). For AUXSPEC, the proofs cover the 300 test problems. SFT, EI and RL indicate supervised fine-tuning, expert iteration and reinforcement learning.

Work	Lang.	Specification validation		Correctness	Manual review	Training
		Reference	Method[1]			
Verina (Ye et al., 2026a)	Lean	Expert spec.	Proof $\vee$ testing	Partial	Benchmark (189)	None
Clever (Thakur et al., 2025)	Lean	Expert spec.	Proof	Partial	Benchmark (161)	None
VeriSpecGen (Ye et al., 2026b)	Lean	Test cases	Proof $\vee$ LLM judge	–	Not reported	SFT
Coins (Yang et al., 2026)	Rocq	Test cases	Proof, LLM judge	–	Benchmark (164)	None
ATLAS (Baksys et al., 2025)	Dafny	Test cases	Verifier $\rightarrow$ LLM judge	✓	Not reported	SFT
Verus-SpecGym (Agarwal et al., 2026)	Verus	Test cases	Verifier $\vee$ execution	–	Not reported	None
LeetProof (Feng et al., 2026)	Lean	Test cases	LLM judge $\rightarrow$ PBT	Partial	Unproved (38)	None
AUXSPEC (ours)	Lean	Test cases	Equivalence proof[2] $\rightarrow$ execution	✓	Test set (300)	SFT, EI, RL

reward signals for reinforcement learning (Wang et al., 2025; Ruan et al., 2026). However, passing test cases does not guarantee correctness, as a solution may still fail on untested inputs (Shi et al., 2026). Formal verification offers a stronger guarantee. Given a specification, which consists of a precondition defining valid inputs and a postcondition characterizing correct outputs (Endres et al., 2024; Chen et al., 2026), a theorem prover like Lean can formally prove that an implementation satisfies its specification for every valid input. This guarantee depends on the specification correctly capturing the natural-language problem, yet producing such specifications at scale remains difficult. Manually curating specifications requires substantial expert effort: Clever and Verina contain only 161 and 189 problems, respectively (Thakur et al., 2025; Ye et al., 2026a). Automated methods can scale further, but may produce incorrect specifications (Dougherty & Mehta, 2025; Bursuc et al., 2026; Dougherty et al., 2026). Scaling specification therefore requires more than generating formal statements. It also requires efficient and scalable feedback on whether generated specifications match the source problem, analogous to the test-case feedback used to train code-generation models.

Training specification generators has been explored in auto-active verification languages, where a verifier, like an SMT solver, provides the reward by checking candidate specifications on test cases (Yan et al., 2026; Huang et al., 2026; Baksys et al., 2025). However, SMT solvers limit the complexity of specifications that can be checked (Yang et al., 2026; Ye et al., 2026a), motivating the use of interactive theorem provers such as Lean. In Lean, prior work checks generated specifications by turning each test case into a theorem to be proved, a method known as symbolic specification testing (Lahiri, 2024; Ye et al., 2026b; Yang et al., 2026). These theorems can require nontrivial proofs, especially when the specification involves unbounded quantifiers. If a proof cannot be found, the test case may be left unresolved or evaluated by an LLM judge. Thus, unlike in code generation (Wang et al., 2025; He et al., 2026), test cases do not directly provide cheap rewards for training specification generators in Lean (Table 1).

To address this issue, we propose AUXSPEC, a framework for training models to generate program specifications in Lean 4. AUXSPEC retains a declarative postcondition that specifies correctness independently of the implementation and pairs it with an executable postcondition for evaluating test cases, a pairing with a long history in formal methods (Section 6). For each synthesized problem, we prove a postcondition equivalence theorem, separately from generation (Section 2), showing that the two postconditions agree for every valid input. This theorem allows us to evaluate any number of test cases by direct execution in Lean, without proving each case separately.

We build a synthesis pipeline to generate Lean specifications and implementations for 3,037 LeetCode problems. For model training, we combine synthesized data from LeetCode and Codeforces with expert iteration to obtain 19.5K samples for supervised fine-tuning (SFT).

We fine-tune Qwen3.5-9B and Qwen3.8-27B on these samples and then apply reinforcement learning (RL) with our execution-based reward on 13,637 tasks over 8,579 further problems. The resulting AUXSPEC-9B and AUXSPEC-27B models can generate and repair both specifications and implementations. Relative to its base model, AUXSPEC-27B improves SpecGen pass@4 from 14.3% to 84.3% on AUXSPEC-Test and from 21.9% to 90.5% on Verina with its original prompts. The corresponding improvements for AUXSPEC-9B are from 0.7% to 82.7% and from 7.0% to 90.5%. Moreover, RL with our execution-based reward improves SpecGen pass@4 over SFT alone for both model sizes (Section 5.4).

Our contributions are threefold.

- **Executable specification format.** We introduce a specification format that pairs a declarative postcondition with an executable Boolean postcondition, obtained through type class resolution in Lean when a decidability instance exists and defined separately otherwise. A postcondition equivalence theorem proves that the two agree on valid inputs (Section 2).
- **AUXSPEC dataset and AUXSPEC-Test.** We release the 3,037 LeetCode problems with specifications, implementations, frozen test cases and postcondition equivalence proofs, including AUXSPEC-Test, an expert-reviewed subset of 300 problems that serves as a benchmark for Lean specification generation (Sections 3 and 3.3).
- **AUXSPEC-9B-RL and AUXSPEC-27B-RL.** Using our execution-based reward, we train and release open-weight models that generate and repair Lean specifications and implementations from natural-language descriptions. AUXSPEC-27B-RL surpasses all frontier models in domain, and both models surpass the prior state of the art among open-source specification models on Verina (Tables 4 and 5).

2 PROBLEM FORMULATION

In our autoformalization setup, we represent each programming problem with a Lean file. This section describes the file format, the generation and repair tasks over its components, and the acceptance criteria used in data synthesis, model training and evaluation.

Format. For each programming problem, we take a natural-language problem description, input constraints, and test cases as input, and aim to produce a Lean file containing a specification, an implementation, two theorem statements, and executable tests (Figures 13 and 14 in Appendix G.2.1). The specification consists of a precondition defining valid inputs, a declarative postcondition characterizing correct outputs independently of how they are computed, and an executable postcondition that expresses the same condition in a form that can be directly evaluated on test cases by Lean. The implementation is an algorithm that computes an output from a given input. The file also contains two theorems: the *postcondition equivalence theorem*, asserting that the two postconditions agree on valid inputs, and the *correctness theorem*, asserting that the implementation satisfies the declarative postcondition. We include the test cases in the Lean file as `#guard` commands that check the precondition, the executable postcondition and the implementation. We illustrate the format with two synthesized problems, *Count Number of Special Subsequences* and *Continuous Subarray Sum* (Appendix G.2.1).

Tasks. We train our models on specification generation (SpecGen) and code generation (CodeGen), as in prior work (Ye et al., 2026a), and their repair variants, SpecGen-Repair and CodeGen-Repair (Table 2). The synthesis pipeline performs these tasks and assembles the generated components into complete Lean files to build the training data and test set.

Table 2. Input and output of each task. Broken is the rejected output that a repair task revises, and Diagnostics are the compiler errors and failed test cases of that output. Test layers check the output (Table 7 in Appendix A.2).

Task	Input					Output	
	Description	Constraints	Cases	Broken	Diagnostics	Component	Test layers
SpecGen	✓	✓	✓			Specification	L1[a] L1[b] L2[a] L2[b]
SpecGen-Repair	✓	✓	✓	✓	✓	Repaired specification	L1[a] L1[b] L2[a] L2[b]
CodeGen	✓		✓			Implementation	L3
CodeGen-Repair	✓		✓	✓	✓	Repaired implementation	L3

Acceptance criteria. We apply a common set of acceptance criteria for data filtering, expert iteration, reinforcement learning and evaluation. A generated output is accepted only if the assembled file compiles, the task-specific test layers pass, and LLM judges approve (Zheng et al., 2023).

Test cases are fixed before synthesis and executed as `#guard` commands, which fail at compilation when the expression is false. The commands are organized into layers by component and type of test (Table 7 in Appendix A.2), and the generated commands must match the frozen test cases. Beyond the test cases, LLM judges check faithfulness (whether the postcondition captures the problem's requirements), non-tautology (whether the postcondition characterizes correct outputs rather than

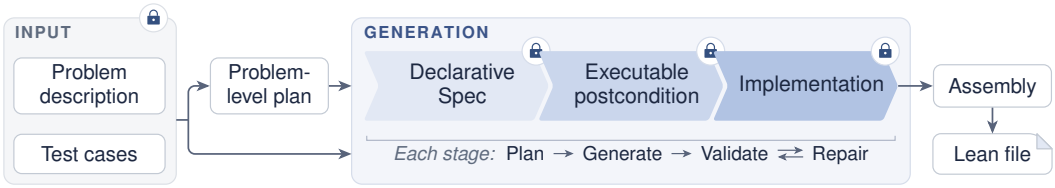


Figure 1. Data synthesis pipeline for AUXSPEC. A problem-level plan guides three sequential generation stages. Each stage plans, generates, validates, and repairs its component before the next stage. Locks mark what stays frozen during synthesis: the inputs and every accepted component. Assembly combines the accepted components with the two theorem statements into a Lean file. Figure 12 in Appendix G shows each stage on a running example.

restating a solution algorithm), completeness (whether the executable postcondition accepts correct outputs on untested inputs) and termination (whether the implementation terminates on valid inputs).

Proofs. The Lean file includes two theorem statements, whose proofs are left as `sorry` throughout data synthesis, training, and evaluation. We instead adopt executable tests and LLM judges as acceptance criteria, and assess the reliability of these checks separately. For 35.5% of the synthesized problems, the executable postcondition is obtained directly through decidability instances and agrees with the declarative postcondition by construction. After synthesis, we attempt to prove postcondition equivalence for all 3,037 problems, and 99.0% are provable without modifying their definitions (Section 3.2), which indicates that the proofs expose errors beyond tests and judges in only 1.0% of the synthesized problems. The remaining problems are repaired and then proved. For model-generated specifications, we also validate acceptance by proof, and the postcondition equivalence theorem holds for 99.2% of accepted outputs of both RL models (Appendix H.4). We note that the proved theorem shows only that the two postconditions agree, but both may encode the same misunderstanding: expert review finds 9 such errors among the 300 AUXSPEC-Test problems (Appendix I.1). Semantic agreement with the problem therefore rests on the test cases, the judges and expert review (Section 3.3).

We do not require correctness proofs for generated implementations either, since such proofs are long (a median of about 160 lines on AUXSPEC-Test, Appendix H.3) and implementation verification is not the focus of this work. We nevertheless prove correctness for all 300 AUXSPEC-Test problems to provide verified references for evaluation (Section 3.3).

3 DATA SYNTHESIS

To train and evaluate models on the tasks defined in Section 2, we synthesize and release 3,037 LeetCode problems in our Lean format. The dataset provides training data and includes a held-out, expert-reviewed set of 300 problems, AUXSPEC-Test, for evaluation. In this section, we describe the synthesis pipeline, report its results and introduce AUXSPEC-Test.

3.1 SYNTHESIS PIPELINE

Our pipeline takes a problem description and test cases as input (Figure 1). We first make a problem-level plan, then generate the declarative specification (the precondition and the declarative postcondition), the executable postcondition, and the implementation in three stages. In each stage, we plan and generate a component, validate it against the acceptance criteria of Section 2 and repair it if validation fails. We separate the stages to prevent the generation of a later component from overfitting the test cases or weakening an accepted specification (Baksys et al., 2025). Finally, we assemble the accepted components into a Lean file in the format of Section 2.

Inputs. We collect 3,109 LeetCode problems (Xia et al., 2025) with typed function signatures, reference solutions, and existing test cases (Table 3). Before synthesis, we construct and freeze positive and negative test cases for each problem (Appendix G.1.1). A positive test case pairs a valid input with the output of the reference solution, which is computed independently of the specifications under synthesis (Wang et al., 2025; Fu et al., 2025; He et al., 2026). A negative test case is either an invalid input or a valid input paired with an incorrect output.

Specification. We first generate the precondition and the declarative postcondition. To make the declarative postcondition executable, we apply `decide` when a decidability instance is available

Table 3. Number of problems and samples for each source. Problems are those used for SFT or RL, and SFT, RL and AUXSPEC-Test share no problem. Difficulty follows the labels of each source, and an LLM judge assigns the 1,484 problems without a label, including all Aizu problems (Appendix C). RL tasks are pairs of a problem and a task.

Source	Problems	Difficulty			SFT		RL		AUXSPEC-Test
		Easy	Medium	Hard	Problems	Samples	Problems	Tasks	Problems
LeetCode	2,761	629	1,442	690	2,701	14,928	60	105	300
Codeforces	6,155	1,898	1,568	2,689	1,292	4,564	4,863	7,790	0
AtCoder	900	329	385	186	0	0	900	1,452	0
Aizu	988	382	391	215	0	0	988	1,545	0
CodeChef	536	187	214	135	0	0	536	920	0
GeeksforGeeks	1,232	749	483	0	0	0	1,232	1,825	0
Total	12,572	4,174	4,483	3,915	3,993	19,492	8,579	13,637	300

and otherwise attempt to synthesize such an instance (Appendix G.2). The result, called the decidable form, agrees with the declarative postcondition by construction. When instance synthesis fails, the declarative postcondition is reformulated into a separate executable postcondition, and the postcondition equivalence theorem, proved after synthesis (Section 3.2), states that the executable postcondition agrees with the declarative postcondition.

Implementation. Given the accepted specification, we plan an algorithm and generate the implementation, where LLM judges check that the implementation terminates on all valid inputs and is fast enough for the stated constraints (Appendix G.3.1).

Repair. If a component fails validation, the model revises it from compiler errors, failed test cases and judge feedback for up to 18 rounds. Each repair round modifies only the faulty code block, leaving other blocks and accepted components unchanged, which prevents repairs from weakening a specification to pass the test cases. The recorded repairs provide the repair samples of Section 4.1.

Models. Based on a pilot comparison of frontier models (Appendix G.8), we select Kimi-K3 (Kimi Team, 2026) for generation and repair, DeepSeek-V4-Pro (DeepSeek-AI, 2026) for planning, and DeepSeek-V4-Flash for the judges.

3.2 RESULTS

The pipeline synthesizes 3,037 of the 3,109 problems (97.7%, Appendix G.5). Only 34.8% of the problems pass all stages at the first attempt, and repair recovers most of the others within 2 rounds (Appendix G.4.1). For 35.5% of the synthesized problems, the declarative postcondition can be made executable by `decide`. These executable postconditions pass validation on the first attempt in 98.1% of cases, compared with 52.4% for those generated separately (Appendix G.2.2). Since tests on small inputs do not detect nontermination or inefficiency, the implementation judge rejects 19.6% of the generated implementations that pass all test cases (Appendix G.3.1).

After synthesis, we attempt to prove the postcondition equivalence theorem for all 3,037 problems (Appendix H), and the attempts expose statement errors in 31 problems (1.0%), all of which had passed the test layers and the judges. Most of these errors (19) are in the declarative postcondition, which execution never evaluates (Table 39 in Appendix H). We repair the 31 problems and then prove the theorem. For the remaining 99.0% of the problems, the theorem is provable as synthesized.

3.3 AUXSPEC-TEST

AUXSPEC-Test is a held-out evaluation benchmark consisting of 300 problems from the synthesized dataset, with 72 easy, 155 medium and 73 hard problems, yielding 834 samples across four tasks (Table 11 in Appendix C.3). We manually review every component against the problem description and identify errors in 21 problems (Appendix I.1). Attempts to prove the correctness theorems reveal one additional error (Appendix H.3). We correct these 22 problems and prove correctness for all 300; the remaining 278 problems (92.7%) require no changes to their definitions.

These reviewed problems serve as references for evaluating model-generated specifications and implementations. For SpecGen, we pair a generated specification with the reference implementation, and for CodeGen, a generated implementation with the reference specification. The repair tasks are evaluated in the same way, and we check the resulting Lean files against the acceptance criteria of

Section 2. Since evaluation relies on test cases and judges rather than comparison with the reference specification, it can accept alternative correct formulations.

4 TRAINING

4.1 TRAINING DATA

Each synthesized problem contributes samples to the tasks of Section 2. The generation tasks take accepted components as targets, and the repair tasks take components rejected during synthesis as inputs. To obtain a starting point for RL that reasons concisely, we rephrase the long and often redundant reasoning recorded during synthesis into a short trace of the planning behind a generated component or the diagnosis behind a repair (Appendix G.6).

Increasing the number of problems improves SpecGen performance more than continued pretraining on a Mathlib corpus (Appendix E.2). We therefore source further problems for supervised fine-tuning (SFT) and reinforcement learning (RL) from Codeforces, AtCoder, Aizu, CodeChef and GeeksforGeeks through TACO (Li et al., 2023) and open-r1/codeforces (Penedo et al., 2025) (Table 3). To eliminate the risk of data contamination, the final SFT and RL data are disjoint from each other and from the evaluation benchmarks (Appendix C).

4.2 SUPERVISED FINE-TUNING

We expand the SFT data to 19.5K training samples over 3 rounds of expert iteration (Polu et al., 2023), each fine-tuning Qwen3.8-27B from the base model on the accumulated samples and adding the outputs that meet the acceptance criteria of Section 2. SpecGen performance rises as the data grow, with no sign of saturation (Appendix E.1). Fine-tuning Qwen3.5-9B (Qwen Team, 2026a) and Qwen3.8-27B (Qwen Team, 2026b) on these samples yields AUXSPEC-9B-SFT and AUXSPEC-27B-SFT, with the settings of Appendix B.

4.3 REINFORCEMENT LEARNING

We further apply RL to the SFT models for 1 epoch on 13,637 SpecGen and CodeGen tasks over 8,579 problems disjoint from the SFT data and AUXSPEC-Test (Table 3), yielding AUXSPEC-9B-RL and AUXSPEC-27B-RL. We ablate the components of our recipe (Appendix F.2) and compare with VeriSpecGen on the same backbone (Appendix F.1).

Reward design. Since the SFT models solve only 10.1% and 18.7% of SpecGen problems at pass@1 (Section 5.2), a binary reward that requires all acceptance criteria induces low reward variance, which slows optimization even when the reward is faithful to the acceptance criteria (Razin et al., 2025). We therefore score each rollout with a 6-stage progressive reward that follows the order of the acceptance criteria of Section 2 (Le et al., 2022; Huang et al., 2026). Failing a stage assigns the rollout the value of that stage, from format (0), compilation (0.1), test execution (0.2), test case alignment (0.1), non-tautology (0.1) to overall quality (0.4), and passing all stages yields 1. With v_k the value of stage k, the reward of rollout y is

$$r(y) = \begin{cases} 1 & \text{if } y \text{ passes all stages,} \\ v_k & \text{if } k \text{ is the first stage that } y \text{ fails.} \end{cases} \tag{1}$$

We value non-tautology at 0.1 to rank specifications that restate the computation without specifying the problem below failed tests. Replaying the training rollouts shows that at this value tautological specifications have a negative mean advantage in every group, whereas at a higher value such as 0.3 the advantage turns positive in groups where no rollout passes (Appendix E.3.2). Two kinds of rollouts that pass all stages receive partial credit instead of 1, enumerative specifications (0.4) and slow implementations (0.5), which evaluation accepts. To prevent reward hacking, a test counts toward the required number of tests only if it calls a declaration of the output, which excludes tests that hold trivially. Falling short of this number reduces the reward to 0.1. Hard-coding answers, by special-casing the test inputs or by copying the given test cases into the implementation, sets the reward to 0, and so does returning a repair unchanged (Appendix E.3.1).

Self-repair. Besides generation, we train on repair tasks built from failed rollouts of the current policy, which avoids distribution shift (Kumar et al., 2025). During training, these failed rollouts

are paired with compiler messages or judge feedback in a buffer, from which 40% of each batch is drawn as single-turn repair tasks (Table 2). Repairs continue for up to 3 rounds, and the reward of a repair rollout is scaled down linearly from 1 in the first round to 0.8 in the last.

Policy optimization. Our RL algorithm extends Decoupled Clip and Dynamic Sampling Policy Optimization (DAPO; Yu et al., 2025) by normalizing the advantage by the standard deviation over the batch instead of the group (Hu et al., 2025) and masking the clipped surrogate $\ell_{i,t}$ of token t in rollout i with $m_{i,t} \in \{0, 1\}$. We maximize

$$\mathcal{J}(\theta) = \mathbb{E}\Big[\frac{\sum_{i,t} m_{i,t}\, \ell_{i,t}(\theta)}{\sum_{i,t} m_{i,t}}\Big], \tag{2}$$

which reduces to the original algorithm with $m_{i,t} = 1$ and group normalization. We classify a rollout as repetitive if a 40-character n-gram occurs at least 20 times and over 40% of n-grams are repeats (Welleck et al., 2020). Repetitive rollouts that reach the length limit receive reward 0, and the mask removes the gradient of such a rollout beyond 200 tokens past the repetition onset. For the 27B model, the loss adds a penalty with weight 0.1 on the probability of the end-of-sequence token inside the thinking segment where it exceeds 1e-6. Appendix B lists the other hyperparameters.

5 Experiments

5.1 Setup

Training settings, including hyperparameters and the Lean environment, are listed in Appendix B, and the sampling settings in Appendix B.2. We report pass@k for $k \in \{1, 4\}$, the fraction of problems with at least one of k samples meeting the acceptance criteria of Section 2 (Appendix D.1). All evaluated models receive the same prompts (Appendix K). The baselines are our backbones, frontier models and prior work on specification generation.

- **Instruct Models:** To isolate the effect of fine-tuning, we evaluate the backbones Qwen3.5-9B (Qwen Team, 2026a) and Qwen3.8-27B (Qwen Team, 2026b).
- **Frontier Models:** Claude-Opus-5 (Anthropic, 2026a), Claude-Sonnet-5 (Anthropic, 2026b), GPT-5.6-Sol and GPT-5.6-Luna (OpenAI, 2026), Kimi-K3 (Kimi Team, 2026), and DeepSeek-V4-Pro and DeepSeek-V4-Flash (DeepSeek-AI, 2026).
- **Prior Work:** We compare against VeriSpecGen (Ye et al., 2026b), which fine-tunes Qwen3-4B and Qwen3-Coder-30B-A3B on refinement trajectories distilled from a closed teacher, using the reported results as the models are not publicly available.

5.2 Main results

We report the performance on the in-domain test set in Table 4. Training improves SpecGen pass@4 from 0.7% to 82.7% for the 9B backbone and from 14.3% to 84.3% for the 27B backbone. AuXSpec-27B-RL surpasses every frontier model in both pass@1 and pass@4 on SpecGen and CodeGen, with 65.7% SpecGen pass@1 (vs. 57.7% for Claude-Opus-5, the strongest frontier model). The 9B RL model also outperforms every frontier model in SpecGen pass@1 (58.6%). On the repair tasks, the 27B RL model attains the best CodeGen-Repair pass@1 (88.1%) and stays within 2 points of the best result on the other metrics, while the 9B RL model surpasses most frontier models in SpecGen-Repair pass@1. To assess false positives of SpecGen acceptance, we prove the postcondition equivalence theorem for 99.2% of accepted RL outputs (Appendix H.4), and expert review confirms 90 of 100 sampled outputs, with all errors in the validity predicate (Appendix I.2).

5.3 Out-of-domain evaluation

To assess generalization beyond the synthesized data, we evaluate on Verina (Ye et al., 2026a), a 189-problem manually curated benchmark with reference specifications, implementations and test cases. Following the original evaluation, generated specifications must be sound and complete with respect to the reference specification on the test cases, and generated implementations must pass all positive test cases. AuXSpec-27B-RL outperforms all baselines in SpecGen pass@1 with 91.9% and trails the best baseline in pass@4 by 1.1 points (Table 5). On CodeGen, both RL models outperform all

Table 4. Main results on AUXSPEC-Test. **Bold** indicates the best result for each metric and underline the second best. Green numbers in parentheses show the gain in points over the backbone of the same size.

Model	SpecGen		CodeGen		SpecGen-Repair		CodeGen-Repair	
	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4
Frontier Models								
Claude-Opus-5	57.7	81.3	93.0	98.3	**83.0**	**90.1**	86.9	95.1
Claude-Sonnet-5	13.8	39.3	54.8	73.0	66.4	81.7	77.2	88.3
GPT-5.6-Sol	29.8	59.7	80.0	96.3	61.8	77.1	82.3	89.3
GPT-5.6-Luna	9.5	28.7	39.9	71.3	49.4	69.5	71.8	87.4
Kimi-K3	43.8	83.7	67.2	94.3	73.5	87.8	81.3	**97.1**
DeepSeek-V4-Pro	21.2	54.3	48.6	72.3	57.6	82.4	76.9	94.2
DeepSeek-V4-Flash	38.4	76.7	45.0	69.7	67.0	88.5	73.3	95.1
Instruct Models								
Qwen3.5-9B	0.2	0.7	8.2	18.0	8.4	26.0	4.9	14.6
Qwen3.8-27B	4.5	14.3	8.9	22.0	44.7	73.3	35.7	64.1
Ours								
AUXSPEC-9B-SFT	10.1 (+9.9)	24.3 (+23.6)	13.1 (+4.9)	28.3 (+10.3)	56.7 (+48.3)	76.3 (+50.3)	59.5 (+54.6)	80.6 (+66.0)
AUXSPEC-9B-RL	58.6 (+58.4)	82.7 (+82.0)	77.1 (+68.9)	96.0 (+78.0)	73.5 (+65.1)	84.7 (+58.7)	74.5 (+69.6)	94.2 (+79.6)
AUXSPEC-27B-SFT	18.7 (+14.2)	40.7 (+26.4)	27.6 (+18.7)	52.7 (+30.7)	73.1 (+28.4)	84.7 (+11.4)	72.3 (+36.6)	87.4 (+23.3)
AUXSPEC-27B-RL	**65.7** (+61.2)	**84.3** (+70.0)	**94.8** (+85.9)	**99.7** (+77.7)	81.9 (+37.2)	89.3 (+16.0)	**88.1** (+52.4)	95.1 (+31.0)

baselines, with pass@1 of 92.6% for the 9B model and 91.8% for the 27B model (vs. 56.6% for the best baseline). The 9B model also attains the highest pass@4 (98.4%). Including the test cases in the prompt improves SpecGen pass@4 of both RL models from 90.5% to 96.8%, and the same test cases improve the pass@4 of frontier models by 5 to 24 points (Table 5).

Table 5. Out-of-domain evaluation on Verina without and with test cases in the prompt. The setting without test cases follows the default of Verina. Bold, underlines and green numbers as in Table 4. † indicates results reported in the original paper (Ye et al., 2026b) and scored by the same Verina evaluator, which are available only without test cases in the prompt. VeriSpecGen-4B and VeriSpecGen-30B-A3B are the SFT No-Test models of VeriSpecGen fine-tuned from Qwen3-4B-Instruct-2507 and Qwen3-Coder-30B-A3B, respectively. Italic values are pass@10 in place of pass@4 and are not underlined.

Model	w/o tests				w/ tests			
	SpecGen		CodeGen		SpecGen		CodeGen	
	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4
Frontier Models								
Claude-Opus-5	77.8	89.9	57.0	78.8	86.2	95.2	56.6	79.4
Claude-Sonnet-5	57.8	85.7	20.2	46.6	64.8	90.5	17.9	41.8
GPT-5.6-Sol	42.1	73.5	16.3	35.4	85.1	**97.9**	20.8	40.7
GPT-5.6-Luna	50.8	86.2	20.2	51.9	67.3	94.7	24.5	58.2
Kimi-K3	69.2	89.4	34.0	69.8	77.6	96.3	34.9	72.5
DeepSeek-V4-Pro	44.7	76.7	32.3	61.4	53.8	82.5	37.6	70.4
DeepSeek-V4-Flash	51.9	84.7	6.5	21.2	60.2	92.6	5.8	18.5
Instruct Models								
Qwen3.5-9B	2.0	7.0	25.3	48.7	7.0	18.5	23.4	49.2
Qwen3.8-27B	7.2	21.9	4.2	15.9	28.8	55.0	3.3	13.2
Prior Work								
VeriSpecGen-4B	30.3†	*45.2*†	35.4†	*66.8*†	–	–	–	–
VeriSpecGen-30B-A3B	47.6†	*63.9*†	58.7†	*84.0*†	–	–	–	–
Ours								
AUXSPEC-9B-SFT	34.9 (+32.9)	65.6 (+58.6)	36.5 (+11.2)	71.4 (+22.7)	37.3 (+30.3)	69.8 (+51.3)	38.0 (+14.6)	70.4 (+21.2)
AUXSPEC-9B-RL	73.7 (+71.7)	**90.5** (+83.5)	**89.8** (+64.5)	**97.4** (+48.7)	87.0 (+80.0)	96.8 (+78.3)	**92.6** (+69.2)	**98.4** (+49.2)
AUXSPEC-27B-SFT	52.9 (+45.7)	77.2 (+55.3)	58.7 (+54.5)	85.2 (+69.3)	57.3 (+28.5)	82.5 (+27.5)	57.8 (+54.5)	86.8 (+73.6)
AUXSPEC-27B-RL	**84.5** (+77.3)	**90.5** (+68.6)	89.7 (+85.5)	93.7 (+77.8)	**91.9** (+63.1)	96.8 (+41.8)	91.8 (+88.5)	93.7 (+80.5)

We further evaluate general capabilities on AIME 25 (Mathematical Association of America, 2025), GPQA-Diamond (Rein et al., 2024), and HumanEval+ and MBPP+ (Liu et al., 2023) with the sampling settings of DeepSeek-R1 (Guo et al., 2025). Supervised fine-tuning lowers every score, and RL recovers about half or more of each drop, with AUXSPEC-9B-RL surpassing its backbone on both code benchmarks (Appendix D.3).

Besides providing additional learning signal during training, the repair tasks enable the model to iteratively refine its specifications based on previous outputs and feedback from the compiler and

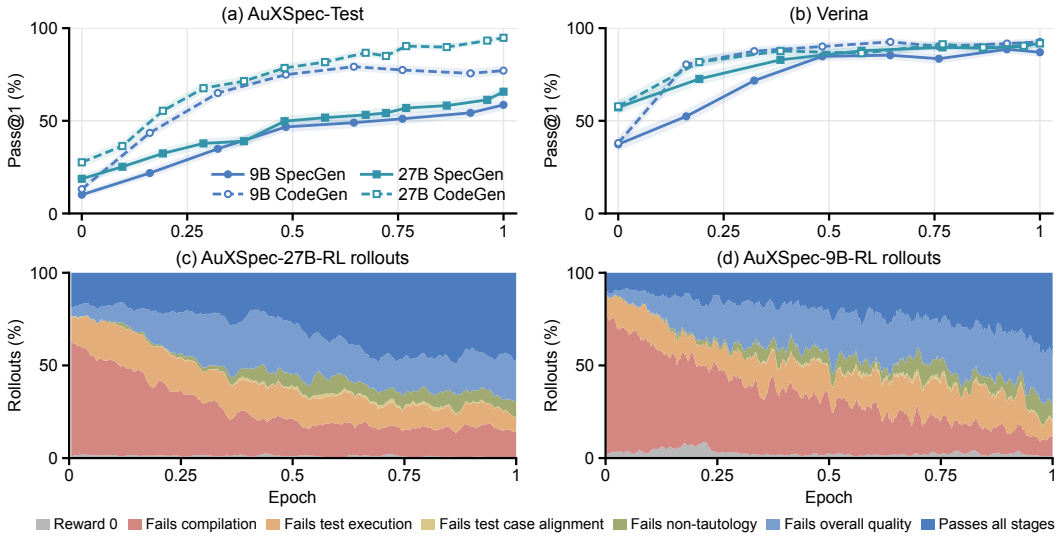


Figure 2. Training dynamics of RL over 1 epoch. (a, b) SpecGen (solid) and CodeGen (dashed) pass@1 of the 9B and 27B models on AUXSPEC-Test and Verina, where epoch 0 is the SFT model and bands show 1 standard error. (c, d) Fraction of rollouts by reward outcome, from the first failed stage to passing all stages, over all 4 tasks with a 5-step moving average.

judges at test time. On ICPC-Eval (Xu et al., 2025b), a benchmark of problems from 11 recent contests of the International Collegiate Programming Contest (ICPC), AUXSPEC-27B-RL solves a higher percentage of problems with refine@16 (82.7%) than with pass@16 (61.3%, Appendix D.4).

5.4 RL ANALYSIS

Reinforcement learning from the SFT models improves SpecGen pass@1 on AUXSPEC-Test from 10.1% to 58.6% for the 9B model and from 18.7% to 65.7% for the 27B model (Table 4 and Figure 2a), and the gains transfer to Verina, where SpecGen pass@1 rises from 37.3% to 87.0% and from 57.3% to 91.9% (Table 5 and Figure 2b). The reward is computed on the synthesized problems only, yet Verina scores each specification against the reference specification on test cases absent from the prompt, so the policy learns faithfulness to the description beyond the frozen test cases.

Gains at pass@1 can also come from sharpening the sampling distribution, in which case the SFT model would catch up at large k (Yue et al., 2025). SpecGen pass@256 of AUXSPEC-27B-SFT (85.7%) stays below pass@8 of AUXSPEC-27B-RL (88.2%), and the SFT model therefore needs more than 32 times the samples of the RL model to reach this pass rate. Within 256 samples, 33 of the 300 problems are solved only by the RL model and 1 only by the SFT model (Appendix E.3.3). The RL model is also token-efficient, requiring 20.3% of the output tokens per solved SpecGen sample of the most economical frontier model (Appendix D.2).

We utilize a progressive reward to assign fine-grained credit to partial progress. A binary reward assigns 0 to every rollout that fails any stage and provides no signal in groups where no sample passes. With fine-grained values, the group-relative advantage can distinguish better samples from worse ones within a group, which raises the probability of the better samples until the group saturates. Replaying the rollouts of the final runs confirms the denser signal. Groups without a passing rollout make up 37% of SpecGen groups for the 27B model and 47% for the 9B model and would be discarded by dynamic sampling under a binary reward (Yu et al., 2025), whereas the progressive reward yields nonzero variance in 87.9% and 86.2% of them (Appendix E.3.1). The rollouts shift accordingly and follow this group-relative credit. In the early stage, few CodeGen groups contain a fast and correct program, where slow but correct programs receive a positive advantage and grow (up to 67.4% of correct CodeGen rollouts of the 27B model). These slow programs then decrease gradually as fast programs appear (0.4% at the end), whereas the 9B model, which solves fewer CodeGen problems, ends with more of them (30.4%). Enumerative specifications show a similar pattern (17.0% to 2.8% of otherwise accepted specifications for the 27B model, Figure 2c).

The selection of the intermediate values relies on the reward dynamics, which we measure by replaying the training rollouts and recording the advantage of each stage in groups of different strength.

For instance, correct but slow implementations are valued at 0.5, at which their advantage is positive only in groups with fewer than two passing rollouts, whereas at 0.8 it would stay positive with six passing rollouts of 16 (Appendix E.3.2). As CodeGen groups without a passing rollout fall from 44% to 0% over training (Figure 4b2), the mean advantage of slow implementations turns from $+0.10$ to -0.17 (steps 1–40 and 161–208). In SpecGen, 29% to 52% of groups remain without a passing rollout in each period (Figure 4b1), so hard SpecGen problems contribute a learning signal throughout training.

6 Related work

Specification generation. Specification autoformalization translates requirements into formal specifications. Prior work studies postcondition generation from natural language (Endres et al., 2024), executable Python specifications (Chen et al., 2026) and user intent in Dafny (Lahiri, 2024).

Training for specification generation. Training for program specification generation has concentrated on auto-active verification languages (Section 1). In Lean and Rocq, prior work checks a specification by proving each test case separately or by searching for counterexamples (Feng et al., 2026), and reinforcement learning from verifier feedback has targeted correctness proofs (Li et al., 2026b) but remains underexplored for specification generation (Baksys et al., 2025). Specification generators there are trained instead on synthesized and repaired examples, where a reference implementation labels tests, a proof attempt checks each labeled test, and LLM judgments guide refinement when proof automation is inconclusive (Ye et al., 2026b). AuXSpec also labels tests with reference solutions, and relies on the test results of the executable postcondition to filter the training data and to define the reinforcement-learning reward (Section 2).

Executable specifications. The executability of specifications has long been debated in formal methods. Executable specifications can be validated against intended behavior by execution (Kemmerer, 1985; Fuchs, 1992), whereas declarative specifications are more expressive and independent of any algorithm (Hayes & Jones, 1989). Runtime assertion checking and property-based testing run contracts and properties as tests (Meyer, 1992; Cheon & Leavens, 2002; Claessen & Hughes, 2000), and proof assistants test conjectures before proving them (Berghofer & Nipkow, 2004; Lean Community, 2024). Boolean reflection pairs a proposition with a Boolean function and a proof that the two agree (Gonthier & Mahboubi, 2010; Paraskevopoulou et al., 2015). Executable counterparts can also be derived automatically from inductive relations in proof assistants, but the derivation covers a restricted class of relations, and users still write executable variants by hand and prove that the two agree (Berghofer et al., 2009; Paraskevopoulou et al., 2022). AuXSpec utilizes both sides of this trade-off, where we let models generate the executable variants, a step that previously required manual curation.

Adjacent settings. Generating specifications for existing implementations utilizes a verifier to check the specification against the code but assumes the code is correct (Chen et al., 2025; Yan et al., 2026). In this setting, reinforcement learning has included executable checks on mutated-output tests in the reward (Huang et al., 2026). Our SpecGen task instead starts from the problem description, input constraints and tests, without access to an implementation. Vericoding benchmarks evaluate the reverse task, generating implementations and proofs from a given formal specification (Zhao et al., 2026; Bursuc et al., 2026).

7 Conclusion

We presented AuXSpec, a framework for training models to generate program specifications in Lean with execution-based feedback. AuXSpec pairs each declarative postcondition with an executable postcondition, and the postcondition equivalence theorem states that both agree on all valid inputs. This allows us to evaluate specifications by directly executing test cases, without proving each case separately. At the same time, we retain the declarative postcondition for formal correctness proofs. With this format, we synthesize and release 3,037 LeetCode problems, each with a proved postcondition equivalence theorem. The dataset includes AuXSpec-Test, a held-out benchmark of 300 expert-reviewed problems, all with proved correctness theorems.

Training on the synthesized data with an execution-based reinforcement-learning reward yields 9B and 27B models, and the 27B model surpasses all frontier models on AUXSPEC-Test and outperforms all baselines in SpecGen pass@1 on the out-of-domain Verina benchmark. Since the models are trained only on synthesized problems, the gains on Verina suggest that training with the reward improves faithfulness to the problem description beyond the test cases used in the reward. We discuss the limitations of this work in Appendix J.

ACKNOWLEDGMENTS

This research was supported by the starting grant "A New Paradigm for Flow and Cut Algorithms" (no. TMSGI2 218022) of the Swiss National Science Foundation and the Swiss AI small project grant "Formal Reasoning in Sparse Data Regimes" (a0093). Sorrachai Yingchareonthawornchai is supported by Dr. Max Rössler, the Walter Haefner Foundation, and the ETH Zürich Foundation.

AI USE STATEMENT

We used generative AI tools to synthesize the dataset and to produce proofs. The pipeline uses LLMs to synthesize, repair, judge and rewrite specifications and implementations (Section 3, Appendices G and K), and proof agents built on LLMs produce the postcondition equivalence and correctness proofs, which the Lean kernel checks (Appendix H). We did not use generative AI tools to develop the conceptual framework or to interpret results, and translation and qualitative data analysis do not apply to this work. We also used generative AI tools to edit the text, create figures, search the literature and format references. We reviewed all AI-assisted work, ran the AI-generated code, checked every number against the script that computed it and every citation against the original publication, and take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

One of the authors carried out the expert review of AUXSPEC-Test, and 2 of the authors reviewed the accepted outputs (Appendix I), and the work involves no other human participants and no personal data. The problems come from LeetCodeDataset (Xia et al., 2025) and TACO (Li et al., 2023), both released under the Apache-2.0 license, with the web-crawled problem statements of TACO under CC BY 4.0. We use them for research with attribution and will release our data and models under compatible licenses. The released models generate specifications and implementations for programming problems, and every accepted implementation remains checkable by proof against the postcondition.

REPRODUCIBILITY STATEMENT

The synthesis pipeline is described in Section 3, with every stage in Appendix G, the counts at each step in Appendix G.5, the models used in Appendix G.8, condensed prompts of the stages and judges in Appendix K, and the proving of the postcondition equivalence and correctness theorems in Appendix H. Training is described in Section 4, with the training and sampling settings in Appendices B and B.2, expert iteration in Appendix E.1, the metrics in Appendix D.1 and the reinforcement-learning reward and ablations in Appendix E.3. We will release the synthesized problems with test cases and postcondition equivalence proofs, AUXSPEC-Test with correctness proofs, the evaluation code and the model weights.

REFERENCES

Anmol Agarwal, Natalie Neamtu, Pranjal Aggarwal, Seungone Kim, Jannis Limperg, Cedric Flamant, Kanna Shimizu, Bryan Parno, and Sean Welleck. Verus-specgym: An agentic environment for evaluating specification autoformalization. *arXiv preprint arXiv:2605.26457*, 2026.

Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. On-policy distillation of language models: Learning from self-generated mistakes. In *International Conference on Learning Representations (ICLR)*, 2024.

Zeyuan Allen-Zhu and Yuanzhi Li. Physics of language models: Part 3.1, knowledge storage and extraction. In *International Conference on Machine Learning (ICML)*, 2024.

Anthropic. Introducing claude opus 5. https://www.anthropic.com/news/claude-opus-5, 2026a.

Anthropic. Introducing claude sonnet 5. https://www.anthropic.com/news/claude-sonnet-5, 2026b.

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.

Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. Llemma: An open language model for mathematics. In *International Conference on Learning Representations (ICLR)*, 2024.

Mantas Baksys, Stefan Zetzsche, Olivier Bouissou, Remi Delmas, Soonho Kong, and Sean B. Holden. Atlas: Automated toolkit for large-scale verified code synthesis. *arXiv preprint arXiv:2512.10173*, 2025.

Stefan Berghofer and Tobias Nipkow. Random testing in Isabelle/HOL. In *IEEE International Conference on Software Engineering and Formal Methods (SEFM)*, pp. 230–239, 2004.

Stefan Berghofer, Lukas Bulwahn, and Florian Haftmann. Turning inductive into equational specifications. In *International Conference on Theorem Proving in Higher Order Logics (TPHOLs)*, pp. 131–146, 2009.

Sergiu Bursuc, Theodore Ehrenborg, Shaowei Lin, Lacramioara Astefanoaei, Ionel Emilian Chiosa, Jure Kukovec, Alok Singh, Oliver Butterley, Adem Bizid, Quinn Dougherty, et al. A benchmark for vericoding: Formally verified program synthesis. *Transactions on Machine Learning Research*, 2026.

Saikat Chakraborty, Gabriel Ebner, Siddharth Bhat, Sarah Fakhoury, Sakina Fatima, Shuvendu Lahiri, and Nikhil Swamy. Towards neural synthesis for SMT-assisted proof-oriented programming. In *IEEE/ACM International Conference on Software Engineering (ICSE)*, 2025.

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.

Tianyu Chen, Shuai Lu, Shan Lu, Yeyun Gong, Chenyuan Yang, Xuheng Li, Md Rakib Hossain Misu, Hao Yu, Nan Duan, Peng Cheng, et al. Automated proof generation for rust code via self-evolution. In *International Conference on Learning Representations (ICLR)*, 2025.

Zaoyu Chen, Jianbo Dai, Boyu Zhu, Jingdong Wang, Huiming Wang, Xin Xu, Haoyang Yuan, Zhijiang Guo, and Xiao-Ming Wu. CodeSpecBench: Benchmarking LLMs for executable behavioral specification generation. *arXiv preprint arXiv:2604.12268*, 2026.

Daixuan Cheng, Shaohan Huang, and Furu Wei. Adapting large language models via reading comprehension. In *International Conference on Learning Representations (ICLR)*, 2024.

Yoonsik Cheon and Gary T. Leavens. A simple and practical approach to unit testing: The JML and JUnit way. In *European Conference on Object-Oriented Programming (ECOOP)*, pp. 231–255, 2002.

Koen Claessen and John Hughes. QuickCheck: A lightweight tool for random testing of Haskell programs. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pp. 268–279, 2000.

Xingyu Dang, Christina Baek, J. Zico Kolter, and Aditi Raghunathan. Assessing diversity collapse in reasoning. In *ICLR 2025 Workshop on Scaling Self-Improving Foundation Models without Human Supervision*, 2025.

DeepSeek-AI. DeepSeek-V4: Towards highly efficient million-token context intelligence. *arXiv preprint arXiv:2606.19348*, 2026.

Quinn Dougherty and Ronak Mehta. Proving the coding interview: A benchmark for formally verified code generation. In *IEEE/ACM International Workshop on Large Language Models for Code (LLM4Code)*, pp. 72–79, 2025.

Quinn Dougherty, Max von Hippel, Simon Henniger, Hazel Shackleton, and Mike Dodds. Fvspec: Real-world property-based tests as lean challenges. *arXiv preprint arXiv:2606.01008*, 2026.

Madeline Endres, Sarah Fakhoury, Saikat Chakraborty, and Shuvendu K. Lahiri. Can large language models transform natural language intent into formal method postconditions? *Proceedings of the ACM on Software Engineering*, 1(FSE):1889–1912, 2024. doi: 10.1145/3660791.

Yueyang Feng, Dipesh Kafle, Vladimir Gladshtein, Vitaly Kurin, George Pîrlea, Qiyuan Zhao, Peter Müller, and Ilya Sergey. Certified program synthesis with a multi-modal verifier. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2026.

Jia Fu, Xinyu Yang, Hongzhi Zhang, Yahui Liu, Jingyuan Zhang, Qi Wang, Fuzheng Zhang, and Guorui Zhou. Klear-CodeTest: Scalable test case generation for code reinforcement learning. *arXiv preprint arXiv:2508.05710*, 2025.

Norbert E. Fuchs. Specifications are (preferably) executable. *Software Engineering Journal*, 7(5): 323–334, 1992.

GasStationManager. SafeVerify: A Lean 4 script for robustly verifying submitted proofs of theorems and implementations of functions. https://github.com/GasStationManager/SafeVerify, 2024.

Georges Gonthier and Assia Mahboubi. An introduction to small scale reflection in Coq. *Journal of Formalized Reasoning*, 3(2):95–152, 2010.

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*, 2023.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 645:633–638, 2025.

Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. Don't stop pretraining: Adapt language models to domains and tasks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.

Ian J. Hayes and Cliff B. Jones. Specifications are not (necessarily) executable. *Software Engineering Journal*, 4(6):330–338, 1989.

Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang Zhang, Jiacheng Xu, Wei Shen, et al. Skywork open reasoner 1 technical report. *arXiv preprint arXiv:2505.22312*, 2025.

Zhongmou He, Yee Man Choi, Kexun Zhang, Ivan Bercovich, Jiabao Ji, Junting Zhou, Dejia Xu, Aidan Zhang, Yixiao Zeng, and Lei Li. HardTestGen: A high-quality RL verifier generation pipeline for LLM algorithmic coding. In *International Conference on Learning Representations (ICLR)*, 2026.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations (ICLR)*, 2020.

Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen. REINFORCE++: Stabilizing critic-free policy optimization with global advantage normalization. *arXiv preprint arXiv:2501.03262*, 2025.

Zhechong Huang, Zhao Zhang, Zeyu Sun, Huifeng Sun, and Yingfei Xiong. Specrl: Reinforcement learning with test-based completeness rewards for formal specification synthesis. In *IEEE Transactions on Software Engineering*, 2026.

Albert Q. Jiang, Sean Welleck, Jin Peng Zhou, Timothée Lacroix, Jiacheng Liu, Wenda Li, Mateja Jamnik, Guillaume Lample, and Yuhuai Wu. Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. In *International Conference on Learning Representations (ICLR)*, 2023.

Zhengbao Jiang, Zhiqing Sun, Weijia Shi, Pedro Rodriguez, Chunting Zhou, Graham Neubig, Xi Victoria Lin, Wen tau Yih, and Srinivasan Iyer. Instruction-tuned language models are better knowledge learners. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.

Woogyeol Jin, Taywon Min, Yongjin Yang, Dennis Wei, Yi Zhou, Swanand Ravindra Kadhe, Nathalie Baracaldo, and Kimin Lee. Entropy-aware on-policy distillation of language models. In *International Conference on Machine Learning (ICML)*, 2026.

Richard A. Kemmerer. Testing formal specifications to detect design errors. *IEEE Transactions on Software Engineering*, SE-11(1):32–43, 1985.

Kimi Team. Kimi k3: Open frontier intelligence. *arXiv preprint arXiv:2607.24653*, 2026.

Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, JD Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, Lei M. Zhang, Kay McKinney, Disha Shrivastava, Cosmin Paduraru, George Tucker, Doina Precup, Feryal Behbahani, and Aleksandra Faust. Training language models to self-correct via reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2025.

Shuvendu K. Lahiri. Evaluating LLM-driven user-intent formalization for verification-aware languages. In *Formal Methods in Computer-Aided Design (FMCAD)*, 2024.

Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Chu Hong Hoi. CodeRL: Mastering code generation through pretrained models and deep reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Lean community. REPL: A read-eval-print loop for Lean 4. https://github.com/leanprover-community/repl, 2023.

Lean Community. Plausible: A property testing framework for Lean 4. https://github.com/leanprover-community/plausible, 2024.

Rongao Li, Jie Fu, Bo-Wen Zhang, Tao Huang, Zhihong Sun, Chen Lyu, Guang Liu, Zhi Jin, and Ge Li. TACO: Topics in algorithmic COde generation dataset, 2023. URL https://arxiv.org/abs/2312.14852.

Yaxuan Li, Yuxin Zuo, Bingxiang He, Jinqian Zhang, Chaojun Xiao, Cheng Qian, Tianyu Yu, Huan ang Gao, Wenkai Yang, Zhiyuan Liu, et al. Rethinking on-policy distillation of large language models: Phenomenology, mechanism, and recipe. *arXiv preprint arXiv:2604.13016*, 2026a.

Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d'Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. Competition-level code generation with AlphaCode. *Science*, 378(6624):1092–1097, 2022.

Zenan Li, Ziran Yang, Deyuan He, Haoyu Zhao, Andrew Zhao, Shange Tang, Kaiyu Yang, Aarti Gupta, Zhendong Su, and Chi Jin. Goedel-code-prover: Hierarchical proof search for open state-of-the-art code verification. In *Conference on Language Modeling (COLM)*, 2026b.

Wenze Lin, Jiale Zhao, Xitai Jiang, Songde Rao, Yining Li, Shenzhi Wang, Bingxiang He, and Gao Huang. On-policy distillation with verifiable reward. *arXiv preprint arXiv:2608.24696*, 2026.

Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by ChatGPT really correct? rigorous evaluation of large language models for code generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong. ProRL: Prolonged reinforcement learning expands reasoning boundaries in large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

Kevin Lu and Thinking Machines Lab. On-policy distillation. https://thinkingmachines.ai/blog/on-policy-distillation/, 2025.

Mathematical Association of America. American invitational mathematics examination (AIME) 2025. https://maa.org/maa-invitational-competitions/, 2025.

Bertrand Meyer. Applying "design by contract". *Computer*, 25(10):40–51, 1992.

Theo X. Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, and Armando Solar-Lezama. Is self-repair a silver bullet for code generation? In *International Conference on Learning Representations (ICLR)*, 2024.

OpenAI. GPT-5.6: Frontier intelligence that scales with your ambition. https://openai.com/index/gpt-5-6/, 2026.

Zoe Paraskevopoulou, Cătălin Hriţcu, Maxime Dénès, Leonidas Lampropoulos, and Benjamin C. Pierce. Foundational property-based testing. In *International Conference on Interactive Theorem Proving (ITP)*, pp. 325–343, 2015.

Zoe Paraskevopoulou, Aaron Eline, and Leonidas Lampropoulos. Computing correctly with inductive relations. In *ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI)*, pp. 966–980, 2022.

Guilherme Penedo, Anton Lozhkov, Hynek Kydlíček, Loubna Ben Allal, Edward Beeching, Agustín Piqueres Lajárin, Quentin Gallouédec, Nathan Habib, Lewis Tunstall, and Leandro von Werra. CodeForces. https://huggingface.co/datasets/open-r1/codeforces, 2025.

Stanislas Polu, Jesse Michael Han, Kunhao Zheng, Mantas Baksys, Igor Babuschkin, and Ilya Sutskever. Formal mathematics statement curriculum learning. In *International Conference on Learning Representations (ICLR)*, 2023.

Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026a. URL https://qwen.ai/blog?id=qwen3.5.

Qwen Team. Qwen3.8-Max: A new bar for coding and cowork, August 2026b. URL https://qwen.ai/blog?id=qwen3.8.

Noam Razin, Zixuan Wang, Hubert Strauss, Stanley Wei, Jason D. Lee, and Sanjeev Arora. What makes a reward model a good teacher? an optimization perspective. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level Google-proof Q&A benchmark. In *Conference on Language Modeling (COLM)*, 2024.

Z. Z. Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanjia Zhao, Liyue Zhang, Zhe Fu, Qihao Zhu, Dejian Yang, et al. DeepSeek-Prover-V2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition. *arXiv preprint arXiv:2504.21801*, 2025.

Chi Ruan, Dongfu Jiang, Huaye Zeng, Ping Nie, and Wenhu Chen. EvolveCoder: Evolving test cases via adversarial verification for code reinforcement learning. *arXiv preprint arXiv:2603.12698*, 2026.

Marco Dos Santos, Hugues de Saxcé, Haiming Wang, Ran Wang, Mantas Baksys, Mert Unsal, Junqi Liu, Zhengying Liu, and Jia Li. Kimina Lean server: A high-performance Lean server for large-scale verification. In *The 5th Workshop on Mathematical Reasoning and AI at the 39th Conference on Neural Information Processing Systems (NeurIPS)*, 2025.

John Schulman. Approximating KL divergence. http://joschu.net/blog/kl-approx.html, 2020.

Vedant Shah, Johan Obando-Ceron, Vineet Jain, Brian Bartoldson, Bhavya Kailkhura, Sarthak Mittal, Glen Berseth, Pablo Samuel Castro, Yoshua Bengio, Esmeralda S. Whitammer, et al. A comedy of estimators: On KL regularization in RL training of LLMs. In *Conference on Language Modeling (COLM)*, 2026.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, et al. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

Jingwei Shi, Xinxiang Yin, Jing Huang, Jinman Zhao, and Shengyu Tao. CodeHacker: Automated test case generation for detecting vulnerabilities in competitive programming solutions. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 2352–2382, 2026.

Zhengxiang Shi and Aldo Lipani. Don't stop pretraining? make prompt-based fine-tuning powerful learner. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Amitayush Thakur, Jasper Lee, George Tsoukalas, Meghana Sistla, Matthew Zhao, Stefan Zetzsche, Greg Durrett, Yisong Yue, and Swarat Chaudhuri. CLEVER: A curated benchmark for formally verified code generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

Chen Wang, Zhaochun Li, Jionghao Bai, Yining Zhang, Hexuan Deng, Ge Lan, and Yue Wang. Distilled reinforcement learning for LLM post-training. *arXiv preprint arXiv:2607.17247*, 2026a.

Haiming Wang, Huajian Xin, Chuanyang Zheng, Zhengying Liu, Qingxing Cao, Yinya Huang, Jing Xiong, Han Shi, Enze Xie, Jian Yin, et al. LEGO-prover: Neural theorem proving with growing libraries. In *International Conference on Learning Representations (ICLR)*, 2024.

Rui Wang, Hongru Wang, Yi Chen, Boyang Xue, Tianqing Fang, Wenhao Yu, and Kam-Fai Wong. Demystifying on-policy distillation: Roles, pathologies, and regulations. *arXiv preprint arXiv:2607.13399*, 2026b.

Zihan Wang, Siyao Liu, Yang Sun, Ming Ding, and Hongyan Li. CodeContests+: High-quality test case generation for competitive programming. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 5576–5600, 2025.

Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. Neural text generation with unlikelihood training. In *International Conference on Learning Representations (ICLR)*, 2020.

Yunhui Xia, Wei Shen, Yan Wang, Jason Klein Liu, Huifeng Sun, Siyue Wu, Jian Hu, and Xiaolong Xu. LeetCodeDataset: A temporal dataset for robust evaluation and efficient training of code LLMs. *arXiv preprint arXiv:2504.14655*, 2025.

Huajian Xin, Z. Z. Ren, Junxiao Song, Zhihong Shao, Wanjia Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiushi Du, et al. Deepseek-prover-v1.5: Harnessing proof assistant feedback for reinforcement learning and monte-carlo tree search. In *International Conference on Learning Representations (ICLR)*, 2025.

Hongling Xu, Qi Zhu, Heyuan Deng, Jinpeng Li, Lu Hou, Yasheng Wang, Lifeng Shang, Ruifeng Xu, and Fei Mi. KDRL: Post-training reasoning LLMs via unified knowledge distillation and reinforcement learning. *arXiv preprint arXiv:2506.02208*, 2025a.

Shiyi Xu, Yiwen Hu, Yingqian Min, Zhipeng Chen, Wayne Xin Zhao, and Ji-Rong Wen. ICPC-Eval: Probing the frontiers of LLM reasoning with competitive programming contests. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025b.

Chuanhao Yan, Fengdi Che, Xuhan Huang, Xu Xu, Xin Li, Yizhi Li, Xingwei Qu, Jingzhe Shi, Chenghua Lin, Yaodong Yang, et al. Re:form – reducing human priors in scalable formal software verification with RL in LLMs: A preliminary study on Dafny. *Transactions on Machine Learning Research*, 2026.

Fanpeng Yang, Xing Li, Shuling Wang, Jie An, Zeyu Sun, Shenghua Feng, Wenhan Wang, Weiyi Wang, Naijun Zhan, and Fanjiang Xu. How powerful are LLMs in generating formal program specifications? In *International Conference on Machine Learning (ICML)*, 2026.

Zitong Yang, Neil Band, Shuangping Li, Emmanuel Candès, and Tatsunori Hashimoto. Synthetic continued pretraining. In *International Conference on Learning Representations (ICLR)*, 2025.

Zhe Ye, Zhengxu Yan, Jingxuan He, Timothe Kasriel, Kaiyu Yang, and Dawn Song. Verina: Benchmarking verifiable code generation. In *International Conference on Learning Representations (ICLR)*, 2026a.

Zhe Ye, Aidan Z. H. Yang, Huangyuan Su, Zhenyu Liao, Samuel Tenka, Zhizhen Qin, Udaya Ghai, Dawn Song, and Soonho Kong. Intent-aligned formal specification synthesis via traceable refinement. *arXiv preprint arXiv:2604.10392*, 2026b.

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. DAPO: An open-source LLM reinforcement learning system at scale. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

Zichao Yu, Chengzhi Yu, Shengze Xu, Yujin Han, Bingqing Jiang, Xu Wang, and Difan Zou. Mismatch matters: On-policy distillation beyond token agreement. *arXiv preprint arXiv:2608.09836*, 2026.

Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. Does reinforcement learning really incentivize reasoning capacity in LLMs beyond the base model? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

Z.ai. Glm-5.2: Built for long-horizon tasks. https://z.ai/blog/glm-5.2, 2026.

Haoyu Zhao, Ziran Yang, Jiawei Li, Deyuan He, Zenan Li, Chi Jin, Venugopal V. Veeravalli, Aarti Gupta, and Sanjeev Arora. Algoveri: An aligned benchmark for verified code generation on classical algorithms. In *International Conference on Machine Learning (ICML)*, 2026.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, et al. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Yichi Zhou, Jianqiu Zhao, Yongxin Zhang, Bohan Wang, Siran Wang, Luoxin Chen, Jiahui Wang, Haowei Chen, Allan Jie, Xinbo Zhang, et al. Solving formal math problems by decomposition and iterative reflection. *arXiv preprint arXiv:2507.15225*, 2025.

A FILE FORMAT

A.1 DECLARATIONS

Each term of the paper is implemented by a declaration (Table 6).

Table 6. Declarations per AUXSPEC problem.

Term	Declaration	Description
Precondition	`isValidInput`	Executable predicate encoding the input constraints
Postcondition	`<algo>_spec`	`Prop` relating input and output
Executable postcondition	`<algo>_check`	`Bool` function of input and output
Postcondition equivalence theorem	`<algo>_check_iff_spec`	`_check` decides `_spec` on valid inputs
Implementation	`<algo>`	Total function computing the output
Correctness theorem	`<algo>_correct`	`<algo>` satisfies `_spec` on valid inputs

A.2 FORMAT

Each file includes precondition, postcondition and executable postcondition with postcondition equivalence theorem, implementation with correctness theorem, and test cases. Helper definitions follow a topological order of dependency. The namespace, the type abbreviations, the printed expressions of the frozen test cases, and both theorem statements are assembled. The test cases are organized into test layers (Table 7).

Table 7. Test layers with the component under test and the error detected.

Layer	Component Under Test	Test Case	Assertion	Error Detected
L1[a]	Precondition	Valid input	Accepts	Overly restrictive
L1[b]		Invalid input	Rejects	Overly permissive
L2[a]	Executable postcondition	Valid input, correct output	Accepts	Incomplete
L2[b]		Valid input, incorrect output	Rejects	Unsound
L3	Implementation	Valid input, expected output	Matches	Incorrect
L4	Executable postcondition on implementation output	Valid input	Accepts	Implementation output rejected

B CONFIGURATION

B.1 TRAINING

Both model sizes are fine-tuned with the settings in Table 8. AUXSPEC-9B-SFT is fine-tuned on an earlier version of the SFT data with 18,489 samples over 4,026 problems. In a round of expert iteration, the model is fine-tuned for 3 epochs on the training samples of that round. Checkpoints are evaluated after each epoch, and the models reported in the main text are the third-epoch checkpoints. Loss is taken on assistant tokens only.

Table 8. Supervised fine-tuning settings.

Setting	9B	27B
Base model	Qwen3.5-9B	Qwen3.8-27B
Hardware	4 GH200	8 GH200
Batch (sequences)	64	
Sequence length (tokens)	16,384	
Learning rate	1e-5, cosine	
Warmup	5%	
Epochs	3	

Table 9. Reinforcement learning settings.

Setting	9B	27B
Start model	AUXSPEC-9B-SFT	AUXSPEC-27B-SFT
Hardware	16 GH200	64 GH200
Prompts per step	32	64
Rollouts per prompt	16	
Prompts per mini-batch	8	
Learning rate	1e-6, constant	
Clip range ϵ_{low} / ϵ_{high}	0.2 / 0.28	
Temperature / top-p	0.7 / 1.0	
Max new tokens	16,384	
Training steps (1 epoch)	434	208

Reinforcement learning follows the settings of Table 9 for 1 epoch of the 13,637 tasks, which takes 434 steps for the 9B model and 208 steps for the 27B model. Rollout and training share the same GPUs, the Lean compilation pool runs on the CPUs of the same nodes, and the LLM judges are called through an external API. Rollouts longer than 12,288 tokens lose reward linearly, up to 1.0

at the limit of 16,384 tokens, following the soft overlong punishment of DAPO (Yu et al., 2025). The optimizer is AdamW with weight decay 0.01 and gradient clipping at 1.0, and the clipped surrogate adds dual clipping with constant 3.0 for negative advantages. The maximum prompt length is 8,192 tokens. A SpecGen rollout needs 12 tests and a CodeGen rollout 3, fewer for problems with fewer test rows. A rollout fails compilation when a declaration does not elaborate, a test cannot be evaluated, compilation exceeds the time limit, the compiler reports an error without a position, or the rollout contains `sorry` outside the two theorem statements, `admit`, `axiom` or `native_decide`. A rollout receives 0.3 when a judge call fails (Table 19). A rollout is repetitive when a 40-character n-gram, taken at a stride of 4 characters, occurs at least 20 times and fewer than 60% of the n-grams are distinct.

B.2 Inference

Expert iteration, the evaluation of our models and the analyses of Appendix E sample with the settings of Table 10. Baselines follow the sampling settings recommended for each model with the same maximum output length, except the long-reasoning models DeepSeek-V4-Flash, DeepSeek-V4-Pro and Kimi-K3, which run without a length limit.

Table 10. Inference settings.

Inference engine	vLLM
Temperature	0.3 (evaluation), 0.7 (expert iteration and analyses)
Top-p	0.95
Max new tokens	16,384
Samples per problem, compile@k and pass@k	4

C Training data

C.1 Source selection

The LeetCode problems combine 2,869 problems of LeetCodeDataset (Xia et al., 2025) with 746 free problems collected from the site after the dataset snapshot. Of the 746 problems, 263 pass a buildability filter that requires at least 3 examples and a Python function signature and excludes database, shell, concurrency and design problems. The seeds are 3,109 problems, 2,864 from LeetCodeDataset and 245 from the site. We take the Codeforces and AtCoder problems from the train split of TACO (Li et al., 2023). For RL, we add problems from the train and test splits of TACO, including Aizu, CodeChef and GeeksforGeeks problems, and 989 Codeforces problems from open-r1/codeforces (Penedo et al., 2025). Problems are admitted with at least 3 tests, a reference solution, and a single accepted answer, since one reference output cannot label a problem with several accepted answers. Stored tests whose inputs violate the constraints of the problem description are removed before use, which drops 1,028 of 90,885 test rows on Codeforces and 2,119 of 22,115 on AtCoder. Difficulty in Table 3 follows the labels of LeetCode and TACO and the ratings of open-r1/codeforces. The 1,484 problems without a label, including all Aizu problems, take the tier that an LLM judge assigns in the conversion audit, where easy problems need direct simulation or a single formula, medium problems a standard technique and hard problems a non-obvious idea or a combination of techniques.

C.2 Problem format

Each Codeforces and AtCoder problem, whose input is given from standard input, is rephrased into a typed function interface and admitted only when the reference solution still passes the stored tests through that interface. Conversion admits 3,490 of the 4,314 Codeforces problems (80.9%). The rejected problems have stored tests that no reference solution reproduces (361), inputs that change under parsing and unparsing (259), several accepted answers (184), or other harness errors (20). Freezing the test cases removes 205 more problems, 201 of which lack 3 invalid inputs for $\mathrm{L1}^{\mathrm{b}}$, and leaves 3,285. Of the 1,440 AtCoder problems, conversion admits 892, an audit for descriptions that drop a stated condition and for examples that violate the stated constraints leaves 677, and freezing leaves 623. Most AtCoder rejections at conversion are inputs that change under parsing and unparsing (352 of 548), which trace to faulty stored tests such as leading zeros or counts that do not match the input.

C.3 Data breakdown

We count the samples of each task for supervised fine-tuning, reinforcement learning, and AuXSpec-Test (Table 11). A reinforcement learning sample is a generation prompt, and repair prompts are built during training from failed rollouts (Section 4.3).

Table 11. Number of samples of each task in each stage.

Task	SFT	RL	AuXSpec-Test
SpecGen	6,528	8,281	300
SpecGen-Repair	3,807	0	131
CodeGen	6,645	5,356	300
CodeGen-Repair	2,512	0	103
Total	19,492	13,637	834

The 19,492 supervised fine-tuning samples are drawn from the synthesis pipeline of Section 3 and expert iteration (Table 12). An expert-iteration repair chain accepted after more than 1 round becomes 1 generation sample with a reasoning trace over the first attempt, feedback and repair.

Table 12. Supervised fine-tuning samples by source.

Source	SpecGen	SpecGen-Repair	CodeGen	CodeGen-Repair	Total
Synthesis pipeline	4,786	2,479	4,475	1,100	12,840
Expert iteration	373	1,328	739	1,412	3,852
Expert iteration, rewritten	1,369	0	1,431	0	2,800
Total	6,528	3,807	6,645	2,512	19,492

The corpus builder also removes samples longer than 16,384 tokens (42), duplicate samples (57), samples beyond 3 per problem and task (11), and 1 sample that is not a whole file.

C.4 Expert iteration details

Each round samples 1 refinement chain per problem and task from a fine-tuned Qwen3.8-27B. A chain generates an output, compiles the output, and repairs the output for up to 8 rounds, or 16 rounds in the last round of expert iteration. A chain is accepted under the criteria of Section 2, and the reasons of each judge rejection enter the next repair round as review text. Expert iteration runs 3 rounds (Table 13). The first 2 rounds sample from the model fine-tuned on the LeetCode data, and the last round from the model fine-tuned on the data of the first 2 rounds and the Codeforces pipeline data. Accepted chains that pass at the first round become generation samples, and later rounds become repair samples or are rewritten into generation samples (Appendix C.3). A scan over all refinement records adds 384 samples from accepted chains that no round collected. The chains on reinforcement learning problems in the last round add no supervised fine-tuning samples.

Table 13. Rounds of expert iteration. A chain is 1 refinement run on 1 problem and task.

Round	Maximum repair rounds	Chains	Accepted chains	SFT samples
LeetCode	8	5,239	2,268	3,610
Codeforces	8	3,778	1,248	2,137
Last round	16	3,208	1,444	521
Scan over all records	–	–	–	384
Total		12,225	4,960	6,652

Expert iteration adds repairs more than first attempts. Of the 2,268 accepted chains of the LeetCode round, 685 pass at the first round, and of the 298 accepted SpecGen chains on reinforcement learning problems in the last round, 290 need at least 1 repair round.

C.5 Decontamination

The training problems pass filters that remove the 300 AuXSpec-Test problems, the problems that Verina links to, and the reinforcement learning problems from supervised fine-tuning. Verina problems are matched through the LeetCode links in the Verina descriptions, where 43 of the 189 problems link to 36 LeetCode problems. The filters remove 21 samples of AuXSpec-Test problems and 166 samples of 32 LeetCode problems that Verina links to, and a recount of the final training data finds no problem of these sets.

D EVALUATION

D.1 METRICS

Compile@k and Pass@k. We follow the standard definition of pass@k (Chen et al., 2021) and define compile@k by analogy. For a problem with k samples, compile@k is 1 when at least one sample compiles, and pass@k is 1 when at least one sample meets the acceptance criteria of Section 2. The reported value is the mean over problems, which equals the unbiased estimator of pass@k (Chen et al., 2021) since exactly k samples are drawn per problem. Test case alignment compares the test lines of a sample with the frozen test cases regardless of layer headings, and Lean code blocks are extracted regardless of the capitalization of the language tag.

D.2 TOKEN EFFICIENCY

We measure output tokens per response and per solved sample (Table 14). AUXSPEC-27B-RL requires the fewest output tokens per solved sample of all 13 models in the 4 generation settings. On SpecGen of AUXSPEC-Test, it requires 80% fewer tokens per solved sample than GPT-5.6-Sol, the most economical frontier model (2.7K vs. 13.2K), while its pass@1 is 35.9 points higher (65.7% vs. 29.8%; Table 4). RL lowers the cost per solved SpecGen sample at both model sizes, as pass@1 grows faster than the response length (19.3K to 2.7K tokens for the 27B model and 36.0K to 11.0K for the 9B model).

RL changes the length of the reasoning trace in opposite directions at the 2 model sizes. Reasoning traces of the 27B model become shorter after RL, whereas reasoning traces of the 9B model become longer (1.6K to 0.4K and 1.7K to 4.6K thinking tokens per SpecGen response). A plausible reading is that the larger model needs less explicit reasoning to state a specification, whereas the smaller model relies on longer reasoning to reach a correct output. Thinking tokens are the tokens of the thinking segment. Tokens of our models and the backbones are counted with the tokenizer of the backbone, and tokens of the frontier models are the output tokens reported by the API, thinking included.

Table 14. Average output tokens per response (Resp.) and output tokens per solved sample (Solved), the total output tokens of a task divided by the samples that meet the acceptance criteria, both in thousands. Responses are capped at 16,384 tokens for our models, the backbones and the Claude and GPT models, at 32,768 for Kimi-K3 and at 65,536 for the DeepSeek models. Underline indicates the lowest value of each column within each group, and **bold** the lowest overall.

	AUXSPEC-Test				Verina			
	SpecGen		CodeGen		SpecGen		CodeGen	
Model	Resp.	Solved	Resp.	Solved	Resp.	Solved	Resp.	Solved
Frontier Models								
Claude-Opus-5	9.7	16.8	4.8	5.2	5.2	6.0	3.9	6.9
Claude-Sonnet-5	10.2	74.4	7.3	13.4	4.4	6.9	9.6	53.7
GPT-5.6-Sol	3.9	13.2	2.8	3.5	2.6	3.0	3.1	14.7
GPT-5.6-Luna	3.4	36.0	2.8	7.0	2.7	4.0	2.5	10.1
Kimi-K3	10.7	24.4	6.7	10.0	5.2	6.7	9.0	25.7
DeepSeek-V4-Pro	27.0	126.9	24.1	49.5	17.1	31.8	20.8	55.2
DeepSeek-V4-Flash	49.6	129.2	49.0	109.0	37.8	62.9	43.2	741.4
Instruct Models								
Qwen3.5-9B	3.2	1,912.1	6.8	82.9	2.6	37.6	2.3	9.7
Qwen3.8-27B	13.4	296.8	15.9	178.6	8.1	28.0	15.5	469.8
Ours								
AUXSPEC-9B-SFT	3.6	36.0	2.2	16.6	4.0	10.6	1.5	4.0
AUXSPEC-9B-RL	6.4	11.0	4.1	5.3	3.9	4.5	1.9	2.0
AUXSPEC-27B-SFT	3.6	19.3	1.9	7.0	3.8	6.7	1.4	2.5
AUXSPEC-27B-RL	**1.8**	**2.7**	**1.7**	**1.8**	**1.6**	**1.7**	**1.1**	**1.2**

D.3 GENERAL CAPABILITIES

To assess general capabilities after specialization, we evaluate the backbones, SFT models and RL models on AIME 25 (Mathematical Association of America, 2025), GPQA-Diamond (Rein et al., 2024), and HumanEval+ and MBPP+ of EvalPlus (Liu et al., 2023). EvalPlus extends HumanEval (Chen et al., 2021) and MBPP (Austin et al., 2021) with more test cases per problem. Sampling

follows the settings of DeepSeek-R1 (Guo et al., 2025) with temperature 0.6, top-p 0.95 and at most 32K output tokens, and pass@1 averages 64 samples per problem on AIME 25, 8 on GPQA-Diamond and 4 on HumanEval+ and MBPP+.

Supervised fine-tuning lowers all 4 scores of both backbones, most on AIME 25, by 36.7 points for the 9B model and 53.1 for the 27B model (Table 15). RL recovers 72% of this drop on AIME 25 at both sizes and about half of the drop on GPQA-Diamond (47% and 52%). On the 2 code benchmarks, AUXSPEC-9B-RL surpasses its backbone (88.6% vs. 86.6% on HumanEval+ and 70.8% vs. 65.3% on MBPP+), and AUXSPEC-27B-RL trails its backbone by 3.2 and 1.5 points.

Table 15. General capabilities of the backbones, SFT models and RL models (pass@1, %). Numbers in parentheses show the change in points from the backbone of the same size.

Model	AIME 25	GPQA-Diamond	HumanEval+	MBPP+
Qwen3.5-9B	62.3	80.2	86.6	65.3
AUXSPEC-9B-SFT	25.6 (−36.7)	62.2 (−18.0)	63.9 (−22.7)	59.9 (−5.4)
AUXSPEC-9B-RL	52.2 (−10.1)	70.7 (−9.5)	88.6 (+2.0)	70.8 (+5.5)
Qwen3.8-27B	84.1	84.5	92.5	80.4
AUXSPEC-27B-SFT	31.0 (−53.1)	62.9 (−21.6)	79.0 (−13.5)	71.1 (−9.3)
AUXSPEC-27B-RL	69.3 (−14.8)	74.2 (−10.3)	89.3 (−3.2)	78.9 (−1.5)

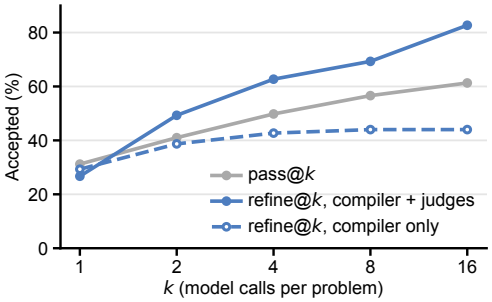
D.4 SELF-REPAIR ON ICPC-EVAL

Besides generation, AUXSPEC models are trained on the repair tasks (Section 2), which revise a rejected output from compiler and judge feedback. To assess the value of this ability, we evaluate AUXSPEC-27B-RL on ICPC-Eval (Xu et al., 2025b), a benchmark of competitive programming problems whose specifications require careful definitions. Of its 118 problems, we exclude 28 problems with multiple correct outputs, which a fixed expected output cannot judge, and 15 problems whose test cases exceed the length range of AUXSPEC-Test, which leaves 75 problems. Following Refine@K of ICPC-Eval and the comparison of self-repair with i.i.d. sampling at an equivalent budget (Olausson et al., 2024), we compare 3 settings within a budget of k model calls per problem:

- **Pass@k** draws k independent samples and is estimated from 16 samples.
- **Refine@k with compiler and judge feedback** draws 1 sample and returns each rejected output with the compiler messages or the judge findings, in the format of the repair tasks, until an output is accepted or k calls are made.
- **Refine@k with compiler feedback only** follows the same chain and stops at the first output that compiles.

All settings sample at temperature 0.6 and top-p 0.95 with at most 32,768 generated tokens, and the acceptance criteria (Section 2) score every output.

With compiler and judge feedback, refine@16 reaches 82.7%, compared with 61.3% for pass@16 at an equivalent budget (Figure 3). With compiler feedback only, more chains compile within 2 rounds than independent samples do (88.0% vs. 79.3% for compile@2), but refine@16 reaches 44.0%, as the chain stops at the first output that compiles and the 42 outputs that the judges reject are never revised.



k	1	2	4	8	16
Pass@k	31.2	41.0	49.8	56.6	61.3
Refine@k, compiler + judges	26.7	**49.3**	**62.7**	**69.3**	**82.7**
Refine@k, compiler only	29.3	38.7	42.7	44.0	44.0

Figure 3. SpecGen results of AUXSPEC-27B-RL on 75 problems of ICPC-Eval, as the percentage of problems with an accepted output within k model calls per problem. Refine@k returns compiler and judge feedback (compiler + judges) or compiler messages only (compiler only). Bold indicates refine@k values above pass@k.

E TRAINING METHOD ANALYSIS

E.1 EXPERT ITERATION

SpecGen pass@4 rises at each stage of the SFT data (Table 16), from 30.7% on the LeetCode data to 39.3% after adding the first round of expert iteration and 47.3% after further adding the Codeforces problems and 2 more rounds.

Table 16. Pass@4 on the in-domain test set after each stage of the SFT data, sampled with the analysis settings of Table 10. Each stage adds data to the previous stage and is fine-tuned from Qwen3.8-27B for 3 epochs.

Training data	Problems	Samples	SpecGen	CodeGen
LeetCode	2,687	7,993	30.7	38.0
+ expert iteration round 1	2,733	11,811	39.3	52.0
+ Codeforces, expert iteration rounds 2 and 3 (AUXSPEC-27B-SFT)	3,993	19,492	47.3	48.7

E.2 CONTINUED PRETRAINING

Motivation. An early supervised fine-tuning run of Qwen3.5-9B compiles 5.05% of the 1,208 SpecGen samples and 3.8% of all 3,596 samples on the held-out problems. Classifying the first compiler error of the failed samples shows that 40.8% misuse Mathlib declarations, through application type mismatches (17.3%), invalid fields or projections (14.7%) and type mismatches (8.7%), and 65.2% once unknown identifiers (11.6%) and missing instances (12.8%) are counted. With retrieved declaration names in 58% of the answers, the failures come from the use of Mathlib.

Corpus. Synthesis augments prompts with Mathlib declarations retrieved for the problem (Appendix G.1.2), and across the 3,109 LeetCode problems the retrievals cover 1,381 distinct declarations, of which 1,298 are found in the pinned Mathlib source. These declarations come from 411 files, and we take all declarations of those files rather than the retrieved ones alone, which yields 26,936 declaration heads. The heads are rewritten with DeepSeek-V4-Flash (DeepSeek-AI, 2026) into English lessons that state the content, the application, the occasion and the contrast of a declaration, which follows the rewriting of raw corpora into reading-comprehension text (Cheng et al., 2024) and textbook-style synthetic corpora (Gunasekar et al., 2023). Together with 8,312 rows of Mathlib source, 2,336 lessons on decidability routes and 3,790 pairs of compiler error and repair, the corpus $\mathcal{M}$ holds 41,374 rows and about 43M tokens.

Training and result. Adding problems to supervised fine-tuning improves pass@4 more than continued pretraining does (Table 17). Continued pretraining minimizes the next-token loss on the corpus,

$$\mathcal{L}_{\text{CPT}}(\theta) = -\,\mathbb{E}_{u\sim\mathcal{M}} \sum_{t} \log \pi_\theta(u_t \mid u_{<t}), \tag{3}$$

for 2 epochs, 92 steps, on Qwen3.5-9B base with learning rate 5e-6 and sequence length 16,384, and supervised fine-tuning follows on 4,826 rows from 1,631 problems. Domain corpora of this kind help on downstream tasks (Gururangan et al., 2020) and formal provers are pretrained on Lean and mathematics before fine-tuning (Azerbayev et al., 2024; Xin et al., 2025), yet the benefit of continued pretraining is not universal (Shi & Lipani, 2023). Here continued pretraining improves SpecGen pass@4 from 8.2 to 9.8 and pass@4 over the 4 tasks from 26.4 to 27.4, whereas adding 1,056 problems, 3,167 rows, to supervised fine-tuning improves the same 2 numbers to 13.7 and 30.2. The corpus loss falls from 1.03 to 0.69 with held-out perplexity 2.03. The corpus is therefore fit while the downstream gain stays at 1.0 point, the distance between document perplexity and downstream use reported for knowledge injection through continued pretraining (Jiang et al., 2024). Content memorized from a corpus becomes usable after many representations of the same content (Allen-Zhu & Li, 2024; Yang et al., 2025), and each of the 26,936 declarations appears in 1 lesson here. Lessons state the meaning of declarations, whereas a compiling specification requires declarations composed into terms that the elaborator accepts, and added problems carry test cases and judges that label samples of that composition. Later rounds therefore put the effort into data.

Table 17. Continued pretraining compared with more fine-tuning data for Qwen3.5-9B, pass@k with 4 samples on an earlier version of the training data (4,826 rows from 1,631 problems). SFT, more data adds 3,167 rows from 1,056 problems. Parentheses show the change from SFT.

Model	AuXSpec-Test						Verina SpecGen	
	SpecGen		CodeGen		All 4 tasks		Basic + Advanced	
	p@1	p@4	p@1	p@4	p@1	p@4	p@1	p@4
Qwen3.5-9B	0.0	0.0	0.0	0.0	2.4	7.3	2.0	7.0
SFT	2.0	8.2	4.3	15.6	13.6	26.4	16.0	35.4
CPT + SFT	2.3 (+0.3)	9.8 (+1.6)	5.9 (+1.6)	16.0 (+0.4)	14.1 (+0.5)	27.4 (+1.0)	16.0 (0.0)	37.6 (+2.2)
SFT, more data	7.0 (+5.0)	13.7 (+5.5)	7.0 (+2.7)	18.8 (+3.2)	18.3 (+4.7)	30.2 (+3.8)	18.3 (+2.3)	43.9 (+8.5)

E.3 Reinforcement learning

The final recipe for AuXSpec-9B-RL and AuXSpec-27B-RL combines the progressive reward of Eq. (1), the self-repair buffer, the overlong penalty and the mask on repetitive rollouts on the clipped surrogate of DAPO with batch-normalized advantages (Section 4.3), and adds the end-of-sequence penalty for the 27B model. The runs on the 9B model develop the recipe, and both models train for 1 epoch, 434 steps for the 9B model and 208 steps for the 27B model, with the settings of Table 9. Both runs remain stable through the epoch, with 0.12% and 0.60% of rollouts repetitive and 0.13% and 0.11% malformed for the 27B and 9B models. The reported models are the final checkpoints, and SpecGen pass@1 on AuXSpec-Test still rises over the last evaluated checkpoints, from 54.3% at step 400 to 58.6% for the 9B model and from 61.3% at step 200 to 65.7% for the 27B model. The subsections below analyze the reward stages, the reward dynamics and pass@k.

E.3.1 Reward stages

The progressive reward keeps most groups without a passing rollout in the gradient. Under a binary reward these groups have equal rewards and are discarded by dynamic sampling (Yu et al., 2025). They make up 37% of SpecGen groups for the 27B model and 47% for the 9B model, and the progressive reward keeps nonzero variance in 87.9% and 86.2% of them (Table 18).

Table 18. Generation groups with nonzero reward variance under the progressive reward and under a binary reward that requires all stages, scored on the rollouts of the whole run of each model. A binary reward has zero variance in every group without a passing rollout.

	27B		9B	
	SpecGen	CodeGen	SpecGen	CodeGen
Groups	8,346	5,384	7,461	5,279
Nonzero variance, progressive (%)	91.6	59.9	91.3	73.0
Nonzero variance, binary (%)	59.4	41.9	50.9	49.3
Groups without a passing rollout	3,048	1,309	3,497	1,743
nonzero variance, progressive (%)	87.9	74.3	86.2	71.8

Table 19 lists the reward and the fraction of rollouts of each outcome. Each rule against reward hacking closes a shortcut that the stages would otherwise reward. A test that calls no declaration of the output, such as `#guard true`, passes without testing the output, and counting only tests that call a declaration keeps the test stages meaningful. An escape hatch, the `sorry` and `admit` tactics, an `axiom` declaration or the `native_decide` tactic, lets an output compile without the checks that later stages assume. Special-casing the test inputs, excluding a test value from the validity predicate and copying the given test cases into the implementation all pass the test stages without solving the problem and receive 0. Altering a given test case lets a rollout pass the test stages against a different expected output and receives 0.1. A test that cannot run would otherwise earn the test execution credit of a test that ran and failed. An unchanged repair would keep the credit of the stages the failed input already passes. The non-tautology stage precedes the overall quality stage, the order in which evaluation runs the judges.

The partial credit separates accepted outputs by form. The non-tautology judge classifies a specification as declarative, enumerative or tautological. A declarative specification states the condition of the problem and is sound and complete: every answer that satisfies the specification is correct, and every correct answer satisfies the specification. An enumerative specification states the same condition over a candidate list that the specification constructs, such as all sublists of the input, instead of

quantifying over the candidate type. A tautological specification replaces the condition with an algorithm for the answer, which is the role of the implementation. The implementation judge marks a correct CodeGen implementation as slow when the implementation takes more than about 10^7 steps on the largest admitted input and a method faster by more than a logarithmic factor exists, known as a slow positive (Li et al., 2022). Appendix K.3 lists both prompts.

Table 19. Reward and fraction of rollouts (%) for each outcome of the progressive reward in SpecGen (left) and CodeGen (right), each with the repair task, over the whole run of each model. [1]Only in the repair tasks, which otherwise share the outcomes and values, with the reward of a repair rollout scaled down from 1 in the first round to 0.8 in the last.

SpecGen				CodeGen			
Outcome	Reward	27B	9B	Outcome	Reward	27B	9B
Malformed format or no code block	0	0.30	1.70	Malformed format or no code block	0	0.22	1.17
Special-cased or copied test cases	0	0.17	0.08	Special-cased or copied test cases	0	0.66	1.41
Unchanged repair[1]	0	0.36	0.67	Unchanged repair[1]	0	0.73	0.59
Fewer tests than required	0.1	0.61	0.52	Fewer tests than required	0.1	0.19	0.18
Fails compilation	0.1	29.16	38.70	Fails compilation	0.1	12.42	17.37
Fails test case alignment	0.1	0.96	0.68	Fails test case alignment	0.1	0.21	0.35
Fails test execution	0.2	17.01	18.12	Fails test execution	0.2	10.72	12.04
Fails non-tautology	0.1	7.67	8.46	No verdict from a judge after retries	0.3	0.00	0.06
No verdict from a judge after retries	0.3	0.00	0.02	Fails overall quality	0.4	0.31	0.57
Fails overall quality	0.4	17.18	10.95	Correct but slow implementation	0.5	21.94	37.67
Enumerative specification	0.4	1.89	2.14	Passes all stages	1.0	52.58	28.58
Passes all stages	1.0	24.68	17.96				
Rollouts		251,232	224,528	Rollouts		112,000	111,792

As training proceeds, compilation failures fall from 49.0% of rollouts in the first 40 steps to 15.5% in the last 40 steps for the 27B model and from 66.0% to 10.7% for the 9B model, while passing rollouts rise from 18.4% to 45.0% and from 9.9% to 38.5% (Figure 2c and 2d).

E.3.2 Reward dynamics

We analyze the group-relative advantage (Shao et al., 2024) before batch normalization. Replaying the training rollouts with the value of one outcome changed shows how the value sets this advantage (Figure 4a1 and 4a2). Failures that attempt to score unfaithfully are shortcuts rather than progress. For instance, tautological specifications restate algorithms instead of specifying problems, and failed test case alignments usually rewrite the given test cases. Such unfaithful attempts should be discouraged in every group. Higher values, though well below the full reward of 1, still assign these failures positive advantages in groups without correct rollouts, whereas replaying at 0.1 shows advantages that are negative and below those of faithful attempts in almost every group (Table 20).

Table 20. Advantage of shortcut failures at the final value and at an alternative higher value, replayed on the training rollouts. For failed test case alignments, the advantage is given as the fraction of SpecGen-Repair rollouts with a positive advantage.

Shortcut Failure	Final Value	Advantage	Alternative Value	Advantage
Tautological specification	0.1	Negative in every group	0.3	Positive in groups without a full-score rollout
Failed test case alignment	0.1	Positive for 0.8% of rollouts	0.2	Positive for 58.3% of rollouts

Based on observation, some attempts make partial progress that is correct in substance but not yet in the perfect form. For instance, 1) enumerative specifications already state the correct condition and only need to be rewritten declaratively, and 2) slow implementations are already correct and only need faster algorithms. Binary judges also classify enumerative specifications poorly. An intermediate reward for partial progress refines sparse signals into fine-grained rewards. When rewards are low, like SpecGen in the early steps, partial progress rewards small improvements of failing attempts, while once correct rollouts are common, declarative specifications and fast implementations are assigned higher advantages.

By replaying the rollouts, we fix intermediate values at which partial progress receives positive advantages in groups with few correct rollouts and negative ones in groups with many. A slow implementation at 0.5 turns negative from 2 passing rollouts of 16, compared with 6 at 0.8, and an enumerative specification at 0.4 turns negative from 3, compared with 14 when valued like a full-score answer. As the policy improves, the advantage of partial progress turns from positive to negative (Figure 4b2 and 4c2). CodeGen groups without a passing rollout fall from 44% in steps 1

to 40 to 0% in steps 161 to 208, and the mean advantage of slow implementations falls from $+0.10$ to -0.17 over the same periods. SpecGen keeps 29% to 52% of groups without a passing rollout in each period, where quality rejections keep a mean advantage near 0 with both signs (Figure 4c1). Partial progress thus provides intermediate supervision that keeps the reward informative throughout training (Table 21).

Table 21. Partial progress during RL. The advantage of partial progress turns negative from the listed fraction, which refers to full-score rollouts in a whole group of 16, replayed on the 27B rollouts. The effect of slow implementations comes from the 27B run, with fractions over 20-step windows of CodeGen rollouts. A repair round from an enumerative specification succeeds when the repaired output is a declarative specification that passes all stages, counted over all repair rounds started from enumerative specifications (336 for the 27B model and 480 for the 9B model).

Partial Progress	Reward	Negative Advantage From	Effect
Enumerative specification	0.4	3/16	Repair rounds turn 74.1% (27B) and 54.6% (9B) of enumerative specifications into declarative ones that pass all stages.
Slow implementation	0.5	2/16	Slow implementations rise from 14.5% of correct CodeGen rollouts in steps 1 to 20 to 67.4% in steps 81 to 100 and fall to 0.4% by the end of the 27B run.

Whether a partial value acts as credit depends heavily on the capability of the policy to sample better rollouts in the same group. For the 9B model, 71.7% of correct but slow CodeGen rollouts fall in groups without a fast correct rollout, where the slow value keeps a positive advantage, compared with 48.4% for the 27B model, and the 9B run does not start repair rounds from slow rollouts. CodeGen pass@1 of AUXSPEC-9B-RL accordingly drops from 77.1% to 44.7% when slow programs count as failures, compared with 94.8% to 88.0% for AUXSPEC-27B-RL.

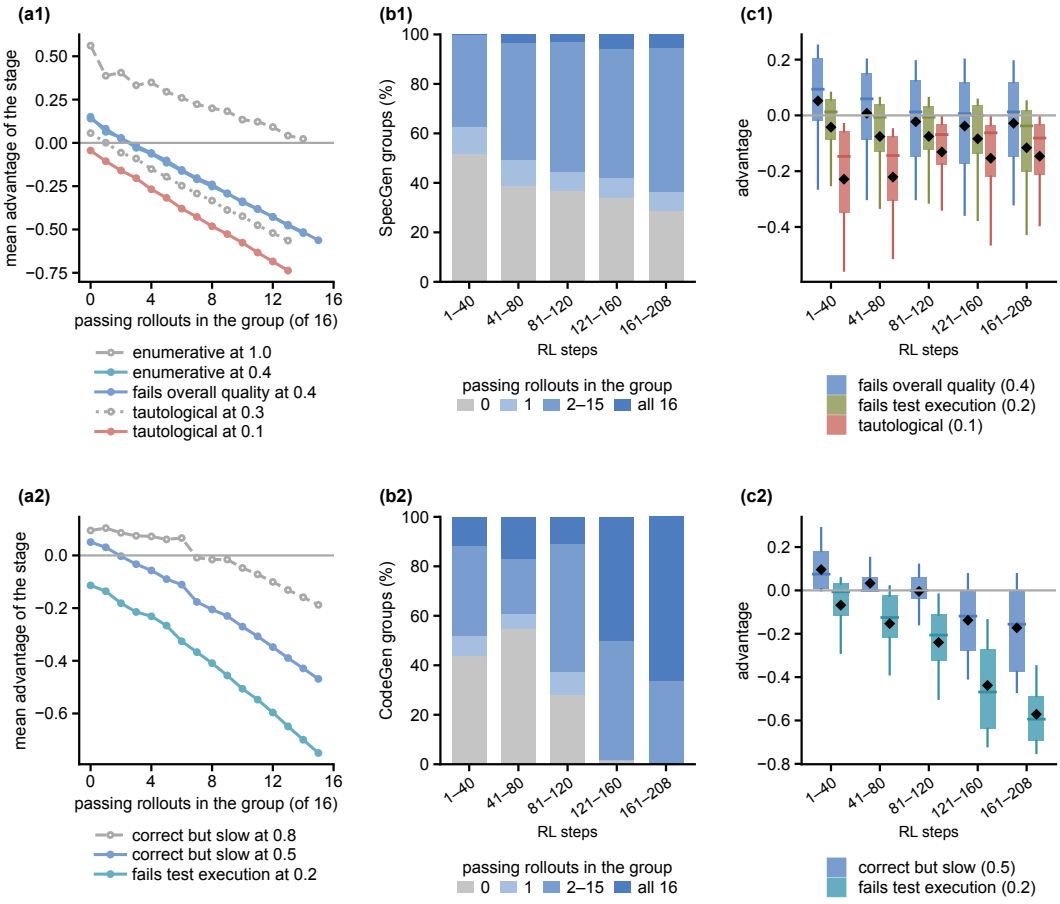


Figure 4. Advantage of stages in the 27B run, SpecGen in the top row (a1 to c1) and CodeGen in the bottom row (a2 to c2). (a) Mean advantage against the number of passing rollouts in a group of 16, replayed on the same rollouts with the value of one stage changed. (b) Groups by the number of passing rollouts in each period of training. (c) Advantage of each stage per period under the final values, with boxes at the 25th and 75th percentiles, whiskers at the 10th and 90th percentiles and diamonds at the mean.

E.3.3 PASS@k

A key question for RL with verifiable rewards is whether it expands the set of problems a model can solve or only raises the probability of solutions that the starting model already samples. Pass@k at large k separates the two, since a starting model that catches up at large k solves the same problems with more samples (Yue et al., 2025). Prolonged RL has been reported to expand the set most where the starting model is weak (Liu et al., 2025).

We compare AUXSPEC-27B-RL with AUXSPEC-27B-SFT, the starting model of RL, at temperature 0.6 and top-p 0.95. Pass@k follows the unbiased estimator (Chen et al., 2021) on 64 samples per problem for the SFT model and 32 for the RL model. Pass@64, pass@128 and pass@256 count a problem as solved when any of 64, 128 or 256 samples passes, and additional samples are drawn only for problems not yet solved.

The RL model stays above the SFT model at every k and reaches 96.3%, compared with 85.7% at $k = 256$, where the SFT model remains below the pass@8 of the RL model (88.2%, Figure 5). The SFT model still gains at $k = 256$, while the gain per doubling of k falls from 5.8 points at $k = 64$ to 4.3 and 2.7 points at $k = 128$ and $k = 256$. The gap narrows from 13.0 to 10.6 points between $k = 128$ and $k = 256$ as the RL model approaches the ceiling, with 11 unsolved problems of which 10 are solved by neither model.

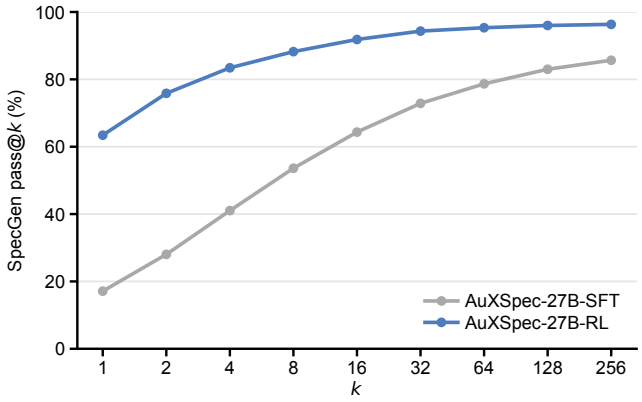


Figure 5. SpecGen pass@k on AuXSpec-Test of AuXSpec-27B-SFT and AuXSpec-27B-RL.

As a result, 33 problems are solved only by the RL model and 1 only by the SFT model within 256 samples (Table 22).

Table 22. SpecGen problems of AuXSpec-Test solved within 256 samples by AuXSpec-27B-SFT and AuXSpec-27B-RL.

	Problems	Fraction (%)
Solved by both models	256	85.3
Solved only by AuXSpec-27B-SFT	1	0.3
Solved only by AuXSpec-27B-RL	33	11.0
Solved by neither model	10	3.3

Appendix E.4 reports RL with a distillation term from AuXSpec-27B-SFT.

E.4 On-policy distillation

To transfer the capability of AuXSpec-27B-SFT, the strongest of our models on AuXSpec-Test, to the smaller AuXSpec-9B-SFT, we apply on-policy distillation (OPD) and OPD within reinforcement learning (Agarwal et al., 2024; Xu et al., 2025a), but neither improves the student model. OPD, with the k3 estimator in the advantage, suffers from mode collapse, where the student model concentrates on the decidable form of the executable postcondition, leading to a compile rate of 2.0% without compiler feedback. SpecGen pass@4 drops from 24.3 to 5.3 while CodeGen changes by at most 3 points. Distillation within reinforcement learning suffers from teacher redundancy, where a teacher model fine-tuned on the same data carries 4% to 7% of the advantage with little information about which samples succeed. We next describe the setup, the failure patterns and the attribution.

Setup. On-policy distillation trains student models on self-generated responses with token-level supervision from teacher models (Agarwal et al., 2024). We select AuXSpec-27B-SFT as the teacher model π_T, the strongest of our models on AuXSpec-Test (Table 4), and AuXSpec-9B-SFT as the student model π_θ, which has the same vocabulary and chat template as the teacher model for a locally served teacher model to score student tokens. For a prompt x and a response $y \sim \pi_\theta(\cdot \mid x)$, the per-token log-ratio

$$\ell_t = \log \pi_\theta(y_t \mid x, y_{<t}) - \log \pi_T(y_t \mid x, y_{<t}) \tag{4}$$

is the k1 estimate of the reverse KL divergence $\mathrm{KL}(\pi_\theta \,\|\, \pi_T)$, and $k_3(t) = e^{-\ell_t} - 1 + \ell_t$ is a non-negative estimate of the same divergence (Schulman, 2020). Both settings enter the policy gradient through the per-token advantage A_t.

- **OPD**: the per-token advantage is $A_t = -k_3(t)$ with no task reward, as in on-policy distillation (Agarwal et al., 2024; Lu & Thinking Machines Lab, 2025), over 200 steps on 12,791 prompts (1 epoch), with student responses and teacher log-probabilities at temperature 1.0.
- **OPD within RL**: the objective adds a distillation term with per-token advantage $-\ell_t$, weighted by $\beta = 0.1$, to the reinforcement learning objective of Appendix E.3, as in joint distillation and reinforcement learning (Xu et al., 2025a) with the k1 estimator in place of k2. Variants keep the distillation term only where the sign agrees with the task advantage (Lin et al., 2026) or drop

the term and reweight the task advantage by clipped teacher ratios (Wang et al., 2026a). Student responses are sampled at temperature 0.7 and teacher log-probabilities at 1.0.

Diagnostics. Besides pass@k, we track token entropy, the fraction of generated text inside repetition loops, the fraction of SpecGen samples in the decidable form of Section 3.1, and teacher informativeness, the mean of $-\ell_t$ on sequences with positive task advantage minus the same mean on sequences with negative advantage (Wang et al., 2026b).

Results. OPD degrades SpecGen and changes CodeGen by at most 3 points (Table 23). SpecGen pass@4 drops from 24.3 to 5.3 on AUXSPEC-Test and from 42.0 to 18.5 on Verina advanced problems, and CodeGen pass@4 from 30.7 to 27.7. In OPD within RL, teacher informativeness stays at +0.004 or below and the teacher term carries 4% to 7% of the total advantage over the first steps. The 3 runs ended at steps 4 to 9 when the judge service ran out of budget, and the measured informativeness led us to stop this line rather than resume.

Table 23. On-policy distillation of AUXSPEC-9B-SFT from AUXSPEC-27B-SFT, pass@k with 4 samples. Parentheses show the change from AUXSPEC-9B-SFT, and the teacher row is a reference. *An earlier checkpoint of AUXSPEC-27B-SFT, used as the teacher.

Model	AUXSPEC-Test				Verina SpecGen			
	SpecGen		CodeGen		All		Advanced	
	p@1	p@4	p@1	p@4	p@1	p@4	p@1	p@4
AUXSPEC-9B-SFT	9.6	24.3	13.5	30.7	30.3	61.4	17.6	42.0
+ OPD	2.3 (−7.3)	5.3 (−19.0)	15.2 (+1.7)	27.7 (−3.0)	29.1 (−1.2)	43.9 (−17.5)	8.3 (−9.3)	18.5 (−23.5)
AUXSPEC-27B-SFT* (teacher)	20.4	47.0	23.1	44.3	49.1	82.5	41.4	79.0

Phenomena. After OPD, 1,147 of 1,200 SpecGen samples (95.6%) fail at compilation, compared with 1,032 (86.0%) before OPD and 875 (72.9%) for the teacher (Figure 6), and OPD also narrows the sample distribution and lets repetition loops grow unpenalized.

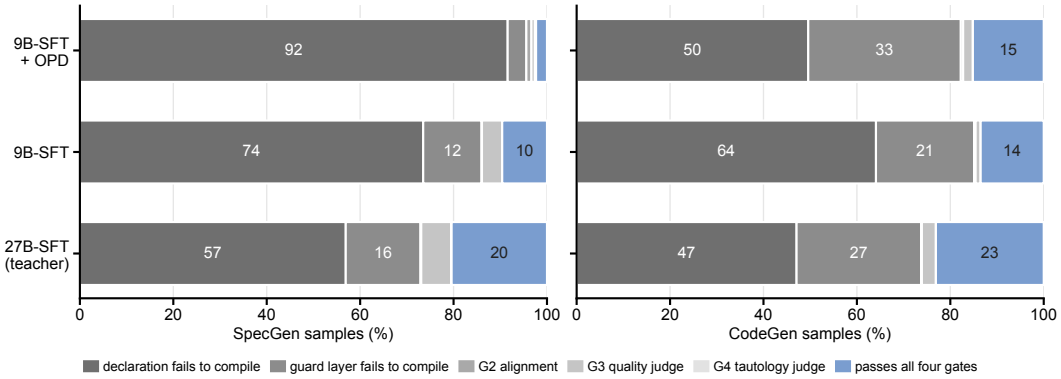


Figure 6. Distribution of SpecGen and CodeGen samples by failure stage, 1,200 samples per model and task. After OPD, 92% of SpecGen samples stop at a declaration that does not compile, while CodeGen samples shift from declaration errors to guard-layer errors with an unchanged pass rate.

Template concentration. The number of SpecGen samples stating the executable postcondition in the decidable form increases from 150 (12.5%) before OPD to 706 of 1,200 (58.8%) after OPD, compared with 114 (9.5%) for the teacher (Figure 7). Within this form, the compile rate falls from 10.7% to 2.0%, compared with 21.9% for the teacher, and failed samples report a missing decidability instance.

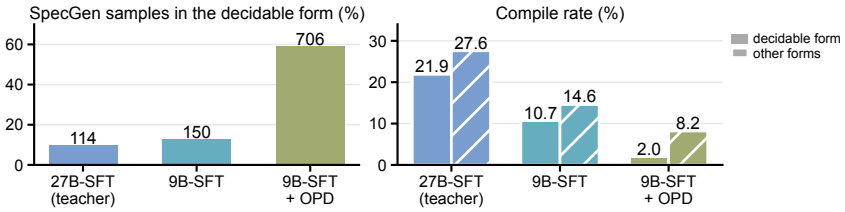


Figure 7. Use of the decidable form in SpecGen samples, with counts out of 1,200 above the bars (left), and compile rates of the decidable form against other forms (right).

Diversity collapse. Token entropy falls from 0.451 to 0.012 (−97.3%) within the first 100 steps (Figure 9), the entropy collapse also reported for reinforcement learning (Yu et al., 2025; He et al., 2025). Near-duplicate sample pairs occur on 40.7% of SpecGen problems, compared with 0.3% before OPD (Figure 8). On Verina basic problems, pass@1 rises from 39.8 to 44.7 while pass@4 falls from 75.9 to 63.0, the pattern of diversity collapse in which gains at small k are paired with losses at larger k (Dang et al., 2025; Yue et al., 2025).

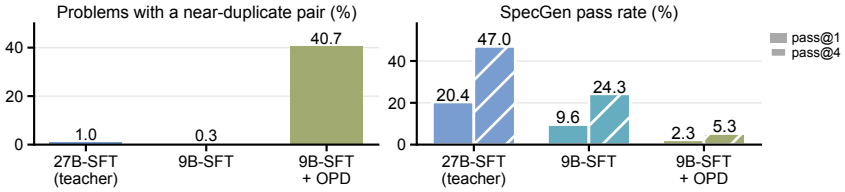


Figure 8. Near-duplicate SpecGen samples per problem (left) and SpecGen pass@1 against pass@4 (right).

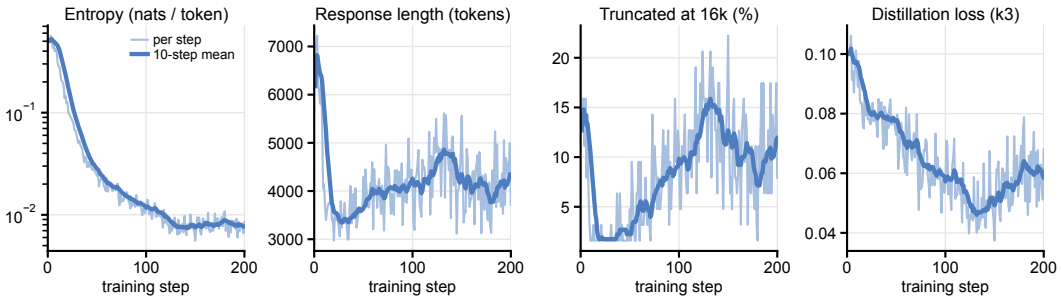


Figure 9. Training dynamics of OPD: token entropy, mean response length, fraction of responses truncated at 16,384 tokens, and distillation loss. Light lines show single steps and dark lines a moving average over 10 steps.

Degenerate agreement. Repetition loops (Holtzman et al., 2020) grow during training, and from step 51 on, looping responses hold 43% to 51% of the generated text. Inside a loop, the teacher ranks the next student token first on 98.9% of tokens, compared with 90.7% before the loop, and the k3 loss per token therefore drops from 0.136 to 0.019 (−86.0%, Figure 10). The same degenerate agreement is reported for on-policy distillation (Yu et al., 2026). Loops account for 1 of the 60 SpecGen problems lost to OPD, where all 4 samples are truncated or loop.

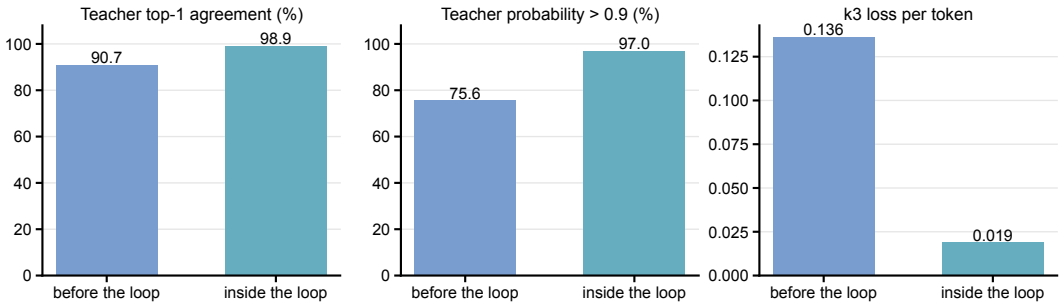


Figure 10. Degenerate agreement inside repetition loops: teacher top-1 agreement, fraction of tokens with teacher probability above 0.9, and k3 loss per token, on 26 looping SpecGen outputs.

Attribution. We attribute the SpecGen drop to estimator bias, which lowers entropy throughout training, to mode collapse, which moves probability onto a form that often fails to compile, and to teacher redundancy, which makes the teacher term uninformative about sample success.

1. *Estimator bias.* With k_3 inside the advantage, the policy gradient is a biased estimate of the reverse KL gradient that also raises the forward KL, and this configuration is reported to collapse training (Shah et al., 2026). The entropy decline of Figure 9 matches this bias.
2. *Mode collapse.* The pipeline sets the decidable form as the first executable postcondition and replaces the form through instance synthesis or a separate definition when compilation fails (Section 3.1), and the form covers 35.5% of synthesized problems as a result. Evaluation scores single samples without compiler feedback. We conjecture that the mode-seeking reverse KL (Agarwal et al., 2024; Jin et al., 2026) moves probability onto this frequent form, including on problems where the form does not compile. The student adopts the form in 54.6% of samples on problems where the teacher never adopts it. OPD therefore amplifies the form beyond teacher usage.
3. *Teacher redundancy.* Teacher models with higher scores help only when offering capabilities beyond the experience of the student (Li et al., 2026a). Teacher and student are fine-tuned on the same data, and although OPD within RL adds the fine-grained reward of reinforcement learning, the teacher term stays uninformative about which samples succeed.

Identical tokenizers and chat templates rule out a format mismatch, and on problems the teacher fails, the student stops at the failing stage of the teacher no more often after OPD than before.

F Ablation Study

F.1 Shared backbone with VeriSpecGen

To compare with VeriSpecGen on a shared backbone, we fix the backbone to Qwen3-4B-Instruct-2507, the backbone of VeriSpecGen-4B, and keep the rest of our pipeline: the same SFT data and configuration (19,492 samples, 3 epochs), the RL recipe of AuXSpec-27B (Section 4.3; 64 problems and 16 samples per step), and the same evaluation protocol (Section 5.3), with the official Verina evaluator at temperature 0.3.

With our prompts and acceptance criteria, the backbone before training scores 0.0% SpecGen pass@1 on both AuXSpec-Test and Verina (Tables 24 and 25). None of its 1,200 SpecGen samples on AuXSpec-Test compiles: the outputs either reach the generation length limit or contain Lean errors such as Lean 3 syntax. VeriSpecGen reports 14.7% for the same backbone with its own prompt and evaluation harness, so the 2 base rows are not directly comparable.

SFT on our data improves generation only slightly (6.9% SpecGen pass@1 on Verina without test cases, close to 9.5% for VeriSpecGen-4B SFT Spec-Only), while the repair tasks improve the most (from 6.5% to 49.0% SpecGen-Repair pass@1 on AuXSpec-Test).

Table 24. In-domain evaluation of the shared-backbone model on AuXSpec-Test, with the protocol of Table 4. The model starts from Qwen3-4B-Instruct-2507, the backbone of VeriSpecGen-4B, and is trained with our SFT data and RL recipe. Base Model indicates Qwen3-4B-Instruct-2507 before training, evaluated with our prompts. VeriSpecGen does not release its models, so the table has no VeriSpecGen rows. Bold indicates the best value of each column.

Checkpoint	SpecGen		CodeGen		SpecGen-Repair		CodeGen-Repair	
	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4
Ours, Qwen3-4B-Instruct-2507								
Base Model	0.0	0.0	2.8	4.0	6.5	8.4	4.1	6.8
SFT	1.2	4.0	7.2	15.3	49.0	72.5	48.5	66.0
RL step 20	3.8	11.3	15.2	34.3	52.1	69.5	51.2	72.8
RL step 40	6.8	17.7	23.8	45.7	54.6	72.5	54.6	74.8
RL step 60	9.3	23.7	34.7	65.0	**56.1**	**77.9**	54.1	76.7
RL step 80	13.7	32.3	42.8	72.0	54.2	75.6	56.8	81.6
RL step 100	17.9	38.3	51.3	75.7	46.6	71.0	56.1	81.6
RL step 120	**24.0**	**43.7**	**57.8**	**81.7**	45.6	76.3	**59.0**	**84.5**

RL accounts for most of the gain. On Verina without test cases, SpecGen pass@1 improves by about 6 points every 20 steps, from 23.0% at step 20 to 54.2% at step 120, and still improves by 3.5 points between steps 100 and 120 (Table 25). From step 40, all 4 scores surpass every VeriSpecGen-

4B model, and from step 100 every VeriSpecGen model including the 30B-A3B models (54.2% and 77.2% SpecGen and CodeGen pass@1 at step 120, compared with at most 47.6% and 61.2%), where our pass@4 is compared with their pass@10. On AUXSPEC-Test, SpecGen pass@1 improves from 1.2% after SFT to 24.0% at step 120 and CodeGen pass@1 from 7.2% to 57.8%, and both still rise (Table 24). Our SFT data contain 19,492 samples, compared with 121,531 (No-Test) or 343,827 (Full) examples for the VeriSpecGen models.

Table 25. Out-of-domain evaluation on Verina with a shared backbone. The setting without test cases follows the default of Verina. We also report the setting with test cases in the prompt, since concrete input-output pairs specify behavior that the description may state ambiguously (Section 5.3). Our model starts from Qwen3-4B-Instruct-2507, the backbone of VeriSpecGen-4B, and is trained with our SFT data and RL recipe; Base Model indicates the backbone before training, evaluated with our prompts. † indicates results reported in the original paper (Ye et al., 2026b) without test cases in the prompt, and italic values are pass@10 in place of pass@4. VeriSpecGen does not release its models, so its results with test cases are not available (–). Among our values without test cases, underline indicates values above every VeriSpecGen-4B model but not above every VeriSpecGen-30B-A3B model, and **bold** indicates values above every VeriSpecGen model, including VeriSpecGen-30B-A3B. Our pass@4 is compared with their pass@10.

Model	w/o tests				w/ tests			
	SpecGen		CodeGen		SpecGen		CodeGen	
	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4
Ours, Qwen3-4B-Instruct-2507								
Base Model	0.0	0.0	0.4	0.5	0.0	0.0	0.0	0.0
SFT	6.9	24.9	21.6	54.5	12.4	35.4	21.8	53.4
RL step 20	23.0	54.5	50.3	79.4	25.3	57.7	51.3	76.2
RL step 40	32.9	59.8	57.9	**85.2**	42.1	68.3	61.9	86.8
RL step 60	39.3	**68.8**	**64.8**	**86.8**	48.8	79.4	70.6	92.1
RL step 80	46.6	**75.7**	**67.3**	**91.5**	59.8	84.1	73.7	92.6
RL step 100	**50.7**	**78.8**	**74.3**	**92.6**	66.7	87.3	81.5	94.7
RL step 120	**54.2**	**84.1**	**77.2**	**95.2**	73.8	89.9	84.0	96.8
VeriSpecGen-4B, Qwen3-4B-Instruct-2507								
Base Model	14.7†	*16.3*†	25.3†	*28.3*†	–	–	–	–
SFT Spec-Only	9.5†	*16.6*†	18.0†	*26.7*†	–	–	–	–
SFT Full	22.8†	*38.8*†	43.5†	*72.7*†	–	–	–	–
SFT No-Test	30.3†	*45.2*†	35.4†	*66.8*†	–	–	–	–
VeriSpecGen-30B-A3B, Qwen3-Coder-30B-A3B								
Base Model	29.4†	*41.7*†	39.7†	*53.5*†	–	–	–	–
SFT Spec-Only	30.4†	*50.0*†	41.5†	*68.4*†	–	–	–	–
SFT Full	43.5†	*58.0*†	61.2†	*79.7*†	–	–	–	–
SFT No-Test	47.6†	*63.9*†	58.7†	*84.0*†	–	–	–	–

F.2 RL RECIPE ABLATION

To isolate the contribution of each component, we construct 3 variants of the RL recipe of AUXSPEC-9B, each removing one component, and train them from AUXSPEC-9B-SFT with the same training prompts, order and seed:

- **w/o test execution** compiles each rollout without its test cases, so no test case is executed, and keeps all judge stages.
- **w/o LLM judges** assigns the top reward to every rollout that passes the test cases.
- **w/o self-repair** fills every batch slot with a new training prompt.

The variant without self-repair therefore consumes training prompts faster, so each variant is compared with AUXSPEC-9B-RL after approximately the same number of training prompts rather than at the same step (Table 26). Evaluation follows Table 4.

Table 26. Component ablations of the 9B recipe on AUXSPEC-Test, with the protocol of Table 4. Each ablation run starts from AUXSPEC-9B-SFT and removes one component, and AUXSPEC-9B-RL is the reference run. Training prompts indicates the number of problems drawn from the RL pool up to each checkpoint. Checkpoints are evaluated every 25 steps up to step 150. Bold indicates the best value of each column and underline the second best.

Model	Step	Training prompts	SpecGen		CodeGen		SpecGen-Repair		CodeGen-Repair	
			pass@1	pass@4	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4
AUXSPEC-9B-SFT	0	0	10.1	24.3	13.1	28.3	56.7	76.3	59.5	80.6
AUXSPEC-9B-RL	70	2,149	21.8	51.3	43.5	80.7	54.3	84.7	55.6	86.4
	140	4,162	34.9	62.0	65.0	89.0	49.6	84.7	55.3	92.2
	210	6,186	**46.7**	**71.7**	**74.9**	**92.0**	67.4	**87.0**	71.1	90.3
w/o test execution	25	961	12.2	28.0	21.8	40.0	59.4	80.9	60.0	77.7
	50	1,833	13.7	30.7	26.0	48.0	62.2	80.2	60.2	82.5
	75	2,604	16.8	35.3	24.0	45.0	63.0	85.5	62.4	85.4
	100	3,362	18.1	37.0	28.9	52.0	62.6	80.2	67.5	88.3
	125	4,111	24.3	45.3	34.4	61.7	66.8	85.5	69.7	90.3
	150	4,796	31.4	53.7	37.6	62.3	65.8	84.7	63.1	86.4
w/o LLM judges	25	830	11.7	29.0	26.0	48.0	62.4	80.9	60.2	76.7
	50	1,626	15.1	32.3	40.2	66.3	62.4	83.2	64.3	90.3
	75	2,346	26.3	50.0	58.3	86.3	68.3	**87.0**	67.0	86.4
	100	3,030	35.8	62.0	67.0	90.0	69.1	**87.0**	68.2	88.3
	125	3,711	31.6	51.0	64.6	84.7	63.4	82.4	69.2	**93.2**
	150	4,556	20.0	35.7	54.8	75.3	43.3	67.9	60.4	84.5
w/o self-repair	25	1,300	13.8	30.3	21.3	42.3	57.1	77.1	61.4	81.6
	50	2,396	14.7	32.3	26.2	53.3	45.8	79.4	51.7	76.7
	75	3,378	20.6	44.3	46.8	75.0	66.4	80.9	63.6	80.6
	100	4,346	23.0	44.3	53.7	81.3	65.8	81.7	64.8	84.5
	125	5,332	33.2	57.3	60.5	88.3	66.6	80.2	**71.4**	90.3
	150	6,334	43.1	67.3	70.2	91.7	**70.8**	86.3	68.9	88.3

LLM judges. Without LLM judges in the reward, the model exploits the reward by restating the given test cases as the specification. We note that the training reward keeps increasing, while SpecGen pass@1 on AUXSPEC-Test reaches its highest value around step 100 and then decreases (Figure 11c1). The gap between the 2 therefore widens over training (84.7% of SpecGen rollouts accepted in steps 141–150 vs. 20.0% pass@1 at step 150, down from 35.8% at step 100). Tautological specifications account for a growing part of the evaluation samples (1.9% at step 50 vs. 58.8% at step 150). On Verina, these specifications pass the test cases given in the prompt but fail without them (88.4% vs. 52.0% SpecGen pass@1 at step 150).

Test execution. Without test execution in the reward, the reward accepts implementations that the test cases would reject. About half of the CodeGen rollouts pass the remaining stages early in training (49.7% vs. 19.7% of CodeGen rollouts accepted in steps 41–50 without and with test execution), while CodeGen pass@1 on AUXSPEC-Test improves slowly (Figure 11b2). After approximately the same number of training prompts the variant falls behind AUXSPEC-9B-RL in both tasks (34.4% vs. 65.0% CodeGen and 24.3% vs. 34.9% SpecGen pass@1 after 4,111 and 4,162 prompts). The gap in CodeGen persists at step 150 even though the variant has drawn more training prompts (37.6% vs. 65.0% after 4,796 vs. 4,162 prompts; Table 26). Repair remains comparable to the other runs (65.8% SpecGen-Repair pass@1 at step 150), so the variant falls behind in generating a correct output from scratch. The judges alone cannot check the behavior of an implementation, and test execution keeps the reward accurate.

Self-repair. Without self-repair every batch slot holds a new training prompt. The variant therefore consumes training prompts faster and outperforms AUXSPEC-9B-RL at the same step (43.1% vs. 34.9% SpecGen pass@1 at steps 150 and 140 after 6,334 vs. 4,162 prompts). AUXSPEC-9B-RL outperforms it in both generation and repair even after fewer training prompts (21.8% vs. 14.7% SpecGen and 54.3% vs. 45.8% SpecGen-Repair pass@1 after 2,149 and 2,396 prompts). After about 6.2K prompts it still outperforms the variant in generation (46.7% vs. 43.1% SpecGen pass@1; Table 26). Training without repair tasks is thus a viable recipe when new problems are plentiful.

The repair tasks are designed for settings in which a model revises its output from feedback over several rounds, such as long-horizon verification and collaboration between agents. Refinement from compiler and judge feedback is part of the training objective, and the model also refines its specifications at test time: on ICPC-Eval, refine@16 reaches 82.7% of the problems, compared with 61.3% for pass@16 (Appendix D.4).

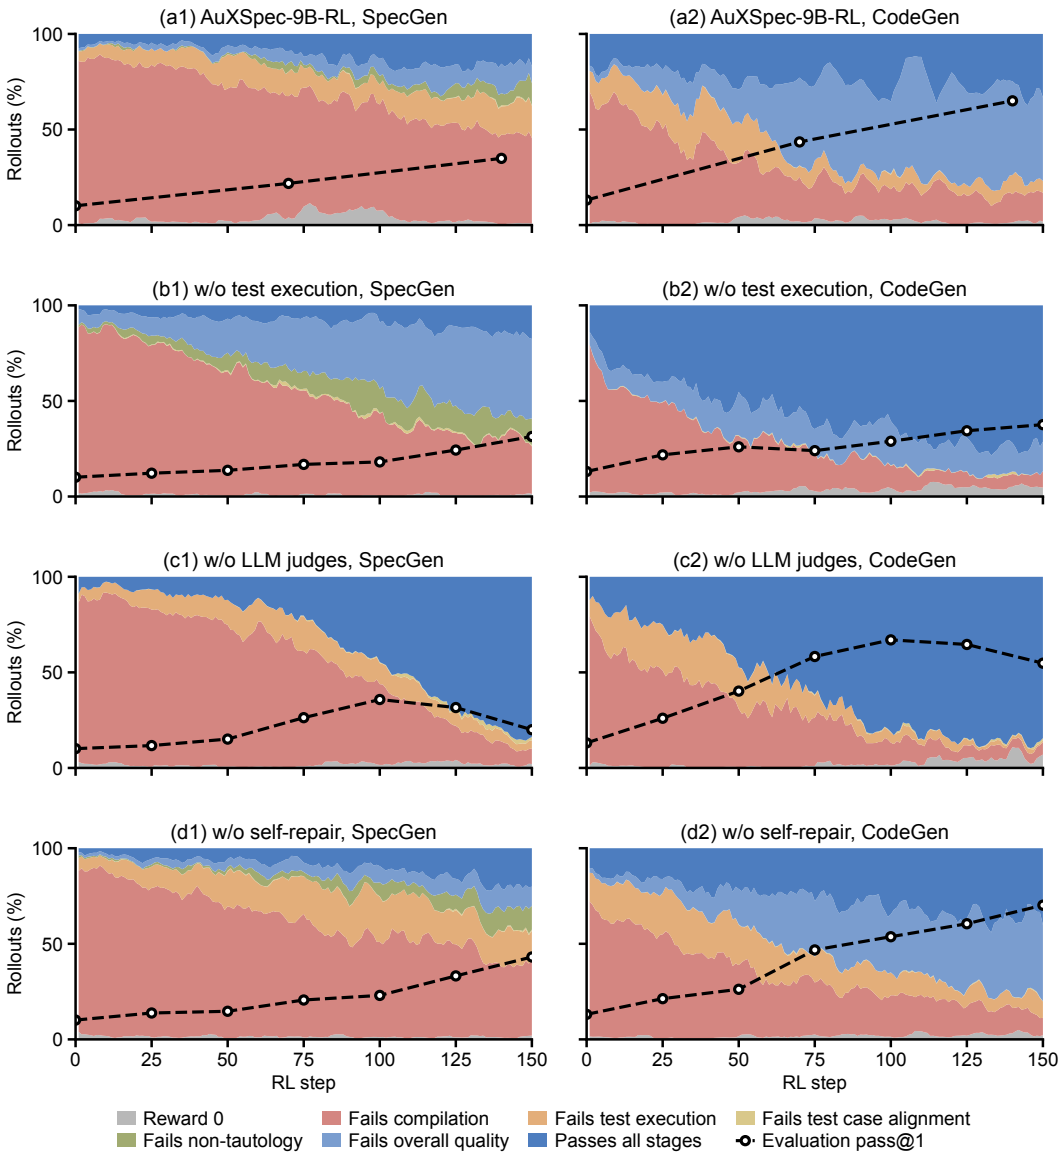


Figure 11. Training rollouts by reward outcome for AUXSPEC-9B-RL and the 3 variants, one row per run, on SpecGen (left) and CodeGen (right), from the first failed stage to passing all stages, with a 5-step moving average, up to step 150. The variant without self-repair draws more training prompts per step than the other runs, and Table 26 compares the runs after approximately the same number of training prompts. Figure 2d shows the full run of AUXSPEC-9B-RL. Each variant drops the bands of its removed stages, so rollouts that those stages would reject count as passing all stages. Dashed lines indicate pass@1 of the same task on AUXSPEC-Test at the evaluated checkpoints, and the first point of each line indicates AUXSPEC-9B-SFT. A widening gap between the top band and the dashed line indicates rewarded outputs that evaluation rejects.

We also report Verina at the same checkpoints, where the variants follow similar trends (Table 27).

Table 27. Component ablations of the 9B recipe on Verina, with and without test cases in the prompt. The setting without test cases follows the default of Verina. We also report the setting with test cases in the prompt, since concrete input-output pairs specify behavior that the description may state ambiguously (Section 5.3). Rows as in Table 26.

Model	Step	Training prompts	SpecGen, w/o tests		CodeGen, w/o tests		SpecGen, w/ tests		CodeGen, w/ tests	
			pass@1	pass@4	pass@1	pass@4	pass@1	pass@4	pass@1	pass@4
AUXSPEC-9B-SFT	0	0	34.9	65.6	36.5	71.4	37.3	69.8	38.0	70.4
AUXSPEC-9B-RL	70	2,149	52.1	82.0	75.8	93.7	52.4	86.2	80.4	97.4
	140	4,162	67.3	91.5	86.4	**97.4**	71.7	92.1	87.6	**98.4**
	210	6,186	73.4	**92.1**	**89.2**	96.8	84.8	**97.4**	**90.1**	96.8
w/o test execution	25	961	47.6	70.9	63.9	84.7	48.8	75.1	63.9	86.2
	50	1,833	52.0	78.8	67.9	85.7	56.6	82.0	66.5	85.7
	75	2,604	55.2	78.8	69.3	87.3	60.3	82.0	71.2	89.4
	100	3,362	58.6	81.0	73.5	91.0	65.5	83.1	73.0	89.9
	125	4,111	65.1	85.7	76.5	91.5	72.8	88.9	77.6	92.6
	150	4,796	69.0	85.7	80.0	92.6	73.9	90.5	80.8	96.3
w/o LLM judges	25	830	41.1	73.0	63.1	83.6	45.2	71.4	64.2	85.2
	50	1,626	48.9	77.8	68.8	91.0	46.7	76.7	70.1	91.0
	75	2,346	55.3	84.1	76.1	94.7	63.0	89.4	76.6	96.3
	100	3,030	63.4	88.4	84.7	95.8	75.0	95.2	83.2	**98.4**
	125	3,711	63.9	87.3	84.0	96.8	82.1	95.8	89.0	97.4
	150	4,556	52.0	79.4	81.5	94.7	**88.4**	96.3	87.7	97.9
w/o self-repair	25	1,300	43.3	73.5	57.1	81.0	50.0	78.8	61.5	83.1
	50	2,396	57.1	84.7	68.5	86.8	58.3	84.1	70.5	90.5
	75	3,378	61.1	84.7	70.9	89.9	63.2	88.4	74.2	92.1
	100	4,346	64.4	86.2	77.1	93.1	65.1	87.8	79.8	93.1
	125	5,332	67.3	85.7	81.0	96.3	78.0	92.6	81.0	92.6
	150	6,334	**73.8**	89.9	83.1	95.8	82.8	94.7	83.3	96.8

G PIPELINE DETAILS

Figure 12 shows the pipeline of Figure 1 on a running example, with the input of each stage and the Lean code that the stage produces. Each stage below ends with an analysis of the stage on the 3,037 synthesized problems.

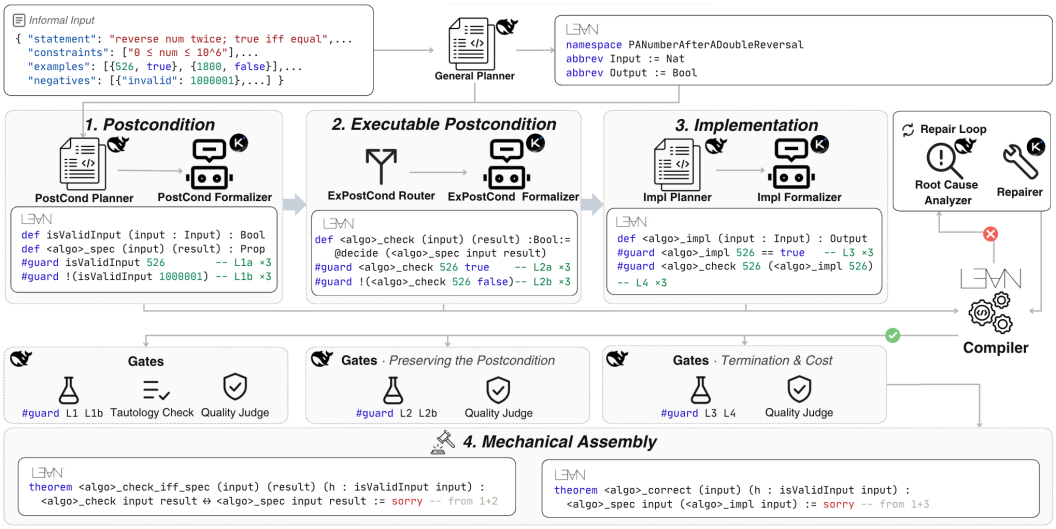


Figure 12. Data synthesis pipeline for AUXSPEC on a running example. Stages (1) to (3) synthesize the postcondition, the executable postcondition, and the implementation under separate plans, and stage (4) assembles the components with the theorem statements.

G.1 PREPROCESSING

G.1.1 TEST CASES

For each of 3,109 problems (719 easy, 1,622 medium, 768 hard), we screen the stored examples and propose negative test cases, invalid inputs and incorrect outputs beside the positive test cases. Invalid inputs are proposed by category (a value just outside a stated bound, a size outside the stated range, a container of the wrong shape, a broken structural invariant, parameters invalid only together, a value outside the stated domain), and incorrect outputs by category (a near miss, a suboptimal feasible answer, an unattainable answer, the answer under a relaxed constraint, an output of the wrong shape, a degenerate answer). Benchmarks store negative test cases beside positive ones (Ye et al., 2026a;b), since a specification that accepts correct outputs on the given examples can still accept incorrect outputs (Endres et al., 2024).

A proposed incorrect output is admitted when the reference solution, run on the proposed valid input, returns an output that is not a reordering of the proposed output. Invalid inputs are proposed together with the violated constraint and are dropped on coinciding with a valid example of the problem. Cases too large for compiled evaluation and cases whose stored answer is null are refused before any model call, with the measured thresholds in Appendix G.5.

The 27,981 test cases, 9 per problem, are assembled into 18 `#guard` commands per problem, one per test case and layer of Table 7, with each positive test case reused at $L1^a$, $L2^a$, L3, and L4. A `#guard` command runs compiled code, and a passed test case relies on the compiler and runtime as trusted base, the same as `native_decide`. For the test cases, we store the provenance (public example or synthesized) and the admission round. A test case judged wrong is replaced by a held-back test case at most once per problem.

Analysis. The frozen test cases, negative test cases included, are shown to the generation model, and the negative test cases act as part of the problem statement rather than as a hidden filter. Of 3,276 specification candidates tested at $L1^a$ and $L1^b$, 8 (0.24%) fail only at $L1^b$, and of 3,195 executable postconditions tested at $L2^a$ and $L2^b$, 4 (0.13%) fail only at $L2^b$. The 4 failures at $L2^b$ have one cause, a lossy output decoder in the specification, such as a tree decoder that drops children below null nodes. The 8 failures at $L1^b$ fall on 5 problems, and in 4 of the 5 problems the frozen invalid

input is in fact valid. Repair narrowed the precondition to reject that input in all 5 problems. Frozen test values keep test cases from adapting to a wrong specification, and these 5 preconditions are the only adaptation in the opposite direction.

Screening. Before any model call, test cases are filtered by size and by null answers. A test case containing a list or string longer than 128 items is refused, since a case is frozen as a literal that compiled evaluation builds item by item. We cap compilation at 600 seconds on the Lean server and observe that a test case exceeding the cap terminates the serving REPL. The threshold was measured over the stored test cases, and any threshold between 102 and 185 items refuses the same test cases. A case whose stored answer is null is refused in the layers that read an answer, since the source dataset stores null wherever the answer is a mutated argument, and a null true answer makes each negative built from the null answer meaningless.

G.1.2 Library retrieval

Models reference library declarations absent from the pinned library version (Chakraborty et al., 2025). Typed queries against Mathlib, for a definition, a theorem, or an instance, are generated together with one sentence stating the use the specification has for the answer, and relevance is judged against that sentence. Search candidates are deduplicated, elaborated with `#check` against the pinned toolchain, and judged for relevance by an LLM given the problem description. Specification synthesis is augmented with the declarations admitted by the relevance judge, 4.2 per problem, each with a checked signature and a one-sentence reason from the LLM. Appendix G.5 reports the counts at each step.

Analysis. Elaboration keeps declarations that exist only in newer Mathlib versions out of specification synthesis (Table 28). 23,608 queries (14,888 definition, 5,052 theorem, 3,668 instance) produced 92,104 candidates. Of the elaborated declarations, 5.4% were absent from the pinned Mathlib, and review dropped 10.0% of the reviewed candidates. Of the 13,191 stored declarations, specification synthesis of the 3,037 problems received 12,876, 4.2 per problem, with 131 problems receiving none. The absent declarations originate in a search index built on the latest Mathlib, while synthesis compiles against the pinned version. Retrieved declarations do not change the first compilation. Specifications that reference a retrieved declaration by the full name compile at the first attempt in 83.2% of 1,871 problems, and the other specifications in 83.9% of 1,032.

Table 28. Library retrieval by step over 3,109 problems and 3,715 rounds. The first 5 rows indicate work across all rounds, and the last row indicates the number of declarations passed to specification synthesis, so a problem retrieved more than once counts only the last round there.

Step	Count	Note
Queries generated	23,608	Definition, theorem, and instance queries
Candidates returned	92,104	Search hits
Declarations elaborated	83,679	Deduplicated, #check against the pinned toolchain
Declarations present	79,126	5.4% of elaborated declarations were absent
Reviewed	17,372	1,729 dropped as irrelevant
Stored	13,191	Last record per problem, 4.2 per problem

G.2 Specifications

Planning. A type plan fixes the declaration names and one type per parameter, with a conversion function wherever the stored value has a different type. The postcondition is then planned regardless of decidability, which keeps computation out of the postcondition. A frozen test value is printed once as the expression used by all test layers to prevent models from altering the value.

Formalization.

Each component is generated separately and compiled with the accepted components. The `#guard` commands of the layers are generated from the expressions the plan fixed, and each command is accepted at the first passing check in the order shape, verbatim expression, token-level expression, and judge review of equivalence. The shape check admits one command per test case, attached to the declaration of that layer and with the polarity of the test case. The expression check compares the tokens of the printed expression, which keeps whitespace from merging 2 values.

Executable postcondition. The executable postcondition is first set to the decidable form of the postcondition, and the file is compiled. When no decidability instance is derived, a model supplies

instances from the compiler diagnostics for up to 3 rounds. When the supplied instances still fail to elaborate, a separate definition is generated. Only a separate definition can disagree with the postcondition, which the L2[a] and L2[b] test cases detect on the tested outputs and the postcondition equivalence theorem excludes on all valid inputs. Appendix G.5 reports the counts by form.

Judges. The postcondition judge is given the problem description, the constraints, and the postcondition source. A rejection quotes a sentence of the description and names an output that the postcondition admits and the problem excludes. The executable postcondition judge is given the constraints, the accepted precondition and postcondition, and the executable postcondition, and a rejection names a valid input with a witness or counterexample that the executable postcondition misses or wrongly admits. The tautology judge rejects a postcondition that computes the output. Problems whose natural statement is the transformation itself are labeled in advance by a majority vote of 3 classifier calls, and the label exempts the problem from the tautology judge. The rejected component is repaired with the named error, or replanned for an error in the plan.

G.2.1 Example problems

Separate executable postcondition. We illustrate the format with problem *1955 Count Number of Special Subsequences* from AuxSpec-Test (Figures 13 and 14). In this problem, we are given an array over $\{0, 1, 2\}$ and need to count, modulo $10^9 + 7$, the subsequences of the form $0^a1^b2^c$ with $a, b, c \geq 1$. The postcondition states the result as the size of the set of index sets whose selected elements take this form. The executable postcondition enumerates all subsets of the indices, a brute-force computation that finishes on the short test inputs. The implementation counts in 1 pass with 3 counters, for subsequences of 0s, of 0s followed by 1s, and of 0s, 1s and 2s, since an element either extends or skips the subsequences counted so far. Since neither specification mentions the counters, the correctness theorem relates the recurrence to the set of index sets.

For this problem, a natural postcondition characterizes the valid index sets S as exactly those that select a subsequence of the form $0^+1^+2^+$ from the input. This characterization is not directly executable, since it ranges over arbitrary finite sets of natural numbers. However, every relevant index set is contained in the input indices. The executable postcondition therefore enumerates $2^{\{0,\ldots,\text{input.length}-1\}}$ and counts the subsets satisfying the same property. The two formulations describe the same relation but are not definitionally identical in Lean, so the postcondition equivalence theorem connects the logical characterization with the bounded enumeration.

Executable postcondition through type class resolution. We illustrate the decidable form with problem *523 Continuous Subarray Sum* from AuxSpec-Test (Figures 15a, 15b and 16). In this problem, we are given an array of non-negative integers and a positive integer k and need to decide whether some contiguous subarray of length at least 2 sums to a multiple of k. The postcondition states that the result is true exactly when there exist 2 indices and an integer multiplier n such that the subarray between the indices sums to $n \cdot k$.

The executable postcondition applies `decide` to the postcondition, which compiles only when type class resolution finds a `Decidable` instance for the unfolded proposition. The existential over natural-number indices and an integer multiplier has no such instance as stated, since the quantifiers range over unbounded domains. The pipeline therefore synthesizes a type class instance that bounds both indices by the length of the array, replaces the multiplier by integer divisibility, and proves the bounded proposition equivalent to the original with `decidable_of_iff`. Type class resolution then assembles the executable postcondition from this instance and the library instances for bounded quantifiers and divisibility, so the executable postcondition agrees with the postcondition by construction (Appendix H.2). The implementation, omitted from the figure, scans prefix sums modulo k with a map from each remainder to the earliest prefix index at which the remainder occurs, and the correctness theorem relates this scan to the existential of the postcondition.

```
import Mathlib

/-!
# Count Number of Special Subsequences

A sequence is special if it consists of a positive number of 0s, followed by a positive number
of 1s, then a positive number of 2s.

Given an array nums (consisting of only integers 0, 1, and 2), return the number of different
subsequences that are special. Since the answer may be very large, return it modulo 10^9 + 7.
-/

abbrev Input := List Int
abbrev Output := Nat

-- Specification helpers
def IsSpecialSubsequence (nums : List Int) (S : Finset Nat) : Prop :=
  (∀ i ∈ S, i < nums.length) ∧
    ∃ a b c : List Int, a ≠ [] ∧ b ≠ [] ∧ c ≠ [] ∧
      (∀ x ∈ a, x = 0) ∧ (∀ x ∈ b, x = 1) ∧ (∀ x ∈ c, x = 2) ∧
      ((nums.zipIdx.filter fun p => p.2 ∈ S).map Prod.fst) = a ++ b ++ c

-- Precondition
def isValidInput (input : Input) : Bool :=
  1 ≤ input.length && input.length ≤ 10^5 &&
    input.all (fun x => x = 0 || x = 1 || x = 2)

-- Postcondition
def numberOfSpecialSubsequences_spec (input : Input) (result : Output) : Prop :=
  ∃ C : Finset (Finset Nat),
    result = C.card % (10^9 + 7) ∧
      ∀ S : Finset Nat, S ∈ C ↔ IsSpecialSubsequence input S

-- Executable postcondition helpers
def specialSplits (l : List Int) : Bool :=
  decide
    (∃ i j : Fin (l.length + 1),
      i ≤ j ∧
        l.take i.val ≠ [] ∧ (l.drop i.val).take (j.val - i.val) ≠ [] ∧ l.drop j.val ≠ [] ∧
          (∀ x ∈ l.take i.val, x = 0) ∧ (∀ x ∈ (l.drop i.val).take (j.val - i.val), x = 1) ∧
            (∀ x ∈ l.drop j.val, x = 2))

def isSpecialSet (nums : List Int) (S : Finset Nat) : Bool :=
  decide (∀ i ∈ S, i < nums.length) &&
    specialSplits ((nums.zipIdx.filter fun p => p.2 ∈ S).map Prod.fst)

-- Executable postcondition
def numberOfSpecialSubsequences_check (input : Input) (result : Output) : Bool :=
  result ==
    ((Finset.range input.length).powerset.filter
    (fun S => isSpecialSet input S = true)).card % (10 ^ 9 + 7)

-- Equivalence theorem
theorem numberOfSpecialSubsequences_check_iff_spec (input : Input) (result : Output)
    (h : isValidInput input) :
    numberOfSpecialSubsequences_check input result
    ↔ numberOfSpecialSubsequences_spec input result := by
  sorry

-- Implementation helpers
structure DPState where
  zeros : Nat
  ones : Nat
  twos : Nat

def dpStep (s : DPState) (x : Int) : DPState :=
  if x = 0 then { s with zeros := (2 * s.zeros + 1) % (10 ^ 9 + 7) }
  else if x = 1 then { s with ones := (2 * s.ones + s.zeros) % (10 ^ 9 + 7) }
  else { s with twos := (2 * s.twos + s.ones) % (10 ^ 9 + 7) }

-- Implementation
def countSpecialSubsequences (input : Input) : Output :=
  (input.foldl dpStep ⟨0, 0, 0⟩).twos

-- Correctness theorem
theorem countSpecialSubsequences_correct (input : Input) (h : isValidInput input) :
    numberOfSpecialSubsequences_spec input (countSpecialSubsequences input) := by
  sorry
```

Figure 13. Declarations of the example problem, without namespace and docstrings (tests in Figure 14).

```
/-! ## Test cases -/

-- L1a
#guard isValidInput [0, 1, 2, 2]
#guard isValidInput [2, 2, 0, 0]
#guard isValidInput [0, 1, 2, 0, 1, 2]

-- L1b
#guard !(isValidInput [])
#guard !(isValidInput [3])
#guard !(isValidInput [(-1)])

-- L2a
#guard numberOfSpecialSubsequences_check [0, 1, 2, 2] 3
#guard numberOfSpecialSubsequences_check [2, 2, 0, 0] 0
#guard numberOfSpecialSubsequences_check [0, 1, 2, 0, 1, 2] 7

-- L2b
#guard !(numberOfSpecialSubsequences_check [0, 1, 2, 2] 2)
#guard !(numberOfSpecialSubsequences_check [2, 2, 0, 0] 1)
#guard !(numberOfSpecialSubsequences_check [0, 1, 2, 0, 1, 2] 0)

-- L3
#guard countSpecialSubsequences [0, 1, 2, 2] == 3
#guard countSpecialSubsequences [2, 2, 0, 0] == 0
#guard countSpecialSubsequences [0, 1, 2, 0, 1, 2] == 7

-- L4
#guard numberOfSpecialSubsequences_check [0, 1, 2, 2] (countSpecialSubsequences [0, 1, 2, 2])
#guard numberOfSpecialSubsequences_check [2, 2, 0, 0] (countSpecialSubsequences [2, 2, 0, 0])
#guard numberOfSpecialSubsequences_check [0, 1, 2, 0, 1, 2] (countSpecialSubsequences [0, 1, 2,
    0, 1, 2])
```

Figure 14. Test cases of the example problem, grouped by test layer.

```
import Mathlib

/-!
# Continuous Subarray Sum

Given an integer array nums and an integer k, return true if nums has a good subarray or false
otherwise.
A good subarray is a subarray where:

its length is at least two, and
the sum of the elements of the subarray is a multiple of k.
-/

abbrev Input := (List Int) × (Nat)
abbrev Output := Bool

-- Specification helpers
def IsGoodSubarray (nums : List Int) (k : Nat) (startIdx endIdx : Nat) (n : Int) : Prop :=
  startIdx + 2 ≤ endIdx ∧ endIdx ≤ nums.length ∧
    ((nums.drop startIdx).take (endIdx - startIdx)).sum = n * (k : Int)

-- Precondition
def isValidInput (input : Input) : Bool :=
  1 ≤ input.1.length && input.1.length ≤ 10^5 &&
    input.1.all (fun x => 0 ≤ x && x ≤ 10^9) &&
    input.1.sum ≤ 2^31 - 1 && 1 ≤ input.2 && input.2 ≤ 2^31 - 1

-- Postcondition
def hasGoodSubarray_spec (input : Input) (result : Output) : Prop :=
  result = true ↔ ∃ (startIdx endIdx : Nat) (n : Int),
    IsGoodSubarray input.1 input.2 startIdx endIdx n
```

Figure 15a. Declarations of the second example problem, part 1.

```
-- Decidability instance
instance decidableGoodSubarrayExists (nums : List Int) (k : Nat) :
    Decidable (∃ startIdx endIdx n, IsGoodSubarray nums k startIdx endIdx n) :=
  decidable_of_iff
    (∃ startIdx < nums.length + 1,
      ∃ endIdx < nums.length + 1,
        startIdx + 2 ≤ endIdx ∧
          (k : Int) | ((nums.drop startIdx).take (endIdx - startIdx)).sum)
    (by
      constructor
      · rintro ⟨startIdx, _, endIdx, hend, hge, hdvd⟩
        rcases hdvd with ⟨n, hn⟩
        refine ⟨startIdx, endIdx, n, hge, ?_, ?_⟩
        · omega
        · rw [mul_comm]
          exact hn
      · rintro ⟨startIdx, endIdx, n, hge, hend, hsum⟩
        refine ⟨startIdx, ?_, endIdx, ?_, hge, ⟨n, ?_⟩⟩
        · omega
        · omega
        · rw [mul_comm]
          exact hsum)

-- Executable postcondition
def hasGoodSubarray_check (input : Input) (result : Output) : Bool :=
  @decide (hasGoodSubarray_spec input result)
    (by unfold hasGoodSubarray_spec; infer_instance)

-- Equivalence theorem
theorem hasGoodSubarray_check_iff_spec (input : Input) (result : Output)
    (h : isValidInput input) :
    hasGoodSubarray_check input result ↔ hasGoodSubarray_spec input result := by
  sorry
```

Figure 15b. Declarations of the second example problem, part 2, whose executable postcondition is `decide` of the postcondition, with the `Decidable` instance found by type class resolution from a synthesized instance. Namespace, docstrings, and the implementation with the correctness theorem are omitted (tests in Figure 16).

```
/-! ## Test cases -/

-- L1a
#guard isValidInput ([23, 2, 4, 6, 7], 6)
#guard isValidInput ([23, 2, 6, 4, 7], 6)
#guard isValidInput ([23, 2, 6, 4, 7], 13)

-- L1b
#guard !(isValidInput ([], 1))
#guard !(isValidInput ([(-1), 1], 1))
#guard !(isValidInput ([1000000001], 1))

-- L2a
#guard hasGoodSubarray_check ([23, 2, 4, 6, 7], 6) true
#guard hasGoodSubarray_check ([23, 2, 6, 4, 7], 6) true
#guard hasGoodSubarray_check ([23, 2, 6, 4, 7], 13) false

-- L2b
#guard !(hasGoodSubarray_check ([23, 2, 4, 6, 7], 6) false)
#guard !(hasGoodSubarray_check ([23, 2, 6, 4, 7], 6) false)
#guard !(hasGoodSubarray_check ([23, 2, 6, 4, 7], 13) true)

-- L3
#guard existsGoodSubarray ([23, 2, 4, 6, 7], 6) == true
#guard existsGoodSubarray ([23, 2, 6, 4, 7], 6) == true
#guard existsGoodSubarray ([23, 2, 6, 4, 7], 13) == false

-- L4
#guard hasGoodSubarray_check ([23, 2, 4, 6, 7], 6) (existsGoodSubarray ([23, 2, 4, 6, 7], 6))
#guard hasGoodSubarray_check ([23, 2, 6, 4, 7], 6) (existsGoodSubarray ([23, 2, 6, 4, 7], 6))
#guard hasGoodSubarray_check ([23, 2, 6, 4, 7], 13) (existsGoodSubarray ([23, 2, 6, 4, 7], 13))
```

Figure 16. Test cases of the second example problem, grouped by test layer.

G.2.2 Analysis

The specification judges rarely reject, and the errors that pass all checks are in the meaning of the specification. Each verdict in Table 29 corresponds to a distinct component that compiles and passes the test layers of the component. The specification judges reject 1.7% to 3.8% of the verdicts, and all rejected problems are accepted after repair. The errors that proving or human review later found had all passed the test layers and the judge of the erroneous component. In 2 problems a correct rejection led to a repair that introduced the error found by review, and in 1 problem the judge accepted on a second reading the defect that the judge had named.

Table 29. Verdicts of the judges on components that compile and pass the test layers of the component. A problem counts as rejected when at least 1 verdict on the problem is a rejection.

Judge	Verdicts	Rejections	Problems rejected	Accepted after rejection
Faithfulness	3,237	123 (3.8%)	112	112
Non-tautology	3,143	69 (2.2%)	61	61
Completeness	3,159	55 (1.7%)	42	42
Implementation	3,667	717 (19.6%)	414	281

The decidable form is reliable where the postcondition quantifies over bounded types (Figure 17). The decidable form serves 1,079 of the 3,037 problems (35.5%). Of the other 1,958 problems, 1,343 have postconditions that quantify over unbounded types, 319 were synthesized before instance synthesis was added, 287 fail instance synthesis for 3 rounds, 6 remain unexplained, and 3 are routed to the decidable form but keep a separate definition in the final file. Where the decidable form applies, the executable postcondition passes all checks at the first attempt in 772 of 787 problems (98.1%), compared with 1,178 of 2,250 (52.4%) for separately generated executable postconditions. Missing instances arise from the decidable form. The pipeline first sets the executable postcondition to `decide` applied to the postcondition, which compiles only when Lean finds a decidability instance for the postcondition. A postcondition that quantifies over an unbounded type (e.g., over all natural numbers or all lists) has no computable instance as stated, and instance synthesis fails, for example on add-digits, where no decidability instance can be synthesized for a postcondition that quantifies over all digit roots d.

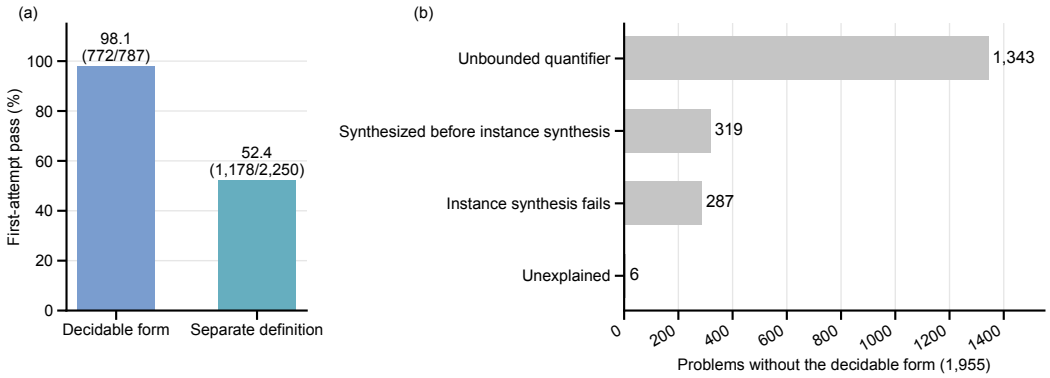


Figure 17. (a) Executable postconditions that pass all checks at the first attempt, by route. (b) Reasons that the decidable form does not apply.

G.3 Implementation

Implementation planning is given the accepted specification as read-only text and selects an algorithm with cost argument, listing specification definitions to call and definitions to replace with reason. The plan is checked before generation for a termination argument based on a decreasing quantity of the input, a stated cost consistent with the approach and the input bounds, a replacement justified by cost or by a different question, and an algorithm that computes the answer, without search over candidates. Beyond the properties in Section 3.1, the implementation judge reports computations already available in the specification.

G.3.1 Analysis

Test cases on small inputs admit false positives (Li et al., 2022), and the implementation judge rejects 717 of the 3,667 implementations (19.6%) that pass all test cases (Figure 18).

- **Fuel parameters (270).** The implementation recurses on a fuel argument, which hides nontermination on inputs beyond the test cases.
- **Cost beyond the input bounds (234).** The implementation is correct but exceeds the cost that the stated input bounds allow, known as a slow positive (Li et al., 2022).
- **Restated executable postconditions (198).** The implementation repeats a computation of the executable postcondition that already runs within the stated input bounds.
- **Replayed test cases (15).** The implementation returns stored outputs for the tested inputs.

The implementation judge is waived after 2 retries, and 133 problems keep an implementation that the judge rejected. Of these, 89 are rejected for a fuel argument that may hide nontermination, 27 for cost beyond the input bounds, 12 for restating the executable postcondition and 5 for replaying test cases. 12 of the 133 are in AuxSpec-Test, and for the 9 among them rejected for fuel, the correctness proofs show that the fuel suffices on all valid inputs.

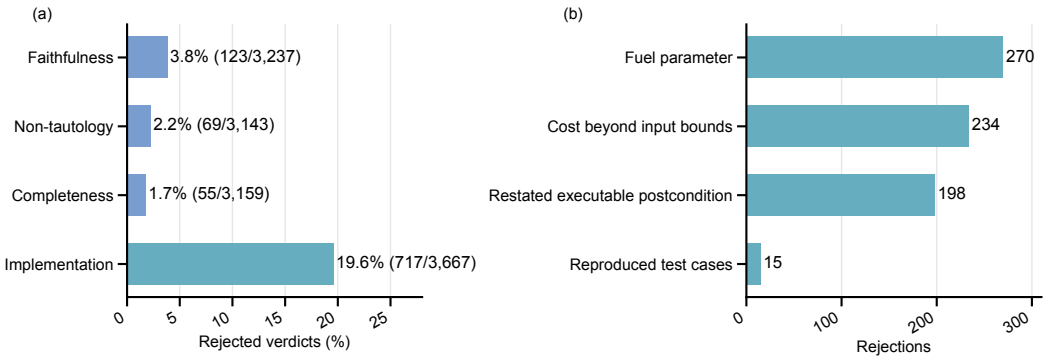


Figure 18. (a) Rejections by each judge on compiled components that pass the test layers. (b) Implementation judge rejections by reason.

G.4 Repair

Compilation runs through the Kimina Lean server (Santos et al., 2025), a REPL-based server (Lean community, 2023). Repair revises rejected components separately for at most 18 rounds per run of the repair loop, which a problem re-enters after replanning. From the second round on, the repair input includes a digest of earlier attempts, and a root-cause analysis runs for errors that line numbers cannot attribute. Repair stops early when a revision repeats an earlier diff or does not change the component. Where a repair attributes the error to a test case, the attribution is reviewed by a judge, and a confirmed attribution replaces the test case with a held-back test case at most once per problem.

G.4.1 Analysis

Repair recovers most first failures within 2 rounds (Figure 19). The fraction of the 3,037 problems accepted within k repair rounds rises from 78.1% at $k = 0$ to 93.9% at $k = 1$ and 97.5% at $k = 2$ for the postcondition, from 67.7% to 91.9% and 97.7% for the executable postcondition, and from 64.3% to 85.0% and 94.4% for the implementation.

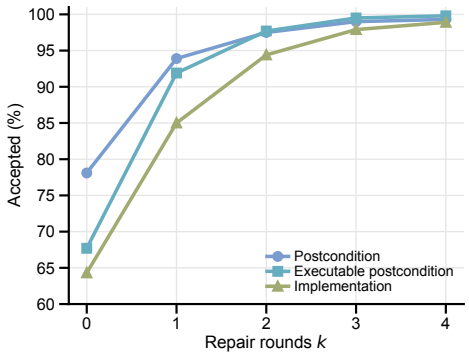


Figure 19. Cumulative acceptance rate of each component over repair rounds on the 3,037 synthesized problems. The implementation depends most on repair, with 64.3% accepted at the first attempt.

The remaining problems took 1,030 postcondition repair rounds, 1,321 executable postcondition rounds and 1,842 implementation rounds (Figure 20).

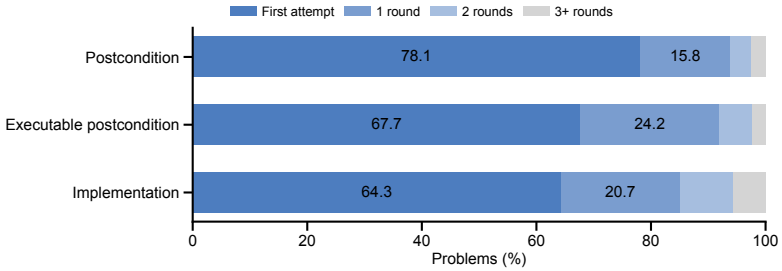


Figure 20. Repair rounds needed by each component on the 3,037 synthesized problems.

1,081 problems (34.8% of the 3,109 seeds) need no repair round for any component, and repair recovers most of the others, most on hard problems, where 22.1% of the seeds pass without repair and 95.7% with repair (Figure 21).

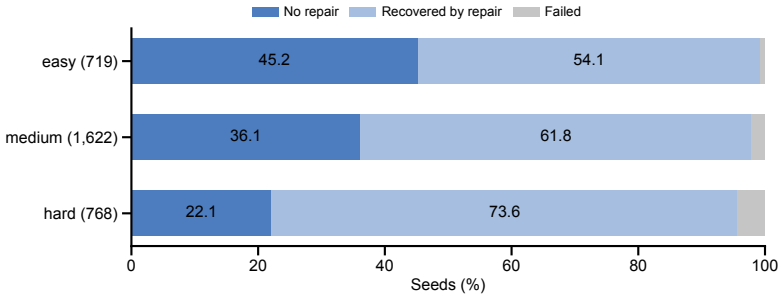


Figure 21. Seeds of each difficulty accepted without repair, recovered by repair, and failed.

Repair never replaces a frozen test case. Repair attributed 113 failures on 96 problems to test cases, and the root-cause analyzer placed the fault in the model-generated `#guard` command in 55 and in the plan in 52, never in a frozen test case. The most frequent case is a negative test case with a negative value for an input that the plan types as a natural number, which the judge resolves by changing the planned type to an integer. Repair rounds go to Lean usage rather than to failed test cases (Figure 22). Of 3,852 recorded rounds, Mathlib API errors take 29.7%, missing instances 23.8% (Appendix G.2.2), type mismatches 12.6%, parse errors 11.3%, proof obligations 10.0%, and failed test cases 3.1%. The fraction of Mathlib API errors matches the misuse of Mathlib declarations that motivates continued pretraining (Appendix E.2).

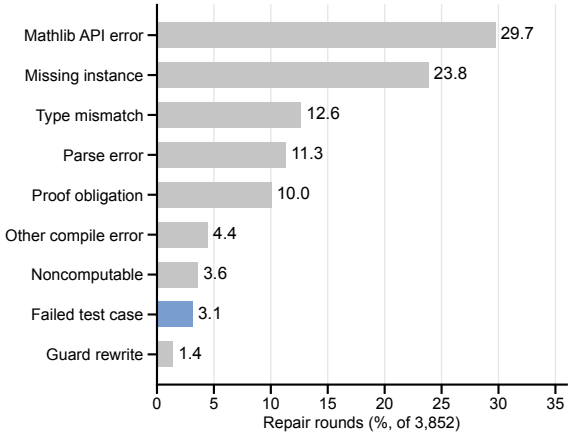


Figure 22. Causes of 3,852 repair rounds. Proof obligations include termination, `omega` goals and index bounds.

G.5 Outcomes

We break down the outcomes of Section 3.2 by difficulty (Table 30).

Table 30. Outcomes of the 3,109 seed problems. Percentages are over the seeds of each difficulty.

Outcome	Problems	Note
Succeeded	3,037	Compiled, all frozen layers pass
Easy	714	Of 719 seeds, 99.3%
Medium	1,588	Of 1,622 seeds, 97.9%
Hard	735	Of 768 seeds, 95.7%
Failed	72	Terminal exits, stalled runs, and later withdrawals

The decidable form of the postcondition is adopted in 1,079 problems (35.5%) and a separate definition in 1,958 problems. The frozen `#guard` commands of 90.0% of the synthesized problems match the planned expressions verbatim, and the other 303 needed token-level matching or a judge.

G.6 Training samples

The reasoning trace of a training sample is rephrased from the reasoning recorded while that problem was synthesized, which is the planning and generation reasoning for the generation tasks and the diagnostics, root cause, and repair reasoning for the repair tasks. A rephrased trace is admitted only when each sentence is grounded in the recorded reasoning, the answer is not stated before being derived, and no pipeline vocabulary appears.

G.7 Cost

Synthesis and supporting experiments cost $2,604 over 226,351 model calls, divided across 9 accounts (Table 31). Accounts follow record keeping, and reasoning-trace rephrasing, model comparisons, and discarded experiments are listed beside retrieval, test case construction, and formalization. Accounts cover different problems, and per-problem costs are not compared across accounts.

Of the 7 accounts with recorded tokens, some are the provider-reported per-call charge as recorded by the pipeline, and the rest are priced from the recorded tokens at the list rate of the provider. The other 2 accounts are estimated from their call counts. List pricing ignores the lower rate for cached input and therefore overstates the cost of the priced accounts. The table lists the models of the accounts, since an account mixes the prices of several models.

The remaining 2 accounts, classification and judging, have a call count and no dollars. One vote is one call in both, and the 99,736 calls are therefore exact. Both accounts recorded calls without tokens and are counted by calls, outside the dollar total.

The dollars and the tokens of the formalization row count different calls. 6,237 of the formalization calls were served from the response cache, at no charge, and these calls make up 12.0% of the

Table 31. Spend by account. The last 2 accounts recorded no tokens, so their tokens and spend are estimated from the exact call counts with the mean usage of DeepSeek-V4-Flash judge calls at list price. The judging calls are an upper bound, since some were served from cache.

Account	Calls	Tokens In	Tokens Out	Models	Spend
Retrieval	7,421	32M	71M	DeepSeek-V4-Flash	$24
Test corpus	43,880	109M	341M	DeepSeek-V4-Flash	$98
Formalization	59,337	182M	290M	Kimi-K3 (generation, repair), DeepSeek-V4-Pro (planning), DeepSeek-V4-Flash (judging)	$2,024
Reasoning traces	12,320	209M	164M	DeepSeek-V4-Flash	$75
Model comparisons	3,183	20M	20M	Claude-Opus-5, DeepSeek-V4-Flash, DeepSeek-V4-Pro, GLM-5.2, Kimi-K3	$184
Pipeline experiments	435	3M	2M	Claude-Opus-5, DeepSeek-V4-Flash	$36
Bake-off	39	0M	0M	GLM-5.2, Kimi-K3	$2
Classification	15,080	36M	68M	DeepSeek-V4-Flash	$24
Judging	84,656	202M	385M	DeepSeek-V4-Flash	$136
Total	226,351				$2,604

tokens in the formalization account. Since tokens include cached calls and dollars exclude cached calls, dollars per token understate the price.

Formalization is the largest account at $2,024, which is $0.65 per problem on average, $0.47 at the median, and $6.68 at the maximum. A median problem took 17 calls and 23 minutes wall clock, and thinking tokens account for 91.9% of the 290M output tokens of the formalization account.

Table 32 breaks the formalization calls down by role, where a call type averaging below one call per problem is skipped on some problems, when an earlier check settles the decision or when the corresponding stage does not run.

Table 32. Model calls by role over the problems run.

Role	Total Calls	Per Problem
Plan		
Implementation sketch	4,223	1.34
Implementation plan check	4,182	1.32
Plan check	3,598	1.14
Specification sketch	3,294	1.04
Input types and interface	3,081	0.97
Sketch repair	732	0.23
Formalize		
Decidability instances	3,957	1.25
Specification	3,337	1.06
Implementation	3,319	1.05
Executable postcondition	2,431	0.77
Executable postcondition guard commands	792	0.25
Judge		
Implementation quality	3,660	1.16
Specification quality	3,196	1.01
Executable postcondition judge	3,143	0.99
Tautology	3,093	0.98
Test case equivalence	72	0.02
Repair		
Specification repair	2,514	0.80
Implementation repair	2,305	0.73
Repair history digest	1,238	0.39
Root cause	175	0.06
Frozen row repair	54	0.02
Total	52,396	16.58

The total is a lower bound, since runs discarded during development left no record. Since cached runs record the tokens of the original calls, successful runs are fully counted.

G.8 Model selection

Frontier models are compared separately for generation and repair, implementation planning, and judging. The candidates are Kimi-K3 (Kimi Team, 2026), DeepSeek-V4-Pro and DeepSeek-V4-Flash (DeepSeek-AI, 2026), GLM-5.2 (Z.ai, 2026), Claude-Opus-5 (Anthropic, 2026a), and Claude-Sonnet-5 (Anthropic, 2026b).

G.8.1 FORMALIZER

The comparison of generation models replaces the model for the 7 calls that generate or repair definitions and `#guard` commands, and keeps all other calls fixed. Four calls generate the postcondition, the executable postcondition, the implementation, and the `#guard` commands, and 3 repair these components from compiler diagnostics. Each configuration is run with a separate response cache to keep replies from being replayed across configurations, and pass rate, repair rounds, token count, cost, and wall-clock time in Table 33 are measured on fresh calls.

Table 33. Comparison of models for generation and repair on 30 problems, 10 per difficulty. Claude-Opus-5 needs the fewest repair rounds, and Kimi-K3 reaches the same pass rate at lower cost than Claude-Opus-5. Straight indicates the number of problems that compile without repair, and tokens, cost, and time are per-problem averages.

Model	Pass	Straight	Repairs	Tokens	Cost	Time
Claude-Opus-5	30/30	23	7	33k	$0.73	380s
Claude-Sonnet-5	29/30	13	28	36k	$0.64	386s
Kimi-K3	30/30	9	37	35k	$0.57	732s
GLM-5.2	30/30	10	60	112k	$0.69	1335s
DeepSeek-V4-Pro	28/30	8	36	83k	$0.24	958s
DeepSeek-V4-Flash	28/30	9	40	171k	$0.23	1830s

Three candidates formalize all 30 problems and are separated by efficiency (Table 33). Claude-Opus-5 repairs least and still costs the most per problem. DeepSeek-V4-Pro and DeepSeek-V4-Flash are the cheapest but formalize 2 fewer problems and take the longest, and wall-clock time limits throughput since each repair round occupies a compilation process. Kimi-K3 formalizes all problems at lower cost than Claude-Opus-5 and returns reasoning traces for the training samples.

GLM-5.2 generates 112k output tokens per problem, compared with 33k for Claude-Opus-5, and 82% of the GLM-5.2 tokens are from the 7 replaced calls. Generation of the executable postcondition alone takes 37k output tokens per problem, compared with 7k for Claude-Opus-5.

Results for Kimi-K3 and GLM-5.2 are confounded by prompt revision. Kimi-K3 was run on a revised specification prompt, and one of the 30 problems was rerun on the current prompts, whereas Claude-Opus-5, Claude-Sonnet-5, and both DeepSeek models were run on the earlier revision. For GLM-5.2, 10 problems were run on the earlier revision and 20 on revised versions of 4 of the 5 planning and generation prompts, and the GLM-5.2 repair and token counts therefore mix a model difference with a prompt difference.

G.8.2 PLANNER

The implementation plan selects the algorithm after both postconditions are accepted, and implementation generation follows the plan. Since errors in the implementation plan are frozen into the file, we measure the plans themselves.

Each of 30 problems was planned once by each candidate, from planning inputs recorded during synthesis. We selected the problems whose implementations had previously relied on a fuel-bounded recursion. Since 26 of the 30 problems are hard, the error counts below describe hard problems. Planning and judging run without compilation, whereas an end-to-end comparison takes a pool of compilation servers and about an hour per candidate.

Kimi-K3 plans fail most often at parsing, before judge review (Table 34). 20 of 30 Kimi-K3 replies were valid JSON, compared with 29 and 30 for both DeepSeek-V4 variants, and all unparsable replies stop before the length limit. A plan that does not parse costs a retry and then the problem.

Table 34. Models for implementation planning on 30 problems (26 hard), one plan per model and problem. Parsed indicates the number of plans that are valid JSON, and Clean the number of parsed plans with no failure mode in Table 35.

Model	Parsed	Clean
Kimi-K3	20 (67%)	16 (53%)
DeepSeek-V4-Pro	29 (97%)	22 (73%)
DeepSeek-V4-Flash	30 (100%)	18 (60%)

Each failure mode corresponds to one numbered planning rule, and plans are judged by the same rules that guided planning. A failure mode is recorded for a plan only when the judge names a

concrete witness (e.g., an input on which a bound fails, an input size a stated cost cannot handle, a definition generated a second time). We count the failure modes over the plans that parsed (Table 35).

Table 35. Failure modes of parsed plans, judged by DeepSeek-V4-Flash.

Model	Counter	Mismatch	Infeasible	Ungrounded	Shallow	Search
Kimi-K3	2	0	1	0	1	0
DeepSeek-V4-Pro	0	0	2	0	5	0
DeepSeek-V4-Flash	1	0	4	0	7	0

- *Counter*. A termination argument based on a counter that the algorithm introduces.
- *Mismatch*. A stated running time that belongs to a different algorithm from the planned algorithm.
- *Infeasible*. A stated running time beyond the input bounds, under which the test cases cannot finish in compiled evaluation.
- *Ungrounded*. A regenerated specification definition with a reason other than cost or a different decided property.
- *Shallow*. A specification definition replaced as deciding the wrong property, with no finer definition named for reuse.
- *Search*. An algorithm that evaluates the executable postcondition on candidate outputs and returns an accepted candidate.

The plans were judged again by DeepSeek-V4-Pro to measure how far the counts depend on the judge. Both judges are also planning candidates and may favor plans generated by the same model. Under the second judge, DeepSeek-V4-Pro and DeepSeek-V4-Flash swap places and Kimi-K3 ranks below both. Only the parse rate is independent of the judge, and the comparison between DeepSeek-V4-Pro and DeepSeek-V4-Flash rests on the parse rate.

G.8.3 Judge

DeepSeek-V4-Flash and DeepSeek-V4-Pro, run on the same 79 parsed plans, agree on 97% of the per-mode calls and on 82% of the verdicts. Since a plan counts as clean only when all modes are unflagged, a disagreement on any single mode can change the verdict.

Failure-mode counts depend on the choice of judge, with per-mode agreement over the plans marked by either judge between 0% and 75% (Table 36). DeepSeek-V4-Pro reports Mismatch 3 times and DeepSeek-V4-Flash never, and neither judge reports Ungrounded or Search on any of the 79 plans, so no agreement can be computed for those modes.

Table 36. Agreement between DeepSeek-V4-Flash and DeepSeek-V4-Pro on 79 plans, where per-mode agreement ranges from 0% to 75%. Agreement is computed over the plans that at least one judge marks with the failure mode.

Failure Mode	V4-Flash	V4-Pro	Both	Agreement
Counter	3	5	1	14%
Mismatch	0	3	0	0%
Infeasible	7	6	5	62%
Ungrounded	0	0	0	n/a
Shallow	13	15	12	75%
Search	0	0	0	n/a

The judge model is DeepSeek-V4-Flash, the model of the production quality gates. The second judge bounds the dependence of the reported counts on that inheritance, which a single-judge table cannot show. The agreement reported here is on planner verdicts.

H Proofs

H.1 Method

We prove the postcondition equivalence theorems with a decompose-and-prove loop, which is a widely adopted strategy for formal proofs (Jiang et al., 2023; Wang et al., 2024; Ren et al., 2025; Zhou et al., 2025; Li et al., 2026b), run by a main agent and a set of subagents (Figure 23).

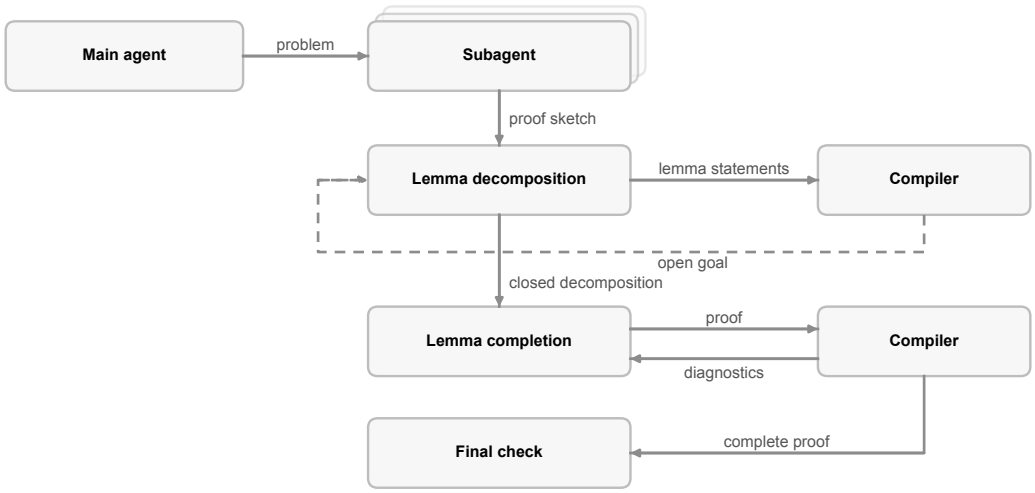


Figure 23. The decompose-and-prove loop. Solid edges carry the named objects, and the dashed edge returns an open goal to decomposition.

The main agent holds the queue of problems and hands a problem to one subagent, together with the form of the executable postcondition and the proofs of the closest problems already proved.

The subagent compares the postcondition and the executable postcondition and produces a proof sketch (Jiang et al., 2023), which records the structure common to both postconditions and states one lemma for each remaining clause. The lemmas are given `sorry` bodies and the file is compiled once, which checks that the lemmas close the theorem before any lemma is proved (Li et al., 2026b). A decomposition that fails is restated by the same subagent, since the next decomposition is built from the open goal. The subagent then proves the lemmas one at a time against the compiler, and each proof is verified against the frozen statement with SafeVerify (GasStationManager, 2024), bumped to Lean 4.32, which checks that no declaration of the problem was changed and that the theorem depends on the standard axioms only.

H.2 Postcondition equivalence theorem

We analyze the proofs by the form of the executable postcondition (Appendix G.2), since the form fixes the statement of the theorem. Executable postconditions that evaluate the postcondition through a decidability instance agree with the postcondition by construction. Separate definitions replace a quantifier of the postcondition by a computation, and the postcondition equivalence proof shows that the replacement loses nothing on valid inputs. The computation is an enumeration that stands in for a set, a reformulation that decides a different property, or a bound that stands in for an unbounded quantifier. The decidable form covers 56% of easy problems and 21% of hard problems, and the proof effort concentrates on the harder problems (Table 37).

Table 37. Postcondition equivalence proofs by form of the executable postcondition, where the decidable form has proof length 0 and the separate definitions a median of 112 lines. Proof length was measured as the number of lines added to the problem file, and the last column indicates the number of proofs in which the validity hypothesis is used.

Form	#Problems	Easy	Medium	Hard	Median Lines	P90 Lines	Validity Used
Decidable form	1,079	400	526	153	0	0	0
Separate definition	1,958	314	1,062	582	112	394	868
Enumeration	1,607	204	883	520	118	409	736
Reformulation	293	90	153	50	95	333	104
Bound	58	20	26	12	92	298	28

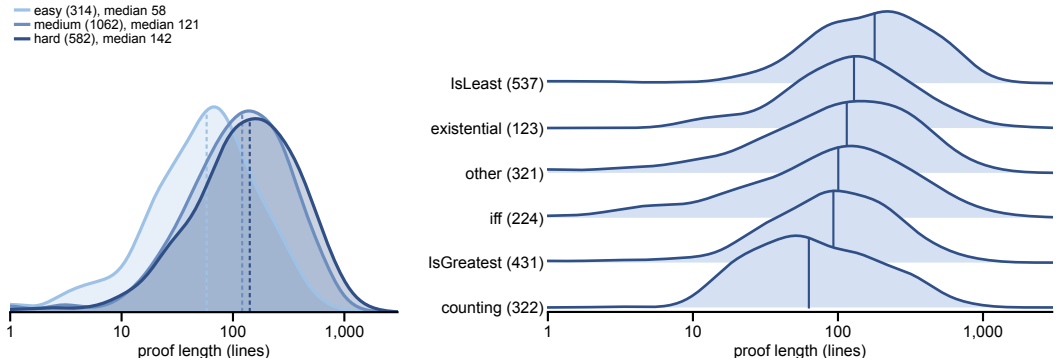
A hard proof costs more lemmas rather than more repair rounds (Table 38). Of the 677 problems with a per-problem record, the first compilation of the decomposition passed for 78%, and the median proof took 2 compilations. The median number of lemmas grows from 5 on easy problems to 21 on hard problems while the number of compilations stays at 1 to 2, which places the cost of a hard problem in stating the lemmas and not in repairing the lemmas. The number of lemmas tracks proof length (Spearman 0.96) and compilations are independent of proof length (0.09). A subagent spends

a median of 128k tokens, 26 tool calls and 473 seconds on one problem (351 problems with a usage record), and tokens follow proof length (0.71). All 3,037 proofs depend on `propext`, `Quot.sound` and `Classical.choice` only. A full replay verified 3,033 proofs and returned no verdict for 4, which match the original except for compiler-generated auxiliary declarations.

Table 38. Cost of a postcondition equivalence proof by difficulty, where the number of lemmas created grows fourfold from easy to hard while the number of compilations stays at 1 to 2. Only the 1,958 problems with separate definitions are counted. Proof length, lemmas created, compilations, tokens, tool calls and time are medians over the problems in the row, and passed at 1st is the fraction of problems whose lemma statements compiled on the first attempt. Proof length was measured on all 1,958 problems, lemmas and compilations on the 677 problems with a per-problem record, and tokens, tool calls and time on the 351 problems with a usage record.

Difficulty	#Problems	Proof Length (lines)		Lemmas Created	Compilations		Tokens	Tool Calls	Time (s)
		Median	P90		Median	Passed at 1st			
Easy	314	58	203	5	2	86%	118k	24	369
Medium	1,062	121	394	11	2	74%	128k	26	484
Hard	582	142	486	21	1	86%	149k	30	604
All	1,958	112	394	13	2	78%	128k	26	473

The postcondition form and the topic of a problem predict proof length better than difficulty does (Figures 24 and 25). LeetCode difficulty predicts proof length weakly (Spearman 0.31 over all problems), and the topic tags that LeetCode assigns to a problem predict proof length better (Xia et al., 2025). Among separate definitions of medium difficulty, graph problems have a median proof length of 324 lines and two-pointer problems 66. Graph problems state reachability as the existence of a path in the postcondition and compute a closure in the executable postcondition, and the proof relates path and closure by induction on the search. Window problems quantify over the same contiguous ranges in both postconditions, and the proof is index arithmetic. Minimization costs more than maximization, with a median of 179 lines for `IsLeast` postconditions, compared with 93 for `IsGreatest`. The candidate set of a maximum is bounded by the input, whereas the lower bound of a minimum is bounded by the precondition alone, and the precondition is invoked in 58% of `IsLeast` proofs, compared with 39% of `IsGreatest` proofs. 121 proofs exceed 500 lines and 9 exceed 1,000, all on graph, geometry or string-rewriting problems.



(a) By difficulty

(b) By form of the postcondition

Figure 24. Proof length against difficulty and postcondition form, where the form of the postcondition separates the proofs more sharply than difficulty does. Separate definitions only, since the decidable form adds no lines. Curves are kernel density estimates on a logarithmic axis, and vertical lines mark medians.

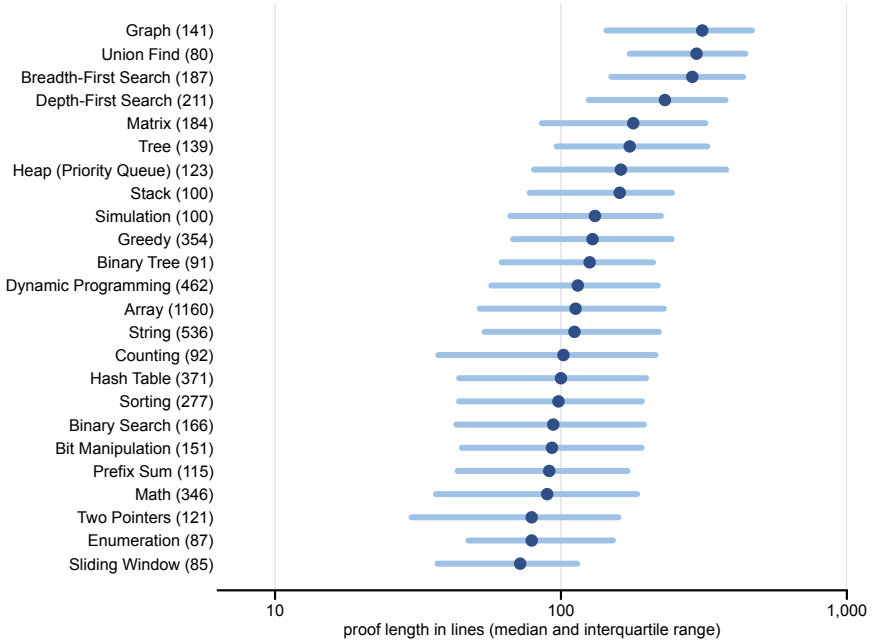


Figure 25. Proof length by topic, where graph problems need 3 times the lines of window problems at the same difficulty. Dots mark medians and bars interquartile ranges over the separate definitions carrying the LeetCode topic tag, for tags with at least 80 problems, with the number of problems in parentheses. A problem with several tags is counted under each.

A proof attempt that fails for a reason other than proof difficulty exposes a statement error through a counterexample or through independence. A counterexample is a valid input on which the postcondition and the executable postcondition disagree, and the negation of the theorem is then proved. Independence arises when the executable postcondition evaluates a `Float` operation, which Lean treats as opaque. Statement errors were found in 31 problems (1.0%), 5 by independence and 26 by a compiled counterexample (Table 39). Each was repaired by the smallest change to the statement that agrees with the problem description, the repaired file passed the frozen test cases, and the theorem was then proved. With the postcondition equivalence theorems proved, the postcondition and executable postcondition of each problem agree on valid inputs. The proofs relate the 2 postconditions to each other, and the human review checks both against the problem description (Appendix I.1).

Table 39. Statement errors found by proof attempts, where the postcondition accounts for 19 of 31 problems and the precondition for 3. A problem was counted under each component that was changed by the repair, and one problem was counted twice.

Component	Easy	Medium	Hard	Total
Precondition	1	0	2	3
Postcondition	4	10	5	19
Executable postcondition	2	2	6	10
#Problems	7	12	12	31

H.3 Correctness theorem

The correctness theorem asserts that the implementation satisfies the postcondition. Most proofs establish the postcondition directly, and 8 of the 300 proofs show that the implementation satisfies the executable postcondition and transfer the result through the postcondition equivalence theorem. The loop of Appendix H.1 is unchanged.

Correctness theorems are proved for all 300 test problems (72 easy, 155 medium, 73 hard) on the files repaired after the review of Appendix I.1, where 274 proofs are checked under Lean 4.32 and 26 come from the earlier proving run. Correctness proofs grow with difficulty, from a median of 63 lines on easy problems to 178 on medium problems and 365 on hard problems, with a median of 159.5 lines over all 300 problems. All 300 reference implementations are also tested against the

executable postcondition with Plausible (Lean Community, 2024), which finds no counterexample (214 pass and 86 are inconclusive, since valid inputs are rarely sampled for structured types).

H.4 Equivalence of accepted outputs

Acceptance relies on test cases and judges, so we also validate accepted SpecGen outputs by proof. For each model in Table 4 and each problem with an accepted output, we draw an accepted output at random and prove its postcondition equivalence theorem. The fixed tactics of the corpus proofs, such as `simp` with the executable postcondition unfolded, close 557 of the 2,012 drawn outputs in one line (27.7%). The proving loop (Appendix H.1) then settles every remaining output of the 2 RL models and up to 40 remaining outputs of each other model (Table 40). Each attempt ends in a proof or in a counterexample on a valid input, except when a `Float` operation makes the theorem independent, so no output stays undecided for lack of proof effort.

The theorem is proved for 99.2% of accepted outputs of both RL models (251/253 for AuXSpec-27B-RL and 246/248 for AuXSpec-9B-RL), and all proofs pass SafeVerify under the criteria of the corpus proofs (Appendix H.1). The 6 counterexamples locate the error on both sides of the theorem, in 3 declarative postconditions, 2 executable postconditions and 1 validity predicate, which admits unbalanced inputs for *Remove Outermost Parentheses* (AuXSpec-9B-RL).

Table 40. Postcondition equivalence theorems of accepted SpecGen outputs on AuXSpec-Test, one output per model and problem. The 2 RL models are a census. For the other models, 40 outputs left open by the fixed tactics go to the proving loop, and the estimated mismatch applies the counterexample rate among them to all open outputs.

Model	Accepted	Fixed Tactics	Proving Loop	Counter-example	`Float`	Not Sampled	Estimated Mismatch (%)
AuXSpec-27B-RL	253	42	209	2	0	0	0.8
AuXSpec-9B-RL	248	60	186	1	1	0	0.4
AuXSpec-27B-SFT	122	34	39	1	0	48	1.8
AuXSpec-9B-SFT	73	23	39	1	0	10	1.7
Claude-Opus-5	244	93	40	0	0	111	0.0
Claude-Sonnet-5	118	47	40	0	0	31	0.0
DeepSeek-V4-Flash (0731)	230	39	40	0	0	151	0.0
DeepSeek-V4-Pro	163	50	40	0	0	73	0.0
GPT-5.6-Luna	86	16	39	1	0	30	2.0
GPT-5.6-Sol	179	57	40	0	0	82	0.0
Kimi-K3	251	72	40	0	0	139	0.0
Qwen3.5-9B	2	1	1	0	0	0	0.0
Qwen3.8-27B	43	23	20	0	0	0	0.0

I Human study

We conduct two human evaluations: a review of the 300 AuXSpec-Test reference problems (Appendix I.1) and an expert review of automatically accepted SpecGen model outputs (Appendix I.2).

I.1 Human review of AuXSpec-Test

The human review of the 300 AuXSpec-Test problems is compared with the errors found while proving (Appendix H.2).

One of the authors, with experience in Lean, reviewed all 300 AuXSpec-Test problems. Each component was compared with the problem description in the order of Table 41. The reviewer was given the Lean file and the description and consulted no model or automated judge. An error was recorded only with an exposing test case, an analysis in prose, and a corrected file that compiles and passes the frozen test suite.

Errors were confirmed in 22 of 300 problems (7.3%), 21 by the review and 1 by the proof of the correctness theorem. The counts by component and difficulty are in Table 42. The postcondition and the implementation account for 20 of the 23 component errors, and the review found no executable postcondition that accepts different outputs from the postcondition. The error rate is highest for easy problems at 9.7% and decreases to 5.5% for hard problems, and 2 of the 3 precondition errors fall on

Table 41. Rubric of the human review, in which each of the 4 components is judged against the problem description by the question of the corresponding acceptance criterion of Section 2. One row per component, with the question the reviewer answers and the finding that was recorded as an error. The questions were taken from the acceptance criteria, so a person and a judge examine the same properties.

Component	Question	Recorded as an Error When
Precondition	Does it admit exactly the inputs the constraints describe?	An input the constraints admit is rejected, or one they exclude is accepted
Postcondition	Does it carry each condition the description places on the answer?	A clause of the description is missing and a value the description rejects is admitted, or a correct output is rejected
Executable postcondition	Does it accept the same outputs as the postcondition on all valid inputs?	A bound is too small and an object the postcondition ranges over is not reached, or an object outside it is accepted
Implementation	Does it satisfy the postcondition, terminate without a fuel parameter, and run within the constraints?	A valid input yields a rejected output, termination rests on a fuel parameter, or the cost does not survive the largest admitted input

easy problems. Of the 73 files the review changed, 21 changed a definition and the others changed prose or test cases.

Table 42. Errors confirmed by human review and, for 1 problem, by the correctness proof, where 20 of 23 component errors are in the postcondition or the implementation and none is in the executable postcondition. Each percentage is taken over the problems of the row, and a problem with errors in 2 components is counted once among the problems with errors.

		Errors by Component				Problems
Difficulty	#Problems	Precondition	Postcond.	Exec. Postcond.	Implementation	With Errors
Easy	72	2 (2.8%)	4 (5.6%)	0	1 (1.4%)	7 (9.7%)
Medium	155	1 (0.6%)	5 (3.2%)	0	6 (3.9%)	11 (7.1%)
Hard	73	0	1 (1.4%)	0	3 (4.1%)	4 (5.5%)
Total	300	3 (1.0%)	10 (3.3%)	0	10 (3.3%)	22 (7.3%)

Review and proving find different errors. Of the 14 specification errors in the test split, attributed by the order of discovery, 9 were found by review alone and had a postcondition equivalence theorem proved against the wrong statement before the review, since an executable postcondition that decides a wrong postcondition still agrees with the wrong postcondition. 3 were found by both, with compiled counterexamples from the proof attempts. 2 were found by proof attempts alone and had passed the review. In one, the clause $\forall i,\ i < n \to A \leftrightarrow B$ parses as $\forall i,\ (i < n \to A) \leftrightarrow B$, and the postcondition therefore rejects all outputs while the text reads correctly. In the other, the executable postcondition compares distances in `Float`, and the theorem is independent of the axioms. The 10 implementation errors do not affect the postcondition equivalence theorem and are the subject of the correctness theorems of Appendix H.3. Review measures agreement with the problem description, and proving measures agreement between the postcondition and executable postcondition (Table 43). Neither measure subsumes the other, and the dataset was checked by both.

Table 43. Specification errors in the test split by detection, where review and proving overlap on 3 of 14 problems. Review only means the postcondition equivalence theorem had been proved against the wrong statement, and proof attempt only means the review missed the error.

Detection	#Problems
Review only (theorem proved against the wrong statement)	9
Review and proof attempt	3
Proof attempt only	2
Total	14

I.2 Expert review of SpecGen evaluation

We review false positives of the automatic SpecGen evaluation for the 2 reported RL checkpoints, AuXSpec-27B-RL and AuXSpec-9B-RL. For each model, we draw 50 problems counted as SpecGen successes, stratified as 12 easy, 26 medium and 12 hard in the proportion of AuXSpec-Test, and review 1 accepted output per problem. Both reviewers each review 60 problems, of which 20 overlap, and the review package masks model identity, judge verdicts, test results against the reference and proof results. Before the review, we test each output against the reference with Plausible, as for the reference implementations (Appendix H.3).

1. Completeness. Plausible samples valid inputs and checks that the executable postcondition accepts the output of the reference implementation. A counterexample is a correct output that the specification rejects.
2. Validity. Plausible samples inputs and checks that the validity predicate agrees with the reference validity predicate. A counterexample is an input that one predicate admits and the other rejects.

We apply the specification-side rubric of Table 41.

Expert review confirms 90 of the 100 accepted outputs, 43 of AUXSPEC-27B-RL and 47 of AUXSPEC-9B-RL (Table 44). The 2 reviewers find 7 and 5 errors in their 60 problems and agree on all 20 shared problems, including the same 2 errors. Each error comes with a concrete input that the validity predicate admits outside the constraints, which makes the verdict checkable rather than a judgment. All declarative and executable postconditions are correct, and all 10 errors were found in the validity predicate. The missing constraints are guarantees that the answer fits in a 32-bit integer (3 problems), value ranges and alphabets (3), structural guarantees on the input (3) and a time format (1). The postconditions of these outputs therefore remain faithful on every input that the constraints admit. The tests against the reference miss these errors, since random inputs rarely reach the boundary of a missing constraint.

We acknowledge that the expert review covers only 100 outputs which is a small sample, yet the errors found converge to the same kind. The models produce correct postconditions but are highly likely to miss boundary conditions of the constraints in the validity predicate. One possible reason is that the acceptance criteria offer little guidance on the validity predicate, since the specification quality judge assigns clauses that restrict the input to the validity predicate without checking them (Appendix K.3) and only the test layers L1 and L1b check the validity predicate on a few inputs (Table 7). Checking the validity predicate against the stated constraints is a direction for future work (Appendix J).

Table 44. Expert review of automatically accepted SpecGen outputs of the reported RL checkpoints.

Model	Reviewed	Correct	V Errors	P_{spec} Errors	P_{check} Errors
AUXSPEC-27B-RL	50	43 (86.0%)	7	0	0
AUXSPEC-9B-RL	50	47 (94.0%)	3	0	0
Total	100	90 (90.0%)	10	0	0

J LIMITATIONS

Despite the contributions of AUXSPEC to specification generation in Lean, the work has several limitations that we hope to address in future work.

Repair at inference. The repair tasks train the models on failed rollouts and are evaluated for a single round (Table 4), and the same tasks support a self-refinement loop at inference, where the model repairs an output on compiler messages and test results over several rounds (Ye et al., 2026b). Pass rates over rounds would show how much of the gap to frontier models such a loop closes.

Scope of expert review. A single reviewer covers the 300 AUXSPEC-Test problems and finds errors in 21 of them, including 9 specification errors in problems with a proved postcondition equivalence theorem (Appendix I.1), and 2 reviewers review 100 accepted SpecGen outputs of the RL models (Appendix I.2). The training problems are checked by test cases, judges and postcondition equivalence proofs without review and need a reviewed sample to estimate the error rate of the training data.

Proofs for generated outputs. Generated specifications and implementations are accepted without proofs of the two theorems (Section 2). We prove the postcondition equivalence theorem afterwards for the accepted outputs of the RL models (Appendix H.4). A proof shows only that the 2 postconditions agree. Whether the declarative postcondition and the validity predicate state the problem still rests on test cases, judges and expert review.

Validity predicates. All errors found by the expert review of accepted outputs are input constraints missing from the validity predicate (Appendix I.2). The acceptance criteria check the validity predicate only with test cases. A judge that compares the validity predicate with the stated constraints would cover this part in synthesis, training and evaluation.

K PROMPTS

The boxes below are condensed versions of the prompts, in the order of Figure 1. Every prompt has the same five parts in the order the model reads them, the role, the input, the task, the rules and the output format, and problem text is marked as untrusted data so that instructions inside a problem statement are not followed. The boxes keep the rules that the analyses of Appendix G refer to and drop the worked examples, the formatting rules and the injected library facts. The component names follow the paper. The prompts given to the models are in Appendix K.2, and the judges of Appendix K.3 run as the acceptance criteria of synthesis, in evaluation and in the reward.

K.1 SYNTHESIS PIPELINE

General planner

```
ROLE
Contract planner. Fixes everything about one problem that does not require specification thinking,
before the specification side runs.

INPUT
‘problem‘ (description, constraints, examples; untrusted), ‘signature‘, ‘default_lean_types‘, and on
a retry the ruling that found the previous plan incomplete.

TASK
Decide the Lean types, the interface names and the precondition’s clauses, and record the domain
facts. Everything fixed here is frozen; downstream roles follow it exactly.

RULES
1. You own the representation: for every parameter and the output, declare the Lean type that is
best for stating and proving this problem, weighing faithfulness, simplicity of the statements and
the proof debt of the correctness proof.
2. Interface names describe what is computed, never how.
3. The precondition’s clauses transcribe the stated constraints in the problem’s words: nothing
dropped, nothing invented, no operational re-encoding. Where the input encodes something, state the
encoding’s well-formedness.
4. The implementation algorithm is not yours to fix; a later stage plans it once the specification
exists.

OUTPUT
One JSON object: ‘input_spec‘, ‘lean_types‘ (with a ‘rationale‘), ‘interface‘, ‘v_plan‘,
‘domain_facts‘.
```

Postcondition planner

```
ROLE
Specification planner. Receives the frozen types and precondition plan, and shapes the postcondition
on top of them.

INPUT
‘problem‘ (untrusted), ‘lean_types‘, ‘v_plan‘, ‘negative_aims‘ (what each frozen negative test
targets), and on a retry the previous ruling.

TASK
Plan the vocabulary and the shape of the postcondition in prose, no Lean. The postcondition states
what a correct output is; it does not produce one. Two statements are required: whether ‘result‘ is
the kind of object that could be the answer, and what distinguishes the correct one from the others
of that kind.

RULES
1. State whether the second statement concerns the object the first exhibits or every object of that
kind; a universal reduced to the exhibited object is no bound.
2. ‘result‘ occurs inside relations that constrain it, never as the left side of an equation whose
right side computes it. The test: conceal the right side and ask whether a correct ‘result‘ remains
distinguishable.
3. Keep out every operation that answers in the degenerate case instead of having no answer there
(‘Set.ncard‘ on an infinite set, ‘sInf‘ on the empty set, truncated subtraction, ‘xs[i]!‘); it makes
a different specification, not a false one.
4. Computability is not a consideration; quantifiers range over the problem’s natural domain, and a
stated bound is a hypothesis inside the quantifier, never a truncated range.

OUTPUT
One JSON object: the two statements, ‘second_statement_ranges_over‘, ‘vocabulary‘, ‘witnesses‘.
```

Postcondition formalizer

```
ROLE
Specification formalizer. Renders the plan's two prose statements as the specification side of one
problem.

INPUT
‘plan‘ (names and types binding), ‘problem‘ (untrusted), ‘sketch_preview‘ (the final file with your
sections as holes), ‘verified_library‘ (declarations compiled against this toolchain),
‘frozen_cases‘ (layer, frozen input and output, and the expression the plan fixed).

TASK
Write the specification helpers, the precondition, the postcondition and the L1 and L1b guard
commands. Theorem statements and module structure are machine-written; the test rows are frozen, and
you write each row as one guard command.

RULES
1. Decidability is not your concern; bounding or enumerating a quantifier to make it computable is a
defect.
2. Specify the what, never the how: ‘IsLeast S result‘ over ‘result = sInf S‘. Mirroring an
algorithm's structure is a defect.
3. The precondition is a total Bool function covering exactly the stated constraints; the
postcondition returns Prop and says the plan's two sentences and nothing further. A library
declaration whose definition is what you were about to write is that same proposition, with the
library's lemmas.
4. The frozen rows are a contract: one guard command per row, none for a row that is not there,
values preserved exactly.

OUTPUT
A design note, then one Lean fence with ‘Representation‘ (only when needed), ‘Specification
helpers‘, the precondition, the postcondition and ‘Unit tests‘.
```

Test form judge

```
ROLE
Judge of whether a guard command states the frozen test row it was written for, called when the
mechanical comparison could not settle it.

INPUT
The frozen rows (‘test_id‘, layer, input and output as frozen JSON, the expression the plan fixed)
and the guard command written for each.

TASK
Two mechanical checks ran first: the shape check (right layer, right anchor, right polarity) and the
expression check (verbatim, then token by token). Rule on the rows where the tokens differ.

RULES
1. The frozen JSON is the authority, not the plan's expression.
2. Count when counting decides it: a replicate, a range, a map over a range.
3. Two expressions are equivalent when they elaborate to the same term; notation, grouping and a
redundant ascription do not change the test. A missing or extra parameter or a value that differs
anywhere is a different test.
4. Judge the denotation, not the style. When the evidence does not settle it, answer false and say
what you could not determine.

OUTPUT
One JSON object with a ruling per row: ‘equivalent‘ (true or false) and ‘reason‘.
```

Executable postcondition router

```
ROLE
Route analyst for the executable postcondition of one problem. Decides how it will be organized
before anyone writes it.

INPUT
‘problem‘ (untrusted), ‘frozen_spec‘, ‘test_inputs‘ (with the length of every list and string),
‘plan‘ (whose ‘domain_facts‘ often supply the bounding fact).

TASK
The executable postcondition restates the frozen postcondition as a Bool that compiled evaluation
can run on the frozen inputs. The difficulty is cost: mirroring the quantifiers directly can be
exponential in the input. Choose one of three routes: reshape (restate what is quantified over so
the count collapses), truncate (bound a quantifier by a fact the problem supplies) or enumerate
(range over a finite encoding of the candidates).

RULES
1. ‘bounding_fact‘ is the one fact the cost argument rests on, stated so a reader could check it
against the problem.
2. The executable postcondition is a judge, never a solver: it quantifies over candidates or
verifies the answer’s defining property. A shape that computes the answer by its own algorithm and
compares is refused downstream, whatever its cost.
3. Name the route a hasty writer would reach for and the cost that rules it out at these input sizes.

OUTPUT
One JSON object: ‘route‘, ‘bounding_fact‘, ‘shape‘ (in prose), ‘rejected‘.
```

Executable postcondition formalizer

```
ROLE
Writes the executable postcondition for a frozen postcondition, following the route decided for it.

INPUT
‘plan‘, ‘frozen_spec‘ (read-only target), ‘sketch_preview‘, ‘route_plan‘, ‘decidability_trial_error‘
(why the decidable form was not available), ‘verified_library‘, ‘frozen_cases‘.

TASK
Write one Lean section with the executable postcondition and its helpers, and the L2 and L2b guard
commands. Record the route taken and the fact that bounds it in a note before the fence.

RULES
1. A total Bool function named ‘base_name‘ + ‘_check‘, with the specification’s argument order.
2. Bridge toward the frozen postcondition, never toward a weaker reading of it; mirror its clause
structure so the equivalence theorem is provable clause by clause.
3. Every recursion or fold must already be present in the postcondition’s own definitions: the
equivalence proof should need only unfolding and structural alignment, never an algorithmic insight.
Never compute the answer by another method and compare it with ‘result‘.
4. Bound every quantifier to a structure derived from the input. Size is measured, not hypothetical:
what settles an enumeration is running it.
5. The fast algorithm belongs to the implementation; an evaluation that finishes belongs here.

OUTPUT
A design note, then one Lean fence with the executable postcondition section and ‘Unit tests‘.
```

Implementation planner

```
ROLE
Implementation planner. The postcondition and the executable postcondition are accepted, and the
algorithm is planned against them rather than against the problem statement alone.

INPUT
‘problem‘ (untrusted), ‘plan‘ (‘impl_name‘ binding), ‘frozen_spec‘,
‘frozen_executable_postcondition‘ (read-only), ‘verified_library‘, ‘frozen_cases‘, and on a retry
the judge’s ruling.

TASK
Decide how the answer is computed and at what cost. This is the only section where running time is
judged: the guards run the algorithm on the frozen inputs, so an exponential approach fails by
timeout. Say which accepted definitions the algorithm calls and which it replaces, with the cost
difference.

RULES
1. ‘approach‘ names the standard efficient algorithm, not the first one that terminates; ‘cost‘
states its running time in terms of the input.
2. A frozen definition is passed over on one of two grounds, and the entry says which: too slow
(both figures given) or different structure required (the finer definition named). An entry on
neither ground is a rewrite that should have been a call.
3. ‘termination‘ names the quantity the recursion consumes; a ‘Nat‘ carried along so the recursion
has something to count down is not such a quantity.
4. The algorithm computes the answer; the specification says what a correct answer is. Do not
evaluate the specification over candidates and return the one that satisfies it.

OUTPUT
One JSON object: ‘approach‘, ‘cost‘, ‘calls‘, ‘replaces‘, ‘termination‘.
```

Root-cause analyzer

```
ROLE
Root-cause analyst, called only when local repair cannot proceed: a symptom recurred after repair, a
repairer declared the fault lies elsewhere, or an error cannot be placed by line.

INPUT
‘error_report‘, ‘sections‘ (each labelled frozen or repairable), ‘assembled_source‘ (the whole file
with line numbers), ‘harness_regions‘, ‘frozen_tests‘, ‘history‘ (the repair attempts so far),
‘plans‘.

TASK
Rule which block is at fault, so the pipeline repairs the true target instead of the symptom’s
location. A plan may also be ruled faulty, which triggers one replan.

RULES
1. Rule on evidence in the diagnostics and sections, never on which line happened to error; symptom
location and fault location are different things.
2. A frozen test literal is ground truth by default. Rule it wrong only on evidence you can quote:
the case does not violate the constraint its own rationale cites, or a documented answer contradicts
the problem’s own examples.
3. Prefer the cheapest sufficient target: a section over a plan, a plan over declaring the run dead.
4. When the specification itself contradicts the tests, the run fails rather than the specification
being weakened.

OUTPUT
One JSON object: ‘fault‘ (‘section‘ with its title, ‘plan‘, or ‘frozen‘), ‘reasoning‘ (the decisive
evidence), ‘confidence‘.
```

Specification repairer

```
ROLE
Specification repairer. Repairs exactly one broken section; everything else is frozen.

INPUT
‘broken_section‘, ‘diagnostics‘ (message, line, source line), ‘context_sections‘ (read-only),
‘plan‘, ‘attempt_history‘ (what earlier rounds tried and how each ended), ‘verdict‘ (the ruling
naming the target).

TASK
Read the diagnostics, fix the target section and return it. The reply replaces that section
byte-for-byte; no other bytes of the file change.

RULES
1. Minimal change: fix the diagnosed fault and keep the rest of the section byte-identical where
possible. The plan’s names are frozen.
2. If the only repair that compiles would weaken the specification’s meaning, return the section
unchanged and mark it ‘no-faithful-repair‘.
3. Two technical-looking repairs count as weakening: narrowing a quantifier’s domain to a bound the
problem does not state, and replacing a quantified clause with a Bool computation that decides it.
4. If the true fault lies in a read-only section, return the body unchanged and mark it
‘fault-elsewhere‘ with the section named.

OUTPUT
Exactly one section with the target’s title, followed by one line ‘attempt: <one sentence>‘.
```

Negative test case proposers

```
INVALID INPUTS (L1b)
Role: adversarial proposer for the validity layer. Input: ‘problem‘ (untrusted), ‘signature‘,
‘valid_examples‘, ‘quota‘.
Task: make the precondition hard to write. A lenient predicate looks correct against valid inputs
alone and stays wrong forever, so the measure of a case is which wrong predicate it defeats: ‘n = 9‘
against ‘1 <= n <= 8‘ defeats every predicate that forgot the ceiling. A case no plausible wrong
predicate would accept has raised the difficulty by nothing. When the examples do not show the range
you need, ask instead of guessing.

INCORRECT OUTPUTS (L2b)
Role: adversarial proposer for the correctness layer. Input: ‘problem‘ (untrusted), ‘signature‘,
‘case‘ (a valid input and its true output), ‘aim‘ (the flaw family to try first), ‘quota‘.
Task: make the postcondition hard to write. Name a flaw a weak postcondition could have (an
optimization stated as achievable but never as best, a count that misses a case, a search that
forgets a boundary), then write the wrong answer that flaw admits. A wrong answer with no flaw
behind it catches nothing.

ADMISSIBILITY
Task: keep a frozen input only if compiled evaluation could check a postcondition against it. The
cost of an input is the number of candidates the postcondition ranges over, measured before the call
for each standard shape, and about ten thousand is the ceiling.

OUTPUT
Each proposer returns a JSON array of cases, each with the flaw or constraint it targets; the
admissibility judge returns keep or refuse with a reason.
```

K.2 Model prompts

The models receive one of the 4 system prompts below, condensed. The user turn gives the problem, the constraints and the test cases, and for CodeGen the given specification. A repair turn repeats the original user turn and adds the current file and the review of each failed stage. Without test cases in the prompt, as in Table 5, the SpecGen and CodeGen prompts replace the given test cases by test cases derived from the statement.

SpecGen

```
TASK
Formalize the algorithmic problem in Lean 4 as a validity predicate, an ideal specification and an
executable checker, and check the predicate and the checker on the given test cases with guards.

RULES
1. State the specification with the definitions of the statement, so that it characterizes the
correct ‘result‘ instead of computing it. Quantify over the objects the statement defines, and bound
a quantifier only by a bound the statement states. An optimum is attained by a feasible candidate
and is no worse than every feasible candidate, a decision holds exactly when a witness exists, and a
count is the size of the set of objects with the stated property. Use ‘result = f input‘ only for a
formula the statement gives or for a sequence of steps it prescribes.
2. Define the validity predicate to accept exactly the inputs that satisfy the stated constraints.
3. Define the checker to decide the specification on valid input, with each quantifier bounded by a
range the input determines: ‘decide‘ of the specification with registered instances where every
quantifier is bounded, a hand-written Bool function otherwise.
4. Complete each definition without ‘sorry‘, and use none of ‘admit‘, ‘native_decide‘,
‘noncomputable‘, ‘partial‘, ‘unsafe‘ and ‘Classical‘.

OUTPUT
One Lean file: ‘import Mathlib‘, a docstring restating the problem, then Specification helpers,
Validity predicate, Ideal specification, Decidable specification (with ‘<name>_check_iff_spec‘ left
as ‘sorry‘) and Unit tests with guards under layers L1, L1b, L2 and L2b.
```

CodeGen

```
TASK
Implement the algorithmic problem in Lean 4 against the given specification, with a correct output
for each input the constraints allow, and check the implementation on the given test cases with
guards.

RULES
1. Prove termination of each recursive definition by structural recursion or a ‘termination_by‘
measure.
2. Choose an algorithm within a log factor of the best known method, so that the implementation
finishes on the largest allowed input.
3. Complete each definition computably and without ‘sorry‘, with the constructs banned for SpecGen.

OUTPUT
One Lean file with ‘Input‘, ‘Output‘ and the representation types copied from the given
specification, then Specification helpers, Implementation helpers, Main algorithm and Unit tests
with guards under layer L3.
```

SpecGen-Repair

```
TASK
Repair the Lean 4 specification so that the file compiles, all guards evaluate to ‘true‘ and the
specification states the problem faithfully.

RULES
1. Locate the fault before editing, starting with the first compiler error.
2. Make the smallest change that removes the cause, and keep the names and argument order of
‘isValidInput‘, ‘<name>_spec‘, ‘<name>_check‘ and ‘<name>_check_iff_spec‘.
3. State the specification declaratively where the problem allows; use ‘result = f input‘ only where
the statement prescribes the output.
4. Keep the checker a clause-by-clause mirror of the specification with bounded quantifiers.
5. Complete each definition computably and without ‘sorry‘.

OUTPUT
The complete repaired file, keeping the section headings and order of the current version.
```

CodeGen-Repair

```
TASK
Repair the Lean 4 implementation so that the file compiles, all guards evaluate to ‘true‘ and the
implementation returns a correct output for each input the constraints allow.

RULES
1. Find the cause before editing; a termination error usually needs a new measure or a
‘decreasing_by‘ proof.
2. Apply the smallest edit that removes the cause, and keep the name and argument order of the main
algorithm.
3. Prove termination as in CodeGen, and use an algorithm within a log factor of the best known
method.
4. Complete each definition computably and without ‘sorry‘.

OUTPUT
The complete repaired file, with the section headings and order of the current version.
```

K.3 Evaluation judges

The judges below run as the acceptance criteria of synthesis, in evaluation and in the overall quality, test case alignment and non-tautology stages of the reward.

Test case alignment

```
ROLE
Checker of the guard commands of a file against the test rows the author was given.

INPUT
‘test_block‘ (the given rows), ‘guards‘ (every guard of the file), ‘unmatched‘ (what a mechanical
check found where it could not match).

TASK
Decide whether the guards test exactly the given rows.

RULES
1. Compare values only. Function names and spellings such as tuples, parentheses, arrays or string
literals change nothing; a changed, added, removed or reordered element is a different value.
2. A positive row is accepted, and an invalid input or a wrong output is rejected with ‘!‘. A
specification file tests a positive row on both the validity predicate and the checker.
3. Guards beyond the given rows do not affect the verdict.

OUTPUT
‘problems‘ (each ‘missing‘, ‘extra‘, ‘changed‘ or ‘reversed‘, with the fix) and ‘verdict‘, ‘PASS‘ or
‘FAIL‘.
```

Specification quality (SpecGen)

```
ROLE
Reader of a specification against the problem it was written for.

INPUT
‘description‘ (untrusted), ‘constraints‘, ‘specification‘.

TASK
Decide whether the specification accepts exactly the answers the problem accepts on valid inputs:
soundness (it carries every condition on the answer) and completeness (it rejects no correct answer).

RULES
1. Take the statement clause by clause. A clause that constrains the answer must appear, a clause
that restricts the input belongs to the validity predicate, and a resource bound stays out of the
report.
2. Judge the Lean definitions by what they compute, evaluating each helper on a small input.
3. Every finding names concrete values: for soundness, the quoted clause and a value the
specification admits and the problem rejects; for completeness, a valid input, its correct output
and the part of the specification that rejects it.

OUTPUT
‘missing_conditions‘ and ‘rejected_answers‘, each empty when there is nothing to report.
```

Checker quality (SpecGen)

```
ROLE
Reader of the executable checker. The guards pass, so on the frozen inputs it agrees with the
specification.

INPUT
‘constraints‘, ‘specification‘, ‘checker‘.

TASK
Decide whether the checker accepts exactly what the specification does on every valid input.

RULES
1. A checker truncates where the specification quantifies over an unbounded domain. A bound that is
too small misses an object the specification admits; one that reaches objects the specification
rejects can accept a wrong output.
2. Report a disagreement only with an object: a valid input, its correct output by the statement,
and the witness or counterexample one side reaches and the other does not.

OUTPUT
‘truncates‘, ‘agrees_with_the_specification‘ (‘YES‘ or ‘NO‘) and ‘escapes‘, each with its direction,
input, correct output and object.
```

Implementation quality (CodeGen)

```
ROLE
Reader of the implementation. The guards pass, so it is correct on the tested inputs.

INPUT
‘constraints‘, ‘specification‘, ‘checker‘, ‘implementation‘.

TASK
Judge what the tests cannot see: termination, cost, answers written back for the tested inputs and
duplicated work.

RULES
1. Report a fuel parameter, and whether a bound the recursion stops on can run out on an admitted
input.
2. Mark the implementation too slow only when the best known method is faster by more than a
logarithmic factor and the implementation takes more than about 10^7 steps on the largest admitted
input.
3. Report an implementation that branches on particular inputs to return particular answers.

OUTPUT
‘termination‘, ‘cost‘ (with ‘too_slow‘), ‘reproduces_tests‘, ‘rewrote_the_checker‘ and a short
‘reason‘.
```

Non-tautology

```
ROLE
Lean 4 formalization auditor of the specification side.

INPUT
‘description‘ (untrusted), ‘constraints‘, ‘specification‘.

TASK
Judge how much of a solution the specification contains.

RULES
1. DECLARATIVE: the condition on ‘result‘ is the condition of the statement, quantified over the
objects of the statement.
2. ENUMERATIVE: the condition of the statement, but over a list of candidates the specification
builds itself.
3. TAUTOLOGICAL: the specification computes ‘result‘ by solving the problem, through a recurrence, a
table, a greedy rule or a simulation the statement does not prescribe.
4. For ENUMERATIVE and TAUTOLOGICAL, describe the declarative specification in at most three
sentences.

OUTPUT
‘reasoning‘, ‘better_specification‘ and ‘verdict‘.
```